Adaptive Reduced System Modeling for Real-time Dynamic Simulations

Stefanos Eleftheriadis*, Mohammad Hashemnezhad*, Savvas Panagi*,
Thibaut Vermeulen†, Geethu Joseph‡, and Petros Aristidou*
* Dept. of Electrical Engineering, and Computer Science and Engineering, Cyprus University of Technology, Limassol, Cyprus
† Réseau de Transport d'Électricité (RTE), Paris, France
‡ Collaborative Research for Energy System Modelling (CRESYM), Brussels, Belgium
Emails: {sc.eleftheriadis, ma.hashemnezhad}@edu.cut.ac.cy, {savvas.panagi, petros.aristidou}@cut.ac.cy,
geethu.joseph@cresym.eu, thibaut.vermeulen@rte-france.com

***Abstract*—Real-time dynamic simulation requires that each simulated time interval complete within the equivalent wall-clock CPU time, a strict budget that large power system models, with hundreds of thousands of state variables, cannot meet. Existing model reduction methods either sacrifice internal network detail, are valid only near a fixed operating point, or cannot adapt to changing contingencies online. This paper proposes an Adaptive Model Selection (AMS) framework that automatically reduces a full-detail dynamic model to a smaller, contingency-specific model by identifying which substations require full fidelity and simplifying the rest. Given the pre-disturbance operating state and event location, a pair of Graph Attention Network (GAT)-based ordinal predictors classify each component's expected dynamic activity level; substations predicted to be inactive are simplified from node-breaker to bus-breaker representation. Applied to a realistic French transmission network model (over 6,000 buses), AMS reduces continuous and discrete model variables by approximately 60%, achieves up to 2.6x simulation speedup, and maintains trajectory errors below the numerical solver tolerance, demonstrating a practical route toward real-time operation of large-scale dynamic simulation.**

***Index Terms*—Adaptive model reduction, Critical area identification, Large-scale power systems, Machine learning**

I. Introduction

Real-time dynamic simulation requires that each simulated time interval complete within the equivalent wall-clock CPU time, a strict budget that large power system models, with hundreds of thousands of state variables, cannot meet. This constraint is critical for operator training and dynamic security assessment [1], where both computational speed and dynamic fidelity must be preserved.

The core challenge is to reduce model size while satisfying three requirements simultaneously: (i) *accuracy*, meaning trajectory errors remain below solver tolerance; (ii) *structure preservation*, meaning internal network topology and component-level models are retained for realistic operator interaction; and (iii) *adaptivity*, meaning the simplification adjusts to the specific operating condition and contingency.

A. Literature Review

Several studies have explored machine learning (ML)-based methods to predict power-system dynamic behavior without full-order simulations. Early work demonstrated that post-disturbance trajectories contain sufficient information for rapid stability assessment [2], and subsequent deep learning and graph-based architectures improved robustness by incorporating network topology into the learning process [3]. More recent approaches shifted toward pre-disturbance information, using steady-state operating conditions and physics-informed features to estimate stability directly from the pre-contingency state [4]. Despite these advances, ML-based methods focus on predicting stability outcomes, typically as binary classifications, and provide limited information on the level of model fidelity required at the component or substation level. Generalization to unseen operating conditions also remains a concern.

In parallel, model reduction has a long history in power-system dynamic equivalencing. Structure-preserving methods such as coherency-based aggregation [5] and system-theoretic approaches such as balanced truncation [6] offer mathematically principled reductions, but are computationally expensive and strongly operating-point dependent. Measurement-based and data-driven dynamic equivalents [7, 8] emerged as more practical alternatives, yet reduced models are generally valid only near the identification point, may degrade under large disturbances, and obscure internal dynamics. This is particularly problematic for operator-training applications, where preserving internal network structure and component-level behaviour is essential. Key Performance Indicators (KPIs) [9, 10] have been used to identify dynamically relevant components for offline analysis, but are rarely applied as predictive targets for adaptive online simplification.

