INDIGO: Orchestration Across Multi-Operator, Multi-Vendor, 5G Networks During Civilian Disasters

John Drogo*, Anya Trumbach*, Sebastian Theiler†, Sreeram Mandava‡, Sanskar Shah‡, Ivan Seskar‡, Martin Skorupski§, K. K. Ramakrishnan¶, Kostas Oikonomou‡, Gil Zussman*, Tracy van Brakle‡
*Department of Electrical Engineering, Columbia University, New York, NY, USA
†Department of Electrical & Systems Engineering, Washington University in St. Louis, St. Louis, MO, USA
¶Department of Computer Science and Engineering, University of California Riverside, Riverside, CA, USA
‡WINLAB, Rutgers University, New Brunswick, NJ, USA
§Highstreet Technologies, Berlin, Germany
{j.drogo, art2187, gil.zussman}@columbia.edu, s.k.theiler@wustl.edu,
{sreeram, seskar, ko56, vanbrakle}@winlab.rutgers.edu, ss4308@scarletmail.rutgers.edu,
martin.skorupski@highstreet-technologies.com, kk@cs.ucr.edu

***Abstract*—During civilian disasters, the communications networks vital for emergency response are often severely damaged. Although individual mobile network operators (MNOs) have developed systems to mitigate such outages, such siloed solutions are limited by a lack of interoperability. Collective surviving resources across multiple MNOs constitute a more resilient resource pool than any individual network. To leverage this resource pool, we developed the Intelligent Networks Designed and Integrated for Globalized Operations (INDIGO) system. INDIGO improves existing solutions by combining multiple MNOs into a virtual network. We leverage open standards such as O-RAN and TM Forum APIs to provide monitoring and management. Control is facilitated by a hierarchical Artificial Intelligence (AI) Planner that quickly explores restoration options and issues configuration commands to create a virtual slice that satisfies mission requirements. We deploy a prototype of INDIGO and emulate an urban network comprised of three MNOs on the NSF PAWR COSMOS testbed. We then evaluate INDIGO's usability and performance through a timed case study of non-subject matter experts responding to a simulated disaster scenario. Each participant was able to create network coverage plans for three disaster response teams in less than six minutes.**

***Index Terms*—5G, Multi-Operator Network Orchestration, O-RAN Alliance, TM Forum, AI Planning**

I. Introduction

Rapid response to civilian disasters is crucial to minimize loss of life. Emergency response teams require careful coordination, often through a mobile network. However, such networks can face severe outages due to infrastructure damage caused by disasters [1]. The loss of communications infrastructure during Hurricane Katrina in 2005 and the failure of more than 90% of cellular sites in Puerto Rico after Hurricane Maria in 2017 are two prominent examples where response coordination was severely hampered by infrastructure damage [2]–[4]. Current disaster response networks often depend on the infrastructure of a single Mobile Network Operator (MNO). However, the collective surviving equipment of all MNOs in a region constitutes a more robust resource pool than any sole operator. The current lack of automated mechanisms to dynamically aggregate equipment across all MNOs implies that this redundancy cannot be efficiently exploited.

To address this, we present the Intelligent Networks Designed and Integrated for Globalized Operations (INDIGO) system [5]. INDIGO is an artificial intelligence-powered cellular network control system that rapidly adapts infrastructure through open, industry-standard APIs to restore critical network access in response to disasters. This enables the creation of a virtual network slice that utilizes available cellular towers and radios across several MNOs. The goal is to reduce outage times from hours to minutes by connecting MNOs through multi-operator orchestration that uses human interpretable plans quickly generated by a hierarchical Artificial Intelligence (AI) Planner. We deployed INDIGO on the NSF PAWR COSMOS testbed and simulated a small urban network to evaluate INDIGO's performance as non-subject matter experts restored communications during a simulated disaster [6].

