

PINNSim: A Simulator for Power System Dynamics based on Physics-Informed Neural Networks

Jochen Stiasny, Spyros Chatzivasileiadis
Division for Power and Energy Systems
Technical University of Denmark
Kgs. Lyngby, Denmark
{jbest, spchatz}@dtu.dk

Baosen Zhang
Electrical and Computer Engineering
University of Washington
Seattle, WA 98195
zhangbao@uw.edu

Abstract—The dynamic behaviour of a power system can be described by a system of differential-algebraic equations. Time-domain simulations are used to simulate the evolution of these dynamics. They often require the use of small time step sizes and therefore become computationally expensive. To accelerate these simulations, we propose a simulator – PINNSim – that allows to take significantly larger time steps. It is based on Physics-Informed Neural Networks (PINNs) for the solution of the dynamics of single components in the power system. To resolve their interaction we employ a scalable root-finding algorithm. We demonstrate PINNSim on a 9-bus system and show the increased time step size compared to a trapezoidal integration rule. We discuss key characteristics of PINNSim and important steps for developing PINNSim into a fully fledged simulator. As such, it could offer the opportunity for significantly increasing time step sizes and thereby accelerating time-domain simulations.

Index Terms—Dynamical systems, Differential-algebraic equations, Physics-Informed Neural Networks, Time-domain simulation

I. INTRODUCTION

Differential equations form a centre piece in the modelling of the dynamic behaviour of power systems. They provide very widely applicable component and system models, however, their solution requires numerical integration methods. These tools for Time-Domain Simulations (TDSs) constitute a workhorse of power system analysis and can reliably deliver accurate solutions. The used algorithms are scalable, versatile, and applicable without case-specific modifications, but they incur a high computational burden [1]. Efforts around improved solvers and parallelisation can reduce this burden, see [2], [3], [4] for an overview. However, a fundamental problem is the small time step size that is required by the schemes to ensure accurate results and numerical stability. Hence many time steps and function evaluations are necessary.

The proposition of Neural Network (NN)-based approaches [5], and in particular Physics-Informed Neural Networks (PINNs) [6], has very different characteristics. In the training process, PINNs learn the solution to the provided differential equations and subsequently provide fast and sufficiently

accurate solutions over long time intervals – addressing two main issues of conventional numerical integration methods. The authors in [7] refer to these NNs for simulating dynamical systems as a direct-solution model.

As such, PINNs have been adopted to power system dynamics in [8] to predict single machine dynamics. In a similar setup, the related methodology of operator learning is applied [9]. Other learning-based methods for simulating dynamics – see [7] for an overview – have also been explored in the context of single components, e.g. in [10]. It is an open question if the promising results for single components can also be achieved when predicting the dynamics of multi-machine systems. For PINNs, [11], [12] explored this question for a 9, 11, and 39-bus system. The authors in [13] predict power system transients with an alternative architecture which relies on a transformation into the frequency domain. The effect of generator outages is predicted in [14] and [15] estimates system dynamics after line faults.

In the above cases, the speed advantage of learning-based methods persists and the accuracy remains sufficiently high for the trained setup. However, all approaches require assumptions about the operating conditions, such as the network topology and the set-points, as well as the disturbance type and location during the dataset generation and the training of the model. As a result these learning-based methods cannot be reliably applied to unseen conditions or faults. If one attempted to include all possible conditions and faults in the learning stage, the resulting problem will not be tractable as the required dataset becomes too high-dimensional [16]. Hence, attempts to scale NN-based TDS approaches need to focus on limiting the assumptions that are required for the learning task such as operating conditions and disturbance characteristics.

To that end, we propose the simulator *PINNSim*. We first learn the solutions of all dynamic components in the system separately, i.e., with entirely independent PINNs. Thereby, the learning tasks are of relatively low dimensionality and hence remain tractable. Meanwhile, the use of PINNs allows us to benefit from fast solutions for large time steps. To then perform a TDS of the multi-machine system, we need to determine the interactions between the components. We achieve this by applying a root-finding algorithm similar to conventional TDS algorithms; at this stage, no learning is required. As a result,

This work was supported by the ERC Project VeriPhIED, funded by the European Research Council, Grant Agreement No: 949899.

the scalability and flexibility of PINNSim become decoupled from the difficulties of learning high-dimensional problems. At the same time, we still can benefit from fast and accurate solutions for larger time step sizes due to the use of PINNs.

In Section II, we present the conceptual motivation behind PINNSim. In Section III, we introduce the power system formulation and the methodology of PINNSim. Section IV describes the setup of the numerical experiments to demonstrate a proof of concept and in Section V, we discuss the results by highlighting key characteristics of PINNSim. Section VI discusses the steps for developing PINNSim from a proof of concept to a fully fledged simulator. Section VII concludes.

II. CONCEPT

To facilitate the presentation of the methodology in Section III, we want to first describe the conceptual idea behind PINNSim. It originates from the problem that we face when solving Differential-Algebraic Equations (DAEs), a form of differential equations where a set of algebraic equations constrains the differential equations. It formulates as

$$\frac{d}{dt}\mathbf{x} = \mathbf{f}(\mathbf{x}, \mathbf{y}) \quad (1a)$$

$$\mathbf{0} = \mathbf{g}(\mathbf{x}, \mathbf{y}) \quad (1b)$$

and we refer to $\mathbf{x}(t)$ as the differential variables and to $\mathbf{y}(t)$ as the algebraic variables. The update function $\mathbf{f}(\mathbf{x}, \mathbf{y})$ and the algebraic relationship $\mathbf{g}(\mathbf{x}, \mathbf{y})$ govern the dynamics of the system¹. Our interest focuses on the particular form of index-1 DAEs or semi-analytical DAEs [17]. This form implies that \mathbf{g} can be differentiated once with respect to time t which is possible when $\frac{\partial \mathbf{g}}{\partial \mathbf{y}}$ is non-singular. Based on (1), we could describe the temporal evolution of \mathbf{x} and \mathbf{y} as

$$\mathbf{x}(t) = \mathbf{x}_0 + \int_{t_0}^t \mathbf{f}(\mathbf{x}, \mathbf{y}) d\tau \quad (2a)$$

$$\mathbf{y}(t) = \mathbf{g}'(\mathbf{x}) \quad (2b)$$

where $\mathbf{x}_0 = \mathbf{x}(t_0)$ represents the initial condition and $\mathbf{g}'(\mathbf{x})$ describes the solution to the algebraic equations given \mathbf{x} . However, usually no analytical expression describes the evolution $\mathbf{x}(t)$ and $\mathbf{y}(t)$ for a given \mathbf{x}_0 . Hence, we revert to numerical integration methods to obtain an approximate solution.

