Enhancing Privacy in Multi-Domain Network Intent Negotiation*

Pedro Martinez-Julia[1], Ved P. Kafle[1], Hitoshi Asaeda[1], Diego Lopez[2], and Antonio Skarmeta[3†]

[1] National Institute of Information and Communications, Technology(NICT), Tokyo, Japan
{pedro,kafle,asaeda}@nict.go.jp
[2] Telefonica I+D, Madrid, Spain
diego.r.lopez@telefonica.com
[3] University of Murcia, Murcia, Spain
skarmeta@um.es

Abstract

In this study, we analyzed the quantity of sensitive knowledge present in data exchanges between agents of multi-domain network intent negotiations. Using information theory, we constructed a model that maps the set of characteristics that define negotiation operations to a value that represents the quantity of sensitive knowledge shared in relation to the number of negotiation goals achieved. We used the model to find a set of characteristic boundaries for which a negotiation process will produce the optimum relation of goals attained per each sensitive knowledge item shared. We validated our model by constructing a negotiation process with such characteristics, and demonstrating it attains the goals while sharing 20 % less knowledge than previous state-of-the-art systems.

1 Introduction

Intent-based networking (IBN) proposes and promotes the specification of network services using natural language or a high-level formal language. These specifications are known as network intents [1]. For network services to be created and operated according to provided network intents, these must be first translated into a totally formal and computer-friendly structure called network service description (NSD).

The NSDs resulting from the resolution of network intents are then managed by an IBN platform. Most common platforms supporting IBN are network function virtualization (NFV) platforms, such as pen Source MANO (OSM) [2] and OSDM [3]. These understand the NSD and deploy the required components to construct the intended network services.

A special case of network intents and NSDs are multi-domain network intents, as discussed in [4], [5], and [6]. These impose the requirement that multiple domains must be involved in the management of both the network intent life-cycle and the network service life-cycle. In this study, we studied the former, namely the life-cycle of a multi-domain network intent, to determine the amount of knowledge that is shared among the domains involved in the network intent resolution process.

We analyzed the message exchanges of the network intent resolution process found in state-of-the-art systems for IBN [7]. We constructed a model that maps the set of characteristics that define such process to a value that represents the quantity of sensitive knowledge shared. We then designed a new system for multi-domain network intent resolution that minimized the amount of shared knowledge according to the model. We validated both the

*Proceedings of the 9th International Conference on Mobile Internet Security (MobiSec'25), Article No. S4, December 16-18, 2025, Sapporo , Japan.
†Corresponding author

model and system in an experimentation platform, showing that the system we propose shares 20 % less knowledge than state-of-the-art systems. This contributes to improving the overall privacy in IBN systems.

In summary, the contributions of this study are as follows:

- Analysis of the multi-domain network intent resolution process present in state-of-the-art IBN systems.
- Construction of a model that obtains the quantity of knowledge shared by the elements involved n the multi-domain resolution process.
- Design of a new system that minimizes the amount of knowledge shared for network intent resolution while retaining the resolution accuracy.

The remainder of this paper is organized as follows. First, in Section 2 we will contextualize our study and will discuss related work. Then, in Section 3 we will discuss the model we constructed and the system we designed, and in Section 4 we will discuss how we evaluated them to demonstrate their properties and benefits. Finally, in Section 5 we will discuss our conclusion and introduce some indications for future work.

2 Background

IBN enables the construction and management of network services from specifications formulated in a high-level formal language or even in a natural language. Resolving multi-domain network intents requires the cooperation of multiple elements from the domains involved in the later deployment process. The process requires the elements to share knowledge, which exposes privacy concerns. This conducted us to the problem we tackled in this study, as described below.

2.1 Problem Statement

In this study, we investigated the effect that the message exchanges of the systems used for resolving multi-domain network intents had on the total amount of knowledge shared among them. Our goals were:

- Formulating a model that maps the qualities of the message exchanges of a multi-domain network intent resolution system to the total amount of knowledge shared in the resolution process.
- Designing a new system whose message exchanges are defined in a way that the model output is minimized.

2.2 Related Work

In general, previous work in multi-domain network intent resolution did not properly study the amount of knowledge shared by the elements involved in the process. Nevertheless, the resolution process was well specified by many previous works, so we covered them to compliment our analysis.

First, we found a survey on the state-of-the-art IBN systems in [8], extending a previous survey presented in [9]. They outline the little coverage of multi-domain network intent resolution, but also remark the structure of the resolution processes found in current systems. Moreover, the work presented in [10] organized multi-domain network intent requirements in hierarchies and delegated their resolution to third parties. The work presented in [11] proposes an intent-based system for orchestrating networks involving multiple

vertical domains—access, transport, and core. It does not rely on a multi-element or multi-agent system. A central element retrieves knowledge from all domains and resolves the intents locally.

