Predictive OMS Switchover towards Proactive Disaster Recovery in 5G Networks

Charles F. Santos[1,2], **Augusto V. Neto**[2], **Ramon R. Fontes**[2], **Roger Immich**[2], **Vicente Sousa Jr.**[2], **Helber W. da Silva**[1,2]

[1]Federal Institue of Rio Grande do Norte (IFRN)
Av. Sen. Salgado Filho, 1559 - Natal - RN - Brazil

[2]Federal University of Rio Grande do Norte (UFRN)
Campus Universitário - Lagoa Nova, Natal - RN - Brazil

charles.santos@ifrn.edu.br, augusto@dimap.ufrn.br, ramon.fontes@imd.ufrn.br

roger@imd.ufrn.br, vicente.sousa@ufrn.br, helber.silva@ifrn.edu.br

Abstract. *The increasing complexity and critical nature of Operations and Maintenance Systems (OMS) in 5G mobile networks necessitate robust Disaster Recovery (DR) solutions to ensure continuous service and minimal downtime. Disaster Recovery Systems (DRS) are essential for maintaining network resilience by facilitating seamless failover and recovery processes. The primary function of switchover enables a DRS to ensure 5G service continuity during unforeseen disasters. This research addresses the limitations of traditional rule-based decision-making techniques that often rely on binary switchover logic, which can be inadequate for the intricate demands of 5G networks. We propose the predictive OMS Switchover (pOM2S), a machine learning (ML)-driven approach that utilizes data on computing and networking metrics to estimate the switchover period for each redundant candidate. In doing so, our solution is able to select the backup OMS that can perform the fastest in a disaster event. Evaluated in a 5G emulation testbed simulating real-world conditions, results demonstrated Random Forest's superior accuracy (MAE: 1.68s, R^2: 0.94) over Linear Regression, Artificial Neural Networks, and Support Vector Machine techniques. Experimental results validate pOM2S's effectiveness in balancing predictive precision and operational practicality, suggesting that the predictive decision-making method relies on a highly accurate model.*

Resumo. *A crescente complexidade e a natureza crítica dos Sistemas de Operações e Manutenção (OMS) em redes móveis 5G exigem soluções robustas de Recuperação de Desastres (DR) para garantir a continuidade do serviço e o mínimo de tempo de inatividade. Os Sistemas de Recuperação de Desastres (DRS) são essenciais para manter a resiliência da rede, facilitando processos de failover e recuperação sem interrupções. A função principal da comutação permite que um DRS assegure a continuidade do serviço 5G durante desastres imprevistos. Esta pesquisa aborda as limitações das técnicas tradicionais de tomada de decisão baseadas em regras que frequentemente dependem de lógica de comutação binária, inadequada para as demandas intrincadas das redes 5G. Propomos a Comutação Preditiva do OMS (pOM2S), uma abordagem orientada por aprendizado de máquina (ML) que utiliza dados sobre métricas* *de computação e rede para estimar o tempo de transição para cada um dos candidatos redundantes. Com essa estimativa, nossa solução é capaz de selecionar o OMS de backup mais rápido em um evento de desastre. Avaliados em um ambiente de emulação 5G que simula condições do mundo real, os resultados demonstraram a precisão superior da Floresta Aleatória (MAE: 1,68s, R^2: 0,94) em comparação com Regressão Linear, Redes Neurais Artificiais e técnicas de Máquinas de Vetores de Suporte. Os resultados experimentais validam a eficácia do pOM2S em equilibrar precisão preditiva e praticidade operacional, sugerindo que o método de tomada de decisão preditiva depende de um modelo altamente preciso.*

1. Introduction

The growing complexity and criticality of Operations and Maintenance Systems (OMS) [Zheng et al. 2020] in 5G mobile networks necessitate robust Disaster Recovery (DR) solutions to ensure uninterrupted services and minimal downtime. Disaster Recovery Systems (DRS) play a pivotal role in maintaining network resilience by orchestrating seamless failover and recovery mechanisms. Switchover serves as the DRS main functionality, seeking to enable 5G service continuity during unexpected disaster scenarios, ranging from natural disasters (earthquakes, floods, or hurricanes) to cyberattacks, hardware failures, power outages, overloaded network conditions, and other examples potentially impacting widespread damage. The DRS monitors redundant OMS instances deployed within the 5G infrastructure, ensuring seamless transition of operations from a compromised (primary) to a fully operational (backup) instance. This capability is critical for maintaining the resilience of 5G networks, as it ensures OMS functionality remains active under adverse conditions, thus preserving network stability and high service availability.

Historically, DR solutions for telecommunication networks have relied heavily on reactive rule-based decision-making techniques. These traditional methods often utilize basic binary switchover logic, where operations shift solely between the primary and the backup OMS instances – a limitation well-documented in previous studies [Lawler et al. 2007]. While effective in simpler contexts, traditional DR solutions often struggle to address the intricate requirements of 5G networks by facing challenges, such as [Zaretalab et al. 2020]: high operational costs, inefficient resource utilization, and limited adaptability to dynamic network conditions. Moreover, reliance on reactive, binary rule-based decision-making processes results in suboptimal performance, especially in handling diverse failure scenarios. These limitations underscore the need for advanced, predictive approaches in DR tailored to the unique challenges of 5G networks.

