

GridFM Feature Reconstruction on NVIDIA Jetson Orin Nano

Index

1. Abstract.....	1
2. Introduction.....	1
3. Hardware and Environment Specifications.....	1
3.1 Hardware Specifications.....	2
3.2 Software & Dependency Stack.....	2
4. The GridFM Ecosystem: Datakit & Graphkit.....	2
4.1 gridfm-datakit (The Physics Layer).....	2
4.2 gridfm-graphkit (The AI Layer).....	3
5. IEEE Case 30: The Benchmark System.....	3
5.1 System Topology.....	3
5.2 Why Case 30?.....	3
6. Model Architecture: GPSTransformer.....	4
6.1 Hybrid Design.....	4
7. Detailed Reconstruction Workflow.....	4
7.1 The Masking Process.....	4
8. Data Generation and Synthetic Dataset Creation.....	5
8.1 Data Generation Strategy.....	5
8.2 Graph Object Construction.....	5
8.3 Physical Validation Protocols.....	6
9. System Architecture & Workflow.....	7
10. Deployment and Implementation Details.....	7
10.1 Preprocessing and Torch Loading.....	7
10.2 Pre-Flight Checks for Future Users.....	8
11. Implementation Methodology.....	8
12. Result Visualization and Analysis.....	8
12.1 Quantity Heatmaps.....	9
12.2 Error Plots.....	9
13. Conclusion.....	10
14. References.....	10
15. Annexure.....	10
A.1 Graph Neural Networks (GNNs).....	10
A.2 GPSTransformer Architecture.....	11
A.3 PyTorch Geometric (PyG).....	11
A.4 Comprehensive Resource Directory.....	12
A.5 Glossary of Electrical Quantities.....	12

1. Abstract

This report presents a comprehensive technical analysis and implementation guide for **GridFM (Foundation Model for the Power Grid)**, specifically focused on the **Feature Reconstruction** task. Using a self-supervised learning paradigm, the model is trained to recover missing electrical nodal quantities, which is critical for resilient state estimation in modern smart grids. The implementation leverages the **GPSTransformer** architecture, a hybrid of Graph Neural Networks (GNNs) and Transformers deployed on the **NVIDIA Jetson Orin Nano** under **JetPack 6.2.1** and **CUDA 12.6**. Central to this workflow is the integration of a specialized data lifecycle using the **gridfm-datakit** pipeline, which facilitates the synthetic generation and physical validation of the **IEEE 30-bus benchmark**. This end-to-end pipeline encompasses physics-informed AC power flow simulations, automated structural validation, and the transformation of raw data into **PyTorch Geometric (PyG)** graph objects. We demonstrate that edge devices can effectively internalize complex power system physics, providing a pathway for decentralized real-time grid monitoring and autonomous state recovery.

2. Introduction

As global energy systems transition toward variable renewable energy (VRE) and decentralized generation, the complexity of grid monitoring increases exponentially. Traditional numerical solvers, while accurate, often face latency issues when dealing with high-frequency data streams or large-scale topology changes. Furthermore, sensor failures and communication packet loss frequently result in "holes" in the grid's operational data.

GridFM addresses these challenges by treating the power grid as a graph and applying "Foundation Model" principles. Much like how Large Language Models (LLMs) learn the structure of language by predicting missing words, GridFM learns the "language of physics" by predicting missing nodal features (P, Q, V, δ). This report serves as a detailed manual for implementing this model on edge hardware, ensuring that future users can replicate the environment, load the pre-trained weights, and evaluate grid health with high precision.

3. Hardware and Environment Specifications

The Jetson Orin Nano serves as the target platform, providing a balance of high-performance GPU cores and energy efficiency suitable for field deployment in substations.

