This document is a preprint of: A. Schosser, F. Burmeister, P. Schulz, M. Khursheed, S. Ma and G. Fettweis, "Advancing Spectrum Anomaly Detection through Digital Twins," in IEEE Communications Magazine, vol. 63, no. 4, pp. 40-46, Apr 2025. DOI:10.1109/MCOM.001.2400221

Advancing Spectrum Anomaly Detection through Digital Twins

Anton Schösser, Friedrich Burmeister, Philipp Schulz, Mohd Danish Khursheed, Sinuo Ma and Gerhard Fettweis

***Abstract*—6th generation (6G) cellular networks are expected to enable various safety-critical use cases, e.g., in the industrial domain, which require flawless operation of the network. Thus, resilience is one of the key requirements for 6G. A particularly critical point is that 6G, as any other wireless technology, is based on the open radio medium, making it susceptible to interference. Especially intentional interference, i.e., jamming, can severely degrade the network operability. Therefore, a new approach for detecting anomalies in the radio spectrum using a digital twin (DT) of the radio environment is presented in this work. This allows the integration of contextual awareness in the anomaly detection process and is thereby superior to state-of-the-art methods for spectrum anomaly detection. We propose a suitable system architecture and discuss the tasks of machine learning (ML) therein, particularly for reducing the computational complexity and to detect anomalies in an unsupervised manner. The feasibility of the approach is demonstrated by ray tracing simulations. The results indicate a strong detection capability in case of an accurate DT and thereby illustrate the potential of DTs to enhance monitoring of wireless networks in the future.**

I. Introduction

Wireless connectivity is often considered as a backbone of factories of the future, enabling a variety of automation and robotics use cases which help to achieve a new level of flexibility, efficiency, and sustainability in production. However, wireless technologies by definition rely on an open transmission medium, making them vulnerable to external interference. Such external interference, for instance caused by jamming or misconfigured devices, might lead to performance degradation or even failures of the wireless communications systems. This can have serious consequences, for instance, high costs due to production downtime, damage to machinery, or even life-threatening situations for workers in human-machine interaction areas. While there has been tremendous growth in technical capabilities, e.g., increased data rates and reduced latency, the resilience of wireless systems has not kept pace with the widespread application areas and is therefore seen at the forefront of potential issues of the 6th generation (6G) of cellular communications [1]. Resilience hereby denotes the ability of a system to maintain (at least a minimal set of) functionality under adverse conditions.

Resilience is particularly crucial in non-public networks (NPNs), where the network serves as an enabler for critical use cases in the industrial domain, such as automation tasks. Seamless operations are essential in this context. However, traditional network monitoring and root cause analysis (RCA), as applied in conventional cellular networks, are not suitable for NPNs due to the need for expert knowledge and extensive manual work [2]. Thus, conventional network monitoring and RCA do not scale easily for such small-scale networks but would instead lead to high costs, counteracting the objective of cost-efficient production.

This gap between state-of-the-art RCA for cellular networks and the requirements with a particular view on NPN underlines the urge for sophisticated and highly automated solutions for wireless network monitoring. An upcoming concept that has received a lot of attention recently for this problem is denoted as network digital twins (DTs) [3]. They are often seen as one of the key enabling technologies towards 6G, not only enhancing network monitoring but also supporting use cases such as data generation for artificial intelligence (AI) and what-if-analysis. An overview of recent works is provided in [3], but many of those focus on the vision, general requirements, or enabled use cases in various areas of wireless communications. Works towards specific implementations of DTs, particularly at the physical layer, are still at the beginning [4]. To advance the research on using DTs in 6G, this paper contributes in the following ways:

- For the use case of spectrum anomaly detection, it is presented how a DT of the radio environment can be utilized to integrate contextual awareness in the anomaly detection process and thereby advance over state-of-the-art methods.
- The proposed approach is supported by results from a ray tracing-based simulation framework. Moreover, the integration of machine learning (ML) is demonstrated for different tasks within the framework.
- This paper shows opportunities for more extensive studies towards DT-based anomaly detection in the spectrum by identifying open research directions in this topic.

II. Background

This work intertwines the topics of anomaly detection in the radio spectrum and DTs. Therefore, some background information on both topics is provided in this section.

A. Spectrum Anomalies

An anomaly is a sample or group of samples that significantly deviates from the majority in a collection, suggesting a different origin [5]. Anomaly detection aims to identify such samples to recognize unexpected changes in a system. However, due to the complexity of modern communication systems and the rarity of anomalies, developing a robust detection mechanism is a challenge.

The authors are with Vodafone Chair Mobile Communications Systems, Technische Universität Dresden, Germany.