B. Research Gap and Motivation

Despite recent advances, no existing method simultaneously satisfies all three requirements for large-scale real-time simulation. Offline reduction methods are not contingency-adaptive and may obscure internal dynamics. ML-based approaches output system-level stability labels rather than component-level fidelity maps. Classical structure-preserving methods are strongly operating-point dependent. Moreover, most approaches are validated only on IEEE benchmark systems

979-8-3315-8909-7/26/$31.00 ©2026 IEEE

[3, 11], leaving scalability to realistic large-scale networks insufficiently addressed.

To address these limitations, this work proposes an Adaptive Model Selection (AMS) framework that combines KPI-driven activity estimation with graph-based learning and adaptive topology simplification. The framework (i) predicts component-level dynamic activity from pre-disturbance operating conditions and event location, and (ii) converts those predictions into substation-level simplification decisions, replacing node-breaker models with simplified bus-breaker representations for substations predicted to be dynamically inactive. Component-level predictions are aggregated to ensure consistent substation-level actions, enabling computational acceleration while preserving full fidelity in dynamically important regions.

The remainder of the paper is organized as follows. Section II presents the methodology, Section III describes the case study, Section IV discusses the results, and Section V concludes the paper.

II. Methodology

This section presents the proposed adaptive model-selection methodology, organized into two phases: an offline learning phase, where component-level dynamic activity targets are constructed and graph-based predictors are trained, and an online adaptation phase, where predicted activity classes are converted into simplification decisions. Following a disturbance, different regions of the power system exhibit heterogeneous activity levels depending on electrical distance, operating conditions, and network interactions; simplifying the model uniformly without identifying these regions risks inaccurate or unstable simulations. The framework addresses this by training graph-based surrogates offline and deploying them online to predict which components are dynamically active, so that only inactive regions are simplified while critical areas retain full fidelity.

A. Step 1: Dataset Construction

A representative dataset is generated by simulating a diverse set of operating conditions and disturbance events. Each simulation produces a post-disturbance trajectory, and together these trajectories cover a broad range of system responses. This diversity is important for training robust predictors that generalize across operating regimes.

For each scenario, components that are disconnected or subjected to control actions (e.g., protection trips or operator interventions) are identified first. These are assigned directly to the highest activity class, indicating that no simplification is permitted, and are excluded from the KPI computation and normalization steps that follow.

For all remaining components, the activity index is computed using the maximum moving-window variance formulation of [9]:

$$\begin{aligned} \mathrm{Var}_i^{(k)} &= \mathrm{Var}(x_i(t)), \quad t \in [t_k,\ t_k + \Delta t_{\mathrm{window}}], \\ t_{k+1} &= t_k + \Delta t_{\mathrm{step}}, \quad S_i = \max_k \mathrm{Var}_i^{(k)}. \end{aligned} \tag{1}$$

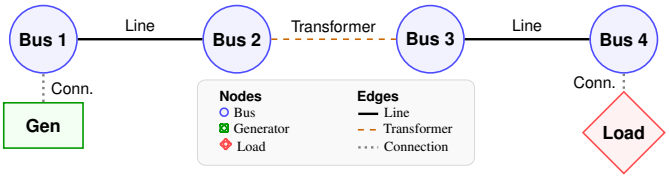

Fig. 1: Graph construction principle used

The variance is computed over a sliding window of length $\Delta t_{\mathrm{window}}$, shifted by Δt_{step} at each step. The activity score S_i is the maximum variance across all windows, capturing the peak dynamic behavior of component i.

To avoid scale bias across heterogeneous physical quantities, normalization is applied separately to each KPI group. A KPI group G collects variables sharing the same physical meaning and units (e.g., all bus voltage magnitudes, or all generator apparent powers). Each group is normalized independently using global min-max scaling over the complete dataset $\mathcal{D}$:

$$\tilde{S}_i^{(G)} = \frac{S_i^{(G)} - \min_{j \in G,\, \mathcal{D}} S_j^{(G)}}{\max_{j \in G,\, \mathcal{D}} S_j^{(G)} - \min_{j \in G,\, \mathcal{D}} S_j^{(G)}}. \tag{2}$$

Global normalization, as opposed to per-scenario normalization, preserves absolute severity information and ensures that activity scores are comparable across different operating conditions and disturbances. In this work, voltage magnitudes form one KPI group and apparent power form another.