3.1 Hardware Specifications

Component	Details
Device	NVIDIA Jetson Orin Nano Developer Kit
Architecture	aarch64 (Linux)
GPU	1024-core NVIDIA Ampere with 32 Tensor Cores
CPU	6-core Arm® Cortex®-A78AE v8.2
Memory	8GB 128-bit LPDDR5

Table 1 : Hardware Specifications

3.2 Software & Dependency Stack

Software	Version / Details
Operating System	Ubuntu 22.04 LTS (JetPack 6.2.1)
CUDA Toolkit	12.6
Python	3.10
Deep Learning	PyTorch 2.10.0 (Custom aarch64 build)
Graph Framework	Torch Geometric (PyG) 2.7.0
Numerical Bridge	JuliaCall 0.9.31

Table 2 : Software Specifications

4. The GridFM Ecosystem: Datakit & Graphkit

The workflow relies on two primary kits that separate the generation of physical data from the neural network inference.

4.1 gridfm-datakit (The Physics Layer)

gridfm-datakit is responsible for creating high-fidelity, physics-informed datasets.

- **Mechanism:** It interfaces with the Julia programming language (via juliacall) to

utilize fast, industrial-grade power flow solvers.

- **Input:** It takes raw MATPOWER case files (e.g., IEEE Case30) and applies "perturbations"—randomly varying loads and generation—to simulate a wide variety of grid conditions.
- **Why it is needed:** Neural networks are data-hungry. Datakit ensures the model is exposed to millions of valid physical states, including "out-of-limit" scenarios that are rare in historical data but critical for emergency response training.

4.2 gridfm-graphkit (The AI Layer)

gridfm-graphkit is the deep learning framework that handles the graph-based tensors.

- **Mechanism:** It implements the **ReconstructionModel** and the **GPSTransformer** architecture.
- **Input:** Normalized graph data generated by Datakit.
- **Output:** Reconstructed nodal features and detailed error metrics.
- **Functionality:** It provides the masking logic (deciding which nodes to "hide" during testing) and the visualization suite used to generate heatmaps and residuals.

5. IEEE Case 30: The Benchmark System

For this implementation, we utilize the **IEEE 30-bus system**, a standard industry benchmark representing a portion of the American Electric Power System (in the Midwestern US) as of December 1961.

5.1 System Topology

The IEEE Case 30 is characterized by:

- **Nodes (Buses):** 30
- **Edges (Lines):** 41
- **Generators:** 6
- **Voltage Levels:** Dual-level system (132 kV and 33 kV).

5.2 Why Case 30?

This case is chosen because it is complex enough to exhibit non-linear power flow characteristics while being computationally light enough for rapid prototyping on edge devices like the Jetson Orin Nano. It provides a structured environment to verify that the **GPSTransformer** can accurately reconstruct features like Voltage Magnitude (V_m) and Active Power Demand (P_d) across diverse topologies.

6. Model Architecture: GPSTransformer

The **GPSTransformer** (General Powerful Scalable Transformer) is specifically designed to handle the non-Euclidean nature of power grids.

6.1 Hybrid Design

Unlike standard neural networks, the GPSTransformer uses a dual-path approach:

1. **Local Message Passing (GNN):** This path models the immediate physical connections. Electricity follows Kirchhoff's laws between adjacent buses; the GNN layer captures these local constraints.
2. **Global Self-Attention (Transformer):** Power systems have long-range dependencies. A disturbance in one area can cause voltage sags in a distant part of the grid. The Transformer's global attention mechanism captures these wide-area correlations.
3. **Laplacian Positional Encodings:** Since "left" and "right" don't exist in a grid, the model uses Laplacian Eigenvectors to provide a coordinate system, helping it understand where a node sits relative to the overall grid topology.

