

Myconet: A Fungi-inspired Model for Superpeer-based Peer-to-Peer Overlay Topologies

Paul L. Snyder and Rachel Greenstadt and Giuseppe Valetto
 Department of Computer Science
 Drexel University
 {pls29,greenie,valetto}@cs.drexel.edu

Abstract—Unstructured peer-to-peer networks can be extremely flexible, but, because of size, complexity, and high variability in peers’ capacity and reliability, it is a continuing challenge to build peer-to-peer systems that are resilient to failure and effectively manage their available resources. We have drawn inspiration from the sophisticated, robust, root-like structures of fungal hyphae to design Myconet, an approach to superpeer overlay construction, which models regular peers as biomass, and superpeers as hyphae that attract and concentrate bio-mass, while maintaining strong inter-connections with one another. Simulations of the Myconet peer-to-peer protocol show promising results in terms of network stabilization, response to catastrophic failure, capacity utilization, and proportion of peers to superpeers, when compared to other unstructured approaches.



Fig. 1. Mycelium growing [5]

I. INTRODUCTION

Peer-to-peer (P2P) networks are highly-decentralized, large-scale distributed systems. These networks can be extremely large, scaling to millions of peers, with levels of network membership that can be highly unstable. Many P2P networks require high levels of self-management, resilience in the face of changing conditions, and central authorization or supervision is impractical or impossible. The inspiration for this work came from the root structures of fungi, *mycelia*, and the realization that these systems display self-organizing properties and can be used as a guiding metaphor for improving overlay networks in peer-to-peer systems. Fungi-inspired models have been previously proposed as a paradigm for pervasive adaptive systems [1], [2].

Overlay networks impose a topology on top of these often chaotic networks, often acting as an enabler for other services (such as search or routing). Overlays may be either structured (as, for example, with deterministically-placed distributed hash tables) or unstructured.

We have developed a biologically-inspired model, called *Myconet*, for the construction of a robust overlay model within unstructured P2P networks with superpeers. Myconet is based on models of fungal growth, and uses the concept of *hyphae* (the root-like structures of some species of fungi) to guide the self-organization of the network topology.

Biologically-inspired metaphors and models have been lately the subject of active research [3], as they hold promise to enable properties, such as resilience, emergent adaptation, and self-organization, that are desirable for large-scale, distributed systems. In Myconet, the robustness and sophistication of

natural hyphal structures inspires interconnection strategies between peers and superpeers, the promotion of regular peers to superpeers, and the incremental aggregation of regular peers around superpeers.

This paper describes the Myconet model for overlay network construction, and reports a set of empirical results obtained from a simulation of that model developed using *PeerSim* [4]. Our results show how Myconet exhibits a set of desirable properties: it spontaneously and quickly achieves high levels of capacity utilization; at the same time, its topology converges towards a nearly optimal proportion of superpeers to regular peers; finally, the network is resilient and recovers quickly in the face of catastrophic network disruption. The Myconet overlay is designed to dynamically maintain a configurable number links between superpeers to facilitate network tasks. Myconet makes a design choice to slightly underutilize a subset of superpeers for the purposes of maintaining greater flexibility and robustness against failure.

We discuss all of these results, and compare them to other works in recent literature that also aim at the construction of robust and efficient superpeer-based overlay topologies. We also report on the engineering lessons that we have learned while developing Myconet and the insights we have gained on the challenges of designing and implementing a biologically-inspired, self-organized system. We conclude by outlining future directions for this line of research and for Myconet.

II. BACKGROUND

Superpeer approaches to peer-to-peer overlays attempt to exploit the heterogeneous capacities of the participating peers to improve performance for the entire network [6], [7]. Superpeers may take on service roles for other peers, such as indexing files, routing data, or forwarding searches. Connections between superpeers serve to reduce the network diameter and make these services more efficient.

Designing superpeer-based overlay topologies on large-scale P2P networks is difficult, as no global view of the network exists. Further, such networks can be extremely dynamic as peers frequently join and leave (whether by failure or deliberate disconnection). The number of peers needed for a particular network is unlikely to be known in advance, so decentralized protocols are needed.

Fungi, the inspiration for our overlay model, are much more than mushrooms and yeast. Many fungi reproduce primarily by vegetative growth; that is, by extending filamentous strands (called *hyphae*) through the soil (or other growth medium) as depicted in Figure 1. These hyphae search for biomass to assimilate, collecting nutrients and water. The hyphae concentrate the biomass and also use it to fuel hyphal growth. The system of hyphae is referred to as a *mycelium*. The mycelium constantly adapts to changing environmental conditions by routing nutrients and biomass to areas of need.

Mycelia have a number of self-organizing properties that have led researchers in the pervasive computing community to propose the use of fungal networks as the inspiration for a new paradigm for decentralized systems [1], [2]. Mycelia grow using decentralized, local interactions from which organization emerges, they are able to adjust to changing conditions or damage by dynamically altering hyphal structures. Their many interconnections make the overall network both efficient at transporting nutrients and robust to environmental stresses and local failures.

These properties match particularly well with the desiderata for unstructured peer-to-peer systems, however the fungal metaphor has not yet been applied in this context. Our approach to superpeer overlay construction, using a protocol we call Myconet, is modeled after hyphal growth patterns.

III. APPROACH

Myconet uses a relatively simple collection of rules and parameters to regulate the growth and maintenance of the overlay. These rules define three different states for superpeers (with we refer to as “hyphae” or “hyphal peers”) between which they dynamically transition in response to changing conditions. This adaptive, multilevel topology is one of the primary contributions of Myconet. The state transitions are illustrated in Figure 2 and will be explained in detail in the rest of this section.

Our rules are loosely inspired by a fungal metaphor, in which regular peers are regarded as “biomass” and the superpeer overlay as hyphae criss-crossing the network. Each of the superpeers maintains a number of links to other superpeers and adjusts the overlay topology according to rules that balance growth or contraction of the overlay with the concentration

of biomass around the hyphae best able to make use of it. We refer to the “neighborhood” of a hyphal superpeer as the collection of biomass peers that are connected to it and the hyphae to which it has hyphal links.

A peer is characterized by a capacity value representing the number of biomass peers that it is able to support. This is a simplification for purposes of simulation, and implementation outside of the simulator will require a more complex evaluation of capacities, in terms of bandwidth, demands of the individual peers, and the utility of a biomass peer being associated with a particular superpeer. In our model, superpeers may be temporarily over capacity; in such a case our protocol will shift peers away from overloaded peers (possibly promoting new superpeers, as discussed below).

Myconet superpeers can be in different states, depending on their level of biomass concentration (the number of associated biomass peers) and number of hyphal links; the conditions under which a peer transitions between states are shown in Figure 2.

