Journal of Multidisciplinary Engineering Science and Technology (JMEST)
ISSN: 2458-9403
Vol. 13 Issue 4, April - 2026

Parametric Geometry Generation and Multi-Scale GPU-Accelerated Thermal Simulation of Air-Insulated Substations

Antonis Papadakis[1,2]*, Sotiris Pigiotis[1], Filotheos Zachariadis[1], Eleni Constantinide[1]
[1]KYAMOS LTD, 37 Polyneikis Street, Strovolos, 2047, Nicosia, Cyprus
[2]Frederick University, 7. Y. Frederickou, Pallouriotisas, 1036, Nicosia, Cyprus

Abstract— An integrated platform is presented for the automated generation of substation geometry and for the multi-scale thermal simulation of air-insulated electrical substations (AIS). The parametric geometry agent is implemented as a FastAPI microservice and produces simulation-ready STL files from engineering specifications through CadQuery-based solid modelling. The component library comprises busbars, breakers, clamps, transformers, routed jumper tubes, and support insulators, organised into three preset configurations that range from a single-junction prototype to a five-bay yard-scale substation. The generated geometry is consumed directly by a GPU-accelerated Lattice Boltzmann solver for coupled thermal-flow simulation. A multi-scale simulation strategy comprising four complementary levels is demonstrated: (i) two-dimensional thin-slab parametric sweeps (40 scenarios, 256² grid) for parametric design exploration; (ii) three-dimensional scaled surrogate campaigns (40 scenarios, 80³ grid) for the generation of AI training data that cover a range of operating regimes at tractable computational cost; (iii) full-scale simulation on real STL geometry (845,000 cells) for production-scale validation; and (iv) realistic multi-component AIS simulation (33 million cells) for engineering demonstration. The platform supports automated parametric design studies of substation thermal behaviour without recourse to manual CAD modelling, thereby establishing a direct pathway from engineering specification to physics-based simulation.

Keywords— Parametric geometry; air-insulated substation; STL generation; multi-scale simulation; Lattice Boltzmann Method; GPU computing; CadQuery; FastAPI; thermal management

I. INTRODUCTION

The design and thermal assessment of air-insulated electrical substations has traditionally relied on manual CAD modelling followed by mesh generation and simulation by the finite-element or finite-volume method. This workflow is time-consuming, requires specialist CAD expertise, and constrains the number of design variants that can practically be evaluated within a given engineering timeframe. For parametric studies in which busbar spacing, conductor sizing, heater placement, or bay configuration must be varied systematically, the manual workflow becomes computationally and organisationally prohibitive.

An alternative approach is proposed in the present work, based on a parametric geometry agent that generates substation CAD models automatically from engineering specifications, coupled with a GPU-accelerated Lattice Boltzmann method (LBM) solver [1], [2] implemented in CUDA [3] that performs the thermal simulation. The complete workflow, extending from specification to thermal results, is automated and supports parametric design exploration without manual intervention.

The present work also addresses the challenge of multi-scale simulation for AI training. The generation of comprehensive training datasets requires a large number of simulations at moderate resolution, whereas validation requires high-resolution simulations performed on geometry derived from real engineering specifications. A four-level simulation strategy is proposed that balances computational cost against physical fidelity at each scale.

The scope of the present manuscript is deliberately focused on the automated parametric geometry pipeline and the multi-scale simulation strategy that this pipeline supports. Quantitative results for the production AI surrogate (the TUNet2D and TUNet3D champion models), full-scale thermal field predictions, and detailed AI-versus-solver accuracy comparisons are reported in the companion paper [4]. Section III.B and Section VI of the present work provide standalone evidence of the solver's physical fidelity, of the suitability of the generated fields for AI-surrogate training, and of the engineering value of the parametric pipeline.

II. PARAMETRIC STL AGENT

A. *Architecture*

The geometry generator is implemented in Python using CadQuery [5] for solid modelling and served as a REST API via FastAPI [6]. The service exposes REST endpoints for health checking, preset listing, preset generation, and custom YAML-based generation (Fig. 1).