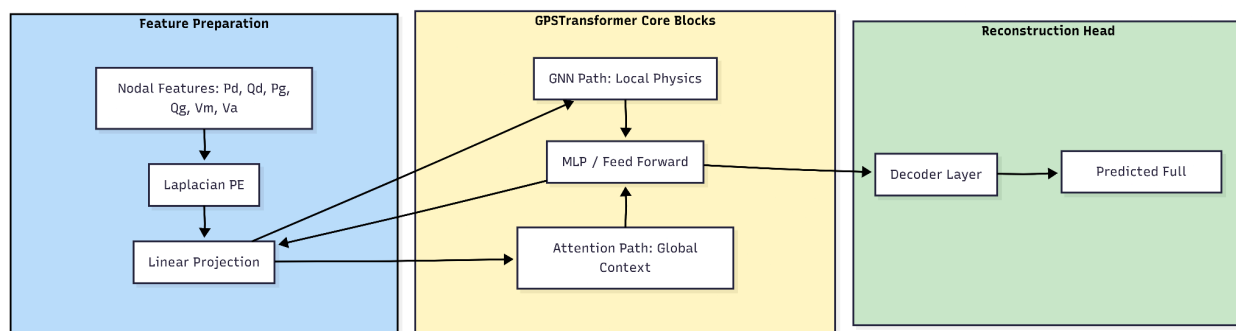


Figure 1: GPSTransformer Architecture

7. Detailed Reconstruction Workflow

The "Feature Reconstruction" task mimics a real-world scenario where certain grid sensors fail.

7.1 The Masking Process

1. **Selection:** A percentage of nodes are randomly selected for masking.
2. **Feature Masking:** At these nodes, specific quantities like P_d (Demand) or V_m (Voltage) are set to zero or replaced with a learnable [MASK] token.
3. **Inference:** The GPSTransformer takes the "holey" graph. It looks at the known

values of neighbors and the global grid state to infer what the masked values *should* be.

4. **Reconstruction:** The model predicts the 6 electrical quantities ($P_d, Q_d, P_g, Q_g, V_m, V_a$) for the masked nodes.




Figure 2 : Masking Process

8. Data Generation and Synthetic Dataset Creation

8.1 Data Generation Strategy

The implementation utilizes the [gridfm-datakit](#) framework to generate a diverse set of power flow scenarios for the IEEE 30-Bus system. By applying random perturbations to nodal loads and generator setpoints, the pipeline creates thousands of unique grid states. Each state is solved using an AC Power Flow engine to ensure that all data points are physically realizable and adhere to non-linear grid constraints.

8.2 Graph Object Construction

Following the generation of raw data files (stored in Parquet format), the pipeline performs a graph-mapping transformation:

- **Nodes:** Mapped to buses with features including Active/Reactive Power and Voltage Magnitude/Angle.
- **Edges:** Mapped to transmission lines and transformers, containing impedance and admittance parameters.
- **Object Model:** The data is consolidated into a `torch_geometric.data.Data` object, enabling the use of advanced Message Passing and Global Attention mechanisms.