Additionally, the work presented in [7] covers multi-domain network intent resolution as a side aspect of their IBN-based network service life-cycle management. It follows a similar approach to the works discussed above, with an element-to-element communication-based negotiation. We used this common process in our analysis to define the reference process and its message exchanges. Other approaches proposed to map the intent resolution problem to other problem resolution and/or inference structures, For instance, the work presented in [12] adopted the virtual network function˜(VNF) placement problem to assign resources to network intents. The work presented in [13] proposed to use hybrid ensemble learning algorithms for network intent resolution in relation to their resource utilization.

Some solutions proposed to address the complexity of multi-domain network intents through the use of declarative programming environments and reasoning engines, such as Prolog [14] [14] and Haskell [15]. However, they do not specify any particular process for communication among elements, so we assumed they relied on the state-of-the-art systems discussed above.

3 Knowledge Sharing for Multi-Domain Network Intent Negotiation

Our approach to resolve the problem introduced above consists of two aspects. On the one hand, we analyzed the sensitive information present in message exchanges done by previous state-of-the-art systems for multi-domain network intent negotiation and constructed a representative model. On the other hand, we constructed a new system that improves privacy by minimizing the amount of sensitive information exchanged according to the model.

3.1 Message Exchanges

Negotiating the deployment of multi-domain network intents involves the owner of the network service that is being deployed—namely, network tenant—, the user interface (UI) of the IBN platform, the intent manager (IM), and the agents that represent each domain in the negotiation. The process is as follows. First, the tenant inputs into the UI a network intent in plain text format, using natural language or a high-level formal language. Second, the UI constructs a knowledge object (KO) that represents the network intent. Each knowledge item present in the KO corresponds to a trivially identified sentence, formal block, etc. The UI does not perform any other analysis.

Third, the UI requests the IM to deploy the intent by providing the network intent knowledge object (NIKO) it constructed before. Fourth, the IM responds to the UI that the NIKO has been registered and proceeds to resolve the intent by interacting with the resolution agents. The first step of this sub-process consists of a message from the IM to the first agent—agent A—requesting to resolve the NIKO. Agent A executes a reasoning step, which consists of the following sub-steps:

1. Adding all elements from the NIKO to a structure named $\mathrm{RKO_A}$.
2. For all item a in $\mathrm{RKO_A}$, item b in $(\mathrm{RKO_A} \cup \mathrm{KB_A})$, and expression $(a \wedge b \implies r)$ in $(\mathrm{RKO_A} \cup \mathrm{KB_A})$, remove a and b from $\mathrm{RKO_A}$, add r to $\mathrm{RKO_A}$.
3. Send the resulting $\mathrm{RKO_A}$ to agent B.

Upon reception, agent B follows the same reasoning sub-steps, considering the received $\mathrm{RKO_A}$ as the input NIKO and sending the resulting $\mathrm{RKO_B}$ to agent C. The process

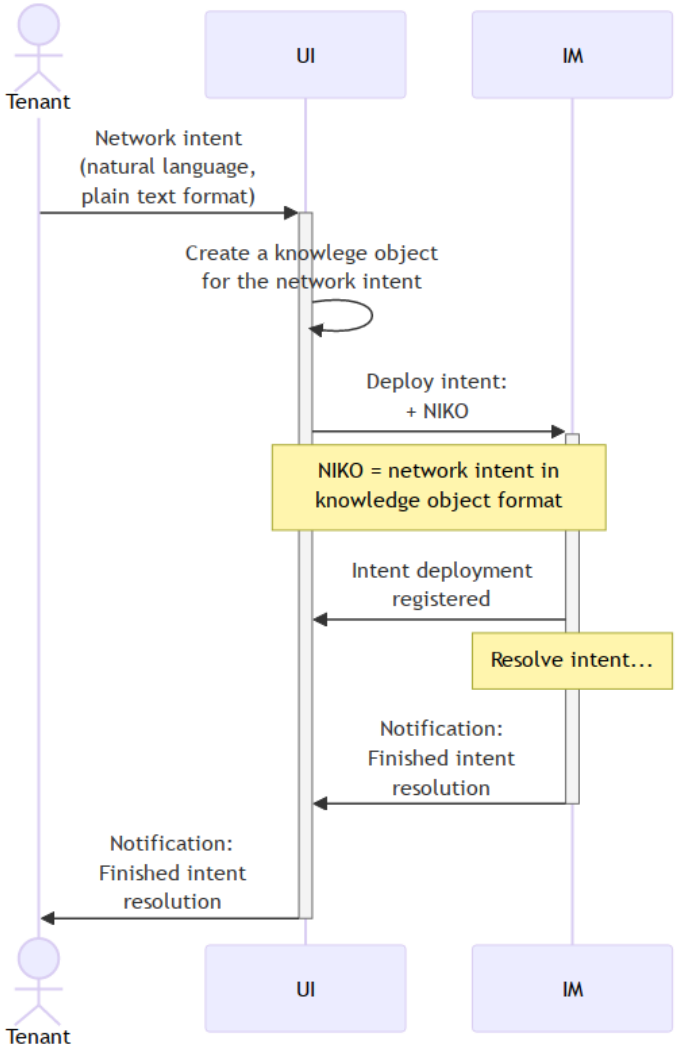

Figure 1: Message Sequence of a Network Intent Deployment Request

continues through all agents. When a final agent finds that there are no further agents to send request to—which means that either the intent has been fully resolved or that all agents have been questioned and the intent cannot be fully resolved—a response with the final KO, represented as RKO_C is sent back agent-by-agent until it reaches the IM. Finally, the IM registers the result, extracts the network service definition (NSD) from it, and notifies the tenant through the UI that the intent has been resolved. The NSD is forwarded to the element of the IBN platform that is responsible of its deployment—the NSD enforcer.