Besides concentrating biomass peers as close as possible to their capacity, superpeers also work towards maintaining a configurable number of connections (C_n) to other superpeers, which we call “hyphal links”, and which serve a number of purposes. First of all, it is well-known that a strongly interconnected network of superpeers improves the efficiency of certain operations performed over the network, such as search [8], [9], [10]). Furthermore, it also strengthens the fabric of the overlay in the event of the failure of a superpeer [11]. In addition to those general benefits, in Myconet hyphal links also provide the means for superpeers to exchange biomass peers with one another, driving the overlay towards higher level of utilization of the superpeers’ capacity. Hyphal links are represented in our system as direct pointers in between superpeers; in the experiments presented in this paper each superpeer maintains 5 of those pointers, that is, $C_n = 5$.

When hyphae have achieved C_n inter-hyphal connections, they will stop creating additional hyphal links unless they again fall below that level. (This may happen because of protocol operations or because a neighbor hypha leaves the network due to churn or failure.) Additionally, there is a small probability p (discussed in the description of immobile peers, below) that an immobile hypha may create a new hyphal link even it already has C_n links, to prevent the network from becoming static or stuck in a local optimum.

Bootstrapping of a Myconet network requires a way for peers to discover other live biomass and hyphal peers. We follow the approach used by (among others) SG-1 [12], and use a simple, randomly-connected, lower-level overlay that maintains a list of known live peers and their states. For purposes of the simulation we used the gossip-based, Simplified Newscast protocol provided by the PeerSim framework to communicate the Myconet protocol status of peers (*biomass*, *extending*, *branching*, or *immobile*). Newscast chooses a random peer from its internal list each round and exchanges peer lists and state information. Myconet leverages this lower-level protocol when it needs to select an arbitrary, non-neighbor hypha. We hope to reduce this reliance on Newscast in our future work.

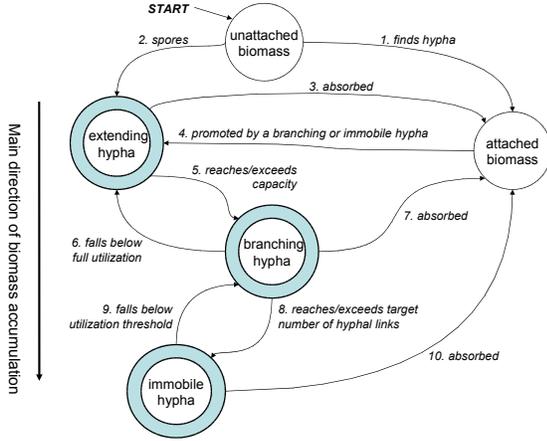


Fig. 2. Myconet protocol state transitions. The transitions are explained by number in Section III. Peers adaptively transition between regular peers (biomass) and three different types of hyphal superpeers depending on changing conditions.

Biomass Peers

In Myconet, each peer executes specific rules depending on its state. When protocol execution begins, all peers are in the biomass state and there are no hyphal peers. A biomass peer b has the goal of becoming attached to one and only one other peer, which must be a hypha. Therefore, if b is disconnected, it will first try to find hypha with available capacity and connect to it by querying the lower-level Newscast protocol (Transition 1 in Figure 2). If no such hypha can be found, the biomass peer will “spore” (Transition 2 in Figure 2), becoming a stand-alone, extending hypha, which will then be able to attract other neighboring biomass peers, as well as to connect to other hyphae. If the biomass peer b is already attached to a hyphal peer, it takes no other action.

Hyphal Peers

Hyphal peers may be in one of three states: *extending*, *branching*, or *immobile*. In accordance with our fungal metaphor, extending hyphae are those that continuously explore the network, foraging for new biomass. (It must be noticed that, as we just described, it is actually the biomass peers that seek out an extending hypha and attach themselves to it, but the result is equivalent.) Maintaining extending hyphal peers allows the network to quickly adapt to changing conditions.

Branching peers are responsible for growing new extending hyphae and building interconnecting links between hyphae. Immobile peers are those that are at or near full utilization and have achieved the ideal number of hyphal links (C_n); they pull biomass from branching and extending peers to keep themselves at full capacity and regrow hyphal connection if they are lost due to peers leaving the network.

As the protocol continues to execute, the “highest-quality” peers (in our simulation model, those with the largest capacity) will converge towards, and ultimately reach, the immobile state. The protocol guides the overlay towards that goal by

using the following general rule: A hyphal peer h always looks for the highest-capacity biomass peer b_{max} that is either a direct client h or a client of one of h ’s neighbor hypha. If b_{max} has a capacity higher than h itself, then b_{max} and h swap roles: b_{max} becomes a hypha, all of h ’s biomass peers and links to other hyphae are transferred to b_{max} , and h reverts to being a biomass peer attached to b_{max} . This rule progressively promotes the highest-capacity peers to superpeer status.

Hyphal peers then follow further rules depending on their current state: extending, branching, or immobile.

Extending Peers: In the case in which an extending hypha h_e is not connected to any branching or immobile hypha, h_e will attempt to form a hyphal link to a random peer of one of those types (random peers are selected by querying the lower-level Newscast protocol). These links ensure that the initial, stand-alone clusters that form around extending hyphae in the bootstrapping stage of the protocol will gradually converge towards a single, larger connected network.

If no immobile or branching peers can be found, h_e will then try to connect to another extending hypha h_{e2} . Whenever two extending peers have become neighbors (whether because of this rule or as a result of other protocol dynamics discussed below), the larger will attempt to “absorb” the smaller (Transition 3 in Figure 2), if it has sufficient unutilized capacity. When h_e absorbs h_{e2} , all of h_{e2} ’s biomass peers are transferred to h_e , all of h_{e2} ’s hyphal links are transferred to h_e , and h_{e2} reverts to being a biomass peer attached to h_e . This rule is the main means by which the number of superpeers in the network is contracted, once again favoring the hyphal peers with higher capacity.

Finally, if the extending hypha has reached or exceeded its biomass capacity, it will become a branching hypha (Transition 5 in Figure 2) and the excess capacity will be handled by the rules for branching peers.

Branching Peers: Branching hyphae are peers that are at or near their ideal biomass capacity but have yet to reach their ideal number of hyphal connections C_n (see above). Their purpose is multifold: they help to regulate the number of extending peers in the network, expanding or contracting their number to adequately handle any biomass peers; they act as a conduit to move biomass between extending and immobile hyphae; and they seek to construct new inter-hyphal connections, which are important for the robustness of the overlay.