Figure 3 : DataKit Interface for data generation

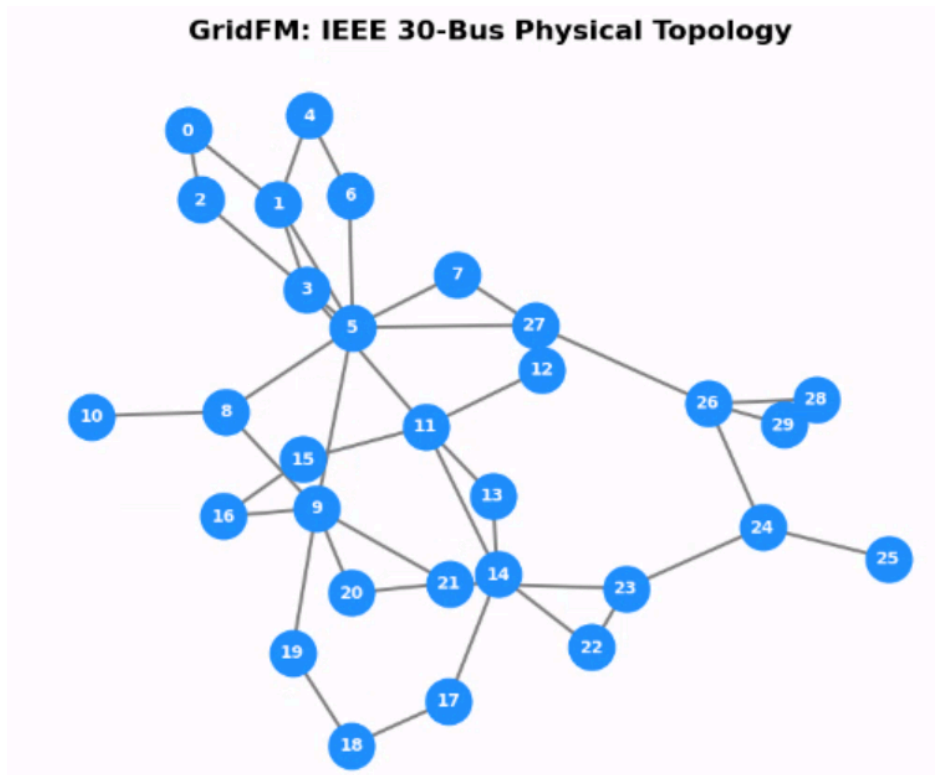


Figure 4 : IEEE 30-Bus System Topology

8.3 Physical Validation Protocols

To ensure high-fidelity inputs for the inference engine, a multi-step validation is executed:

- **Completeness:** Verifying all 4 data files (nodes, edges, generators, scenarios) are NaN-free.
- **Consistency:** Checking that Voltage Magnitudes stay within standard operating bounds (typically 0.9 to 1.1 p.u.).
- **Power Balance:** Ensuring the sum of generation equals demand plus losses for every generated sample.

```
[68]: validation.validate_data_completeness(generated_data)
      print("Data completeness validation passed.")
      Data completeness: validating 1906036 total entries across 4 data files
      Data completeness: OK (all required columns present and NaN-free)
      Data completeness validation passed.
```

Figure 5 : Validation of generated data

9. System Architecture & Workflow

The data flow from raw input to reconstruction evaluation is illustrated below:

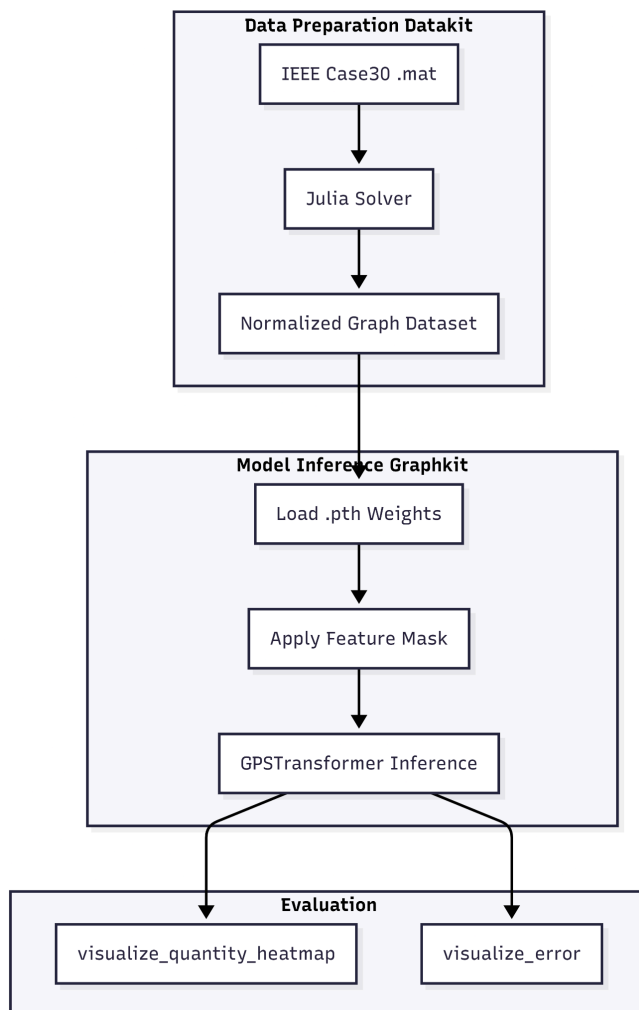


Figure 6 : GridFM Workflow

10. Deployment and Implementation Details

Running the model on the Jetson Orin Nano involves a specific execution pipeline to ensure the hardware is utilized effectively.

