

SPLAYNET : Towards Nearby Self-Adjusting Wireless Networks

Gaddipathi Bharathi*¹, Parchuri Kaneeja²

*¹Associate Professor, Department of MCA , St. Mary's Group of Institutions, Guntur, Andhra Pradesh, India.

²PG Student ,Department of MCA , St. Mary's Group of Institutions, Guntur, Andhra Pradesh, India

ABSTRACT

This paper starts the investigation of locally self-changing systems: organizes whose topology adjusts progressively and in a decentralized way, to the correspondence design . Our vision can be viewed as a dispersed speculation of the self-altering datastructures presented by Sleator and Tarjan [22]: rather than their spread trees which powerfully upgrade the query costs from a solitary hub (in particular the tree root), we try to limit the steering cost between discretionary correspondence combines in the system. As an initial step, we consider disseminated twofold hunt trees (BSTs), which are appealing for their help of insatiable steering. We present a straightforward model which catches the key tradeoff between the advantages and expenses of self-changing systems. We introduce the SplayNet calculation and formally dissect its execution, and demonstrate its optimality in particular contextual investigations. We additionally present lower bound methods in light of interim cuts and edge development, to think about the restrictions of any request enhanced system. At last, we stretch out our investigation to multi-tree systems, and highlight an interesting contrast amongst great and disseminated spread trees.

Keywords : BST, Self-Adjusting Wireless Networks, SPLAYNET

I. INTRODUCTION

In the 1980s, Sleator and Tarjan [22] proposed an engaging new worldview to outline effective Binary Search Tree (BST) data structures: as opposed to advancing customary measurements, for example, the inquiry tree profundity in the most pessimistic scenario, their spread data structure self-changes with its utilization design, moving even more oftentimes got to components nearer to the root. A characteristic execution metric to assess a self-changing framework is the amortized taken a toll: the "normal cost" for a most pessimistic scenario arrangement of operations (of a specific class).

Since this fundamental work, self-altering data structures have been considered seriously, and different more productive self-changing data structures, for example, Tango BSTs [7] or multi-spread trees [23] have been proposed. Specifically, the famous Dynamic Optimality guess [7] keeps on confounding scientists:

the guess asserts that spread trees execute and some other paired inquiry tree calculation up to a steady factor.

Rather than these adaptable exemplary data structures, the present disseminated data structures and systems are still enhanced

Toward static measurements, for example, the breadth or the length of the longest course: the self-altering worldview has not overflowed to appropriate arranges yet.

We, in this paper, start the investigation of a circulated general-inaction of self-streamlining data structures. This is a non-insignificant speculation of the exemplary spread tree idea: While in class-sic BSTs, a query ask for dependably begins from a similar hub, the tree root, disseminated data structures and systems, for example, skip diagrams [2], [13] need to help directing solicitations between subjective combines (or

companions) of conveying hubs; at the end of the day, both the source and additionally the goal of the solicitations wind up noticeably factor. Figure 1 shows the distinction amongst exemplary and appropriated parallel pursuit trees.

In this paper, we ask: Can we receive comparable rewards from self-modifying whole systems, by adaptively diminishing the separation between every now and again imparting hubs?

As an initial step, we investigate completely decentralized and self-modifying Binary Search Tree systems: in these systems, hubs are orchestrated in a paired tree, which regards hub identifiers. A BST topology is appealing as it bolsters avaricious directing: a hub can choose locally to which port to forward a demand given its goal address.

A. Our Contributions

This paper makes the accompanying commitments.

1) We start the investigation of self-changing appropriated reports structures and present a formal model likewise. Our model is straightforward yet catches the basic tradeoffs between the advantages of self-modifications (to be specific shorter steering ways) and their expenses (specifically reconfigurations).

2) We present a self-altering disseminated BST called Splay Net. Splay Net is a characteristic speculation of the exemplary spread tree calculation, which "spreads" communication accomplices to their normal predecessor. Splay Net is completely decentralized as in every topological change and additionally directing is neighbourhood.

3) We formally dissect the execution of Splay Net (as far as amortized costs). Specifically, we demonstrate that the general cost is upper limited by the exact entropies of the sources and goals in the correspondence design; a basic lower bound takes after from restrictive experimental entropies. We likewise demonstrate the optimality of our approach in particular contextual investigations, e.g., when the communication design takes after an item circulation. At last, we additionally display a dynamic programming calculation to ideally take care of the disconnected issue variation in polynomial time.