The interactions between the network tenant and the IM, using the UI as intermediate, process are illustrated in Figure 1. The interactions involving the IM and the agents are illustrated in Figure 2.

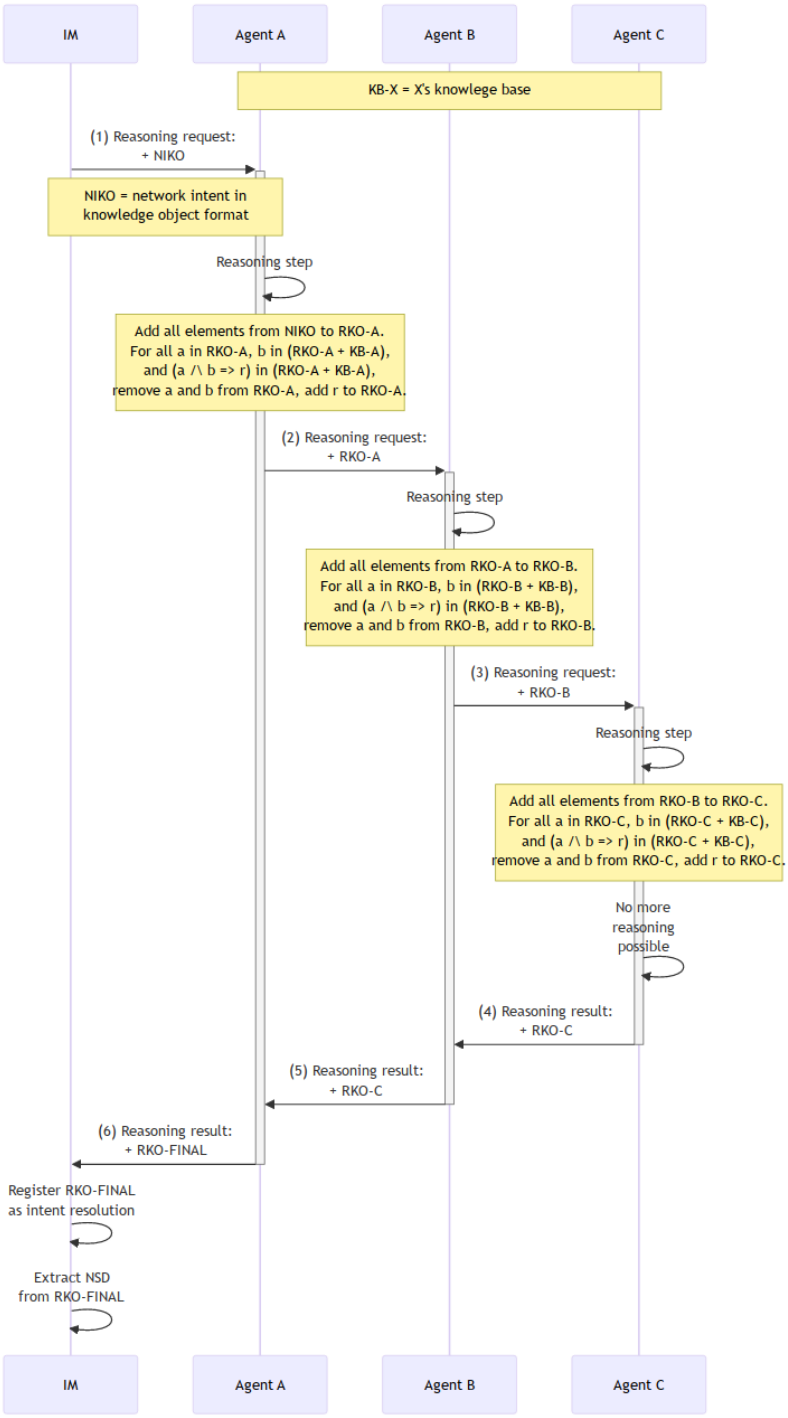

Figure 2: Message Sequence of the Resolution of a Network Intent Across Multiple Domains

3.2 Model Formulation and Analysis

Our goal was to minimize the amount of information shared among agents while attaining the resolution objective. We decided to construct a model based on information theory that obtains such amount from the interactions—message exchanges—described above.