The normalized scores $u \in [0, 1]$ are then discretized into K ordinal classes, each representing an increasing level of dynamic severity. The highest class is reserved for disconnected or controlled components and is assigned independently of the KPI discretization.

B. Step 2: Power System Graph Representation

Each scenario is represented as a homogeneous graph, illustrated in Fig. 1. Nodes correspond to buses, generators, and loads; edges correspond to transmission lines, transformers, and connection links. This representation encodes both network topology and electrical interactions in a unified structure suitable for graph-based learning.

A homogeneous formulation is adopted for simplicity and training stability. Type-aware architectures such as R-GCNs (Relational Graph Convolutional Networks) or heterogeneous GATs can define separate message-passing rules per component type, but at the cost of increased model complexity and parameter count. Here, node and edge types are instead encoded directly as input features alongside electrical quantities and impedance-based distances, allowing the model to distinguish component roles without requiring separate embedding spaces. Type-aware extensions are left for future work.

Each node carries a 7-dimensional feature vector $\mathbf{x}_i = [t_i,\ v_i,\ \theta_i,\ p_i,\ q_i,\ f_i,\ d_i]$, where t_i is the node type, (v_i, θ_i) the voltage magnitude and angle, (p_i, q_i) the active and reactive power injections, f_i a binary event indicator, and d_i a graph-distance feature. The graph-distance feature d_i is

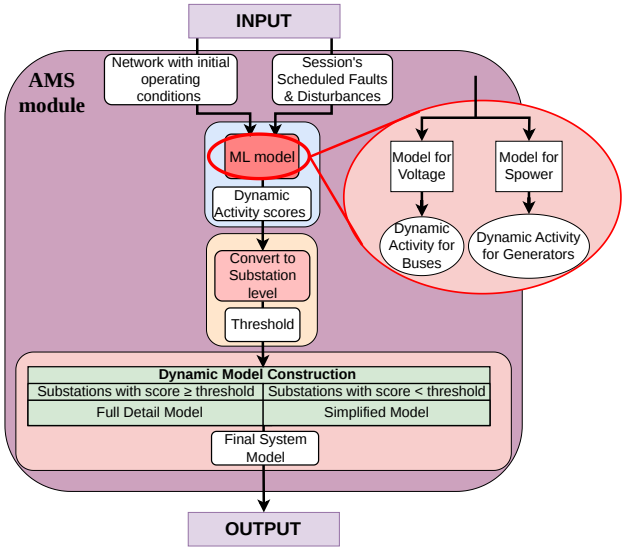

Fig. 2: AMS methodology diagram with parallel Voltage/S_{power} models and substation-level conversion

the shortest impedance-weighted distance from the node to the disturbance location, computed via Dijkstra's algorithm. A logarithmic transform is applied to reduce skewness and improve numerical stability [12]. This feature is well-suited to graph message passing, as it reflects how electrical influence propagates along network paths. Each edge carries a 4-dimensional feature vector $\mathbf{e}_{ij} = [c_{ij},\ f_{ij},\ r_{ij},\ x_{ij}]$, encoding the edge type c_{ij}, a binary event indicator f_{ij}, and the branch resistance and reactance (r_{ij}, x_{ij}).

C. Step 3: Training Architecture and Ordinal Regression Formulation

Two GAT-based predictors are trained in parallel: a bus-level voltage model and a generator-level apparent-power (S_{power}) model. Both share the same edge-aware message-passing backbone (Fig. 3), but differ in their output heads and training masks. The voltage model is trained exclusively on bus nodes, while the S_{power} model is trained exclusively on generator nodes. This masked supervision confines each model to a physically meaningful subset of the system. Through edge-aware message passing, each node aggregates both local electrical conditions and non-local network interactions, producing embeddings that reflect topological structure and electrical coupling. The two models produce independent node-level activity predictions, which are later aggregated into substation-level decisions as shown in Fig. 2.

