

Hybrid Symbolic-Numeric Framework for Power System Modeling and Analysis

Hantao Cui , Member, IEEE, Fangxing Li , Fellow, IEEE, and Kevin Tomsovic , Fellow, IEEE

Abstract—With the recent proliferation of open-source packages for computing, power system differential-algebraic equation (DAE) modeling and simulation are being revisited to reduce the programming efforts. Existing open-source tools require manual efforts to develop code for numerical equations, sparse Jacobians, and discontinuous components. This paper proposes a hybrid symbolic-numeric framework, exemplified by an open-source Python-based library ANDES, which consists of a symbolic layer for descriptive modeling and a numeric layer for vector-based numerical computation. This method enables the implementation of DAE models by mixing and matching modeling components, through which models are described. In the framework, a rich set of discontinuous components and standard transfer function blocks are provided besides essential modeling elements for rapid modeling. ANDES can automatically generate robust and fast numerical simulation code, as well as and high-quality documentation. Case studies present a) two implementations of turbine governor model TGOV1, b) power flow computation time break down for MATPOWER systems, c) validation of time-domain simulation with commercial software using three test systems with a variety of models, and d) the full eigenvalue analysis for Kundur’s system. Validation shows that ANDES closely matches the commercial tool DSATools for power flow, time-domain simulation, and eigenvalue analysis.

Index Terms—Power systems, open-source, DAE modeling, symbolic calculation, time-domain simulation, ANDES library.

I. INTRODUCTION

POWER system modeling and transient simulation is a widely studied yet challenging topic. Digital computer-based simulation has been dominating in the industry and academia with both closed-source tools [1] and open-source tools [2]–[11] widely used. Although simulation software comes with a set of built-in models, users will likely need to customize models for new devices or control algorithms.

To develop new models for simulation software is to implement the model equations in a program that can interact with the predefined software architecture. In general, there are two

Manuscript received February 20, 2020; revised June 24, 2020 and July 30, 2020; accepted August 11, 2020. Date of publication August 17, 2020; date of current version February 19, 2021. This work was supported in part by the Engineering Research Center Program of the National Science Foundation and the Department of Energy under NSF Award EEC-1041877 and in part by the CURENT Industry Partnership Program. Paper no. TPWRS-00276-2020. (Corresponding author: Fangxing Li.)

The authors are with the Department of Electrical Engineering and Computer Science, The University of Tennessee, Knoxville, TN 37996 USA (e-mail: hcui7@utk.edu; fli6@utk.edu; tomsovic@utk.edu).

Color versions of one or more of the figures in this article are available online at <https://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TPWRS.2020.3017019

approaches to implement user-defined models (UDMs): programmatically or through a graphical user interface (GUI) [12], [13], which is usually not available in open-source tools due to complexity and lack of return. Still, open-source tools are crucial for scientific research, but they require programming proficiency to develop new models on top of a deep understanding of the tool [14].

Two advanced UDM solutions exist in open-source tools: Dome cards [5] and the Function Mockup Unit (FMU) support in GridDyn [8]. Dome cards are plain-text files containing model descriptions in the card protocol. Using a symbolic library under the hood, Dome uses cards to generate intermediate code that can be modified into final models. Although cards are flexible, they do not live with the simulation code, and manual tweaks are often required. On the other hand, FMU is compiled directly from Modelica, an equation-based modeling language. Modelica libraries such as OpenIPSL [15] have been developed for power system simulation. Although FMU has excellent speed and interoperability through the Functional Mockup Interface (FMI), it has seen few adoptions in power system tools due to path dependence¹ and, technically, data structure.²

This work proposes a hybrid symbolic-numeric method aiming to reduce the efforts for modeling differential-algebraic equation (DAE) in power systems while maintaining numerical performance with the help of a symbolic toolbox. The proposed method can be applied to major programming languages. An implementation has been open-sourced as the ANDES library written in Python, a scripting language suitable for power systems research and rapid prototyping. Different from Dome cards, symbolically defined models are part of the library and distributed with the program. Main contributions are as follows:

- 1) The proposed hybrid symbolic-numeric method allows simple scripting of DAE models with descriptive equation strings instead of hard-coded implementations.
- 2) ANDES is the first open-source power system tool that enables writing models from block diagrams using modular discontinuous components and modeling blocks (such as transfer functions and proportional-integral controllers).
- 3) The library can generate efficient and robust numerical code from descriptive models for fast simulation.

¹Most of the widely used commercial tools today have a vast library of built-in models, which started to accumulate long before FMU was invented.

²In Modelica/FMU, models are written separately and used combinatorially. Some implementations even require the precompilation of all possible combinations.

- 4) By preserving numerical interfaces, it can accommodate models that are much easier to implement in the traditional numerical way.

The prior works on symbolic modeling and our advancements are discussed in the following. Decades ago, symbolic approaches to power flow modeling [16], optimization [17], [18], and device transients modeling [19]–[21] were introduced. The pioneering works well proved the concept but exposed a remaining issue: scalability. Namely, symbolic equations must be written for each device instance rather than each model type [20]. For large systems, a massive number of repetitive symbolic equations need to be created, which are difficult to maintain and solve. Besides, any system topology change requires manual modification to equations and is thus prone to errors. In contrast, the proposed library models the abstract model type in the symbolic layer, agnostic to test systems. Therefore, the computation time to process symbolic equations scales to the number (and the complexity) of model types, not the number of devices in any particular test case. In the generated code, vectorization is utilized for speed, thus equations of all devices of the same type are updated in the same function calls.

This paper is organized as follows. Section II discusses the motivations and design philosophy of the work. Section III and Section IV explain the techniques for the symbolic and numeric layers with sufficient examples. Section V presents case studies, including two implementations of the TGOV1 model, power flow for MATPOWER systems [22], time-domain simulation verification with DSATools TSAT using three test systems with a variety of models, and full eigenvalue analysis. Section VI concludes the proposed work.

II. MOTIVATIONS AND DESIGN PHILOSOPHY

The overarching goal of the proposed hybrid symbolic-numeric method and its implementation in ANDES for power system modeling and analysis is to *make modeling as simple as describing equations and make simulations as fast as using crafted code*. Simplifying DAE modeling renders the library easy to use and modify for research and education. Maintaining a fast simulation speed makes the library capable of running large-scale studies. As discussed, a purely symbolic approach will not scale to large systems, and a purely numerical approach will not reduce the programming efforts. Therefore, a hybrid approach is proposed to take advantage of symbolic and numeric approaches in one library.

The design philosophy is two-fold: 1) to enable descriptive modeling using provided modeling elements and blocks, and 2) enable robust and fast numerical simulation through code generation and vectorization. The first item can be realized in the symbolic layer in which model developers can mix and match parameters, variables, discrete components to describe DAE models. The second item can be realized through code generation from symbolically defined equations and coordination of the numerical functions.

