Title	OSPF ベースの AS 内経路最適化に向けたデジタルツインによる経路制御支援
Author(s)	西村, 優典
Citation	
Issue Date	2026-03
Type	Thesis or Dissertation
Text version	author
URL	https://hdl.handle.net/10119/20418
Rights	
Description	Supervisor:宇多　仁, 先端科学技術研究科, 修士(情報科学)

Digital Twin-based Routing Support for OSPF-Based Intra-AS Route Optimization

2410139 Nishimura Yusuke

In large-scale network environments, primarily autonomous systems (AS), automation is achieved using IGPs such as OSPF. In recent years, management has become increasingly complex, requiring high reliability in its operation.

By default, OSPF link costs are calculated based on the bandwidth of each interface. However, in real networks, there are many situations where operators need to reflect their own intentions, such as "bypassing a specific node for maintenance" or "prioritizing a specific link," from a traffic engineering perspective.

Currently, achieving these intentions relies on trial and error based on operator experience, requiring significant effort. Furthermore, changing a single link cost can affect the traffic distribution ratio across the entire network, carrying a significant risk of unintended traffic concentrations or communication outages. It is also difficult to analyze the cause of problems when they occur or to reproduce past conditions for verification. Network simulators and emulators are commonly used as existing verification methods. However, these often rely on fixed configuration information. As a result, there are limitations to faithfully reproducing and predicting real-time, large-scale traffic distribution flowing within an actual network.

To address these challenges, this study proposes the creation of a "digital twin" environment that faithfully recreates an identical network in a virtual space based on dynamic operational data from the actual network. This method enables "what-if analysis," which allows for proactive and quantitative verification of the impact of configuration changes on the entire network without affecting the physical environment. It also accepts abstract operator intent as input and converts it into link costs, aiming to enable verification and prediction of their impact in the digital twin.

The proposed system consists of four independent modules, each with its own function, to reflect the dynamic changes of the real network and realize simulations that reflect operational intent.

The first module, the "Collector Module," uses SNMP to obtain OSPF link state databases (LSDBs) from actual routers. The obtained data is used to generate a network topology based on graph theory. Other methods estimate neighbor relationships by collecting interface information (iftables) and ARP tables from each router. However, our method reconstructs a detailed L3 topology, including neighbor relationships, metrics, and multi-access segments, which these methods cannot capture.

The second module, the "Processor Module," normalizes traffic data obtained from NetFlow/sFlow. Destination and source addresses in the flow information are mapped to nodes within the topology. If an address is not within the topology, it is determined that the communication is with an AS outside the AS, and that address is mapped to an ASBR.

The third module, the "Simulator Module," performs shortest path calculations (SPF) based on NetrkX to reproduce OSPF behavior. It supports traffic superposition and ECMP modeling, calculating the load on each link and recreating the current network state in a virtual space.

The fourth module, the "Visualizer Module," visualizes the results of the calculations made so far as a heat map using logarithmic normalization. This intuitively presents network congestion, potential bottlenecks, and changes in traffic flow over time.

The core result of this research is the implementation of an "intent resolver" that allows operators to achieve their goals without complex cost calculations, and injects threads into the digital twin. Operators only need to declaratively describe abstract threads, such as "what state they want to achieve," rather than "how to configure it."

This system defines six types of operational intents to be resolved. First, as intents related to structural constraints, we implemented "Node Drain" for hosting and "Link Blocking" for physical disconnection. Next, as intents related to route guidance, we implemented "Low Latency Guidance," which prioritizes specific paths while minimizing side effects, and "Load Balancing," which automatically adjusts cumulative costs to establish ECMP. Furthermore, we implemented "Flow Reduction," which gradually adjusts flow rates, and "Safety Range Analysis," which calculates the boundary value below which configuration changes do not affect other routes. These algorithms apply the theory of sensitivity analysis in shortest path problems, enabling the calculation of configuration values ⊠⊠ based on mathematical evidence.

We also focused on the order in which the derived link costs are set to ensure the safety of network operations. If cost changes are applied in the wrong order, temporary routing loops, black holes, and other failures may occur during OSPF convergence. This research proposes the automatic generation of "safe sequences" to minimize these risks. First, a route to which traffic can escape is established by performing a cost-decreasing operation. After that, traffic is removed from the original route by performing a cost-increasing operation. In this way, the sequence is determined automatically. This maximizes availability during large-scale configuration change processes and reduces the risk of service interruptions caused by human error.

To verify the effectiveness of the proposed method, we evaluated it using actual data from the WIDE project's backbone network. In verifying the comprehensiveness of topology construction, the proposed LSDB analysis method identified 45 routers and 116 links. This was a significantly higher detection rate than conventional CLI-based and config-based methods. This demonstrated the proposed method's ability to accurately capture configuration errors and dynamic segment changes. The traffic estimation accuracy was compared with actual measurements from the Zabbix operations monitoring system. We confirmed extremely high reproducibility, with an error of 0.08 % at edge nodes. For backbone links, we observed an underestimation of approximately 30 % due to sampling bias in flow information and L2 overhead. However, the relative load relationships were consistent with actual measurements, demonstrating their usefulness for the practical purpose of identifying congestion points.

Furthermore, we confirmed the usefulness of the intent resolver using multiple intents created using data from the WIDE project. In the "load balancing" scenario, we successfully split approximately 963.5 Mbps of traffic concentrated on a specific path equally between two paths, each at 481.76 Mbps, through automatic cost adjustment. In the "flow reduction" scenario, we also identified an optimal intermediate cost value through sensitivity analysis that diverted only bursty traffic while maintaining nearby reachability. These results demonstrate that the proposed method abstracts complex metric calculations and dramatically improves operational efficiency and safety.

This research has demonstrated that intent-based control using digital twins will be an extremely effective foundation for next-generation network operation automation. In the future, we plan to introduce a statistical complementary model for traffic sampling, integrate it with other routing protocols, and expand it to multi-area OSPF.