1) Types of Anomalies: Anomalies of the radio spectrum can be on the one hand missing signals when a signal is expected, for example a faulty device [6]. Since the lost link can be identified by the network, this is a trivial case. On the other hand, anomalies can be unexpected signals which cause additional interference. This excludes the variety of interference types, such as inter-symbol interference, inter-carrier interference, or inter-cell interference, which are already taken into account by the design of a wireless communications system. We distinguish two types of unexpected interference: intentional and unintentional. Adding intentional interference to hinder legitimate users' communication via specific channels is referred to as jamming. Due to safety-critical radio applications in the future combined with the availability of inexpensive yet powerful software defined radios, jamming and anti-jamming strategies have received strong research interest throughout the last years [7].

Typically, jammers for communications systems are categorized into naive and smart ones. Naive jammers transmit signals which are not protocol-specific, for example, single or multi-tone jammer, narrow- or wideband noise, or sweep signals. Such jammers are neither difficult to detect nor energy-efficient, which is important for mobile jammers. To improve the stealthiness of naive jammers, the signals can be sent with a random duty cycle.

Unlike naive jammers, smart jammers are designed for a specific protocol. This increases the complexity but at the same time allows unprecedented levels of stealthiness and energy efficiency. One example is a deceptive jammer which transmits a legitimate-like signal and thereby blocks the communications channel, a strategy used to jam decentralized Wi-Fi networks [7]. Establishing a smart jammer for cellular networks typically requires cell synchronization. Once this is achieved, a variety of jamming attacks can be launched, for instance, preventing a user from connecting to the network by jamming the random access channel [7].

Next to jamming, also unintentional sources of interference must also be considered to achieve a comprehensive view of the problem. Such unintentional sources could be misconfigured transmitters (TXs), radio frequency (RF) leakage from cable plants, violation of out-of-band radiation limits, etc. [8]. Moreover, in unlicensed bands, a spectrum policy might be enforced, e.g., it is not allowed to operate a mobile Wi-Fi hotspot in a factory hall. Even though it would be a legitimate use of the band, the additional interference caused by such a network could severely degrade the performance of the industrial wireless network and therefore needs to be detected.

2) State of the Art for Anomaly Detection: The variety of approaches to detect anomalies in the spectrum is almost as wide as the variety of potential spectrum anomalies itself. In the following, we distinguish two categories: general spectrum anomaly detection and jammer detection in particular. Both approaches are typically realized with (sequences of) spectrograms, but there are also methods based on IQ samples or other features. As the relations in the spectrum are rather complex and depend on many factors (e.g., traffic pattern, link adaption, and RF propagation effects), ML is mostly superior in detecting those anomalies compared to conventional methods such as peak detectors [8]. Particularly, unsupervised learning with autoencoder-based architectures has received significant attention recently [6]. The approach has proven to be effective for typical jammers or simple anomalies, but more complex anomalies or devices hiding in legitimate transmit patterns (e.g., deceptive jammers) are rather difficult to capture.

Jammer detection in the spectrum is usually realized using supervised learning to classify spectrograms, providing for each jammer class and normal operations spectrogram examples in the training process. The classification of the jammer type, potentially complemented with the extraction of further features such as bandwidth or center frequency gives indications of how potential countermeasures should be implemented. However, the remaining question is whether it is reasonable to assume that all kinds of jamming types are known during training. To develop more general and yet highly sensitive detection approaches, the algorithms must be able to deal with unknown jammer types equally well in the operational phase.

B. Digital Twins of Communications Networks

A DT refers to a virtual representation of an object or process. It is moreover characterized by its bidirectional connection with its real-world counterpart, the so-called physical twin (PT) [3]. With respect to 6G, DTs have been identified as technology with an astonishing potential for a wide variety of applications, e.g., network planning, monitoring, and optimization or data collection for ML [9]. Therefore, network DTs have received significant attention from the research community and are also in the focus of the standardization bodies, e.g., with ITU-T Recommendation Y.3090 or 3GPP TR 28.915. It is foreseen, that DTs tracking down even to the physical wireless environment will operate in 6G, which has been highlighted recently from a conceptual point of view [4] as well as it has been demonstrated from an implementation point of view [10]. However, to the best of our knowledge, the application of a ray tracing-based DT for spectrum monitoring and anomaly detection, which is proposed in this paper, has not been discussed so far.

III. Proposed Resilience Concept

To meet the demand for a general yet sensitive spectrum anomaly detection system with a high level of sensitivity for different kinds of anomalies, it is beneficial to not only consider the radio spectrum but also integrate further available information in the detection procedure. Simply put, it is easier to detect deviations from normal operation the more information is available to monitor a system. Types of additional information could be:

- Positions of legitimate TXs. In 5th generation (5G) terminology, we denote user equipments (UEs) as well as gNodeBs as legitimate TXs.
- Device orientation and antenna pattern of legitimate TXs in case of directional antennas.