Branching peers adjust the number of extending peers by attempting to maintain one and only one connected extending hypha. If a branching peer h_b does not have a link to an extending neighbor (as is the case after it is first promoted to branching status), it will choose its largest biomass peer and promote it to the extending hyphal state (Transition 4 in Figure 2). Branching peers also help contract the overall number of superpeers in the network: if the branching hypha h_b is connected to two (or more) extending peers, it will pick two of them, h_{e1} and h_{e2} . h_b then connects the extending peers h_{e1} and h_{e2} to each other. h_b maintains its connection to the higher-capacity peer h_{e1} , but severs its connection to h_{e2} . The rules for extending hyphae will then be triggered, leading to the collapse of the two peers in the next round (Transition 3

in Figure 2).

Branching peers also work to pull biomass from extending peers; similarly, immobile peers pull biomass preferentially from branching peers. In this way, biomass peers tend to gradually aggregate around those high-capacity superpeers that have established themselves in the network over a period of time. Notice that since promotion to branching hyphal status occurs only after an extending hypha reaches full utilization, pulling of biomass occurs only when a branching peer has lost some of its own biomass. That may occur because the biomass has been pulled from them by connected immobile peers or because some of client biomass peers have left the network.

Whenever a branching peer has fallen below full utilization, it tries to obtain new biomass from its neighbors. If branching peer h_b is of larger capacity than a neighbor h_n and h_b 's unused capacity is greater than the number of biomass peers attached to h_n , h_b will absorb h_n outright (Transition 7 in Figure 2). All of h_n 's biomass peers attach to h_b , all of h_n 's hyphal links are transferred to h_b , and h_n becomes a biomass peer of h_b .

If h_b is still not at full capacity after attempting to absorb neighboring hyphae, it then checks whether any of its connected extending peers have biomass peers; if they do, it will transfer biomass peers to itself until it is at capacity, or until no more biomass is available. In the latter case, if h_b is still under-utilized, it will drop back down to extending status (Transition 6 in Figure 2).

If, as a result of the processes described above, a branching peer exceeds its capacity (that is, if it has too many connected biomass peers), it will push the excess biomass down to a neighboring extending peer.

A branching peer also seeks to reinforce the fabric of the overlay by growing links to randomly-selected, existing hyphae. Random hyphae are chosen by querying the peer list maintained by the lower-level Newscast protocol. These cross-connections add resilience to the network in case of the failure of one or more superpeers. Whenever a branching peer has accumulated enough hyphal links such that it reaches or exceeds the set parameter C_n , it promotes itself to immobile status (Transition 8 in Figure 2).

Immobile Peers: Immobile peers have achieved their ideal number of hyphal connections and are have connected biomass peers sufficient to fully-utilize their capacity. If an immobile peer falls under its biomass capacity, it will attempt to absorb biomass from a connected branching or extending hypha. In this way, immobile peers (which have proven to be of high capacity and stable) continue to maintain high levels of utilization, and client biomass peers are moved to these peers in preference to other superpeers.

If two extending hyphae are connected to an immobile peer h_i , h_i will create a direct connection between them, so that the lower-capacity peer is absorbed by the higher-capacity one. (This rule is identical to the rule for branching peers.) Similarly, if h_i has both a branching peer and an extending peer connected to it, it transfers the extending peer to the branching peer.

Next, if an immobile peer h_i is of larger capacity than a neighbor h_s and h_i 's unused capacity is greater than the

number of biomass peers attached to h_s , h_i will absorb h_s outright (Transition 10 in Figure 2). All of h_s 's biomass peers become attached to h_i , all of h_s 's hyphal links are transferred to h_i , and h_s becomes a biomass peer of h_i .

h_i then checks to see if it is still under its biomass capacity. If it is, it will attempt to absorb biomass from neighbor branching or extending peers. It transfers biomass from such peers until it has reached full utilization or no further biomass peers are available. If, after this process, the immobile peer h_i is over its biomass capacity, it will transfer the excess to a neighbor with available capacity (immobile, branching, or extending, in order of preference). If h_i does not have any under-capacity neighbors, it will promote its largest-capacity biomass peer to extending status, and transfer the excess capacity to it (Transition 4 in Figure 2).

If h_i drops below the number of hyphal links C_n that Myconet is trying to maintain, whether because any of the rules above or because of peers leaving the network, it will randomly form new links to another existing hyphal peer. (As with branching peers, these are selected from the list maintained by Newscast.) Also, to prevent the topology from becoming stagnant, there is a small probability p that h_i may form another hyphal connection even if it already has C_n neighbor hyphae. (In our simulation, we used $p = 0.05$.)

If, because of any of the above rules (or because other peers have chosen to connect to it), h_i has more than C_n connections to other hyphae, it will randomly drop extra hyphal connections until it gets back at the C_n level.

Finally, if h_i falls below a certain utilization threshold u_i and is unable to regain the lost biomass, (in our experiments we set the threshold equal to 80% of h_i 's capacity), it will demote itself to become a branching hypha (Transition 9 in Figure 2).

IV. EVALUATION

We have developed the Myconet protocol on top of the Java-based PeerSim simulator [4]. We have also leveraged PeerSim for collecting data pertaining to our evaluation, focusing on the following major aspects: 1) how efficiently the network self-configures into an overlay with a stable number of superpeers; 2) how quickly the overlay reaches certain target levels of utilization of the capacity provided by the set of all peers, that is, how quickly the protocol comes up with a superpeer topology that covers some given percentages of all other peers' capacity; 3) how well the overlay approximates the minimum theoretical number of superpeers needed within the network to achieve 100% utilization; 4) how robust the overlay is, that is, how well it resists and recovers from major disruptions of the network, as in the case of a catastrophic scenario in which a large percentage of the superpeers are eliminated.

The above-mentioned metrics are standard for the evaluation of overlay topologies for unstructured P2P networks. Whenever possible, we not only present our results, but also compare them to those reported in other recent research works with concerns are similar to ours, namely the SG-1 system described by Montresor [6] and the ERASP system described by Liu *et al.* [7]. We have chosen those as our benchmarks

since they have reported significant results using or more of the metrics above and presented them in a consistent way, which makes easy for us to perform similar experiments within our approach. Whenever feasible, we have configured our environment to match that used by these previous works to facilitate comparison.

It is worth mentioning that while other protocols have methods of maintaining connections between superpeers in the overlay, convergence and disruption to these interconnections is not measured in their experiments [6], [7]. In our experiments, the number of immobile peers represents those peers that have established $C_n = 5$ hyphal connections, and the number of branching and extending peers represent peers with fewer connections at any given round of our simulations. The level of interconnection is important because it improves the overall efficiency of the tasks the overlay will be asked to perform for the network (i.e. searching, etc).

