2017 IEEE 11th International Conference on Semantic Computing

Topic Distribution Constant Diameter Overlay Design Algorithm (TD-CD-ODA)

Sina Layazali
Department of Computer Engineering
Cankaya University, Ankara, Turkey
Email: sinalayazali@gmail.com

Kasim Oztoprak
Department of Computer Engineering
KTO Karatay University, Konya, Turkey
Email: kasim.oztoprak@karatay.edu.tr

Erdogan Dogdu
Department of Computer Engineering
TOBB ETU, Ankara, Turkey
Email: erdogandogdu@gmail.com

***Abstract*—Publish/subscribe communication systems, where nodes subscribe to many different topics of interest, are becoming increasingly more common in application domains such as social networks, Internet of Things, etc. Designing overlay networks that connect the nodes subscribed to each distinct topic is hence a fundamental problem in these systems. For scalability and efficiency, it is important to keep the maximum node degree of the overlay in the publish/subscribe system low. Ideally one would like to be able not only to keep the maximum node degree of the overlay low, but also to ensure that the network has low diameter. We address this problem by presenting Topic Distribution Constant Diameter Overlay Design Algorithm (TD-CD-ODA) that achieves a minimal maximum node degree in a low-diameter setting. We have shown experimentally that the algorithm performs well in both targets in comparison to the other overlay design algorithms.**

I. Introduction

Publish/subsribe (pub/sub) is a popular communication paradigm to deliver messages or notifications to users in a distributed environment. With the advances in Internet and Web communication, domain areas such as social networks, Internet of Things, and many related applications all utilize pub/sub communication in some ways [1]. Therefore, designing efficient and scalable pub/sub architecture is still an important and relevant research area.

In publish/subscribe (pub/sub) systems, publishers and subscribers interact in a decoupled fashion. They use logical channels for delivering messages according to the nodes' subscription to the services of interest. Publishers publish their messages through logical channels that deliver the messages to the nodes that subscribed to the respective services. Pub/sub systems can be either topic-based or content-based. In a topic-based pub/sub system, messages are published to "topics", where each topic is uniquely associated with a logical channel. Subscribers in a topic-based system will receive all messages published to the topics to which they subscribe. The publisher is responsible for defining the classes of messages to which subscribers can subscribe. In a content-based system, messages are only delivered to a subscriber if the attributes of those messages match constraints defined by the subscriber; each logical channel is characterized by a subset of these attributes. In publish/subscribe (pub/sub) systems, publishers and subscribers interact in a decoupled fashion. They use logical channels for delivering messages according to the nodes' subscription to the services of interest.

Low node degrees are desirable in practice for scalability and also due to bandwidth constraints. Nodes with a high number of adjacent links will have to manage all these links and the traffic going through each of the links[2][3][4][5][6]. The node degrees and the number of edges required by a topic-connected overlay network will be low if the node subscriptions are well-correlated. In this case, by connecting two nodes with many coincident topics, one can satisfy connectivity of many topics for those two nodes with just one edge. Several recent empirical studies suggest that correlated workloads are indeed common in practice [7][8]. In this paper, we focus on building overlay networks with low maximum node degrees and also keep the diameter low as far as possible. The importance of minimizing the maximum node degree has been recognized in some network domains, such as that of survivable network design [9] and that of establishing connectivity in wireless networks [10]. As in many other systems, a space-time trade-off exists for pub/sub systems: On one hand, one would like the total time taken by a topic-based broadcast (which directly depends on the diameter of each topic-connected sub-network) to be as small as possible; on the other hand, for memory and node bandwidth considerations, one would like to keep the maximum node degree low. Those two measures are often conflicting. Most of the current solutions adopted in practice actually fail at maintaining all mentioned factors that makes an overlay network more scalable. We propose a low diameter algorithm for constructing topic-connected networks with low maximum node degree. In fact, the results show that our algorithm improves the results presented in earlier works.

In this paper we attack the Low-Diameter Topic-Connected Overlay (LD-TCO) problem. The problem is defined as follows: Given a collection of nodes *V*, a set of topics *T*, and the node-interest assignment function *Interest*, connect the nodes in *V* into a topic-connected overlay network *G* that has a low diameter and the least possible maximum node degree.