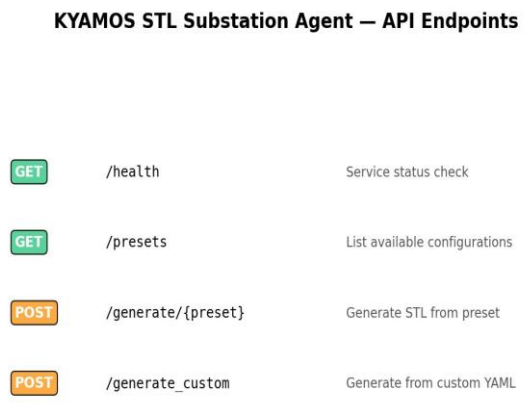

Fig. 1. KYAMOS STL Substation Agent — API endpoints.

B. *Component Library*

The generator provides parametric implementations of standard AIS components:

TABLE I. Component library.

Component	Modelling Approach	Key Parameters
Busbars	Hollow cylinders (outer and inner circles, extruded)	Outer diameter, wall thickness, length, phase spacing
Breakers	Solid cylinders inline with bus	Diameter, length, bay position
Clamps	Rectangular blocks with cylindrical saddle cut	Width, height, gap, bus outer diameter
Transformers	Box solids	Width, height, depth, position
Jumpers	Multi-segment polyline tubes	Outer diameter, wall, waypoints, routing
Insulators	Cylinders with flanges	Diameter, height

Busbars are modelled as hollow tubes rather than solid cylinders because HV rigid busbars are physically hollow in accordance with standard substation practice [7]. At AC power frequencies the current concentrates in a thin outer annulus due to the skin effect, so the inner material of an equivalent solid bar would carry negligible current; hollow tubes therefore provide comparable ampacity at substantially lower mass and cost. The hollow geometry must be represented in the simulation in order to capture the correct outer surface area, the correct convective boundary condition, and the correct conductor thermal mass. Clamps are modelled as rectangular blocks with cylindrical saddle cuts because this geometry captures the realistic thermal contact area between the support and the bus conductor, which controls the effective heat-loss path from the energised bus into the supporting structure. Transformers are modelled as simple rectangular boxes because, for simulation of the air domain surrounding the equipment, the external envelope and surface area dominate the thermal exchange with the surrounding air; internal structure and cooling-fin detail influence the transformer's own thermal state but have only minimal effect on the adjacent air-plume behaviour. Jumpers are modelled as hollow tubes following multi-segment polylines because the geometric path of the jumper, including its changes in orientation between bay equipment, determines the local convective heat-transfer coefficient along each segment, a horizontal segment sits within its own rising thermal plume, whereas a vertical segment does not, and because the jumper also provides a thermal conduction path between adjacent bays.
External hardware, including bushings, radiators, support brackets, wheels, and control cabinets, is intentionally omitted from the component library. For the thermal simulation use case, these features generate negligible Joule heating themselves and present a small thermal mass compared with the primary conductor and transformer envelopes, so their omission represents an acceptable fidelity-cost trade-off. The procedural generator can be extended with these components without changes to the HTTP service contract or the existing preset configurations, should a later study require radiative-cooling dominance or detailed mechanical assessment.

C. *Preset Configurations*

TABLE II. Preset configurations.

Preset	Components	Domain	Use Case
Simple	1 bus, 1 breaker, 1 trafo	~5 m × 3 m × 3 m	Solver testing
ais_demo	3-phase, 2 bays, 2 trafos	34 m × 13 m × 6 m	Engineering demo
ais_grande	3-phase, 5 bays, configurable	80 m × 36 m × 14.5 m	Yard-scale design

D. *YAML Specification Format*

Each preset is driven by a YAML configuration that maps directly onto CadQuery modelling parameters. An excerpt of the ais_grande configuration illustrates the convention:

```
yard:
  bus_x_span: 72.0          # meters
  bus_height_z: 7.5
  phase_y: [-4.5, 0.0, 4.5]   # three-phase lateral offsets
  bay_centers_x: [-24.0, -12.0, 0.0, 12.0, 24.0]
  transformer_center: [0.0, -10.0, 2.25]
busbar:
  od: 0.12                  # outer diameter (HV tubular busbar)
  wall: 0.012
```

```
breaker:
  diameter: 0.12
  length: 1.80
  center_z: 1.70
jumper:
  od: 0.07
  style: elbow              # vertical drop then horizontal
  vertical_drop: 2.5
bay_targets: [line, feeder, transformer, feeder, line]
export:
  out_dir: stl
  write_unions: true
```

Each parameter carries a direct geometric meaning in metres or integer indices. Modification of the substation layout requires alteration of this file alone: phase spacing, bay count, transformer location, conductor sizing, and jumper routing style are each exposed as top-level fields. Design variations within the envelope of a preset do not require modification of the underlying CadQuery code. The Agent validates the submitted YAML, executes the corresponding generator, and returns the STL, STEP, and manifest outputs within a single HTTP transaction, typically completing in one to five seconds depending on preset complexity.

