

BRISA: Combining Efficiency and Reliability in Epidemic Data Dissemination

Miguel Matos*, Valerio Schiavoni[†], Pascal Felber[†], Rui Oliveira*, Etienne Riviere[†]

*HASLab - High-Assurance Software Lab, INESC TEC & U. Minho, Portugal. Email: {miguelmatos, rco}@di.uminho.pt

[†]University of Neuchâtel, Switzerland. Email: first.last@unine.ch

Abstract—There is an increasing demand for efficient and robust systems able to cope with today's global needs for intensive data dissemination, e.g., media content or news feeds. Unfortunately, traditional approaches tend to focus on one end of the efficiency/robustness design spectrum, by either leveraging rigid structures such as trees to achieve efficient distribution, or using loosely-coupled epidemic protocols to obtain robustness.

In this paper we present BRISA, a hybrid approach combining the robustness of epidemic-based dissemination with the efficiency of tree-based structured approaches. This is achieved by having dissemination structures such as trees implicitly emerge from an underlying epidemic substrate by a judicious selection of links. These links are chosen with local knowledge only and in such a way that the completeness of data dissemination is not compromised, i.e., the resulting structure covers all nodes. Failures are treated as an integral part of the system as the dissemination structures can be promptly compensated and repaired thanks to the underlying epidemic substrate.

Besides presenting the protocol design, we conduct an extensive evaluation in a real environment, analyzing the effectiveness of the structure creation mechanism and its robustness under faults and churn. Results confirm BRISA as an efficient and robust approach to data dissemination in the large scale.

I. INTRODUCTION

We live in a digital era whose foundations rely on the production, dissemination, and consumption of data. The rate at which content is produced is ever increasing [16], putting pressure on dissemination systems able to deliver the data to its intended consumers. Examples include the distribution of digital media (e.g., music, news feeds) on the Internet [14] or software updates in a datacenter infrastructure [33].

On account of its importance, significant research has been dedicated to conceiving efficient and robust data dissemination systems [5], [7], [8], [11], [26]. Unfortunately, both design vectors, efficiency and robustness, are often addressed disjointly: either by a highly efficient structure based on trees or by a highly robust unstructured gossip-based approach.

Disseminating data using trees is attractive as, once formed, trees enable the efficient delivery of data to all participants without having to deal with duplicate transmissions [9], [37]. However, under churn and faults, the rigid structure that makes the tree efficient must be rebuilt constantly, hindering robust dissemination and continuity of service.

On the other hand, gossip-based dissemination systems rely on redundancy instead of structure to offer guarantees on the delivery of data to all participants [5], [11]. Gossip-based dissemination was initially proposed in the context of database replica synchronization in the ClearingHouse project [10]. The

transmission of several copies to random nodes enables gossip-based systems to be oblivious to faults and churn, as messages will be received through different paths. The cost is increased bandwidth and processor usage due to duplicates.

Gossip-based principles have also been used to build robust and scalable membership [15], [18], [21] and failure detection [30] services. As long as (1) the graph induced by the (partial) views offered by the membership service is connected and (2) all nodes have at least one incoming link, dissemination can trivially be achieved by flooding. As with traditional gossip-based dissemination, this approach yields a large number of duplicates due to link redundancy.

a) Contributions: In this paper we present BRISA, an efficient, robust and scalable data broadcast mechanism. BRISA leverages the resilience and scalability of a gossip substrate to build dissemination trees that are correct, i.e., cover all nodes, by construction. Trees are built in a distributed fashion with local knowledge only and with minimal overhead. BRISA has been designed in a way that upon failures or churn, trees are easily and rapidly compensated and repaired using the underlying gossip-based overlay management and dissemination protocol. BRISA is extensively evaluated by means of deployments of a prototype on a cluster and on PlanetLab.

b) Roadmap: The remaining of this paper is organized as follows. Section II describes the design of BRISA and Section III presents the experimental evaluation. Then, Section IV presents some perspectives, Section V discusses related work and finally Section VI concludes the paper.

II. BRISA

In this section, we first present the assumptions made by BRISA, and in particular the nature of the underlying peer sampling service and the guarantees it provides. Then, we introduce the key design principles of the protocol and in particular how the dissemination structures are constructed. Finally, we show how BRISA deals with dynamism, generalize the construction of dissemination structures with desirable efficiency/robustness criteria and discuss some optimizations.

A. Peer Sampling Service Layer

We assume an underlying peer sampling service (PSS) [18] that provides each node with a *view*, i.e., a set of non-faulty nodes chosen at random from the entire network. The objective of the PSS is to create views such that the overlay obtained by connecting each node to its neighbors (the nodes in its




Fig. 1. HyParView [21]: views maintenance.

view) is *connected*. That is, every node is able to transitively reach every other node in the network even under high rates of churn and failures [15], [21]. Moreover, the PSS is expected to rapidly replace failed nodes from the views and, for the sake of load balancing, provide similarly sized views and evenly distribute the number of incoming links at each node.

The update of the views can be either continuous (*proactive* peer sampling) or happen only when a node fails or a new one joins the system (*reactive* peer sampling). In the *proactive* case, nodes periodically share their views with their neighbors regardless of the actual need to replace failed entries, resulting in the view being a continuous stream of node samples from the network. A typical example of this class is Cyclon [35]. In the *reactive* method, the view is kept unchanged unless some of its entries need to be updated, i.e., for replacing a failed node or for accommodating a new node. Typical examples include Scamp [15], Araneola [28] and HyParView [21].

In this paper we rely on a reactive PSS, specifically on the HyParView protocol [21]. The motivation for this choice comes from the additional stability of reactive approaches, which simplifies the process of creating efficient and correct dissemination structures. In short, HyParView maintains two views at each node: a *passive* view and an *active* view (see Figure 1). Only the active view containing the node's neighbors is exposed to the application. The passive view is maintained in a proactive manner by periodic exchanges and sharing of the passive views with the node's neighbors. The entries in the active view are managed in a reactive manner: a neighbor in this view only changes upon failure, or for accommodating a newly joined node. An opened TCP connection is maintained with each of the nodes in the active view for communication efficiency, in particular, latency. Due to the limited size of the active view, efficient heartbeat-based fault detection can be used for all neighbors. Upon detection of a failed neighbor, a replacement node is selected from the passive view and moved to the active view. When the active view is full and a new node attempts to join, a random node is removed from the active view to accommodate the joiner. In order to avoid chain reactions due to the massive number of joins in the bootstrap process when nodes have complete active views (node A's view size is full so it removes node B, B also removes A from its view and promotes a node C from its active view, C must add B to its view and thus remove an existing one as its active view is already full, removing D and




Fig. 2. Distribution of duplicates per message for each node for 500 messages in a 512 nodes HyParView network for various active view sizes.

so on so forth), we allow the active view size to grow past the configured value by a given *expansion factor*. Nodes evictions do not result in replacements when the view size is between the target view size and this size times the expansion factor. We used an *expansion factor* of 2 throughout the evaluation. The impact on the actual view sizes is limited as shown later in the analysis of the degree distribution (Figure 7).