In this paper, we provide a preliminary outline of the architecture and components of INDIGO, a system designed to orchestrate network services across multiple independent 5G operators. The key contributions include:

1) A modular design, extensible through rApps, that leverages open, industry-standard APIs from the TM Forum and O-RAN Alliance for interoperability across MNOs and equipment vendors [7], [8].
2) The application of hierarchical AI planning in the Mission rApp to translate high-level mission needs into a concrete set of network slice configurations.
3) A Mission User Interface (MUI) that visualizes network state and enables users to request virtual slice creation.

In Section II, we review the related work and the foundational technologies that enable our approach. In Section III, we detail the INDIGO architecture and its core components. In Section IV, we discuss our preliminary implementation of INDIGO, and in Section V we walk through a case study of participants using this implementation to respond to a simulated disaster. In Section VI, we provide concluding remarks and discuss future work.

While the INDIGO prototype demonstrates the benefits of AI-enabled multi-operator orchestration, this work does not imply support or represent the official position of any public safety, government, or industry partner.

II. Related Work

INDIGO builds on the convergence of public safety communications, networking technologies, and open standards to support emergency response communications. We summarize these developments in this section.

A. Emergency Response Networks

APCO Project 25 (P25) is a Land Mobile Radio (LMR) system introduced in the United States (US) to solve interoperability issues between emergency responder radio equipment. It remains in use in the US and internationally, although TETRA, a competing LMR standard, is used more broadly outside the US [9]–[11]. Several countries have established national networks around TETRA for emergency responders, including BOSNet in Germany and Airwave in the United Kingdom [12]. Given the limited spectral bandwidth allocated to these radio systems, most LMR deployments support low data rates and are primarily restricted to voice communications. Additionally, coordination is required to deconflict multiple concurrent users.

The limitations of LMR systems became apparent during the September 11, 2001 attacks. The rapid increase in communications between agencies responding to the disaster resulted in outages [10], [11]. In response, the US Congress began investing in dedicated nation-wide mobile networks to support emergency response teams [11], [13], [14].

Although nationwide single MNO systems can provide emergency network coverage, they consolidate public safety communications under a single operator. In the event of an outage, restoration depends on this MNO, with timelines and effectiveness varying with the cause and the maturity of the specific MNO's system. The Federal Communications Commission analysis of a 2020 T-Mobile outage, which led to the failure of thousands of individuals' 911 calls, underscored the need for robust and interoperable recovery mechanisms [15]. In contrast to single MNO networks, decentralized network solutions have been proposed, such as the creation of ad hoc networks over WiFi Direct links [16]. However, such decentralized networks have not achieved the widespread adoption necessary for public safety systems.

B. 5G Networking

5G networking introduced significant architectural changes over previous mobile technologies enabling ultra-reliable and low-latency communications (URLLC), extreme mobile broadband, and massive machine-type communications [17, pp. 1–6]. Network slicing was introduced in 5G to support URLLC and similar applications by logically partitioning a single network into multiple isolated End-To-End (E2E) virtual networks [18]. Each slice is managed by an MNO and can be configured with desired performance characteristics, including bandwidth and latency [17, pp. 41–50] [19]–[22].

Several MNOs have begun adopting O-RAN and deploying Service Management and Orchestration (SMO) software. SMO assists with the coordination of network slice lifecycles and the management of the Radio Access Network (RAN) in a single operator domain. An MNO may use SMO to automate network functions, such as resource management and service provisioning [23].

C. Network Interoperability

The TM Forum and O-RAN Alliance are developing standards that are well-positioned to overcome the management barriers between MNOs. O-RAN seeks to disaggregate traditionally monolithic RAN units by defining a modular framework and introducing the RAN Intelligent Controller (RIC) [8]. The RIC is a host platform for third-party network management applications known as xApps, for real time management, and rApps, for non-real time management. This architecture enables data-driven control loops for traffic steering, resource allocation, and network slicing, while simplifying the integration of artificial intelligence and machine learning into the management and control layers [20], [24], [25].

The TM Forum expands this architecture by providing a suite of REST-based open APIs for business and operational functions such as service and slice resource catalogs, ordering, configuration, and inventory. These APIs provide a standard method of interoperator communication and are currently the most mature open APIs implemented in industry [26].