4) We acquaint novel lower bound procedures with think about the confinements of self-changing systems. These methods depend on interim cuts and edge extension, and may. Be of unbiased interest and find programs beyond the putting studied on this paper.

Five) in the end, we provoke the dialogue of greater complex self-adjusting networks, particularly topologies inclusive of multiple trees. We make the interesting remark that in contrast to traditional data structures in which the self-adjustment blessings of multiple timber is restricted, in a distributed placing, a single extra best can on occasion reduce the amortized fee dramatically.

In summary, our work shows that even as a few algorithmic concepts of conventional splay trees may be generalized to networks, the distributed placing calls for new analytical tools. Furthermore, our consequences spotlight that self-adjustment benefits can certainly be reaped also within the context of networks; for multi-tree networks, those benefits can even be drastically better than in traditional data structures.

In trendy, we regard our look at as a first step, and trust that our model and outcomes open a rich location for future studies.

B. Paper Organisation

The imminent segment ii introduces our formal model and presents the reader with the essential heritage. section iii describes an offline set of rules to compute foremost static disbursed bests and presents the splay net method. phase iv derives entropy-based totally upper and decrease bounds on the performance of splaynets, and segment v studies the locality and convergence properties of the splaynet algorithm in specific eventualities. phase vi then derives stepped forward lower bounds which permit us to show the optimality of splaynets in additional situations. in phase vii we initiate the dialogue of datastructures primarily based on multiple bsts. after reviewing associated paintings in section viii, we conclude our contribution in segment ix. within the appendix, a few extra technical information are supplied.

II. Version and history

We take into account a set of n nodes (or peers) $v = \{1, \dots, n\}$ interacting in line with a positive communication sample. The sample is modeled by way of $t = (0, 1, \dots, m-1)$: a chain of m conversation requests wherein $t = (u, v) \in V \times V$, with supply u to destination v , henceforth now and again denoted with the aid of $\text{src}(t)$ and $\text{dst}(t)$, respectively. Our goal is to find a verbal exchange community g which connects the nodes v in step with the conversation sample: additionally, g have to be selected from a certain circle of relatives of favored topologies g , as an example, the set of tree topologies (the point of interest of this paper), expander graphs, or low-diameter networks, and so forth. every topology $g \in \mathcal{G}$ is a graph $g = (V; E)$. we distinguish trouble versions: (1) a static variation where g can be optimized in the direction of the verbal exchange pattern inside the feel that it may make the most, e.g., lengthy-term characteristics of t , however, g is constant and can't change over time. (2) a self-adjusting variation where g may be adapted through the years.

Typically, it's miles suited that networks are adjusted easily, and we're interested by nearby transformations: converting communication pattern ends in "nearby" adjustments of the communication graph through the years.

2. as cited above, this paper focuses on a placing in which g represents the set of binary search trees (bsts), for this reason-forth sometimes clearly called tree networks. except their simplicity, bsts are appealing for their low node degree and the opportunity to path regionally: given a destination identifier (or cope with), each node can determine domestically whether to forward the packet to its left toddler, its proper baby, or its determine; see appendix a for info.




Figure 1: (a) Classic BST versus (b) conveyed BST: Classic spread tree datastructures enhance the separation of the components from the root (the query cost), while in an appropriated datastructure,

correspondence happens between self-assertive hubs (the associates). Additionally take note of that in a BST organize, asks for likewise travel upwards in the tree; by the by, as we will see, steering choices are totally nearby.

The neighborhood changes of tree systems are called ro-tations. Casually, a revolution in an arranged twofold hunt tree switches up to three contiguousness connections, while keeping subtrees in place. Note that it is conceivable to change any parallel hunt tree into some other twofold pursuit tree by a succession of neighborhood changes (e.g., by enlistment over the subtree roots).

For our formal examination, we consider a disentangled syn-chronous show where initial a correspondence ask for arrives, at that point nearby system changes can be performed, lastly, the demand is fulfilled (i.e., the movement steered). In this paper, we are regularly keen on a setting where the solicitations are drawn indiscriminately from a settled however obscure correspondence lattice. Solidly, we will in some cases see the correspondence asks for as instigating a demand chart $R(\cdot) = (V(\cdot); E(\cdot))$ over the vertices V ; the edges $E(\cdot)$ of $R(\cdot)$ are commented on with recurrence data. (At the point when clear from the unique circumstance, we will regularly discard in $R(\cdot)$, $V(\cdot)$, $E(\cdot)$, and basically compose R, V, E .)