Fig. 1 shows the overview of the proposed hybrid symbolic-numeric framework with the upper part for hybrid modeling and the lower part for numerical simulation. This framework can

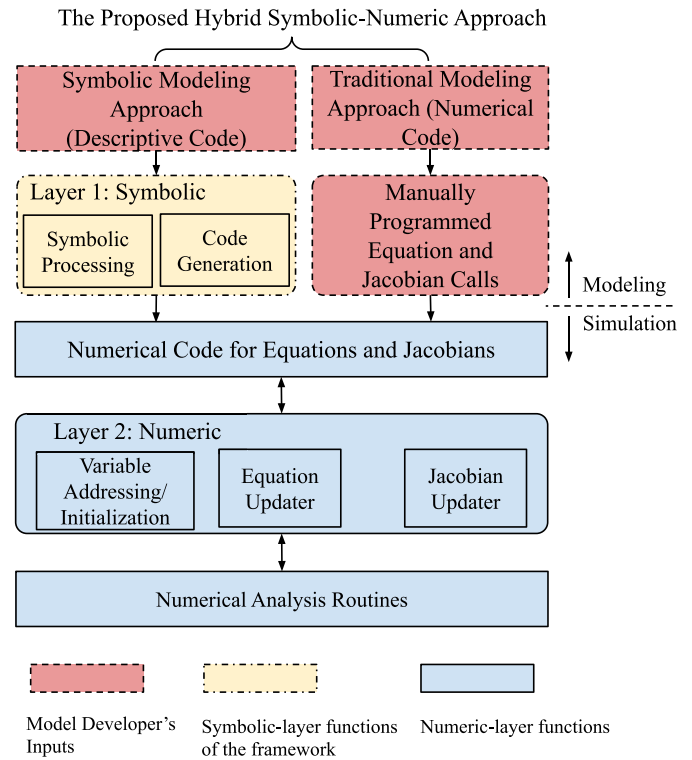


Fig. 1. Overview of the hybrid symbolic-numeric approach for modeling and simulation. Red boxes with dotted border indicate the required manual efforts.

accommodate two modeling approaches: 1) the proposed symbolic modeling approach using descriptive code, and 2) the traditional numerical modeling approach. The symbolic approach is recommended due to simplicity and robustness because less programming is needed. The symbolic layer can automatically generate symbolic equations and Jacobians, which, altogether, will be generated into loadable numerical code [23]. It ensures the same models will be used for simulation and documentation to achieve consistency between description and simulation. Alternatively, the traditional numerical modeling approach can be used if a model cannot be easily implemented in the symbolic modeling approach.

The lower part of Fig. 1 shows the numeric layer in the proposed framework for simulation. This layer organizes numerical code for equations and Jacobians, which include those generated by the symbolic layer and the manually written ones, to provide interface methods for addressing, initialization, and equation evaluation. Power system cases are loaded, and vector operations are utilized for optimal performance in a scripting environment. Routine developers can develop specific numerical routines by calling the provided interface methods in specific orders.

The two-layer hybrid architecture also benefits the end-users who are not looking to develop models but instead use the library as a simulation tool. Procedures in the symbolic layer only need to be executed once by the end-user, and the generated code will be serialized to disk for future reuse. In terms of simulation performance, the proposed framework is on par with pure numerical libraries, since all computations in the numerical layer use vector operations.

```

1 class Shunt(Model):
2     def __init__(self):
3         self.bus = IdxParam(info='bus_index')
4         self.g = NumParam(info='conductance', unit='pu')
5         self.b = NumParam(info='susceptance', unit='pu')
6         self.a = ExtAlgeb(model='Bus', indexer=self.bus,
7                           src='a', e_str='g*v*v')
8         self.v = ExtAlgeb(model='Bus', indexer=self.bus,
9                           src='v', e_str='-b*v*v')

```

Listing 1: Shunt model for power flow (imports are omitted for simplicity).

III. SYMBOLIC MODELING FRAMEWORK

This section describes the implementation of the symbolic layer for the proposed library. The symbolic layer covers class-based declarative modeling, symbolic processing, code generation, and automated documentation. Methods discussed in this section are exemplified in the Python language with the SymPy library but can be extended to other environments.

A. Basic Modeling Elements

The proposed library starts by observing that all DAE models can be described with a few categories of basic modeling elements. Such categories include parameters, variables, discrete components, and services:

- 1) *Parameters* are typically externally supplied data for defining specific devices.
- 2) *Variables* either differential or algebraic, are the unknowns to be solved in the DAE system. Each variable is associated with values and an equation.
- 3) *Discrete Components* describe the discontinuities, such as limits, associated with variables.
- 4) *Services* are assisting types for simplifying expressions or fulfilling supplementary actions.

The framework provides the above categories of modeling elements that can be instantiated to describe DAE models. Modeling elements are containers in both symbolic and numeric layers. In the symbolic layer, modeling elements are containers for metadata, such as name, description, unit, and equation strings. In the numeric layer, they provide storage for associated numerical data, such as values and addresses.

B. Classes for Descriptive Modeling

Python classes are the top-level containers to describe models. A class for a DAE model can be created by defining class member attributes using the provided modeling elements. The idea is best explained with a simple example, such as a constant shunt capacitor model for power flow given by

$$\begin{aligned} p_h &= -gv_h^2 \\ q_h &= bv_h^2 \end{aligned} \quad (1)$$

where h is the connected bus index, v is the bus voltage, p and q are the power injections, and g and b are the conductance and susceptance, respectively. The implementation for the Shunt model is given in Listing 1 with the following remarks:

- 1) Lines 3–5 declares parameters `bus`, `g`, and `b` for bus index, shunt conductance, and susceptance value.




Fig. 2. A typical PSS final output limiter.

- 2) Lines 6–9 declares external algebraic variables `a` and `v` for voltage phase and magnitude at the buses whose indices are `bus`.
- 3) Lines 7 and 9 declares the active power load (v^2g) and reactive power load ($-v^2b$) on the power balance equations associated with `a` and `v`.
- 4) The `e_str` equation strings contain `a`, `v`, `b` and `b` strings, which are declared data attributes of the class.

It is important to note that Listing 1 is an abstract Shunt model rather than just one particular Shunt device. The Shunt model will host all Shunt devices of the same kind through vectorization so that only one invocation is needed for each equation. An excellent discussion on this design choice can be found in Chapter 9.2 of [14].

Like a compiler, the underlying symbolic library requires a list of symbols to process equation strings. The base class `Model` handles the bookkeeping of member attributes for all derived models. Models can automatically capture the names and attributes instances to the corresponding storage in the declaration sequence based on attribute type. In Python, this is achieved by overloading the `__setattr__()` protocol, which is invoked every time an attribute is assigned. Therefore, the captured names will be converted to symbols for equation processing. The approach allows us to keep the class definition concise while automatically performs the bookkeeping.

Therefore, the efforts to develop DAE models have been reduced. All that required is to set up correct element containers and describe the mathematical equations.

C. Discrete Components

Discrete components such as limiters and deadlocks are common in practical models but are intricate to implement. They often require manipulating equations and Jacobian, which, if not implemented correctly, can halt simulations. In existing tools, discrete components are implemented *ad hoc* and require manual efforts to be ported from one model to another.