The distribution of peer capacities may have significant effects on the behavior of a protocol. We follow SG-1 and ERASP by testing Myconet under two differing capacity distributions. While Liu *et al.* do not specify the exact distributions used in their tests of ERASP, Montresor reports results for SG-1 using (1) a uniform random distribution in the range $[1, c_{max}]$ and (2) a power-law distribution such that the probability of a peer n having a given capacity x are $P[c_n = x] = x^{-\alpha}$, with $1 \leq x \leq c_{max}$, $c_{max} = 500$, $\alpha = 2$, and a network size of 100,000 peers. [6]. We have used these parameters in our experiments except where noted. The results shown in the Figures in the remainder of this Section represent average values, over twenty-five experimental runs.

We did note that, though the number of superpeers selected by Myconet under the two distributions were slightly different (reflecting the differing capacities available in the network), the shape of the curves for both were extremely similar. We have included graphs for both distributions in Figures 3 and 4, but have chosen to omit the graphs for the uniform distribution for other experiments to avoid redundancy.

With respect to our first concern, that is, the ability of Myconet to converge to stable and efficient configuration, we measure the time (i.e., the number of rounds in our simulation) that Myconet takes to converge towards a stable number of superpeers. As discussed in the description of our approach, the bootstrapping configuration of Myconet is such that at time 0 all peers can potentially act as superpeers, as they either grow a new extending hypha or find an hypha to connect to.

In Figures 3 and 4, we show the dynamics of the total number of superpeers in a network of 10,000 peers. We also show the number of hyphae in each of the three Myconet protocol stages *extending*, *branching*, and *immobile*. As can be seen, the number of extending peers initially spikes very high as the overlay bootstraps. This is followed by a similar, smaller spike in the number of branching peers as the higher-capacity peers are promoted to branching status and absorb their smaller neighbors. Over the next rounds, the largest peers are promoted to become immobile hyphae and the total number of superpeers in the network drops as the overlay consolidates biomass peers around those immobile peers.

As no new peers are being introduced into the network in

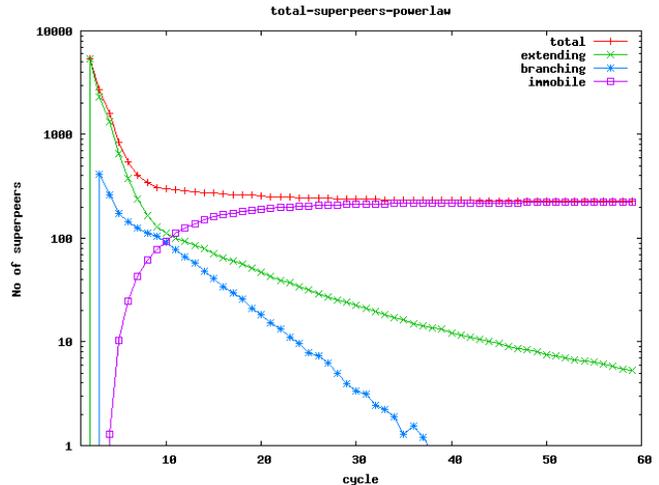


Fig. 3. Superpeers in network at each round (power-law). The network quickly converges to around 200 interconnected superpeers.

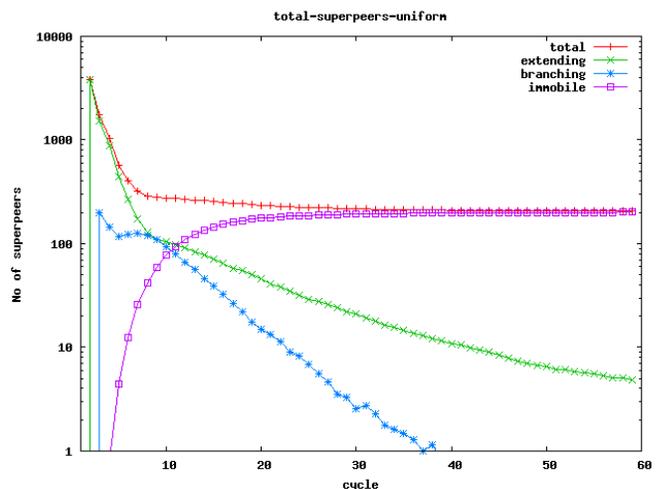


Fig. 4. Superpeers in network at each round (uniform). The network quickly converges to around 200 interconnected superpeers.

these experiments, the number of branching peers drops to very low levels and the extending hyphae are slowly promoted or absorbed. Peers attached to the remaining extending (and lower-capacity) superpeers are pulled in to become clients of larger and stabler immobile hyphae. Some extending hyphae remain active in the network; this is part of Myconet's means of continuously looking for further biomass and adjusting in response to churn and failures.

Figure 5 shows the utilization trends of the superpeers for the power-law distribution (results for the uniform distribution were similar). At the beginning of the simulation, the superpeers are designated largely at random and most will be low capacity. The network quickly returns the lower-capacity peers to biomass status and the higher-capacity peers become superpeers (these superpeers will be initially underutilized, as the dip in Figure 5 shows). Then, we can see how the overlay quickly escalates branching and immobile superpeers to full utilization. Looking at the set of all superpeers, that is,

taking extending hyphae into account, the overall utilization is somewhat lower, though it quickly climbs into the 80% range and then rises more slowly, as extending superpeers are promoted or absorbed. This occurs by design, as some underutilized extending hyphae serve as connection points for any loose biomass in the network. When compared with SG-1 and ERASP, the rate of convergence to 80% is similar. Myconet prioritizes quickly reaching capacity of the highest capacity peers while maintaining extending hyphal peers to accommodate growth, so its convergence to 95% utilization is slower, but still occurs by round 30-35. Myconet’s slightly lower utilization figures are due primarily to the extending hyphae maintained. As no new peers join the network, the remaining extending nodes are slowly reduced, increasing the overall utilization figures.

To compare our approach to a theoretical optimum, we calculate the minimum number of superpeers required for a given network and capacity distribution by repeatedly selecting the largest-capacity peers until the total capacity of the selected peers is equal to or greater than the total number of peers in the network. Figure 6 depicts the number of superpeers selected by Myconet after 60 simulation rounds for network sizes from 10^3 to 10^6 , and graphs this against the calculated optimal number of peers (for clarity, only results for the power-law capacity distribution are shown as the curves for the uniform distribution are very similar).