An important aspect of HyParView is that links with neighbors are *bidirectional*. If node A has node B in its active view, then B also has A as its neighbor. In a connected overlay, using bidirectional links allows us to ensure that messages disseminated by flooding will reach all the nodes in the system without requiring anti-entropy mechanisms [10] (where nodes periodically poll other nodes for the content they have missed). A node receiving a message *for the first time* from a neighbor simply propagates it to all its other neighbors.

Flooding is ensured to reach all nodes as long as no node in the system has an active view with only failed nodes. The larger the active views the smaller the chances for this to occur. However, the larger the view, the larger the number of relayed messages and consequently the number of duplicate receptions. As a concrete example, Figure 2 presents the cumulative distribution function (CDF) of the number of duplicates during the dissemination of 500 messages over a 512 nodes HyParView network for different view sizes. We observe that, as the size of the view grows, nodes quickly receive large amounts of duplicate messages. For instance, half of the nodes receive more than one duplicate with a view size of 4, while they receive more than 7 duplicates with a view size of 10.

BRISA develops on top of HyParView. It takes advantage of the connectivity guarantee that can tolerate up to 80% node failures [21] to bootstrap efficient dissemination structures that eliminate (or considerably reduce) the number of duplicates, while keeping the robustness offered by the underlying PSS.

B. BRISA: Rationale

The objective of BRISA is to support the efficient, robust and scalable dissemination of a stream of messages from one source to the entire network. Efficiency relates primarily to the limitation of duplicate receptions that waste bandwidth and processor resources. Additional criteria that we consider




Fig. 3. Reception of a duplicate and deactivation of one link, for a tree BRISA structure. Depending on the parent selection strategy, the deactivated link can be the previous parent or the node sending the duplicate.

for efficiency are: the reduction of the end-to-end delay (dissemination time from the source to the last receiver) and network efficiency (ratio between the delay for receiving a message through BRISA as compared to a hypothetical direct communication from the source). Robustness relates to fault tolerance: dissemination should progress despite the inactivity of some nodes (failure or disconnection) and the system should be able to rapidly detect and mask such faults. Finally, BRISA scales to very large networks, because the view size is kept small and under strict control by the PSS, preventing the load at any node to grow linearly with the system size.

The main idea behind BRISA stems from the observation that it is the *possibility* of receiving messages through multiple paths that makes gossip-based approaches robust, not necessarily the actual transmissions. Our goal is to limit or even eliminate duplicates that are received by any node in the network, but nonetheless maintain the *possibility* of receiving the messages from multiple paths.

BRISA automatically derives dissemination structures on top of the undirected HyParView overlay. The structure that emerges from the views are oriented and can either be a tree, by restricting the inbound neighbors of every node to a single node (parent), or a directed acyclic graph (DAG) by allowing multiple parents for each node. The creation of the structure is performed by local and unilateral decisions made by the nodes about the set of neighbors that should be *active* and actually relay inbound traffic and those that should be *inactive*. In the case of a tree the reception of duplicates is effectively eliminated; in a DAG, it is significantly reduced.

The resulting dissemination structure must ensure complete disseminations, i.e. that all nodes receive all messages. To that end, we must ensure that it does not contain a non-connected sub-graph that would not receive the message from the other components of the structure. This property is ensured by enforcing the absence of *cycles*. In fact avoiding cycles is the main concern when determining the set of active and inactive neighbors of a node. In the following sections we first describe how the emergence of trees is achieved in BRISA, and then generalize the approach to DAGs.

C. Emergence of the Dissemination Structure

The bootstrap of the BRISA dissemination structure starts by sending the first message of the stream to all neighbors of the source node. Nodes receiving the message for the first time

simply forward it to all the nodes in their active view. This flooding operation reaches all nodes, given the connected and bidirectional nature of the overlay provided by HyParView. Note that the bootstrap can also be done by injecting an empty message (without payload) in the system if the initial flood of an application message poses bandwidth concerns.

The links used for the initial flooding dissemination form a graph that serves as the basis for the construction of the BRISA dissemination structure. A tree structure emerges by having each node autonomously prune out *all but one* of its inbound links.

During the initial flood, nodes receive the message from a number of different neighbors. Out of these sources, each node autonomously selects one as its parent in the dissemination structure. Future messages in the stream will then be received only from the parent node. The selection is achieved by the use of a *link deactivation* mechanism and follows one of the selection strategies presented in Section II-E.

It is important to notice that deactivating a link from a node does not imply removing the corresponding entry from the HyParView active view. The overlay constructed by the PSS remains available and is used both as a provision of nodes for reparations upon failure, or as a fallback for dissemination when reparation is temporarily not possible.

Figure 3 presents the principle of the link deactivation mechanism for constructing a tree. Initially, links from nodes *X*, *Y*, and *Z* belong to node *A*'s view and are all *active*. The first reception of a message from *X* results in *A* considering *X* as its parent for the corresponding stream. A subsequent reception of a duplicate from node *Y* triggers the link deactivation mechanism. The two possibilities consist in either deactivating the link from *X* or the link from *Y*.

There are three guiding principles for deciding which link to deactivate. First, the construction process must avoid cycles. Second, it must seek to meet the target number of parents for each node (one for the tree structures, more when generalizing to DAGs). Finally, when both conditions are met, the parent selection strategy chooses the new parent based on different criteria for shaping the emerging structure (Section II-E).

D. Preventing Cycles

A mandatory condition for selecting a parent node is that it does not yield a cycle in the dissemination structure. That is, the potential parent of a node *N* does not receive the stream




Fig. 4. Avoiding creating a cycle for a tree, by checking that the node N is not in the dissemination path to the potential parent.

directly or indirectly from N itself. For a tree this implies that the parent of N must not appear in the sub-tree rooted at N .

To verify this condition each node piggybacks on the messages it relays all node identifiers in the path from the stream's source. When selecting its parent, a node N rejects those candidates whose message path to the stream's source includes N itself. This is illustrated by Figure 4, where grey nodes are not eligible as parents of node N . It is important to note that the overhead of path embedding is minimal and very attractive when compared to probabilistic inclusion structure such as Bloom filters [6]. As a matter of fact, the size of the embedded path is bounded by the tree height, which is roughly $O(\log_b(N))$ where N is the system size and b the active view size. For instance, in a system with 1×10^6 nodes with an active view size of 8, the average tree height is $\log_8(1 \times 10^6) \approx 7$. This bounds the maximum metadata size a message needs to carry which, assuming a 48 bit ip:port pair as unique identifier, is only 336 ($7 * 48$) bits. On the other hand, a Bloom filter in the same scenario with a false positive probability of 1×10^{-6} requires a filter with 28 755 176 bits [6]. Taking into account the metadata size required, the fact that path embedding is exact (false positive probability is zero) and the computational overhead associated with Bloom filters (which requires computing several hashes), path embedding presents many advantages over Bloom filters.