Our research pursues to overcome these limitations by engaging 5G operators to harness the advancements of telco-cloud facilities (network virtualization [Mijumbi et al. 2016], cloud/edge-native architecture [Group 2020], Software-Defined Networking (SDN) [Haleplidis et al. 2015], and others), so provisioning multiple OMS instances (multi-redundancy) at the network core infrastructure to foresee enhancing resilience. Although promising, this comes at the cost of requiring DRS to stride beyond the traditional binary rule-based decision-making logic for incorporating OMS multi-redundancy awareness and implementing a more efficient switchover approach. One of the primary difficulties lies in selecting the most suitable backup OMS instance to take over operations from a compromised primary OMS timely to prevent service and networking quality degradations [Zunino et al. 2024].

Ensuring seamless and rapid switchover to yield a full transition requires sophisticated algorithms, capable of evaluating the multiple capabilities of all OMS backup candidates simultaneously and in real time, so that accomplishing the transfering of data and services timely. The highly dynamic nature of 5G networks, characterized by fluctuating traffic loads, variable connectivity, and diverse failure scenarios, complicates the DR decision-making process further. Thus, we claim that such decision must consider real-time network analysis on latency, throughput, and resource availability, while also factoring in historical data to derive the reliability and performance of potential backups.

In this research, we move from the traditional binary rule-based decision-making paradigm towards statistic-based techniques to tackle the challenges of predictive OMS switchover in 5G DR with backup options multi-reduncancy settings. With this goal in mind, we investigate Machine Learning (ML) techniques to provide a method for selecting the optimal OMS backup option, among multi-redundant candidates, based on cost predictions of each OMS backup candidate for transitioning from the primary OMS. Considering the dynamic and resource-intensive nature of the underlying 5G networking infrastructure, along with the complexity of multi-redundant options, the method takes features at both computing and networking levels to provide high-accurate decisions. The lack of a related literature solution motivates our research, departing from the premise that accurate switchover cost prediction is vital for optimizing resource allocation, maintaining uninterrupted service, and minimizing 5G service disruptions.

The methodology adopted in this research involves an extensive analysis of state-of-the-art ML techniques relevant for estimating the switchover costs of OMS in DR use cases. A systematic evaluation was conducted within a testbed environment that emulates real-world technologies and scenarios characteristic of 5G networks and DR settings, aimed at identifying the most suitable ML model. Subsequently, we implemented the selected ML model in a novel OMS selection method, termed predictive OMS Switchover (pOM2S), specifically tailored for 5G networks. This proposal transcends existing rule-based reactive solutions by leveraging, for the first time in literature, ML models to analyze and estimate OMS switchover costs based on real-time computing and networking metrics. Furthermore, pOM2S determines the optimal OMS backup instance, among multiple redundant candidates, which will assume the primary role in the 5G infrastructure. Tests and experiments were conducted within a 5G emulated testbed, aiming for comprehensive assessments and analyses under conditions closely resembling real network environments.

The contributions of this work are threefold: (i) the design and implementation of a predictive switchover cost mechanism for multi-redundant OMSs in a 5G DRS architecture, (ii) a detailed comparative analysis of ML techniques to identify the optimal model for this use case, and (iii) validation of the pOM2S proposed solution's effectiveness through experimental results obtained from an emulation testbed. By addressing the critical challenge of predictive switchover cost estimation, this work provides a foundational step towards more resilient and efficient DR strategies for 5G and beyond networks. This particularly strives to yield a proactive management approach, aiming to enhance DR decision-making by anticipating switchover procedures for timely accomplishment.

The remainder of this paper is organized as follows. Section 2 reviews related work and highlights existing challenges in DR for 5G networks. Section 3 describes the proposed methodology and ML model selection process. Section 4 presents the testbed setup and experimental results, followed by a discussion of findings. Finally, Section 5 concludes the paper with insights and future directions for research in this domain.

2. Related Work

This section examines existing research footprints on DRS solutions in the 5G mobile network landscape. Our comparative analysis focuses on multiple optiomns redundancy, statistical decision-making, and OMS switchover as critical functional requirements to ensure the effectiveness and reliability of DRS approaches in 5G networks.

In [Zhu et al. 2022], authors propose a framework for keeping the balance between the cost and the reliability of the 5G User Data Management (UDM), providing a backup-based DRS approach. This solution employs a rule-based decision-making logic, by undertaking database synchronization from the main 5G UDM to the 4G HSS (Home Subscriber Server) assigned to serve as the disaster backup for 5G. Although novel, this approach traditionally employs a binary switchover scheme, and it is particularly focused on the UDM function-level resilience, instead of the whole 5G infrastructure. In the same way, [Leiter et al. 2022] introduced a mechanism atop the Open Network Automation Platform (ONAP), which proceeds to redeploy the 5G core User Plane Function (UPF) from one Kubernetes namespace to another. Authors of [Xie et al. 2023] also proposed a DR solution with particular focus at the UPF private 5G networking level, but operating under the ”1+N” (multi-redundant) DRS decision-making approach, enforcing the hypothesis that conventional ”1+1” (primary and backup DR) is low-resilient, thus limiting reliability and efficiency aspects.

