To Split or not to Split? A Simulation Study on the Network Convergence Duration of Multi-Area OSPF

Theofilos Sachinidis, Anastasios C. Politis and Constantinos S. Hilas
Dept. of Computer, Informatics and Telecommunications Engineering
International Hellenic University
Serres, Greece
{theosach,anpol,chilas}@ihu.gr

***Abstract*—Internetworks with a large number of routing devices that utilise the Open Shortest Path First (OSPF) routing protocol, benefit greatly if they are designed in a hierarchical manner. The multi-area feature of OSPF leads to reduced network overhead, low processing and memory requirements and smaller routing tables. This, however, comes at the cost of increased configuration complexity, especially at the Area Border Routers. The threshold after which splitting the OSPF domain into multiple areas, is an ongoing debate among network administrators. A well known rule-of-thumb is recommended by large vendors of networking devices setting that threshold to 50 routing devices. The objective of this preliminary work is to investigate that threshold by conducting a simulation-based study using the OPNET network simulation tool. To assess the network performance, we select the network convergence duration as the quantitative Key Performance Indicator (KPI) of our study. Results indicate that the above mentioned threshold does not appear to be appropriate and a revision of that empirical rule should be considered.**

***Keywords*—multi-area OSPF; network convergence duration; network simulation; OSPF; routing protocol**

I. Introduction

OSPF [1] is an Interior Gateway Protocol (IGP) that constitutes the routing protocol of choice by network administrators. Even though several studies [2], [3], [4] have demonstrated that OSPF exhibits inferior performance as compared to its basic competitor (the Cisco's proprietary Enhanced Interior Gateway Routing Protocol), it is widely deployed due to its non-proprietary specification. It belongs to the class of link-state routing protocols and its most popular version is OSPF version 2 for the IPv4 protocol (OSPF version 3 is available to support the IPv6 protocol).

OSPF provides a multi-area operation which is particularly useful in topologies involving a large number of routing devices or devices with reduced processing capabilities. This feature has certain advantages over a flat OSPF network, including low memory and processing requirements by network devices as well as lower network overhead [5]. However, maintaining a multi-area OSPF network comes at the cost of high configuration complexity [6].

Large vendors of routing devices, such as Cisco, suggest that when an OSPF domain includes more than 50 routing devices, the application of multi-area OSPF should be considered [7]. This threshold is not firm and exists as a recommendation to facilitate the network design and planning phase.

Once the application of multi-area OSPF is decided, the number of routing devices that should be included in a single area is another planning challenge that must be met by network architects. As before, vendor-specific recommendations exist to facilitate such approaches proposing a maximum of 50 routing devices per OSPF area [8]. On the other hand, research papers suggest a more dynamic approach stipulating that number to be equal to $\sqrt{n}$, where n is the number of routing devices in the OSPF domain [6].

In practice, however, network engineers follow a laissez-faire approach based mostly on experience. Hence, it is not uncommon to come across real large OSPF networks that operate in the single-area regime.

In this preliminary work we attempt to investigate the above mentioned design thresholds by performing a series of simulation scenarios. Since memory and processing requirements in modern routing devices does not impose a critical performance limit, we focus on the network convergence duration as the quantitative KPI of our investigation.

The rest of the paper is structured as follows. In Section II, a brief overview of the OSPF protocol and its multi-area functionality is provided. Section III describes the simulation model details that is used in this work and in Section IV the simulation results are presented. The paper is concluded with concluding remarks and future work in Section V.

II. OSPF Background

In this Section we provide a brief overview of the OSPF operation and its ability to operate in a hierarchical mode (multi-area OSPF).

A. OSPF Overview

The purpose of any routing protocol is the construction of the routing table with information about known networks and subnets, network topology, neighbour routers, etc. OSPF achieves that goal by operating in steps:

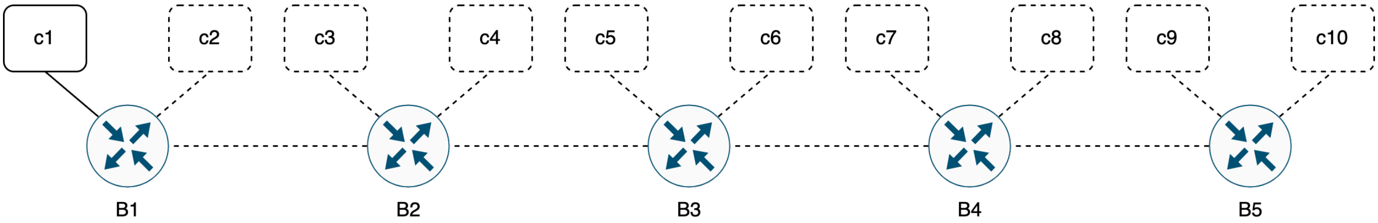

Fig. 1. The network topology for the simulation scenarios.

