MILP-Based Optimal Handover Decisions: A Benchmark for Mobility Management Algorithms

Johannes Voigt, *Graduate Student Member, IEEE*, and Peter M. Rost, *Senior Member, IEEE*

***Abstract*—With the growing use of autonomous vehicles and autonomous aerial vehicles (AAVs), the requirements for reliable and seamless mobility in mobile networks are increasing. Although many handover optimization methods have been proposed, fair comparison is difficult because they rely on differing assumptions and there is no common benchmark. Therefore, it remains unclear how much optimization potential remains. We present a genie-aided benchmark based on mixed-integer linear programming (MILP) that uses precomputed signal quality to determine optimal handover decisions for given user trajectories in arbitrary network layouts. The framework captures key mobility mechanisms, including handover interruptions and radio link failures, and quantifies how far practical algorithms operate from the achievable performance limit.**

***Index Terms*—Handover, mixed-integer linear programming, mobility management, mobile network optimization.**

I. Introduction

CELLULAR networks were designed to provide mobile users with connectivity. Although the introduction of 5G NR has led to an increase in use cases with static users, the focus of future networks remains on mobile users. Autonomous vehicles with strict reliability requirements and autonomous aerial vehicles (AAVs), which require seamless vertical handover (HO), create new challenges for mobility management. At the same time, heterogeneous network (HetNet) structures with small cells, some of which are turned on only temporarily, make network parameterization more challenging and lead to more frequent HOs of mobile users.

Several problems are directly associated with HO. First, regular HOs lead to short connection interruptions, called handover interruption times (HITs), which are required to establish a connection to the target cell. Second, radio link failures (RLFs) can lead to handover failures. Reduction of HIT, improved reliability, and robust mobility support for user equipment (UE) speeds up to 500 km/h are key objectives for next-generation mobile networks [1].

Many concepts have been proposed to improve the HO performance, including radio resource control (RRC) parameter optimization and learning-based methods [2]. However, these approaches have in common that they lack a theoretical upper bound for the HO metrics considered. The remaining optimization potential is therefore usually unknown and prevents a fair and quantitative comparison of mobility management methods.

Received 11 March 2026; revised 4 May 2026; accepted 4 June 2026. Date of publication 8 June 2026; date of current version 16 June 2026. Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – 556070992. The associate editor coordinating the review of this letter and approving it for publication was K. Khawam. *(Corresponding author: Johannes Voigt.)*

The authors are with the Communications Engineering Laboratory (CEL), Karlsruhe Institute of Technology (KIT), 76187 Karlsruhe, Germany (e-mail: johannes.voigt@kit.edu).

Digital Object Identifier 10.1109/LCOMM.2026.3701321

A. Related Work

Recent work has considered predictive and learning-based optimization of HO decisions. In [3], predictive HO decisions are made using online convex optimization, while [4] applies deep reinforcement learning to learn HO policies in dynamic environments. However, these approaches do not provide a global optimality benchmark, and the achievable performance gap remains unclear.

To address this limitation, in [5], the HO process in heterogeneous networks is analyzed probabilistically, and an upper bound for the achievable capacity is derived for a specific scenario using non-causal knowledge. This enables a more structured comparison of HO algorithms, but remains limited to the setting considered.

In [6], a genie-aided approach is proposed that represents the HO problem as a mixed-integer linear program (MILP), where optimal HOs maximize the average signal-to-noise ratio. By introducing a penalty for late-delivered packets, the authors implicitly consider HIT. However, RRC counters and timers, explicit modeling of RLFs and constraints on the frequency of HOs due to preparation delays are not considered. In contrast, the proposed formulation directly integrates these mechanisms while providing a global optimality benchmark.

B. Contributions

In this work, we present a method to compute a trajectory-specific upper bound on the performance of mobility management algorithms. It serves as a benchmark for the evaluation of HO strategies, such as event-based heuristics and machine learning-based approaches, and is not intended for real-time deployment. Additionally, the resulting optimal solutions can be used as labeled data for supervised learning.

The proposed optimization problem is formulated as a MILP and enables HO optimization for arbitrary network layouts and user trajectories. In the formulation, we explicitly model radio link monitoring (RLM)/RLF counters and timers, HO preparation and execution delays, as well as RLF interruptions. The objective function characterizes the Pareto frontier between average throughput and link reliability. As in prior work, we use a genie-aided approach in which signal conditions over the entire time horizon are known a priori. The resulting solution is therefore conditioned on precomputed signal quality and does not model the impact of HOs on interference or cell load.

To facilitate the use of the method in further research and to allow an easy comparison of HO approaches, the MILP implementation[1] and the dataset/results[2] are publicly available.

