Machine-Learning-Based Predictive Handover

Ahmed Masri*, Teemu Veijalainen†, Henrik Martikainen*, Stephen Mwanje†, Janne Ali-Tolppa†, Márton Kajó‡
*Magister Solutions, email: {ahmed.masri, henrik.martikainen}@magister.fi
†Nokia Bell Labs, email: {teemu.veijalainen, janne.ali-tolppa, stephen.mwanje}@nokia-bell-labs.com
‡Technical University of Munich, email: kajo@net.in.tum.de

***Abstract*—Good mobility performance is critical in cellular networks for ensuring that each user is always connected to the best possible cell and that the handovers are executed timely to minimize radio link failures while avoiding unnecessary handovers. A potential new approach to this challenge is to learn the local radio conditions and adapt and customize the mobility to them. This work proposes and investigates a machine learning method for learning the optimal time and destination for handovers in 5G radio networks, as well as how to use the learned model to trigger handovers based on the predicted radio conditions. The complete solution is analyzed and compared to the state of the art mobility methods to evaluate its performance in reducing the system total outage.**

***Index Terms*—5G, Supervised Learning, Predictive Handover, Neural Networks, Long-Short Term Memory**

I. Introduction

A major driver for Fifth-Generation New Radio (5G-NR) is to provide Ultra-Reliable Low-Latency Communication (URLLC) [1], required, for example, in reliable factory automation, remote control, smart grids, self-driving cars and any critical tasks that rely on close-to-real-time communication. Handovers are a major cause of interruptions in mobile communications, so for URLLC services it is critical that these are minimized.

Thus, URLLC communication requires a very low Mobility Interruption Time (MIT), defined by the 3GPP as the time during which a user terminal cannot exchange user plane packets as it is handed over from one cell or base station to another [2]. For a handover (HO) event, the MIT (T_{MIT}) consists of two parts as given by Eq. 1:

$$T_{MIT} = (1 - P_{HOF}) * T_{HIT} + P_{HOF} * T_{HOF}, \qquad (1)$$

where, T_{HIT} is the interruption time during a successful HO and the T_{HOF} is the interruption time in case of a Handover Failure (HOF) or a Radio Link Failure (RLF). A HOF occurs, if a User Equipment (UE) is handed over too early or to a wrong target cell, and RLF when a HO is triggered too late or not at all, leading to corresponding possible options of recovery [1]. The value P_{HOF} is the probability of either a HOF or a RLF occurring during a HO. The total accumulative MIT experienced by a UE is the product of T_{MIT} and the number of HOs.

The total T_{MIT} can be reduced by reducing the P_{HOF} or the number of unnecessary HOs, especially the ping pongs. In LTE networks the typical T_{HIT} is reported to be about $50\,ms$ [3] [4], while the T_{HOF} ranges from a few hundred milliseconds to several seconds. As the T_{HOF} has higher impact on the total T_{MIT} than T_{HIT}, it follows that reducing the P_{HOF} will contribute more to reducing the T_{MIT}, while keeping in mind that increasing the likeliness of handovers could cause many unnecessarily ping-pongs that may accumulate and then increase T_{MIT} just the same.

Fig. 1a shows the current reactive HO mechanisms. HO execution is delayed using the Time-To-Trigger (TTT) (e.g. $200-300$ ms) and the cell-pair-specific offset (e.g. $1-3$ dB), which however may cause an RLFs if the delay is long. On the other hand, shorter TTT and smaller offset may lead to too early triggering and/or triggering HOs to sub-optimal targets. Classical solutions, such as Mobility Robustness Optimization (MRO), adjust HO parameters based on mobility Key Performance Indicators (KPIs). However, these solutions provide generic cell-pair-level configurations that do not account for differences in on UE speed, trajectory across the cell border and temporal changes in the radio propagation environment. On the other hand, the management effort required to use the HO parameters to configure more customized HO behavior that is fit the 5G URLLC use cases is prohibitive, since it requires that the mobility parameters be configured for each speed and trajectory crossing the cell border.