E. *Output Format*

The generator produces individual STL files per component category, each tagged with a region_id that maps to the solver's material system. A manifest.json file provides component metadata. A STEP assembly file is generated for CAD preview. All outputs feed directly into the solver's voxelizer without format conversion.

F. *Generated Geometry Example*

Fig. 2 shows the three-dimensional CAD rendering of the ais_grande preset: a yard-scale air-insulated substation comprising three parallel phase busbars, five breaker bays, a central transformer, routed jumper connections, and support insulators. The geometry spans 80 m × 36 m × 14.5 m and is generated from a YAML specification file in a few seconds.

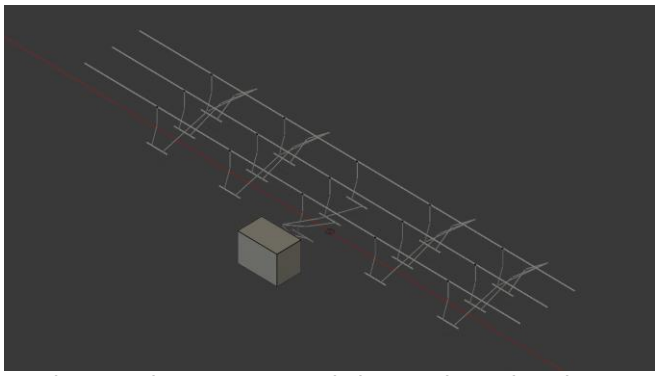

Fig. 2. CadQuery-generated ais_grande yard: 3-phase busbars, 5 bays, transformer, jumper connections.

G. *Dual-Client Architecture*

The Agent is exercised through two independent client paths that share a common backend service. A standalone wxWidgets 3.2.8 desktop client (Fig. 3) provides operator-workstation access for the composition of YAML configurations, the dispatch of preset and custom generation requests, and the inspection of returned STL, STEP, and manifest file lists. A browser-based WebGL viewer provides an interactive three-dimensional preview of the generated geometry, supporting rotation, zoom, and preset switching; this client is intended for design-review workflows in which visual confirmation of the requested layout is required before an expensive simulation is committed. Both clients invoke the same HTTP endpoints. Geometry files are written to disk by the service and referenced by path in the response, rather than transmitted inline, which keeps the HTTP layer lightweight. This separation of client logic from service logic permits the addition of further interfaces, such as voice-driven or schema-driven assistants, without modification of the generator or its contract.

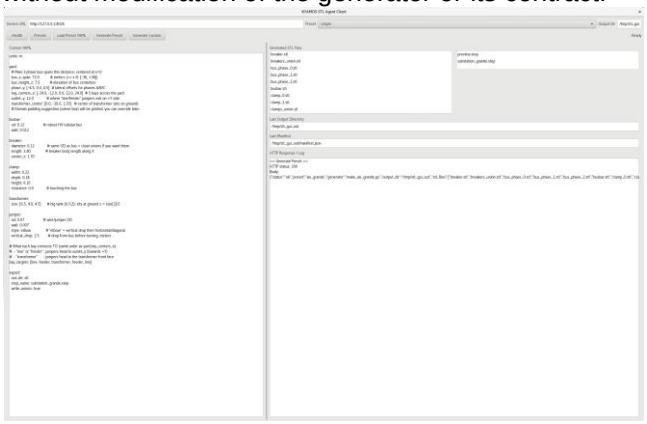

Fig. 3. KYAMOS STL Substation Agent wxWidgets desktop client exercising the ais_grande preset against the service; the HTTP log pane confirms a 200 response and the STL and STEP file lists show the generated geometry written to the configured output directory.