To resolve the non-trivial integration operation and the implicit relationship \mathbf{g} , numerical schemes often restrict the functional form of \mathbf{x} to a certain approximation $\hat{\mathbf{x}}$. For instance, Runge-Kutta (RK) methods assume a polynomial form of $\hat{\mathbf{x}}(t) = \mathbf{x}_0 + a_1(t - t_0) + a_2(t - t_0)^2 + \dots$ as they match the Taylor expansion up to a certain degree by construction. The different RK schemes prescribe the order of the scheme and the computation of the coefficients. When algebraic variables are present, they have to be interfaced with the approximation of the differential variable to incorporate their interaction with each other. Simultaneous and partitioned integration methods describe such routines [1] and are used

¹For notational clarity, we formulate an autonomous, unforced system. The conceptual idea can also accommodate non-autonomous and forced systems.

in many variations. The accuracy of these constructions is dependent on the order of the used integration scheme and potential ‘‘interface’’ errors. These considerations and aspects of numerical stability limit the usable time step size.

By choosing a different functional form for $\hat{\mathbf{x}}$, like Fourier-series based or around the Adomian decomposition [18], the practical time step size might be increased. Due to their high flexibility of their functional form, PINNs, and NNs in general as suggested in [5], can allow significantly larger time steps. In fact, we can even choose a functional form of $\hat{\mathbf{y}}(t)$ and approximate $\hat{\mathbf{x}} = \text{PINN}(t, \hat{\mathbf{y}})$ in dependence of it. However, this great approximation flexibility of PINNs comes with the challenge of generalising well across the entire domain of interest, i.e., being accurate over the entire domain and not only on the training dataset. When the dimensionality of \mathbf{x} and \mathbf{y} in (1) increases, the training of a PINN to approximate a wide range of solutions becomes increasingly difficult and eventually intractable. With PINNSim, we avoid this problem by exploiting the structure of the power system specific DAEs. The structure allows a decomposition of (1) into multiple smaller sub-problems which remain tractable from a learning perspective. At the same time, the use of PINNs allows for large and accurate time steps of these sub-problems. To simulate the entire system, we need to align the sub-problems’ solutions by enforcing the algebraic relationship $\mathbf{g}(\mathbf{x}, \mathbf{y})$.

III. METHODOLOGY

This section describes the problem formulation and its decomposition for PINNSim in Section III-A. We then describe the main elements of the algorithm in Sections III-B to III-D and lastly the entire algorithm in Section III-E.

A. Problem formulation and solution approach

The key relation for power system dynamics is described by the current balance, i.e., Kirchhoff’s current law, and it needs to hold at all times t and for each of the n buses in the network. We distinguish between the complex current injections stemming from the network $\bar{\mathbf{i}}^N \in \mathbb{C}^n$ and from connected components $\bar{\mathbf{i}}^C \in \mathbb{C}^n$. The former can be described by the algebraic relationship (neglecting electro-magnetic transients)

$$\bar{\mathbf{i}}^N = \bar{\mathbf{Y}}\bar{\mathbf{v}}. \quad (3)$$

where $\bar{\mathbf{v}} \in \mathbb{C}^n$ represents the complex voltages at the buses and $\bar{\mathbf{Y}} \in \mathbb{C}^{n \times n}$ the complex admittance matrix. The component current injections \bar{i}_i^C at bus i can stem from a *dynamic component*, i.e., their behaviour is governed by differential equations,

$$\bar{i}_i^C = h_i(\mathbf{x}_i, \bar{v}_i) \quad (4a)$$

$$\frac{d}{dt}\mathbf{x}_i = \mathbf{f}_i(\mathbf{x}_i, \bar{v}_i, \mathbf{u}_i) \quad (4b)$$

where the current injection depends on the state vector $\mathbf{x}_i \in \mathbb{R}^p$ and the voltage \bar{v}_i of the bus i to which the component is

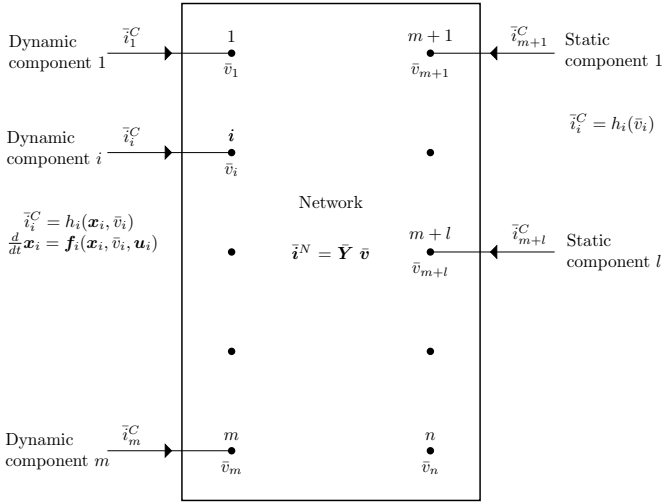


Fig. 1. Structure of the DAEs that govern the power system dynamics with current injections \bar{i}^C from the components, e.g., generators and loads, and current flows in the network \bar{i}^N . Adapted from [19].

connected². The update function f_i can furthermore depend on control inputs u_i . For *static components* the current injections becomes a function of the local voltage $\bar{i}_i^C = h_i(\bar{v}_i)$. If no component is connected to bus i , then $\bar{i}_i^C = 0$. If multiple components, indexed by k , were connected to the same bus, their current injections $\bar{i}_i^C = \sum_k \bar{i}_{i,k}^C$ are summed up and the resulting current injection then depends on the states of the connected dynamic components $x_{i,k}$ and the local voltage \bar{v}_i .

Figure 1 schematically depicts this structure with m dynamic components, l static component, and n buses. When we formulate the current balance

$$\mathbf{0} = \bar{i}^C(\mathbf{x}(t), \bar{\mathbf{v}}(t)) - \bar{i}^N(\bar{\mathbf{v}}(t)), \quad (5)$$

and consider the system's structure, we note that the component currents \bar{i}^C have only a *local* dependency, i.e., to their corresponding state vectors x_i and the local voltage \bar{v}_i . Their interaction arises from the coupling through the network currents \bar{i}^N , which in turn only depends on all complex voltages $\bar{\mathbf{v}}$.