[1]Source code available at: https://github.com/kit-cel/handover-optim-milp

[2]Dataset/results at IEEE DataPort: https://dx.doi.org/10.21227/9zk6-vb30

II. Preliminaries

A. Mobility Management in 3 GPP Cellular Networks

In modern cellular networks such as 5G NR, mobility is implemented using an event-based protocol. If predefined conditions are met, e. g., the A3 event [7] is triggered and measurement reports are sent to the serving cell. These conditions are based on parameters such as offsets, hysteresis, and trigger delays (time-to-trigger (TTT)), and are evaluated using time-averaged measurements of the reference signal received power (RSRP). Once the serving cell has decided to perform a HO, the target cell prepares the resources required for the UE in the HO preparation phase. After the HO command is sent to the UE and successfully received, the UE connects to the target cell via the random-access channel and experiences a short service interruption of approximately 40–80 ms [8], [9].

During operation in the RRC Connected state, the signal conditions are continuously monitored by RLM, based on RRC parameters configured by the serving cell. Counters such as N310 and N311 count *out-of-sync* and *in-sync* events, respectively, indicating physical-layer problems caused by deteriorating radio conditions, i. e., when the signal-to-interference-plus-noise ratio (SINR) falls below or rises above predefined quality of service (QoS) thresholds (Q_{out}/Q_{in}). If the out-of-sync condition persists for a time T_{310}, an RLF is declared, and the UE initiates a time-consuming connection re-establishment procedure, resulting in a long service interruption on the order of 10^3 ms [10]. Such events may also occur when HOs are triggered too early or too late.

The modeling of radio link monitoring and RLF within the optimization problem, as well as in the reference simulations, is based on the parameters and definitions specified in [8].

B. Mixed-Integer Linear Programming

The proposed optimization problem is formulated as a MILP with binary and integer decision variables, linear constraints, and a linear objective function. Binary operations, indicator functions, and conditional constraints are represented using standard linear encodings and big-M formulations [11]. Bilinear constraints involving binary and bounded integer variables are linearized using McCormick envelope formulations [12]. Limiting functions, such as minimum or maximum operators, are reformulated using auxiliary variables and linear constraints. The resulting problem is a pure MILP that can be solved to global optimality using a branch-and-cut solver, e. g., the Gurobi optimizer [13].

III. Problem Formulation

Bold uppercase letters denote matrices, bold lowercase letters denote vectors, and lowercase letters denote scalars.

A. System Model

We consider a single UE moving through an arbitrary cellular network consisting of $N_{\text{cell}} \in \mathbb{N}$ cells, where $\mathcal{B} = \{1, \dots, N_{\text{cell}}\}$ denotes the set of candidate target cells. Handover decisions are made per UE and can therefore be analyzed on a user-specific level. Accordingly, we focus on single-UE trajectories and assume precomputed signal quality. Extensions to multi-UE scenarios would require modeling the coupling between users via interference and cell load, which is beyond the scope of this benchmark formulation.

Corresponding to the sampling interval in [14], time is discretized into slots of duration $t_{\text{res}} = 40$ ms. Let $t \in \mathcal{T} = \{0, \dots, T-1\}$ denote the discrete time index, where $T \in \mathbb{N}$ is the time horizon. The total time interval is therefore $T \cdot t_{\text{res}}$.

For each UE, wideband RSRP and SINR are simulated for all N_{cell} cells and time steps $t \in \mathcal{T}$. For optimization, we assume that the measured values over the entire horizon T are known a priori. The simulated SINR is denoted by $\mathbf{\Gamma} \in \mathbb{R}^{N_{\text{cell}} \times T}$, where $\boldsymbol{\gamma}_t = (\gamma_{1,t}, \dots, \gamma_{N_{\text{cell}},t})^\mathsf{T}$ contains the SINR of all candidate target cells at the time step t.

The counters and timers described in Section II-A are incorporated into the optimization as constants. We define $N_{310}, N_{311} \in \mathbb{N}_0$ and $T_{310} = T_{310,\text{ms}}/t_{\text{res}} \in \mathbb{N}_0$, expressed in discrete time steps. The QoS thresholds are set to $Q_{\text{in}} = -6$ dB and $Q_{\text{out}} = -8$ dB according to [8].

B. Mobility Model and Timing Assumptions

The UE movement and radio conditions are assumed to be unaffected by HO decisions. Hence, the UE trajectory is unaffected by HOs or RLFs, allowing the precomputation of $\mathbf{\Gamma}$. All timing-related mechanisms are represented in discrete time and aligned with t_{res}, such that there is no stochastic uncertainty within the optimization. Consequently, the optimization is performed offline over a finite time horizon and is genie-aided, assuming full knowledge of future radio conditions.