III. SOLVER VALIDATION AND SURROGATE FEASIBILITY

A. *Quantitative Benchmark Validation*

The Lattice Boltzmann solver underlying all four simulation scales has been validated against five independent benchmarks that span flow-only, thermal-only, and coupled thermal-flow physics. The solver is based on the D3Q19 multi-relaxation-time (MRT) collision scheme [2] with buoyancy coupling through the Guo forcing term [8]. Table III summarises the benchmark set, the physics verified by each case, and the measured error metrics. All benchmarks pass with errors below 4 %; the coupled-flow benchmarks (Poiseuille flow and the de Vahl Davis cavity at Rayleigh number $Ra = \beta g \Delta T L^3/(\nu\alpha) = 10^4$) pass with errors below 0.5 %, which are among the stricter quantitative criteria available in the open literature for coupled buoyancy flows in this regime. Detailed

descriptions of the benchmark test cases are provided in the companion paper [4].

TABLE III. Solver benchmark validation results.

Benchmark	Physics verified	Metric	Result
Poiseuille channel flow	Flow profile + no-slip wall boundary condition	NRMSE	3.4%
Thermal source 1D diffusion	Heat source + diffusion across MPI rank boundary	L2 norm	1.73×10^{-4}
Heated cavity at Ra = 0	Pure conduction (Nu = 1 reference)	Nu error	0.008%
Heated cavity at Ra = 10^4	Coupled buoyancy (de Vahl Davis reference [9])	Nu error	0.4%
MPI consistency	Multi-rank vs single-rank field equivalence	Max Δ	0 (identical)

B. Canonical Two-Dimensional Conduction Test

The benchmarks in §III.A validate the physical fidelity of the solver. In this subsection a canonical two-dimensional conduction test serves a dual purpose: it further validates the steady-state solver output on a curved solid–fluid interface, and it demonstrates that the generated thermal fields are of suitable quality for downstream AI-surrogate training. Prior to the application of the solver to complex substation geometries, this test has been used for validation. A heated circular inclusion at dimensionless temperature 50, embedded in a uniform background field at dimensionless temperature 30, is simulated on a 40 × 40 two-dimensional grid. The reported temperature values are non-dimensional, normalised by a reference temperature, and do not represent thermodynamic absolute temperatures. The test evaluates the steady-state thermal field topology and the behaviour of the solver in the neighbourhood of a curved solid–fluid interface, independently of the absolute temperature scale.

Fig. 4 presents the ground-truth temperature field at steady state: a smooth circular thermal distribution centred on the heated inclusion, consistent with the analytical solution for two-dimensional conduction around a circular heat source in an otherwise uniform medium. Fig. 5 presents the prediction obtained from a UNet demonstrator retrained on this 2D dataset for the present revision. The model reproduces the overall location and approximate extent of the heated inclusion. The thermal gradient across the solid–fluid boundary is rendered more diffusely than in the ground-truth field, reflecting the smoothing bias characteristic of convolutional architectures trained with L2 loss on small two-dimensional datasets. The canonical test in this section serves to validate the solver output (Fig. 4); quantitative performance of the project's production AI surrogate (the TUNet2D champion) on the substation thermal simulation campaigns is reported in the companion paper [4].

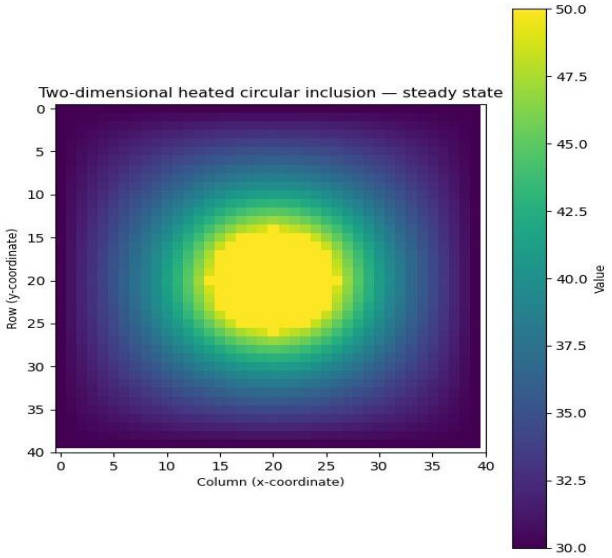

Fig. 4. Solver output: steady-state temperature field (dimensionless, normalised by a reference temperature) for the two-dimensional heated circular inclusion (ground truth).