Solidly, the hub set V of R is given by the arrangement of hubs taking part in t , i.e., $V = \{v : \exists t \in T, v \in t\}$, and the arrangement of coordinated edges E is given by $E = \{e \in E : \exists t \in T, e \in t\}$. The weight $w(e)$ of each coordinated edge $e = (u, v) \in E$ is the recurrence $f(u, v)$ of the demand from u to v in T . In the accompanying, we will once in a while basically compose $w(u, v)$ to indicate the weight $w(e)$ of edge e . For instance, in a few situations the correspondence design between the hubs V may shape a tree (e.g., a multicast tree), an entire diagram, or an arrangement of separated parts (e.g., portraying a bunched correspondence design).

let a be an algorithm that given the request t and the graph $g \in \mathcal{G}$ at time t , transforms the present day graph (thru neighborhood ameliorations) to $g_{t+1} \in \mathcal{G}$ at time $t + 1$. we can use the notation $a = ?$ to refer to a static (i.e., non-adjusting) "al-gorithm" which does no longer exchange the conversation community over time.

We are inquisitive about the essential tradeoff between the benefits of self-adjusting algorithms (i.e., shorter routing paths) and their charges (specifically reconfiguration costs). We introduce a maximum simple, linear price model that captures this tradeoff. Concretely, we denote the fee of the community differences at time t by way of $(a; g_t; t)$, and we denote the quantity of rotations executed to exchange g_t to g_{t+1} ; while a is clear from the context, we will without a doubt write t . We denote with $dg(\cdot)$ the distance function between nodes in g , i.e., for 2 nodes $v; u \in V$ we define $dg(u; v)$ to be the wide variety of edges of a shortest course between u and v in g , and we count on messages are routed alongside the shortest paths. For a given collection of communication requests, the cost for an algorithm is given via the wide variety of changes and the distance of the conversation requests plus one (i.e., additionally a request $(u; u)$ comes at a minimum cost of one unit).

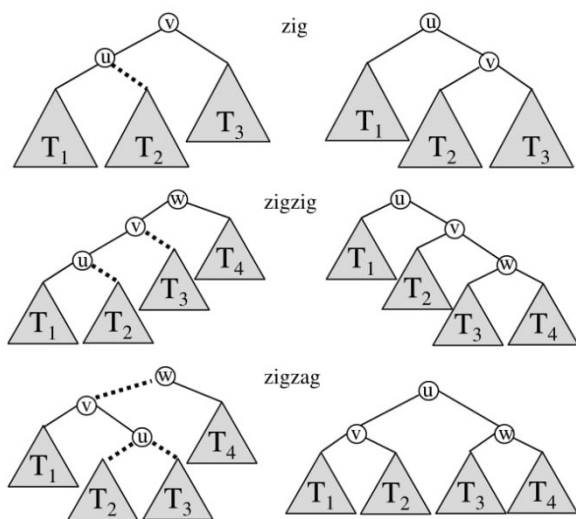


Fig. 2: Basic rotations of splay trees. The dashed bold line indicate adjacency relationships which are not maintained during the operation.

In the subsequent formal definitions, we need the average and amortized cost. For an algorithm A and given an initial community g_0 with node distance function $d(\cdot)$ and a sequence $\sigma = (\sigma_0; \sigma_1; \dots; \sigma_{m-1})$ of communication requests through the years, we outline the (average) price of A as:

The amortized value of A is defined because the worst feasible fee of A , i.e., $\max_{g_0} \text{fee}(A; g_0; \sigma)$.