We present an algorithm Topic Distribution Constant Diameter Overlay Design Algorithm *(TD-CD-ODA)* for LD-TCO problem. *TD-CD-ODA* algorithm is a greedy algorithm which relies on repeatedly adding a star of nodes that would lead to a small number of total edges in the final overlay network. Our experimental results show that our algorithm improves

978-1-5090-4284-5/17 $31.00 © 2017 IEEE
DOI 10.1109/ICSC.2017.88

on the previous algorithms presented in [6] and [11]. We also prove that previously presented algorithms may produce overlay networks of high maximum degree. The approximation ratio on the diameter obtained by the previously presented algorithms may be as bad as $\theta(n)$. As we mentioned earlier, keeping the maximum node degree and the network diameter low are key to the design of scalable and efficient topic-based pub/sub systems. The main contribution of this work is therefore improving the previous results.

The rest of the paper is organized as follows: In section II we present the problem formally, then in section III we present the related work. Section IV presents our Topic Distribution Constant Diameter Overlay Design Algorithm *(TD-CD-ODA)*. Our simulation results on the performance of the *TD-CD-ODA* algorithm are presented in Section V. We conclude the paper, also presenting some future work, in Section VI.

II. Low-Diameter Topic-Connected Overlay Network Problem

Let V be a set of nodes, T be a set of topics. and $n = |V|$, and the interest function *Interest* is defined as $Interest : V \times T \rightarrow \{0, 1\}$. For a node $v \in V$ and topic $t \in T$, $Interest(v, t) = 1$ if and only if node v is subscribed to topic t, and $Interest(v, t) = 0$ otherwise. For the set of nodes V, an overlay network $G(V, E)$ is an undirected graph on the node set V with edge set $E \subset V \times V$. For a topic $t \in T$, let $V_t = \{v \in V | Interest(v, t) = 1\}$ be the set of nodes interested in topic t. Given a topic $t \in T$ and an overlay network $G(V, E)$, the number of topic-connected components of G for topic t is equal to the number of connected components of the subgraph of G induced by V_t. An overlay network G is topic-connected if and only if it has one topic-connected component for each topic $t \in T$. The diameter of a topic is the length of the longest shortest path between the nodes subscribed to this topic. The diameter of a graph is the maximum topic diameter. The degree of a node v in an overlay network $G(V, E)$ is equal to the total number of edges adjacent to v in G.

III. Related Work

In order to design an efficient pub/sub system, an effective publication routing protocol can play a vital role on the system's performance. Therefore, the quality of a constructed overlay can be assessed based on the complexity of the routing scheme applied [6]. If all nodes subscribe to the same topic $m \in M$, the complexity would be intuitively minimized. In this case, the topic can be arranged into a dissemination tree that includes the following features [5] [12]:

- The tree comprises merely the nodes which are interested in topic m,
- The diameter of the tree would be lower.

Greedy Merge (GM) algorithm [5] solves the first issue and the second issue is solved by Low Diameter Publish/Subscribe Overlay Algorithms.

In the existing low diameter algorithms for constructing a topic-based pub/sub overlay network including the number of nodes $n \in N$ interested in topic m, the overlays are constructed in a *decentralized* manner. Hence, the nodes, which are subscribed to similar topics, do not need to depend on other nodes to forward their messages through the network. As mentioned before, most of the approaches failed to decrease the maximum degree and the diameter of the overlay simultaneously. Although GM algorithm is proposed to solve this problem, the algorithm prepares a low maximum-node-degree and requires the lowest number of links to construct an overlay, it suffers from a high rate of diameter and a lack of scalability [5]. The Constant Diameter Overlay Design Algorithm (CD-ODA) [6] establish a topic-based pub/sub network in which each node interested in a variety of topics is assured to have a low diameter to forward its messages. However, the maximum degree resulting from these algorithms is considerably high.

The diameter and maximum degree provided by a network can play a vital role in efficient routing [13][9][10]. Hence, designing a network with a minimum diameter and low maximum degree can noticeably improve the simplicity and performance of the network.