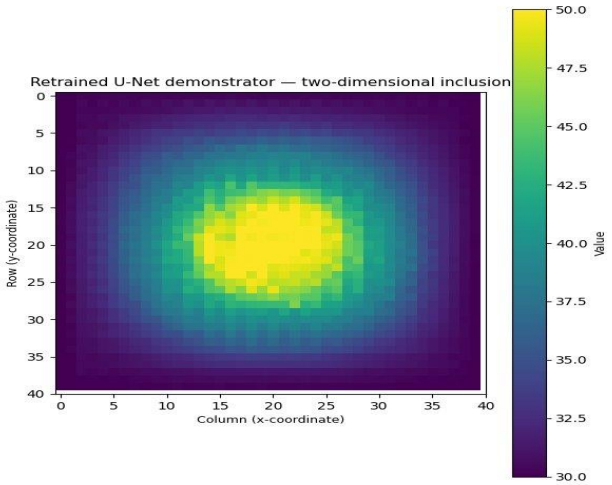

Fig. 5. AI prediction: retrained UNet demonstrator output for the same two-dimensional heated-inclusion test case. Colorbar values are dimensionless, matching the colorbar scale of Fig. 4.

IV. MULTI-SCALE SIMULATION STRATEGY

Four simulation scales are employed, each fulfilling a distinct role in the overall design and AI-training workflow:

A. Level 1: 2D Thin-Slab Parametric Sweeps

Using NZ = 3 with periodic boundary conditions in the z-direction, the three-dimensional solver produces

effectively two-dimensional results. This configuration supports parametric exploration at low computational cost: 40 scenarios complete in approximately one hour on a single GPU. The two-dimensional campaign spans 5 inlet velocities, 4 source amplification levels, and 2 lateral boundary conditions, covering the Richardson number range from buoyancy-dominated natural convection to flow-dominated forced convection. A representative Level 1 output is shown in Fig. 6 for the amplified forced-convection regime, with the temperature rise field on the left and the flow-speed field on the right. The 2D validation example presented in Section III.B uses a UNet demonstrator retrained on a canonical test; the project's production 2D substation surrogate is the separate TUNet2D champion model, which is trained on the thin-slab campaign itself and whose quantitative results are reported in [4].

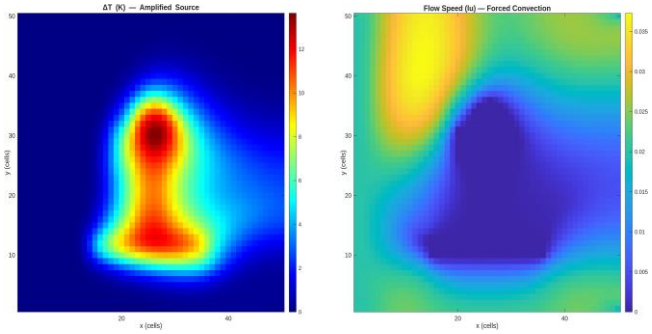

Fig. 6. Level 1 output: 2D amplified forced convection — temperature rise ΔT in K (left) and flow speed in lattice units (lu) (right).

B. Level 2: 3D Scaled Surrogate

An STL-informed analytical surrogate preserves the component layout at 80^3 resolution. Direct STL voxelization at reduced scale is infeasible because of thin-feature aliasing: the 12 cm busbars become sub-grid at coarse dx. The surrogate ensures continuity of the solid regions while maintaining the physical proportions of the substation. A 40-scenario campaign at this scale completes in approximately 30 minutes and produces the training dataset used by the three-dimensional AI model.

Thin-Feature Aliasing Analysis

The resolution-to-feature ratio is a critical constraint in direct STL voxelization. At the realistic AIS campaign resolution of dx = 3 cm, the 12 cm busbars span four cells across their diameter, at the lower edge of the ratio needed for stable thermal gradient computation at the boundary layer. At coarser resolutions, the picture deteriorates rapidly: a 5 cm voxel yields 2-3 cells across, sufficient for phase-level thermal mass but inadequate for boundary-layer heat flux resolution; an 8 cm voxel yields 1-2 cells, at which point busbars become effectively sub-grid and are lost under the union of boundary conditions applied to their neighbouring cells. Clamps, which are only 10 cm thick, disappear entirely at resolutions coarser than 5 cm. This aliasing pattern motivated the analytical surrogate approach at the 80^3 scaled resolution: direct STL voxelization at that grid size would render the 12 cm busbars as 0.6–1.2 cells, a non-resolvable configuration. The STL-informed surrogate instead reconstructs geometrically proportional components with guaranteed continuous solid regions, preserving the thermal-gradient topology of each component class regardless of absolute grid resolution, at the cost of losing exact CAD placement of individual equipment items.