The detection of cycles is not only done during the initial flooding phase: a node that detects a cycle from a parent simply makes the link from that parent inactive and selects a new parent using the regular selection mechanism or the fallback to flooding as we describe later in this section.

E. Parent Selection Strategies

From N 's eligible parents (that is, those not having N in the path followed by the messages from the stream's source), BRISA selects one according to the following strategies:

1. First-come first-picked. The node sending the first received message is selected as parent, all subsequent duplicates received trigger the deactivation of the incoming link.

2. Delay-aware. This strategy considers the round-trip time between N and the candidate nodes. The one with the lowest delay is selected as parent. We leverage the periodic keep-alive messages that are exchanged by the nodes in the active views

at the HyParView level to measure round-trip times.

A simple optimization is available when building a dissemination tree using the first-come first-picked strategy: the deactivation of links can be symmetric. Supposing node A receives a message first from node B and then from node C , A will pick the link from B and send a deactivate message to C . But it can further mark its outgoing link to C as inactive as A knows it will not be eligible as parent for C , as C already received the message first.

We discuss the possibility of more elaborate parent selection strategies in Section IV.

F. Handling Dynamism

The insertion and removal of nodes in the system is handled by the underlying HyParView PSS. A new node joins by contacting an existing system node. The new node is provided with an active view with the size of that of its contact point, and is inserted in the active views of the associated nodes. BRISA automatically marks links to new nodes as active. As a result, the joining node will have all its inbound links marked as active and will receive its first message multiple times. All that remains is to select its parent(s) according to the mechanism discussed previously.

The detection of node failures is also performed at the level of the active view, by exchanging periodic keep-alive messages over the established TCP connections, or when a node fails to acknowledge the reception of a transmission (as detected by the TCP flow control for that link). When a node notices that one of its neighbors is removed from its active view (due to a failure), it first checks if that neighbor was a parent. If that is the case, the node needs to quickly find a replacement parent using one of two strategies. It first attempts a *soft repair* by trying to select as parent one of the current neighbors. A simple approach is to reactivate all its inbound links and proceed with the normal parent selection process. This can however be optimized by leveraging the keep-alive messages used for monitoring the active view at the PSS level and piggyback up-to-date information required by the parent selection procedure. If a suitable parent is found then its inbound link is directly re-activated. Note that this mechanism uses local knowledge only and requires a single message exchange being thus very fast and efficient. Further, as our evaluation shows, almost all repairs can be done using this strategy leading to minimal disruptions in the structure.

If no replacement parent exists in the active view, we resort to a *hard repair* that uses the underlying flooding approach for rebuilding part of the dissemination structure. The orphan node first re-activates all its incoming links and considers itself a fresh node by forgetting its position in the cycle detection mechanism. This allows the orphan node to take any of its active view nodes as a parent as in the initial dissemination.

To ensure the tree remains connected, it is necessary to rebuild the incoming links for a part of the structure rooted at the orphan node. To support this, the orphan node propagates a re-activation order to all its current children. These children




Fig. 5. Avoiding creating a cycle for a DAG, by checking that the level of the potential parent is less than or equal to the level of the node.

proceed with the same operation, which conducts to a part of the structure being re-bootstrapped by the structure emerging on top of flooding. We note that this re-activation order does not need to be propagated all the way down the sub-tree rooted at the orphan node, and is in practice limited to a small subset of the structure. Indeed, nodes stop re-activating their links and propagating the re-activation order as soon as they can select a suitable parent in their active view that can replace the one that sent the re-activation message (this previous parent node will then receive subsequent message and may become a child, effectively exchanging roles). The number of nodes that need to re-activate their links for temporary flooding is independent of the position of the original orphan in the tree: the depth of the re-activation solely depends only on nodes in the sub-tree finding suitable replacement parents, which is independent of the position of the original orphaned node.

Finally, nodes can compensate message loss during the parent recovery process by directly asking its new found parent to send the missing ones. Since parent recovery is quick (Section III-C) the number of messages each parent needs to buffer is small. Nonetheless more complex approaches such as [19] could still be used to ensure nodes buffer messages for long enough to allow recovery.

G. Generalized Dissemination Structures

To enhance service continuity under failures and churn, BRISA can generalize the tree structure to directed acyclic graphics (DAGs) by having each node being served by several parents. This way, a node that sees one of its parents fail can seamlessly keep receiving the flow without the need to first undergo the selection of a new parent. This is attained at the cost of handling a controlled level of message duplicates.

The establishment of a DAG basically involves making a number $p > 1$ of inbound links active in such a way that cycles are avoided. The technique to prevent cycles we used for trees is however unfeasible in the case of DAGs due to the amount of control information required to be exchanged. Indeed, a node in the n^{th} level of the tree requires a set of n node identifiers to define the path from the stream source to itself, while for a DAG with p parents per node this set at level n could reach $p^{n+1} - 1$ should all paths be non-overlapping.

Conversely, for DAGs, we use an approximate quantitative approach that does not include the nodes identifiers but just the *depth* each node is in the DAG as illustrated by Figure 5.

The source of the message stream is at depth 0 and every message carries its sender's depth encoded by a single integer. Initially, the depth of a node N is undefined and, upon reception of its first message from a node with depth $i - 1$, N places itself at depth i . From then on, N can select parents, and thus receive messages, from nodes at any depth not greater than i . Should N receive a message from a node at depth i (its current depth) then N moves to depth $i + 1$ and immediately updates its downstream children nodes accordingly.

Similar to the technique we used for trees, it is clear that any node M served directly or transitively by node N will be at a depth strictly greater than N . Therefore, M cannot become a parent of N and yield a cycle.

As mentioned, the technique is however approximate. As illustrated in Figure 5 it can yield false negatives by discarding valid potential parents. Any two paths (rooted at S) are likely to be labeled similarly with respect to depths. Since the tagging is purely quantitative, a node from one path may be dismissed as a potential parent of a node in another path despite the paths being causally unrelated. An alternative is to rely on Bloom filters to maintain the set of nodes that need to be excluded for the parent selection process. However, as for trees this a costly technique when compared to the simplicity and efficiency of depth encoding. In our experiments, nodes always obtained the desired number of parents, thus we consider this approach an attractive alternative when compared to the cost of both an exact predictor (path embedding) and of a probabilistic one (Bloom filters).

After determining the set of potential parents with the above strategy all that remains is selecting the *best* ones by using the parent selection strategies presented above.

III. EVALUATION

In this section we evaluate BRISA on two different testbeds: (1) a local cluster of 15 computers equipped each with 2.2 GHz Core 2 Duo CPU and 2 GB of RAM and connected by a 1 Gbps switched network, supporting up to 512 BRISA nodes and (2) a slice of up to 200 nodes on the global-scale PlanetLab [1] testbed. The prototype¹ leverages Splay [23], an integrated system for the development, deployment and evaluation of distributed applications.

The evaluation is focused on the aspects that drove BRISA's design: efficiency and robustness. For each experiment and unless otherwise stated, we bootstrap the system with the specified number of nodes using the first-come first-picked strategy with an expansion factor of two, randomly choose a node to be the source across all the experiment and then have it inject 500 messages at a rate of 5 per second, taking measurements as appropriate. Due to the generic nature of BRISA, the message payload is an opaque random bit string with the specified size.