Entropy and Empirical Entropy

The entropy of the communication pattern turns out to be a useful parameter to evaluate the performance of self-adjusting SplayNets. For a discrete random variable X with possible values $f(x); x$, the entropy $H(X)$ of X is $n \geq 1$ that $P_i \geq 0$ defined as $H(X) = -\sum_{i=1}^n p(x_i) \log_2 p(x_i)$ where $p(x_i)$ is the probability X takes the value x_i . Note that, $0 \log 1$ is considered as 0. For a joint distribution over $X; Y$, the joint entropy is defined as $H(X; Y) = -\sum_{i,j} p(x_i; y_j) \log_2 p(x_i; y_j)$. Also $p(x_i; y_j)$ definition of the conditional entropy $H(X|Y)$: recall the $n \geq 1$ $P_j H(X|Y = y_j) = \sum_{j=1}^n p(y_j) H(X|Y = y_j)$. sequence of communications is revealed over Since the P

Since the P time and may not be chosen from a fixed probability distribution, we are often interested in the empirical entropy of, i.e., the entropy implied by the communication frequency $f(x_1); \dots; f(x_n)$ be the empirical

Let $H(\cdot)$ = entropy measure of the frequency distribution of the communication sources (origins) occurring in the communication sequence, i.e., $f(x_i)$ is the frequency with which a node x_i appears as a source in the sequence, i.e., $f(x_i) =$ and analogously, the empirical conditional entropies $H(X|Y)$ and $H(Y|X)$.

B. Splay Trees

Our work can be regarded as a distributed generalization of splay trees, binary search trees whose topology adapts to the lookup sequence. Indeed, assuming that all requests originate from the same node, the SplayNet problem becomes equivalent to the classic splay tree problem. In the following, we hence briefly review the concept of splay trees.

For a node set V with unique identifiers (IDs) or values, we consider the family B of the set of all binary search trees over the IDs of V . Let $s = (v_0; v_1; \dots; v_{m-1})$, $v_i \in V$, be a sequence of lookup requests. In the classic offline problem (i.e., for algorithm $A = ?$), the goal is to find the best search tree $T \in B$ that minimizes the cost $\text{Cost}(T; s)$.

We will make use of the following two well-known properties of optimal BSTs

III. SPLAY NETWORKS

We first acquaint ourselves with the distributed BST model by studying optimal static topologies. Subsequently, we present the SplayNet approach and introduce the simple SPLAYNET algorithm to self-adjust the network.

A. Optimal Static Distributed BST

Given a certain communication pattern or “guest graph”, the optimal BST can be computed in polynomial time using a dynamic programming approach. The main insight needed is that the problem for the entire tree can be decomposed into optimal subproblems for smaller trees, and that the demand towards a given node in a subtree can be decoupled from nodes outside a given subtree: the precise topological structure of the nodes outside a subtree does not matter.

Concretely, the optimal static tree T which minimizes the sum of the weighted node distances

B. Self-Adjusting Distributed BSTs

Let us now consider self-adjusting BST networks. The SplayNet algorithm presented in this paper is a natural generalization of the classic splay tree algorithm. It is based on a double splay strategy: similarly to classic splay trees, SplayNet aggressively moves communicating nodes together; however, rather than splaying nodes to the root of the BST, locality is preserved in the sense that the source and the destination node are only rotated to their common ancestor (of the subtree covering the communication partners).

Concretely, consider a communication request $(u; v)$ from node u to node v , and let $T(u; v)$ denote the lowest common ancestor of u and v in the current network T . For an arbitrary node x , let $T(x)$ be the subtree rooted at x . When a request $(u; v)$ occurs, SplayNet first simply splays u to the lowest common ancestor $T(u; v)$ of u and v , using the classic splay operations Zig, ZigZig, ZigZag from [22] (see Figure 2). We assume that the splay function returns the tree resulting from these operations. Subsequently, the idea is to splay the destination node v to the child of the lowest common ancestor $T^0(u; v)$ of u and v in the resulting tree T^0 . Observe that this common ancestor is u itself ($u = T^0(u; v)$), i.e., we define the double splay algorithm SplayNet to splay v such that it becomes a child of u . (Note that the child is uniquely

defined: if $u > v$, v will be the left, if $u < v$, the right child of u .)

The formal algorithm listing of SplayNet is shown in Algorithm 2.

Algorithm 2 Algorithm SplayNet

-
1. (* upon request $(u; v)$ in T^*)
 2. $w := T(u; v)$
 3. $T^0 := \text{splay } u \text{ to root of } T(w)$
 4. splay v to the child of $T^0(u)$
-

CASE STUDIES: LOCALITY AND OPTIMALITY

The section provides insights into the properties of our algorithms in different specific settings. In particular, using different case studies, we show that our approach sometimes exhibits a desirable local convergence to the optimal network. The section also serves the purpose of introducing some analytical tools that are useful to study SplayNets.

A. Cluster Scenario

We first study a scenario which shows the locality properties of SplayNet. In this scenario, requests are clustered, and so is the resulting tree of SplayNet.