The authors of [Al-Essa and Abdulbaki 2016] explore redundancy at the network infrastructure level to provide a backup DR site. Assuming a fully deployed network infrastructure, their proposal follows a ”1+1” redundancy scheme and employs a rule-based binary decision-making logic for switchover. In [Liang et al. 2022], a solution leveraging system metrics for failure prediction is proposed, using a DRS architecture to proactively initiate hierarchical recovery actions before failures occur. As disaster prediction is the paper’s primary focus, the solution relies on a single local backup information system instance to assume the primary role. The cloud service selection problem, achieved by cloud users’ difficulty in selecting a Content Service Provider (CSP) fitting their Quality of Service (QoS) requirements, is addressed in [Saha et al. 2021]. The solution, called H-MCDM, is proposed to tackle the cloud selection problem by ranking, through the ANP algorithm, a set of CSPs based on insights about benefit, opportunities, cost, and risk. In [Abdelaziz et al. 2023], in turn, harnesses the Multiple Criteria Decision Making (MCDM) algorithm to select CSPs based on QoS parameters (e.g., CPU, memory, and disk performance). Although providing selection algorithms based on multi-criteria, these approaches disconsider key network-level parameters, which significantly impact switchover cost and efficiency in DRS at the OMS level.

As summarized in Table 1, existing approaches exhibit significant gaps to cater with the functional requirements we elicitated for a reliable and efficient DRS solutions in 5G networks. Our comparative analysis review that related works focus on

Related Work	Multiple Options Redundancy	Statistical Decision-Making	OMS-level Switchover	Network-level Knowledge
[Zhu et al. 2022]	–	–	–	✓
[Leiter et al. 2022]	–	–	–	✓
[Xie et al. 2023]	✓	–	–	✓
[Al-Essa and Abdulbaki 2016]	–	–	–	✓
[Liang et al. 2022]	–	-	✓	–
[Saha et al. 2021]	✓	✓	–	–
[Abdelaziz et al. 2023]	✓	✓	–	–
This work	✓	✓	✓	✓

Table 1. Comparative Analysis of Disaster Recovery Solutions in 5G Networks.

rule-based decision-making approaches for either binary backup selection or recovering specific network functions in the 5G architecture, while overlooking the complexities of multi-redundant options. In response, this paper aims to design a novel OMS switchover method that adopts a statistic-based decision-making to enhance DR mechanisms by minimizing—or even neutralizing—resource downtime via more sophisticated control-plane strategies.

3. ML Model Selection

The design of ML-based models for the pOM2S proposal foresees paving the way towards a DR plan that anticipates the essential procedures for switching from the primary OMS to the most suitable backup OMS among a set of multi-redundant candidates. The methodology aims not only to identify the best-performing ML model for cost estimation, but also to provide insights into integrating ML within disaster recovery solutions in 5G networks.

In this study, we focus on developing the pOM2S approach within a 5G network infrastructure that includes multi-redundant OMS backup options. The pOM2S proposal introduces an innovative decision-making strategy designed to identify the most suitable OMS backup instance from a pool of candidates. The OMS backup instance selection method relies on the prediction of the most minimized switchover cost, thereby yielding an optimal transition. The cost prediction relates to estimating the total time required to execute the OMS switchover, which involves a sequence of the following steps: (1) retrieving all functional data from the primary OMS, (2) transferring the functional data from the primary OMS instance to the backup, (3) restoring the incoming primary OMS functional data, and (4) initializing the new OMS service with succeeding data synchronization.

With this in mind, we begin by examining selected popular ML algorithms that provide the prospect to effectively model and predict the targeting switchover cost. Our literature analysis came up with Linear Regression (LR), Support Vector Machine (SVM), Artificial Neural Network (ANN), and Random Forest (RF). During the evaluation phase, we undergo rigorous testing to assess the performance that each selected ML technique takes in accurately predicting switchover costs. The fundamental metrics used to evaluate and compare the performance of the ML models are accuracy, precision, recall, and F1-score, along with the computational efficiency. The aim is to understand how well these ML models perform in predicting costs associated with transitioning between different

OMS backup options in a multi-redundant setup, as well as to identify the most suitable technique tailored to the specific requirements of the pOM2S solution.

3.1. Dataset Building

Building on our research objective, we begin with the assumption that both the primary OMS and backup candidates are up and running in the 5G network, enabling us to determine the precise amount of OMS-specific data required for a successful switchover. To develop ML-based models, we created a dataset incorporating features relevant for predicting switchover costs. A testbed was implemented to simulate diverse switchover scenarios between OMSs, replicating various conditions, including workload variations in backup infrastructures and network pattern fluctuations in a controlled environment. Table 2 presents the dataset structure, defining the features that make ML models training possible to estimate the time required to perform all essential operations for transferring all OMS data from the primary instance to a backup option.

	Attributes of candidate instances computing load		Network condition attributes between primary OMS and a candidate instance			Attribute Class
Attribute	CPU Load	Memory Load	Latency	Latency Variation	Packet Loss	Switchover period
Range	0%-100%	0%-70%	0ms-15ms	0ms-5ms	0%-2%	$\Delta(m)$

Table 2. Structure of the dataset built from the testbed.