For each agent A_j, we obtained models for the quantity of information Q_j^x of the input NIKO ($x = \text{in}$), the output NIKO ($x = \text{out}$), the newly added items ($x = \text{new}$), and net shared ($x = \text{sha}$). We constructed the models in base of the following considerations:

- The input NIKO of agent A_j is m_j, the output NIKO of agent A_j is m'_j, which is also the input NIKO of agent A_{j+1}.
- Each m_j contains the set of knowledge items present in the NIKO of the request received by agent A_j.
- Each m'_j contains the set of knowledge items present in the NIKO of the request sent by agent A_j to agent A'_j.
- When agent A_{j+1} exists, $m'_j = m_{j+1}$.
- The knowledge base of agent A_j is represented by K_j.
- The agents are sorted by the position in the message exchange sequence shown in Figure 2—the operation of agent A_{j+1} comes after the operation of agent A_j.

The resulting models are as follows:

$$Q_j^{\text{in}} = \sum_{p \in m_j} I(p) - \sum_{\substack{p,\ q \in m_j \\ p \neq q}} I(p;q)$$

$$Q_j^{\text{out}} = \sum_{p \in m'_j} I(p) - \sum_{\substack{p,\ q \in m'_j \\ p \neq q}} I(p;q)$$

$$Q_j^{\text{new}} = Q_j^{\text{out}} - Q_j^{\text{in}}$$

$$Q_j^{\text{new}} = \sum_{p \in m'_j \cap m_j^{\complement}} I(p) - \sum_{\substack{p,\ q \in m'_j \cap m_j^{\complement} \\ p \neq q}} I(p;q)$$

$$\begin{aligned} Q_j^{\text{sha}} = & \sum_{\substack{p,\ q \in m_j \\ r \notin m_j \\ r \in m'_j \\ (p \wedge q \implies r) \in K_j}} (I(r) - I(r;p) - I(r;q)) + \dots \\ \dots + & \sum_{\substack{p \in m_j \\ q,\ r \notin m_j \\ r \in m'_j \\ (p \wedge q \implies r),\ q \in K_j}} (I(r) - I(r;p) + I(r;q)) \end{aligned}$$

Assuming that the multi-domain system is able to fully resolve all input network intents, we formulated that the probability of $(p \wedge q \implies r) \in K_j$ for any $p,\ q \in m_j$ is less than the probability of $(p \wedge q \implies r') \in K_{j+1}$ for $p,\ q \in m'_j$, which are the same $p,\ q$ present in the request received by agent A_j, now received by agent A_{j+1}. Note that $Q_j^{\text{sha}} = 0$ if, for all r, we have that $r \in m_j$.

We applied such probabilities to the information quantities presented above. Thus, we had that:

$$Q_j^{\text{sha}} \leq Q_{j+1}^{\text{sha}}$$

$$Q^{\text{sha}} = \sum_{j \in \text{ agents}} Q_j^{\text{sha}}$$

From these expressions, we obtained a key outcome: The later and farther a rule gets activated, the more probable it is for the overall system to reveal a higher amount of internal knowledge—so the less overall privacy is attained.

Minimizing the amount of internal knowledge that gets externalized at some point required a modification of the message sequence to made sure that as many rules as possible were activated as earlier as possible. Therefore, we introduced a new round of exchanges to the sequence. In this round, each agent receives the original intent and expresses its ability to resolve all items. Once the round finishes, the original intent is sent to the agent that declared to have higher ability to resolve all items. The resulting sequence diagram is shown in Figure 3.