The proposed library provides discrete components that are readily usable for describing DAE models. Discrete components can export binary flags, which are evaluated in the numerical layer, to indicate the discontinuous status. Flags can be used in equations to construct piece-wise equations with the benefit of not manipulating Jacobian matrices since discrete flags are preserved as variables in the corresponding derivative equations.

For example, a hard limiter takes an algebraic variable and two limit parameters as inputs and exports three flags, `z_l`, `z_u`, and `z_v` to indicate within, at the lower, and at the upper limits. As a use case, consider a typical power system stabilizer (PSS) output limiter shown in Figure 2, where the final output V_{out} depends on the terminal voltage V_t and the given limits V_{CL} and V_{CU} . The output limiter can be conveniently implemented

```

1 self.OL = HardLimit (u=self.Vt, lower=self.VCL,
2                     upper=self.VCU)
3 self.Vout = Algeb(e_str='vss*OL_zi-Vout',
4                  v_str='vss*OL_zi')

```

Listing 2: Stabilizer output limiter implementation.




Fig. 3. Illustration of Reduce and Repeat services for COI.

as in Listing 2, where Lines 1-2 creates a hard limiter called OL that exports flags `OL_zi`, `OL_zl` and `OL_zu`. Line 3 utilizes `OL_zi` to construct the output variable `Vout` with its equation through `e_str` and the initial value equation through `v_str`.

D. Services

While the descriptive equation modeling is robust and straightforward, one needs to realize the limitation: vector operations are limited to arithmetic calculations. Descriptive equations cannot handle programmatic operations such as conditions and loops. Services are helper types to overcome such limitations by allowing computing and storing values outside the DAE system using user-defined functions. They are custom-computed but used in the same way as variables.

An illustrative example is the calculation of inertia weights in the center-of-inertia (COI) model. As shown in Figure 3, each COI device links to a number of generators (stored in `syn`), retrieves their inertia H_s , and needs to compute the weights on the rotor speed for each linked generator. A numerical program can quickly sum up the inertia and divide each inertia by the sum. However, since element-wise vector operations do not allow summation, the proposed library introduces two service types, one to reduce H_s into H_t using a summation function and the other to repeat the sum H_t into the same shape as H_s . The element-wise division H_s/H_t can be performed thereafter.

E. Modeling Blocks

In addition to descriptive equation modeling, the library allows us to write models directly from transfer function diagrams. A similar concept was reported in the InterPSS controller modeling language (CML) [24], which utilizes the Java Annotation feature to provide a scripting environment for controller prototyping. ANDES allows the composition of modeling elements into reusable modeling blocks, which can exports variables with equation templates. Modeling blocks are instantiated as class member attributes like variables. Upon instantiation, variable name placeholders in equation strings will be substituted




Fig. 4. A chained transfer functions example.

```

1 self.LG = Lag(u=self.u, T=self.T, K=self.K)
2 self.LL = LeadLag(u=self.LG_y, T1=self.T1,
3                  T2=self.T2)

```

Listing 3: Implementation of chained lag and lead-lag transfer functions.

with the actual names. About 20 commonly used proportional-integral controllers and transfer functions, some with limiters, have been implemented.

For example, the chained transfer functions in Figure 4 can be implemented in barely two lines of self-explanatory code, as given by Listing 3. Internally, model elements with their equations tailored with the instance name will be exported and captured by the hosting model class. Block outputs are always named the block instance name with an underscore and letter `y`. In this example, Line 1 exports a differential variable named `LG_y`, which is passed as an input in Line 2. Similarly, the output of the lead-lag instance is accessible as `LL_y`.

Modeling blocks can save efforts to reimplement equations in different models and improve readability. In the meantime, modeling blocks can be mixed with custom descriptive equations using the exported variables whenever flexibility is needed.

F. Symbolic Processing and Code Generation

The symbolic processor converts the metadata, namely, equation strings, into symbolic expressions for symbolic Jacobian derivation, code generation, and documentation. These functionalities are part of the base `Model` class and will be inherited by all derived models. An external symbolic library is utilized to generate the symbolic expressions and Jacobian matrices for each model with the following steps:

- 1) Prepare all variable symbols into a vector `xy` in the declaration order so that each variable has a stable index.
- 2) Convert each equation string to a symbolic expression (using `sympy.sympify`).
- 3) Group differential and algebraic expressions into two vectors, `f` and `g`, respectively, in the declaration order.
- 4) Derive the expression vectors with respect to the ordered variable vector to obtain Jacobian matrices $\frac{df}{dxy}$ and $\frac{dg}{dxy}$ (using `sympy.Matrix.jacobian`).
- 5) Convert the Jacobian matrices to sparse to obtain non-zero triplets (`row`, `column`, `value`), where `row` is the index of the equation in the equation array, `column` the variable index, and `value` the derivative expression.

The following performance characteristics are relevant. Symbolic processing is executed over each model, and thus the processing time scales linearly to the number of models. Each model only has a few to tens of equations; thus, the processing time is fast. The processing is done before loading any system and is test-case independent.

The symbolic processing for `Shunt` is illustrated in Equations (2) to (4). The Jacobian derivation and triplet conversion shown in Equations (4) are automated with the symbolic library.

$$\mathbf{xy} = [a, v] \quad (2)$$

$$\mathbf{g} = [v^2g, -v^2b] \quad (3)$$

$$\frac{d\mathbf{g}}{d\mathbf{xy}} = \underbrace{\begin{bmatrix} 0 & 2vg \\ 0 & -2vb \end{bmatrix}}_{\text{dense}} \rightarrow \underbrace{\begin{bmatrix} (0, 1, 2vg) \\ (1, 1, -2vb) \end{bmatrix}}_{\text{sparse (row, column, value)}} \quad (4)$$

Code generation generates and stores numerical functions that are executable and will return the values of expressions. The code generation feature of the external symbolic library is utilized in the following steps:

- 1) Generate numerical functions for each initialization, differential and algebraic equations, and each element in Jacobian Matrices (using `sympy.lambdify`).
- 2) For each Jacobian matrix, store the equation index `row`, variable index `column`, and the anonymous function for `value` correspondingly in lists.

It is important to note that `row` and `column` are local to each model and only depends on the number of declared variables. The following remarks are relevant.

- 1) In terms of performance, the generated numerical functions use the efficient NumPy library for vectorial computation and thus runs as fast as manually crafted code.
- 2) The overhead for symbolic processing and code generation can be eliminated by reusing the generated code through efficient serialization and de-serialization.
- 3) The library also takes manually written numerical function calls, as long as indices are provided and functions have the same signature as the generated code. This feature can be helpful to reuse existing numerical code.

At this point, executable numerical code is obtained from the symbolically described DAE models.

G. Documentation

Code documentation is essential for disseminating open-source research but is often underappreciated. The situation is understandable because maintaining documentation can take as much as, if not more than, the development efforts. All the existing power system simulators rely on manual efforts to document the implemented models.