III. Architecture Overview

The INDIGO system consists of three layers of software, as illustrated in Fig. 1. The rApp layer provides a user interface and produces mission plans, the Multi-Operator Service Management and Orchestration (MO-SMO) layer converts mission plans into TM Forum API calls, and the Single Operator SMOs (SO-SMOs) represent the infrastructure of the individual MNOs.

A. rApps

The top software layer of INDIGO consists of O-RAN compliant rApps that run on the MO-SMO's Non-Real Time (Non-RT) RIC to realize different RAN automation and management use cases. Our initial Mission rApp hosts the MUI, which allows emergency responders to request a virtual network over a selected geographic area. This request is routed through the AI Planner module inside the Mission rApp to develop a plan for the creation of a multi-operator network slice and is sent to the MO-SMO for implementation.

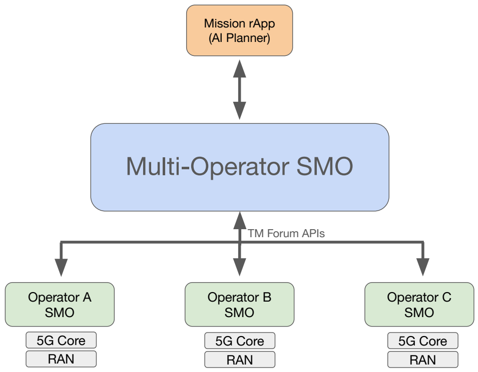

Fig. 1: Overall System Architecture: INDIGO consists of a three tiered software stack. The top level consists of the Mission rApp, which presents a Communications Unit Leader with an overview of current network status and allows them to create virtual network slice requests. The rApp's AI Planner then explores the network state and creates a mission plan that is sent to the MO-SMO in the middle layer. The MO-SMO converts this plan to a sequence of TM Forum API calls for the underlying infrastructure. The calls are processed by the SMOs on the bottom layer creating a virtual slice that restores communication.

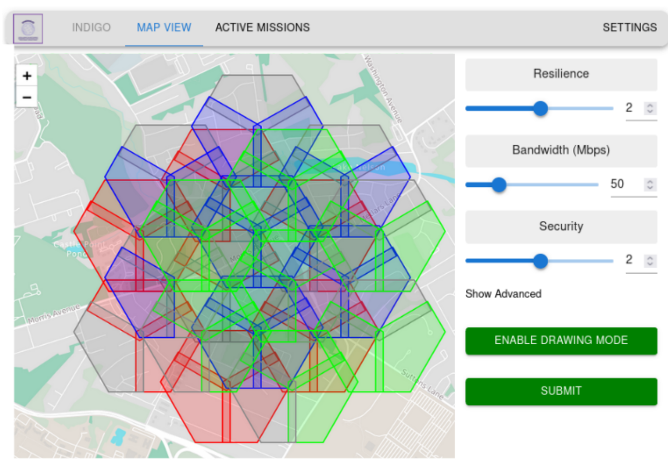

Fig. 2: Mission User Interface (MUI) Map View: MNO coverage is indicated in the map on the left of the screen by hexagons. The colored hexagons (red, blue, and green) correspond to active coverage of a specific MNO, while light gray regions indicate cellular tower failure. Users can adjust Resilience, Bandwidth, and Security parameters with the right-sidebar (Loss and Latency can also be adjusted by clicking the "Show Advanced" button). Users can toggle the ability to draw a boundary for a virtual slice with the "Enable Drawing Mode" button and submit a slice creation request to the AI Planner with the "Submit" button.

Mission User Interface (MUI) - The MUI is a web application that provides an interface for users to interact with INDIGO. The streamlined layout is designed to minimize cognitive load in high-stress operational scenarios. A screenshot of the MUI map view is shown in Fig. 2 which displays information on the available infrastructure of the MNOs and their current coverage area. The interface provides a map-based drawing tool that enables users to define the geographic boundaries of the requested network slice. Additional settings for how the plan is generated can be modified in the Settings page, located in the upper right corner. Upon submission, the MUI backend validates the requested network slice specifications and sends a slice creation request to the AI Planner for plan generation.