- *Neighbour Discovery*: Each routing node is responsible to discover neighbouring nodes that are OSPF-capable and belong to the same OSPF area. Neighbouring OSPF routers maintain their adjacency by exchanging special Hello packets periodically.
- *Routing Information Propagation*: A router creates and disseminates advertisements containing local information such as links and their status, associated costs, directly connected subnets, IP addresses etc. These advertisements are called Link State Advertisements (LSAs). A node that receives an LSA from its neighbour must forward it to all its other neighbouring routers by using a procedure known as flooding.
- *Topology Mapping*: Each node stores the LSAs that receives from all other devices in the same OSPF domain in the Link State Database (LSDB). Eventually, every router in the same OSPF area will have an identical LSDB. The contents of LSDB provide a complete map of the OSPF area and its details (link states, subnets etc.).
- Shortest Path Discovery: Based on the LSDB, a device executes the Shortest Path Algorithm (SPF), which is essentially the Dijkstra's algorithm. The outcome is a spanning tree containing the least-cost paths from the device to every subnet served by the OSPF area.
- Routing Table Construction: The paths created by the SPF are inserted in the routing table of each device.

The above procedure is performed during the initialisation phase of the network and leads to the steady-state operation of the network, which is known as convergence.

When a topology change occurs (e.g., a directly connected link fails) the routing device that discovers the deviation, creates a Link State Update (LSU) that is propagated to all devices in the OSPF area by using the flooding procedure. When the transient is absorbed by the system (i.e., all routing tables are updated), the network is re-converged.

B. Multi-Area OSPF

When an OSPF network grows, the convergence and re-convergence durations can be significantly increased due to the flooding process and the time required to execute the SPF algorithm. To avoid such performance degradation, the OSPF specification provides the option to group routing devices in smaller domains (areas). All nodes belonging to the same area exchange detailed LSAs and have identical LSDBs, as described previously. However, the LSAs forwarded from one area to another are less specific with regard to the topology structure of the originating OSPF domain. This prevents the receiving device from re-executing the SPF algorithm and allowing it to directly insert that summary information into its routing table.

Breaking down a network into a multi-area OSPF system involves a configuration procedure which may be complex in large internetworks. This procedure must follow a few rules:

- Each area must have a unique identification number.
- A backbone area (Area 0) must always be present in the design and must be contiguous.
- Non-backbone areas must be connected to the backbone area, by, at least, one Area Border Router (ABR).

The backbone area is responsible for distributing LSAs originating from a non-backbone area to all other OSPF domains.

III. Simulation Model

We created a modular OSPF topology, as shown in Fig. 1, in the OPNET network simulation tool. The general form of this topology includes five backbone routing devices ($B1$ through $B5$) and a series of clusters ($c1$ through $c10$) with each of them comprising of ten routing devices. The structure of the cluster is depicted in Fig. 2. The device indicated as R in the figure represents the demarcation point between the backbone and the cluster.

In every simulation scenario we gradually add one cluster, leading, ultimately, to an OSPF domain with ten clusters of ten routing devices totalling 100 routing devices (excluding backbone routers). All the devices are based on the Cisco 4500 model which is available in the simulator. The model includes six 10BaseT Ethernet ports, which are used to provide connectivity among the routing devices. The simulation details are summarised in Table I.

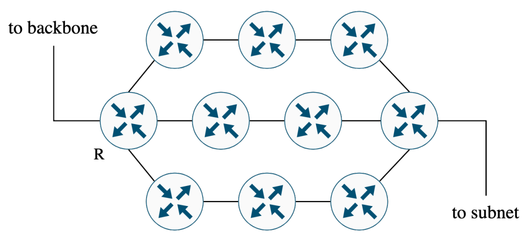

Fig. 2. Structure of a cluster.