CD-ODA starts with the overlay $G(V, \phi)$, a disconnected graph. At each iteration of CD-ODA, a node u that has the maximum number of neighbours with non-empty interest intersection is chosen. The number of neighbors of node u with shared interests is equal to:

$$n_u = |\{v \in V | \exists t \in T; Interest(v, t) = Interest(u, t) = 1, u \neq v\}| \quad (1)$$

The algorithm then adds an edge between the chosen node and each of its neighbors (with shared interests) to the graph and the topics of shared interests are removed from the set of topics.

In [14], [6] and [11] three new algorithms based on CD-ODA, Constant Diameter Overlay Design Algorithm I (CD-ODA-I), Constant Diameter Overlay Design Algorithm II (CD-ODA II), and 2-Diameter Overlay Design Algorithm (2D-ODA), are presented. All algorithms follow the greedy approach described below.

Constant Diameter Overlay Design Algorithm I (CD-ODA-I):

The algorithm starts with the overlay $G(V, \phi)$. At each iteration of CD-ODA-I, a node u that has the maximum weight neighbors is chosen. The neighbor weight of node z is equal to:

$$W_z = \sum_{t \epsilon T} |\{v \in V | \, Interest(v, t) = Interest(z, t) = 1\}| \quad (2)$$

An edge between u and each of its neighbors is then added, and the topics in the added node¢s interest assignment are removed from the set of topics.

Constant Diameter Overlay Design Algorithm II (CD-ODA-II):

CD-ODA-II starts with the overlay network $G(V, \phi)$. At each iteration of CD-ODA-II, a node u which has the maximum connection density, d_u is chosen. The connection density d_u of a node u, is given by

$$d_u = \frac{\sum_{t \epsilon T} |\{v \in V| \, Int(v,t) = Int(u,t) = 1\}|}{|\{v \in V| \, \exists t \in T, Int(v,t) = Int(u,t) = 1\}|} \qquad (3)$$

Note that $d_u = w_u/n_u$. An edge is added between the node u with maximum node density and each of its neighbors. And the topics in this node's interest assignment are removed from the set of topics.

2-Diameter Overlay Design Algorithm (2D-ODA):

This algorithm is a greedy algorithm and generates a star for each topic and hence each topic-connected component will have a diameter of two [11].

2D-ODA starts with the overlay network $G(V, E)$ and topic set T. At each iteration of the 2D-ODA, first for each node u a topic $t_1 \in T$ and $Int(u, t_1) = 1$ is chosen such that for node u topic t_1 has maximum topic density. The topic density $td(u, t_1)$ of a node u for a topic t_1 where $Int(u, t_1) = 1$ on an overlay network $G(V, E)$ is given by:

$$td(u,t_1) = \frac{\sum_{t \epsilon T} |\{v \in V| \, Int(v,t) = Int(u,t) = 1\}|}{|\{v \in V| \, Int(v,t_1) = Int(u,t_1), (u,v) \notin E\}|} \qquad (4)$$

At each iteration of the 2D-ODA, node u with the maximum topic density for some topic t_1 is chosen, and node u with its star structure for topic t_1 with nodes v where $Int(u, v)$=1 are added to the network. Topic t_1 is then removed from the topic set T.

The overlay networks resulting from [13] and [15] are not required to be topic-connected. In [16], [12] and [17], topic-connected overlay networks are constructed, but they make no attempt to minimize the maximum node degree. The first papers to directly consider node degrees when building topic-connected pub/sub systems were [5] and [6], as we mentioned above. Some of the high level ideas and proof techniques of [5], [6], [14] and [11] have their roots in techniques used for the classical Set-Cover problem. We benefit from some of the ideas in [5], [6], [14] and [11] and also build upon their constructions for Set-Cover.

In this work, we address the problem of minimizing both the maximum degree and achieving a low diameter in topic-connected pub/sub overlay network design. As we mentioned earlier, minimizing both diameter and maximum node degree is important in many network domains, such as survivable network design [9] and establishing connectivity in wireless networks [10].