Activity prediction is formulated as an ordinal regression problem rather than standard multi-class classification. The activity classes represent increasing levels of dynamic severity, and the ordering between classes carries physical meaning: a prediction error of one class is qualitatively less harmful than an error of several classes. Preserving this monotonic structure improves robustness to intermediate misclassifications. The

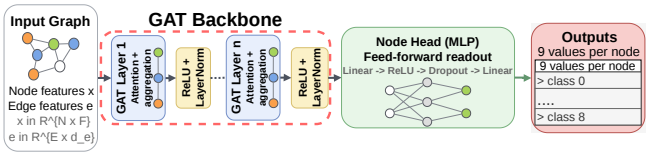

Fig. 3: Graph attention forward architecture for node-level ordinal classification

CORAL framework [13] is adopted for this purpose. CORAL (COntinuous RAnk Logits) decomposes the K-class problem into $K-1$ ordered binary tasks, each corresponding to a severity threshold, and enforces consistent rank ordering across predictions. The GAT encoder captures spatial dependencies in the network, while CORAL enforces the ordered severity boundaries. An additional asymmetric penalty term further biases the model toward conservative predictions in safety-critical conditions.

The training objective combines weighted binary cross-entropy (BCE) over the ordinal levels with a penalty for underestimating severe events. Let $\mathbf{z}_i \in \mathbb{R}^{K-1}$ denote the logits for sample i and $\mathbf{t}_i \in \{0,1\}^{K-1}$ the corresponding CORAL encoding. The objective is:

$$\mathcal{L} = \frac{1}{B(K-1)} \sum_{i=1}^{B} \sum_{k=0}^{K-2} \text{BCE}(z_{i,k}, t_{i,k}; w_k) + \lambda \mathcal{L}_{\text{pen}}, \quad (3)$$

where $w_k = \frac{N_k^-}{\max(N_k^+, 1)}$ corrects for ordinal class imbalance, and $\lambda \geq 0$ controls the strength of the underestimation penalty. The penalty is computed from the predicted severity level $\hat{y}_i$, obtained by aggregating sigmoid probabilities over the ordinal logits. The underestimation error for sample i is $\delta_i = \max(0, y_i - \hat{y}_i)$, and the penalty term is:

$$\mathcal{L}_{\text{pen}} = \frac{1}{\max(|H|, 1)} \sum_{i \in H} \delta_i^2, \qquad H = \{i : y_i \geq \tau\}. \quad (4)$$

The set H contains only samples with true severity at or above threshold τ, focusing the penalty on conditions where underestimation carries the greatest operational risk.

Model selection and hyperparameter tuning prioritize high-severity detection. In addition to early stopping based on validation loss, the decision threshold ϵ used during simplification is selected on the validation set to maximize high-severity recall while maintaining acceptable overall performance. Learning-rate scheduling is driven by a composite validation metric emphasizing recall and F1-score for high-severity classes. Feature scaling uses training-set statistics only, applied separately to node and edge features, to prevent information leakage across dataset splits.

D. Step 4: Online Inference and Adaptive Simplification

During operation, the trained models are deployed for adaptive simplification, as illustrated in Fig. 2. The current system state, including topology, operating conditions, and out-of-service components, is obtained from Supervisory Control

and Data Acquisition (SCADA) or Phasor Measurement Unit (PMU) measurements and converted into a graph using the same construction as during training, with features normalized using training-set statistics. The voltage model predicts activity classes for bus nodes and the S_{power} model predicts classes for generator nodes through masked inference; predictions are aggregated to the substation level by taking the maximum predicted class among all associated buses and generators. Finally, a threshold $\epsilon \in \{0, \ldots, K\}$ governs the simplification decision: substations with predicted class $\geq \epsilon$ are retained in full node-breaker detail, while all others are converted to simplified bus-breaker representations, focusing computational effort on dynamically active regions.