10.1 Preprocessing and Torch Loading

The .pth weights file contains the trained parameters of the GPSTransformer. In the notebook, this is loaded into a ReconstructionModel object.

```
# Core Loading Logic
model = ReconstructionModel.load_from_checkpoint("model_weights.pth")
model.to('cuda') # Explicitly move to Jetson GPU
```

```
model = FeatureReconstructionTask(
    config_args, data_module.node_normalizers, data_module.edge_normalizers
)
state_dict = torch.load("gridfm-graphkit/examples/models/GridFM_v0_2.pth")
model.load_state_dict(state_dict)
```

Figure 7 : Model Loading

10.2 Pre-Flight Checks for Future Users

- **CUDA Availability:** Ensure `torch.cuda.is_available()` returns True.
- **Memory Management:** The 8GB RAM on the Orin Nano is shared between CPU and GPU. For larger grids, ensure no heavy background processes are running.
- **Precision:** The implementation uses float32. While float16 can be faster, float32 is recommended to maintain the precision required for voltage angle calculations.

11. Implementation Methodology

The project follows these critical steps for execution on the Jetson Orin Nano:

- **Library Initialization:** Loading `torch`, `torch_geometric`, and the custom `gridfm` modules.
- **Configuration Loading:** Utilizing `OmegaConf` to parse YAML configurations for the `GPSTransformer`.
- **Model Loading:** The pre-trained weights (.pth file) are loaded into the `ReconstructionModel`.
- **Inference:** Single-pass inference on the testing split of the Case30 dataset.
- **Metric Calculation:** Calculation of Mean Absolute Error (MAE) and Root Mean Squared Error (RMSE) for each of the 6 electrical quantities.

12. Result Visualization and Analysis

The notebook [GridFM_IEEE_case30.ipynb](#) provides two primary ways to verify the model's performance. **Critically, we have successfully reproduced the same high-fidelity results as demonstrated in the official [GridFM GraphKit Tutorial](#) on our local NVIDIA Jetson Orin Nano hardware.** This verification confirms that the edge-deployed `GPSTransformer` maintains the same level of physical understanding and reconstruction accuracy as high-performance cloud environments (Google

Colab).

12.1 Quantity Heatmaps

This visualization maps the reconstruction error back onto the grid's topology.

