TRAISIM

Training Simulator for Power System Operators

S. Panagi[1] | S. Eleftheriadis[1] | T. Vermeulen[2] | G. Joseph[3]
P. Aristidou[1]

[1]Cyprus University of Technology [2]Réseau de Transport d'Électricité [3]Collaborative Research for Energy System Modelling

MOTIVATION

Why Real-Time Operator Training?

Modern grids are complex **cyber-physical infrastructures** integrating renewables, distributed resources, and advanced protection. **TSOs** must make real-time decisions amid intricate model interactions and emerging AI-support tools.

Existing platforms rely on **static or simplified modules**, reduced-order models, or pre-computed scenario databases — most **closed-source** and vendor-locked, limiting transparency and excluding smaller TSOs and academia. Operators face **cascading failures, renewable uncertainty, and multi-layer control** that demand immersive, real-time simulation with full DAE solvers retaining operational fidelity.

Core Challenge

Simulate a **6,000+ bus** network faster than real-time on **commodity hardware**, within the **1–2 s** SCADA refresh constraint, with full **model-level** fidelity.

Simulation Framework

An **open-source** platform on **Dynawo**: algorithmic transparency, no license costs, deployable on **commodity x86 hardware** in existing control centres.

ARCHITECTURE

Operator Training Platform

Six tightly integrated subsystems coordinate to deliver a realistic, real-time control-room experience for power system operators.

Trainer supervises & injects | **Game Master AI** faults · scenarios | **HMI** control-room UI

↕ scenario events & operator actions ↕

Orchestrator time-sync · SCADA ⟷ **Info Flow** latency · comm. ⟷ **Automata** protection FSMs

↑ real-time system state updates ↑

Physical Simulator
Dynawo · DynaWaltz DAE Solver
RTE French Grid · 6,000+ buses · up to 320k variables

KINSOL Newton | KLU sparse LU | Adaptive timestep

HEADROOM METRIC
$\delta t = T_{\text{sync}} - T_{\text{sol}}$
$\delta t > \delta t_{\text{min}} \geq 0$ required

SIMULATION STEP
$T_{\text{sync}} = 1\,\text{s}$
physical simulator

SCADA RATE
1–2 s
end-to-end refresh

TEST CASE

RTE French Transmission Network

Full model of France from 63 kV to 400 kV with detailed dynamic components.

7,075 Branches	**335** Sync. Generators
3,039 Dynamic Loads	**839** Load Tap Changers

THREE MODEL FIDELITY LEVELS

- **Model 1:** Simplified topology · 80k variables · ✓ Real-time
- **Model 2:** Breakers tracked · 210k variables · △ Marginal
- **Model 3:** All switches · 320k variables · × Violations

14 contingency scenarios: Loss of Busbar Section · Loss of Line · Loss of Generator

BENCHMARKING RESULTS — YEAR 1

Computational Performance & Solver Tuning

Worst-Case Solving Time T_{sol} (s) vs. 1 s Budget

PSCC 2026 — "Towards an Open-Source Real-Time Operator-Training Platform: Analysis of Computational Efficiency" — submitted & accepted, Limassol, June 8–12, 2026

KINSOL Time Distribution — Model 3 (LG Event)

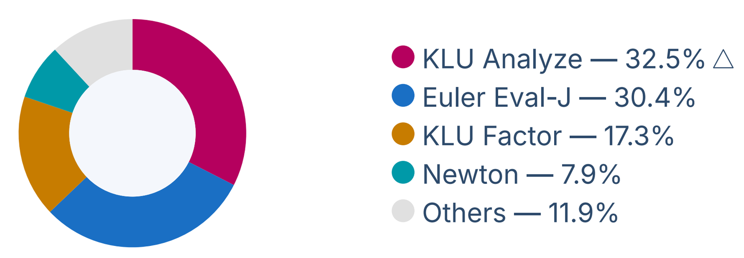

Root cause: KLU Analyze (symbolic factorization) re-analyzes the full Jacobian sparsity structure on every topology change, scaling **2.3× longer** in Model 3 vs Model 2.

Key Findings

- ✓ **Model 1 — Fully Real-Time**
  0 violations across all 3 scenario types on laptop hardware
- ✓ **Tuning: 2× Speedup**
  fnormtolAlgJ + msbsetAlgJ tuning eliminates violations for **Model 2 LBS & LL**
- × **Model 3 — Persistent Issues**
  Jacobian structure analysis overhead critical; parameter tuning insufficient
- ! **300% Variable Explosion**
  Full-switch model: 80k→320k variables, driving KLU Analyze to **dominant bottleneck**