The proposed library can automatically document the implemented equations for DAE models developed using declarative classes. Human-friendly equations can be generated from symbolic expressions by substituting in LATEX-formatted variable strings. The documentation feature completes the symbolic layer to ensure the same models are used for simulation and documentation. To the best knowledge of the authors, the proposed library is the first in power system tools capable of generating equation documentation directly from source code. For interested readers, the documentation is available online [25], and the model documentation is under Section “Model References”.

TABLE I
MODELING ELEMENTS AND THEIR NUMERICAL ATTRIBUTES

	Value (<i>v</i>)	Address (<i>a</i>)	Equation (<i>e</i>)	Flags
Parameter	✓			
Variable	✓	✓	✓	
Discrete				✓
Service	✓			

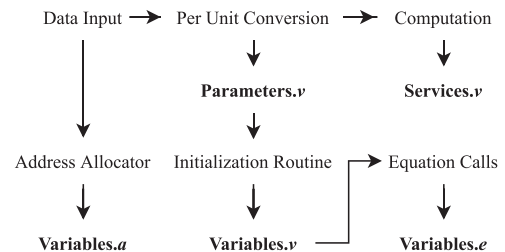


Fig. 5. Data flow paths for setting up and the numerical storage.

IV. NUMERIC LAYER IMPLEMENTATION

The numeric layer establishes data structure for vector operations, and dispatches generated numerical code for the procedures in numerical simulation, such as setting initial values, updating equations, and building Jacobian matrices.

A. Data Structure and Vector Storage

The numeric layer relies on arrays and sparse matrices to properly store data associated with declared elements. Numerical values belonging to a modeling element instance are stored in the instance attributes. Depending on the type, an element instance may contain member attributes for addresses, values, and equations. Table I shows the supported attributes of the element types. Each address, value, and equation value attributes are stored as an array with its length equal to the number of devices. For example, if a particular system contains three `Shunt` devices, attributes `b` and `g` will each contain a value array `v` with a length of three.

Numerical arrays are updated at different phases in simulations, as outlined in Fig. 5. Parameter values are set after loading the data file and converting it to per unit under the system base. Variable addresses are allocated after loading the test system, values set by initialization calls, and equation values updated by equation calls. Service values are updated in multiple phases — some are computed when accessed for the first time, and others are computed after parameters are set. Discrete flags are updated before or after equation updates, depending on the discrete type.

B. Variable Initialization

Variable initialization routine sets variable initial values before a routine starts. It includes setting the starting point for power flow and initializing the rest of the variables for dynamic routines. Although power flow initialization is simple, there could be value conflicts depending on the input data format. For example, default initial bus voltages are set by buses and overwritten by PV-generators. The library uses an additional

flag to indicate if the values from one model overwrite the shared variables at the end.

Variable initialization for dynamics is mathematically a root-finding problem for the DAE system with all derivatives zeroed out. Two approaches can be used: sequential or iterative. Variables with an explicit solution can be initialized sequentially, while those without must be solved iteratively.

The library provides three entry points for initialization. First, an explicit-form equation can be specified for each variable if it can be initialized sequentially. A common technique is to set the initialization equation for a service that calculates from other services. Second, an optional, implicit-form equation with an initial value can be specified for each variable. All implicit equations will be gathered and solved iteratively from the given initial value. Third, a placeholder function is available if one decides to write numerical code. For best practice, sequential initialization should be used whenever possible. For convergence consideration, the initial values for the iterative initializers need to be carefully selected.

C. Numerical Equation Evaluation

After loading a test case and counting the total number of variables, four numerical arrays are created to hold all variables and equations. Each variable in a model receives an array of addresses indexing into the corresponding DAE array. The same addresses can be used to access the corresponding numerical values of variables and equations.

It is worth noting that the power system data structure introduces external variables for one model to link to another. As shown in Listing 1, the Shunt model creates two external algebraic variables, a and v , for linking to Bus devices with the indices given by `bus`. Variable addresses of the linked Bus devices will be assigned to Shunt so that Shunt has access to the Bus phase angles and voltages.

Memory copying of arrays imposes a significant overhead in numerical simulation. As a solution, all internal variables are assigned contiguous addresses so that a no-copy array view can be stored locally in each model. External variables are not guaranteed to link to contiguous devices, so their variables and equations are stored in local arrays and merged into the DAE arrays after evaluation. Although this implementation is specific to *NumPy*, the general rule applies to avoid memory copying, especially in computation-intensive programs. However, one needs to realize the downside of this approach — it rules out the possibility of parallelizing equation updates across models. Since parallel equation updates are difficult in Python due to global interpreter lock (GIL), this shared-memory sequential evaluation approach will give the best performance.

Steps to update equations for each model are as follows.

- 1) Copy external variables from DAE arrays to model.
- 2) Call generated numerical functions using local values as inputs and store the outputs locally.
- 3) Update equation values for equation-dependent limiters such as anti-windup limiters.
- 4) Merge local external equation values to DAE equations.

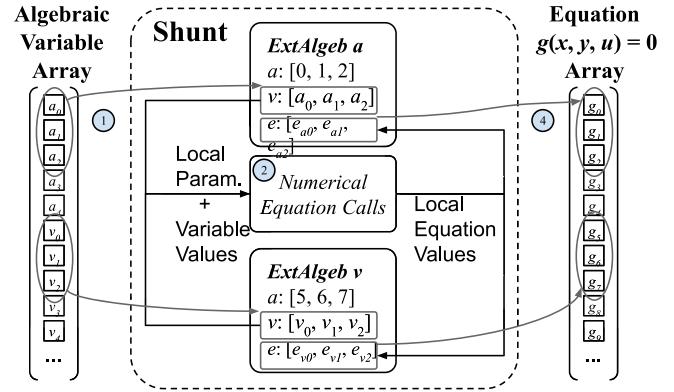


Fig. 6. Illustration of the equation update procedure (without Step 3).

This procedure (without step 3) is illustrated in Fig. 6. Note that step 3 is needed for models with equation-dependent limiters. Anti-windup limiters, for example, check the equation values to update the limiter status. Step 3 updates limiter status and sets the differential equation values to zero for the binding anti-windup limiters.

D. Incremental Jacobian Building

Building Jacobian matrices involve steps to fill in sparse Jacobian matrices incrementally and efficiently. It is especially relevant for implicit numerical integration routine since Jacobian updates take up the most overhead. This subsection discusses how the Jacobian indexing is done with the local variable indices (from the symbolic layer) and variable addresses (assigned in the numeric layer).

It is worth noting the difference between the local variable indices and the assigned variable addresses. A local variable index is a scalar number based on the sequence of declaration and is independent of test cases. Variable addresses are assigned as arrays after loading a specific test case. Local indices are used to look up corresponding addresses in order to determine the positions of the values.