We start with a preliminary study, in Section III-A, on the structural properties of the dissemination structures created

¹The prototype along with instructions to reproduce the experiments is available at <http://www.splay-project.org/splay/ipdps2012/brisa.zip>




Fig. 6. Depth distribution for 512 nodes (first-come first-picked strategy).




Fig. 7. Degree distribution for 512 nodes (first-come first-picked strategy).

by BRISA as those properties impose well-known bounds in resource usage and dissemination time. Then, in Section III-B we inspect the network properties of BRISA, namely bandwidth consumption and routing delays and analyze the results according to the structural properties. Next we evaluate the behavior of BRISA under churn in Section III-C. Finally, in Section III-D, we compare BRISA with other approaches.

A. Structural properties

We first study the shape of the structures generated by BRISA, namely trees and DAGs with 2 parents. The shape (depth and degree), imposes constraints on latency and on the distribution of the dissemination effort. Results for each configuration are obtained after building the respective structure and letting it stabilize. The reason for using this basic strategy is twofold: i) a naive strategy helps to better understand the basic behavior of BRISA thus serving as a baseline for more elaborate strategies and ii) the limited number of physical nodes hides significative differences on the observation of structural shape properties changes that better emerge in the large scale. Depth places a lower bound on the dissemination time due to the cost of traversing several intermediate nodes and thus should be kept as low as possible. Figure 6 presents the depth distribution in a universe with 512 nodes. As expected, larger views allow nodes to have more children thus reducing maximum depth. The larger depths in DAGs are because depth measures the maximum distance, i.e. the longest path from the root to the node, which increases with the extra number of links. The steep curves hint that the structures built by BRISA are fairly balanced, i.e., do not degenerate into long chain even with a simplistic strategy thus

preserving desirable properties for dissemination. An analysis of the degree distribution confirms this observation.

The degree of a node in BRISA is given by the number of outgoing links and thus bounds the message copies a node receives/sends. This is directly related to the dissemination effort and as such degree distribution should be as narrow as possible indicating an evenly distributed load. When analyzing the degree distribution presented in Figure 7 three main observations arise. First DAGs are more effective than trees in having a greater share of the nodes contribute to the dissemination effort (nodes with degree zero are leaves). This is due to the additional number of parents that reduces the chance of having all outgoing links deactivated. Secondly, degree distribution is also highly affected by the view size provided by the PSS: higher values lead to shallower trees thus resulting in more leaves, while lower values lead to deeper trees due to the limitation imposed by the view sizes. Such relation between degree and depth can be observed in Figure 8, which depicts sample trees obtained by BRISA. As a matter of fact, despite using a simple strategy, the resulting trees are fairly balanced which is essential for efficient dissemination. Finally, despite using an *expansion factor* of 2 the number of nodes with degree higher than the configured value remains small as hinted in Section II-A.

B. Network properties

In this section we focus on the network properties of the dissemination structures obtained by BRISA.

First, we analyze the routing delay of dissemination. To this end, we use the cumulative round trip times, taken at each hop, from the root to a given node. When compared against the round trip time of direct communication between the root and that node, it indicates the effectiveness of BRISA in building dissemination structures with low end-to-end delays, an essential property for a dissemination system. The ratio between the first and second measurements gives the stretch factor. However, due to PlanetLab asymmetries that deter direct communication between some nodes we instead present the cumulative distribution of the raw results in Figure 9. Not surprisingly, the flooding strategy yields the worst results due mainly to the heavy load imposed on the network. In this non-structural metric, the effects of a delay-aware strategy become clear when compared to the simplistic first-come first-pick: for instance 40% of the nodes reduce the routing delays to half.

Next, we focus on the bandwidth usage. This measures the dissemination effort and is directly influenced by depth and degree distribution. Figure 10 and Figure 11 depict download and upload bandwidth usage, respectively, for payload sizes of 1, 10, 50 and 100 KB. We used stacked bars with decaying shades of grey for representing a distribution using a set of percentiles. For instance, the medium shade of grey gives the median value (half of the node below that value, the other half above), while the lighter shade gives the 90th percentile: 90% of the nodes are associated with a lower bandwidth.

As expected, trees are more frugal with respect to download as nodes receive exactly one copy of each message whereas




Fig. 8. Sample tree shape for 100 nodes with a HyParView active view size of 4 on the left and 8 on the right. Expansion factor is 1.




Fig. 9. Routing delays distribution on PlanetLab for a 150 nodes network. Structure is a tree with view size 4. Message size is 1KB×200 messages.




Fig. 11. Bandwidth usage for a 512 nodes network, upload.




Fig. 10. Bandwidth usage for a 512 nodes network, download

in DAGs nodes receive two copies (one for each parent). For each structure, the increase in bandwidth usage for the different view sizes is due to the PSS. The small difference, negligible when compared to application messages, hints at a low overhead service. The differences in the percentiles for the DAG are related to the depth of nodes (Figure 6) as nodes at lower depths may not be able to find additional parents and thus receive messages only from a single parent.

For upload, results are naturally similar. DAGs require more links and consequently nodes will have to relay messages to more neighbors, increasing upload bandwidth usage. The differences between percentiles for a given configuration are explained by the degree distribution (Figure 7) as nodes with higher degrees need to upload more.

C. Robustness

We now focus on the behavior of BRISA under continuous churn in order to assess its robustness. Each experiment is associated with a synthetic churn trace based on the churn support module of Splay. The synthetic description is given in Listing 1 and proceeds as follows: first we bootstrap the system and let it stabilize. After, we induce churn at rate X by having X percent nodes fail at random and X percent new nodes join the system during each minute.

```

from 1s to N s join N
at 1000s set replacement ratio to 100%
from 1000s to 1600s const churn X% each 60s
at 1600s stop

```

Listing 1. Churn trace generation script

Table I depicts the results obtained for networks with 128 and 512 nodes. For simplicity we ensure that the source node does not fail. However, we note that the failure of source node would only produce a negligible impact in the presented results. In fact only the direct children of the source (a small number limited by the view size) would experience the effect of a parent failure.

We defined the following metrics:

- **Parents lost per minute:** rate at which nodes lost any of their parents;
- **Orphans per minute:** rate at which nodes lost all parents (implying disconnections);
- **Percentage of soft repairs:** upon disconnections, how many nodes successfully repair their incoming links using the *soft repair* mechanism;
- **Percentage of hard repairs:** upon disconnections, how many nodes required using the *hard repair* mechanism.

