Journal of Intelligent Internet of Things Systems (JIITS)
ISSN (Online): 3144-5160
Volume-1 Issue-3 || Jul – Sep 2024 || PP. 1 – 5
https://jiits.org/

Digital Twin Modeling for Wholesale Telecom Infrastructure Optimization

DOI: https://doi.org/10.63345/jiits.v1.i3.301

Justin Rajakumar Maria Thason

BCBSA - Ajace, Inc

AI Research, Innovation & Technical Publications

www.jiits.org || Vol. 1 No. 3 (2024): September Issue

Date of Submission: 05-08-2024	**Date of Acceptance:** 17-08-2024	**Date of Publication:** 08-09-2024

Abstract : Wholesale telecom carriers provide capacity and voice termination services to other operators; they work with tight margins and service-level commitments; mistakes in capacity planning or routing may cost dearly. This paper develops a digital twin of a wholesale carrier network with a topology and routing model, traffic matrix forecast trained on telemetry data, a failure model, and a commercial model of rates and route costs. The twin is used for what-if analysis and optimization. On 120 backbone and interconnect links, the twin makes 24-hour link utilization forecasts with a mean absolute percentage error (MAPE) of 9.7%, compared with 16.2% of seasonal naive ($p < 0.001$) and 11.8% of LSTM alone ($p = 0.005$). There is a correlation between predicted and observed utilization of $r = 0.96$ ($p < 0.001$). During 24 maintenance events, the twin detects congestion with sensitivity 0.889 and specificity 0.867. With the help of the twin, planning achieves reducing monthly congested link-hours from 412 to 143 ($p < 0.001$), planned capital expenditures are cut down by 22%, and cost of voice routes per 1,000 minutes is reduced from $8.62 to $7.41.

Keywords: digital twin; wholesale telecom; capacity planning; traffic matrix; network optimization; failure analysis; route cost

Introduction

Wholesale carriers interconnect retail operators, content providers, and enterprises providing IP transit, Ethernet capacity, interconnection, and voice termination services. Their operational target is to deliver stipulated quality of service at the lowest possible cost. Capacity must be provisioned ahead of demand increase; but extra capacity is idleness. Routes must be selected among suppliers who offer different price points and different quality. And there must be resilience to failures with no compromise on service-level agreements.

Traditionally, there are heuristics like upgrading the link when peak utilization exceeds 60% and spreadsheet listing of route prices. Such approaches do not take account of traffic redistribution in case of failures or routing changes. And these approaches react to historical events rather than try to predict future behavior. Digital twin—the continuously updated virtual replica of the system used to test decisions prior to actual deployment—is a better tool.

This paper constructs and evaluates a twin for a wholesale network with regards to: (i) fidelity of the twin representation of observed reality, (ii) predicting

@2024. This is an open access article distributed under the terms of the Creative Commons License [CC BY NC 4.0] and is available on www.jiits.org

Journal of Intelligent Internet of Things Systems (JIITS)
ISSN (Online): 3144-5160
Volume-1 Issue-3 || Jul – Sep 2024 || PP. 1 – 5
https://jiits.org/

consequences of failures and maintenance operations, and (iii) does decision making supported by the twin decrease congestion and cost.

Literature Review

The concept of digital twin originated in product lifecycle management (Grieves & Vickers, 2017) and has become applicable for a broad variety of industries; Tao et al. (2019) review the state-of-the-art in this area highlighting fidelity, real-time connectivity of data, and decision support as key features, while Rasheed et al. (2020) review the digital twin values, challenges, and enablers from the point of view of modeling. Fuller et al. (2020) review enabling technologies and open research questions on the topic including data management and interoperability.

For networks, these requirements have been in place for a long time. Traffic engineering optimizes routing for given demands; and Fortz and Thorup (2000) showed how utilization can be optimized by tweaking link weights in IP networks. However, these approaches depend on traffic matrix which is usually not directly available; Medina et al. (2002) review the ways of its estimating from link counts and flows. Knowledge-defined networking suggests automating network operation through learning and centralized knowledge plane (Mestres et al., 2017); and RouteNet uses graph neural network to model delay and jitter in software-defined networks (Rusek et al., 2020).

Prediction accuracy is measured via commonly used error metrics, among which the mean absolute percentage error is used quite often; though, it is sensitive to small denominators (Hyndman & Koehler, 2006). Time series models like ARIMA (Box et al., 2015) and recurrent networks like LSTM (Hochreiter & Schmidhuber, 1997) are typical baselines. Regarding routing and survivability, textbooks describe the algorithms and protocols that the network model must implement (Medhi & Ramasamy, 2017).