For a generic triplet (`row`, `column`, `value(*args)`) where `row` and `column` are two scalars for the local indices, and `value(*args)` is the numerical function for the Jacobian value with `args` being a list of local values. Recall that `value` is the derivative of the `row`-th equation with respect to the `column`-th variable. Jacobian values, which have the same length as the `row` and `column` addresses, should be summed at the positions defined by the case-specific addresses for the `row`-th equation and the `col`-th variables.

Fig. 7 illustrates the process with three Shunt devices as an example. There are two Jacobian triplets from the symbolic layer to be placed at local indices (0, 1) and (1, 1). In the numeric layer, the zeroth variable a is assigned addresses $[0, 1, 2]$ and the first variable v is assigned address $[5, 6, 7]$. Evaluate the numerical function $2vg$ to obtain the Jacobian elements, for example, $[0.002, 0.002, 0.002]$. Next, these elements will be summed up at positions with the row number equal to the addresses of a ($[0, 1, 2]$) and the column number equal to the address of v

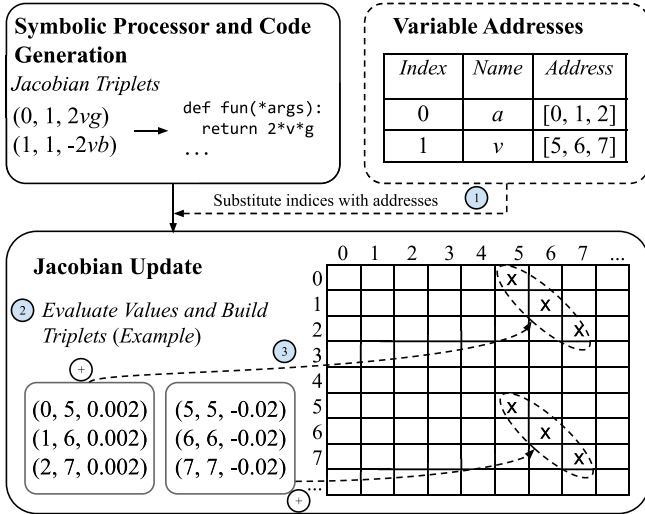


Fig. 7. Illustration of the Jacobian update procedure.

([5, 6, 7]). Repeat the process until all elements from all models are added.

For performance consideration, the library implements a two-step process that builds the sparsity pattern for one time and then fills in the values repeatedly. It is known that incrementally building sparse matrices can be time-consuming if repeated memory allocation is needed. By using the addresses of elements, zero-filled sparsity pattern matrices can be constructed. The memory for the non-zero elements is pre-allocated, and in-place modifications can apply. This technique is especially relevant for high-level languages without direct memory access.

V. CASE STUDIES

For verification and demonstration, this section presents a model implementation, power flow calculation, time-domain simulation, and eigenvalue analysis. The implementation of turbine governor model TGOV1 is demonstrated with source code developed in the proposed library. Next, power flow results are reported with their time breakdown analyzed. Further, time-domain simulation and eigenvalue analysis are verified against DSATools 19.0.

All subsequent studies are performed in CPython 3.7.7 with ANDES 1.0.3, SymPy 1.5.1, NumPy 1.18.4, and CVXOPT 1.2.5 on an AMD Ryzen 7 2700X CPU running Debian 10. In addition, a custom C-based routine is used for fast in-place sparse matrix addition.

A. Example Model: TGOV1

The TGOV1 turbine governor model [26] (shown in Fig. 8) is used as a practical example with sufficient complexity to demonstrate the proposed work. This model is composed of a lead-lag transfer function and a first-order lag transfer function with an anti-windup limiter. The corresponding differential equations and algebraic equations are given in (5) and (6).

$$\begin{bmatrix} \dot{x}_{LG} \\ \dot{x}_{LL} \end{bmatrix} = \begin{bmatrix} z_{i,lim}^{LG} (P_d - x_{LG}) / T_1 \\ (x_{LG} - x_{LL}) / T_3 \end{bmatrix} \quad (5)$$




Fig. 8. The control diagram of TGOV1 turbine governor.

```

1 def __init__(self):
2     # 1. Declare parameters from case file inputs.
3     self.R = NumParam(info='Turbine_governor_droop',
4                       non_zero=True, ipower=True)
5     # Other parameters are omitted to conserve space.
6
7     # 2. Declare external variables from generators.
8     self.omega = ExtState(src='omega', model='SynGen',
9                          indexer=self.syn)
10    self.tm = ExtAlgeb(src='tm', indexer=self.syn,
11                    model='SynGen', e_str='u*(pout-tm0)')
12
13    # 3. Declare services for temporary values.
14    self.G = ConstService(e_str='u/R')
15    self.tm0 = ExtService(src='tm',
16                       model='SynGen', indexer=self.syn)
17
18    # 4. Declare variables and equations.
19    self.pref = Algeb(v_str='tm0*R',
20                   e_str='tm0*R-pref')
21    self.wd = Algeb(e_str='(1-omega)-wd')
22    self.pd = Algeb(v_str='tm0',
23                  e_str='G*(wd+pref)-pd')
24    self.LG_y = State(v_str='pd',
25                    e_str='LG_lim_zi*(pd-LG_y)/T1')
26    self.LG_lim = AntiWindup(u=self.LG_y,
27                          lower=self.VMIN,
28                          upper=self.VMAX)
29    self.LL_x = State(v_str='LG_y',
30                    e_str='(LG_y-LL_x)/T3')
31    self.LL_y = Algeb(v_str='LG_y',
32                    e_str='T2/T3*(LG_y-LL_x)+LL_x-LL_y')
33    self.pout = Algeb(v_str='tm0',
34                    e_str='(LL_y+Dt*wd)-pout')
    
```

Listing 4: Implementation of the TGOV1 Model.

$$\begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} = \begin{bmatrix} (1 - \omega) - \omega_d \\ R \times \tau_{m0} - P_{ref} \\ (P_{ref} + \omega_d) / R - P_d \\ D_t \omega_d + y_{LL} - P_{OUT} \\ \frac{T_2}{T_3} (x_{LG} - x_{LL}) + x_{LL} - y_{LL} \\ u (P_{OUT} - \tau_{m0}) \end{bmatrix} \quad (6)$$

where LG and LL denote the lag block and the lead-lag block, \dot{x}_{LG} and \dot{x}_{LL} are the internal states, y_{LL} is the lead-lag output, ω the generator speed, ω_d the generator under-speed, P_d the droop output, τ_{m0} the steady-state torque input, and P_{OUT} the turbine output that will be summed at the generator.

An implementation of the TGOV1 model using descriptive equations is given in Listing 4. It consists of four types of declarations: parameters, external variables, initial external values, and internal variables and equations. Parameters are declared with special properties for data consistency and per-unit conversion. For example, Line 4 specifies that the droop parameter R must be non-zero and is an inverse-of-power per-unit quantity in

```

1 self.GA = Gain(u='wd+pref', K=self.G)
2 self.LG = LagAntiWindup(u=self.GA_y, T=self.T1,
3     K=1, lower=self.VMIN, upper=self.VMAX)
4 self.LL = LeadLag(u=self.LG_y, T1=self.T2,
5     T2=self.T3)

```

Listing 5: Block implementation of the three transfer functions in TGOV1.