Dimensionless Invariance under Geometric Scaling

The scaled surrogate campaign does not claim full dynamic similarity with the realistic AIS case at a single operating point. For the characteristic groups [10] $Ra = \beta g \cdot \Delta T \cdot L^3/(\nu \cdot \alpha)$, $Re = U \cdot L/\nu$, $Pr = \nu/\alpha$, and $Gr = Ra/Pr$, holding air properties (ν, α, β, g) and the temperature rise ΔT fixed while reducing the absolute length scale L causes Ra and Gr to scale as L^3 and is therefore not preserved between the two scales. Re can be matched by compensatory adjustment of U, and Pr is a fluid-property ratio that is exactly preserved for air at atmospheric conditions. The surrogate is therefore not a strict dimensionless replica of the realistic case.

Instead, the design purpose of the scaled campaign is different: the 40-scenario sweep (5 characteristic velocities × 4 source-strength levels × 2 lateral boundary conditions) spans a wide range of (Ra, Re) operating regimes, from buoyancy-dominated natural convection ($Ri \to \infty$) to fully forced flow ($Ri \ll 1$). The AI model trained on this sweep therefore sees the parametric structure of thermal response across the full operating envelope of interest, at approximately 1/100 the simulation cost of equivalent coverage on the realistic AIS grid. Fine-tuning on a small number of realistic AIS scenarios (cropped 128^3 subvolumes at dx = 3 cm) then adapts the learned representation to the production geometry and its specific operating point.

C. Level 3: Full-Scale Real STL

A compact real-geometry STL, derived from a single-bay substation section with external dimensions 2.4 m × 2.1 m × 2.2 m, is voxelized at dx = 3 cm, producing 845,000 cells. This compact real STL is distinct from the yard-scale geometries produced by the Agent (Table II), and is used here specifically to validate solver behaviour on real-world CAD at production grid resolution without the computational cost of a full 33-million-cell run. Three representative cases have been executed, including natural convection and forced convection with amplified source. Thermal plumes from individual busbar and conductor components are visualised and exhibit the expected qualitative physical behaviour: buoyant plume rise, conduction through clamp supports, and transformer envelope heating at production resolution. Quantitative validation against substation field measurements is identified as future work.

D. Level 4: Realistic Multi-Component AIS

The ais_demo geometry, cropped to a single bay, was voxelized at 33 million cells and simulated with QVOL_SCALE = 1000. A total of 20 scenarios spanning inlet velocity, source amplification, and lateral boundary condition variations were executed on this geometry, each run to 10,000 time steps; the hero case reported here is representative of the campaign.

A global override of THERMAL_OMEGA = 1.25 was imposed in order to maintain thermal diffusion at fine resolution. The simulation ran stably to 10,000 time steps on this geometry. The thermal field and the AI prediction results obtained on this geometry are reported in the companion paper [4].

V. COMPUTATIONAL PERFORMANCE

TABLE IV. Computational cost by simulation scale.

Scale	Grid	GPUs	Time / scenario
2D thin-slab (40 scenarios)	256 × 256 × 3	1	~1.5 min
3D scaled surrogate (40 scenarios)	80^3	2	~45 s
Full-scale real STL (3 cases)	101×89×94	4	~5 min
Realistic AIS (20 scenarios)	267×467×267	1 (32GB)	~60 min

VI. PARAMETRIC DESIGN CASE STUDY

A representative design study is presented to demonstrate the practical utility of the parametric pipeline. The study examines the effect of phase spacing on the thermal behaviour of an ais_demo-class substation. The underlying engineering question is whether the baseline phase spacing of 4.5 m, typical of HV AIS practice, can be reduced for compact urban deployment without compromising the thermal rating envelope. Five configurations are evaluated, with phase spacings of 3.0, 3.5, 4.0, 4.5, and 5.0 m, at fixed conductor sizing and loading.