- Assigned radio resources (radio resource management (RRM)).
- Physical environment.

We refer to the consideration of additional information as the integration of contextual awareness in the anomaly detection process. This is achieved through a DT which bundles the collected information and generates an accurate representation of the PT, with a particular focus on the spectrum for the scope of this work. Due to their nature, spectrum anomalies are not represented in the DT and lead to unexpected deviations between the PT and the DT. Thus, we propose to identify anomalies based on such unexpected deviations. To achieve this, the overall system architecture must be designed to provide the DT with the relevant information. Once the detection is successful, countermeasures can be initiated to maintain the network operability. Note that the proposed concept targets NPNs, which cover limited areas and are therefore limited in complexity, and operate in a licensed band. Based on this, we assume that all legitimate TXs are known by the network.

A. System Architecture

The system architecture of the proposed concept is shown in Fig. 1. As shown there, a central unit (CU) collects and processes the aforementioned types of information to run the DT and subsequently does the anomaly detection by comparing the perceived and expected spectrum occupation.

The PT consists of the physical environment (e.g., walls, machines, other objects) and all radio devices. Information about the physical environment can be monitored due to cameras, LiDAR systems of robots, or by joint communications and sensing (JCAS), a promising technology foreseen to be realized in the context of 6G.

Furthermore, there are distributed sensing units (SUs) which continuously monitor the radio spectrum under consideration. They are placed at fixed locations and are equipped with a wired connection to the CU to allow an information flow even in the case of external disturbances. Distributed SUs improve the coverage of the anomaly detection system compared to a single SU. Moreover, in the case of heavy jamming, the low-noise amplifier (LNA) of one or more SUs can get saturated. This makes it easy to detect the anomaly, but difficult to derive further information, e.g., localizing the jammer. With several SUs installed, it is less likely that all of them are saturated, and further information for countermeasures can still be derived. In general, the system architecture allows the integration of mobile SUs or feedback from regular devices which wirelessly share spectrum information with the CU. However, a frequent yet precise feedback might be a burden to the available bandwidth [3].

As 6G standardization is just starting, we give possible origins of the required information according to the 5G standard and the architecture specified by the O-RAN alliance. The information about the intended utilization of the radio spectrum (RRM) must be provided by the scheduler within the central communication entity of the network (gNodeB in Fig. 1). As specified by O-RAN, an xApp might serve as an interface to access the RRM information from the distributed unit (DU). Various approaches are available to to achieve awareness on the location of legitimate TXs. In the context of 5G, localization information can be obtained via requests to the access and mobility function (AMF). Moreover, highly accurate radio-based localization techniques, foreseen to be enabled by 6G, could be employed. In addition, robotic devices typically localize themselves in the environment, e.g., due to LiDAR, and could communicate their positions to the network. However, regardless of how the localization is performed, the position information will always be prone to errors.

With all this information, the DT can continuously be updated and infer an accurate expectation of how the spectrum occupation should look like under normal operation. For this, either ray tracing can be used, which is expected to achieve real-time capability in the near future [4], or ray tracing can be replaced by ML in a computationally efficient way as has already been demonstrated in the literature [11]. As shown in Fig. 1, the DT basically reflects the PT. However, before the detection the jammer is not known and therefore not represented in the DT. This leads to deviations between the PT and the DT, which are used to detect the anomaly. More specifically, spectrum occupancy information can simply be received signal strength (RSS) values for a specific band, as employed in Sec. IV and shown in Fig. 1, or spectrograms in a more sophisticated implementation. The observed spectrum occupation, perceived by the SUs, is then compared against the expected spectrum occupation based on the DT as described in the next section.

B. Workflow

Since the DT will never be a perfect replica of the PT, there will always be deviations between the PT and the DT. Thus, a method is needed to determine whether the deviations come from inaccuracies in the DT model, such as TX localization errors or geometry simplifications affecting ray tracing results, or whether they indicate the presence of an additional interference source, meaning a spectrum anomaly. For this task, we propose an unsupervised learning approach that employs an algorithm to generalize the statistics of normal training data and subsequently identifies anomalies based on deviations from these statistics. In the following, we refer to this as anomaly detection algorithm.