PINNSim exploits this structure as follows: We assume a voltage evolution $\hat{\bar{\mathbf{v}}}(t)$ – we denote all approximations with the hat symbol $\hat{\cdot}$ – based on which we evaluate the resulting network currents $\hat{\bar{i}}^N$ and the component currents $\hat{\bar{i}}^C$. For the latter, we furthermore require the state evolution $\hat{\mathbf{x}}_i$ of each component; we approximate the necessary temporal integration with PINNs. When comparing the resulting currents

$$\hat{\bar{i}}^C(\hat{\mathbf{x}}(t), \hat{\bar{\mathbf{v}}}(t)) - \hat{\bar{i}}^N(\hat{\bar{\mathbf{v}}}(t)) \quad (6)$$

the result will not equal $\mathbf{0}$ by default, i.e., the current balance (5) is not obeyed. However, the mismatch indicates the quality

²For notational ease, we assume that component i is connected to bus i . If component i was connected to bus j , \bar{v}_i would be replaced by \bar{v}_j . If a component is connected to multiple buses, all of the corresponding voltages will be included in $h_i(\cdot)$ and $f_i(\cdot)$ in (4a) and (4b).

of the approximation. By adjusting the assumed voltage evolution $\hat{\bar{\mathbf{v}}}(t)$, we iteratively reduce this mismatch. The resulting voltage $\hat{\bar{\mathbf{v}}}(t)$ and state $\hat{\mathbf{x}}(t)$ evolutions constitute the solution of PINNSim to the system of DAEs.

B. Parametrisation of the voltage evolution $\hat{\bar{v}}_i(t)$

We describe the evolution of the complex voltage in polar form $\bar{v}_i(t) = V_i(t)e^{j\theta_i(t)}$. To approximate $\bar{v}_i(t)$, we express the voltage magnitude V_i and the voltage angle θ_i as a power-series with respect to time up to power r as suggested in [20]

$$\hat{\bar{v}}_i(t) = \left(\sum_{k=0}^r V_{k,i}(t-t_0)^k \right) e^{j(\sum_{k=0}^r \theta_{k,i}(t-t_0)^k)}. \quad (7)$$

The coefficients $V_{0,i}, V_{1,i}, \dots, V_{r,i}$ and $\theta_{0,i}, \theta_{1,i}, \dots, \theta_{r,i}$ form the parameters which we will later on update to improve the approximation. The vector Ξ_i collects all parameters at bus i

$$\Xi_i = [V_{0,i}, \theta_{0,i}, \dots, V_{r,i}, \theta_{r,i}], \quad \Xi_i \in \mathbb{R}^{2(r+1)}. \quad (8)$$

This parametrisation is repeated for all n buses in the system. We concatenate all Ξ_i in the vector $\Xi \in \mathbb{R}^{2(r+1)n}$.

C. Solving component dynamics with PINNs

The exact solution for the evolution of the differential variables $x_i(t)$ of a single dynamic component i can be obtained by integration of (4b)

$$x_i(t) = x_{0,i} + \int_{t_0}^t f_i(x_i(\tau), \bar{v}_i(\tau), u_i) d\tau. \quad (9)$$

As there usually exists no explicit analytical solution to (9), we need to approximate the solution. As motivated earlier, using PINNs allows us to accurately approximate (9) over large time steps $\Delta t = t - t_0$ while being fast to evaluate. The approximation depends on the time step size Δt , the initial state $x_{0,i}$, potentially the control input u_i , and the voltage profile \bar{v}_i . As we cannot train the PINN with an arbitrary \bar{v}_i , we restrict the learning to the same form as in (7), i.e., $\bar{v}_i \approx \hat{\bar{v}}_i(\Delta t, \Xi_i)$. Hence, the approximation of (9) will also depend on Ξ_i

$$\hat{x}_i(t) = \text{PINN}(\Delta t, x_{0,i}, \Xi_i, u_i). \quad (10)$$

The functional form of the PINN is a standard feed-forward NN with K hidden layers parametrised by the weight matrices \mathbf{W}^k and bias vectors \mathbf{b}^k and a non-linear activation function σ

$$z_0 = [\Delta t, x_{0,i}, \Xi_i, u_i] \quad (11a)$$

$$z_{k+1} = \sigma(\mathbf{W}^{k+1} z_k + \mathbf{b}^{k+1}), \quad \forall k = 0, 1, \dots, K-1 \quad (11b)$$

$$\hat{x}_i = x_{0,i} + \Delta t(\mathbf{W}^K z_K + \mathbf{b}^K). \quad (11c)$$

We adjust the last layer to enforce that the initial condition $x_{0,i}$ is met if $\Delta t = 0$, i.e., $t = t_0$. Thereby we ensure numerical consistency and can improve inaccuracies related to recursive application of the PINN [5]. To train a PINN, we optimise

$$\min_{\mathbf{W}^1, \mathbf{b}^1, \dots, \mathbf{W}^K, \mathbf{b}^K} \mathcal{L}_x + \alpha \mathcal{L}_c, \quad (12)$$

where \mathcal{L}_x evaluates the prediction error based on a dataset \mathcal{D} of simulated data points

$$\mathcal{L}_x = \frac{1}{|\mathcal{D}|} \sum_{j=1}^{|\mathcal{D}|} \left\| \mathbf{x}_i^{(j)} - \hat{\mathbf{x}}_i^{(j)} \right\|_2^2 \quad (13)$$

and a physics-based loss term \mathcal{L}_c evaluated on *collocation* points which do not require any simulation

$$\mathcal{L}_c = \frac{1}{|\mathcal{D}_c|} \sum_{j=1}^{|\mathcal{D}_c|} \left\| \frac{d}{dt} \hat{\mathbf{x}}_i^{(j)} - \mathbf{f} \left(\hat{\mathbf{x}}_i^{(j)}, \hat{v}_i^{(j)}, \mathbf{u}_i^{(j)} \right) \right\|_2^2. \quad (14)$$

The two loss terms are weighted with the hyperparameter α . For a more detailed explanation of PINNs, we refer to [11]. Based on the evolution of the state $\hat{\mathbf{x}}_i(t)$ and the voltage $\hat{v}_i(t)$, we can calculate the current injection \hat{i}_i^C by evaluating

$$\hat{i}_i^C(t) = h_i(\hat{\mathbf{x}}_i(t), \hat{v}_i(t)). \quad (15)$$

For each of the m dynamic components, we will train a separate PINN and all training can be performed in advance of executing PINNSim.