IV. Topic Distribution Constant Diameter Overlay Design Algorithm

In this section we present our overlay design algorithm *TD-CD-ODA* for the Low-Diameter-TCO problem. *TD-CD-ODA* is a greedy algorithm. Our algorithm generates a star for each topic and hence each topic-connected component will have low diameter. *TD-CD-ODA* starts with the overlay network $G(V, \phi)$ and topic set T. The algorithm tries to reach an overlay network with a minimal diameter and a minimal maximum node degree. At each iteration of *TD-CD-ODA*, a node v that has the minimum topic distribution t_v is chosen. Topic distribution of node u interested in topic t is given in the Algorithm 1 below.

Algorithm 1 TD-CD-ODA

```
Input: V (set of nodes), T (set of topics), Interest(V → T), subscriptionsize
while T is NOT empty do
    For each topic tϵT, n_t ← (the total number of nodes interested in topic t)
    For each node v calculate its topic distribution
        q_v = (Σ_{tϵT} |n_t| for Interest(v,t) = 1) / subscription size
    Find node v with minimum q_v
    Connect node v to other nodes u such that there exists a topic t ∈ T with Interest(v,t) = Interest(u,t) = 1
    Remove all topic t ∈ T such that Interest(v,t) = 1
end while
```

$$q_v = \frac{\sum_{t \epsilon T} |n_t| \;\; for \; Interest(v,t) = 1}{subscription \; size}$$

V. Experimental results

We implemented the GM [5], CD-ODA [6], CD-ODA-I [14], CD-ODA-II [14], 2D-ODA [11], and our TD-CD-ODA algorithms in C++. These algorithms are compared according to the average and maximum node degree in the resulting tables. The diameter is always taken as 2 for the CD-ODA, CD-ODA-I, CD-ODA-II, 2D-ODA, and TD-CD-ODA algorithms. The diameter may be $\theta(n)$, where n is number of nodes, for the GM algorithm. When we compare the results CD-ODA, CD-ODA-I, CD-ODA-II, 2D-ODA, and TD-CD-ODA algorithms, our algorithm TD-CD-ODA algorithm reduces the maximum node degree of the presented overlay networks. Since CD-ODA-II has the best performance in comparison to CD-ODA and CD-ODA-I, we eliminated their results from tables and discussions.

A. Average and Maximum Node Degree with Varying Number of Nodes

For the first experiment, the number of nodes varies between 200 to 400. We fixed the number of topics at 100 and the number of subscriptions at s=10. Topic interests of each node are randomly selected with a uniform distribution. This experimental setting is similar to previous studies in [5], [6], [7], and [11]. Table I compares the resulting average node degrees for the algorithms. The average node degree decreases for GM, varies for CD-ODA II and 2D-ODA, and slightly increases for TD-CD-ODA when the number of nodes increases.

Since more nodes are added throughout the network, the possibility of joining nodes with a higher correlation will increase. Therefore, with a lower number of edges, more nodes are connected and the maximum node degree of the constant diameter algorithms will increase as depicted in Table II.

The most important point about the GM algorithm is that, occasionally a node, which has the highest correlation with

TABLE I
Average Node Degree for Different Number of Nodes

Nodes	GM	CD-ODA II	2D-ODA	TD-CD-ODA
200	6.56	13.58	14.38	17.35
250	6.48	13.56	15.13	17.64
300	6.10	13.06	13.64	18.00
350	6.08	13.02	13.34	17.90
400	5.94	12.93	13.29	17.98

TABLE II
Maximum Node Degree for Different Number of Nodes

Nodes	GM	CD-ODA II	2D-ODA	TD-CD-ODA
200	20	131	128	114
250	21	176	153	146
300	22	209	191	183
350	23	231	240	204
400	25	259	256	240

TABLE III
Average Node Degree for Different Number of Topics

Topics	GM	CD-ODA II	2D-ODA	TD-CD-ODA
200	9.41	14.81	15.02	17.16
250	10.34	14.89	14.82	16.90
300	10.91	14.79	14.49	16.42
350	11.66	14.65	14.21	16.00
400	11.99	14.36	13.90	15.55

TABLE IV
Maximum Node Degree for Different Number of Topics

Topics	GM	CD-ODA II	2D-ODA	TD-CD-ODA
200	30	90	80	62
250	28	74	53	51
300	27	62	49	42
350	32	58	43	39
400	35	51	44	34

many other nodes, will appear in the network. This node increases the maximum node degree considerably as it is connected to many nodes in a star topology (Table II). Although GM algorithm has low maximum and average node degree, but it has higher diameter.