For the HO preparation time $T_{\text{prep}} \in \mathbb{N}_0$, the execution time $T_{\text{exec}} \in \mathbb{N}_0$, and the RLF recovery time $T_{\text{RLF}} \in \mathbb{N}_0$, constant values are assumed to be multiples of t_{res}. The total duration of a HO is $T_{\text{HO}} = T_{\text{prep}} + T_{\text{exec}}$. The choice of these constants directly influences the optimization result:

- Setting $T_{\text{prep}} = 0$ corresponds to negligible HO preparation latency, e. g., due to pre-preparation of HO commands as in conditional handover (CHO).
- Setting $T_{\text{exec}} = 0$ models an idealized seamless HO and abstracts mechanisms such as make-before-break or HO with dual active protocol stack (DAPS). In this case, apart from unavoidable RLFs or QoS interruptions, the optimal cell assignment corresponds to selecting the best cell with respect to $\boldsymbol{\gamma}_t$ at each time step.

The model captures the timing effects of HOs relevant for connectivity and throughput, but does not explicitly model protocol signaling or beam-level procedures. To reduce the complexity of the MILP, the best received beam per UE and cell is assumed.

C. Optimization Problem and Objectives

The assignment of a UE to one of the candidate serving cells in $\mathcal{B}$ at each time step $t \in \mathcal{T}$ is described by the indicator matrix $\boldsymbol{X} \in \mathcal{X} = \{0,1\}^{N_{\text{cell}} \times T}$, similar to [6]. The column vector $\boldsymbol{x}_t = (x_{1,t}, \dots, x_{N_{\text{cell}},t})^\mathsf{T}$ contains the decision variables at time t. Note that by definition, $\boldsymbol{X}$ is constrained to be binary. If $x_{i,t} = 1$, the UE is connected to cell i. To ensure that a

UE is connected to exactly one cell at time t, the assignment vector $\boldsymbol{x}_t$ is constrained to have exactly one entry equal to 1, i. e., $\sum_{n=1}^{N_{\text{cell}}} x_{n,t} = 1$.

To indicate the time steps during which service is temporarily interrupted, we introduce binary indicators for ongoing HO execution $i^{\text{HO,run}} \in \{0,1\}^T$ and RLF recovery $i^{\text{RLF,run}} \in \{0,1\}^T$. Optimization of $\boldsymbol{X}$ is subject to constraints that enforce valid HO timing, correct RLM/RLF timing behavior (N310, N311, T310), RLF detection, and the corresponding interruption indicators. The complete formulation of the auxiliary variables and constraints is provided in the Appendix.

For users requiring maximum throughput, e. g., enhanced mobile broadband (eMBB), we define the achievable rate averaged over time (average delivered rate)

$$\begin{aligned}\mathcal{L}_{\text{r}}(\boldsymbol{X}) &= \frac{1}{T}\sum_{t=0}^{T-1} \log_2\left(1+\boldsymbol{\gamma}_t^{\mathsf{T}}\boldsymbol{x}_t\right)\cdot\left(1-i_t^{\text{HO,run}}-i_t^{\text{RLF,run}}\right)\\ &= \frac{1}{T}\sum_{t=0}^{T-1} r_t\left(1-i_t^{\text{out}}\right), \end{aligned} \tag{1}$$

where the one-hot vector $\boldsymbol{x}_t$ selects the serving-cell SINR, and $r_t = \log_2\left(1+\boldsymbol{\gamma}_t^{\mathsf{T}}\boldsymbol{x}_t\right)$ denotes the corresponding instantaneous achievable rate. In addition, $i_t^{\text{out}} = i_t^{\text{HO,run}} + i_t^{\text{RLF,run}}$ indicates if the UE is out of service due to HO or RLF. Note that $i_t^{\text{HO,run}} + i_t^{\text{RLF,run}} \le 1$ holds and $i_t^{\text{out}} \in \{0,1\}$ is guaranteed. Thus, no rate contribution is counted during link interruptions.

For users with high reliability requirements, e. g., autonomous vehicles and AAVs, we maximize the average connected time

$$\begin{aligned}\mathcal{L}_{\text{c}}(\boldsymbol{X}) &= \frac{1}{T}\sum_{t=0}^{T-1}\left(1-i_t^{\text{HO,run}}-i_t^{\text{RLF,run}}\right)\\ &= 1-\frac{1}{T}\sum_{t=0}^{T-1} i_t^{\text{out}}, \end{aligned} \tag{2}$$

by minimizing the number of HOs and therefore the total HITs. Note that $\boldsymbol{x}_t$ does not appear directly in the objective, but is used indirectly in the constraints that define the states of the indicator variables $i_t^{\text{HO,run}}$ and $i_t^{\text{RLF,run}}$.