III. Case Study

All simulations are performed using Dynaωo [14], an open-source dynamic simulation environment developed jointly by RTE and academic–industrial partners. Dynaωo provides direct access to solver internals and detailed per-step timing data, making it well suited for evaluating real-time performance. Its availability with a validated French transmission system model further motivates its use.

A. French Power System Dynamic Model

The test network includes more than 6000 buses, 7075 transmission branches, and 335 synchronous generators, each equipped with detailed models for excitation systems, automatic voltage regulators, power system stabilizers, turbines, and speed governors. The network also contains 3039 dynamically modeled loads, including voltage-control devices, voltage-dependent loads, and specialized load models. It spans voltage levels from 63 kV to 400 kV (excluding the eastern 63 kV network), comprises 4009 substations, and includes neighboring substations for boundary consistency. Long-term dynamics are captured through 839 load tap changers distributed over 2138 transformers, with current-limiting and tap-blocking devices. In total, the model contains approximately 319 880 continuous variables, 334 848 discrete variables, and 276 151 root functions, making it substantially larger than the IEEE benchmark systems used in most related work.

B. Scenario Generation

The dataset is built by combining ten operating points (OPs) with a large set of *N-1* contingencies, as listed in Table I. The contingency set contains approximately 1500 unique events, covering disconnections of transmission lines, generators, and other grid components. A filtering step removes contingencies that are inapplicable to a given operating point, retaining only those where the target component exists and is in service. Multi-element disturbances and cascading failures are left for future work.

Each scenario is simulated independently, enabling straightforward parallel execution. After contingency filtering, the final dataset contains 10 360 dynamic simulations. Scenarios are randomly partitioned at the scenario level into 80% training, 10% validation, and 10% test sets, where each scenario corresponds to a unique operating point–contingency pair. The same split is used for both the voltage and S_{power} models, and all reported metrics are evaluated on the 10% test set.

TABLE I: Operating points considered in this work.

OP	Description
1	SCADA snapshot (peak load, ∼80 GW).
2	SCADA snapshot (stressed condition, France–Spain interface).
3	SCADA snapshot (minimum load, ∼30 GW).
4–5	Variants of OP 1 via regional load scaling.
6–10	Variants of OP 2 via uniform and regional scaling.

C. Class Construction and Labeling Strategy

The normalized activity space is uniformly discretized into nine bins, producing classes 0–8. Class 9 is reserved for manually disconnected or controlled components (discrete events) and assigned independently of the KPI discretization. The 10-class scheme is chosen to capture fine-grained variations in dynamic severity: a finer ordinal resolution allows the model to distinguish intermediate activity levels, which is important for setting conservative simplification thresholds. The same labeling strategy is applied to both voltage and apparent power KPIs, ensuring consistency between the two prediction targets.

D. Simplification Configuration

For all experiments, the simplification threshold is set to $\epsilon = 1$, so only substations assigned to Class 0 are simplified; all others retain the full node-breaker representation. This is the most conservative operating point of the framework.

A heuristic simplification strategy, called 'Regional', is used as a baseline. It preserves the directly affected substation and all substations within a third-order graph neighborhood in full detail, converting all others to bus-breaker representations. Unlike AMS, the Regional strategy uses only topological proximity and requires no learned model, making it a transparent reference for assessing the accuracy–performance trade-off.

Fidelity is measured using Dynamic Time Warping (DTW) [15], computed on voltage trajectories relative to the full node-breaker reference simulation. DTW accommodates minor temporal shifts arising from numerical discretization and is well established for electromechanical transient comparison. Speedup and DTW results are reported for three representative contingencies from the test set: a transmission line disconnection, a busbar section disconnection, and a generator disconnection, selected to cover different disturbance categories.

IV. Results

To illustrate the ordinal classification behaviour, Fig. 4 shows representative apparent power and voltage time series across activity classes. The class index reflects increasing system dynamics: lower classes correspond to near steady-state conditions, while higher classes capture stronger fluctuations and transient events. This ordinal structure is explicitly exploited in the proposed learning framework (Section II).