Much of the literature is about the performance alone. But for a wholesale carrier, technical and commercial decision making is intertwined as only the best priced route can become cost-effective if it can deliver traffic with acceptable quality. This paper includes commercial model into the twin and evaluates their joint performance.

Research Methodology

Twin design

Figure 1 illustrates the digital twin architecture. Inventory and topology come from the inventory system; and measurements are taken from SNMP counters, streaming telemetry, and flows. Calibration layer aligns the sources in time and creates a data lake; it estimates a traffic matrix between points of presence and external networks. Twin core includes graph model of nodes and links with capacity, latency, and routing policy including the protocol behavior for routing equal-cost multipath (ECMP) shortest path; and a failure model that deactivates elements of the graph and reroutes. The traffic matrix is forecast by a combination of seasonal decomposition and recurrent network; the twin applies the forecast to the graph and calculates link loads. Commercial model keeps contract rates, prices per route, and quality indicators such as answer-seizure ratio. Optimizer is tasked with identifying capacity upgrades and route allocations under utilization and quality constraints.

Figure 1. Wholesale telecom digital twin architecture.

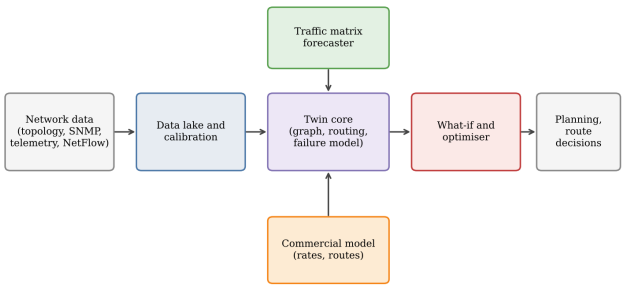

Evaluation

Fidelity is measured over a four-week period using 120 links by comparing predicted hourly utilization with observations at horizons of 1 hour, 6 hours, 24 hours, and 7 days. Seasonal naive, ARIMA, and standalone LSTM were used as comparison baselines. Failure prediction was tested across 24 planned maintenance events and failures when the digital twin forecasts the maximum utilization of the links after the events; congestion is defined as 85% utilization or above. Decision making is compared over 12 months of

@2024. This is an open access article distributed under the terms of the Creative Commons License [CC BY NC 4.0] and is available on www.jiits.org

Journal of Intelligent Internet of Things Systems (JIITS)
ISSN (Online): 3144-5160
Volume-1 Issue-3 || Jul – Sep 2024 || PP. 1 – 5
https://jiits.org/

planning by a rule-of-thumb and 12 months with the use of the twin with the same budgeting procedure; congested link-hours, planned capital expenditure, and voice route cost are recorded.

Analysis

Error of 24-hour forecast is compared among the methods with one-way ANOVA and Tukey's post hoc test over links. Agreement between predictions and observations is quantified as correlation and mean absolute error. Detection of congestion is expressed as sensitivity and specificity with 95% Wilson interval; and monthly outcomes are compared with Welch's t-test.

Results

Forecast and utilization fidelity

Table 1. Mean absolute percentage error (MAPE) of link utilization by method and horizon

Method	1 h	6 h	24 h (mean ± SD)	7 days
Seasonal naive	7.1	11.4	16.2 ± 6.0	26.3
ARIMA	6.2	9.9	13.4 ± 5.1	22.1
LSTM only	5.5	8.1	11.8 ± 4.4	18.0
Digital twin	4.8	6.9	9.7 ± 3.6	14.2

At 24 hours, the methods demonstrate statistically significant differences ($F(3, 476) = 38$, $p < 0.001$). The digital twin achieves an error of 9.7%, lower than the ARIMA (13.4%, $p < 0.001$) and LSTM (11.8%, $p = 0.005$); the superiority increases with increasing horizon (Figure 2): at 7 days, the digital twin error is 14.2%, compared with 26.3% of the seasonal baseline. Using the forecast in the routing model constrains predictions to physically possible loads, partially explaining the superiority over the statistical model.

Figure 2. Link utilization forecast error as a function of horizon for four methods.