The testbed for dataset construction was designed based on the following premises:

1. The disaster recovery plan for the pOM2S must balance cost efficiency with OMS resource requirements. This implies that the multi-redundancy model employs warm-standby candidate OMSs without data synchronization from the primary OMS.
2. Activation of an OMS backup via control messages and data transfer depends on computational load metrics (e.g., CPU, memory, and storage utilization) of target infrastructures.
3. OMS data/resource synchronization is potentially subject to network conditions between primary and backup instances, requiring collection of network performance metrics (e.g., latency, jitter, packet loss).

Figure 1 illustrates the elements of the tested environment used to generate the required database. This representation also shows the main steps performed in each test round, consisting of (1) starting and (2) completing the switchover time calculation.

Between the time measuring stages (1) and (2), Figure 1 depicts the steps involved in the switchover time measuring, namely: (a) retrieving all functional data from the primary OMS, (b) transferring the functional data from the primary OMS instance to the backup, (c) restoring the incoming primary OMS functional data, and (d) initializing the new OMS service with succeeding data synchronization. The total execution time for each step (a-d), constitutes a value $\Delta(m)$ measured in seconds. Each instance represents a test round, where the switchover time, denoted by $\Delta(m)$, is the average of three transfering executions. This attribute is related to another categorized as "computational load of OMS candidate" and "network conditions between the primary and candidate OMSs." For each instance, the value of each attribute varies within specified ranges.

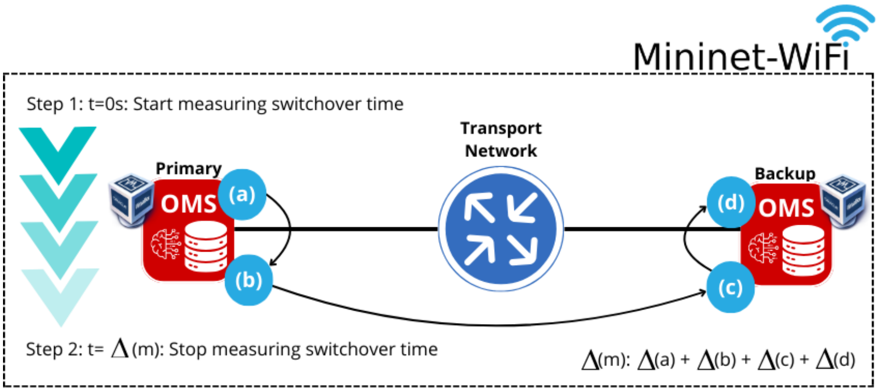

Figure 1. Testbed used for dataset building.

These ranges consider the characteristics of the testbed designed for this research. The instances of the primary and backup OMS candidates consist of virtual machines (VMs) created using Oracle VM VirtualBox[1]. Each VM runs on Linux Ubuntu 20.04, equipped with two Intel(R) Core(TM) i7-1265U 12^{th} Generation virtual CPUs (vCPUs), and 2GB of RAM. In this proof of concept, the performance function from the FCAPS methodology [Chang and Lin 2021] is be considered, meaning that an OMS instance will utilize Prometheus[2] and Grafana[3] for gathering the target monitoring metrics.

Considering that the OMS backup candidate instance may share resources with others independent and isolated services, the OMS instance is virtualized using Docker containers[4]. In this scenario, the OMS monitors an infrastructure emulated by Containernet branching into the Mininet-WiFi[5] tool. The Mininet-WiFi operates on a Dell Latitude 5430 laptop, which shares the same CPU family as the virtual machines and 16GB of RAM.

The volume of monitoring data (i.e., OMS data) generated by Prometheus is 100MB in the experiments, which is transferred to the OMS backup instance via Secure Copy Protocol (SCP)[6]. To control network conditions according to those specified in Table 2, the TC [7] utility is employed. At the OMS backup instance, CPU and RAM loads are stressed by using the stress-ng[8] tool. After running the experiments, a dataset was built comprising 26,990 lines, each representing a test iteration within the virtualized experimental environment, and 6 columns for the independent variables.

3.2. Emulated Testbed Settings

The emulated testbed, initially used for dataset building, was scaled up to include a primary OMS instance and three OMS backup instances for experimental validation. Figure 2 depicts the topology layout and settings of the emulated testbed used in the evaluation experiments.

[1]https://www.virtualbox.org
[2]https://prometheus.io
[3]https://grafana.com
[4]https://www.docker.com/
[5]https://mininet-wifi.github.io/containernet/
[6]https://www.openssh.com/txt/release-8.0
[7]https://man7.org/linux/man-pages/man8/tc.8.html
[8]https://manpages.org/stress-ng