For a network with 1,000 peers (with $c_{max} = 500$ and other parameters as discussed above), on average only 3 superpeers are required to cover all peers, and Myconet converges to that number. Note that the last extending peer will never be promoted as there are insufficient peers in the network to saturate its capacity. Myconet scales smoothly up to 10^6 peers, the largest number with which we tested our simulation. At this network size, the Myconet overlay averaged only 51 peers over the optimum after 60 rounds, which is quite reasonable for a network of one million peers.

To test the self-healing ability of Myconet, we simulated failure conditions by removing large percentages (30% and 50%) of the superpeers at round 30 of the simulation. (An equal number of new, unattached peers are introduced at the same time to keep the total number of peers in the simulation at 10^5). As can be seen in Figures 8, 9, 11, and 12, following a catastrophic event the number of superpeers spikes while overall utilization drops by roughly the percentage of superpeers killed. The network then quickly reconverges to pre-event levels. (These figures depict the behavior for a network with a power-law capacity distribution; the behavior for the uniform distribution is similar.)

We also tested the effects of specifically targeting superpeers in the *immobile* protocol state. Figures 10 and 13 show the behavior of Myconet when 80% of immobile hyphae are removed at round 30. These peers are replaced by peers whose capacities are randomly drawn from the power law distribution, greatly reducing the overall capacity of the network. While the spikes are somewhat higher in these tests, the network again quickly converges to the new optimum which is still near the pre-event number of superpeers and utilization percentages. The reconverged number of superpeers

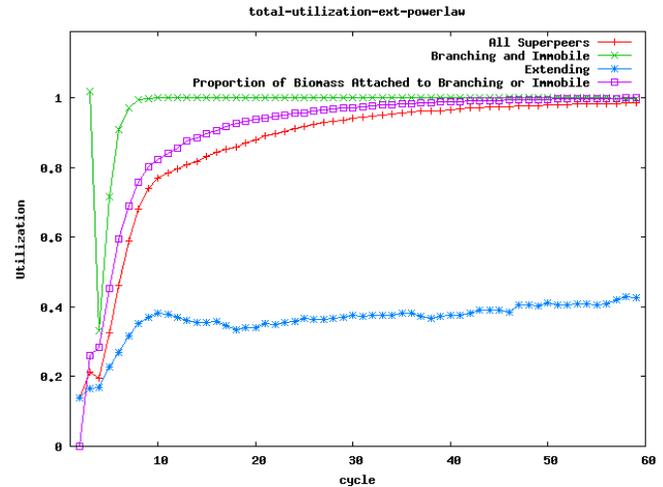


Fig. 5. Superpeer utilization at each round. Branching and immobile peers quickly reach full utilization.

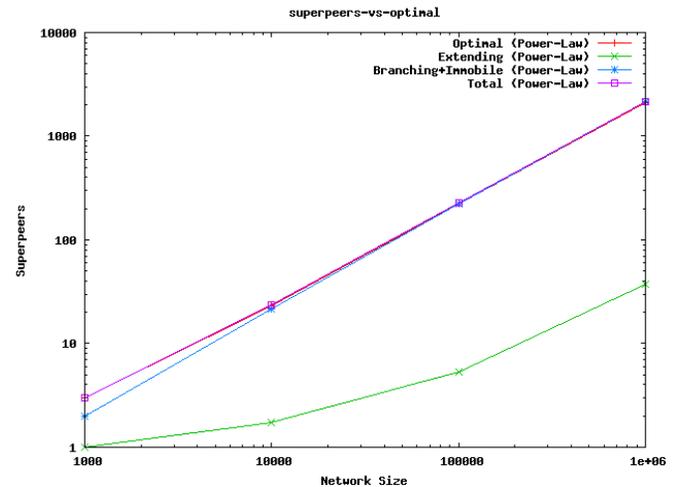


Fig. 6. Superpeers in network after 60 rounds vs. theoretical optimum (power-law). The total number of superpeers closely tracks the optimum. Results for a uniform capacity network were similar. (Detailed numbers are shown in Figure 7.)

Nodes	Extend	Branching & Immobile	Total	Optimum
10^3	1.00	2.00	3.00	3.00
10^4	1.72	21.64	23.36	22.84
10^5	5.38	221.62	227.00	221.75
10^6	37.44	2134.56	2172.00	2120.88

Fig. 7. Average superpeers in network after 60 rounds vs. theoretical optimum (power-law). (This table supplements the graph in Figure 6.)

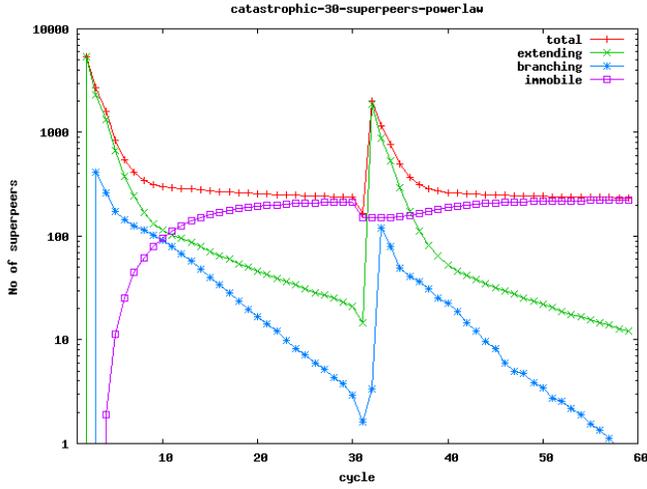


Fig. 8. Number of superpeers with removal of 30% of all superpeers at round 30 (power-law)

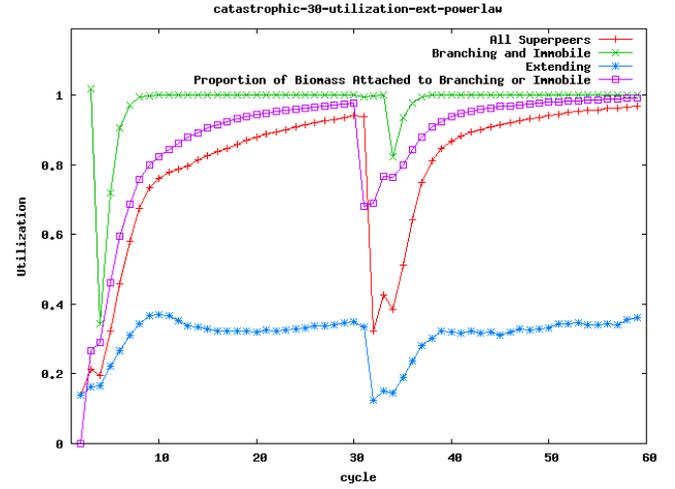


Fig. 11. Utilization with removal of 30% of all superpeers at round 30 (power-law)