EXTENSION 2026 — WORKPLAN

Four Work Packages

WP1 · Solver Optimization

- Profile & reduce KLU Analyze call frequency; skip re-analysis when discrete events preserve matrix sparsity
- Incremental Jacobian structure updates: warm-start BTF permutations & cache sparsity patterns for known post-event topologies
- Thread-level parallelism for Euler Evaluation-J & KLU Factor via OpenMP
- Evaluate drop-in KLU replacements for parallel numeric factorization

WP2 · Adaptive Model Management

- AMS module: GATv2 with multi-task prediction heads → per-component activity scores & right-sized models
- Dynamic fidelity rules per training scenario; Ward-equivalent fallback for latent sub-networks below activity threshold
- Expand training set with high-RES operating points & frequency-centric contingencies
- Full prototype & accuracy–performance trade-off evaluation using DTW-based metrics across 3 KPI groups

WP3 · Orchestrator Co-Design

- Slack-time API between orchestrator & simulator: signal available headroom & anticipated-event queue from Game Master
- Dead-time offloading: pre-compute & cache KLU Analyze for anticipated topological changes during non-critical intervals
- Adaptive solver aggressiveness: dynamically tighten/ relax tolerances based on real-time headroom budget
- Elastic refresh policies for worst-case spikes with time-borrowing across consecutive steps

WP4 · Validation & Dissemination

- Recommendations report with implementation roadmap for upstream Dynawo integration
- Open-source release of solver patches & AMS module on SPS-L GitHub repositories
- Three peer-reviewed papers: benchmarking (PSCC 2026), solver optimisation (POSYDYS) & AMS methodology (journal)

EXTENSION 2026 — ADAPTIVE MODEL SELECTION (AMS)

AI-Driven Fidelity Control

The AMS module predicts **component-level dynamic activity** before each training session, selecting only the fidelity needed for the active scenario. A **Graph Attention Network (GATv2)** learns which network components participate actively in disturbance propagation, enabling hybrid models that retain full detail where needed and simplify elsewhere.

AMS INTEGRATION WITH TRAINING PLATFORM

Game Master Orchestrator

Initial OP + Contingencies ↓

AMS Module GATv2 + MLP heads ⟵ **IIDM + DYD** graph construction

Activity scores [0,1] per component ↓

Full Detail score $> \varepsilon$ + **Simplified** score $< \varepsilon$ → **Hybrid IIDM+DYD** right-sized model

↓ fed to Physical Simulator (Dynawo)

DYNAMIC ACTIVITY KPI

$$Var_i^k = \text{Var}(x_i(t),\ t \in [t_k,\ t_k{+}\Delta])$$
$$S_i = \max_k\ Var_i^k$$
$$\tilde{S}_i = \text{MinMax}(S_i) \ \in [0,\ 1]$$

3 KPI groups independently normalized: **Generator Spower · Frequency · Voltage**

EARLY RESULTS

Continuous Variables	−75.7% [74.5–76.4%]
Discrete Variables	−69.2% [67.5–69.7%]
Root Functions	−37.4% [37.1–37.6%]
Power Prediction Accuracy	99.35% (+0.64% / -0.008%)
Voltage Prediction Accuracy	99.71% (+0.29% / -0.006%)

Training data: ~12,000 dynamic simulations across multiple operating points and N-1 contingencies. GATv2 with node prediction heads.