TABLE I. SIMULATION DETAILS

Characteristic	Setting
Link-Layer Technology	Ethernet (homogeneous)
Routing Device Model	Cisco 4500
Number of routing devices	Variable
Data Rate	10 Mbps
Simulation Duration	800 sec
Link Failures	1
Link Failure Time Point	500th sec
Measured Statistic	Network Convergence Duration (sec)

At some point during the simulation run we intentionally fail a link in order to produce a topology change and trigger the exchange of LSAs between network devices. Based on that topology change we collect the convergence duration statistic.

A. Remarks on the simulation model

In our simulation model we aimed for network homogeneity, in terms of link-layer technology. To this direction we selected the Ethernet technology throughout the simulated network. The reason for this approach is based on our desire to avoid difficulties in result interpretation that may be caused by LSA fragmentation. A large number of devices leads to larger LSA sizes which might require fragmentation if they exceed the Maximum Transfer Unit (MTU) of the link they need to traverse. Hence, having different link-layer technologies in our network may lead to network fragmentation of the LSAs exchanged between routing devices in some parts of the network while not manifesting at all in others.

It would be, also, preferable to conduct our simulation scenarios by using devices with fast ethernet interfaces in order to provide results representing more realistic topologies. However, investigating the available routing devices models in the simulation tool, we have concluded that there are no available models with more than two fast ethernet interfaces. Hence, in order to construct a homogeneous topology, in terms of available data rate and MTU, the Cisco 4500 router model was selected. Lastly, even though link data rate will impact the measured KPI (i.e., the network convergence duration), it is logical to assume that the conclusions drawn from the KPI measurements can be generalised to include any topology with data rate homogeneity.

IV. SIMULATION RESULTS

As mentioned earlier in this paper, the KPI selected to assess the performance of the OSPF domain, is the network convergence duration. For this metric we collect multiple measurements for different simulation seed values. The mean of these values is used to represent the simulation results.

The network convergence duration versus the number of routing devices included in the network topology, is depicted in Fig. 3. Two settings are visible: one for 10 routing devices per OSPF area and one for 20.

Firstly, it is visible that for a flat OSPF network the convergence duration, after a topology change, is increasing as the number of routing devices climbs up. Grouping the network devices into OSPF areas clearly reduces the convergence duration which seems to be fairly constant for both approaches (10 routers/area and 20 routers/area).

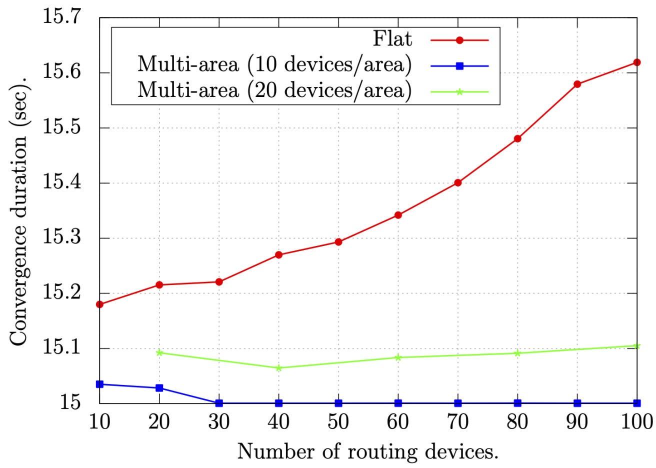

Fig. 3. Convergence duration (sec) versus the number of routing devices for a single-area and a multi-area OSPF network.

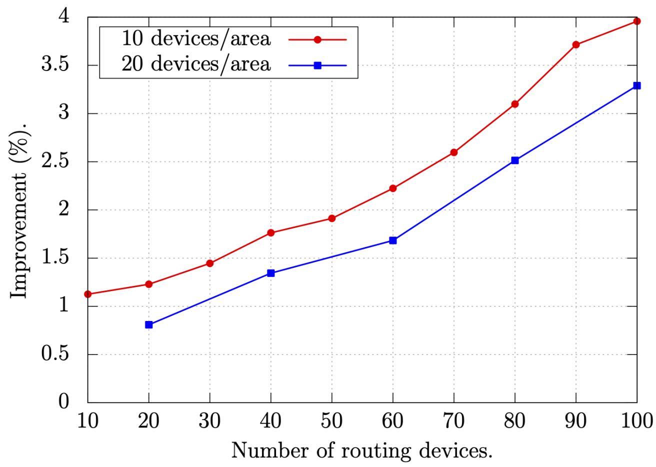

Fig. 4. Convergence duration improvement (%) versus the number of routing devices for a multi-area OSPF network with 10 devices/area and 20 devices/area.

