DOE Research Performance Final Report – Cover Page

- Submitted to Department of Energy, Office of Electricity Delivery and Energy Reliability
- Award Number: DE-OE0000624
- Title: High Fidelity, "Faster than Real-Time" Simulator for Predicting Power System Dynamic Behavior
- PI: Alexander J. Flueck, Associate Professor, flueck@iit.edu, 312-567-3625
- DUNS Number: 042084434
- Recipient: Illinois Institute of Technology, 3300 South Federal Street, Chicago, IL 60616-3793
- Consortium members:
  - Alstom Grid (now GE Energy)
  - Argonne National Laboratory
  - Electrocon International Inc.
  - McCoy Energy
- Project start date: 1 Oct 2012
- Project end date: 28 Feb 2017

Submitted by Alexander J. Flueck, PhD

Acknowledgment: This material is based upon work supported by the Department of Energy under Award Number(s) DE-OE0000624.

Executive summary

The "High Fidelity, Faster than Real-Time Simulator for Predicting Power System Dynamic Behavior" was designed and developed by Illinois Institute of Technology with critical contributions from Electrocon International, Argonne National Laboratory, Alstom Grid and McCoy Energy. Also essential to the project were our two utility partners: Commonwealth Edison and AltaLink. The project was a success due to several major breakthroughs in the area of large-scale power system dynamics simulation, including (1) a validated faster than real-time simulation of both stable and unstable transient dynamics in a large-scale positive sequence transmission grid model, (2) a three-phase unbalanced simulation platform for modeling new grid devices, such as independently controlled single-phase static var compensators (SVCs), (3) the world's first high fidelity three-phase unbalanced dynamics and protection simulator based on Electrocon's CAPE program, and (4) a first-of-its-kind implementation of a single-phase induction motor model with stall capability. These four project goals are summarized here and presented in the report with additional details

Faster than real-time dynamics simulation

The prototype dynamics simulator achieved an impressive 2x faster-than-real-time speedup during stable dynamics for an 18,946 bus case with 1,263 detailed generators and 1,018 Type 1 direct current exciters. In addition, the faster-than-real-time simulator achieved a 1.68 speedup during a mix of stable and unstable dynamics for a 20 second simulation. However, to run faster-than-real-time is not enough. The simulation results need to be accurate. Therefore, an industry standard simulator was used in the benchmarking process to verify the accuracy of the faster-than-real-time simulator.

Unbalanced three-phase network modeling

The unbalanced three-phase network modeling is a major breakthrough for transmission dynamics analysis. It enables the investigation of network imbalances, such as those due to untransposed lines or unbalanced faults. It also enables power system engineers to perform model validation on a wider set of recorded disturbances, which are typically unbalanced. In addition, unbalanced three-phase modeling will enable a more accurate simulation of the contributions from and the impacts on future devices with a power electronics interface, such as static var compensators.

Comprehensive dynamics and protection simulation

The TS3ph-CAPE co-simulation framework introduced the world's first high fidelity three-phase unbalanced dynamics and protection simulator based on the CAPE simulation engine. In addition to the comprehensive dynamics and protection system modeling capabilities, the co-simulation framework is ideal for studying cascading failures. Protective devices have played a key role in past cascading failures, but there remains significant uncertainty as to when and what type of relay might assert a trip signal to a breaker. Rather than model protective devices on an ad-hoc basis, which is a current industry practice, the co-simulation framework enables a transmission planning engineer to simulate actual relay responses with currently deployed field settings.

Single-phase induction motor modeling with stall capability

The successful implementation and demonstration of the single-phase induction motor model with stall capability is a major breakthrough. The stall dynamics are critical to modeling the phenomenon known as fault induced delayed voltage recovery. The single-phase induction motor dynamics also play a critical role in slower time scale voltage collapse. Therefore, the first-of-its-kind implementation of a single-phase induction motor model with stall capability will allow engineers to simulate and design mitigation strategies against voltage collapse and cascading outages.

Modeling and computational challenges

Several challenges were encountered throughout the project. A brief summary is provided here.