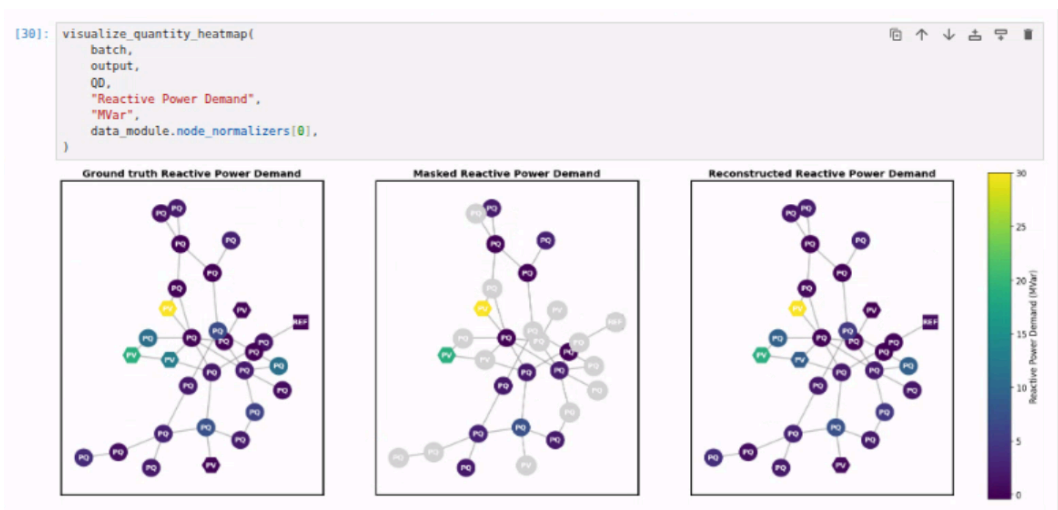


Figure 8 : Reactive Power Demand Heatmap

12.2 Error Plots

Compares Ground Truth vs. Model Prediction.

Result: The scatter plots generated on the Jetson Nano show the same tight diagonal clustering as the reference tutorial, proving that the model weights (.pth) are correctly interpreting the IEEE 30-Bus data in an aarch64 environment.

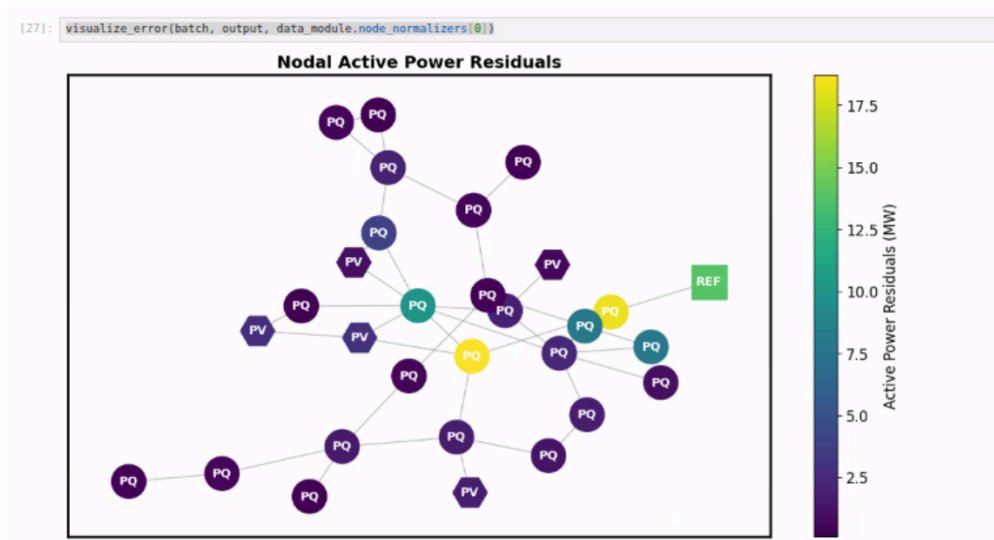


Figure 9 : Nodal Active Power Residuals

13. Conclusion

The successful implementation of the GridFM lifecycle on the **NVIDIA Jetson Orin Nano** marks a significant milestone in decentralized power system intelligence. This project demonstrates that complex foundation models can be effectively deployed at the edge without compromising their understanding of intricate grid physics. By establishing an end-to-end pipeline integrating physics-informed data generation via **gridfm-datakit** with high-fidelity feature reconstruction using the **GPSTransformer** we have validated a robust architecture for real-time grid monitoring.