A practical system typically requires a trade-off between throughput (1) and reliability (2). Thus, we maximize the time-averaged achievable rate, subject to the QoS requirement of a maximum tolerable outage rate $\varepsilon \in [0,1]$:

$$\bar{r}^* = \max_{\boldsymbol{X}\in\mathcal{X}} \mathcal{L}_{\text{r}}(\boldsymbol{X}), \quad \text{s.t.}\ \frac{1}{T}\sum_{t=0}^{T-1} i_t^{\text{out}} < \varepsilon. \tag{3}$$

To obtain a MILP formulation, we introduce the Lagrange multiplier $\lambda \ge 0$ to relax the hard outage constraint. The additive constant term $\lambda\varepsilon$ does not affect the optimum and can therefore be omitted:

$$\begin{aligned}\max_{\boldsymbol{X}\in\mathcal{X}} \mathcal{L}(\boldsymbol{X},\lambda) &= \max_{\boldsymbol{X}\in\mathcal{X}} \frac{1}{T}\sum_{t=0}^{T-1} r_t\left(1-i_t^{\text{out}}\right)\\ &\quad -\lambda\left(\frac{1}{T}\sum_{t=0}^{T-1} i_t^{\text{out}}-\varepsilon\right)\\ &= \max_{\boldsymbol{X}\in\mathcal{X}} \mathcal{L}_{\text{r}} - \lambda\left(1-\mathcal{L}_{\text{c}}\right). \end{aligned} \tag{4}$$

The design parameter λ (in bit/s/Hz) specifies how much rate is sacrificed to avoid one time step of outage and therefore corresponds directly to the Pareto frontier between rate and reliability under the modeled system assumptions.

D. Problem Classification and MILP Reformulation

Due to the bilinear coupling between $\boldsymbol{x}_t$ and i_t^{out}, the original formulation of (4) is quadratic and non-convex even under continuous relaxation. Following Sec. II-B, all bilinear terms are linearized using auxiliary variables. Since $\boldsymbol{\gamma}_t$ is constant and $\boldsymbol{x}_t$ is a one-hot vector, we precompute $\rho_{n,t} = \log_2(1+\gamma_{n,t})$ and write $r_t = \boldsymbol{\rho}_t^{\mathsf{T}}\boldsymbol{x}_t$, which is linear in $\boldsymbol{x}_t$. Hence, the reformulated problem is a pure MILP. The resulting MILP is solved using Gurobi [13] with its branch-and-cut framework, which certifies global optimality upon convergence. Empirically, the problem formulation remains tractable for moderate values of N_{cell} and T.

IV. Results

For evaluation, we simulate a 3GPP-based setup with a hexagonal wrap-around layout of 7 base stations (BSs), each with 3 cells, using the configuration in [9]. To obtain statistically robust results, we perform Monte Carlo simulations with $M = 600$ UE trajectories of length 60 s, i. e., 10 h in total. Each UE is randomly dropped in the simulated area and moves in a random direction at speed $v = 60$ km/h.

As a reference, we use the A3 event-driven mobility procedure defined in Sec. II-A and evaluate $K = 2480$ combinations of the parameters TTT, offset, and hysteresis. The A3 baseline uses time-averaged RSRP for event triggering, including TTT. Thus, TTT is part of the A3 reference simulation, whereas the MILP benchmark optimizes cell assignments directly subject to (4). For each configuration k, simulation results are aggregated over time and the M simulated UEs, yielding:

$$L_{\text{r},k} = \frac{1}{M}\sum_{m=1}^{M}\frac{1}{T}\sum_{t=0}^{T-1} r_{m,t}\left(1-i_{m,t}^{\text{out}}\right), \tag{5}$$

$$L_{\text{c},k} = 1-\frac{1}{M}\sum_{m=1}^{M}\frac{1}{T}\sum_{t=0}^{T-1} i_{m,t}^{\text{out}}. \tag{6}$$

For a given Lagrange multiplier λ, we assign to each A3 configuration the score

$$s_k(\lambda) = L_{\text{r},k} - \lambda\left(1-L_{\text{c},k}\right), \tag{7}$$

which uses the same rate–connectivity trade-off as in objective (4). Let $Q_s(q,\lambda)$ denote the empirical q-th percentile of the scores $\{s_k(\lambda)\}_{k=1}^K$. The corresponding upper-tail set is $\mathcal{S}_q(\lambda) = \{k \mid s_k(\lambda) \ge Q_s(q,\lambda)\}$. Instead of selecting a single configuration, we summarize the best-performing configurations by averaging across the selected parameter sets:

$$(\bar{L}_{\text{c},q}(\lambda), \bar{L}_{\text{r},q}(\lambda)) = \frac{1}{|\mathcal{S}_q(\lambda)|}\sum_{k\in\mathcal{S}_q(\lambda)} (L_{\text{c},k}, L_{\text{r},k}), \tag{8}$$

representing the average performance of all configurations whose score is at least the empirical q-th percentile in (7).