As expected the rate at which parents are lost is higher for DAGs than trees due to the larger number of parents of the former. Nonetheless DAGs are much more robust with nodes being seldom fully disconnected. For instance, with a churn rate of 5% per minute, which implies half of the nodes leaving the system within the ten minutes of the experiment, only 17 nodes on an universe of 512 got disconnected (1.7 per minute * 10). Of those, all but one were able to recover using the soft repair, which simply implies activating a link to a new parent. Moreover, the time required for hard repairs, studied in the next section, is very low meaning that despite disconnections nodes are able to promptly repair their connectivity with minimal effort. Finally, quick parent recovery also allows nodes to quickly recover

lost messages thus ensuring that all application messages are effectively delivered. Such recovery capabilities under high churn, combined with efficient dissemination structures that are correct by design made BRISA a promising substrate for efficient and robust dissemination in the large scale.

D. Comparison with existing approaches

In this section we compare BRISA's bandwidth usage, structure construction time, dissemination latency and parent recovery delays with several approaches. Those protocols were chosen as representatives of different points in the efficiency/robustness design spectrum. The comparison is done against a BRISA tree with an HyperView active view size of 4. In order to assess the inherent overhead of each approach and for fairness reasons the other approaches were implemented and evaluated in the same environment as BRISA and configured with equivalent settings. We considered the following:

a) *SimpleGossip*: This approach lies on the robustness end of the spectrum. We use Cyclon [35] as the PSS. Due to its proactive nature we use a combination of rumor mongering (push) to infect most of the nodes and anti-entropy (pull) to ensure completeness [10]. Rumor mongering follows an infect and die strategy with a fanout of $\ln(N)$, where N is the system size and anti-entropy exchanges updates with a single random node with a frequency that is the double of the message creation ratio.

b) *SimpleTree*: Oppositely, this approach lies on the efficiency side of the design spectrum. We consider a tree created randomly with the help of a centralized node. The only criteria for a node joining the tree is to connect to a parent that joined earlier in the past, which avoids creating a cycle in a similar manner to the one used in TAG. This parent is provided by the centralized node that randomly picks any of the previously joined nodes as a parent for a newly joined node. Dissemination is done by pushing the messages immediately through tree links thus minimizing latency.

c) *TAG*: For this approach which tries to achieve both robustness and efficiency we use TAG [27]. As BRISA, TAG maintains a tree and a gossip-based overlay to combine the efficiency of trees and resilience of gossip. Nodes are further organized in a linked list structure sorted by joining time, with nodes maintaining information about their predecessors/successors up to two hops away. New nodes traverse this list backwards until an application specific condition is met. In the traversal, nodes pick k random peers to form the gossip overlay and join the tree by choosing a suitable parent. Upon parent failures, nodes update the linked list and traverse it to find a new parent and thus restore the tree. With respect to dissemination, TAG uses a pull-based approach with nodes pulling content both from the tree and from gossip neighbors. For instance, as TAG is designed with media streaming in mind, gossip partners can be used to pre-fetch segments (messages) that are still ahead of the play time. Because TAG relies on pull we expect increased dissemination latency due to the additional roundtrips and pull period. We chose to compare BRISA against TAG due to its proximity in terms




Fig. 12. Bandwidth usage for a 512 nodes network.

of goals and general approach (combining tree efficiency and gossip robustness) and the differences in its design choices (e.g., tree construction and a pull-based approach). We believe this choice will allow us to better assess the merits of each approach in the following evaluation scenarios.

Bandwidth usage We first focus on the bandwidth usage of each protocol by considering two metrics: *stabilization bandwidth* and *dissemination bandwidth*.

Stabilization bandwidth is the bandwidth used to bootstrap the protocol including the construction/emergence of the overlay and tree structures and is measured until stabilization. After stabilization we consider the *dissemination bandwidth* as the bandwidth associated with message disseminations and subsequent management overhead. Once the structure stabilizes, we inject messages with payload sizes from 0 to 20 KB in a network of 512 nodes. This differentiation allows us to clearly observe the overhead imposed in each phase. As SimpleGossip does not use any structure we represent all the bandwidth consumed under *dissemination bandwidth*.

Figure 12 presents bandwidth consumption averaged over all nodes. As expected, TAG and BRISA are comparable and the actual cost is dominated by the sending of data among peers rather than the management cost of bootstrapping the dissemination structures. The smaller management overhead of SimpleTree is due to the fact that only a single communication step with the centralized node is needed while the other protocols require inter-node communications. The small extra bandwidth cost for TAG and BRISA when compared to SimpleTree is from the maintenance of the PSS layer and dissemination structures that are key to the performance in terms of delays and robustness as we explore later. For the smaller message sizes, SimpleGossip is comparable with both BRISA and TAG due to the absence of structure management and because Cyclon does not use explicit fault detection mechanisms. However, this is quickly offset for larger message sizes due to the excessive number of duplicates SimpleGossip relays which leads to high network usage.

Structure Construction Time In this experiment we measure the time necessary to bootstrap the dissemination structures both on the cluster environment and on PlanetLab. Due to the absence of structure of SimpleGossip and the construction simplicity of SimpleTree, they are not considered in the

Churn conditions		Parents lost/min.	Orphans/min.	% Soft repairs	% Hard repairs
128 Nodes					
Churn rate: X=3%	Tree	2.3	2.3	87.0	13.0
(38 leaves & 38 joins /10 min.)	DAG, 2 parents	4.0	0.2	92.5	7.5
Churn rate: X=5%	Tree	3.4	3.4	79.4	20.6
(64 leaves & 64 joins /10 min.)	DAG, 2 parents	7.0	0.3	90.0	10.0
512 Nodes					
Churn rate: X=3%	Tree	22.2	22.2	88.2	11.8
(154 leaves & 154 joins /10 min.)	DAG, 2 parents	36.8	2.3	94	6
Churn rate: X=5%	Tree	22.2	22.2	87.7	12.3
(256 leaves & 256 joins /10 min.)	DAG, 2 parents	32.3	1.7	94.1	5.9

TABLE I
IMPACT OF CHURN IN BRISA FOR A 128 AND 512 NODE NETWORKS WITH ACTIVE VIEW SIZE 4.




Fig. 13. Construction time for 512 (on cluster) and 200 (PlanetLab) nodes network.

Protocol	Latency (seconds)	Overhead
SimpleTree	100,025	-
Brisa	106,587	+6%
SimpleGossip	128,23	+28%
TAG	200,476	+100%

TABLE II
DISSEMINATION LATENCY FOR A 512 NODES NETWORK FOR 500 MESSAGES OF 1KB.

experiment. For BRISA we consider the time elapsed since a node sends the first deactivation message until all its inbound links except one are deactivated. In the case of TAG we use the time since a node joins the list until it settles its position on that list. Results are presented in Figure 13. It is interesting to observe that in absolute terms (x scale is logarithmic) TAG is marginally faster than BRISA on the cluster but much slower on PlanetLab. This is because the construction mechanism happens at once in TAG by traversing the list, whereas in BRISA it is triggered by the reception of messages. As BRISA keeps the connection to its neighbors open, in the adverse environment of PlanetLab, the traversal cost of TAG (i.e. creating a connection to a node, exchanging messages, tearing it down and proceeding to the next node) easily outweighs the time BRISA needs to wait for the reception of the messages from all its neighbors.