Successes in using Machine Learning (ML) solutions for mobile network automation have motivated its use for automation in 5G technology [5], [6], as it offers solutions for achieving the required adaptability with low management cost. We have studied the possibility of using ML to learn the mobility policy based on UE measurements and predefined performance objective, instead of using the predefined measurement events of baseline HO. This allows the mobility to be dynamically optimized for each UE and situation while considering the predicted radio conditions. To evaluate this, we developed a predictive classifier, which takes UE measurements and determines the optimal point in time and the target cell to HO, in order to minimize the total MIT. The concept was evaluated using simulations in a challenging industrial 5G network setup.

In the rest of the paper, we review the related work in Section II and describe the predictive ML-based HO method in Section III, giving a comprehensive comparisons between our predictive HO method against baseline HO procedures and classical MRO [7]. We then respectively describe the test environment and results of an implementation of the predictive HO idea in Sections IV and V, before we end with concluding remarks and an outlook to future work in section VI.

978-3-903176-32-4

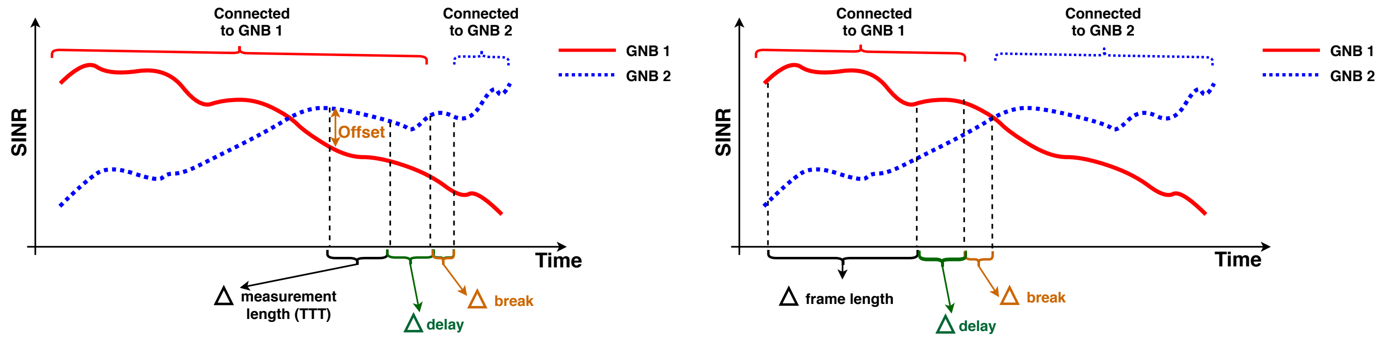

(a) Reactive HO

(b) Predictive HO

Fig. 1. Different HO mechanisms

II. Related Work

Several publications have applied ML approaches to optimizing the performance of HOs without changing the reactive HO regime. Most of these works use Reinforcement Learning (RL) ML techniques, in which a model is trained online during interaction with environment, without the need for prepared training data. For example, authors in [8] applied a RL technique called ϵ-greedy Q-learning toward learning the optimal HO policy, which maximizes the future throughput expected under the locations and velocities of a pedestrian. Their work depends on slow moving users and a human tracking module, e.g. an RGB-D camera, which makes their RL model customized for such non common scenario. However, RL techniques have one main drawback, i.e. that the learn-by-experience approach requires a trial and error period, which cannot be tolerated by most real-life environments [9]. For this reason, we focus our work in this paper on supervised ML, in which training data is collected and an offline model trained, after which an online prediction using the trained model can be executed [10].

Authors in [11] proposed an approach toward improving the conditional HO technique by predicting the target cells to be prepared for a possible upcoming HO. Their results show a promising improvement towards reducing the MIT. However, the approach was tested with HO triggering still using the legacy LTE HO methods, in a simplistic environment that is not representative of typical 5G URLLC conditions.

Authors in [12] proposed to assist HOs in mmWave vehicular networks, by using historical HO data and a simple ML algorithm to determine the internal relationship between a vehicles' status information when requesting HOs and the final HO decisions. However, their proposal optimized for vehicular networks in a specific scenario, and cannot be generalized for other challenging 5G scenarios.