D. Updating the voltage profile

The approximations $\hat{v}_i(t)$ and $\hat{\mathbf{x}}_i(t)$ are both continuous functions and therefore the same holds for $\hat{i}_i^C(t)$ and $\hat{i}_i^N(t)$. The current balance (5) will only be fulfilled, if $\hat{i}_i^C(t) = \hat{i}_i^N(t)$ for the entire time step and at all buses. We express this notion by requiring that the norm $\left\| \hat{i}_i^C - \hat{i}_i^N \right\|$ shall be 0. This norm

$$\left\| \hat{i}_i^C - \hat{i}_i^N \right\| = \sqrt{\langle \hat{i}_i^C - \hat{i}_i^N, \hat{i}_i^C - \hat{i}_i^N \rangle} \quad (16)$$

is induced by the inner product $\langle \bar{a}, \bar{b} \rangle$ between two complex functions $\bar{a}(t), \bar{b}(t)$ over the interval $[t_0, t_0 + \Delta t]$

$$\langle \bar{a}, \bar{b} \rangle = \int_{t_0}^{t_0 + \Delta t} \bar{a}(t) \bar{b}^*(t) dt. \quad (17)$$

As \hat{i}_i^C and \hat{i}_i^N depend on the parametrised voltages $\hat{v}(t, \Xi)$, we subsequently aim to find a parametrisation Ξ that minimises $\left\| \hat{i}_i^C - \hat{i}_i^N \right\|$. To this end, we formulate the following optimisation problem for the entire system as the sum of the norms (we square (16) to avoid the calculation of the square root)

$$\min_{\Xi} \sum_{i=1}^n \left\| \hat{i}_i^C - \hat{i}_i^N \right\|^2. \quad (18)$$

The above expression involves the integration

$$\left\| \hat{i}_i^C - \hat{i}_i^N \right\|^2 = \int_{t_0}^{t_0 + \Delta t} \Re \left(\hat{i}_i^C - \hat{i}_i^N \right)^2 + \Im \left(\hat{i}_i^C - \hat{i}_i^N \right)^2 dt$$

which we approximate with the Midpoint rule, i.e., we split the interval $[t_0, t_0 + \Delta t]$ into s equally sized intervals of width $\frac{\Delta t}{s}$ and sum the function value at the middle of these intervals

$$\left\| \hat{i}_i^C - \hat{i}_i^N \right\|^2 \approx \frac{\Delta t}{s} \sum_{j=1}^s \Re \left(\hat{i}_i^C(t_j) - \hat{i}_i^N(t_j) \right)^2 + \Im \left(\hat{i}_i^C(t_j) - \hat{i}_i^N(t_j) \right)^2. \quad (19)$$

Thereby, we can approximate the optimisation in (18) as

$$\min_{\Xi} \frac{\Delta t}{s} \boldsymbol{\rho}^\top \boldsymbol{\rho} \quad (20)$$

where $\boldsymbol{\rho}$ collects all summands in a vector

$$\boldsymbol{\rho} = \begin{bmatrix} \Re \left(\hat{i}_1^C(t_1) - \hat{i}_1^N(t_1) \right) \\ \Im \left(\hat{i}_1^C(t_1) - \hat{i}_1^N(t_1) \right) \\ \vdots \\ \Re \left(\hat{i}_n^C(t_s) - \hat{i}_n^N(t_s) \right) \\ \Im \left(\hat{i}_n^C(t_s) - \hat{i}_n^N(t_s) \right) \end{bmatrix}, \quad \boldsymbol{\rho} \in \mathbb{R}^{2ns} \quad (21)$$

We solve the above non-linear least square problem by iteratively updating the parameters $\Xi^{(k+1)} = \Xi^{(k)} + \Delta \Xi$. To determine $\Delta \Xi$, we compute the Jacobian

$$\mathbf{J} = \frac{\partial \boldsymbol{\rho}}{\partial \Xi} \Big|_{\Xi^{(k)}}, \quad \mathbf{J} \in \mathbb{R}^{2ns \times 2(r+1)n} \quad (22)$$

at the values $\Xi^{(k)}$ and solve the linear problem

$$(\mathbf{J}^\top \mathbf{J}) \Delta \Xi = -\mathbf{J}^\top \boldsymbol{\rho}. \quad (23)$$

To calculate \mathbf{J} in (22), we use Automatic Differentiation (AD) [21] as all computations for \hat{i}_i^C and \hat{i}_i^N are expressed as explicit functions. While the size of \mathbf{J} can become large, it has a very sparse structure that can be exploited in its construction and when solving (23) – herein lies the key to the scalability of PINNSim as it closely resembles the structure of conventional integration schemes such as the trapezoidal method.

E. PINNSim: The full time-stepping simulator

Algorithm 1 shows the integration of the previous sections into the full PINNSim algorithm that is used for the computation of a time step. Its accuracy depends on the approximation

Algorithm 1 PINNSim - single time step

Require: $\mathbf{x}_0, t_0, t_{\max}, \Delta t, s, r, \Delta \Xi^{\max}, k^{\max}$