However, as presented in Fig. 4, the improvement on the measured KPI by utilising OSPF areas is clearly small. For example, for a 100 routing devices network (105 including the backbone routers) there is a 4% improvement in the network convergence duration when each area includes 10 routing devices. That improvement becomes even lower when the OSPF areas include 20 routing devices.

From the above observations, we can come to the following conclusions:

- The threshold of 50 routing devices after which the application of multi-area OSPF is recommended does not seem to be confirmed by our simulation study. Based on Fig. 4, the improvement on network convergence duration when a 50-router network is split in multiple areas, barely reaches 2 %. Hence, a revision of that number may be

appropriate.
- Given that an OSPF domain is split in areas, it is best to keep the number of routers per area as small as possible. This however will come at the cost of increased complexity.

V. Conclusion

In this preliminary work, we studied the performance obtained of the multi-area feature of the OSPF routing protocol, in an attempt to confirm or dispute the well known threshold of routing devices after which the application of that feature is recommended. To this direction, its performance is compared with that of a single-area (flat) network by means of simulation. The metric chosen for our study is the convergence duration, since modern routing devices do not impose any limitations regarding memory and/or processing requirements. The results obtained indicate that the suggestion of 50 routing devices as an area-splitting threshold does not seem to be confirmed, since the performance improvement it offers is subtle. Furthermore, based on the current work, it is apparent that when planning a multi-area OSPF network, the number of routing devices per area should be as low as possible, in order to gain the highest performance improvement in terms of convergence duration.

Several aspects of the current study are planned as future work by extending our simulation model. For example, the unequal distribution of network devices in OSPF areas, the simulation of networks with a larger number of links among routing devices in order to consider link redundancy or a rule of thumb for the partitioning of large networks, are some of the issues that are considered as future research.

References

[1] J. Moy, "OSPF Version 2',' STD 54, RFC 2328, doi: 10.17487/RFC2328, April 1998.

[2] A. G. Biradar, "A Comparative Study on Routing Protocols: RIP, OSPF and EIGRP and Their Analysis Using GNS-3," 2020 5th IEEE International Conference on Recent Advances and Innovations in Engineering (ICRAIE), Jaipur, India, 2020, pp. 1-5, doi: 10.1109/ICRAIE51050.2020.9358327.

[3] S. U. Masruroh, K. H. P. Widya, A. Fiade and I. R. Julia, "Performance Evaluation DMVPN Using Routing Protocol RIP, OSPF, And EIGRP," 2018 6th International Conference on Cyber and IT Service Management (CITSM), Parapat, Indonesia, 2018, pp. 1-6, doi: 10.1109/CITSM.2018.8674051.

[4] M. Athira, L. Abrahami and R. G. Sangeetha, "Study on network performance of interior gateway protocols — RIP, EIGRP and OSPF," 2017 International Conference on Nextgen Electronic Technologies: Silicon to Software (ICNETS2), Chennai, India, 2017, pp. 344-348, doi: 10.1109/ICNETS2.2017.8067958.

[5] J. Borthakur, "A comparison study of single area OSPF Network to multiple area OSPF Network implementation in a Campus area Network," 2022 13th International Conference on Computing Communication and Networking Technologies (ICCCNT), Kharagpur, India, 2022, pp. 1-6, doi: 10.1109/ICCCNT54827.2022.9984627.

[6] K. Manousakis and A. J. McAuley, "Using stochastic approximation to design OSPF routing areas that satisfy multiple and diverse end-to-end performance requirements," 2008 6th International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks and Workshops, Berlin, Germany, 2008, pp. 394-402, doi: 10.1109/WIOPT.2008.4586097.

[7] W. Odom, CCNA 200-301 Official Cert Guide, Volume 1. Hoboken, NJ: Cisco Press, 2020.

[8] J. Tiso, Designing Cisco Network Service Architectures (ARCH). Indianapolis, IN: Cisco Press, 2012.