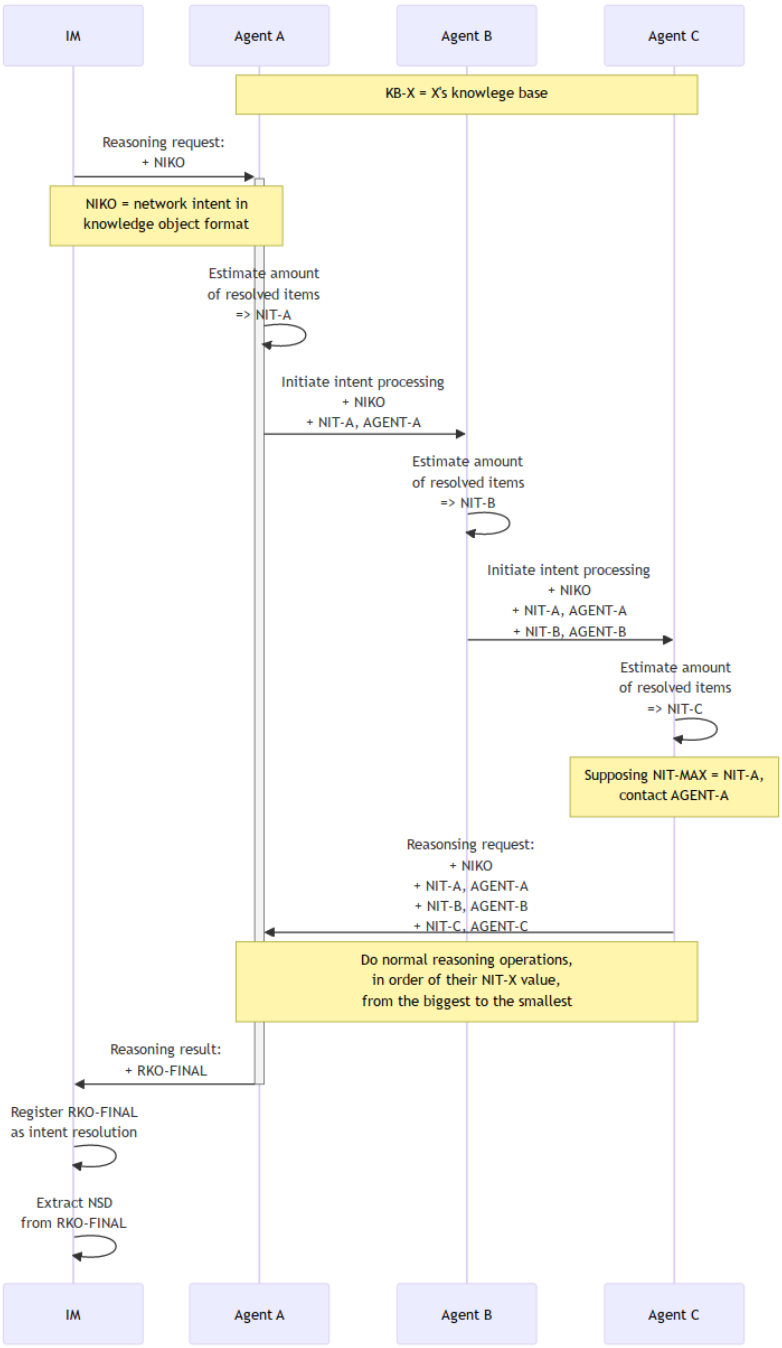

Figure 3: Adjusted Message Sequence of the Resolution of a Network Intent Across Multiple Domains

4 Evaluation

We evaluated our proposal by implementing a multi-domain negotiation system that followed the optimum sequence diagram supported by the new model. Our implementation and an implementation of a reference system for multi-domain intent negotiation were deployed in a real platform that comprised three separated domains. For each implementation of the negotiation system, we deployed several multi-domain network services defined by multi-domain network intents, and measured the amount of knowledge exchanged for resolving the network intents.

4.1 Platform and System Implementation

The platform we used to evaluate our proposal consisted of three separated domains: one in NICT premises (Tokyo, Japan), one in Telefonica premises (Madrid, Spain), and the other in UMU premises (Murcia, Spain). The distances between NICT domain and Telefonica and UMU domains were more than 10000 km, whereas the distance between Telefonica and UMU domains was about 500 km.

We connected the domains through a virtual private network based on WireGuard [1] over the Internet. The control planes were interconnected using iCPN [16], which offers adequate functions for efficient management of monitoring data, as well as knowledge and intent-specific information. The underlying components required to realize the network service functions were instantiated through OpenStack [17] and some components of OpenSource MANO (OSM) [2]. The intent manager interfaced with the agents on each domain.

The most relevant components of the overall system and their relations are depicted in Figure 5 and a schema of the deployed network service is shown in Figure 4.