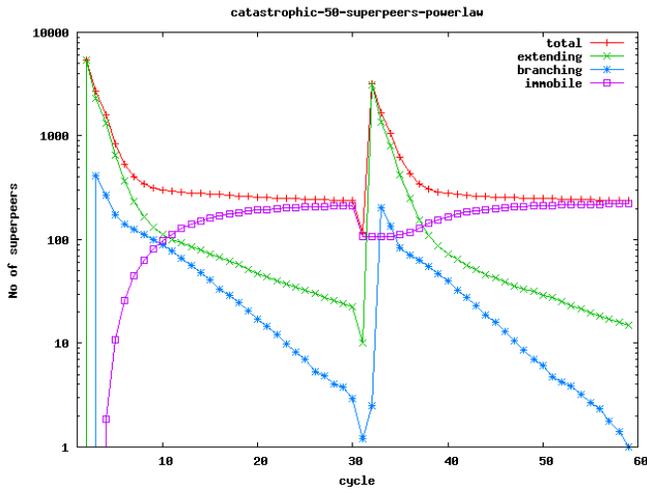


Fig. 9. Number of superpeers with removal of 50% of all superpeers at round 30 (power-law)

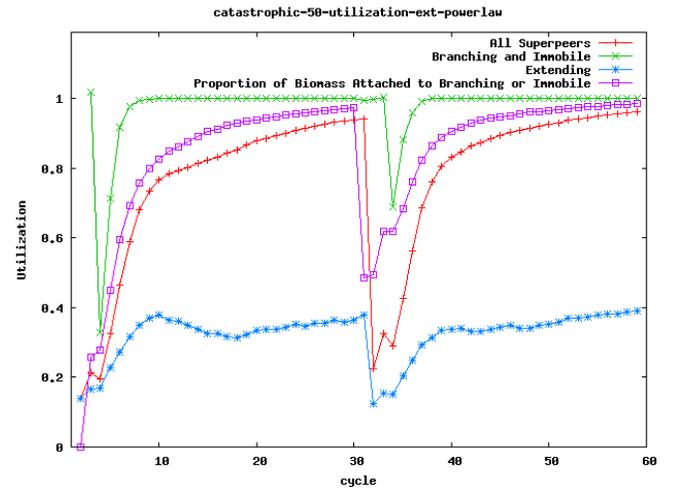


Fig. 12. Utilization with removal of 50% of all superpeers at round 30 (power-law)

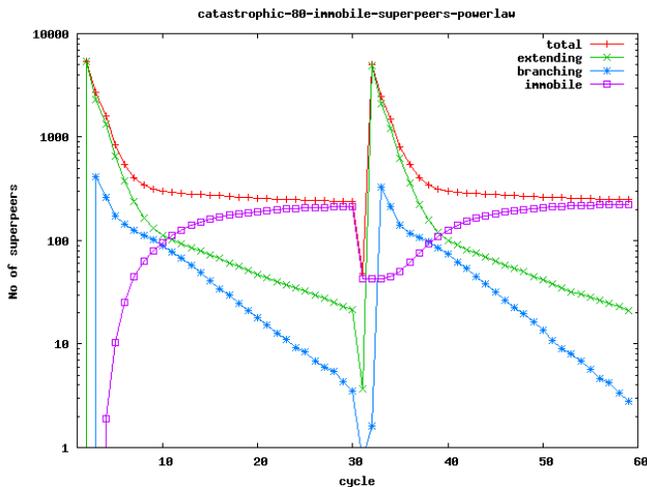


Fig. 10. Number of superpeers with removal of 80% of immobile peers at round 30 (power-law)

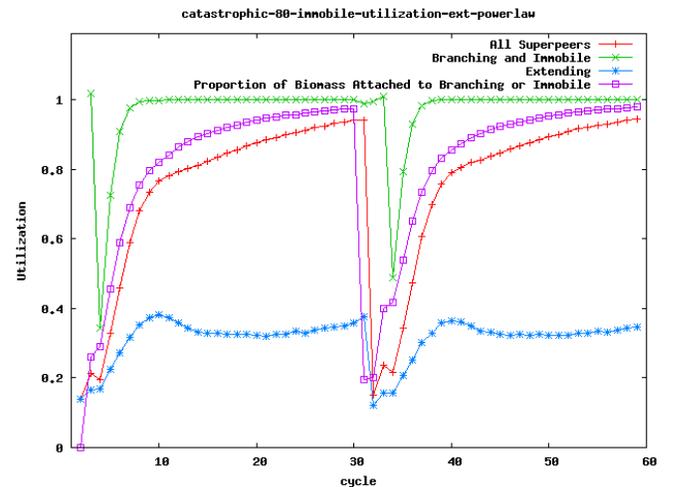


Fig. 13. Utilization with removal of 80% of immobile peers at round 30 (power-law)

is naturally slightly higher, as the killed immobile peers were the highest-capacity peers in the network, and more of the remaining, smaller peers must be promoted to make up the difference.

Our results for these self-healing experiments are particularly interesting when compared to the behavior of our benchmarks, SG-1 and ERASP. The SG-1 paper documents the behavior of the protocol when 50% of the superpeers are killed. For uniform distributions, SG-1 returns to approximately pre-removal numbers of superpeer in ten to fifteen rounds, while the power-law distribution reconverges to a higher number. (Whereas the superpeers count appears to be between 100 and 200 superpeers prior to the removal, the network seem to resettle at a number somewhat over 400 following the event.)

The ERASP paper documents the behavior of their protocol following the removal of 30% of the superpeers for an unspecified capacity distribution. Scenarios are examined where client peers have either one or two parent superpeers. The introduction of a second parent superpeer greatly decreases the spike in the number of superpeers immediately following the removal, but, in both cases, the network reconverges to a measurably higher number of superpeers, although it is difficult to determine the quantitative difference from the data provided in the paper.

In contrast, Myconet converges to pre-event levels in all of our experiments. Our hypothesis about this phenomenon, which is not explained in the literature we have examined, is that in the case of those other systems, the superpeers that survive the catastrophic events might not be significantly re-engaged during the recovery of the overall network, which may cause situations in which the topology after the catastrophic event becomes less than optimal. In Myconet, instead, all types of superpeers, primarily extending hyphae, but also branching and immobile hyphae, participate actively in the recovery phase: for example, according to the rules of the protocol, immobile nodes can be demoted, at least temporarily, or they be could absorbed by nodes with higher capacity.

V. LESSONS LEARNED

The process of developing and refining the Myconet protocol taught us a number of lessons regarding the idiosyncratic challenges and surprises one can encounter when striving to design emergent behavior in biologically-inspired systems.