AI Planner - The Mission rApp contains an AI Planner for plan generation. AI planning, a subfield of artificial intelligence, emphasizes and exploits human experts' knowledge about a problem domain, rather than learning about it. It is a broad discipline with applications in many diverse fields, including networking [27]. Here, the AI Planner has at its disposal a snapshot of the state of the operator networks collected via the MO-SMO. It uses that information to create a network slice in the region identified by the user. When there are failures in the operator networks, the AI Planner coordinates the sharing of resources between the operators through the MO-SMO. The planner design emphasizes the *speed* and *understandability* of decisions. The user may accept, modify, or reject the plans it formulates. These plans are similar to what human experts would produce if they had enough time but are generated in seconds.

The AI Planner is *hierarchical*, which implies that it recursively decomposes tasks into simpler components [28]. The AI Planner can be thought of as a human strategist or manager for the creation of multi-operator network slices using the available resources. This manager has at their disposal a team of 'experts,' the operator SMOs and rApps, each of whom knows how to accomplish certain tasks. In addition to the state of the networks, the planner embodies higher-level expert strategies on how to decompose tasks into simpler sub-tasks. Given the top-level task of creating a network slice with specified parameters, there are many methods for decomposing it, each of which applies under a set of conditions; the same applies to the sub-tasks into which this task is recursively decomposed. The planner searches for applicable decomposition methods until the given tasks are reduced to primitive actions. Primitive actions are tasks that a 'team member' (in this case, an operator SMO) can perform without further guidance.

B. *Multi-Operator SMO (MO-SMO)*

The middle layer is the MO-SMO platform, which coordinates across the SO-SMOs. Based on its knowledge of the operators' network topologies, the MO-SMO orchestrates the multi-operator network by translating the slice creation plan provided by the AI Planner into specific TM Forum API calls for each operator SMO to execute. In the current implementation, the capabilities of the MO-SMO are focused on the construction of TM Forum API calls to accomplish the requested network plan. Actual execution and monitoring of these calls to the SO-SMOs is currently simulated.

C. *Operator SMOs*

The bottom layer consists of the SO-SMOs and the physical infrastructure they manage. To be compatible with the INDIGO architecture, SO-SMOs must provide:

1) *Topology information exposure and management*. This enables the MO-SMO to obtain and use a complete view of the single operator's orderable network slice subnets, including information about their availability and status.

2) *E2E slice or sub-slice ordering*. This includes the operator's ability to create, activate, and configure the slice's RAN, Core, and 5G transport components.

The infrastructure of the SO-SMOs is currently simulated in our setup and acknowledges a requested TM Forum API call while returning static state information. This allows us to verify the plan generation and ensure that the resulting TM Forum API calls sent to the simulated SMOs adhere to the open standards.

IV. Implementation

We implemented INDIGO on virtual machines (VMs) in the NSF PAWR COSMOS testbed to emulate the cellular network infrastructure of an urban environment [6]. The MUI consists of front-end and back-end Podman containers which communicate to the AI Planner container on a separate VM. The AI Planner is colocated with the MO-SMO and communicates to the MUI back-end through JSON messages sent over TCP connections. The MO-SMO simulates acknowledgments of instructions sent to the SO-SMOs after generating TM Forum API calls. The SO-SMOs are configured to represent a network consisting of three MNOs with 63 radio units (RUs) deployed across 21 tower locations. The SO-SMOs report pre-defined status messages simulating a partially degraded network.

V. Case Study in System Usage

A. Case Study Overview

To evaluate INDIGO's ease of use and response time, we conducted a trial with four non-subject matter experts, users who had not previously worked with INDIGO, role-playing as Communications Unit Leaders (COMLs) in a simulated disaster scenario. The COMLs were given four minutes to read a manual on using INDIGO. COMLs were then briefed for three minutes on a simulated network scenario in which widespread power outages and physical damage to cell towers had degraded the three regional MNOs following a natural disaster. In total, 15 of the 63 radio units had failed.