In a manual CAD workflow, each configuration requires approximately two hours of modelling effort (parametric model update, consistency checking, and export to simulation-ready STL) before the thermal simulation itself can begin. Five configurations would therefore consume a full working day of CAD effort, or approximately ten hours, before computational simulation could commence. By contrast, with the Agent, the study requires modification of a single field in the YAML configuration (phase_y: [-X, 0.0, +X]) and the dispatch of five preset-modified generation requests. The complete set of geometries is then available in approximately thirty seconds. Execution of each configuration at Level 1 (two-dimensional thin-slab) consumes approximately 1.5 minutes per scenario and provides immediate insight into the relative thermal performance across the design space; the full design exploration completes in under ten minutes, including geometry generation.

The progression from Level 1 insight to engineering confidence follows the multi-scale strategy described in Section IV. The two-dimensional sweep identifies the most thermally marginal configuration, typically the tightest phase spacing, which is subsequently promoted to Level 2 (three-dimensional scaled, approximately 45 s) for three-dimensional verification; the most promising candidate is then promoted to Level 4 (realistic AIS at 33 million cells, approximately 60 min) for a production-grade thermal assessment. The end-to-end study, from engineering question to production-quality answer, completes in approximately 75 minutes of wall-clock time. Such a cycle would be prohibitive in a manual-CAD workflow and demonstrates the direct engineering value of automated parametric geometry generation.

The Level 1 sweep results for this study are summarised qualitatively as follows. At fixed conductor sizing and loading, the peak conductor temperature rise increases monotonically as phase spacing is reduced from 5.0 m to 3.0 m, with the tightest configuration exhibiting the highest centre-bay ΔT and the most confined buoyant plume topology. The 4.5 m baseline shows clear thermal margin relative to the 3.0 m and 3.5 m configurations, while the 4.0 m configuration sits near the boundary of the thermally comfortable envelope. These qualitative observations are obtained within the Level 1 sweep in under ten minutes of wall-clock time and serve as the screening input for the Level 2 and Level 4 promotion steps described above. Absolute quantitative results depend on the numerical-stabilisation settings discussed in Section VII.A and on the downstream AI surrogate evaluation reported in [4]; the Level 1 sweep is presented here as a qualitative design-screening tool rather than as a calibrated ampacity calculation.

VII. DISCUSSION

The multi-scale strategy supports efficient use of computational resources. The two-dimensional sweeps provide physical insight at low computational cost. The three-dimensional scaled surrogate generates the training-dataset volume required by the AI model. The full-scale STL simulation validates the solver on geometry derived from real engineering specifications. The realistic AIS simulation demonstrates engineering-scale capability and produces physically plausible thermal patterns at full resolution. While established LBM frameworks such as Palabos [11] are available for general-purpose fluid simulation, the present work employs an in-house GPU-resident implementation tailored to the multi-region thermal modelling and the Joule-heating source term required for substation applications.

Numerical Stabilization Choices

Two numerical overrides warrant explicit justification. First, the source amplification factor QVOL_SCALE = 1000 used in the realistic AIS campaign is a multiplier on the physical volumetric Joule heating source term applied to heater-region cells. It is chosen so as to produce a characteristic temperature rise ΔT in the 20–25 K range over the 10,000-step simulation window, which is sufficient to elevate the thermal field cleanly above the LBM noise floor at the 3 cm grid resolution. Lower values of QVOL_SCALE produce proportionally smaller ΔT while preserving the qualitative plume topology: tests at QVOL_SCALE = 100 yield ΔT in the 2–3 K range with the same spatial hot-zone locations and the same three-phase thermal ordering. The parameter is therefore a signal-to-noise tradeoff rather than a controller of the physical regime. Absolute temperature magnitudes reported in subvolume figures depend directly on this setting, whereas the relative spatial structure does not.

Second, the thermal relaxation cap THERMAL_OMEGA = 1.25 is imposed because, at dx = 3 cm with the flow-dominated time step, the physical thermal diffusivity of air maps to a lattice relaxation parameter ω_T that approaches the LBM stability limit of 2.0. A cap below 2.0 is necessary for numerical stability; the value 1.25 provides a margin while permitting meaningful thermal diffusion. The cap effectively imposes a maximum lattice thermal diffusivity and is therefore a known source of deviation between the simulated and the physical transient thermal time scales, while leaving the steady-state spatial heat-distribution pattern largely intact. A dual-time-step scheme, in which the flow and thermal advancements are decoupled and governed by their respective stability constraints, is identified as the principled long-term remedy and is planned as future work.