Many biologically-inspired systems operate according to simple rules with a concept of locality; that is, elements typically interact only with other elements that are in their immediate “neighborhood”. It is well-known that the emergent consequences of a particular local rule at the global level are not always obvious. We experienced this during the development of Myconet, and observed that the behavior of the system could be quite sensitive to small changes in rules or parameters. For example, for the overall behavior and performance of the Myconet protocol, one of the most significant steps was the addition of the rather simple rule whereby the larger of two connected, extending peers demotes the smaller one and absorbs all of its biomass peers. This

rule proved to be instrumental for the rapid consolidation of superpeers, the efficient convergence towards a stable number of superpeers, as well as high levels of capacity utilization.

Such sensitivity to small, local changes is coupled with the current lack of a systematic discipline and best practices for testing emergent systems. We felt our trial-and-error process was truly a “Galilean” loop of observation, hypothesis formulation, and experimentation. That highlights the importance of inserting parameterization and modularity in the system, allowing the exploration of a wide range of possible settings, often along multiple design dimensions. This is also where the value of a powerful and flexible simulation tool in the designer’s arsenal becomes most evident.

A further challenge was observing and regulating the behavior of the system at very large scales. While we were able to view the complete behavior of the network and all node interactions at smaller network sizes, direct observation was not feasible for networks of 10^4 nodes or larger. This was crucial because, while in the process of experimenting with and tuning the protocol rules, the emergent behavior of the system could vary significantly at different scales. The most useful tools for working with very large scales were extensive logging and instrumentation, which allowed us to gather detailed statistics, and validate the behavior of the system at both local and global levels.

Besides being prepared to overcome those challenges, a designer of bio-inspired systems should also be ready to recognize, accept and embrace the “surprises” that the emergent nature of those systems will reveal. In the case of Myconet, although we had early on a clear set of objectives, toward which we worked in a goal-oriented fashion, in the end we came to realize that certain design characteristics, which we now regard as significant contributions of the protocol, have revealed themselves incrementally during the design process.

A major example is the role of extending hyphae. By continuously growing some new extending hyphae, Myconet trades some speed in network convergence for an approach where continual adaptation and adjustment are built-in. The consequence is that we introduce a slight measure of non-determinism in the network configuration: in Myconet, extending nodes are always “on the prowl” for new biomass, and through their relationship with extending nodes, all other types of hyphal peers are at all times open to self-optimization opportunities. As a result, the robustness and ability of the system to self-heal are enhanced. Our protocol converges to a number of superpeers that is very close to optimal, and returns right to this level following a catastrophic disruption to the network.

Another example is how our various superpeer types effectively codify in the protocol what in other systems is only a heuristic: that is, the fact that “veteran” nodes often make good superpeers. In Myconet, we observed that a node typically tends to be promoted to immobile status when it has high capacity *and* has participated in the network for enough time. That is not a heuristic, but an intrinsic aspect of how the rules handle nodes, by promoting them through the different stages of the protocol.

VI. RELATED WORK

Approaches for peer-to-peer overlay networks are surveyed in [13], which discusses both structured and unstructured models. Structured approaches include Chord, Pastry, Tapestry, and Kademia. As we are focusing on unstructured overlays, we will not delve deeply into this area.

Unstructured approaches include such well-known file-sharing networks as BitTorrent [14], eDonkey [15], KaZaa [16], and Gnutella [17]. Many of these networks employ superpeers to facilitate and/or speed up a variety of critical functionality, such as search of information and routing of data.

Since it is often unfeasible to pre-determine the topology and the quantity of superpeers that can adequately service a P2P network, because of its dynamism and the lack of global structural information, techniques for self-organization are commonly applied to this problem.

First of all, several approaches exist that tackle a generalization of this problem, that is, “topology management”. Those approaches aim at manipulating generic network topologies (not necessarily, or exclusively, P2P networks) in a self-adaptive way. For instance, Zweig and Zimmerman present an approach that is able to autonomically respond to node failures in a network by switching between a scale-free and a normal or Poisson topology [18]. In [19], local information on neighborhood structure is used to drive the network incrementally closer to a given set of requirements about its overall topology by perturbing the neighborhoods in ways that are consistent with the goals set by those requirements.

Another work that aims at P2P topology management is T-MAN [12]. In T-MAN, a gossip protocol is used to diffuse information that gradually evolves the topology of the P2P network towards a preferred one, defined by means of a ranking function predicated on inter-node distance. T-MAN shows how a structured overlay such as a ring or torus can be efficiently imposed on an unstructured network. SwapLinks [20] also examines overlay construction, but focuses on building random (single-level, non-superpeer) graphs and enabling random node selection.

Although generic topology management approaches could conceivably be used to address the specific issue of creating of a superpeer-based overlay in P2P networks, more interesting for our purposes are those self-organizing approaches that focus upon building a superpeer topology with nearly-optimal utilization of the collective peers’ capacity and network robustness, such as the aforementioned SG-1 [6] and ERASP [7].

SG-1 and ERASP also build and maintain two-level (superpeer) hierarchies; however Myconet’s superpeers are further differentiated into extending, branching, and immobile peers. Myconet’s overlay differs from both SG-1 and ERASP by introducing dynamic rules for adjusting the number of interconnections between superpeers to reinforce the overlay and enhance self-healing. ERASP specifically examines the effects of allowing leaf peers to connect to multiple superpeers. Like Myconet, ERASP uses a concept of the time that a node has participated in the network as a factor in determining its reliability.

Other published superpeer overlay approaches include [8] which proposes “clusters” of superpeers, where superpeers in each cluster act as parents for the same group of clients and uses a simple splitting method when a superpeer exceeds its capacity. [21] and [22] focus on a separate problem, that of locality in superpeer selection; these approaches use Yao-Graphs to organize superpeer coverage based on the underlying network topology.

Fungi-inspired mechanisms have been previously examined for application to communication problems. At the PerAda Summer School workshop in Rimini, Paechter *et al.* recently proposed a novel biologically-inspired model based on the behavior of fungal colonies. [1] Their initial work applies a relatively low-level fungal model to routing in a small telecommunications network [2]. Our approach is more loosely inspired by fungal models and applies the metaphor at a higher level.

VII. CONCLUSIONS

Myconet demonstrates that the fungal metaphor is powerful when applied to peer-to-peer overlays and that it can be used to create a self-organizing peer-to-peer overlay that provides advantages over other approaches in terms of its robustness to failure. We therefore urge other researchers not to overlook Kingdom Fungi when building biologically inspired systems.