The S_{power} and voltage models are evaluated on 1036 test scenarios. Prediction outcomes at node and substation level are

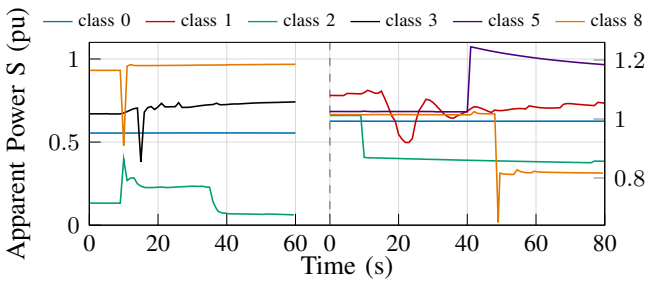

Fig. 4: Representative apparent power and voltage time-series examples across ordinal activity classes.

TABLE II: Prediction outcomes at node and substation level

Level / Model	Correct	Under	Over	MAE
Spower	79.47% (445432/560475)	0.01% (75/560475)	20.51% (114968/560475)	1.8467
Voltage	93.26% (5392523/5782301)	0.01% (419/5782301)	6.73% (389359/5782301)	0.6017
Substation	94.26% (4431815/4701914)	0.01% (339/4701914)	5.74% (269760/4701914)	0.5361

reported in Table II. At node level, the voltage model achieves 93.26% accuracy and the S_{power} model achieves 79.47%, with underestimation below 0.01% in both cases. The near-zero underestimation rate is a direct consequence of the asymmetric loss, which penalizes underestimation more strongly than overestimation, enforcing a conservative prediction bias. This is intentional: in safety-critical applications such as operator training, underestimating the activity of a substation risks hiding critical dynamic conditions, whereas overestimation leads only to conservative but safe simplification decisions.

Accordingly, the dominant error type is overestimation: 20.51% for S_{power} and 6.73% for voltage. The lower accuracy of the S_{power} model is partly attributable to its restriction to generator nodes, which reduces the number of labeled training samples and introduces more heterogeneous dynamics compared to bus-level voltage prediction. The Mean Absolute Error (MAE) column in Table II quantifies the average magnitude of ordinal errors, accounting for both under- and overestimation. At substation level, max-based aggregation improves accuracy to 94.26% while maintaining 0.01% underestimation; MAE also decreases to 0.5361, reflecting reduced local prediction noise and improved robustness after aggregation.

The impact on model complexity is summarized in Table III. The present analysis is restricted to substation-level topology reduction (node-breaker to bus-breaker); simplification of generator, load, and control component models is left for future work. In the reduced configuration, substations with $\hat{y}_{\text{sub}} < \epsilon$ are simplified by removing internal switching devices. Across the test set, AMS achieves average reductions of 63.54% in continuous variables, 58.25% in discrete variables, and 31.69% in root functions. The comparatively smaller reduction in root functions reflects their indirect dependence on network topology rather than a direct correspondence to substation switching devices.

TABLE III: Reduction in system size (continuous variables, discrete variables, root functions)

System	System Size	System Size (%)
Full	(319880, 334848, 276151)	(100.00%, 100.00%, 100.00%)
Regional	(102328, 136081, 225216)	(27.45%, 34.30%, 64.18%)
AMS	(137962, 169460, 240838)	(37.01%, 42.71%, 68.63%)

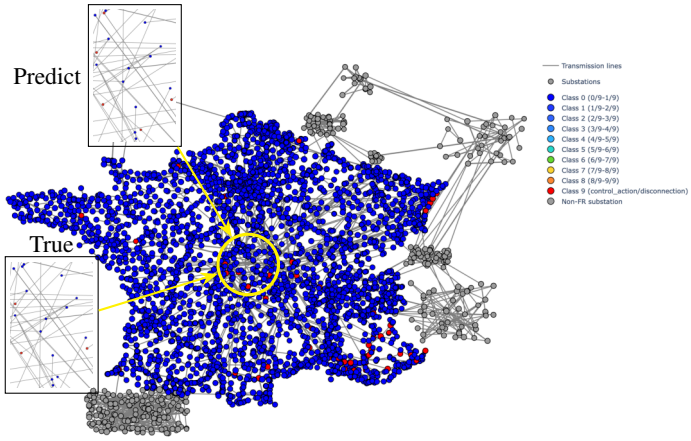

Fig. 5: Comparison between predicted and ground-truth substation classifications for a generator disconnection event.