The deployment of the proposed framework consists of two phases. In the first one, referred to as the offline phase, a database of samples from normal operations is collected to derive the statistics of deviations between the PT and the DT under normal conditions, i.e., there is no unexpected interference. This approach allows to detect anomalies without making any assumptions about them in contrast to supervised approaches, where prior knowledge of anomaly characteristics is required. Moreover, collected measurements can be used to refine the ray tracing model and thereby achieve a better correspondence between PT and DT. Already in this phase, common outlier detection methods must be used to ensure clean training data, as a contamination of the data in this

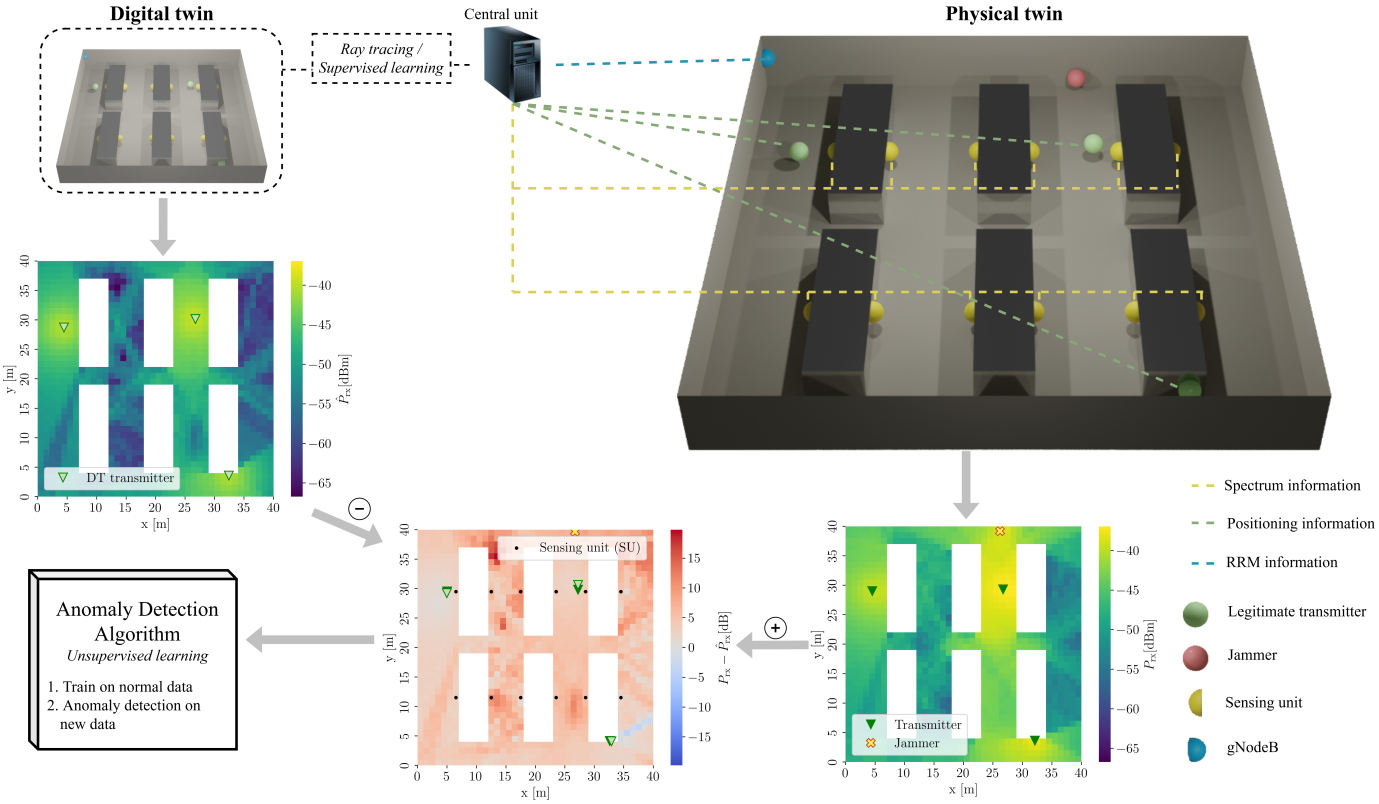

Fig. 1: Proposed system architecture visualized at the example of the scenario used for simulation in Sec. IV. The spectrum occupation is indicated exemplarily by the RSS which is also used for the simulation.

phase with anomalies would poison the model and degrade the performance of anomaly detection in the online phase.

After the offline phase, i.e., when sufficient data are collected and the anomaly detection algorithm is trained, the online phase starts. The DT continuously mirrors the PT and the data from comparing PT and DT are processed with the anomaly detection algorithm to identify spectrum anomalies. Online learning can be considered in this phase for further refining the anomaly detection algorithm, but this also poses the danger of an attacker poisoning the algorithm over time.

C. Potential Countermeasures

To maintain network operability in case a spectrum anomaly is detected, different countermeasures can be initiated, depending on the type of anomaly, and the technical capabilities of the communications systems.

If the interference is limited to a narrow band, this information might be taken into account by RRM, i.e., subcarriers affected by the anomaly are not allocated. In the case of wideband interference, RRM can no longer mitigate the interference. However, if the radio system is supporting various frequency bands, a change to a distant frequency band might be an option. If the jammer transmits directional, i.e., it targets a base station (BS), e.g., by beamforming, in the spatial domain, triggering a hand-over of legitimate users to neighboring BSs is a valid option.