Using the proposed MILP, we optimize (4) for different λ and obtain the optimal cell assignments for each UE, while the reference is evaluated using (8). Figure 1 shows the rate–outage trade-off. For the objective in (4), the MILP traces the Pareto frontier and thus defines an upper bound on HO performance, whereas the red dotted curve shows the Pareto front over all A3 parameter configurations. The gap indicates the remaining optimization potential that cannot be closed by

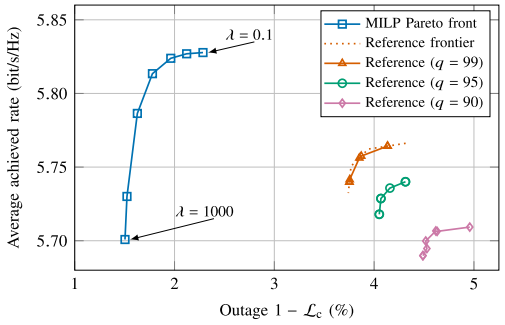

Fig. 1. Rate–outage frontiers. The blue curve shows the MILP solution for varying λ. The red dotted curve represents the Pareto front of all A3 parameter settings. The marked curves show the average performance of configurations whose scores exceed the 90th, 95th, and 99th percentiles according to (7).

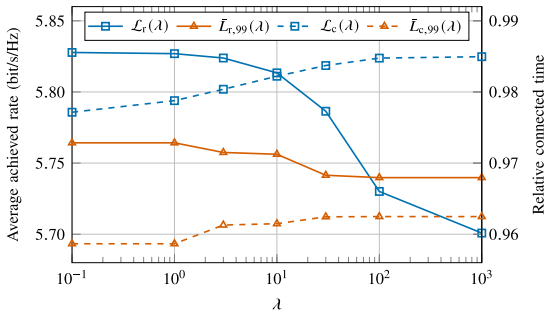

Fig. 2. Trade-off between average achieved rate (left axis, solid lines) and relative connected time (right axis, dashed lines) versus Lagrange multiplier λ. The reference curve shows the average performance of parameter configurations whose scores lie above the 99th percentile according to (7).

tuning the A3 parameters alone. Since the benchmark is agnostic to the underlying decision mechanism, it can also evaluate advanced HO protocols, including learning-based methods. In particular, a given policy can be applied to the same signal trajectories, and its performance can be compared to the Pareto frontier, thereby quantifying its gap under identical conditions.

In addition, the distance between the average performance of $\mathcal{S}_q(\lambda)$ for $q \in \{90, 95, 99\}$ and the reference front indicates the sensitivity to parameter settings. Suboptimal configurations increase outage mainly due to higher RLF rates measured in events/UE/s, rising from about 2.6×10^{-3} (MILP) to 1.3×10^{-2} (reference front) and up to 2.1×10^{-2} for $q = 90$, i. e., a factor of about 5–8. The HO rate also increases from about 0.20 to 0.32, thereby increasing HIT and reducing the average rate.

Figure 2 shows the rate–connectivity trade-off. Larger values of λ penalize outage more strongly, trading rate for connectivity. Compared to the reference, the MILP reduces outage by approximately 2 percentage points across all λ through fewer and better-timed HOs. As λ increases, the MILP reduces HIT from about 2.0% to 1.6% and the ping-pong (PP) rate drops to nearly zero, whereas A3 reduces HIT only from about 3.1% to 2.6% at the cost of a higher RLF rate. This explains the remaining performance gap and the limited ability of A3 to realize UE-specific HO behavior.

Due to the exponential worst-case complexity in N_{cell} and T, runtime is an important consideration. Across all runs, the median runtime is 96 s, while some instances require up to 36 h, indicating a strongly right-skewed distribution. This is primarily caused by trajectory-dependent combinatorial complexity, where multiple serving cell candidates remain ambiguous over long periods of time. As an offline benchmark, these runtimes are acceptable and can be further reduced by pruning candidate cells, shorter horizons (rolling-horizon), and warm-starting from heuristic solutions. All optimizations were performed on a compute node equipped with two AMD EPYC 9555 CPUs and 256 GB RAM. No solver time limit was set.