Figure 5 compares predicted and ground-truth substation classifications for a generator disconnection event. The main plot shows the network colored by activity class, with red indicating highly dynamic substations that require full-detail modeling. The zoomed areas focus on the fault region: the lower panel shows the true values and the upper panel shows the model predictions. The classifier correctly identifies most affected substations; the residual errors are predominantly overestimations, consistent with the conservative design.

Real-time performance results are reported in Table IV and Fig. 6. The full node-breaker model incurs multiple real-time violations per contingency, confirming that it cannot meet the real-time execution budget without reduction. The Regional model eliminates all violations but at a significant accuracy cost: DTW errors are on the order of 10^{-2}, three orders of magnitude above the solver tolerance of 10^{-4}, indicating non-negligible trajectory deviation. AMS achieves a more balanced outcome: only one timing violation per contingency remains, solver time per step is substantially reduced relative

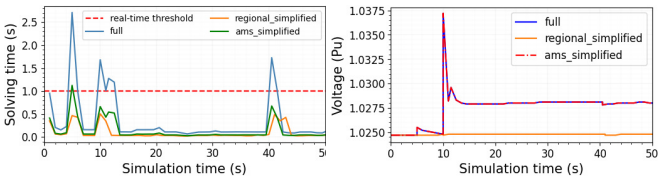

Fig. 6: (Left) Solver time comparison. (Right) Dynamic response across model configurations.

TABLE IV: Real-time performance and voltage fidelity comparison across scenarios and models

Scenario	Model	RT Viol.	$\|\|T_{sol}\|\|_\infty$	DTW	Speedup
Line Disc.	Full	8	2.7245	ref	ref
	Regional	0	0.5056	0.035	5.39x
	AMS	1	1.1268	1.55×10^{-5}	2.42x
Busbar Sect.	Full	7	2.7321	ref	ref
	Regional	0	0.5132	0.036	5.32x
	AMS	1	1.1036	1.53×10^{-5}	2.48x
Generator	Full	23	2.6736	ref	ref
	Regional	0	0.4547	0.056	5.88x
	AMS	1	1.0390	6.2932×10^{-5}	2.57x

to the full model, and DTW errors remain around 10^{-5}, well below solver tolerance. This confirms that the trajectory deviation introduced by AMS simplification is negligible from a numerical standpoint.

Overall, AMS reduces timing violations and preserves high-fidelity dynamics, but does not yet fully guarantee real-time compliance under severe transients. Its key advantage over the Regional strategy lies in selectivity: rather than applying a uniform topological rule, AMS focuses full-detail modeling on the substations predicted to be dynamically active, concentrating computational effort where it is most needed.

V. Conclusion

This paper presents the Adaptive Model Selection (AMS) framework, which uses graph-based ordinal predictors to guide substation-level simplification for real-time dynamic simulation. In the current formulation, simplification is restricted to substations; extensions to generators, loads, and control system models are left for future work.

The results illustrate a clear trade-off between computational efficiency and simulation accuracy across the three strategies evaluated. The full node-breaker model preserves complete dynamic fidelity but fails to meet the real-time execution budget during contingencies. The Regional strategy achieves faster execution by simplifying all topologically distant substations, but at the cost of significant trajectory deviation. AMS occupies the middle ground: it operates close to the real-time budget, with only isolated timing violations during severe transients, while keeping DTW errors below the numerical solver tolerance. This confirms that the trajectory deviation introduced by AMS simplification is negligible in practice. By selectively preserving full-detail models only where dynamic activity is predicted, AMS delivers substantial computational savings without meaningful loss of accuracy, supporting its suitability as a practical approach to real-time simulation of large-scale power systems.