B. Average and Maximum Node Degree with Varying Number of Topics

For the second set of experiments, the number of nodes is kept fixed at 200, the number of topics are varied from 200 to 400, and the number of subscriptions is fixed at s=10. Each node's topic interests are randomly selected with a uniform distribution. This experimental setting is also similar to previous studies [5], [6], [7] and [11]. Table III gives the comparison results for the algorithms for the average degree metric. The average node degree of the graph increases as the number of topics increases for GM and CD-ODA II algorithms, since the edges will have lower correlations and the number of nodes with low correlation increases. And, the average node degree decreases for 2D-ODA and TD-CD-ODA algorithms as the number of topics increases (Table III).

When the number of topics is increased, the overlay will face two different conditions. First, the correlation between nodes will become lower, so more edges will be used to connect the nodes. Second, the number of nodes, which lacks neighbors, will also increase. This has affected the maximum degree of all constant diameter algorithms. As can be seen in Table IV, the maximum node degree of all algorithms are decreasing, except GM algorithm, as the topics set includes more different topics. It shows that GM algorithm is strongly affected by the condition 1. Hence, maximum node degree jumped from 28 to 35 Table IV.

C. Average Node Degree with Varying Subscription Sizes

For the third set of experiments, the number of nodes and the number of topics are fixed at 250 and 100 respectively. The subscription size varies between 15 and 35, and each node's topics are randomly selected with a uniform distribution. Table V gives the comparison results for the algorithms for the average node degree metric. The average node degree of the overlay network decreases when the subscription size increases for all algorithms except TD-CD-ODA since all algorithms can find edges with higher correlation. 2D-ODA improves CD-ODA-II on average as subscription size increased to 35 (Table V). In TD-CD-ODA the average node degree increases as the subscription size increases since we keep the diameter low.

On the other hand, when subscription size is growing, the correlation between nodes will increase. Therefore, the GM algorithm can find many pairs of nodes with shared interest, which consequently reduces the total number of topic connected components dramatically (Table VI). Hence, the maximum node degree decreases as the subscription size increases for GM only. For all other algorithms maximum node

TABLE V
Average Node Degree for Different Subscription Size

Subs.Size	GM	CD-ODA II	2D-ODA	TD-CD-ODA
15	9.63	20.92	23.25	29.49
20	9.32	21.15	23.21	29.81
25	8.92	21.48	24.36	32.07
30	8.61	20.63	22.35	33.63
35	8.22	19.46	20.15	33.98

TABLE VI
Maximum Node Degree for Different Subscription Size

Sub.Size	GM	CD-ODA II	2D-ODA	TD-CD-ODA
15	23	143	137	123
20	21	180	173	166
25	20	197	190	189
30	21	199	199	189
35	17	199	199	190

degree increases as the subscription size increases due to low-diameter requirement. TD-CD-ODA shows a lower maximum node degree in comparison to the other ODA algorithms in all cases (Table VI). This is due to the selection process in TD-CD-ODA, in each step of the iteration a node with the lowest topic interest distribution is chosen. Therefore, TD-CD-ODA algorithm distributes the topic-interested clusters better in comparison to the other ODA algorithms and therefore results in lower maximum node degree.

VI. Conclusions and Future Work

In this paper, we studied the optimization problem (2D-TCO) that constructs a practical and scalable overlay network for publish/subscribe communication with many topics. We presented a topic-connected overlay network design algorithm (TD-CD-ODA) which has lower maximum degree and diameter of 2. We have shown experimentally that our algorithm improves the average node degree compared to other well-known low-diameter algorithms. As for future work, we would like to build upon our TD-CD-ODA algorithm, by formally and experimentally evaluating the hardness of obtaining a topic-connected overlay design algorithm which achieves a "good" trade-off between low diameter, low average, and maximum node degree.

References

[1] Y. Zhao, K. Kim, and N. Venkatasubramanian, "Dynatops: a dynamic topic-based publish/subscribe architecture," in *Proceedings of the 7th ACM international conference on Distributed event-based systems*. ACM, 2013, pp. 75–86.