V. Conclusion

This letter introduced a genie-aided MILP-based benchmark for HO optimization that explicitly models HIT, RLF recovery, and application-dependent objectives. With full knowledge of the radio conditions, the method provides an upper bound for the HO performance and achieves the Pareto frontier between the mean rate and connected time.

In the setup considered, a consistent performance gap between the benchmark and A3 baselines highlights the remaining optimization potential and the need for novel HO protocols to close this gap. Despite an exponential worst-case complexity, the median runtime remains tractable for moderate problem sizes. In general, the proposed benchmark enables a quantitative comparison of mobility management algorithms and can serve as a reference for algorithms approaching the achievable rate–reliability trade-off. Possible extensions include multi-user, load-aware scenarios, where HO decisions are coupled through shared network resources.

Appendix
MILP Formulation Details

A. Auxiliary Variable Definitions

- Precomputed in-sync and out-of-sync matrices:
$$\begin{aligned}\boldsymbol{Q}^{\text{out}} &= \left(\boldsymbol{q}_0^{\text{out}}\, \boldsymbol{q}_1^{\text{out}}\, \cdots\, \boldsymbol{q}_{T-1}^{\text{out}}\right)\\ &= 1\{\boldsymbol{\Gamma} \le Q_{\text{out}}\} \in \{0,1\}^{N_{\text{cell}}\times T}\\ \boldsymbol{Q}^{\text{in}} &= \left(\boldsymbol{q}_0^{\text{in}}\, \boldsymbol{q}_1^{\text{in}}\, \cdots\, \boldsymbol{q}_{T-1}^{\text{in}}\right)\\ &= 1\{\boldsymbol{\Gamma} \ge Q_{\text{in}}\} \in \{0,1\}^{N_{\text{cell}}\times T}.\end{aligned}$$
- Sync status of the serving cell: $\boldsymbol{i}^{\text{in}}, \boldsymbol{i}^{\text{out}} \in \{0,1\}^T$
- N310/N311 counters and flags for all $t \in \mathcal{T}$:
$$\begin{aligned} i_t^{310} &\in \{0,1\}, & i_t^{311} &\in \{0,1\},\\ u_t^{310} &\in \{0,\ldots,N_{310}+1\}, & u_t^{311} &\in \{0,\ldots,N_{311}+1\},\\ u_t^{310,\text{sat}} &\in \{0,\ldots,N_{310}\}, & u_t^{311,\text{sat}} &\in \{0,\ldots,N_{311}\},\\ c_t^{310} &\in \{0,\ldots,N_{310}\}, & c_t^{311} &\in \{0,\ldots,N_{311}\}.\end{aligned}$$
- T310-related variables for all $t \in \mathcal{T}$: remaining time variable $\tau_t \in \{0,\ldots,T_{310}\}$, and T310 state indicators:
$$i_t^{\text{T310,start}}, i_t^{\text{T310,cancel}}, i_t^{\text{T310,exp}}, i_t^{\text{T310,stop}}, i_t^{\text{T310,run}} \in \{0,1\}.$$
- RLF occurrence/recovery: $\boldsymbol{i}^{\text{RLF,start}}, \boldsymbol{i}^{\text{RLF,run}} \in \{0,1\}^T$.
- Variable $\boldsymbol{i}^{\text{chg}} \in \{0,1\}^T$, where $i_t^{\text{chg}} = 1$ indicates a change in the serving cell from step $t-1$ to step t ($\boldsymbol{x}_{t-1} \neq \boldsymbol{x}_t$).
- The variables $\boldsymbol{i}^{\text{HO,run}}, \boldsymbol{i}^{\text{HO,end}} \in \{0,1\}^T$ indicate an ongoing HO execution and the time at which a HO is complete.

B. Constraints

1) Q-in & Q-Out Indicators for the Selected Serving Cell:

$$i_t^{\text{in}} = \boldsymbol{x}_t^{\mathsf{T}} \boldsymbol{q}_t^{\text{in}} \in \{0,1\}, \quad i_t^{\text{out}} = \boldsymbol{x}_t^{\mathsf{T}} \boldsymbol{q}_t^{\text{out}} \in \{0,1\} \quad \forall t. \tag{9}$$

2) RLF Start and Running Logic: For $t = 0$, the indicators are initialized to $i_0^{\text{RLF,start}} = 0$ and $i_0^{\text{RLF,run}} = 0$, and for $t \geq 1$:

$$i_t^{\text{RLF,start}} = 1\{\tau_{t-1} = 1\} \cdot (1 - i_{t-1}^{\text{RLF,run}}), \tag{10a}$$

$$i_t^{\text{RLF,run}} \geq 1\left\{ \sum_{\kappa=0}^{\min\{t,\, T_{\text{RLF}}-1\}} i_{t-\kappa}^{\text{RLF,start}} \geq 1 \right\}, \tag{10b}$$

$$i_t^{\text{RLF,run}} \leq i_{t-1}^{\text{RLF,run}} + 1\left\{ \sum_{\kappa=0}^{\min\{t,\, T_{\text{RLF}}-1\}} i_{t-\kappa}^{\text{RLF,start}} \geq 1 \right\}. \tag{10c}$$