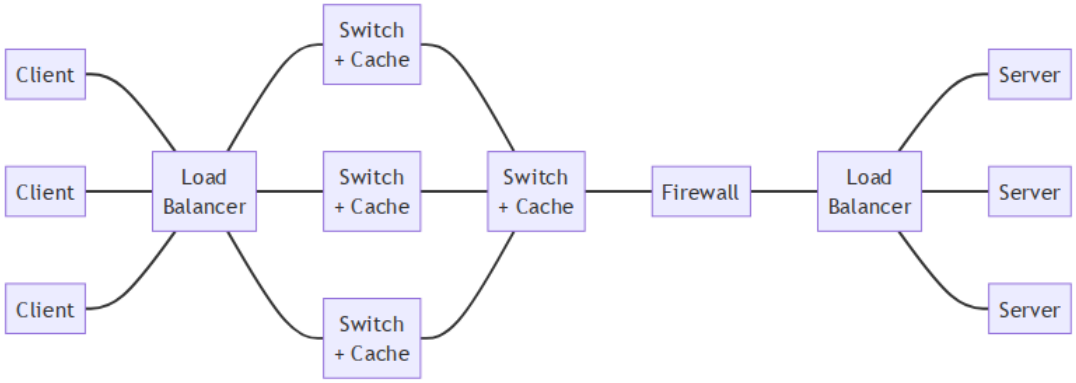

Figure 4: Network Service

[1] https://www.wireguard.com

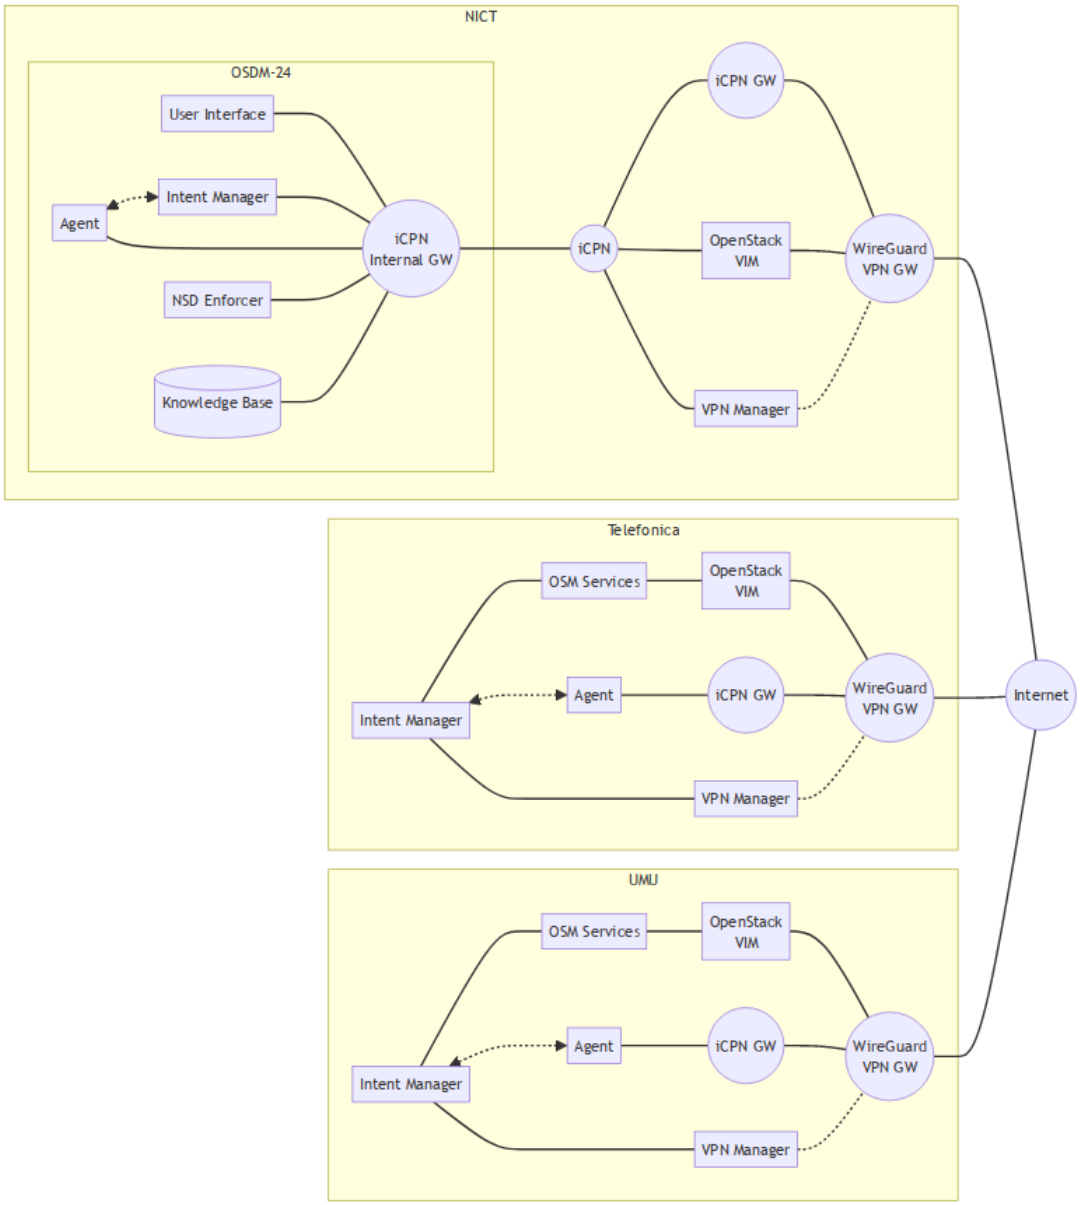

Figure 5: Experimentation Scenario

4.2 Experiment Execution and Results

To validate the outcomes of our model and the corresponding new system, whose design is guided by the model, we carried out an experiment that consisted of using different configurations for the agents of the three domains. Each configuration had a different knowledge base assigned to each agent. We prepared three knowledge bases, each with a different quantity of knowledge items—namely, a different level of knowledge to resolve network intents.

We had six different configurations, five of them were obtained by using the previous system and permuting the knowledge bases assigned to each agent, so that each agent had a different amount of knowledge items—we had a small knowledge base of 1009 bytes, a medium knowledge base of 3419 bytes, and a large knowledge base of 5039 bytes. The sixth configuration used our system and an arbitrary assignation of knowledge base to each agent—a straight forward assignation of knowledge bases sorted by size. The properties of each configuration are summarized in Table 1.

From the experiment execution we measured, on the one hand, the total amount of data exchanged by all agents and subtracted the basic amount of data exchanged by empty knowledge exchanges, which is estimated from the number of network intents (15), the number of concepts in each network intent (4), the base size of an exchange (1000 bytes), and the links (4, which is A→B, B→C, C→B, and B→A). The result number shows that, with the previous system, the amount of knowledge shared falls between 135 to 185 KiB. The mean of all exchanges is around 170 KiB. Meanwhile, our new system only exposed 135 KiB, which is 35 KiB less than the average and equals the best case of the previous system. It is a reduction of 20 % of shared knowledge by our system with respect to an average system. These results are represented in Figure 6.