Initialise: $k = 0, \Xi^{(0)}, \Delta \Xi^{(0)}$

- 1: **while** $\Delta \Xi^{(k)} > \Delta \Xi^{\max}$ and $k < k^{\max}$ **do**
 - 2: **for** component $i = 1, \dots, m$ & query points $j = 1, \dots, s$ **do**
 - 3: Predict state with PINN $\hat{\mathbf{x}}_i(t_j, \mathbf{x}_0, i, \Xi_i^{(k)})$
 - 4: Calculate component injections $\hat{i}_i^C = h_i(t_j, \hat{\mathbf{x}}_i, \Xi_i^{(k)})$
 - 5: Calculate contribution of \hat{i}_i^C to \mathbf{J}
 - 6: **end for**
 - 7: **for** query points $j = 1, \dots, s$ **do**
 - 8: Calculate network injections $\hat{i}_i^N(t_j, \Xi^{(k)})$
 - 9: Calculate contribution of \hat{i}_i^N to \mathbf{J}
 - 10: **end for**
 - 11: Assemble \mathbf{J} and solve $(\mathbf{J}^\top \mathbf{J}) \Delta \Xi^{(k+1)} = -\mathbf{J}^\top \boldsymbol{\rho}$
 - 12: Update iteration $\Xi^{(k+1)} = \Xi^{(k)} + \Delta \Xi^{(k)}, k = k + 1$
 - 13: **end while**
 - 14: **return** Trajectory across time step $t, \hat{\mathbf{x}}(t), \hat{v}(t)$ with $\Xi^{(\text{final})}$
-

quality of $\hat{\mathbf{x}}(t)$ and $\hat{v}(t)$ and on the tolerance settings for $\Delta \Xi$ and the maximum number of iterations k^{\max} . Therefore, requirements on the resulting tolerance limit the suitable time step size Δt . By repeatedly applying Algorithm 1, we obtain a time-stepping scheme that then allows the simulation of dynamics beyond Δt .

IV. NUMERICAL EXPERIMENTS

This section presents the setup of the numerical experiments to illustrate a proof of concept for PINNSim. The chosen dynamical model describes power system dynamics, we note though, that we use a simplified model as the focus lies on the numerical aspects rather than obtaining insights into the transient phenomena.

A. Power system modelling

As an example for a dynamic component, we consider a two-axis generator model as modelled in [19].

$$\begin{bmatrix} T'_{do} \\ T'_{qo} \\ 1 \\ 2H \end{bmatrix} \frac{d}{dt} \begin{bmatrix} E'_q \\ E'_d \\ \delta \\ \Delta\omega \end{bmatrix} = \begin{bmatrix} -E'_q - (X_d - X'_d)I_d + E_{fd} \\ -E'_d + (X_q - X'_q)I_q \\ \omega_s \Delta\omega \\ P_m - E'_d I_d - E'_q I_q - (X'_q - X'_d)I_d I_q - D \Delta\omega \end{bmatrix} \quad (24a)$$

$$\begin{bmatrix} I_d \\ I_q \end{bmatrix} = \begin{bmatrix} R_s & -X'_q \\ X'_d & R_s \end{bmatrix}^{-1} \begin{bmatrix} E'_d - V \sin(\delta - \theta) \\ E'_q - V \cos(\delta - \theta) \end{bmatrix} \quad (24b)$$

$$\bar{i}^C = (I_D + jI_Q) = (I_d + jI_q)e^{j(\delta - \pi/2)}. \quad (24c)$$

(24a) corresponds to (4b) and (24b) and (24c) to (4a). For this study, we simplify the model above to a classical machine model by setting the reactances to $X'_q = X'_d$ and $X_q = X'_d$ and then finding the integral manifold such that the internal voltages E'_q and E'_d remain constant at E'_{q0} and $E'_{d0} = 0$. For more details we refer to [19]. Now, the rotor angle δ and the frequency deviation $\Delta\omega$ form the state \mathbf{x}_i , the magnitude V and angle θ of the terminal voltage form \bar{v}_i , and the mechanical power P_m and the excitation voltage E_{fd} form the control input \mathbf{u}_i . More detailed components models could include higher order electro-mechanical modes, governor dynamics for P_m and exciter dynamics E_{fd} . Similarly, inverter-based resources or voltage dependent loads could be included.

The results in Section V will demonstrate how PINNSim can effectively increase the allowable time step size. We observe this characteristic already for the following simple setup. We consider the IEEE 9-bus system described in [19, pp. 164–167] with three generators (all modelled as classical machines as above) with parameters from Table I. The initial conditions $\mathbf{x}_{0,i}$ and control inputs \mathbf{u} are determined from assuming an equilibrium state for the load flow case in [19]. To perturb the system, we reduce the mechanical power P_m of generator 1 to 50% of its initial value and then observe the resulting trajectory³.

B. Implementation and NN training

The entire simulator is implemented in PyTorch [22] as we require the functionality of AD for the computation of the Jacobian \mathbf{J} and the training of the PINN models. Each PINN consists of three hidden layers with 32 neurons and

³Here, we apply a disturbance that leaves the power flow at the time of the event unchanged. In contrast, short-circuits and topology changes alter the power flow as the algebraic variables change instantaneously. The resulting jumps can be challenging for simulators. Treating these cases with PINNSim will be addressed in future work.

TABLE I
GENERATOR PARAMETERS AND SET POINTS (ALL IN p.u.)

Gen.	H	D	X_d	X'_d	R_s	P_m	E_{fd}
1	23.64	2.364	0.146	0.0608	0.0	0.71	1.08
2	6.4	1.28	0.8958	0.1969	0.0	1.612	1.32
3	3.01	0.903	1.3125	0.1813	0.0	0.859	1.04

applies a tanh activation function. We employ a Xavier-normalised initialisation [23] for the PINN parameters and subsequently optimise them using a L-BFGS optimiser [24] for 2000 epochs. For each generator the training dataset \mathcal{D} comprises 2500 simulated points, 500 of which form a validation dataset. The collocation dataset \mathcal{D}_c has size 5000. For all datasets we sample from the input domain which consists of the prediction time step Δt , the initial condition $\mathbf{x}_{0,i}$, and a voltage parametrisations Ξ_i with $r = 2$. The bounds of the resulting nine-dimensional input domain are shown in Table II. They are chosen such that all conditions encountered in the simulation are covered. We use the equilibrium values $\delta_i^{eq}, \theta_i^{eq}, V_i^{eq}$ from the power flow solution in the definition of the domain, but it can also be defined without this information.

The numerical simulations of the system of DAEs that serve as ground truth are performed using Assimulo [25]. As a direct comparison to PINNSim, we implement a trapezoidal integration scheme, see [26]. The timing of the simulations is conducted on a AMD Ryzen 7 PRO (1.9 GHz, 8 cores) and 16GB RAM with the implementation of PINNSim that is publicly available [27].