Definition 3 (Cluster Scenario). In a cluster scenario the communication pattern partitions the nodes into k contiguous and disjoint intervals $I_1 [I_2 [I_3 [\dots [I_k$ where nodes within an interval I_j have consecutive numbers and where communication only happens between node pairs in the interval. In particular, a request $(u; v)$ implies that u and v belong to the same interval: $(u; v) \in I_j$. A maximal cluster scenario is a cluster scenario where no intervals can be divided into two intervals. (a) for an example of a maximal cluster scenario. For this case we are able to prove the following.

Theorem 8. In a maximal cluster scenario, SplayNet features the following two properties:

- 1) SplayNet will eventually construct a tree network in which for any communication pair $(u; v) \in I_j$, for any $j \in \{1, \dots, k\}$, u and v are connected by a local path which only includes nodes from I_j .
- 2) Once this local routing property is established, it will never be violated again.

Proof: For any BST T and an interval I_j let r_j be the node such that the subtree $T(r_j)$ is the smallest size sub-tree where all the nodes of I_j are in $T(r_j)$; we denote this tree by $T(I_j)$. Let $out(T(I_j))$ denote the number of requests $(u; v) \in R(\cdot)$ s.t. either u or v are in $T(I_j)$, but not both.

For the (maximal) cluster scenario and a BST T , we consider the potential function $\Phi = \sum_{j=1}^k out(T(I_j))$. We will prove the theorem by first showing that when $\Phi = 0$, for P

any $(u; v) \in I_j, j \in \{1, \dots, k\}$, u and v are connected by a local path: a path which only includes nodes from I_j ; remains zero after such a request. Subsequently, we will show

Illustration for proof of Claim 1 and Theorem 8. $T(r_j)$ is the tree rooted at r_j (and includes T_1 and T_2). that when $\Phi > 0$, SplayNet cannot increase Φ , and there exists a request in which Φ will reduce the potential.

We start with the following claim.

Claim 1. Consider I_j s.t. $out(T(I_j)) = 0$. Then this property is invariant for all future requests in and $out(T(I_j))$ will always remain 0.

Proof: A generic situation for $T(I_j)$ is illustrated Each node not in I_j is in one of 3 possible locations: in one of the (possibly empty) subtrees T_1 ; T_2 ; or T_3 . Clearly, T_1 contains only IDs smaller than I_j and must be attached to the smallest ID in I_j , and T_2 contains only IDs larger than I_j and must be attached to the largest ID in I_j . T_3 can contain either smaller or larger IDs depending on whether r_j is a left or a right son of its parent. Since $out(T(I_j)) = 0$, there are no requests involving T_3 so all requests remain within $T(I_j)$ and so $out(T(I_j))$ will remain 0 for all future requests.

Now observe that if $\Phi = 0$, every $out(T(I_j)) = 0$, and so will remain zero also in the future.

Next we claim that no request can increase Φ , and that if $\Phi > 0$, there is a request in that will decrease the potential function. Consider $T(I_j)$ and any request $(u; v)$. If $(u; v) \in I_j$ then $out(T(I_j))$ does not change. If $u; v \in T_i$ for $i = 1; 2$ or 3 , then again $out(T(I_j))$ does not change. If $u \in T_1$ or $u \in T_2$ and $v \in T_3$, then the lowest common ancestor of $u; v$ is in T_3 and after splaying u and v , $out(T(I_j))$ will decrease by one. Therefore Φ can only decrease. Now if $\Phi > 0$, there is some j for which $out(T(I_j)) > 0$. Take the request $(u; v)$ that is a witness; after

this request $out(T(I_j))$ will decrease by one, so will decrease. So overall given a cluster \mathcal{C} , Φ will decrease to zero. To conclude the theorem we show that if $\Phi = 0$ then for any communication pair $(u; v) \in I_j$, for any $j \in \{1, \dots, k\}$, u and v are connected by a local path which only includes nodes from I_j . Assume by contradiction that this is not the case and there is a j and $(u; v) \in I_j$ (assume w.l.o.g. that $u < v$) s.t. the path between them visits a node $w \in I_i$. Also assume w.l.o.g. that $w < u$. Now it must be the case that either $T(r_j) \cap T(r_i)$ or $T(r_i) \cap T(r_j)$. Let's assume w.l.o.g. that $T(r_j) \cap T(r_i)$ so it must be that case that $r_j < w$. Since it is a maximal cluster scenario there must be a path in from a node $x \in I_i$, $x > w$ to a node $y \in I_i$, $y < r_i$ (otherwise the cluster I_i can be divided into two clusters).