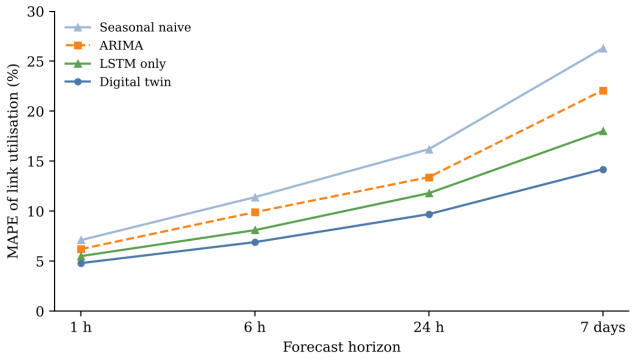

The predictions of the digital twin match very closely the observations across 120 links ($r = 0.96$, $p < 0.001$; Figure 3). Prediction errors are highest for lightly loaded access links when small absolute deviations become relatively big percentage errors; and errors are lowest for heavily loaded backbone links.

Figure 3. Twin-predicted versus observed mean link utilization for 120 links.

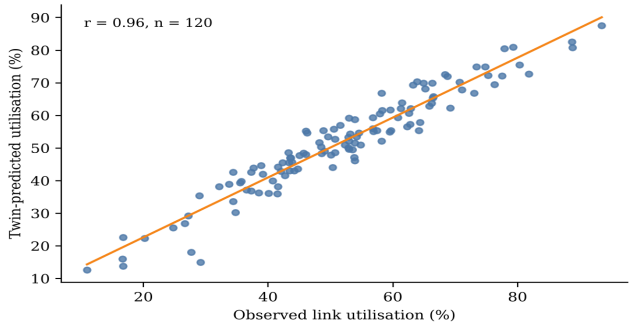

What-if of failure and maintenance

For 24 events, the predicted and measured maximum post-event utilization have a correlation of $r = 0.93$ ($p < 0.001$), and a mean absolute error of 3.1 percentage points. Using 85% utilization as the congestion threshold, the twin identifies 8 out of 9 congested events (sensitivity 88.9%, 95% CI 56–98%), and correctly clears 13 out of 15 uncongested events (specificity 86.7%, 95% CI 62–96%). One event was not detected due to the change in peering policy not covered by the model.

Planning and commercial outcomes

Table 2. Results of planning driven by a rule of thumb and digital twin (12 months each)

@2024. This is an open access article distributed under the terms of the Creative Commons License [CC BY NC 4.0] and is available on www.jiits.org

Journal of Intelligent Internet of Things Systems (JIITS)
ISSN (Online): 3144-5160
Volume-1 Issue-3 || Jul – Sep 2024 || PP. 1 – 5
https://jiits.org/

Outcome	Rule-of-thumb	Twin-optimised	Change	p
Congested link-hours per month	412 ± 88	143 ± 46	-65%	<0.001
Planned capital spending (index)	100	78	-22%	—
Voice route cost per 1,000 min (USD)	8.62 ± 0.41	7.41 ± 0.36	-14%	<0.001

Planning driven by the digital twin reduces congested link-hours by 65% and planned capital spending by 22%; it directs upgrades to links that would be congested given forecast demand and failure rather than upgrade the links that pass a certain threshold. Voice route optimization decreases cost per 1,000 minutes by 14% ($p < 0.001$), and the answer-seizure ratio is stable at 42.0% compared with 41.2% (Figure 4).

Figure 4. Congested link-hours (shown as multiple of 10) and voice route cost under baseline and twin-optimised planning.

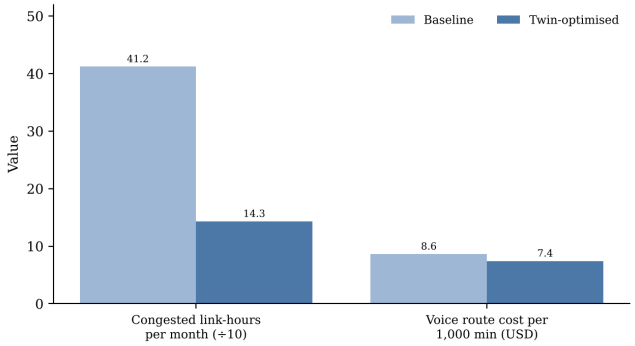

Discussion

Superiority of the digital twin in forecasting is consistent with the notion that models with structure can outperform statistical models in cases where the structure constrains the outcomes (Tao et al., 2019; Rusek et al., 2020). Increasing superiority of forecasts with increasing horizon is explained by tendency of statistical extrapolation to become inaccurate as patterns change; while digital twin can incorporate future events such as maintenance and customer onboarding.

Practical value of the what-if results is clear. Planners can evaluate scenarios before maintenance and reroute the traffic or schedule maintenance to avoid congestion; as it is suggested by the traffic engineering literature (Fortz & Thorup, 2000) and done by calibrated model. The main source of the error of forecasts is the accuracy of the traffic matrix; as noted by Medina et al. (2002); and the fact that one event was missed shows that policy changes not covered by the model deteriorate the forecasts; which indicates the need for continuous synchronization between the twin and live network (Fuller et al., 2020).