TABLE II
INPUT DOMAIN OF THE TRAINING DATASET

Δt	$\delta_{0,i} - \delta_i^{eq}$	$\Delta\omega_{0,i}$
[0.0, 0.3] s	$[-\pi, \pi]$ rad	$[-0.86, 0.86]$ Hz
$\theta_{0,i} - \theta_i^{eq}$	$\theta_{1,i}$	$\theta_{2,i}$
$[-\pi, \pi]$ rad	$[-0.3, 0.3]$ rad s ⁻¹	$[-0.8, 0.8]$ rad/s ²
$V_{0,i} - V_i^{eq}$	$V_{1,i}$	$V_{2,i}$
$[-0.1, 0.3]$ p.u.	$[-0.4, 0.4]$ p.u. s ⁻¹	$[-0.5, 0.5]$ pu/s ²

V. RESULTS

The potential for accelerating TDSs with PINNSim relies on increasing the time step size and thereby reducing the total number of time steps compared to established methods such as the trapezoidal method. We demonstrate this behaviour in Section V-A on the test case. In Section V-B, we then illustrate what allows PINNSim these larger time steps and describe in Section V-C how the involved computational cost compare.

A. Accuracy of PINNSim for a full TDS

First, we consider a simulation of the described test case over 2.5 s. PINNSim will return trajectories for all states $\mathbf{x}_i(t)$

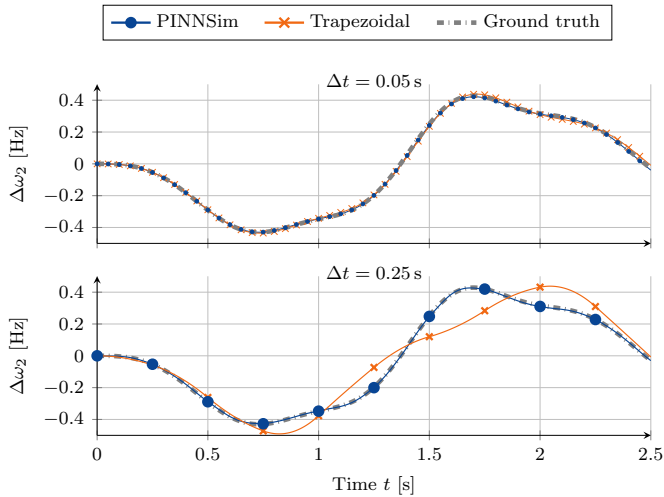


Fig. 2. The panels show the simulated trajectories of the frequency deviation at machine 2 $\Delta\omega_2$ with a trapezoidal and PINNSim time stepping scheme for two time step sizes. The markers indicate the values at the time steps. The curves within a time step stem from the prediction of the PINNs for PINNSim and from a quadratic interpolation for the trapezoidal integration. The ground truth stems from a highly accurate integration scheme.

and voltages $\bar{v}_i(t)$. For clarity of the presentation, we will subsequently only focus on one state, namely the frequency deviation at machine 2, i.e., $\Delta\omega_2$. Figure 2 shows the resulting trajectory of $\Delta\omega_2$. The grey dash-dotted line represents the ground truth solution of the power system dynamics which stems from the Assimulo solver with a tolerance setting of 10^{-12} . The time-stepping schemes, i.e., PINNSim and trapezoidal method, return the values at the end of each time step, indicated by the markers. To obtain the intermediate values we either query the PINNs (for PINNSim) or apply a quadratic interpolation (for the trapezoidal method). While PINNSim and the trapezoidal rule accurately capture the dynamics for a time step size of $\Delta t = 0.05$ s, the trapezoidal rule fails to track the evolution for a larger time step size of $\Delta t = 0.25$ s. In contrast, PINNSim captures the state evolution accurately.

We test this behaviour for more time step sizes and plot in

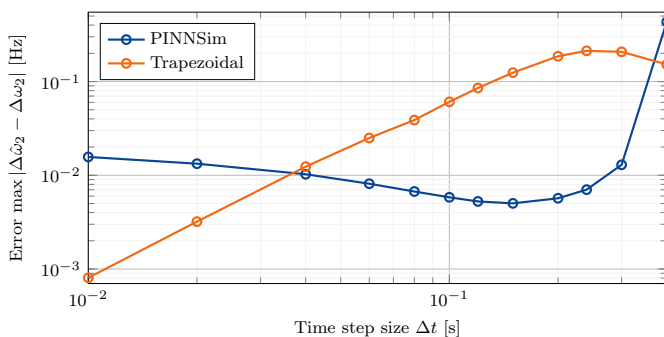


Fig. 3. Comparison of the maximum error of $\Delta\omega_2$ over a trajectory of 2.4 s for a range of time step sizes Δt .

Fig. 3 the maximum error within the simulation time interval. The trapezoidal rule performs well for very small time step sizes, but the errors quickly become very large - for reference, the overall variation of $\Delta\omega_2$ is around 0.8 Hz. PINNSim leads to an almost constant error characteristic, mostly with less than 0.01 Hz error. The slightly increasing error towards smaller time step sizes arises due to accumulating errors. The larger errors for $\Delta t > 0.3$ s are expected, as we trained the PINNs only up to this time step size.

B. Accuracy of PINNSim on a single time step

In the following, we consider a single time step and show the influencing factors on the accuracy. To obtain a range of system states, we simulate the test case for 10 s with an accurate solver and then extract 200 instances, i.e., every 0.05 s.

1) *Voltage parametrisation:* We use these 200 instances as initial values to test the performance of PINNSim on a single time step for different values of Δt , the voltage profile order r and the number of query points s . We evaluate these time




Fig. 4. Maximum error for single time step predictions of different length Δt starting from 200 different initial conditions. Comparison between the trapezoidal rule and PINNSim with different voltage scheme orders (r) and number of query points s . Larger r and s , both improve the accuracy of PINNSim.

steps with respect to the maximum error of $\Delta\omega_2$ at the end of the time step. Figure 4 reports the results. A main observation is that PINNSim with $r = 1$ (linear voltage evolution) and $s = r + 1$ has a comparable relationship between accuracy and time step size as the trapezoidal rule. However, when using PINNSim with $r = 2$, the time step size can be significantly increased without incurring much larger errors. For very small time step sizes PINNSim offers no benefits over the trapezoidal rule as the error decreases slower when reducing Δt . We can improve this characteristic for PINNSim by evaluating (19) on more points, i.e., choosing a larger value for s , shown as the dashed lines. The improved performance, in particular for $r = 1$, originates from the better approximation of (18). However, for large time steps, here for around $\Delta t > 0.2$ s, we additionally require $r = 2$ for accurate predictions; the voltage evolution becomes too non-linear, hence, the linear voltage approximation ($r = 1$) is insufficient. For $r = 2$, the

limitation in time step size arises as the training domain of the PINNs was restricted to $\Delta t \leq 0.3$ s.