If a severe interference level is reached where no communication is possible, countermeasures, e.g., removing the interference source or proactive counter-jamming the receiver of a smart jammer, are required. For this, localization of the anomaly source is required. Compared to the localization of a cooperative transmitter, the localization of the anomaly source poses additional challenges. This is mainly the unknown characteristics of the of the anomalous transmitter, i.e., waveform, transmit power, antenna pattern etc. are unknown. Moreover, the anomalous signal is potentially superimposed by legitimate signals. In addition, factory environments, which are considered one of the main use cases for NPNs, are typically characterized by non-line of sight (NLOS) conditions and strong multi-path effects. Approaches to overcome those challenges might be exploiting the contextual awareness which is provided by the DT and using supervised ML which has shown promising results with respect to the simultaneous blind localization of multiple radio transmitters in such challenging scenarios [12].

IV. Proof of Concept

After presenting the proposed resilience concept, a proof of concept using ray tracing simulations is implemented in this section.[1]

A. Scenario and System Model

The simulation is oriented towards an NPN deployed in a factory hall. The model of the hall with six production lines, simplified as metal cuboids, is shown in Fig. 1. We assume

[1]The code can be found at https://github.com/akdd11/advancing-spectrum-anomaly-detection.

TABLE I: Simulation parameters

Parameter		Value
Geometric properties		
Factory hall		40 m × 40 m × 5 m
		Concrete
Production lines		14 m × 4 m × 3 m
		Metal
Radio properties		
Carrier frequency	f_c	3.7 GHz
Number of legitimate TXs	N_{leg}	3
Transmit power of legitimate TXs	$P_{\text{tx, leg}}$	20 dBm
Number of jammers	N_{jam}	0 or 1
Transmit power of jammer	$P_{\text{tx, jam}}$	20 dBm
Number of sensing units	N_{SU}	12
Antenna patterns		isotropic
Noise floor	$P_{\text{rx, noise}}$	−100 dBm
Material properties		ITU-R P.2040-2

three legitimate TXs, positioned at random locations inside the factory hall. Further parameters follow specifications for NPNs in Germany. In case a jammer is present, it is exemplarily assumed that it mimics a legitimate TX and therefore has the same properties. For the scope of this work, we assume that all legitimate TXs as well as the jammer are located at a height of 1.5 m. The Python module *Sionna* is employed for ray tracing [13].

To perceive the spectrum occupation, 12 SUs (which seems a number causing reasonable complexity when considering the monitoring area of 1600 m^2) are distributed at fixed locations as shown in Fig. 1. In a previous work it has been shown that further increasing the number of SUs compensates for inaccuracies in the DT model and improves the performance [14]. They are distributed regularly over the area and are placed at the edges of the production lines to ensure wiring. In a first scenario, the height of the SUs is set to the same height as the TX and jammer, i.e., 1.5 m, whereas in a second scenario a height of 5 m is considered as if they are mounted at the ceiling. Mounting the SUs at the ceiling could improve the general detection performance, as large parts of the area are covered under line-of-sight (LOS) conditions, whereas jammers with a directional antenna tilted to the ground are potentially more difficult to detect.

For the scope of this work, the following assumptions are used. The legitimate TXs use different subbands within the monitored frequency band. The legitimate signals are superimposed with the jamming signal in case a jammer is present. The SUs measure the wideband received power in the monitored band, which we refer to as received signal strength (RSS) $P_{\text{rx},j}$, with j denoting the index of the SU. The simulation parameters are summarized in Table I.

In the next step, the DT of the radio environment is generated. As the work is simulation-based, the same ray tracing framework as previously explained is employed. Inaccuracies are modeled to take the deviations between the DT and the PT into account. First, the localization inaccuracy is implemented by adding a random offset to the TX location, which is normally distributed with $\mathcal{N}(0,\, 0.37\,\text{m})$ in magnitude and uniformly distributed in the angle [15]. The RSS values estimated by the DT with taking into account the TX localization error are denoted as $\hat{P}_{\text{rx},j}$. Further deviations between the PT and the DT are modeled for each SU with an error term ϵ_j with a log-normal distribution. This term reflects for instance errors in the material properties or geometric simplifications of the physical environment. This results in a difference vector $\boldsymbol{\Delta}$ in the logarithmic domain defined by its entries

$$\Delta_j = P_{\text{rx},j} - \hat{P}_{\text{rx},j} + \epsilon_j, \tag{1}$$

which serves as input for the anomaly detection algorithms. Note, that jammers are not reflected in the DT, as they are not known initially, but the task is to detect them by identifying significant deviations in the statistics of $\boldsymbol{\Delta}$.