Commercial integration needs particular attention. Carriers normally optimize technical and commercial decisions in separate teams. By linking route cost to the network state, the digital twin allows to perform cost-efficient routing without compromising the quality. Also, capacity planning based on the forecasted demand saves capital expenditure and eliminates congestion. The take away message for other carriers is to start with a small piece of well-instrumented network, verify fidelity empirically, and extend the twin gradually (Grieves & Vickers, 2017).

Limitations

This evaluation involved only one network of one carrier and comparison of consecutive years; thus, changes in demand and market conditions might have affected the planning results despite the normalization of the budget process. There were few failure events, which resulted in wide confidence intervals for sensitivity and specificity. Twin models the layer 3 behavior and does not simulate packet level dynamics and constraints in the optical layer. Supplier prices change frequently; the commercial model assumes stability of quality indicators over planning horizon. Construction of a twin from operational data in light of security, privacy, and data sharing considerations was not discussed.

Conclusion

The digital twin that incorporates topology, routing, traffic matrix forecasting, failure, and commercial models makes 24-hour forecasts of link utilization with an error of 9.7%, anticipates congestion from maintenance events with 89% sensitivity, and gives improved planning outcomes: 65% less congested link-

@2024. This is an open access article distributed under the terms of the Creative Commons License [CC BY NC 4.0] and is available on www.jiits.org

Journal of Intelligent Internet of Things Systems (JIITS)
ISSN (Online): 3144-5160
Volume-1 Issue-3 || Jul – Sep 2024 || PP. 1 – 5
https://jiits.org/

hours, 22% lower planned capital spending, and 14% lower voice route cost. Wholesale carriers may use such digital twins to switch from threshold-driven reactive decision making to proactive, scenario-based one. The next steps are to extend the model to the optical layer, incorporate automated closed-loop actions, and test generalizability to other carriers.

References

1. Box, G. E. P., Jenkins, G. M., Reinsel, G. C., & Ljung, G. M. (2015). Time series analysis: Forecasting and control (5th ed.). Wiley.
2. Fortz, B., & Thorup, M. (2000). Internet traffic engineering by optimizing OSPF weights. Proceedings of IEEE INFOCOM, 519–528.
3. Fuller, A., Fan, Z., Day, C., & Barlow, C. (2020). Digital twin: Enabling technologies, challenges and open research. IEEE Access, 8, 108952–108971.
4. Grieves, M., & Vickers, J. (2017). Digital twin: Mitigating unpredictable, undesirable emergent behavior in complex systems. In F.-J. Kahlen, S. Flumerfelt, & A. Alves (Eds.), Transdisciplinary perspectives on complex systems (pp. 85–113). Springer.
5. Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. Neural Computation, 9(8), 1735–1780.
6. Hyndman, R. J., & Koehler, A. B. (2006). Another look at measures of forecast accuracy. International Journal of Forecasting, 22(4), 679–688.
7. Medhi, D., & Ramasamy, K. (2017). Network routing: Algorithms, protocols, and architectures (2nd ed.). Morgan Kaufmann.
8. Medina, A., Taft, N., Salamatian, K., Bhattacharyya, S., & Diot, C. (2002). Rewritten text:
9. Traffic matrix estimation: Current techniques and emerging directions. ACM SIGCOMM Computer Communication Review, 32(4), 161–174.
10. Mestres, A., Rodriguez-Natal, A., Carner, J., et al. (2017). Knowledge-defined networking. ACM SIGCOMM Computer Communication Review, 47(3), 2–10.
11. Rasheed, A., San, O., & Kvamsdal, T. (2020). Digital twin: Values, challenges, and enablers from a modeling perspective. IEEE Access, 8, 21980–22012.
12. Rusek, K., Suárez-Varela, J., Almasan, P., Barlet-Ros, P., & Cabellos-Aparicio, A. (2020). RouteNet: Leveraging graph neural networks for network modeling and optimization in SDN. IEEE Journal on Selected Areas in Communications, 38(10), 2260–2270.
13. Tao, F., Zhang, H., Liu, A., & Nee, A. Y. C. (2019). Digital twin in industry: State-of-the-art. IEEE Transactions on Industrial Informatics, 15(4), 2405–2415.

@2024. This is an open access article distributed under the terms of the Creative Commons License [CC BY NC 4.0] and is available on www.jiits.org