Dissemination Latency We consider dissemination latency as the time elapsed between the reception of the first and last message among the set of all messages. When studied along with bandwidth usage it highlights the tradeoffs of each approach. The message payload is 1 KB and the ideal dissemination latency is 100 seconds (500 messages at 5 per

second). Table II presents the results averaged over all nodes. As SimpleTree is very close to the ideal value we use it as a baseline of comparison for the other approaches. Latency for TAG is significantly higher than the other approaches. This is mainly because TAG uses a pull-based approach to get updates, while the others rely on push. We note however that this is a characteristic that pertains to pull approaches in general and not TAG in particular. The delays for BRISA are similar to the ones for the SimpleTree, with a small variation that we account for the extra context switching and physical machines sharing on our cluster. Differences in practice are expected to be minimal with a SimpleTree, and largely in favor of BRISA when using a delay-aware selection strategy as previously illustrated by Figure 9. Somehow surprisingly, SimpleGossip performs worse than BRISA and SimpleTree. This is because the overhead of dealing with duplicates and eventual omissions that needs to be compensated by the slower anti-entropy mechanism.

Parent recovery delay Our last comparison considers the robustness of BRISA and TAG. As SimpleTree does not consider dynamic scenarios, and SimpleGossip does not maintain any structure both approaches are ignored in this experiment. We apply for both protocols the same churn conditions as described in Section III-C, with a churn rate of 3% and focus on the parent recovery delay for hard repairs in both cases. In BRISA this corresponds to the case where no immediate replacement neighbor is available and the underlying gossip layer is used. In TAG this corresponds to the case where the linked list is broken (i.e., two consecutive simultaneous node failures) and the node needs to be re-inserted into the structure. Figure 14 depicts the results in a 128 nodes network. We note that BRISA, while yielding a similar bandwidth cost, and better dissemination delays, also outperforms TAG regarding robustness in two ways: i) the number of hard repairs almost doubles with TAG (not shown) in the same churn conditions and ii) the delay for recovery is twice as fast for BRISA. This means that both the disruption of dissemination happens less often with BRISA, and that the effect of such disruptions is less than what is experienced with TAG.

IV. PERSPECTIVES

The work on BRISA brings several interesting perspectives that we highlight in this section.

Parent selection strategies: BRISA's parent selection is a flexible mechanism to build dissemination structures with different




Fig. 14. Parent recovery delays for a 128 nodes network with active view size 4, 3% continuous churn conditions.

performance criteria as highlighted in the evaluation. With this in mind there are several other strategies that can be employed to select the *best* parent. We highlight some possibilities: i) *gerontocratic*: which takes into account the uptime of the candidate nodes and selects the one with the highest value. This is based on the observation that the higher the uptime of a node, the more likely it is to remain available [4], ii) *heterogeneity-aware*: which considers the available bandwidth at candidate nodes and iii) *load-balancing*: which selects parents according to load and is to some extent the dual of the gerontocratic strategy by spreading the dissemination effort to newer nodes. It is important to note that the effectiveness of those strategies is effectively limited by two factors, the bounded size of the PSS view and its reactive nature. Adjusting the view size according to node capacities such as in [3] and using a proactive PSS such as Cyclon to improve the diversity of potential parents opens interesting challenges and possibilities.

Another possibility requiring further exploration is the use of DAGs with different number of parents based on several criteria such as available inbound bandwidth or parent reliability (as in the gerontocratic strategy).

Multiple Trees and Multiple Parents: Currently, BRISA is focused only on the construction of single trees or DAGs. The DAG approach can be used as a starting point to efficiently support multiple trees. For instance the parent(s) to which to relay the upward message can be chosen on a message basis thus enabling different strategies on a per-source basis. Other possibility is to choose the number of parents according to some criteria like the gerontocratic strategy for reliability. The main advantage of this approach is little to no overhead to support multiple trees/sources with different tradeoffs.

Stream splitting: When using a DAG configuration, nodes will have more than one parent. Currently, this is used to mask parent failures for applications sensitive to small interruptions due to the effect of the repair mechanism. An interesting perspective is to use this capability to obtain different parts of the stream from different parents, for instance a node with two parents can obtain half of the stream of messages from each parent. This idea is used in SplitStream [7] to enable fault-tolerance and load balancing, but in a quite rigid manner as all nodes need to participate in all trees. When combining the possibility of a varying number of parents (as in the

node heterogeneity strategy) with the possibility of obtaining different parts of the stream from different parents nodes could effectively improve inbound and outbound bandwidth usage.

V. RELATED WORK

Existing approaches to large-scale data dissemination cover two main design domains: overlay management and application-level multicast. In the following we present existing work in this design space and compare it to our approach.

Scribe [8] is an application-level multicast layer that builds dissemination trees by aggregating reverse paths to a rendezvous node in the Pastry [31] distributed hash table. Unlike BRISA, where we assume that all nodes are interested in all messages, Scribe supports group membership management by having each node subscribe to group(s) it is interested in. Yet, the load of dissemination is shared by non-members of the groups that must act as interior (forwarding) nodes in the dissemination trees. Unlike gossip-based dissemination, where the failure of a node has little impact on the functioning of the system, Scribe's *rendezvous* nodes can represent single points of failure and bottlenecks in the system. BRISA also considers emerging a dissemination structure from an existing overlay, but can leverage the gossip dissemination layer as a fallback for robustness. We note that group membership can be implemented in BRISA by maintaining on each node separate views for its subscribed groups, as done in the TERA publish/subscribe system [2]. These group specific views can themselves be constructed by the means of a gossip-based clustering protocol [17].

SplitStream [7] is a high-bandwidth dissemination layer that builds on top of Scribe [8] and Pastry [31]. In order to balance the load of dissemination, it constructs multiple trees that are used for sending alternate pieces of a stream; nodes that participate as a leaf in one tree participate as an interior node in the other(s), thus balancing the in- and out-degrees of nodes.

Chunkyspread [34] also builds multiple dissemination trees, rooted at a single source node. These trees are built on top of an unstructured overlay and not on a DHT. They are used to parallelize the dissemination process by pushing different parts of the data in each tree. Cycles in the trees are avoided by using a technique derived from Bloom filters, whereas BRISA relies on simpler mechanism based on the path or the number of hops from the source. Chunkyspread trees can be constructed by taking into account latency and load metrics that are also considered in the different BRISA parent selection strategies.

In Bullet [20], a stream of data is also pushed through a tree structure. Different data blocks are intentionally disseminated to different branches of the tree, taking into account the bandwidth limits of participating nodes. Bullet complements this tree with a gossip-like layer that allows repairing the faults in the dissemination (missed updates). This mechanism takes the form of a mesh that is used to locate peers with missing items, in a way similar to a PSS. In this sense, Bullet is based on a design choice that is opposite to ours: BRISA complements a robust dissemination layer (the PSS) with an efficient but failure-prone structure (tree/DAG), while Bullet

complements a tree with gossip-style dissemination to support failures. Rappel [29] is another example of a dissemination service that combines a tree structure for dissemination with a gossip-based service for optimization. In the case of Rappel, the gossip-based layer is used to locate suitable peers based on interest-affinity and network distances, and not as a fallback mechanism for dissemination.