On the other hand, to confirm that out proposal retains the same properties than previous work, we measured the time used by each configuration to resolve the provided network intents ($n = 14$). As expected, we confirmed that, the less data exchanged, the less time used to resolve the network intents. This is consistent because all configurations used the same algorithm. The only change was the content of the knowledge base of the agents. Nevertheless, our system took more time than other systems with comparable amounts of data exchanges because our system introduces the initial negotiation. As a result, our system required comparable time to the systems that exchanged the most amount of data, which confirms it is still feasible for the intent resolution purpose. The times are shown in Figure 7.

Table 1: Sizes of Knowledge Base Per Agent.

Configuration	Agent A	Agent B	Agent C
C-1	Small	Large	Medium
C-2	Small	Medium	Large
C-3	Large	Medium	Small
C-4	Large	Small	Medium
C-5	Medium	Large	Small
C-6	Medium	Small	Large
Proposal	Small	Medium	Large

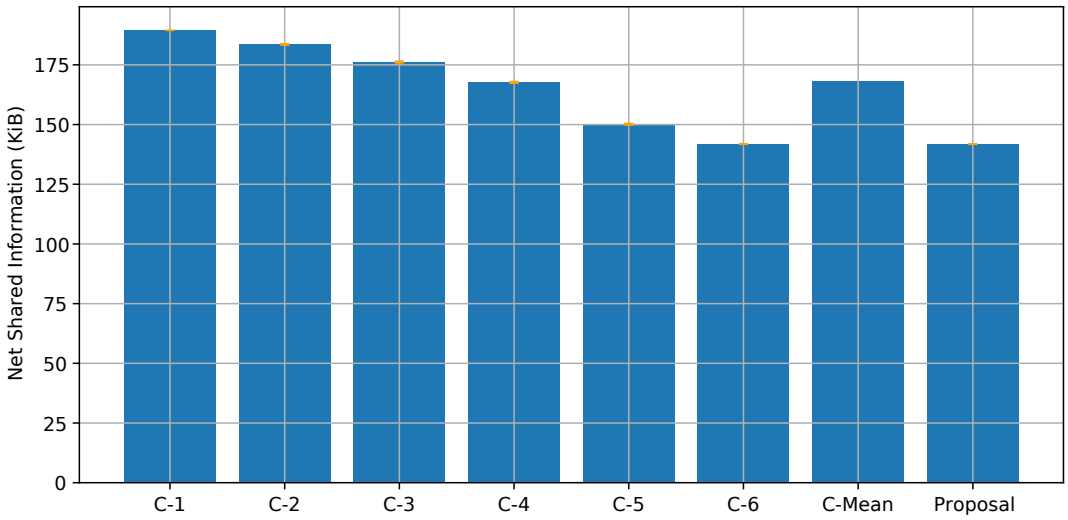

Figure 6: Total amount of knowledge shared by all domains by different intent resolution systems

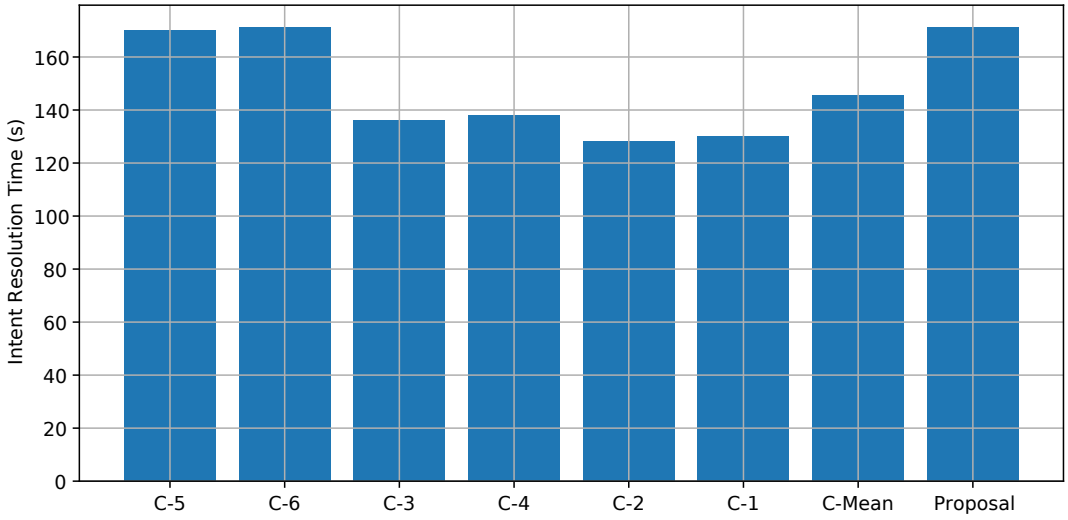

Figure 7: Total time elapsed by each configuration to resolve all network intents

5 Conclusion

In this study, we analyzed the knowledge shared in state-of-the-art systems for network intent resolution. We constructed a model that quantifies the amount of knowledge shared by the systems from their message exchanges. Using the model, we designed a system that minimizes the amount of knowledge shared, while keeping the effectiveness of network intent resolution and ensuring the performance is comparable with the current systems. The resulting system reduced the amount of knowledge shared in 20 % with respect to an average previous state-of-the-art system.