3) N310 Logic: The following constraints apply to all t, except for initialization $u_0^{310} = i_0^{\text{out}}$:

$$b_t^{310,\text{reset}} = (1 - i_t^{\text{RLF,run}}) \cdot (1 - i_t^{\text{HO,run}}), \tag{11a}$$

$$u_t^{310} = b_t^{310,\text{reset}} \cdot \left(c_{t-1}^{310} + i_t^{\text{out}} \cdot (1 - i_{t-1}^{\text{T310,run}}) \right), \tag{11b}$$

$$u_t^{310,\text{sat}} = \min\left(u_t^{310},\, N_{310}\right), \tag{11c}$$

$$c_t^{310} = b_t^{310,\text{reset}} \cdot (1 - i_t^{\text{T310,run}}) \cdot u_t^{310,\text{sat}}, \tag{11d}$$

$$i_t^{310} = 1\{u_t^{310} \geq N_{310}\}. \tag{11e}$$

4) N311 Logic: The following constraints apply to all t, except for initializations $u_0^{311} = 0$ and $c_0^{311} = 0$:

$$b_t^{311,\text{reset}} = (1 - i_t^{\text{RLF,run}}) \cdot (1 - i_t^{\text{HO,run}}), \tag{12a}$$

$$u_t^{311} = b_t^{311,\text{reset}} \cdot \left(c_{t-1}^{311} + i_t^{\text{in}} \cdot i_{t-1}^{\text{T310,run}} \right), \tag{12b}$$

$$u_t^{311,\text{sat}} = \min\left(u_t^{311},\, N_{311}\right), \tag{12c}$$

$$c_t^{311} = b_t^{311,\text{reset}} \cdot (1 - i_t^{\text{T310,stop}}) \cdot i_t^{\text{T310,run}} \cdot u_t^{311,\text{sat}}, \tag{12d}$$

$$i_t^{311} = 1\{u_t^{311} \geq N_{311}\}. \tag{12e}$$

5) T310 Logic: For $t = 0$, $i_t^{\text{T310,start}}$, $i_t^{\text{T310,cancel}}$, $i_t^{\text{T310,exp}}$, and $i_t^{\text{T310,stop}}$ are initialized to 0; for $t \geq 1$ the following applies:

$$i_t^{\text{T310,start}} = i_t^{310} \cdot (1 - i_{t-1}^{\text{T310,run}}) \cdot (1 - i_t^{\text{RLF,run}}), \tag{13a}$$

$$i_t^{\text{T310,cancel}} = i_{t-1}^{\text{T310,run}} \cdot i_t^{311}, \tag{13b}$$

$$i_t^{\text{T310,exp}} = i_{t-1}^{\text{T310,run}} \cdot i_t^{\text{RLF,run}}, \tag{13c}$$

$$i_t^{\text{T310,stop}} = \max\left(i_t^{\text{T310,cancel}},\, i_t^{\text{T310,exp}} \right). \tag{13d}$$

With $\tau_0 = 0$, the remaining time is updated for $t \geq 1$ as:

$$\tau_t = (1 - i_t^{\text{RLF,run}})(1 - i_t^{\text{T310,stop}}) \Big[T_{310} \cdot i_t^{\text{T310,start}} + (\tau_{t-1} - i_{t-1}^{\text{T310,run}})(1 - i_t^{\text{T310,start}}) \Big], \tag{13e}$$

which decrements by 1 while active. T310 is active if $\tau_t \geq 1$:

$$i_t^{\text{T310,run}} = 1\{\tau_t \geq 1\} \forall t. \tag{13f}$$

6) Cell-Change Detection and Constraints:

$$i_0^{\text{chg}} = 0, \qquad i_t^{\text{chg}} = 1 - \boldsymbol{x}_{t-1}^{\mathsf{T}} \boldsymbol{x}_t \quad \forall t \geq 1, \tag{14a}$$

with at most one cell change in any possible HO window:

$$\sum_{\kappa=0}^{\min(T_{\text{HO}}, T-1-t)} i_{t+\kappa}^{\text{chg}} \leq 1 \quad \forall t. \tag{14b}$$