The COMLs were asked to use INDIGO to respond to the disaster scenario, after being given the following context: An Incident Commander for an Urban Search & Rescue (US&R) team requires reliable Push-to-Talk over Cellular voice links and the ability to stream live video from helmet cameras to the Incident Command Post to assess structural damage in real time. The current network performance is insufficient to support this video requirement.

B. Case Study Results

The flow of mission responses is summarized in Fig. 3. To simplify the case study, we simulated request acknowledgments from the SO-SMOs and automatically approved the returned response from the AI Planner. Without a physical SMO layer, we were unable to assess the time it would take for the underlying SMO infrastructure to implement the determined plan. Instead, we timed from when a user opened the MUI (Fig. 3a) to when the MO-SMO issued TM Forum calls to create a virtual slice for the US&R team (Fig. 3d).

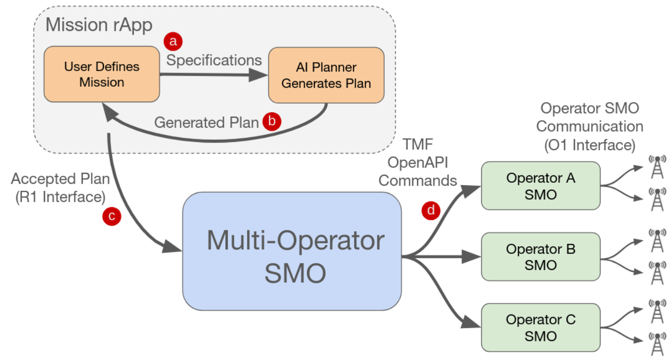

Fig. 3: Mission Flow Diagram: The data flow in INDIGO follows four steps. (a) First, the user defines the slice boundary by drawing a polygon and specifying the required network parameters. (b) The plan is generated by the AI Planner to satisfy the submitted requirements. (c) The plan is sent to the MO-SMO and translated into a sequence of service orders. (d) The MO-SMO sends individual operator SMOs service orders using TM Forum APIs.

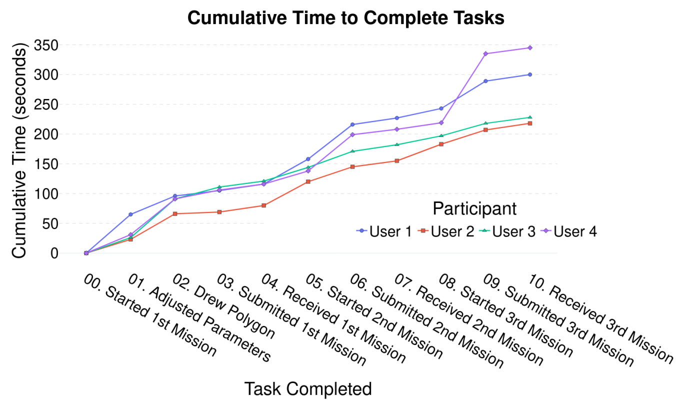

Fig. 4: Case Study Time Responses: Four non-expert users responded to a simulated disaster scenario comprised of three missions. The cumulative response times for the completion of the three missions for the participants is shown. Key events are indicated horizontally. We measured the time for users to identify the outage, request a virtual slice, and receive the adaptation plan from the AI Planner. All users were able to respond to the scenarios within 6 minutes. This suggests the MUI and AI Planner provides an intuitive system for non-subject matter experts to understand and respond to network outages.

This allowed our case study to observe the time for a COML to assess the network state via the MUI and request a virtual slice (Fig. 3a), the AI Planner to construct a plan consisting of configuration steps to create this slice (Fig. 3b), and the MO-SMO to translate the plan into a sequence of service orders (Fig. 3c) dispatched to the individual SMOs (Fig. 3d).

The MUI interactions were recorded to determine the response times. These cumulative response times are shown in Fig. 4 for each of the COMLs in the three consecutive US&R team missions. All COMLs successfully responded to the scenarios within 6 minutes. This shows that the system was able to provide context to non-subject matter experts, identify a mission plan, and start restoration on the order of minutes.