MON [26] relies on a mechanism similar to BRISA to construct spanning trees and DAGs on top of an unstructured overlay. The goal of MON is to manage large-scale infrastructures such as PlanetLab, by using the resulting trees/DAGs to disseminate management commands. Therefore, sessions in MON are intended to be short-lived and the protocol does not provide any support for dynamism in the population of peers. To disseminate data, MON relies on a pull strategy, where nodes can download content simultaneously from multiple parents, if available. This approach eliminates duplicates, as it is the receiver that decides which pieces to receive. However it requires nodes to maintain knowledge of the data blocks/messages present at each parent.

The work presented in [36] stems from an observation similar to ours that even though gossip-based dissemination is attractive due to robustness, achieving completeness requires large fanouts resulting in high overhead. The authors thus propose a hybrid approach that uses gossip-based dissemination with fanouts lower enough to infect most of the population, and ensures completeness by relying on a ring structure that encompasses all nodes. Gossip is used for the bulk dissemination of data, still resulting in many duplicates, as opposed to BRISA, where most of the dissemination happens on the dissemination structure with a controlled number of duplicates. Similarly, in [24], [25] a Chord-like ring overlay is combined with a push mechanism to disseminate messages over a spanning tree optimized for minimal latency. BRISA instead builds on top of an unstructured overlay, and it offers a wider set of options for the tree construction.

In [12] the authors propose an alternative approach to tree repair based on proactive principles. Each node computes alternative parents for its children that can be used upon failures. This minimizes disruptions as nodes known beforehand the new parent they need to connect to. Further it can cope to some extent with multiple concurrent failures and strictly control node degrees, a major goal of the authors. Due to this restriction, tree shape tends to degenerate to a chain overtime penalizing end-to-end delay. BRISA uses a notion similar to the alternative parents without however having the tree degenerate into a chain. This is because [12] only considers potential parents in the failed node subtree while BRISA can consider any node as long as it passes the cycle detection mechanism.

GoCast [32] builds a dissemination tree embedded on a gossip-based overlay that takes into account network proximity to improve end-to-end latency. The tree is built using a traditional Distance Vector Multicast Routing Protocol (DVMRP) and used to push messages as in BRISA. Message identifiers are advertised through the overlay links as in PlumTree [22] and used to recover missing messages due to tree disruptions

that, contrary to BRISA, imposes additional network overhead. Most strikingly this recovery information is not used to repair the tree, which relies solely on DVMRP and thus presents scalability problems due to the overhead of periodic floods to rebuild the tree. Furthermore, BRISA is able to adjust to different performance criteria but could nonetheless take advantage of the network-proximity offered by Gocast's overlay. TAG, the protocol we use in the direct comparison with BRISA also falls into this class due to the use of a tree and a gossip-based overlay. More details can be found in Section III-D. PlumTree [22] also relies on the detection of duplicates and subsequent deactivation of links to build an embedded spanning tree on an underlying unstructured overlay. However, inactive links are still used in a "lazy push" approach, by announcing the message identifier instead of the full payload. These announcements are used to repair the tree: when an announcement for an unknown message is received, the protocol starts a timer. If the timer expires before the reception of the payload the tree repair mechanism is triggered. This approach is highly sensitive to variations in network latency, which lead to unnecessary message recoveries as observed in [13]. BRISA does not separate the dissemination of the header and payload, the dissemination is deterministic, and faults are detected thanks to the underlying PSS layer, which avoids sending periodic probe messages at the level of the dissemination layer. Further, the generic construction mechanism can build trees and DAGs according to different criteria, which is not possible in PlumTree. Due to the use of message advertisements to manage faults both PlumTree and Gocast fall in an undesirable tradeoff: either advertisements are sent sparingly to conserve bandwidth with an impact on recovery time, or advertisements are eagerly sent imposing a constant management overhead in the system.

Thicket [13] uses the same principles of PlumTree to build multiple dissemination trees on top of an unstructured overlay. The goal is to provide similar functionality to SplitStream by balancing the number of trees where a node is interior and also by splitting the content among trees to improve fault-tolerance. The mechanism used to build trees imposes several constraints that do not ensure the resulting tree is connected by design. This is addressed with a tree repair mechanism based on missing messages, as in PlumTree, that requires periodic exchanges of received messages among neighbors which is also used to handle joins and leaves. In contrast, BRISA builds connected trees by design, despite controlled fanouts, and deals with joins and failures with a simple and lightweight mechanism that is triggered only when failures happen. The support for multiple trees in Thicket is based on the premise of load balancing and fault-tolerance by leveraging on network coding techniques. The cost however is a linear growth in the number of links with respect to the number of trees, which poses scalability concerns. BRISA instead provides fault-tolerance and load balancing at the dissemination level by building generic dissemination structures in the form of DAGs. Although not considered in this paper, this opens the possibility of using more advanced load-balancing techniques

as presented in Section IV.

VI. CONCLUSIONS

In this paper, we presented the design and evaluation of BRISA, a data dissemination system that combines the robustness of gossip-based protocols and the efficiency of structured overlays. BRISA automatically emerges efficient dissemination structures from the flooding-based distribution of the first message in a stream. The construction of efficient dissemination structures exploits the path diversity that naturally exists in gossip- and flooding-based dissemination, while avoiding the high level of duplicate reception these mechanisms typically yield. The support for robustness comes from the ability of the underlying gossip layer to rapidly provide replacement nodes upon failures, and by acting as a fallback for reliability to support continuous service upon failures. Therefore, BRISA bridges the gap between reliable but costly gossip-based dissemination and efficient but failure prone tree-based dissemination. We evaluated BRISA with a prototype deployed on a cluster and on PlanetLab. The experiments and comparisons to related work confirmed BRISA as a robust and efficient system for data-intensive applications.

ACKNOWLEDGMENTS

This work was supported in part by the Swiss National Foundation under agreement number 200021-127271/1 and by the Portuguese Science Foundation (FCT) grants SFRH/BD/62380/2009 and PTDC/EIA-CCO/115570/2009.