The Myconet protocol effectively constructs and maintains a **strongly interconnected**, decentralized superpeer overlay that scales to at least 10^6 peers. This overlay quickly converges to an optimal number of superpeers and high levels of capacity utilization. Myconet’s greatest strength is its ability to self-heal the interconnected overlay, quickly repairing the damage from catastrophic events in which 30-50% of the peers are removed from the network. Even when 80% of the highest capacity, most stable peers are removed, Myconet is able to quickly reconfigure its overlay to reflect the new capacity of the network.

Unlike other unstructured approaches, Myconet dynamically adjusts the interconnections between superpeers, increasing efficiency and resilience to the loss of peers.

Myconet’s results are promising enough to warrant further investigation. In particular, we are curious how varying the level of interconnection (C_n) effects the efficiency and robustness of the network, the communications efficiency of the Myconet protocol, and exploring Myconet’s response to churn in the network. Preliminary results concerning churn are promising. We also would like to move away from the round-based simulation to an actual protocol implementation, and investigate how the Myconet topology affects the performance of peer-to-peer protocols run on top of it.

The computational biology literature has developed a number of mathematical models of fungal growth. (See, e.g., [23], [24], [25].) While we initially considered using a more detailed model, we found that our simplified design displayed emergent properties of the sort that we were hoping to achieve. We are considering exploring the behavior of more rigorously biological models as we develop Myconet.

VIII. BIBLIOGRAPHY

REFERENCES

- [1] B. Paechter, E. Hart, and R. Falconer, "Emergent computing for pervasive adaptation," First PerAda Summer School on Pervasive Adaptation, 2008.
- [2] R. Falconer, J. Bown, E. Hart, and J. Timmis, "A new paradigm for specknets: Inspiration from fungal colonies," in *Proceedings of the First Perada Workshop on Pervasive Adaptation*, 2008.
- [3] L. De Castro, *Fundamentals of Natural Computing: Basic Concepts, Algorithms, And Applications*. Chapman & Hall/CRC, 2006.
- [4] M. Jelasity, A. Montresor, G. Jesi, and S. Voulgaris, "PeerSim: A Peer-to-Peer Simulator," URL: <http://peersim.sourceforge.net>.
- [5] myriorama, "mycelium?" <http://www.flickr.com/photos/myriorama/101120710/>, 2004.
- [6] A. Montresor, "A robust protocol for building superpeer overlay topologies," in *Proceedings of the 4th International Conference on Peer-to-Peer Computing*. Zurich, Switzerland: IEEE, Aug. 2004, CONFERENCE, pp. 202–209.
- [7] W. Liu, J. Yu, J. Song, X. Lan, and B. Cao, "ERASP: An Efficient and Robust Adaptive Superpeer Overlay Network," *Lecture notes in computer science*, vol. 4976, p. 468, 2008.
- [8] B. Yang and H. Garcia-Molina, "Designing a super-peer network," in *Proceedings of the 19th International Conference on Data Engineering (ICDE)*, 2003, p. 49.
- [9] P. Garbacki, D. Epema, and M. Van Steen, "A two-level semantic caching scheme for super-peer networks," in *Web Content Caching and Distribution, 2005. WCW 2005. 10th International Workshop on*, 2005, pp. 47–55.
- [10] Y. Yang, R. Dunlap, M. Rexroad, and B. Cooper, "Performance of full text search in structured and unstructured peer-to-peer systems," in *IEEE INFOCOM*. IEEE Press, 2006, pp. 2658–2669.
- [11] B. Mitra, F. Peruani, S. Ghose, and N. Ganguly, "Analyzing the vulnerability of superpeer networks against attack," in *Proceedings of the 14th ACM conference on Computer and communications security*. ACM New York, NY, USA, 2007, pp. 225–234.
- [12] M. Jelasity and O. Babaoglu, "T-Man: Fast gossip-based construction of large-scale overlay topologies," *University of Bologna, Department of Computer Science, UBLCS-2004-7, Bologna, Italy*.
- [13] K. Lua, J. Crowcroft, M. Pias, R. Sharma, and S. Lim, "A survey and comparison of peer-to-peer overlay network schemes," *Communications Surveys & Tutorials, IEEE*, pp. 72–93, 2004.
- [14] B. Cohen, "BitTorrent protocol specification," in *First Workshop on Economics of Peer-to-Peer Systems (P2P03)*.
- [15] O. Heckmann and A. Bock, "The eDonkey 2000 Protocol," Tech. Rep. KOM-TR-08-2002, Multimedia Communications Lab, Darmstadt University of Technology, Dec. 2002, Tech. Rep.
- [16] J. Liang, R. Kumar, and K. Ross, "Understanding KaZaA," *Manuscript, Polytechnic Univ*, 2004.
- [17] M. Ripeanu, "Peer-to-Peer Architecture Case Study: Gnutella Network," in *Proceedings of International Conference on Peer-to-peer Computing*, vol. 101. Sweden: IEEE Computer Press, 2001.
- [18] K. A. Zweig and K. Zimmermann, "Wanderer between the worlds - self-organized network stability in attack and random failure scenarios," in *Conference on Self-adaptive and Self-organizing Systems (SASO)*, 2008.
- [19] V. Matossian and M. Parashar, "Towards autonomic control of network topologies," in *1st IEEE Workshop on Modeling Autonomic Communication Environments (MACE), 2nd International Week on Management of Networks and Services (Manweek 2006)*, 2006.
- [20] V. Vishnumurthy and P. Francis, "On Heterogeneous Overlay Construction and Random Node Selection in Unstructured P2P Networks," in *Proc. INFOCOM 2006*, 2006.
- [21] M. Kleis, E. Lua, and X. Zhou, "Hierarchical Peer-to-Peer Networks using Lightweight SuperPeer Topologies," in *Proceedings of the 10th IEEE Symposium on Computers and Communications (ISCC 2005)*, 2005, pp. 143–148.
- [22] E. Lua and X. Zhou, "Network-aware SuperPeers-Peers Geometric Overlay Network," in *Computer Communications and Networks, 2007. ICCCN 2007. Proceedings of 16th International Conference on*, 2007, pp. 141–148.
- [23] A. Meškauskas, M. Fricker, and D. Moore, "Simulating colonial growth of fungi with the Neighbour-Sensing model of hyphal growth," *Mycological Research*, vol. 108, no. 11, pp. 1241–1256, 2004.
- [24] G. Boswell, H. Jacobs, K. Ritz, G. Gadd, and F. Davidson, "The Development of Fungal Networks in Complex Environments," *Bulletin of Mathematical Biology*, vol. 69, no. 2, pp. 605–634, 2007.
- [25] P. Gerlee and A. Anderson, "Stability analysis of a hybrid cellular automaton model of cell colony growth," *Physical Review E*, vol. 75, no. 5, p. 51911, 2007.