TABLE II
TIME BREAKDOWN (IN SECONDS) FOR MATPOWER TEST CASES

Name	Total Iterations	Solve Equations	Update Equations	Build Jacobians	Total
300	6	0.002	0.002	0.008	0.016
1354pegase	6	0.006	0.002	0.020	0.038
2736sp	5	0.012	0.003	0.033	0.061
6515rte	5	0.036	0.006	0.091	0.189
9241pegase	7	0.092	0.013	0.232	0.421
ACTIVSg10k	5	0.065	0.008	0.120	0.250
ACTIVSg25k	8	0.281	0.028	0.526	0.982

device base MVA. External variable ω is retrieved for calculation and τ_m for power feedback to generators. Note that the equation associated with τ_m replaces the steady-state constant torque τ_{m0} with the turbine output P_{OUT} . The initial value of the mechanical torque is retrieved for variable initialization. Finally, differential and algebraic variables are declared, followed by the mathematical equations in (5)–(6) written in a descriptive format, making it convenient to understand and troubleshoot.

Alternatively, modeling blocks can be used to model part of TGOV1 directly from the transfer function diagram. That is, lines 21–32 in Listing 4 can be simplified into Listing 5, which is highly readable and similar to using a visual modeling tool in a scripting manner. Note that variable `pd` have been replaced with `GA_y` in Listing 5, but the rest remain the same. Modeling users can readily utilize blocks such as `Gain`, `LagAntiWindup` and `LeadLag` without having to reimplement the underlying standard equations.

B. Power Flow Calculation

ANDES implements a Newton-Raphson method for power flow calculation as the first proof of concept. Models for bus, PQ, PV, transmission line, and shunt are developed, and a full Newton-Raphson routine is implemented using the direct sparse linear solver KLU.³ Unlike conventional power flow packages, the symbolically implemented line model does not implement an admittance matrix, although it is feasible to do so numerically. Instead, vector computation of line injections into buses are used to maintain generality across models.

The power flow routine is benchmarked using test systems from MATPOWER 7.0. With the same settings and start points, ANDES is able to solve the cases listed in Table II and obtain identical results to that from MATPOWER. Note that the actual ANDES computation time is about 10% shorter than these reported in the table since the profiler was turned on to obtain the time breakdown.

³KLU is not shipped with CVXOPT but is available through an add-on package `cvxoptklu` (compilation required).




Fig. 9. The speed of generators on Buses 1 and 3.




Fig. 10. Terminal voltages on Buses 1 and 3.

The time breakdown exposes some interesting facts. Updating the numerical equations and solving the linear equations is relatively fast and takes up less than 30% of the time. About half of the time is consumed for filling in Jacobian elements, even though an efficient C-based routine is used to modify values in place. The Jacobian time, however, can be reduced by implementing a dishonest algorithm that avoids updating Jacobians at every iteration step.

C. Time-Domain Numerical Integration

To validate the numerical simulation results, ANDES is compared with the commercial package DSATools TSAT using Kundur's two area system, IEEE 14-bus system and Northeast Power Coordinating Council (NPCC) 140-bus system. All PQ loads are converted to constant impedance after power flow calculation. The implicit trapezoidal method is used with a fixed step size of 1/30 second.

The Kundur's system has four generators [27] in GENROU models [28], each with an EXDC2 exciter and a TGOV1 turbine governor. Parameters of the system are listed in the Appendix. At $t = 2s$, one of the two lines between Bus 8 and Bus 9 is disconnected. The simulation takes 1.2 seconds to complete. Generator speed, terminal voltage, and excitation voltage following a line trip event are compared. Simulation results are depicted in Fig. 9–Fig. 11. Clearly, the proposed hybrid symbolic-numeric library achieves almost the same time-domain simulation results.




Fig. 11. Excitation voltages of generators on Buses 1 and 3.




Fig. 14. NPCC 140-bus system rotor speed comparison.




Fig. 12. IEEE 14-bus system rotor speed comparison.




Fig. 15. NPCC 140-bus system voltage comparison.

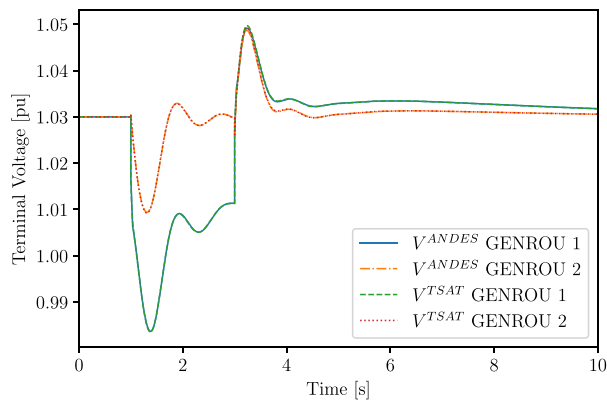


Fig. 13. IEEE 14-bus system voltage comparison.

The modified IEEE 14-bus system for validation uses a variety of models implemented in the hybrid symbolic-numeric framework. These models include generator model GENROU, exciter models ESST3A and EXST1, turbine governor models TGOV1 and IEEEG1, and PSS models ST2CUT and IEEEEST. An extreme scenario that opens line 1-2 at 1 second and reconnects it after 2 seconds is used to trigger nonlinearity. The simulation takes 4.1 seconds to complete. Generator rotor speeds and terminal voltages in Fig. 12 and Fig. 13 show perfect matches

with TSAT. The successful validation of ANDES using this system confirms the correct implementation of all the above models using the proposed framework.

The NPCC 140-bus system (with generator models GENCLS and GENROU, exciter models IEEEEX1 and turbine governor models TGOV1) is studied. The simulation takes 2.5 seconds to complete. The rotor speed and voltage plots in Fig. 14 and Fig. 15 also show perfect match.

It is also important to note that even commercial software does not always agree with each other, especially in large systems, due to factors such as unpublished implementation details and automatic parameter corrections. Nevertheless, the discussed verification provides satisfactory results to prove the proposed concept using the above three test systems.

D. Eigenvalue Analysis

Lastly, the numerical routine for eigenvalue analysis is developed by reusing existing eigenvalue programs. Eigenvalues of the state matrix obtained after the time-domain initialization are plotted in Fig. 16. Two dotted lines in the figure are the loci with 5% damping. Also, the first three eigenvalues ranked by damping ratio (ζ) are compared between ANDES and DSATools SSAT in Table III. The comparison shows that the numerical eigenvalue




Fig. 16. Relevant eigenvalues in the S -domain for Kundur's system.

TABLE III
EIGENVALUE RESULTS COMPARISON FOR KUNDUR'S SYSTEM

	ANDES		SSAT	
	Eigenvalue	ζ (%)	Eigenvalue	ζ (%)
#1	$-0.192 \pm j4.225$	4.53	$-0.192 \pm j4.221$	4.55
#2	$-0.656 \pm j7.086$	9.22	$-0.657 \pm j7.083$	9.23
#3	$-0.653 \pm j6.834$	9.50	$-0.653 \pm j6.832$	9.51

analysis routine in ANDES can obtain very close results to the commercial software SSAT.