The flow of data throughout the system can be broken down into two phases. The Mission Definition phase and the Automated Planning phase. The Mission Definition phase refers to the time it takes a COML to analyze the MUI map display and issue a request to INDIGO to address the outage. The Automated Planning phase refers to the time it takes for

```
((!assign-subnet-to-mo-smo access area2 smo-operator-a #:|DR-smo|
    ((bw . 18) (lat . 2) (loss . 0.1) (res . 2) (sec . 2)))
(!assign-subnet-to-mo-smo transport area2 smo-operator-a #:|DR-smo| nil)
...
(!re-direct 25fc6ce4-9ed7-5022-bcbb-5198cb8141d9 smo-operator-a c9e64b8b-6ef1-57cc-b298-570f03b9471c
    smo-operator-c)
(!re-direct 3030e90b-0e79-5cb8-8bd6-d97f5686a988 smo-operator-a e86eff8f-6629-5176-93e7-b388e7acae50
    smo-operator-c)
```

Listing 1: Excerpt of Mission Plan from Case Study: The plans generated by the AI Planner are human readable configuration steps that achieve the mission goal. This example plan was generated in the case study and shows a request to create a subnet from Operator A with a bandwidth of 18 Mbps, latency target of 2ms, packet loss tolerance of 10%, and a scalar value representing the relative importance of resilience and security for the new slice. The plan then redirects traffic from two radio units from Operator C to this newly created subnet to construct a virtual slice.

INDIGO to produce a network restoration plan and send it to both the COML for review and MO-SMO for execution.

Mission Definition - At T+0 seconds, the COMLs were presented with a map on the MUI that showed the available network infrastructure of three participating partner operators. The coverage in the affected area was colored gray, indicating that the corresponding cellular towers had been damaged. The COMLs adjusted parameters as required by the briefing in an average of 36.3 seconds. They drew a polygon on the map, defining the operational area, and submitted the mission after an average of T+97.8 seconds (Fig. 3a).

Automated Planning - Once the MUI request was submitted, the Mission rApp forwarded the mission requirements to the AI Planner. The AI Planner used its inventory of available network resources and individual MNO topologies to generate a plan that satisfied the mission requirements (Fig. 3b). In this simulation, while no individual operator could meet the mission requirements, the AI Planner was able to combine the resources of the three operators to create a functional network slice. An excerpt of the plan produced by the AI Planner is given in Listing 1. The COMLs received the formatted version of this plan on average 10.5 seconds after submission. The plan generated by the AI Planner was displayed in the MUI for review by the COML. In total, an average of 108.3 seconds elapsed between when users opened the INDIGO MUI and when they received the generated plan.

Second and Third US&R Teams - After receiving and reviewing the generated plan, the COMLs were informed a second US&R team was being deployed to a different location and subsequently asked to generate a new coverage slice. The COMLs completed this second task in an average of 42.8 seconds and received the generated plan 10.3 seconds later. Generating a coverage slice for a third US&R team took the COMLs an average of 51.8 seconds to submit, and a plan was received 10.5 seconds later. COMLs were not asked to change the bandwidth, resilience, and security parameters for the second and third missions. On average, COMLs established coverage for all three US&R teams within T+272.8 seconds.

VI. Conclusion and Future Work

INDIGO addresses the challenge of maintaining a robust network for emergency communications during disasters that cause widespread damage to cellular infrastructure. Our solution alleviates the bandwidth limitations of P25 and other LMR-based systems and takes advantage of additional network resources which are inaccessible to single MNO approaches. The AI Planner and MUI provide a simple means for interacting with INDIGO and adapting the network, even for users without extensive technical training. This enables INDIGO to develop plans for connecting responder teams within minutes.