The whole process is depicted in Fig. 1 with an example radio map of real-world RSS values – corresponding to the visualized scenario – on the right side, indicating moreover the true locations of the legitimate TXs and the jammer. The CU processes the collected information using the DT to generate an estimation of the RSS values, shown on the left side. The difference between the true and estimated RSS values is shown in the middle, where the color reflects the difference of the estimated RSS values. There, also the slight discrepancies between the true and estimated TX location can be seen. The PT values of the RSS and therefore also difference between the observed and estimated RSS is only known at the SU locations, even though here the whole map is plotted for visualization purposes.

B. Anomaly Detection Algorithms

Two algorithms for anomaly detection are compared in this work. Their task is to generalize the statistics of data from normal operations (in the following referred to as normal data) and detect anomalies in new data subsequently. The explanation of the algorithms is kept short here, a more elaborated explanation is provided in [14]. Both algorithms use the vector $\boldsymbol{\Delta}$ which reflects the differences between the PT and DT RSS values at SU locations as input.

1) Local outlier factor (LOF): LOF is a common score-based algorithm for unsupervised anomaly detection. The LOF score is calculated by setting the local density of a point in relation to the local densities of its neighboring points. Points with a significantly lower density are regarded as anomaly. Here, we apply LOF for novelty detection, i.e., the training samples (all assumed to be normal) are used as references and the LOF for a test sample is calculated with respect to the training samples only.

2) Energy detector (ED): The ED is usually applied in signal detection tasks. Here, it is employed on the difference vector $\boldsymbol{\Delta}$ to distinguish whether deviations can be attributed to modeling inaccuracies in the DT or whether the deviations are so big that they are suggesting that an anomalous TX is present. For this, the mean of $\boldsymbol{\Delta}$ is compared against a threshold Δ_{th}, and if the mean is greater than the threshold, an anomaly is assumed. As an analytical derivation of Δ_{th} with the given problem formulation is infeasible, we derive it empirically from the training data analyzing the distribution of

the mean values of $\mathbf{\Delta}$ in the train set and specifying a target false positive rate (FPR).

C. Supervised ML as Computationally-Efficient Replacement for Ray Tracing

Even though accurate real-time radio-frequency ray tracing is foreseen in the future [4], we additionally compare the ray tracing-enabled DT against a ML model learned for path loss estimation (PLE), which offers significantly improved computational efficiency required to meet the latency constraints of the anomaly detection system. The feasibility of learning ML models for PLE has been shown in the literature several times [11], therefore our focus is not on achievable accuracy but rather on its application in a DT-based anomaly detection system. For this reason, the comments on the topic will be kept short and only problem-specific aspects will be highlighted.

The task of the ML model here is to estimate the path loss between the TX position, which is the input, and the SU. Our experiments have shown that a multi-output regressor, i.e., one separate model for each SU, leads to the best results, and already a simple random forest model yields sufficient estimation accuracy for the examined problem (see Section IV-D). The approach is as follows. First, 5000 path loss maps for random TX positions are generated. Those are used to train the ML model, which is considered as offline phase. Please note that a static environment is considered here, but ML for ray tracing has also shown strong performance with varying environments. Once the DT is in operation, the ML model is used for inference. While computing a path loss map for one TX using ray tracing takes 1.95 s on the used PC and configuration, inferring the path loss between the TX and the twelve SUs takes 0.02 s. It shall be emphasized that the focus of this work is not on a low-latency implementation and we therefore see huge potential for improved computational efficiency in the future.

D. Simulation Results

To gain insights into the performance of our proposed approach, we first generate 2800 normal (i.e., without jammer) realizations of the scenario as described in Sec. IV-A. Hereby, each realization is represented by one vector $\mathbf{\Delta}$. Those normal vectors are used to train the anomaly detection algorithm. For testing, 600 additional realizations are generated, with half of them being anomalies. The anomaly detection performance is evaluated employing the receiver operating characteristic (ROC), which is not affected by the proportion of anomalies in the test set. To obtain those curves, the anomaly detection threshold is varied. In particular, for LOF, the score above which a sample is considered an anomaly is varied, whereas for the ED the threshold Δ_{th} is varied. The procedure is repeated five times to improve the robustness of the results.