B. Non-crossing Matching Scenario

SplayNet sometimes achieves an optimal performance by converging to the optimal static network implied by $R(\cdot)$. We have already encountered an example where SplayNets are asymptotically optimal, namely if describes a product distribution (cf Corollary 1). In the following, we will examine other optimal scenarios in more detail.

For a request $(u; v)$ in \mathcal{C} , let $I(u;v)$ denote the interval $[\min(u; v); \max(u; v)]$.

Definition 4 (Non-Crossing Matching Scenario). In a non-crossing matching scenario, it holds for the communication pattern that for any two pairs $(u_1; v_1)$ and $(u_2; v_2)$ in \mathcal{C} , either:

- 1) $I(u_1;v_1) \cap I(u_2;v_2) = \emptyset$ or $I(u_2;v_2) \cap I(u_1;v_1) = \emptyset$, or
- 2) $I(u_1;v_1) \supset I(u_2;v_2)$

It follows from the definition that the request graph $R(\mathcal{C})$ must describe a matching, and for each request $(u; v)$ there are no other requests that enter or leave (i.e., cross) the

interval $I(u;v)$. provide an example of a non-crossing matching scenario. If $R(\mathcal{C})$ describes a non-crossing matching scenario, SplayNet will converge to an optimal solution. Formally:

Theorem 9. In a non-crossing matching scenario \mathcal{C} , SplayNet will eventually converge to a tree where any communication pair $fu; vg \in \mathcal{C}$ is adjacent.

Proof: To any request pair $p_i = (u_i; v_i)$ we assign a level $L(p_i)$ depending on the number of pairs $p_j = (u_j; v_j)$ it is nested in, i.e., $L(p_i) = \sum_{p_j : p_i \subset p_j} L(p_j)$. Pairs with level 0 are the outmost pairs, pairs at level 1 are nested in a single other pair and so on. We will prove by induction on the level that all pairs become (and stay) adjacent. Our induction hypothesis is that pairs at level i coverage and we will show that once that happens, pairs at level $i + 1$ will converge next. First we show that pairs at level 0 converge. Notice that pairs do not intersect (i.e., cross), so the pairs at level 0 represent a clustered scenario. By Theorem 8, after the convergence of clusters, requests within each interval (of pair at level 0) will remain only within the interval. After the cluster convergence, we claim that after a pair $p_i = (u_i; v_i)$ of level 0 communicated, it will stay adjacent forever. To see this, consider any other pair $p_j = (u_j; v_j)$. If p_j is outside the cluster of p_i we are done. Let p_j be inside the interval of p_i , i.e., $I(u_j; v_j) \subset I(u_i; v_i)$. By the definition of SplayNet, after a request $(u_i; v_i)$, v_i is the right (resp. left) child of u_i if $u_i < v_i$ (resp. $u_i > v_i$). We will show that this implies that the nodes $u_j; v_j$ must be both in the same subtree, and can hence not change the adjacency relationship of u_i and v_i anymore. The claim follows by case distinction: (1) If $u_i < v_i$, the left

C. Multicast Scenario

To give one more example where SplayNet is optimal, we consider the multicast tree scenario. For this scenario, we may assume that an overlay network of a binary search tree structure is constructed to facilitate in-network duplication. The same tree may be used by many multicast sources and different receivers, hence the local (hop-by-hop) source and destination pairs (i.e., communication requests) are the endpoints of the tree edges

IV. CONCLUSION

We regard our work as a first step towards the layout of novel distributed datastructures and networks which adapt dynamically to the demand. For you to cognize on the fundamental tradeoff among advantage and cost of self-changes, we purposefully offered our model in a widespread and summary form, and many extra and application-particular elements want to be addressed before our technique can be examined inside the wild. The primary theoretical

simplification made in this paper regards the limit to the tree topology, and the generalization to extra complex and redundant networks is an open query. Furthermore, further to [22], we have focused at the amortized expenses of splaynets, and an thrilling path for future research regards the examine of the conceivable competitive ratio under arbitrary communique styles.