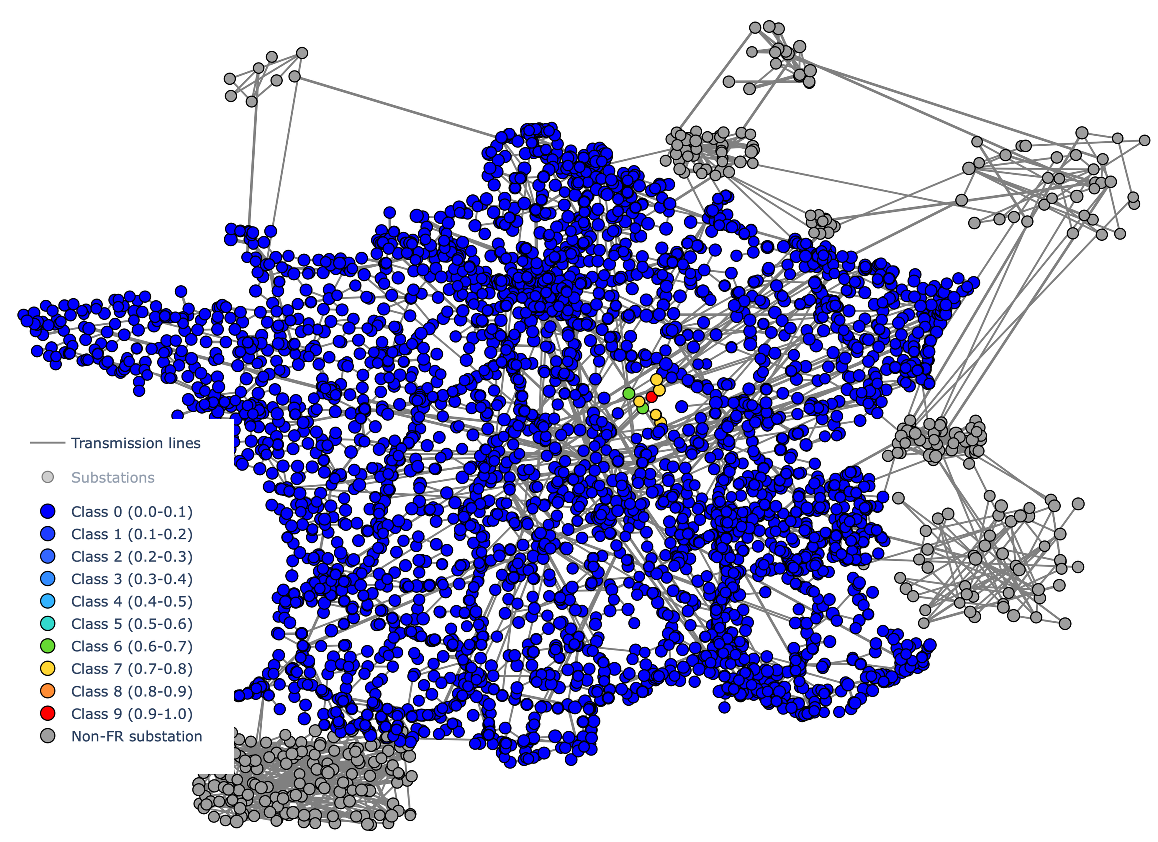

PROJECT TIMELINE & DELIVERABLES

Phase 1 — Benchmarking & Profiling ✓

JAN–APR 2025
Task 1: Scenario Execution ✓
Benchmarking methodology on RTE 6 k-bus French grid · 1 s step validated as best speed–accuracy compromise · Headroom metric $\delta t = T_{\text{sync}} - T_{\text{sol}}$ established · 14 contingency scenarios catalogued (LBS, LL, LG)

MAY–AUG 2025
Task 2: Bottleneck Analysis ✓
Code-level profiling on worst-case event steps · KLU analyze identified as primary hotspot (32.5% of KINSOL time, 2.3× longer in M3 vs M2) · parameter tuning yields 2× speedup on topology events

SEP–DEC 2025
Task 3: Advanced Assessment ✓
3-model fidelity comparison: 80 k → 210 k → 320 k variables (300% explosion) · Tuning eliminates M2 violations for LBS/LL but M3 remains infeasible · PSCC 2026 paper submitted & accepted · 1st & 2nd Technical Reports delivered

Phase 2 — WP2 Adaptive Model Management ⚙

JAN–MAR 2026
WP2 Baseline & First AMS Results ✓
~12 k dynamic simulations across 10 operating points & N-1 contingencies · GATv2 + MLP heads trained on IIDM-derived graphs · 75% variable reduction · Frequency KPI issue identified: global vs. localized behavior → task delayed

APR–JUN 2026
WP2 Model Refinement & Expansion ⚙
Split prediction into 3 models (S-power, voltage, frequency) · Retrain with separate frequency model using global features (inertia, lost power) · Add high-RES & winter worst-case operating points from RTE · Hyperparameter tuning & DTW-based evaluation across 3 KPI groups

JUL–SEP 2026
AMS Consolidation & Dissemination ⏳
Full AMS prototype: IIDM/DYD → GATv2 → hybrid right-sized model fed to Dynawo · Accuracy–performance trade-off validated on extended scenario set · Journal paper drafted on AMS methodology

Phase 2 — WP1 / WP3 / WP4 Status ⚙

APR–JUN 2026
WP1 — Solver Optimization ⚙
Optimization roadmap on GitHub (quick wins → advanced) · Phase 0: event-severity classification & adaptive factorization control (KLU numerical-only refactor) · Large-scale profiling over 12 k scenarios underway · Target: ~30% overall speedup

APR–SEP 2026
WP4 — Validation & Dissemination ⚙
PSCC 2026 presentation (Limassol, Jun 8–12) · POSYDYS solver paper in progress · AMS journal paper planning initiated · Final integrated benchmark planned against baseline

APR–SEP 2026
WP3 — Orchestrator Co-Design ⏳
Deferred until WP1 & WP2 prototypes stabilize · Slack-time API spec & pre-computation caching logic to follow · Will leverage WP1 KLU-Analyze cache for anticipated-event offloading

Sustainable Power Systems Lab | Cyprus University of Technology | Rte | CRESYM

Project: TRAISIM | **Funded by:** CRESYM | **Code:** github.com/SPS-L/ | **Web:** sps-lab.org/ | **Contact:** info@sps-lab.org

CRESROADS 2026