The work in [13] focuses on multi-user, multi-step trajectory prediction using the Long-Short Term Memory (LSTM) supervised ML technique. Predicting the user's future location provides important information towards reducing T_{HOF}. However, their achieved user location prediction, in the order of dozens of seconds to a few minutes, is too long for industrial environments, e.g. to predict the quick jumps between cells within the few hundreds of milliseconds when user is moving between machines with heavy shadowing effects.

III. ML Classification Based Predictive HO

A. The Predictive HO Concept

The aim of the predictive HO is to improve mobility performance over state-of-the-art mobility methods (including MRO) by learning and optimizing the triggering of HOs for a particular mobility environment. As shown in Fig. 2a, the model takes as input UE Reference Signal Received Power (RSRP) measurements from K specific cells. The model implicitly finger-prints the RSRP signals and learns to predict the probability that a given cell will have the best RSRP by a future time instance J. This fingerprint is specific to the geographical area. A HO to a cell C is recommended if cell C is not the serving cell and has the highest predicted probability, thus, the model learns a K-class classification, to be learned by a LSTM neural network.

B. Label Generation

A key aspect in classification problems is how the labels are obtained. In our approach, the labels are estimated offline, from recorded Signal-to-Interference-plus-Noise Ratio (SINR) measurements. Fig. 2b, demonstrates the labeling approach. Each cell is assumed to generate interference for the target candidate cell and the corresponding SINR is mapped to a capacity estimate, similar to Shannon's capacity equation, so that the cell with the highest expected capacity over the samples from $t+i$ to $t+M$ is chosen as the best.

SINR is used to evaluate the spectral efficiency of each cell at a given point in time, where the spectral efficiency is integrated over the labeling window. Additional penalties are defined for handovers, such as UE receiving zero capacity during the interruption, which ensures to only trigger the necessary HOs and subsequently limit the ping-pongs. Weights are used to balance the tradeoff between spectral efficiency and outage from handovers or failures. Different labeling methods are a candidate for further study.

At time t, the input for the model is N previous measurements from the K cells, which we call the input frame and represents the samples that will be available to the model at inference. The label for this input is calculated using the M

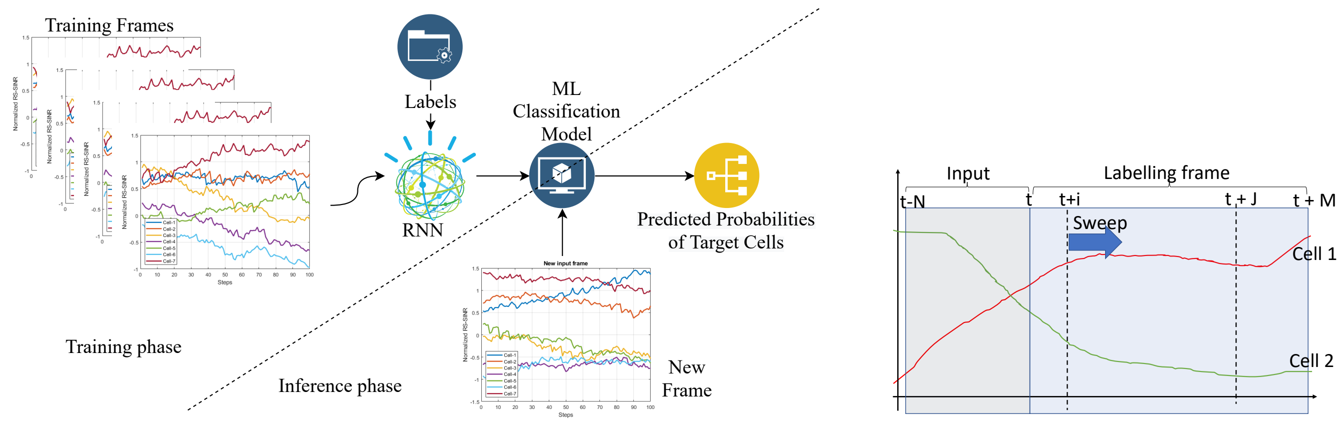

(a) Predictive HO concept (b) Label generation

Fig. 2. Predictive HO

samples after the input frame, which make the labeling frame. Goodness of a HO is estimated for each point i within the M samples by using samples from i to $i+M$. Connection to the serving cell is assumed until $t+i$ after which we then evaluate, using the remaining $M-i$ samples, what would happen if HO is triggered at i towards each candidate cell L.