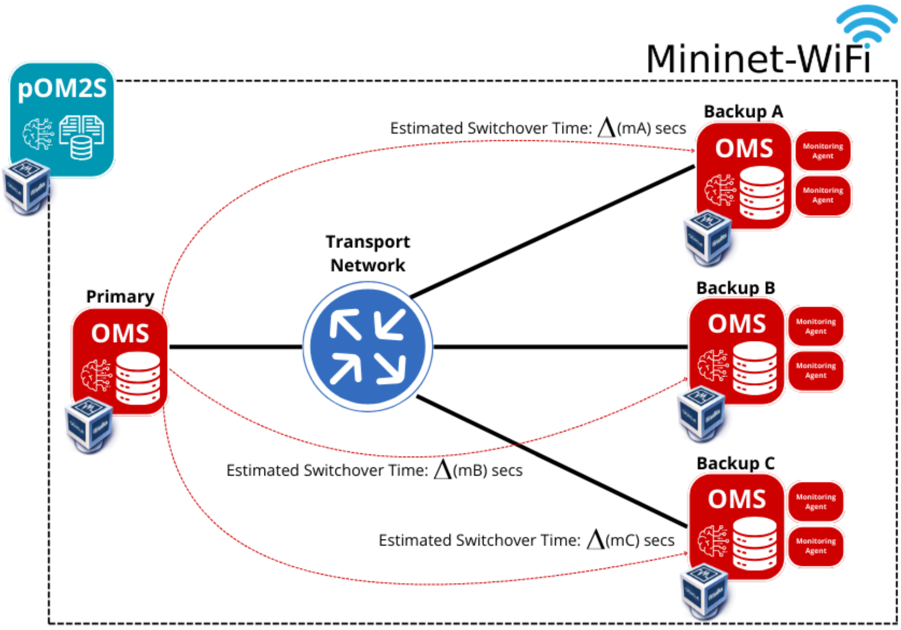

Figure 2. Testbed scaled up: primary OMS and three OMS backup candidates.

While the topology's size and traffic patterns were constrained by the computational resources of the Dell Latitude 5430 laptop (e.g., CPU, memory, and network emulation capabilities), the scaled environment remained sufficiently representative to evaluate ML model performance and validate the pOM2S system. Despite hardware limitations, the testbed replicated critical 5G network dynamics—including workload variations and failover scenarios—ensuring methodological robustness for assessing predictive accuracy and switchover efficiency under controlled yet realistic conditions.

3.3. ML Model Performance Evaluation

In order to train the ML models to predict the switchover period, the Scikit-learn library[9] was used to implement the following algorithms: Linear Regression (LR), Support Vector Machine (SVM), Artificial Neural Network (ANN), and Random Forest (RF). The training process employed 10-fold cross-validation to ensure robust model generalization by evaluating performance across diverse data partitions. After training, the models were evaluated using standard regression metrics: Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), and R^2.

The Mean Absolute Error (MAE) is a widely used metric to quantify the average magnitude of errors between predicted and actual values. A lower MAE indicates a closer fit between model predictions and real outcomes, reflecting higher predictive accuracy. As Figure 3 depicts, RF's lower MAE suggests that, on average, its predictions deviate less from actual switchover times compared to other models. This translates to a more reliable prediction, minimizing discrepancies that could lead to inefficient resource allocation or suboptimal failover responses in the OMS. As a numerical analysis, a true switchover time of 53.13 seconds could yield predictions within the range of 51.37 to 54.89 seconds under RF.

[9] https://scikit-learn.org/

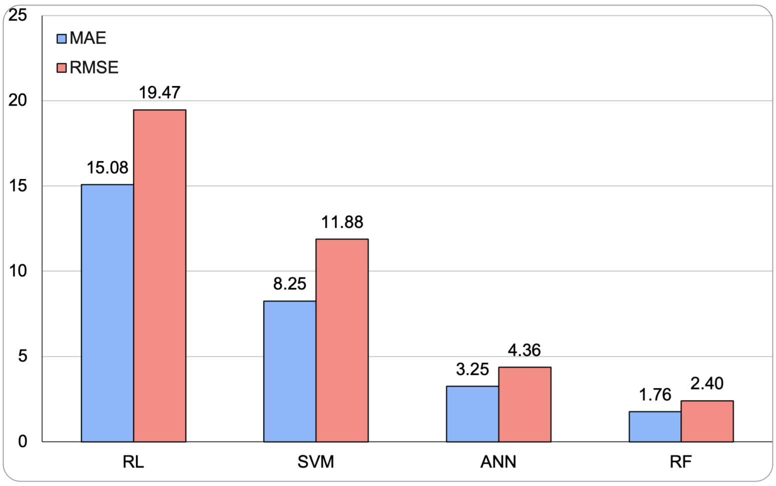

Figure 3. Impact of each ML algorithms in averaging MAE and RSME.

The Root Mean Square Error (RMSE), another fundamental regression evaluation metric, also quantifies prediction errors but assigns greater weight to larger deviations due to the squaring of differences. RMSE penalizes larger errors more than MAE, making it a critical metric for assessing models where occasional large deviations are costly. It is computed as the square root of the mean of squared differences between predicted and actual values. Figure 3 illustrates RMSE values for each algorithm and provides a comparative analysis relative to MAE.

The superior RMSE performance of RF indicates that it not only maintains accuracy across most predictions, but also effectively limits extreme outliers in estimated switchover times. The numerical analysis reveals that the LR model exhibits the highest absolute difference between MAE and RMSE (4.39), while SVM records the largest relative increase in RMSE (43.96%), indicating a higher occurrence of significant prediction errors. In contrast, ANN and RF consistently deliver the most reliable RMSE values, with an average increase of 35.31% over MAE, reinforcing their robustness in switchover time estimation.

Also known as the coefficient of determination, the R^2 Score is a key evaluation metric in regression analysis that quantifies the proportion of variance in the dependent variable explained by the independent variables. In this context, the dependent variable corresponds to the estimated switchover time, while the independent variables are the performance KPIs collected through monitoring. The R^2 Score ranges from 0 to 1, where 1 indicates a perfect fit, meaning the model captures all variance in the data, and 0 signifies no explanatory power, implying the model fails to account for any variance. Figure 4 presents the impacting R^2 Score behavior in the model performance evaluation experiments.