2) *PINN accuracy*: The performance of PINNSim relies in large parts on how well each PINN approximates the integration in (9) for a given voltage profile $\hat{v}_i(t, \Xi_i)$. Figure 5 illustrates the results on a test dataset of 4000 points for generator 2. Both panels show the same results, i.e., the maximum and median error on the test dataset, but on a logarithmic and linear x-axis. The logarithmic axis clearly shows that for small time step sizes, the error of PINNs decreases as we included a time dependency in the final layer (11c). The linear x-axis gives a more intuitive understanding of how much further PINNs can predict the dynamics accurately. Only when reaching the limit of the training domain, the accuracy deteriorates. Revisiting the error plots in Figs. 3 and 4, we can find that it is the particular error characteristic of PINNs that allows the large time steps of PINNSim and why it outperforms the trapezoidal method.

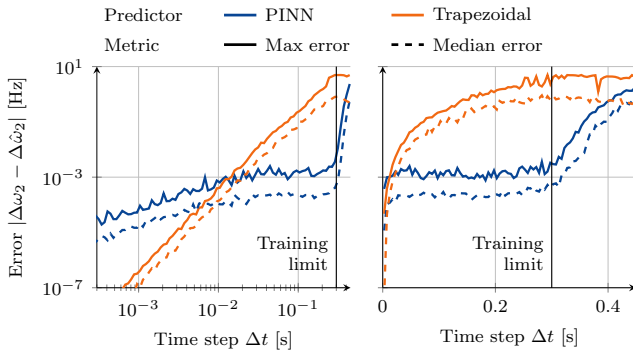


Fig. 5. Error characteristics of the PINN for machine 2 on a test dataset. Both plots show the same results but on a logarithmic x-axis (left) and a linear x-axis (right) to highlight the accuracy of PINNs over large time steps.

C. Computation cost of a time step

The larger time step sizes of PINNSim come at an increased computational cost compared to the trapezoidal rule. We first analyse this cost of PINNSim from a perspective of scalability and then provide a comparison of the total run-time based on the implementation of this proof of concept.

In Algorithm 1, the evaluation of the inner loop (lines 2-8) requires the calculation of network and component current injections \hat{i}^N , \hat{i}^C and their sensitivity to Ξ . For the network currents, this requires merely matrix-vector products and for the component currents the evaluation of the PINNs. A single evaluation (pass) of a PINN requires on the order of 1 μ s, e.g., to compute the current injection. The sensitivity calculations require only two additional passes thanks to AD. As all PINNs could be evaluated in parallel, these computations can easily be scaled to large system. The “expensive” computation arises in the solution of the linear system of equations in line 9 of Algorithm 1, but its sparsity and known structure can be exploited. A closely related calculation is required for the trapezoidal rule.




Fig. 6. Convergence of the objective value in (19) for different number of query points s , voltage profile orders r and time step sizes Δt .

The overall cost then scales linearly with the number of iterations k per time step, i.e., line 1-10 in Algorithm 1, until convergence is reached. We show in Fig. 6 the value of the objective (19) over the iterations for two time step sizes and different values of r . In the left panel, the least square problem is fully determined ($s = r + 1$), hence the objective value continues to decrease. In the right panel, the five additional query points render the problem over-determined, hence the algorithm converges at a non-zero objective value. In either case, we observe that larger time step sizes require more iterations, while the order of the voltage profiles primarily affects the magnitude of the objective value and less its convergence.

Lastly, we provide a comparison of the absolute run-times for the trapezoidal method and PINNSim in Table III. We report the run-time for different combinations of the time step size Δt , the number of iterations k per time step, and the number of query points s for PINNSim. We observe that one faces a trade-off between speed and accuracy with either method. The best values are highlighted and we note that the time step size and the number of iterations k have the biggest influence on this trade-off. We want to stress though, that neither method is optimised for speed, hence the presented results are only indicative for the implementation in [27].

TABLE III
COMPARISON OF TOTAL RUN-TIME FOR SIMULATING 2.4 s

Simulator	Δt [s]	k	s	run-time [s]	$\max \omega_2 - \hat{\omega}_2 $ [Hz]
Trapezoidal	0.04	1	—	0.47	1.34×10^{-2}
		2	—	0.90	1.24×10^{-2}
	0.06	2	—	0.59	2.50×10^{-2}
PINNSim	0.20	5	3	0.72	1.59×10^{-2}
		5	7	0.78	0.98×10^{-2}
	6	3	0.88	0.88×10^{-2}	
	0.24	5	4	0.62	2.52×10^{-2}
		5	7	0.67	1.48×10^{-2}

As stated in [11], [16], the cost of training the PINNs has to be considered as well. Improving the efficiency and performance of the training, e.g., through advanced PINN architectures or hyperparameter tuning, however, can be decoupled from the analysis of PINNSim.

VI. DISCUSSION

The presented results shall serve as a proof of concept of PINNSim as a novel time stepping simulator. We demonstrated that PINNSim outperforms the trapezoidal rule on the metric of the allowable time step size. The following describes four aspects that are necessary to turn PINNSim from a proof of concept into a fully fledged simulator for power system dynamics and to realise the potential acceleration of TDSs.

- 1) *Speed*: Numerical methods require a high level of optimisation in the implementation to become competitive. For PINNSim this concerns primarily the optimisation of the calculation of the residual ρ , the Jacobian \mathbf{J} , and the solution of the sparse linear system in (23). The focus should therefore lie on controlling memory allocations, utilising parallelisation, and exploiting sparsity patterns for these computations.
- 2) *Accuracy*: The accuracy of PINNSim hinges around accurately learned PINNs. We envision that this process can be highly standardised, so that PINNs can be trained reliably to high accuracy and with desirable error characteristics for a wide range of dynamic components.
- 3) *Scale*: The analysis of the computational complexity of PINNSim suggests its scalability but it remains to be shown in practice. Furthermore, we need to investigate the accuracy and convergence properties of the voltage update scheme for larger systems.
- 4) *Power system scenarios*: In this work, we have considered the dynamics of classical machine models under a set-point change. Future work should include more detailed dynamical models and apply PINNSim to short-circuits, topology changes, and other numerically more demanding setups to demonstrate its adequacy for the simulation of power system dynamics.