VI. CONCLUSION

In conclusion, this paper presents a hybrid symbolic-numeric library for DAE-based power system modeling and numerical simulation. This paper presented the design philosophy for a two-layer library that brings together the advantages of symbolic and numeric approaches. The symbolic layer is case-independent and handles descriptive modeling, symbolic processing, code generation, and automated documentation. The numeric layer organizes the generated code for case-dependent initialization, equation update, and Jacobian update. The simplicity of modeling using the proposed library is demonstrated with a TGOV1 turbine governor model. The library is verified for power flow calculation against MATPOWER, and the computation time is analyzed. It is also verified for time-domain simulation using Kundur's system, IEEE 14-bus system, and NPCC system with a variety of dynamic models. The reference implementation in the ANDES library can obtain very close results for time-domain simulation and eigenvalue analysis to DSATools.

APPENDIX
MODEL PARAMETERS FOR KUNDUR'S
TWO AREA SYSTEM

TABLE IV
BUS DATA

uid	idx	Vn	v0	a0	area
0	1	20	1.000	0.570	1
1	2	20	0.998	0.369	1
2	3	20	0.963	0.185	2
3	4	20	0.817	0.462	2
4	5	230	0.979	0.480	1
5	6	230	0.958	0.284	1
6	7	230	0.936	0.127	1
7	8	230	0.879	-0.081	2
8	9	230	0.891	0.094	2
9	10	230	0.830	0.337	2

TABLE V
LINE DATA

uid	idx	bus1	bus2	r	x	b	tap	phi
0	Line_0	5	6	0.005	0.050	0.075	1	0
1	Line_1	5	6	0.005	0.050	0.075	1	0
2	Line_2	6	7	0.002	0.020	0.030	1	0
3	Line_3	6	7	0.002	0.020	0.030	1	0
4	Line_4	7	8	0.022	0.220	0.330	1	0
5	Line_5	7	8	0.022	0.220	0.330	1	0
6	Line_6	7	8	0.022	0.220	0.330	1	0
7	Line_7	8	9	0.002	0.020	0.030	1	0
8	Line_8	8	9	0.002	0.020	0.030	1	0
9	Line_9	9	10	0.005	0.050	0.075	1	0
10	Line_10	9	10	0.005	0.050	0.075	1	0
11	Line_11	1	5	0.001	0.012	0.000	1	0
12	Line_12	2	6	0.001	0.012	0.000	1	0
13	Line_13	3	9	0.001	0.012	0.000	1	0
14	Line_14	4	10	0.001	0.012	0.000	1	0

TABLE VI
PQ DATA

uid	idx	bus	p0	q0
0	PQ_0	7	11.59	-0.735
1	PQ_1	8	15.75	-0.899

TABLE VII
PV DATA

uid	idx	bus	p0	q0	v0	ra	xs
0	2	2	7	3.0	1	0	0.25
1	3	3	7	5.5	1	0	0.25
2	4	4	7	-1.0	1	0	0.25

TABLE VIII
SLACK DATA

uid	idx	bus	p0	q0	v0	ra	xs	a0
0	1	1	7.459	1.436	1	0	0.25	0.57

TABLE IX
GENROU DATA

uid	idx	bus	gen	D	M	xl	xq	xd	xd1	xd2	xq1	xq2	Td10	Td20	Tq10	Tq20
0	1	1	1	0	13.00	0.06	1.7	1.8	0.3	0.25	0.55	0.25	8	0.03	0.4	0.05
1	2	2	2	0	13.00	0.06	1.7	1.8	0.3	0.25	0.55	0.25	8	0.03	0.4	0.05
2	3	3	3	0	12.35	0.06	1.7	1.8	0.3	0.25	0.55	0.25	8	0.03	0.4	0.05
3	4	4	4	0	12.35	0.06	1.7	1.8	0.3	0.25	0.55	0.25	8	0.03	0.4	0.05

TABLE X
EXDC2 DATA

uid	idx	syn	TR	TA	TC	TB	TE	TF1	KF1	KA	KE	VRMAX	VRMIN
0	1	1	0.02	0.02	1	1	0.83	1.246	0.075	20	1	5.2	-4.16
1	2	2	0.02	0.02	1	1	0.83	1.246	0.075	20	1	5.2	-4.16
2	3	3	0.02	0.02	1	1	0.83	1.246	0.075	20	1	5.2	-4.16
3	4	4	0.02	0.02	1	1	0.83	1.246	0.075	20	1	5.2	-4.16

TABLE XI
TGOV1 DATA

uid	idx	syn	R	VMAX	VMIN	T1	T2	T3	Dt
0	1	1	0.05	33	0.4	0.49	2.1	7	0
1	2	2	0.05	33	0.4	0.49	2.1	7	0
2	3	3	0.05	33	0.4	0.49	2.1	7	0
3	4	4	0.05	33	0.4	0.49	2.1	7	0

ACKNOWLEDGMENT

The authors would like to thank Nicholas West for developing the C-based routine for fast in-place sparse matrix addition.