Several conclusions can be drawn from the ROC curves shown Fig. 2, with a height of the SUs of 1.5 m in Fig. 2a and 5 m in Fig. 2b, respectively. First, one can see that the approach achieves a strong anomaly detection performance in the case of an accurate model underlying the DT, which degrades with

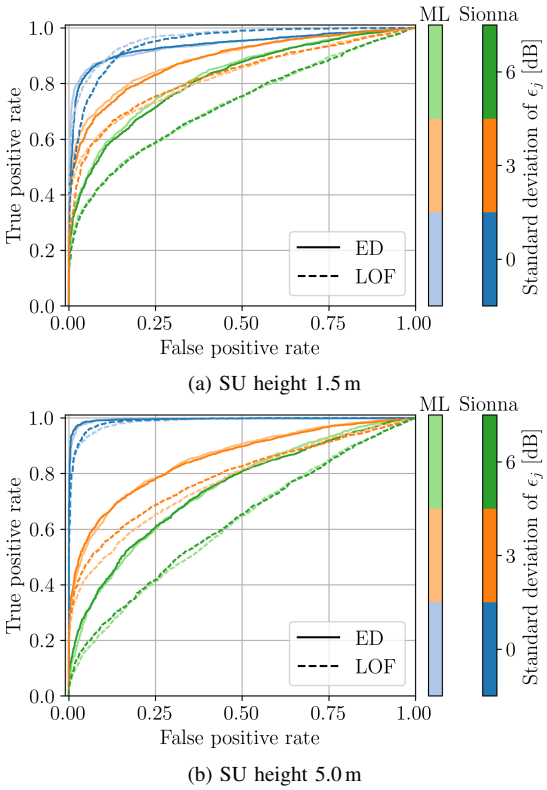

(a) SU height 1.5 m

(b) SU height 5.0 m

Fig. 2: ROC curves for different heights of the SUs

an increasing level of model inaccuracy. Therefore, a focus when applying this concept should be on an accurate model for the DT. Regarding the anomaly detection algorithms, which identify the anomalies based on the input vector $\mathbf{\Delta}$, ED outperforms LOF in almost every considered case. This is probably because ED is more problem-tailored than the general algorithm LOF. Analyzing the effect of the method used for PLE, one can say that the performance of the ML-based DT is comparable to the ray tracing-based DT. Thus, using ML in DTs covering radio propagation seems to be a reasonable approach to reduce computational complexity, particularly in cases with a high number of TXs. Finally, we will discuss the impact of SU height. As shown in Fig. 2b, placing the SUs at the ceiling together with a precise DT leads to superior performance. However, with increasing modeling inaccuracies, the detection performance drops even below that of SUs at 1.5 m. This could be for example caused by a higher portion of LOS connection for higher SUs, and thus the jammer's impact on the RSS in the log-domain gets relatively smaller compared to the error variance. These findings suggest placing SUs at various heights, which we plan to explore further.

V. Research Directions

The proposed concept shows promising anomaly detection performance in the presented simulation framework, motivating further research in the direction of DTs for physical layer monitoring. Sophistication of the approach in more complex scenarios with numerous heterogeneous devices should be a focus. Moreover, fine-grained resolution

in time and frequency needs to be investigated. While these aspects are particularly important for physical layer DTs, for enhanced anomaly detection integrating higher layers in the DT appears promising. Some research on these layers has already been conducted, making their combination particularly intriguing. Considering DTs not only for the presented use case but also for other physical layer tasks, further experimental work needs to be carried out to achieve reliable insights on the accuracy that can be achieved by ray tracing.

As discussed in Section III-C, jamming mitigation methods might not suffice in case of severe jamming but the elimination of the jammer is required, which in turn requires its localization. Thus, we propose further research towards the localization of non-cooperative TXs in complex environments. One key aspect hereby could be the design of ML methods and corresponding training sets which achieve localization of a variety of uncooperative TXs or how the DT can be exploited for such localization tasks.

VI. Conclusion

In this paper, we proposed to employ a DT of the radio environment to detect anomalies in the spectrum. Particularly in the case of an accurate model underlying the DT, the proposed approach shows a strong detection performance in the investigated scenario. This is achieved by integrating contextual awareness in the detection process, i.e., knowledge of legitimate TXs and the physical environment. Moreover, the application of ML in the framework has been discussed, for computationally efficient handling physical layer DTs as well as for identifying anomalies from data in an unsupervised manner. Hence, the proposed approach offers a viable option for improving network resilience within NPNs, while also contributing towards network DTs for 6G. In the future, our work will focus on extending the anomaly detection framework to the time-frequency domain to enable a benchmark comparison with other approaches and study the detection performance particularly for smart jammers.

Acknowledgment

This work was partially supported by the Federal Ministry of Education and Research, Germany (BMBF) as part of the projects 6G-CampuSens (16KISK207) and 6G-life (16KISK001K) as well as by the European Commission through the Horizon Europe/JU SNS project Hexa-X-II (Grant Agreement no. 101095759). The authors alone are responsible for the content of the paper.