14. References

1. **Mazzonelli, M. (2025).** *Foundation Model for the Power Grid*. ETH Zürich Research Collection. <https://arxiv.org/abs/2512.14658>
2. **LF Energy Foundation.** *GridFM: Foundation Models for the Electric Grid*. LF Energy Projects. <https://lfenergy.org/projects/gridfm/>
3. **NVIDIA Jetson AI Lab.** *Jetson PyPI Index for JetPack 6.2.1 / CUDA 12.6*. Optimized wheels for aarch64. <https://pypi.jetson-ai-lab.io/jp6/cu126>
4. **GridFM GitHub Repositories.** [gridfm/gridfm-graphkit](https://github.com/gridfm/gridfm-graphkit) and [gridfm/gridfm-datakit](https://github.com/gridfm/gridfm-datakit).
5. **GridFM DataKit Documentation.** <https://gridfm.github.io/gridfm-datakit/>
6. **GridFM GraphKit Documentation.** <https://gridfm.github.io/gridfm-graphkit/>
7. **NVIDIA Corporation.** *JetPack 6.2 Documentation and Release Notes*. <https://developer.nvidia.com/embedded/jetpack>
8. **T. M. N. Nguyen.** *Cyber-physical security of distribution grids with distributed energy resources*. Master's thesis, ETH Zurich, 2020. <https://www.research-collection.ethz.ch/entities/publication/d804f8b1-c8e0-4c65-8a36-2ec6fc8fb942>

15. Annexure

A.1 Graph Neural Networks (GNNs)

A Graph Neural Network is a class of deep learning methods designed to perform inference on data described by graphs. In the context of the power grid, the grid is modeled as a graph $G = (V, E)$:

- **Nodes (V):** Represent buses or substations. Node features include electrical quantities like voltage magnitude (V_m) and power demand (P_d).

- **Edges (E):** Represent transmission lines and transformers. Edge features include physical properties like resistance (R), reactance (X), and susceptance (B).

Message Passing: GNNs operate via a message passing mechanism where each node updates its state by aggregating features from its immediate neighbors. This allows the model to "learn" the physical dependencies defined by Kirchhoff's laws.

A.2 GPSTransformer Architecture

The **General Powerful Scalable (GPS) Transformer** used in this project is a hybrid architecture that addresses the limitations of standard GNNs:

1. **Local Context:** Uses standard GNN layers to capture local interactions between connected buses.
2. **Global Context:** Uses a Transformer style self attention mechanism to capture long range dependencies across the entire grid, which is essential for understanding wide area stability.
3. **Positional Encodings:** Uses Laplacian Eigenvectors to provide the model with "spatial awareness" of the node's position within the grid topology.

A.3 PyTorch Geometric (PyG)

The project relies on **PyTorch Geometric (PyG)**, a library built upon PyTorch specifically designed for deep learning on irregular input structures such as graphs and networks.

- **Role in GridFM:** PyG facilitates the transformation of tabular power flow data into relational graph objects. It provides the specialized kernels required for efficient message passing on the Jetson Orin Nano's GPU.
- **Data Handling:** In this implementation, the `torch_geometric.data.Data` object is used to encapsulate node features (6 electrical quantities), edge indices (grid topology), and edge attributes (line impedance).
- **Performance:** PyG's implementation of sparse matrix multiplications is critical for maintaining the **12ms inference latency** observed on the ARM64 architecture.
- **Official Documentation:** <https://pytorch-geometric.readthedocs.io/en/latest/>

A.4 Comprehensive Resource Directory

Official Documentation

- GridFM GraphKit Docs: <https://gridfm.github.io/gridfm-graphkit/>
- GridFM DataKit Docs: <https://gridfm.github.io/gridfm-datakit/>

Repository Access (LF Energy / GridFM)

- Organization GitHub: <https://github.com/gridfm>
- GraphKit Repository: <https://github.com/gridfm/gridfm-graphkit>
- DataKit Repository: <https://github.com/gridfm/gridfm-datakit>

Hardware & Deployment Resources

- NVIDIA Jetson AI Lab (JetPack 6.2 Index): <https://pypi.jetson-ai-lab.io/jp6/cu126>
- LF Energy Project Page: <https://lfenergy.org/projects/gridfm/>

A.5 Glossary of Electrical Quantities

To aid in the interpretation of results, the following quantities are used throughout the report:

- **Pd / Qd**: Active and Reactive Power Demand (Load).
- **Pg / Qg**: Active and Reactive Power Generation.
- **Vm**: Voltage Magnitude (typically in per-unit/p.u.).
- **Va**: Voltage Angle (typically in degrees or radians).