This paper presents a preliminary implementation of the INDIGO system. Future work to expand INDIGO will consist of improvements to the existing software and the addition of new rApps. The integration of continuous SMO monitoring into the AI Planner would allow the mission plan to automatically adapt as the situation develops. An extension for managing authentication would increase INDIGO's security posture. Similarly, a resilience rApp would improve system flexibility by managing multiple communication paths, including satellite and drone deployments. The presented case study examines the response times of users interacting with the MUI and the timing of the AI Planner to generate a mission plan for a scenario. Future work should consider the SO-SMOs' timing as they execute the mission plan's steps. In addition, we look to replace the SO-SMO simulations with hardware systems to analyze performance in an E2E deployment.

VII. Acknowledgments

This work was supported by the NSF Convergence Accelerator Cohort Phase 2 Award No. 49100422C0020. We would like to acknowledge Sasha Isler and Riju Dey of Columbia University who contributed to the MUI implementation.

References

[1] S. Hansson, K. Orru, A. Siibak, A. Bäck, M. Krüger, F. Gabel, and C. Morsut, "Communication-related vulnerability to disasters: A heuristic framework," *Int. J. of Disaster Risk Reduction*, vol. 51, p. 101931, Dec. 2020.

[2] L. Comfort and T. Haase, "The impact of Hurricane Katrina on communications infrastructure," *Public Works Management & Policy*, vol. 10, pp. 328–343, Apr. 2006.

[3] Federal Communications Commission, "Communications status report for areas impacted by Hurricane Maria," Federal Communications Commission, Washington, DC, Tech. Rep., Sep. 2017. [Online]. Available: https://transition.fcc.gov/Daily_Releases/Daily_Business/2017/db0927/DOC-346943A1.pdf