REFERENCES

- [1] V. Jalili-Marandi, F. J. Ayres, E. Ghahremani, J. Belanger, and V. Lapointe, "A real-time dynamic simulation tool for transmission and distribution power systems," in *Proc. IEEE Power Energy Soc. General Meeting*, 2013, pp. 1–5.
- [2] J. H. Chow and K. W. Cheung, "A toolbox for power system dynamics and control engineering education and research," *IEEE Trans. Power Syst.*, vol. 7, no. 4, pp. 1559–1564, Nov. 1992.
- [3] E. Zhou, "Object-oriented programming, c++ and power system simulation," *IEEE Trans. Power Syst.*, vol. 11, no. 1, pp. 206–215, Feb. 1996.
- [4] F. Milano, "An open source power system analysis toolbox," *IEEE Trans. Power Syst.*, vol. 20, no. 3, pp. 1199–1206, Aug. 2005.
- [5] F. Milano, "A Python-based software tool for power system analysis," in *Proc. IEEE Power Energy Soc. General Meeting*, Vancouver, BC, Canada, 2013, pp. 1–5.
- [6] S. Cole and R. Belmans, "Matdyn, a new matlab-based toolbox for power system dynamic simulation," *IEEE Trans. Power Syst.*, vol. 26, no. 3, pp. 1129–1136, Aug. 2011.
- [7] M. Zhou and Q. Huang, "InterPSS: A new generation power system simulation engine," 2017. [Online]. Available: <http://arxiv.org/abs/1711.10875>
- [8] P. Top, Y. Qin, and L. Min, "Integration of functional mock-up units into a dynamic power systems simulation tool," in *Proc. IEEE Power Energy Soc. General Meeting*, 2016, pp. 1–5.
- [9] H. Cui and F. Li, "ANDES: A Python-based cyber-physical power system simulation tool," in *Proc. North Amer. Power Symp.*, 2018, pp. 1–5.
- [10] H. Cui, F. Li, and K. Tomsovic, "Cyber-physical system testbed for power system monitoring and wide-area control verification," *IET Energy Syst. Integration*, vol. 2, no. 1, pp. 32–39, 2019.
- [11] F. Li, K. Tomsovic, and H. Cui, "A large-scale testbed as a virtual power grid: For closed-loop controls in research and testing," *IEEE Power Energy Mag.*, vol. 18, no. 2, pp. 60–68, Mar./Apr. 2020.
- [12] Siemens, "PSS/E graphical model builder," 2017. [Online]. Available: <https://assets.new.siemens.com/siemens/assets/api/uuid:d089003aea3cc71b97b762b5c3fac134ff2f752c/09-pti-ff-en-sw-gmb-1611.pdf>
- [13] PowerTech, "DSATools," 2020. [Online]. Available: <https://www.powertechlabs.com/dsatools-services>
- [14] F. Milano, *Power System Modelling and Scripting*. Berlin, Germany: Springer, 2010.
- [15] M. Baudette, M. Castro, T. Rabuzin, J. Lavenius, T. Bogodorova, and L. Vanfretti, "OpenIPSL: Open-instance power system library—Update 1.5 to 'iTesla Power Systems Library (IPSL): A modelica library for phasor time-domain simulations'," *SoftwareX*, vol. 7, pp. 34–36, Jan.–Jun. 2018, doi: [10.1016/j.softx.2018.01.002](https://doi.org/10.1016/j.softx.2018.01.002).
- [16] F. L. Alvarado and Y. Liu, "General purpose symbolic simulation tools for electric networks," *IEEE Trans. Power Syst.*, vol. 3, no. 2, pp. 689–697, May 1988.
- [17] I. Džafić, F. L. Alvarado, M. Glavić, and S. Tešnjak, "A component based approach to power system applications development," in *Proc. 14th Power Syst. Comput. Conf.*, 2002, pp. 24–28.
- [18] I. Džafić, M. Glavić, and S. Tešnjak, "A component-based power system model-driven architecture," *IEEE Trans. Power Syst.*, vol. 19, no. 4, pp. 2109–2110, Nov. 2004.
- [19] F. L. Alvarado, C. A. Cañizares, A. Keyhani, and B. Coates, "Instructional use of declarative languages for the study of machine transients," *IEEE Power Eng. Rev.*, vol. 6, no. 1, pp. 407–413, Feb. 1991.
- [20] F. L. Alvarado, C. A. Cañizares, and J. Mahseredjian, "Symbolically-assisted power system simulation," *Int. J. Elect. Power Energy Syst.*, vol. 18, no. 7, pp. 405–408, 1996.
- [21] W. Gao, E. V. Solodovnik, and R. A. Dougal, "Symbolically aided model development for an induction machine in virtual test bed," *IEEE Trans. Energy Convers.*, vol. 19, no. 1, pp. 125–135, Mar. 2004.
- [22] R. D. Zimmerman, C. E. Murillo-Sánchez, and R. J. Thomas, "MATPOWER: Steady-state operations, planning, and analysis tools for power systems research and education," *IEEE Trans. Power Syst.*, vol. 26, no. 1, pp. 12–19, Feb. 2011.
- [23] A. Meurer *et al.*, "SymPy: Symbolic computing in python," *PeerJ. Comput. Sci.*, vol. 3, 2017, Art. no. e103, doi: [10.7717/peerj-cs.103](https://doi.org/10.7717/peerj-cs.103).
- [24] M. Zhou, "Interpss controller modeling language," 2012. [Online]. Available: <https://docs.google.com/document/d/1zvME4YBibCbEswVgS0PcQJdeA9ESMt9JAoBy7AGrr7c/preview>
- [25] H. Cui, "ANDES documentation," 2020. [Online]. Available: <https://andes.readthedocs.io>

- [26] Powerworld, "Governor TGOV1 model reference," 2020. [Online], Available: https://www.powerworld.com/WebHelp/Content/TransientModels_PDF/Generator/Governor/Governor%20TGOV1.pdf
- [27] P. Kundur, *Power System Stability And Control*. New York, NY, USA: McGraw-Hill, 1994.
- [28] M. Zhang, M. Baudette, J. Lavenius, S. Løvlund, and L. Vanfretti, "Mod- elica implementation and software-to-software validation of power system component models commonly used by nordic TSOs for dynamic simu- lations," in *Proc. 56th Conf. Simul. Modelling*, Linköping Univ., Sweden, 2015, pp. 105–112.



Hantao Cui (Member, IEEE) received the bachelor's and master's degrees in electrical engineering from Southeast University, China, in 2011 and 2013, respectively and the Ph.D. degree from the University of Tennessee, in 2018. Prior to obtaining the Ph.D. degree, he was appointed as a full-time Research Associate with CURENT in 2017. He is currently a Research Assistant Professor with CURENT and the Department of Electrical Engineering and Computer Science, University of Tennessee, Knoxville.

He serves as an Associate Editor for the *Journal of Modern Power Systems and Clean Energy*. He is recognized as an Outstanding Reviewer for 2019 by the IEEE TRANSACTIONS ON POWER SYSTEMS and is Author of two Essential Science Indicators (ESI) highly cited papers. His research interests include power system modeling and simulation, high-performance computing, and software-hardware solutions for scientific computing.



Fangxing Li (Fellow, IEEE) is also known as Fran Li. He received the B.S.E.E. and M.S.E.E. degrees from Southeast University, Nanjing, China, in 1994 and 1997, respectively, and the Ph.D. degree from Virginia Tech, Blacksburg, VA, USA, in 2001. He is currently the James McConnell Professor with the University of Tennessee, Knoxville, TN, USA.

His research interests include deep learning in power systems, renewable energy integration, demand response, power market and power system computing. Since 2020, he has been the Editor-In-Chief for the IEEE OPEN ACCESS JOURNAL OF POWER AND ENERGY and the Chair of IEEE/PES Power System Operation, Planning and Economics (PSOPE) Committee.



Kevin Tomsovic (Fellow, IEEE) received the B.S. degree from Michigan Technological University, Houghton, MI, USA, in 1982, and the M.S. and Ph.D. degrees from the University of Washington, Seattle, WA, USA, in 1984 and 1987, respectively, all in electrical engineering.

He is currently the CTI Professor with the Department of Electrical Engineering and Computer Science (EECS), University of Tennessee, Knoxville, TN, USA, where he also directs the NSF/DOE Engineering Research Center CURENT. He was the EECS Department Head from 2008 to 2013. He was on the Faculty of Washington State University, Pullman, WA, USA, from 1992 to 2008. He was the Advanced Technology for Electrical Energy Chair with Kumamoto University, Kumamoto, Japan, from 1999 to 2000, and was the NSF Program Director with the Electrical and Communications Systems Division of the Engineering Directorate from 2004 to 2006.