The limitation implied by these two choices is that the absolute temperature values and absolute thermal time constants reported for the realistic AIS case should be read as indicative rather than calibrated; they depend on both settings above and would shift under different but equally reasonable alternative choices. Qualitative behaviour, comprising the three-phase thermal ordering, the plume topology, the relative hotspot locations, and the effectiveness of the parametric pipeline for comparing design variants, is robust to both settings within their sensible operational ranges.

The parametric STL Agent supports design studies that would be computationally prohibitive under a manual CAD workflow. Variations in phase spacing, bay configuration, conductor sizing, and transformer placement require only modification of the YAML specification, rather than reconstruction of the underlying CAD model. This capability is of particular value for thermal-clearance assessment during substation design, in which tens of configurations may require evaluation within the available engineering timeframe.

VIII. CONCLUSIONS

An integrated platform for automated parametric geometry generation and multi-scale GPU-accelerated thermal simulation of air-insulated substations has been presented. The KYAMOS STL Agent produces simulation-ready geometry from engineering specifications within a few seconds per configuration. The multi-scale simulation strategy provides four complementary levels of fidelity: two-dimensional thin-slab sweeps for parametric exploration, three-dimensional scaled surrogate campaigns for AI training data generation, full-scale STL simulation for production validation, and realistic multi-component AIS simulation at 33 million cells for engineering demonstration. The platform establishes a direct pathway from engineering specification to physics-based simulation and supports automated thermal design assessment without the requirement of manual CAD modelling or mesh generation.

ACKNOWLEDGMENT

This work is part of the **REALISATION-THERMAL-AI** Project (Grant number: COM-CONCEPT-ENERGY/0624/0163), which is funded by the EU Recovery and Resilience Facility of the European Union - NextGenerationEU, and the Republic of Cyprus through the Research and Innovation Foundation within the framework of the «RESTART 2016-2020» Programmes for Research, under the Component 6.1 «REPowerEU» of the Cyprus Recovery and Resilience Plan. Computations were performed on the KYAMOS V100 InfiniBand cluster.

REFERENCES

[1] S. Succi, The Lattice Boltzmann Equation for Fluid Dynamics and Beyond. Oxford Univ. Press, 2001.
[2] P. Lallemand and L.-S. Luo, "Theory of the lattice Boltzmann method: dispersion, dissipation, isotropy, Galilean invariance, and stability," Phys. Rev. E, vol. 61, no. 6, pp. 6546–6562, 2000.
[3] NVIDIA Corporation, "CUDA C++ Programming Guide," docs.nvidia.com/cuda, 2024.
[4] A. Papadakis, S. Pigiotis, F. Zachariadis, E. Constantinide, "GPU-Accelerated Lattice Boltzmann Thermal Simulation of Electrical Substations with AI-Based Real-Time Prediction," JMEST, 2026 (companion paper).
[5] CadQuery Development Team, "CadQuery: A Python parametric CAD scripting framework," github.com/CadQuery/cadquery, 2024.
[6] S. Ramírez, "FastAPI: A modern, fast web framework for building APIs," fastapi.tiangolo.com, 2024.
[7] IEEE Std 605-2023, IEEE Guide for Bus Design in Air Insulated Substations, Institute of Electrical and Electronics Engineers, 2023.
[8] Z. Guo, C. Zheng, and B. Shi, "Discrete lattice effects on the forcing term in the lattice Boltzmann method," Phys. Rev. E, vol. 65, p. 046308, 2002.
[9] G. de Vahl Davis, "Natural convection of air in a square cavity: a bench mark numerical solution," Int.

J. Numer. Meth. Fluids, vol. 3, no. 3, pp. 249–264, 1983.
[10] A. Bejan, Convection Heat Transfer, 4th ed. Hoboken, NJ, USA: Wiley, 2013.
[11] J. Latt, O. Malaspinas, D. Kontaxakis, A. Parmigiani, D. Lagrava, F. Brogi, M. B. Belgacem, Y. Thorimbert, S. Leclaire, S. Li, F. Marson, J. Lemus, C. Kotsalos, R. Conradin, C. Coreixas, R. Petkantchin, F. Raynaud, J. Beny, and B. Chopard, "Palabos: Parallel Lattice Boltzmann Solver," Comput. Math. Appl., vol. 81, pp. 334–350, 2021.