[4] F. Aros-Vera, S. Gillian, A. Rehmar, and L. Rehmar, "Increasing the resilience of critical infrastructure networks through the strategic location of microgrids: A case study of Hurricane Maria in Puerto Rico," *Int. J. of Disaster Risk Reduction*, vol. 55, p. 102055, Mar. 2021.
[5] INDIGO, "INDIGO," July 2025. [Online]. Available: https://www.indigo.cosmos-lab.org
[6] D. Raychaudhuri, I. Seskar, G. Zussman, T. Korakis, D. Kilper, T. Chen, J. Kolodziejski, M. Sherman, Z. Kostic, X. Gu, H. Krishnaswamy, S. Maheshwari, P. Skrimponis, and C. Gutterman, "Challenge: COSMOS: A city-scale programmable testbed for experimentation with advanced wireless," in *ACM MobiCom*, Apr 2020.
[7] V. Mochalov, N. Bratchenko, G. Linets, and S. Yakovlev, "Distributed management systems for infocommunication networks: A model based on TM Forum frameworx," *Computers*, vol. 8, no. 2, June 2019.
[8] M. Polese, L. Bonati, S. D'Oro, S. Basagni, and T. Melodia, "Understanding O-RAN: Architecture, interfaces, algorithms, security, and research challenges," *IEEE Commun. Surveys Tuts.*, vol. 25, no. 2, pp. 1376–1411, Jan. 2023.
[9] F. Frosali, F. Gei, D. Marabissi, L. Micciullo, and E. Lezaack, "Interoperability for public safety networks," in *Wireless Public Safety Networks 1*, D. Câmara and N. Nikaein, Eds. Amsterdam, Netherlands: Elsevier, Feb. 2015, ch. 5, pp. 127–162.
[10] A. Paulson and T. Schwengler, "A review of public safety communications, from LMR to voice over LTE (VoLT E)," in *IEEE PIMRC*, Sep. 2013.
[11] A. Yarali, Ed., *Public Safety Networks from LTE to 5G*. Wiley-IEEE Press, Dec. 2019.
[12] Kable Business Intelligence Limited, "First responder solutions in the UK and internationally," Kable Business Intelligence Limited, London, UK, Tech. Rep., Sep. 2016. [Online]. Available: https://www.nao.org.uk/wp-content/uploads/2016/09/First-Responder-Solutions-in-the-UK-and-Internationally.pdf
[13] B. Harrison, J. Dimarogonas, J. Catlin, R. H. Donohue, T. Goughnour, J. S. Hollywood, J. Mastbaum, K. Van Abel, and J. Balagna, "Broadband communications prioritization and interoperability guidance for law enforcement: Critical considerations in the transition to the Public Safety Broadband Network," RAND Corporation, Santa Monica, CA, Tech. Rep., Aug. 2022. [Online]. Available: https://www.rand.org/pubs/research_reports/RRA2019-1.html
[14] A. Kumbhar, F. Koohifar, I. Güvenç, and B. Mueller, "A survey on legacy and emerging technologies for public safety communications," *IEEE Commun. Surveys Tuts.*, vol. 19, no. 1, pp. 97–124, Sep. 2017.
[15] Public Safety and Homeland Security Bureau, "June 15, 2020 T-Mobile Network Outage Report," Federal Communications Commission, Washington, DC, Tech. Rep., Oct. 2020. [Online]. Available: https://docs.fcc.gov/public/attachments/DOC-367699A1.pdf
[16] M. Ilbeigi, A. Morteza, and R. Ehsani, "An infrastructure-less emergency communication system: A blockchain-based framework," *J. of Computing in Civil Engineering*, vol. 36, no. 2, p. 04021041, Mar. 2022.
[17] L. Peterson and O. Sunay, *5G Mobile Networks: A Systems Approach*. Cham: Springer International Publishing, 2020.
[18] C. Gutterman, E. Grinshpun, S. Sharma, and G. Zussman, "RAN resource usage prediction for a 5G slice broker," in *ACM MobiHoc*, July 2019.
[19] P. Rost, C. Mannweiler, D. S. Michalopoulos, C. Sartori, V. Sciancalepore, N. Sastry, O. Holland, S. Tayade, B. Han, D. Bega, D. Aziz, and H. Bakker, "Network slicing to enable scalability and flexibility in 5G mobile networks," *IEEE Commun. Mag.*, vol. 55, no. 5, pp. 72–79, May 2017.
[20] J. Dai, L. Li, R. Safavinejad, S. Mahboob, H. Chen, V. V. Ratnam, H. Wang, J. Zhang, and L. Liu, "O-RAN-Enabled intelligent network slicing to meet Service-Level Agreement (SLA)," *IEEE Trans. Mobile Comput.*, vol. 24, no. 2, pp. 890–906, Oct. 2025.
[21] V. Q. Rodriguez, F. Guillemin, and A. Boubendir, "5G E2E network slicing management with ONAP," in *Proc. 23rd Conf. Innov. Clouds, Internet Netw. Workshops (ICIN)*, Feb. 2020.
[22] S. Zahoor, I. Ahmad, M. T. B. Othman, A. Mamoon, A. U. Rehman, M. Shafiq, and H. Hamam, "Comprehensive analysis of network slicing for the developing commercial needs and networking challenges," *Sensors*, vol. 22, no. 17, p. 6623, Sep. 2022.
[23] O-RAN Working Group 4, "Control, user and synchronization plane specification architecture description," O-RAN, Alfter, Germany, Feb. 2025. [Online]. Available: https://specifications.o-ran.org/download?id=881
[24] O-RAN Working Group 2, "R1 Interface: Use Cases and Requirements," O-RAN, Alfter, Germany, Feb. 2025. [Online]. Available: https://specifications.o-ran.org/download?id=870
[25] ——, "Non-RT RIC & A1/R1 Interface: Use Cases and Requirements," O-RAN, Alfter, Germany, June 2025. [Online]. Available: https://specifications.o-ran.org/download?id=876
[26] "TMF630 REST API Design Guidelines," TM Forum, May 2021. [Online]. Available: https://www.tmforum.org/resources/specifications/tmf630-rest-api-design-guidelines-4-2-0/
[27] S. J. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach*, 3rd ed. Upper Saddle River, NJ: Pearson Education, Inc., Dec. 2009.
[28] M. Ghallab, D. Nau, and P. Traverso, *Automated Planning: Theory and Practice*, 1st ed. San Francisco, CA: Morgan Kaufmann, May 2004.