References

[1] H. Tataria *et al.*, "6G wireless systems: Vision, requirements, challenges, insights, and opportunities," *Proceedings of the IEEE*, vol. 109, no. 7, pp. 1166–1199, 2021.
[2] R. Shafin *et al.*, "Artificial intelligence-enabled cellular networks: A critical path to beyond-5G and 6G," *IEEE Wireless Communications*, vol. 27, no. 2, pp. 212–217, 2020.
[3] L. Bariah, H. Sari, and M. Debbah, "Digital twin-empowered communications: A new frontier of wireless networks," *IEEE Communications Magazine*, vol. 61, no. 12, pp. 24–36, 2023.
[4] A. Alkhateeb, S. Jiang, and G. Charan, "Real-time digital twins: Vision and research directions for 6G and beyond," *IEEE Communications Magazine*, vol. 61, no. 11, pp. 128–134, 2023.
[5] L. Akoglu, H. Tong, and D. Koutra, "Graph based anomaly detection and description: a survey," *Data Mining and Knowledge Discovery*, vol. 29, no. 3, p. 626–688, 2014.
[6] S. Rajendran *et al.*, "Unsupervised wireless spectrum anomaly detection with interpretable features," *IEEE Transactions on Cognitive Communications and Networking*, vol. 5, no. 3, pp. 637–647, 2019.
[7] H. Pirayesh and H. Zeng, "Jamming attacks and anti-jamming strategies in wireless networks: A comprehensive survey," *IEEE Communications Surveys & Tutorials*, vol. 24, no. 2, pp. 767–809, 2022.
[8] Z. Li *et al.*, "Scaling deep learning models for spectrum anomaly detection," in *Proceedings of the Twentieth ACM International Symposium on Mobile Ad Hoc Networking and Computing*. Catania, Italy: ACM, 2019.
[9] X. Lin *et al.*, "6G digital twin networks: From theory to practice," *IEEE Communications Magazine*, vol. 61, no. 11, pp. 72–78, 2023.
[10] P. Testolina *et al.*, "Boston twin: the boston digital twin for ray-tracing in 6G networks," in *Proceedings of the 15th ACM Multimedia Systems Conference*, ser. MMSys '24. New York, NY, USA: ACM, 2024.
[11] C. Pyo, H. Sawada, and T. Matsumura, "A deep learning-based indoor radio estimation method driven by 2.4 GHz ray-tracing data," *IEEE Access*, vol. 11, pp. 138 215–138 228, 2023.
[12] I. B. F. de Almeida *et al.*, "Blind transmitter localization using deep learning: A scalability study," in *2023 Wireless Communications and Networking Conference (WCNC)*, Glasgow, United Kingdom, 2023.
[13] J. Hoydis *et al.*, "Sionna: An open-source library for next-generation physical layer research," *arXiv preprint*, 2022.
[14] A. Krause *et al.*, "Digital twin of the radio environment: A novel approach for anomaly detection in wireless networks," in *2023 IEEE Globecom Workshops: 3rd Workshop on Sustainable and Resilient Industrial Networks*, Kuala Lumpur, Malaysia, 2023.
[15] Y. Wang, K. Zhao, and Z. Zheng, "An improved 3D indoor positioning study with ray tracing modeling for 6G systems," *Mobile Networks and Applications*, vol. 28, no. 3, p. 1162–1175, 2023.

Anton Schösser received his Dipl.-Ing. degree in electrical engineering in 2022 from TU Dresden, and is currently pursuing his Ph.D. at the Vodafone Chair Mobile Communications Systems at TU Dresden. His research interests include spectrum monitoring and ML for wireless communications.

Friedrich Burmeister received his Dipl.-Ing. degree in electrical engineering in 2020 from TU Dresden, and is currently pursuing his Ph.D. at the Vodafone Chair Mobile Communications Systems at TU Dresden. His research interests include industrial radio channel measurements and ML for wireless communications.

Philipp Schulz received the M.Sc. degree in mathematics and the Ph.D. (Dr.-Ing.) degree in electrical engineering from TU Dresden, Germany, in 2014 and 2020, respectively. There, he is currently a Research Group Leader with the Vodafone Chair Mobile Communications Systems and focuses on resilient wireless communications.

Mohd Danish Khursheed received his Bachelors in Electronics Engineering from Aligarh Muslim University in 2022 and is currently pursuing his Master in Nanoelectronic Systems at TU Dresden. His research interests include Machine Learning and Digital ASIC Design.

Sinuo Ma received his Dipl.-Ing. degree in electrical engineering in 2022 from TU Dresden, and is currently pursuing his Ph.D. at the Vodafone Chair Mobile Communications Systems at TU Dresden. His research interests include communication signal processing and ML for wireless communications.

Gerhard Fettweis [F'09] earned a Ph.D. at RWTH Aachen in 1990. After a postdoc at IBM Research, San Jose, CA, he joined TCSI, Berkeley, CA. Since 1994, he is Vodafone Chair Professor at TU Dresden. Since 2018, he is founding director of the Barkhausen Institute.