The RF model achieves the highest R^2 Score among evaluated algorithms, demonstrating its superior ability to capture variability in switchover time data. This strong correlation between input features and actual switchover duration underscores RF's reliability for real-world deployment. While the ANN approaches RF's performance, LR and SVM yield only moderate accuracy, revealing limitations in their predictive capabilities. These shortcomings could lead to elevated error rates and suboptimal switchover

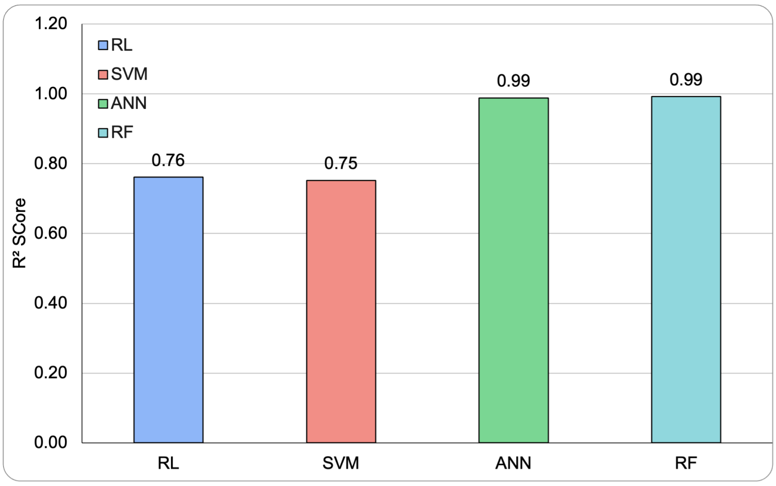

Figure 4. R^2 Score that each ML algorithms impact.

outcomes due to uncertainty in transition time estimates. Figure 5 provide the prediction time behavior that the ML models impact in the testbed.

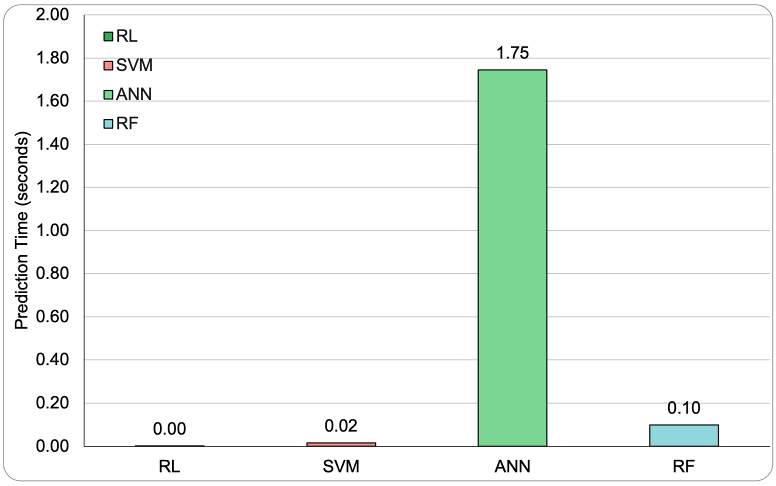

Figure 5. Impacting switchover prediction time behavior of the ML models.

The outcomes of Figure 5, comparing the time required for models to generate predictions on new data—a critical factor for real-time decision-making. The ANN exhibited the worst performance, requiring 1.74 seconds to predict OMS switchover times despite its near-top accuracy. In contrast, LR and SVM delivered the fastest predictions at 1 ms and 15 ms, respectively. While RF performed slightly slower (98 ms), its prediction time remains practical given its superior accuracy.

Beyond accuracy metrics, the RF algorithm is also noted for striking a balance between computational efficiency and predictive accuracy. Unlike more complex deep learning models that might yield marginally better results at the cost of significantly higher computation time, RF provides a practical trade-off—offering strong predictive performance without excessive computational overhead. This efficiency is crucial in real-time or near-real-time applications, such as dynamic network management, where fast and accurate switchover time estimations are essential for minimizing service disruptions and optimizing resource utilization. The reasonable prediction time outcomes of RF suggest that it can be integrated effectively into pOM2S' switchover decision-making processes, ensuring rapid and reliable OMS transitions.

4. Assesment of the pOM2S Effectiveness

After concluding that the RF model yielded higher performance compared to the LR, ANN, and SVM techniques, we proceeded to assess the effectiveness of pOM2S in terms of the accuracy of the ML model when predicting the switchover time for each OMS backup candidate. For the switchover time experiments, we run trials without considering disaster occurrences (since the focus of this paper is on the switchover cost prediction method, not disaster prediction) for each OMS candidate instance following the criteria established in Section 3.1. Three switchover experiments (Exp. #1, Exp. #2, and Exp. #3) are performed per OMS backup candidate instance to identify potential deviations and variations in the process. Table 3 compares the predicted switchover time ($\Delta(m)$) by the pOM2S system against actual exprimental results for three OMS backup candidates.