[2] C. Chen, Y. Tock, H.-A. Jacobsen, and R. Vitenberg, "Weighted overlay design for topic-based publish/subscribe systems on geo-distributed data centers," in *Distributed Computing Systems (ICDCS), 2015 IEEE 35th International Conference on*. IEEE, 2015, pp. 474–485.

[3] C. Chen, H.-A. Jacobsen, and R. Vitenberg, "Divide and conquer algorithms for publish/subscribe overlay design," in *Distributed Computing Systems (ICDCS), 2010 IEEE 30th International Conference on*. IEEE, 2010, pp. 622–633.

[4] C. Chen, R. Vitenberg, and H.-A. Jacobsen, "Constructing fault-tolerant overlay networks for topic-based publish/subscribe," 2013.

[5] G. Chockler, R. Melamed, Y. Tock, and R. Vitenberg, "Constructing scalable overlays for pub-sub with many topics," in *Proceedings of the Twenty-sixth Annual ACM Symposium on Principles of Distributed Computing*, ser. PODC '07. New York, NY, USA: ACM, 2007, pp. 109–118.

[6] M. Onus and A. W. Richa, "Minimum maximum-degree publish-subscribe overlay network design," *IEEE/ACM Transactions on Networking (TON)*, vol. 19, no. 5, pp. 1331–1343, 2011.

[7] H. Liu, V. Ramasubramanian, and E. G. Sirer, "Client behavior and feed characteristics of rss, a publish-subscribe system for web micronews," in *Proceedings of the 5th ACM SIGCOMM conference on Internet Measurement*. USENIX Association, 2005, pp. 3–3.

[8] Y. Tock, N. Naaman, A. Harpaz, and G. Gershinsky, "Hierarchical clustering of message flows in a multicast data dissemination system." in *IASTED PDCS*, vol. 5, 2005.

[9] L. C. Lau, J. S. Naor, M. R. Salavatipour, and M. Singh, "Survivable network design with degree or order constraints," in *Proceedings of the thirty-ninth annual ACM symposium on Theory of computing - STOC '07*, 2007, pp. 651–660.

[10] E. De Santis, F. Grandoni, and A. Panconesi, "Fast low degree connectivity of ad-hoc networks via percolation," in *European Symposium on Algorithms*. Springer, 2007, pp. 206–217.

[11] M. Onus and A. W. Richa, "Parameterized maximum and average degree approximation in topic-based publish-subscribe overlay network design," *Computer Networks*, vol. 94, pp. 307 – 317, 2016.

[12] G. Chockler, R. Melamed, Y. Tock, and R. Vitenberg, "Spidercast: a scalable interest-aware overlay for topic-based pub/sub communication," in *Proceedings of the 2007 inaugural international conference on Distributed event-based systems*. ACM, 2007, pp. 14–25.

[13] R. Baldoni, R. Beraldi, V. Quema, L. Querzoni, and S. Tucci-Piergiovanni, "Tera: Topic-based event routing for peer-to-peer architectures," in *Proceedings of the 2007 Inaugural International Conference on Distributed Event-based Systems*, ser. DEBS '07. New York, NY, USA: ACM, 2007, pp. 2–13.

[14] M. Onus and A. W. Richa, "Parameterized maximum and average degree approximation in topic-based publish-subscribe overlay network design," in *Distributed Computing Systems (ICDCS), 2010 IEEE 30th International Conference on*. IEEE, 2010, pp. 644–652.

[15] R. Chand and P. Felber, "Semantic peer-to-peer overlays for publish/subscribe networks," in *European Conference on Parallel Processing*. Springer, 2005, pp. 1194–1204.

[16] M. Castro, P. Druschel, A.-M. Kermarrec, and A. I. Rowstron, "Scribe: A large-scale and decentralized application-level multicast infrastructure," *IEEE Journal on Selected Areas in communications*, vol. 20, no. 8, pp. 1489–1499, 2002.

[17] S. Voulgaris, E. Rivière, A.-m. Kermarrec, and M. V. Steen, "Sub-2-Sub: Self-Organizing Content-Based Publish and Subscribe for Dynamic and Large Scale Collaborative Networks," p. 16, 2005.