V. REFERENCES

- [1]. B. Allen and I. Munro. Self-organizing binary search trees. *J. ACM*, 25:526-535, 1978.
- [2]. J. Aspnes and G. Shah. Skip graphs. *ACM Transactions on Algorithms (TALG)*, 3(4):37-es, 2007.
- [3]. C. Avin, M. Borokhovich, and S. Schmid. Obst: A self-adjusting peer-to-peer overlay based on multiple bsts. In *Peer-to-Peer Computing (P2P)*, 2013 IEEE Thirteenth International Conference on, pages 1-5. IEEE, 2013.
- [4]. C. Avin, M. Borokhovich, and S. Schmid. Obst: A self-adjusting peer-to-peer overlay based on multiple bsts. In *Proc. 13th IEEE International Conference on Peer-to-Peer Computing (P2P)*, September 2013.
- [5]. C. Avin, B. Haeupler, Z. Lotker, C. Scheideler, and S. Schmid. Locally self-adjusting tree networks. In *Parallel & Distributed Processing (IPDPS)*, 2013 IEEE 27th International Symposium on, pages 395-406. IEEE, 2013.
- [6]. C. Avin, B. Haeupler, Z. Lotker, C. Scheideler, and S. Schmid. Locally self-adjusting tree networks. In *Proc. 27th IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, May 2013.
- [7]. E. D. Demaine, D. Harmon, J. Iacono, and M. Patrascu. Dynamic optimality - almost. In *Proc. 45th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 484-490, 2004.
- [8]. J. Díaz, J. Petit, and M. Serna. A survey of graph layout problems. *ACM Comput. Surv.*, 34(3):313-356, 2002.
- [9]. J. Díaz, J. Petit, and M. Serna. A survey of graph layout problems. *ACM Comput. Surv.*, 34(3):313-356, 2002.
- [10]. H. Farvaresh and M. Sepehri. A branch and bound algorithm for bi-level discrete network

- design problem. *Networks and Spatial Economics*, 13(1):67-106, 2013.
- [11]. U. Feige and J. Lee. An improved approximation ratio for the minimum linear arrangement problem. *Information Processing Letters*, 101(1):26-29, 2007.
- [12]. L. H. Harper. Optimal assignment of numbers to vertices. *J. SIAM*, (12):131-135, 1964.
- [13]. N. J. A. Harvey, M. B. Jones, S. Saroiu, M. Theimer, and A. Wolman. Skipnet: A scalable overlay network with practical locality properties. In *Proc. 4th Conference on USENIX Symposium on Internet Technologies and Systems (USITS)*, 2003.
- [14]. S. Hoory, N. Linial, and A. Wigderson. Expander graphs and their applications. *Bulletin of the American Mathematical Society*, 43(4):439-562, 2006.
- [15]. D. Huffman. A method for the construction of minimum-redundancy codes. *Proceedings of the IRE*, 40(9):1098-1101, 1952.
- [16]. D. Knuth. Optimum binary search trees. *Acta informatica*, 1(1):14-25, 1971.
- [17]. K. Mehlhorn. Nearly optimal binary search trees. *Acta Informatica*, 5(4):287-295, 1975.
- [18]. G. Mitchison and R. Durbin. Optimal numberings of an $n \times n$ array. *SIAM J. Algebraic Discrete Methods*, 7(4):571-582, 1986.
- [19]. R. Ravi, A. Agrawal, and P. N. Klein. Ordering problems approximated: Single-processor scheduling and interval graph completion. In *Proc. International Colloquium on Automata, Languages and Programming (ICALP)*, 1991.
- [20]. S. Schmid, C. Avin, C. Scheideler, and B. Haeupler. Brief announcement: Splaynets (towards self-adjusting distributed data structures). In *Proc. 26th International Symposium on Distributed Computing (DISC)*, 2012.
- [21]. A. Sinclair and M. Jerrum. Approximate counting, uniform generation and rapidly mixing markov chains. *Inf. Comput.*, 82(1):93-133, 1989.
- [22]. D. Sleator and R. Tarjan. Self-adjusting binary search trees. *Journal of the ACM (JACM)*, 32(3):652-686, 1985.
- [23]. C. C. Wang, J. Derryberry, and D. D. Sleator. $O(\log \log n)$ -competitive dynamic binary search trees. In *Proc. 17th Annual ACM-SIAM Symposium on Discrete Algorithm (SODA)*, pages 374-383, 2006.