C. Filtering Classification Decisions using a Dynamic Threshold

An unnecessary HO may generate e.g. an RLF if HO is triggered towards a weak cell, or ping-pong if HO is triggered towards sub-optimal cell, after which the classifier triggers HO back to the optimal cell. The typical way of controlling such false classifications is to add a threshold to the output (after the softmax layer), acting as a confidence minimum which the classifier needs to reach before a handover is triggered. Our proposed method, called Dynamic Confidence Threshold (DCT), dynamically adjusts this threshold during runtime to balance between false positive and false negative HO triggers.

DCT is a simple threshold-setting scheme, which works on the exponential moving average of the measured RSRP-based Signal-to-Interference-plus-Noise Ratio (RSSINR), and adjusts the threshold to higher or lower values by predefined steps. DCT helps to filter out many of the inaccuracies of the classification as shown in the results section. Moreover, since the labeling is based on the estimate of HO performance, the threshold can be fine-tuned online to have better network performance. For example, if the classifier produces too many HOs, the threshold can be increased to filter more of the uncertain situations.

IV. Evaluation Environment and Scenario

The objective of the simulations was to evaluate the performance of the predictive HO model both with and without the classification threshold (DCT) and to compare that performance to the legacy solutions - the baseline with A3 triggering and when applying MRO. Specifically, the study was focused on URLLC services requiring that an industrial environment is considered.

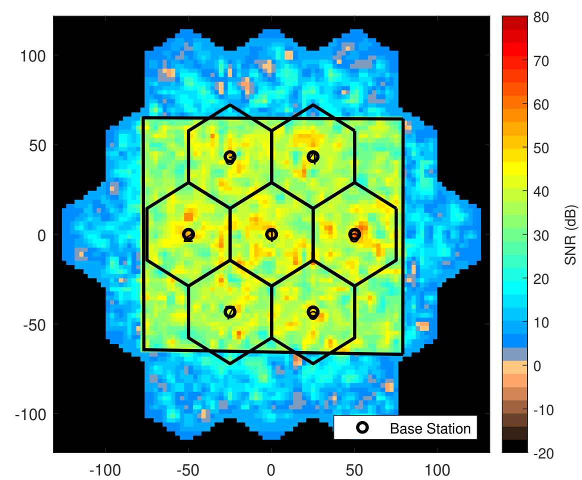

Fig. 3. Signal-to-Noise Ratio (SNR) map of the factory environment

A. Factory Environment and Propagation Model

We consider a radio network in a factory having M micro cells with an inter-site distance of 50 m and N mobile users. For simplicity, the walls have no reflection effects on signaling, but emulate the internal factory structure, an environment with obstacles and heavy clutter, with path-loss given by Eq. 2:

$$PL = PL_0 + 10n log_{10}(\frac{d}{d_0}) + S, \tag{2}$$

where, $PL_0 = 80.84$ dBm is the free space path loss at reference distance $d_0 = 15$ m. The path loss exponent $n = 1.69$ and the shadowing S is a Gaussian-distributed random variable with zero mean and standard deviation, $\sigma = 6.62$ dB. The shadowing correlation and correlation distance are respectively set to 0.5 and 5 m.

The operating frequency is set to 2.4 GHz and the cells' transmission power set to 30 dBm for the whole transmission bandwidth of 100 MHz. Worth noting from the factory's SNR distribution in Fig. 3 is that the scenario is not limited by noise since we target an interference limited system. For a challenging and information-rich input to the neural network model, instead of the absolute RSRP values, we consider the worst-case of the user's signal level which is the fast-fading-