	$\Delta(m)$ predicted	$\Delta(m)$ real		
Candidate Instance	pOM2S	Exp. #1	Exp.#2	Exp.#3
OMS Backup A	90.88s	87.83s	92.74s	91.43s
OMS Backup B	74.43s	76.50s	74.07s	73.21s
OMS Backup C	52.00s	49.36s	53.08s	50.32s

Table 3. Predicted switchover time of pOM2S for each OMS backup candidate instance.

For OMS Backup A, the predicted switchover time of 90.88 seconds closely aligns with the experimental results, which ranged from 87.83 seconds (Exp. #1) to 92.74 seconds (Exp. #2), with an average deviation of 1.84 seconds. While the prediction slightly overestimated the actual time in Exp. #1 (by 3.05 seconds), it underestimated the duration in Exp. #2 (by 1.86 seconds) and Exp. #3 (by 0.55 seconds), reflecting minor systemic variability in scenarios involving Backup A. OMS Backup B, in turn, exhibited the highest prediction accuracy, with a predicted $\Delta(m)$ of 74.43 seconds and actual times of 76.50 seconds (Exp. #1), 74.07 seconds (Exp. #2), and 73.21 seconds (Exp. #3). The average deviation of 1.20 seconds underscores the model's reliability for this candidate, with the largest discrepancy being a 2.07-second overestimation in Exp. #1. This consistency suggests that Backup B's operational environment, potentially characterized by stable resource availability or predictable network behavior, enables precise modeling, making it a robust candidate for time-sensitive failover operations.

In contrast, OMS Backup C displayed the greatest variability, with actual switchover times ranging from 49.36 seconds (Exp. #1) to 53.08 seconds (Exp. #2), deviating from the predicted 52.00 seconds by an average of 1.76 seconds. The 3.72-second spread between the fastest and slowest experimental results highlights potential instability in Backup C's infrastructure, such as fluctuating network latency or resource contention. Notably, the model underestimated the time in Exp. #1 (by 2.64 seconds) and overestimated it in Exp. #2 (by 1.08 seconds), indicating that environmental factors specific to Backup C may introduce challenges for predictive accuracy.

Across all candidates, the pOM2S system achieved an average prediction error of 1.60 seconds (calculated as the mean of individual deviations), validating its overall effectiveness. However, the results reveal context-dependent limitations: while Backup B's performance demonstrates near-optimal alignment with predictions, Backup C's variability suggests the need for additional calibration to account for dynamic network or computational conditions. Furthermore, the occasional overestimation for Backup A (e.g., Exp. #1) and underestimation for Backup C (e.g., Exp. #1) imply that the model could benefit from incorporating real-time adjustments for transient variables, such as sudden workload spikes or network congestion. These insights underscore the importance of tailoring DR strategies to the unique operational profiles of backup instances, balancing predictive precision with adaptive mechanisms to mitigate uncertainty in heterogeneous 5G environments. Table 4 complements the pOM2S effectiveness assessment through experimental deviations between predicted and actual switchover times ($\Delta(m)$) that each OMS backup candidate take in the testbed.

	Difference to real $\Delta(m)$		
Candidate Instance	Exp. #1	Exp.#2	Exp.#3
OMS Backup A	$+1.86s$	$-3.05s$	$+0.55s$
OMS Backup B	$+2.07s$	$-0.36s$	$-1.22s$
OMS Backup C	$-2.64s$	$+1.08s$	$-1.68s$

Table 4. Experimental deviations of predicted and actual switchover times (s).

Table 4 numerical analysis reveal distinct patterns across the three OMS backup candidates. For OMS Backup A, the pOM2S system overestimated switchover duration in Exp. #1 by 1.86 seconds but significantly underestimated it in Exp. #2 by 3.05 seconds, followed by a minor overestimation of 0.55 seconds in Exp. #3. This inconsistency (spanning a 3.60-second range between the largest over) and underestimation suggests situational variability in Backup A's operational environment, such as transient workload spikes or intermittent network instability that the model fails to fully capture.

OMS Backup B demonstrated the most stable performance, with deviations tightly clustered between +2.07 seconds (overestimation in Exp. #1) and -1.22 seconds (underestimation in Exp. #3). The average absolute deviation of 1.22 seconds across trials highlights the model's reliability for this candidate, likely attributable to predictable resource utilization patterns or stable network conditions. Notably, the -1.22-second underestimation in Exp. #3 remains within acceptable margins for time-critical failover operations, reinforcing Backup B's suitability for high-priority scenarios.

In contrast, OMS Backup C exhibited the largest systemic underestimation, with deviations of -2.64 seconds (Exp. #1) and -1.68 seconds (Exp. #3), interrupted by a +1.08-second overestimation in Exp. #2. The 3.72-second spread between the most extreme deviations underscores inherent instability in Backup C's infrastructure, potentially stemming from variable network latency or resource contention. The persistent underestimation trend (average deviation: -1.08 seconds) raises concerns about the model's ability to account for Backup C's operational unpredictability, which could lead to premature failover decisions in practice.

Wrapping up the study, the absolute discrepancy between predicted switchover times $\Delta(m)$ and actual experimental results remains minimal. Across experiments, deviations ranged from 0.36 seconds faster than predicted (smallest difference), to 3.05 seconds faster (largest difference) with six trials completing switchovers earlier than anticipated. The average deviation of 1.68 seconds aligns with the MAE and RMSE performance evaluation metrics reported in Section 3.3, thus validating the pOM2S RF-based model's reliability.