Future work to further reduce the amount of knowledge shared in network intent resolutions will consist of studying the possibility of encoding the knowledge bases in a format that binds the traceability of knowledge. It will reduce the possibility of an agent to forward received knowledge to other agents, so it is expected another enhancement in privacy for IBN. This will be reflected in our model by considering the impact of obfuscation and the possibility of hidden information, analyzing the content of message exchanges from the point of view of information theory.

References

[1] C. Li, O. Havel, A. Olariu, P. Martinez-Julia, J. Nobre, and D. Lopez, "Intent Classification," 2022, https://rfc-editor.org/rfc/rfc9316.txt.

[2] OSM, "OSM Release Five Technical Overview," 2019, OSM White Paper .

[3] P. Martinez-Julia, V. P. Kafle, and H. Asaeda, "Application of category theory to network service fault detection," *IEEE Open Journal of the Communications Society*, vol. 5, pp. 4417–4443, 2024.

[4] A. Hermosilla, J. Gallego-Madrid, P. Martinez-Julia, V. Kafle, K. Trantzas, C. Tranoris, R. Direito, D. Gomes, J. Ortiz, S. Denazis *et al.*, "Deployment of 5g network applications over multidomain and dynamic platforms," in *2022 IEEE Future Networks World Forum (FNWF)*, 2022, pp. 276–281.

[5] A. Hermosilla, J. Gallego-Madrid, P. Martinez-Julia, J. Ortiz, V. P. Kafle, and A. Skarmeta, "Advancing 5g network applications lifecycle security: An ml-driven approach," *Computer Modeling in Engineering & Sciences*, 2024.

[6] P. Martinez-Julia, A. Hermosilla, V. P. Kafle, and H. Asaeda, "Distributed ai-based network intent resolution for reliable security service orchestration," in *8th International Conference on Mobile Internet Security (MobiSec 2024)*, Sapporo, Japan, 2024.

[7] N. F. S. de Sousa and C. E. Rothenberg, "Clara: Closed loop-based zero-touch network management framework," *2021 IEEE Conference on Network Function Virtualization and Software Defined Networks (NFV-SDN)*, pp. 110–115, 2021.

[8] A. Leivadeas and M. Falkner, "A survey on intent-based networking," *IEEE Communications Surveys and Tutorials*, vol. 25, pp. 625–655, 2023.

[9] L. Velasco, M. Signorelli, O. D. Dios, C. Papagianni, R. Bifulco, J. Olmos, S. Pryor, G. Carrozzo, J. Schulz-Zander, M. Bennis, R. Martínez, F. Cugini, C. Salvadori, V. Lefebvre, L. Valcarenghi, and M. Ruiz, "End-to-end intent-based networking," *IEEE Communications Magazine*, vol. 59, pp. 106–112, 2021.

[10] F. Christou, "Decentralized intent-driven coordination of multi-domain ip-optical networks," *2022 18th International Conference on Network and Service Management (CNSM)*, pp. 359–363, 2022.

[11] T. Khan, K. Abbas, A. Rafique, M. Afaq, and W.-C. Song, "Generic intent-based networking platform for e2e network slice orchestration & lifecycle management," *2020*

21st Asia-Pacific Network Operations and Management Symposium (APNOMS), pp. 49–54, 2020.

[12] A. Leivadeas and M. Falkner, "Vnf placement problem: A multi-tenant intent-based networking approach," *2021 24th Conference on Innovation in Clouds, Internet and Networks and Workshops (ICIN)*, pp. 143–150, 2021.

[13] K. Abbas, T. Khan, M. Afaq, and W.-C. Song, "Ensemble learning-based network data analytics for network slice orchestration and management: An intent-based networking mechanism," *NOMS 2022-2022 IEEE/IFIP Network Operations and Management Symposium*, pp. 1–5, 2022.

[14] J. Massa, S. Forti, F. Paganelli, P. Dazzi, and A. Brogi, "Declarative provisioning of virtual network function chains in intent-based networks," in *2023 IEEE 9th International Conference on Network Softwarization (NetSoft)*, 2023, pp. 522–527.

[15] D. Borsatti, W. Cerroni, and S. Clayman, "From category theory to functional programming: A formal representation of intent," in *2022 IEEE 8th International Conference on Network Softwarization (NetSoft)*, vol. , no. , 2022, pp. 31–36.

[16] P. Martinez-Julia, V. P. Kafle, and H. Asaeda, "iCPN: Scalable control plane for the network service automation system," in *Proceedings of the 2024 IFIP/IEEE Network Operations and Management Symposium (NOMS)*. Washington, DC, USA: IEEE, 2024, pp. 1–5.

[17] "The OpenStack Project," 2017, http://www.openstack.org/.