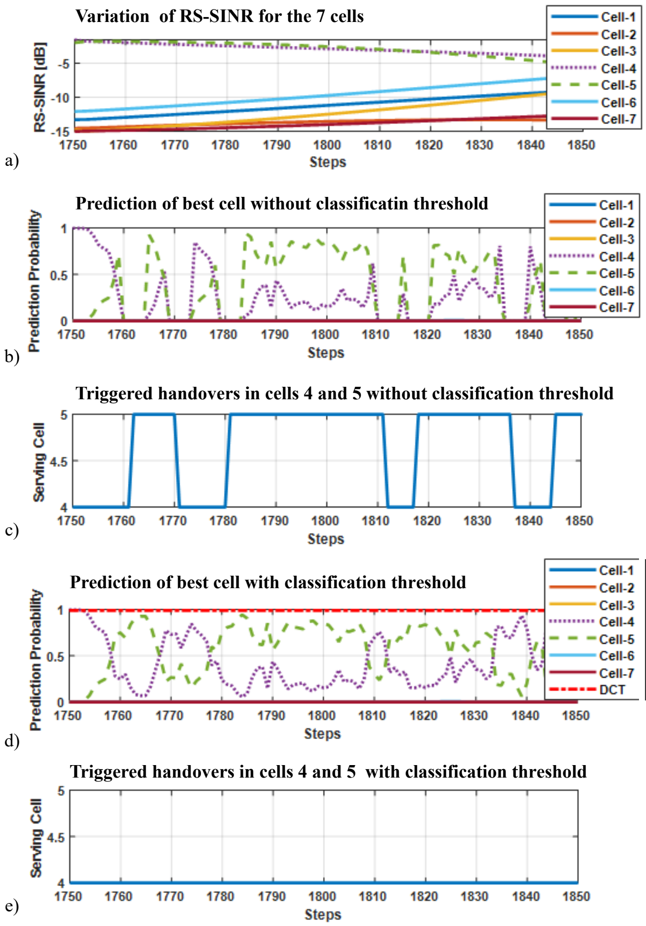

Fig. 4. prediction performance for a user along the edge of cells 4 and 5 without DCT

filtered RSSINR. Full buffer traffic model is used, i.e., users always have data to transmit as long they are in connected state. The users move using a random walk.

Whenever a HO is triggered, it is followed by a prediction break period of 5 time steps (equivalent to 50 ms) within which the HO is assumed to be executed and the model is allowed time to prepare for the new cell. Correspondingly, no new predictions are undertaken in this timeframe.

V. Results and Analysis

Using the above system model, we evaluated the performance of our classification based predictive HO with and without the classification threshold against the baseline with A3 event HO triggering and a finely tuned and converged MRO from [7].

Fig. 4 shows the performance of predictive HO with and without the classification threshold in response to changes in the signal of the serving cell. Initially cell 4 is the serving cell but, as seen in Fig. 4a, both cells 4 and 5 have almost equal SINR values indicating that the UE may be moving along the cell boarder. In Fig. 4b, we see that the predictive HO solution becomes indecisive about the optimal cell and triggers multiple HOs between the two cells. This is evident in Fig. 4c. Applying the classification threshold (DCT) with a higher

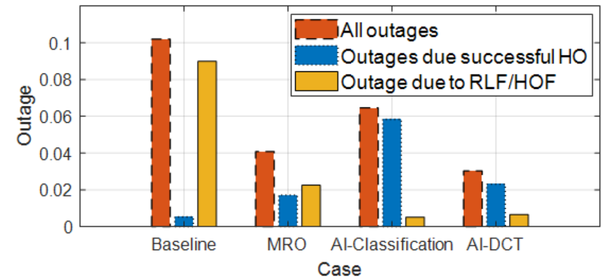

(a) Outages

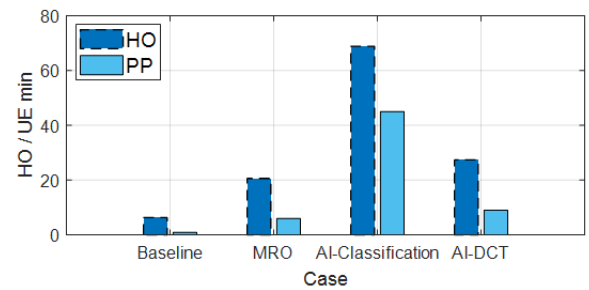

(b) HOs

Fig. 5. Comparisons between baseline, MRO and AI-DCT

threshold value, as shown in Fig. 4d, guides the model to only trigger HO when another cell is predicted to be better than the serving cell with a significant confidence. This reduces the number of unnecessary HOs, as shown in Fig. 4, where all the ping-pongs between cells 4 and 5 are eliminated.