REFERENCES

- [1] Planetlab. <http://www.planet-lab.org>.
- [2] R. Baldoni, R. Beraldi, V. Quema, L. Querzoni, and S. Tucci-Piergiovanni. Tera: topic-based event routing for peer-to-peer architectures. In *International conference on Distributed event-based systems*, pages 2–13. ACM Press, 2007.
- [3] O. Beaumont, L. Eyraud-Dubois, and S. K. Agrawal. Broadcasting on large scale heterogeneous platforms under the bounded multi-port model. In *International Symposium on Parallel & Distributed Processing*, 2010.
- [4] R. Bhagwan, S. Savage, and G. Voelker. Understanding availability. In *IPTPS: International Workshop on Peer-to-Peer Systems*, 2003.
- [5] K. Birman, M. Hayden, O. Ozkasap, Z. Xiao, M. Budiu, and Y. Minsky. Bimodal Multicast. *ACM Transactions on Computer Systems*, 17(2):41–88, 1999.
- [6] B. Bloom. Space/time trade-offs in hash coding with allowable errors. *Communications of the ACM*, 13(7):422–426, 1970.
- [7] M. Castro, P. Druschel, A.-M. Kermarrec, A. Nandi, A. Rowstron, and A. Singh. Splitstream: high-bandwidth multicast in cooperative environments. In *Symposium on Operating systems principles*, 2003.
- [8] M. Castro, P. Druschel, A.-M. Kermarrec, and A. Rowstron. SCRIBE: A large-scale and decentralized application-level multicast infrastructure. *IEEE Journal on Selected Areas in Communications*, 20:2002, 2002.
- [9] Y. Chu, S. Rao, S. Seshan, and H. Zhang. A case for end system multicast. *Selected Areas in Communications, IEEE Journal on*, 20(8):1456–1471, Oct. 2002.
- [10] A. Demers, D. Greene, C. Hauser, W. Irish, J. Larson, S. Shenker, H. Sturgis, D. Swinehart, and D. Terry. Epidemic algorithms for replicated database maintenance. In *PODC: Symposium on Principles of distributed computing*, pages 1–12, 1987.
- [11] P. Eugster, R. Guerraoui, S. Handurukande, P. Kouznetsov, and A.-M. Kermarrec. Lightweight probabilistic broadcast. *ACM Transactions on Computer Systems*, 21(4):341–374, 2003.
- [12] Z. Fei and M. Yang. A proactive tree recovery mechanism for resilient overlay multicast. *IEEE/ACM Trans. Netw.*, 15:173–186, February 2007.
- [13] M. Ferreira, J. Leita, and L. Rodrigues. Thicket: A Protocol for Building and Maintaining Multiple Trees in a P2P Overlay. *Symposium on Reliable Distributed Systems*, 0:293–302, 2010.
- [14] D. Frey, R. Guerraoui, A.-M. Kermarrec, M. Monod, and V. Quema. Stretching gossip with live streaming. In *International Conference on Dependable Systems Networks*, pages 259–264, 2009.
- [15] A. Ganesh, A.-M. Kermarrec, and L. Massoulié. Scamp: Peer-to-Peer Lightweight Membership Service for Large-Scale Group Communication. In *Networked Group Communication*, pages 44–55, 2001.
- [16] J. Gantz. The Diverse and Exploding Digital Universe. Technical report, IDC White Paper - sponsored by EMC, 2008.
- [17] M. Jelasity, A. Montresor, and O. Babaoglu. T-man: Gossip-based fast overlay topology construction. *Computer Networks*, 53(13), aug 2009.
- [18] M. Jelasity, S. Voulgaris, R. Guerraoui, A.-M. Kermarrec, and M. van Steen. Gossip-based peer sampling. *ACM Transactions on Computer Systems*, 25(3), aug 2007.
- [19] B. Kaldehofe. Buffer management in probabilistic peer-to-peer communication protocols. In *Proceedings of the 22nd International Symposium on Reliable Distributed Systems*, pages 76–85, 2003.
- [20] D. Kostic, A. Rodriguez, J. Albrecht, and A. Vahdat. Bullet: High bandwidth data dissemination using an overlay mesh. Oct. 2003.
- [21] J. Leitão, J. Pereira, and L. Rodrigues. HyParView: A membership protocol for reliable gossip-based broadcast. In *International Conference on Dependable Systems and Networks*, pages 419–428, 2007.
- [22] J. a. Leitão, J. Pereira, and L. Rodrigues. Epidemic Broadcast Trees. *SRDS: International Symposium on Reliable Distributed Systems*, 2007.
- [23] L. Leonini, E. Rivière, and P. Felber. SPLAY: Distributed systems evaluation made simple (or how to turn ideas into live systems in a breeze). In *NSDI: Symposium on Networked Systems Design and Implementation*, pages 185–198. USENIX, 2009.
- [24] Z. Li, G. Xie, K. Hwang, and Z. Li. Churn-resilient protocol for massive data dissemination in p2p networks. *IEEE Transactions on Parallel and Distributed Systems*, 22:1342–1349, 2011.
- [25] Z. Li, G. Xie, and Z. Li. Towards reliable and efficient data dissemination in heterogeneous peer-to-peer systems. In *Parallel and Distributed Processing, 2008. IPDPS 2008. IEEE International Symposium on*, pages 1–12. IEEE, 2008.
- [26] J. Liang, S. Ko, I. Gupta, and K. Nahrstedt. MON : On-demand Overlays for Distributed System Management. *WORLDS: Workshop on Real, Large Distributed Systems*, pages 13–18, 2005.
- [27] J. Liu and M. Zhou. Tree-assisted gossiping for overlay video distribution. *Multimedia Tools Appl.*, 29:211–232, June 2006.
- [28] R. Melamed and I. Keidar. Araneola: A scalable reliable multicast system for dynamic environments. In *International Symposium on Network Computing and Applications*, pages 5–14. IEEE, 2004.
- [29] J. A. Patel, E. Rivière, I. Gupta, and A.-M. Kermarrec. Rappel: Exploiting interest and network locality to improve fairness in publish-subscribe systems. *Computer Networks*, 53(13):2304–2320, 2009.
- [30] R. V. Renesse, Y. Minsky, and M. Hayden. A gossip-style failure detection service. *Middleware Conference*, pages 55–70, 2009.
- [31] A. Rowstron and P. Druschel. Pastry: Scalable, decentralized object location and routing for large-scale peer-to-peer systems. In *Lecture Notes in Computer Science*, volume 2218, pages 329–350, 2001.
- [32] C. Tang, R. N. Chang, and C. Ward. Gocast: Gossip-enhanced overlay multicast for fast and dependable group communication. In *DSN’05: Proceedings of the 2005 International Conference on Dependable Systems and Networks*, pages 140–149, Washington, DC, USA, 2005. IEEE.
- [33] Twitter Engineering. Murder: Fast datacenter code deploys using BitTorrent. <http://t.co/uo5rEN4>.
- [34] V. Venkataraman, K. Yoshida, and P. Francis. Chunkspread: Heterogeneous Unstructured Tree-Based Peer-to-Peer Multicast. *International Conference on Network Protocols*, pages 2–11, 2006.
- [35] S. Voulgaris, D. Gavidia, and M. van Steen. Cyclon: Inexpensive membership management for unstructured p2p overlays. *Journal of Network and Systems Management*, 13(2), June 2005.
- [36] S. Voulgaris and M. van Steen. Hybrid dissemination: Adding determinism to probabilistic multicasting in large-scale p2p systems. In *International conference on Middleware*, 2007.
- [37] S. Zhuang, B. Zhao, A. Joseph, R. Katz, and J. Kubiawicz. Bayeux: An architecture for scalable and fault-tolerant wide-area data dissemination. In *International Workshop on Network and Operating Systems Support for Digital Audio and Video*, pages 11–20. ACM Press, 2001.