No cell change is possible while T310 is active, during RLF recovery, or when the target cell is out-of-sync (SINR $< Q_{\text{out}}$):

$$i_t^{\text{chg}} \leq 1 - \max\left(i_{t-1}^{\text{T310,run}},\, i_t^{\text{RLF,run}},\, i_t^{\text{out}} \right) \qquad \forall t \geq 1. \tag{14c}$$

Further, changes are allowed only directly after RLF recovery (re-establishment) or if conditions for a HO are fulfilled:

$$i_t^{\text{chg}} \leq \max\left(i_{t-1}^{\text{RLF,run}}, 1\left\{ \sum_{\kappa=0}^{\min\{t,\, T_{\text{HO}}\}} i_{t-\kappa}^{\text{RLF,run}} = 0 \right\} \right) \quad \forall t \geq 1. \tag{14d}$$

7) Handover Detection & Running Indicator: HO execution ends at t iff the serving cell changes at t, and no RLF occurred:

$$i_t^{\text{HO,end}} = i_t^{\text{chg}} \cdot 1\left\{ \sum_{\kappa=0}^{\min\{t,\, T_{\text{HO}}\}} i_{t-\kappa}^{\text{RLF,run}} = 0 \right\} \forall t. \tag{15a}$$

The HO execution is ongoing at t if an HO execution end flag is in the next T_{exec} steps:

$$i_t^{\text{HO,run}} = 1\left\{ \sum_{\kappa=1}^{\min\{T_{\text{exec}},\, T-1-t\}} i_{t+\kappa}^{\text{HO,end}} \geq 1 \right\} \forall t. \tag{15b}$$

REFERENCES

[1] *Study on Scenarios and Requirements for Next Generation Access Technologies*, document TR 38.913, 3GPP, May 2024. [Online]. Available: http://www.3gpp.org/DynaReport/38913.htm

[2] S. Alraih, R. Nordin, A. Abu-Samah, I. Shayea, and N. F. Abdullah, "A survey on handover optimization in beyond 5G mobile networks: Challenges and solutions," *IEEE Access*, vol. 11, pp. 59317–59345, 2023.

[3] V. Rajabi, J. Fink, M. Kasparick, and S. Stańczak, "Predictive handover optimization," in *Proc. Int. Workshop Smart Antennas (WSA)*, 2024, pp. 91–96.

[4] Z. Wang, L. Li, Y. Xu, H. Tian, and S. Cui, "Handover control in wireless systems via asynchronous multiuser deep reinforcement learning," *IEEE Internet Things J.*, vol. 5, no. 6, pp. 4296–4307, Dec. 2018.

[5] I. Pappalardo, A. Zanella, and M. Zorzi, "Upper bound analysis of the handover performance in HetNets," *IEEE Commun. Lett.*, vol. 21, no. 2, pp. 418–421, Feb. 2017.

[6] D. da Silva Brilhante, J. F. de Rezende, and N. Marchetti, "Handover optimisation for high-capacity low-latency 5G NR mmWave communication," *Ad Hoc Netw.*, vol. 153, Feb. 2024, Art. no. 103328.

[7] *5G; NR; Radio Resource Control (RRC); Protocol Specification*, document TS 38.331, Aug. 2024. [Online]. Available: http://www.3gpp.org/DynaReport/38331.htm

[8] *Evolved Universal Terrestrial Radio Access (E-UTRA); Mobility Enhancements in Heterogeneous Networks*, document TR 36.839, Dec. 2012. [Online]. Available: http://www.3gpp.org/DynaReport/36839.htm

[9] A. Gündogan, A. Badalioglu, P. Spapis, and A. Awada, "On the modelling and performance analysis of lower layer mobility in 5G-advanced," in *Proc. IEEE Wireless Commun. Netw. Conf. (WCNC)*, Mar. 2023, pp. 1–6.

[10] H.-S. Park, Y. Lee, T.-J. Kim, B.-C. Kim, and J.-Y. Lee, "Faster recovery from radio link failure during handover," *IEEE Commun. Lett.*, vol. 24, no. 8, pp. 1835–1839, Aug. 2020.

[11] G. L. Nemhauser and L. A. Wolsey, *Integer and Combinatorial Optimization*. Hoboken, NJ, USA: Wiley, 1988.

[12] G. P. McCormick, "Computability of global solutions to factorable nonconvex programs: Part I—Convex underestimating problems," *Math. Program.*, vol. 10, no. 1, pp. 147–175, Dec. 1976.

[13] Gurobi Optimization, LLC. (2026). *Gurobi Optimizer Reference Manual*. [Online]. Available: https://www.gurobi.com

[14] *Study on Artificial Intelligence (AI)/Machine Learning (ML) for Mobility in NR*, document TR 38.744, Nov. 2025. [Online]. Available: http://www.3gpp.org/DynaReport/38744.htm