5. Conclusion and Future Work

This study underscores the strategic importance of advanced switchover methods for DRS in fortifying the resilience of 5G networks, to ensure uninterrupted OMS service continuity. Traditional rule-based approaches approaches, limited by binary "1+1" switchover decision-making, struggle to afford reliable resilience due to the high-complexity of 5G networks dynamic. On the lacking of statistic-based DRS switchover solutions in scientific literature, we put forward the pOM2S method which harnesses ML-driven predictive analytics to predict switchover costs using real-time computing/network metrics, enabling optimal backup candidate selection among multi-redundant options. The premise is to pave the way for proactive DRS approaches, capable of anticipating seamless switchover and quickly responding to disaster events, thereby mitigating service disruptions and ensuring uninterrupted connectivity. The tests and experiments were conducted within a 5G emulated testbed, evaluating the performance of popular ML-based techniques (LR, RF, SVM , and ANN) in carrying out predicitive switchover of multi-redundant OMS backup candidates. Outcomes demonstrated RF's superior accuracy (MAE: 1.68s, R^2: 0.94), where experimental results validate pOM2S's effectiveness in balancing predictive precision and operational practicality, suggesting that the predictive decision-making method relies on a highly accurate model.

Moving forward, several enhancements will be explored to further refine pOM2S' performance. Advancing the ML models to improve predictive accuracy and responsiveness remains a priority. Additionally, addressing the scalability challenges of managing lar-scale OMS multi-redundancy will be a key focus. Finally, real-world prototype trials are considered to validate pOM2S' efficacy in operational 5G deployments, along with verifying its practical viability for high-stakes applications.

Acknowledgment

This research was carried out under the LANCE (Leading Advanced Technologies Center of Excellence) facilities, and partially funded by the *Coordenação de Aperfeiçoamento de Pessoal de Nível Superior - Brasil* (CAPES) - Finance Code 001. Moreover, this work was partially supported by the Sao Paulo Research Foundation (FAPESP), grant 2021/00199-8, CPE SMARTNESS.

References

Abdelaziz, A. S., Harb, H., Zaghloul, A., and Salem, A. (2023). An enhanced mcdm model for cloud service provider selection. *International Journal of Advanced Computer Science and Applications*, 14(2).

Al-Essa, H. A. and Abdulbaki, A. A. (2016). Disaster recovery datacenter's architecture on network replication solution. In *2016 European Modelling Symposium (EMS)*, pages 175–180.

Chang, W.-C. and Lin, F. J. (2021). Coordinated management of 5g core slices by mano and oss/bss. *Journal of Computer and Communications*, 9(6):52–72.

Group, G.-P. S. N. W. (2020). Cloud native and 5g verticals services. Technical report, 5G PPP.

Haleplidis, E., Pentikousis, K., Denazis, S., Salim, J. H., Meyer, D., and Koufopavlou, O. (2015). Software-Defined Networking (SDN): Layers and Architecture Terminology. RFC 7426.

Lawler, C. M., Harper, M. A., and Thornton, M. A. (2007). Components and analysis of disaster tolerant computing. In *2007 IEEE International Performance, Computing, and Communications Conference*, pages 380–386.

Leiter, A., Hegyi, A., Galambosi, N., Lami, E., and Fazekas, P. (2022). Automatic failover of 5g container-based user plane function by onap closed-loop orchestration. In *NOMS 2022-2022 IEEE/IFIP Network Operations and Management Symposium*, pages 1–2.

Liang, M., Gao, F., and Shi, M. (2022). A failure detection method of remote disaster recovery and backup system. In *ITM Web of Conferences*, volume 47, page 01002. EDP Sciences.

Mijumbi, R., Serrat, J., Gorricho, J., Bouten, N., De Turck, F., and Boutaba, R. (2016). Network function virtualization: State-of-the-art and research challenges. *IEEE Communications Surveys Tutorials*.

Saha, M., Panda, S. K., and Panigrahi, S. (2021). A hybrid multi-criteria decision making algorithm for cloud service selection. *International Journal of Information Technology*, 13(4):1417–1422.

Xie, H., Du, S., Zhang, Q., Yang, Y., and Wang, H. (2023). Application of upf disaster tolerant networking in power private network. In *2023 IEEE 6th Information Technology,Networking,Electronic and Automation Control Conference (ITNEC)*, volume 6, pages 811–815.

Zaretalab, A., Hajipour, V., and Tavana, M. (2020). Redundancy allocation problem with multi-state component systems and reliable supplier selection. *Reliability Engineering & System Safety*, 193:106629.

Zheng, Y., Fu, Y., Zhang, Y., Hao, S., and Cai, Y. (2020). Research and practice of intelligent operation and maintenance system for essential load guarantee based on 5g and iot technology. In *2020 4th International Conference on Power and Energy Engineering (ICPEE)*, pages 53–57.

Zhu, H., Li, J., Hu, J., and Li, W. (2022). Failure-aware and automated disaster backup in the 5g core network. In *2022 International Communication Engineering and Cloud Computing Conference (CECCC)*, pages 48–53.

Zunino, C., Cena, G., Scanzio, S., and Valenzano, A. (2024). Adaptive seamless redundancy to achieve highly dependable mqtt communication. *IEEE Transactions on Industrial Informatics*, 20(1):984–994.