Fig. 5 presents the complete comparison between the ML and non-ML-based solutions. Based on the system outage (derived as the product of count for the specific HO event and the respective T_{MIT}), we observe that the predictive HO with DCT outperforms all other techniques including the finely tuned and converged MRO. The big reduction in outage is a result of the significant reduction in the number of RLFs (leading to reduction in T_{HOF}) from 8 RLF/min in the baseline and 2 RLF/min for the MRO algorithm to less than 0.5 RLF/min when using the ML-based predicted HO. For the default predictive MRO algorithm, the reduction in RLF-triggered outage comes with increased outage due to unnecessary HOs, which the DCT resolves by reducing the per-minute number of HOs triggered.

VI. Conclusion

In this paper, we have investigated a predictive HO method, which utilizes machine learning to learn the best mobility behavior. The solution has been evaluated in a complex industrial environment, where it shows the improved performance in reducing the number of radio link failures and total outage from UE mobility compared to traditional HO triggering as well as the finely-tuned MRO. The work will be extended on two fronts. First, we intend to evaluate the performance of the algorithms in other, non-industrial. Moreover, we will also evaluate alternative implementations, e.g. using convolutional neural networks which may fingerprint the signals much in the same way as fingerprinting images.

References

[1] H. Park, Y. Lee, T. Kim, B. Kim, and J. Lee, "Handover mechanism in nr for ultra-reliable low-latency communications," *IEEE Network*, vol. 32, no. 2, pp. 41–47, 2018.

[2] G. T. 38.913, "Study on scenarios and requirements for next generation access technologies," *Tech. Rep. v14.3.0. Release 14*, 2017.

[3] G. T. 36.881, "Evolved universal terrestrial radio access (e-utra); study on latency reduction techniques for lte (release 14)," 2016.

[4] G. T. 36.300, "Evolved universal terrestrial radio access (e-utra) and evolved universal terrestrial radio access network (e-utran); overall description; stage 2 (release 14)," 2017.

[5] M. E. Morocho Cayamcela and W. Lim, "Artificial intelligence in 5g technology: A survey," in *2018 International Conference on Information and Communication Technology Convergence (ICTC)*, 2018, pp. 860–865.

[6] M. E. Morocho-Cayamcela, H. Lee, and W. Lim, "Machine learning for 5g/b5g mobile and wireless communications: Potential, limitations, and future directions," *IEEE Access*, vol. 7, pp. 137 184–137 206, 2019.

[7] I. Viering, B. Wegmann, A. Lobinger, A. Awada, and H. Martikainen, "Mobility robustness optimization beyond doppler effect and wss assumption," in *2011 8th International Symposium on Wireless Communication Systems*, 2011.

[8] Y. Koda, K. Yamamoto, T. Nishio, and M. Morikura, "Reinforcement learning based predictive handover for pedestrian-aware mmwave networks," in *IEEE INFOCOM 2018 - IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*, 2018, pp. 692–697.

[9] A. Nandy and M. Biswas, "Reinforcement learning basics," in *Reinforcement Learning*. Springer, 2018.

[10] M. Cord and P. Cunningham, "Supervised learning," in *Machine Learning Techniques for Multimedia: Case Studies on Organization and Retrieval (Cognitive Technologies)*. Springer, 2008.

[11] C. Lee, H. Cho, S. Song, and J. Chung, "Prediction-based conditional handover for 5g mm-wave networks: A deep-learning approach," *IEEE Vehicular Technology Magazine*, vol. 15, no. 1, pp. 54–62, 2020.

[12] L. Yan, H. Ding, L. Zhang, J. Liu, X. Fang, Y. Fang, M. Xiao, and X. Huang, "Machine learning-based handovers for sub-6 ghz and mmwave integrated vehicular networks," *IEEE Transactions on Wireless Communications*, vol. 18, no. 10, pp. 4873–4885, 2019.

[13] C. Wang, L. Ma, R. Li, T. S. Durrani, and H. Zhang, "Exploring trajectory prediction through machine learning methods," *IEEE Access*, vol. 7, pp. 101 441–101 452, 2019.