VII. CONCLUSION

With PINNSim, we introduced a novel approach for time-domain simulations that allows the integration of PINNs. As a result we achieve accurate results with significantly larger time steps than with the common trapezoidal method. Hence, we require fewer time steps which could result in a significant acceleration of time-domain simulations in power systems. By design, we only require the training of PINNs for single dynamical components and thereby enable the scalability of PINNSim to large systems. Additionally, all calculations at run-time are highly scalable and parallelisable, necessary requirements to develop PINNSim into a powerful simulator.

REFERENCES

- [1] B. Stott, "Power system dynamic response calculations," *Proceedings of the IEEE*, vol. 67, no. 2, pp. 219–241, Feb. 1979.
- [2] Y. Liu and K. Sun, "Solving power system differential algebraic equations using differential transformation," *IEEE Transactions on Power Systems*, vol. 35, no. 3, pp. 2289–2299, May 2020.
- [3] G. Gurrula, A. Dimitrovski, S. Pannala, S. Simunovic, and M. Starke, "Parareal in Time for Fast Power System Dynamic Simulations," *IEEE Transactions on Power Systems*, vol. 31, no. 3, pp. 1820–1830, 2016.
- [4] P. Aristidou, "Time-domain simulation of large electric power systems using domain-decomposition and parallel processing methods," Ph.D. dissertation, Université de Liège, Liège, Belgium, 2015.
- [5] I. Lagaris, A. Likas, and D. Fotiadis, "Artificial neural networks for solving ordinary and partial differential equations," *IEEE Transactions on Neural Networks*, vol. 9, no. 5, pp. 987–1000, Sep. 1998.
- [6] M. Raissi, P. Perdikaris, and G. E. Karniadakis, "Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations," *Journal of Computational Physics*, vol. 378, no. C, Nov. 2018.
- [7] C. Legaard *et al.*, "Constructing Neural Network Based Models for Simulating Dynamical Systems," *ACM Computing Surveys*, vol. 55, no. 11, pp. 1–34, Nov. 2023.
- [8] G. S. Misyris, A. Venzke, and S. Chatzivasilieadis, "Physics-Informed Neural Networks for Power Systems," in *2020 IEEE Power & Energy Society General Meeting (PESGM)*, Montreal, QC, Canada, pp. 1–5.
- [9] C. Moya, G. Lin, T. Zhao, and M. Yue, "On Approximating the Dynamic Response of Synchronous Generators via Operator Learning: A Step Towards Building Deep Operator-based Power Grid Simulators," Jan. 2023, arXiv:2301.12538.
- [10] T. Xiao, Y. Chen, S. Huang, T. He, and H. Guan, "Feasibility Study of Neural ODE and DAE Modules for Power System Dynamic Component Modeling," *IEEE Transactions on Power Systems*, pp. 1–13, 2022.
- [11] J. Stiasny and S. Chatzivasilieadis, "Physics-informed neural networks for time-domain simulations: Accuracy, computational cost, and flexibility," *Electric Power Systems Research*, vol. 224, p. 109748, Nov. 2023.
- [12] C. Moya and G. Lin, "DAE-PINN: a physics-informed neural network model for simulating differential algebraic equations with application to power networks," *Neural Computing and Applications*, vol. 35, no. 5, pp. 3789–3804, Feb. 2023.
- [13] W. Cui, W. Yang, and B. Zhang, "A Frequency Domain Approach to Predict Power System Transients," *IEEE Transactions on Power Systems*, vol. 39, no. 1, pp. 465–477, Jan. 2024.
- [14] C. Roberts *et al.*, "Continuous-time echo state networks for predicting power system dynamics," *Electric Power Systems Research*, vol. 212, p. 108562, Nov. 2022.
- [15] J. Li, M. Yue, Y. Zhao, and G. Lin, "Machine-learning-based online transient analysis via iterative computation of generator dynamics," in *2020 IEEE International Conference on Communications, Control, and Computing Technologies for Smart Grids (SmartGridComm)*, pp. 1–6.
- [16] J. Stiasny, "Physics-Informed Neural Networks for Power System Dynamics," Ph.D. dissertation, Technical University of Denmark, 2023.
- [17] K. E. Brenan, S. L. Campbell, and L. R. Petzold, *Numerical Solution of Initial-Value Problems in Differential-Algebraic Equations*. Society for Industrial and Applied Mathematics, Jan. 1995.
- [18] G. Adomian, "A review of the decomposition method in applied mathematics," *Journal of Mathematical Analysis and Applications*, vol. 135, no. 2, pp. 501–544, Nov. 1988.
- [19] P. W. Sauer and M. A. Pai, *Power system dynamics and stability*. Upper Saddle River, N.J: Prentice Hall, 1998.
- [20] B. Wang, N. Duan, and K. Sun, "A Time–Power Series-Based Semi-Analytical Approach for Power System Simulation," *IEEE Transactions on Power Systems*, vol. 34, no. 2, pp. 841–851, Mar. 2019.
- [21] A. G. Baydin, B. A. Pearlmutter, A. A. Radul, and J. M. Siskind, "Automatic Differentiation in Machine Learning: a Survey," *Journal of Machine Learning Research*, vol. 18, no. 153, pp. 1–43, 2018.
- [22] A. Paszke *et al.*, "PyTorch: An imperative style, high-performance deep learning library," in *Advances in neural information processing systems*, H. Wallach, H. Larochelle, A. Beygelzimer, F. dAlché Buc, E. Fox, and R. Garnett, Eds., vol. 32. Curran Associates, Inc., 2019.
- [23] X. Glorot and Y. Bengio, "Understanding the difficulty of training deep feedforward neural networks," in *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics. JMLR Workshop and Conference Proceedings*, 2010, pp. 249–256.
- [24] D. C. Liu and J. Nocedal, "On the limited memory BFGS method for large scale optimization," *Mathematical Programming*, vol. 45, no. 1–3, pp. 503–528, Aug. 1989.
- [25] C. Andersson, C. Führer, and J. Åkesson, "Assimulo: A unified framework for ODE solvers," *Mathematics and Computers in Simulation*, vol. 116, pp. 26–43, Oct. 2015.
- [26] F. Milano, *Power System Modelling and Scripting*, ser. Power Systems. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010.
- [27] J. Stiasny, "PINNSim implementation." [Online]. Available: <https://github.com/jbesty>