Several directions remain open for further development. First, the simplification framework will be extended beyond substations to include generators, loads, and control system models, enabling a more complete system-wide reduction. A key objective of this extension is to further reduce computational cost and achieve consistent real-time compliance even under severe transients. Second, the underestimation rate, while already near zero, will be further reduced through improved asymmetric loss formulations and refined threshold selection strategies that more rigorously enforce conservative predictions. Third, the current framework is limited to single-element contingencies; handling cascading and multi-event disturbances will require richer event representations capable of capturing the spatial propagation and interaction effects of simultaneous or sequential faults. Finally, incorporating additional system-level operating descriptors into the graph representation may improve robustness and generalization across a wider range of grid conditions and topological configurations.

Acknowledgment

This work has received funding from the European Commission under the Horizon Europe programme, Grant Agreements No. 101136119 (TwinEU project) and No. 101120278 (DENSE). Additional support was provided through the TRAISIM project, which is funded by CRESYM.

References

[1] S. Panagi, T. Vermeulen, G. Joseph, and P. Aristidou, "Towards an open-source real-time operator-training platform: Analysis of computational efficiency," in *Proceedings of the XXIV Power Systems Computation Conference (PSCC)*. Limassol, Cyprus: IEEE, 2026, to appear.

[2] F. Gomez, A. Rajapakse, U. Annakkage, and I. T. Fernando, "Support vector machine-based algorithm for post-fault transient stability status prediction using synchronized measurements," *IEEE Transactions on Power Systems*, 2011.

[3] J. Huang, L. Guan, Y. Su, H. Yao, M. Guo, and Z. Zhong, "Recurrent graph convolutional network-based multi-task transient stability assessment framework in power system," *IEEE Access*, 2020.

[4] S. Song, S.-W. Min, and S. Jung, "Steady-state data-driven dynamic stability assessment in the energy management system," *Research Square (Research Square)*, 2024.

[5] R. Podmore, "Identification of coherent generators for dynamic equivalents," *IEEE Transactions on Power Apparatus and Systems*, 1978.

[6] M. Safonov and R. Chiang, "A schur method for balanced-truncation model reduction," 1989.

[7] Z. Jiang, X. Zhang, Y. Xue, A. G. Tarditi, and Y. Liu, "Measurement-based power system dynamic model reduction using arx equivalents," 2019.

[8] N. Tong, Z. Jiang, S. You, L. Zhu, X. Deng, Y. Xue, and Y. Liu, "Dynamic equivalence of large-scale power systems based on boundary measurements," *American Control Conference*, 2020.

[9] P. Aristidou and T. V. Cutsem, "Parallel computing and localization techniques for faster power system dynamic simulations," *""*, 2014.

[10] M. Sajjadi, T. Xia, M. Xiong, K. Sun, A. Hoke, B. Wang, and J. Tan, "A participation factor-based approach for defining the emt model boundary for power system simulations with inverter-based resources," *IECON 2020 The 46th Annual Conference of the IEEE Industrial Electronics Society*, 2024.

[11] P. Sun, L. Huo, S. Liang, and X. Chen, "Fast transient stability prediction using grid-informed temporal and topological embedding deep neural network," *ArXiv*, 2022.

[12] R. West, "Best practice in statistics: The use of log transformation," *Annals of Clinical Biochemistry*, 2021.

[13] W. Cao, V. Mirjalili, and S. Raschka, "Rank consistent ordinal regression for neural networks with application to age estimation," *Pattern Recognition Letters*, 2020.

[14] "Dynawo: A comprehensive power system dynamic simulator," 2025, supported by the Linux Foundation Energy initiative. [Online]. Available: https://lfenergy.org/projects/dyna%CF%89o

[15] D. Fabozzi and T. Cutsem, "Assessing the proximity of time evolutions through dynamic time warping," 2011.