- Computational complexity: Detailed generators, detailed controllers and detailed induction motor models contain nonlinear elements due to nonlinear physical phenomena and controller limits, including both windup and non-windup limits. Incorporating these details will increase the time required to complete a simulation.
- Weak scalability due to tightly coupled model equations: Power grid models do not lend themselves to naturally parallel decompositions as the network structure is highly meshed, similar to computational fluid dynamics. Yet at the same time, the network graphs are unstructured, which is unlike computational fluid dynamics problems. This feature of power grids will increase the time required to complete a simulation.
- Memory bandwidth limits for large sparse matrices: Modern computer architectures exploit data locality, which is a desirable feature in most computational problems, with a hierarchy of caches between the processing core and main memory. Unfortunately, power grid simulations depend on sparse matrix representations. Hence, indirect addressing is required, which significantly weakens the data locality properties of large-scale grid models. This issue will increase the time required to complete a simulation.
- Detailed protection models require extensive communication (i.e., collect voltages, currents and breaker status signals) and extensive computation (i.e., determine whether or not a relay element will assert that a specific condition has been met, and evaluate breaker operation logic) to determine protection actions. The extra communication and extra computation will increase the time required to complete a simulation.
- Three-phase unbalanced models of large-scale power systems may not yield faster than real time simulation results, so a tradeoff between modeling detail and simulation performance must be considered.

The project was a success with several major breakthroughs in the area of large-scale power system dynamics simulation, including

- A validated faster than real-time simulation of both stable and unstable transient dynamics in a large-scale positive sequence transmission grid model,
- A three-phase unbalanced simulation platform for modeling new grid devices, such as independent phase controllers for single-phase static var compensators,
- The world's first high fidelity three-phase unbalanced dynamics and protection simulator based on Electrocon's CAPE program, and
- A first-of-its-kind implementation of a single-phase induction motor model with stall capability.

Beyond the current project, there are ongoing research and development initiatives that were launched from this project:

- CAPE is currently being redesigned with all of the knowledge gained during this project, to implement all of the efficiency improvements that were considered to be out of scope. This implementation will include x64, multiple threads, and multiple processes.
- The research prototype TS3ph-CAPE simulator is being used to explore cascading outages.

In the long-term, the simulator capabilities could form the backbone of the newly conceived hybrid real-time protection and control architecture that will coordinate local controls, wide-area measurements, wide-area controls and advanced real-time prediction capabilities. The nation's citizens will benefit in several ways, including (1) less down time from power outages due to the faster-than-real-time simulator's predictive capability, (2) higher levels of reliability due to the detailed dynamics plus protection simulation capability, and (3) more resiliency due to the three-phase unbalanced simulator's ability to model three-phase and single-phase networks and devices.

Accomplishments

This section provides a comparison of the team's actual accomplishments versus the original milestones of the project.

Milestone Title and Description	Planned End Date	Actual End Date	Verification Method	Comments (progress toward achieving milestone, explanation of deviation from plan, etc.)
Task 1 - Update Project Management Plan (PMP)	Ongoing	1/17/13	Submitted to EEREPMC	Complete: any proposed modification will be reviewed with DOE.
Task 2 - Develop testing plan, collect test data cases, create test harness prototype and test these using the simulator with TS; design CAPE API	3/31/13	3/31/13	Testing plan shared; utility data tested with TS3ph; CAPE API reviewed by IIT	Complete: testing plan document reviewed by team and approved; utility data obtained by IIT; open test case (ESCA60 from Alstom Grid) established; new TS routines implemented; CAPE API reviewed and approved; CAPE API implementation initiated.
Task 3 - Identify and recruit members for the UAG	3/31/13	3/31/13	NDAs signed	Complete: Utility partners each signed an NDA; Utility Advisory Group meetings were scheduled and conducted.
Task 4 - Submit for publication technical papers describing the mathematical or computational methodologies in publicly available, peer reviewed journals/reports	3/31/2013, 9/30/2014, 9/30/2015	Ongoing	Papers accepted for publication	Several papers have been published based on Dr. Abhyankar's earlier PhD research at IIT. One conference paper was accepted for IREP 2017. Two panel sessions at IEEE PES GM 2016 in Boston, MA were held. One full paper has been accepted for the special issue of IEEE Transactions on Smart Grid. One panel session was held at IEEE PES GM 2014 in DC.
Task 5 - Develop, optimize and test the simulator (PETSc, TS, realistic data, UAG, CAPE, Rosenbrock)	9/30/14	9/30/14	Testing completed; results provided in previous quarterly repors, Year 2 report and during Year 2 site visit.	Complete: IIT and Alstom Grid verified operation of TS3ph against TSAT for three-phase and single-line-to-ground faults. IIT and Electrocon tested TS3ph-CAPE for three-phase bus fault. ANL performed tests with mulit-rate schemes. ANL tested TS3ph with PETSc variable-step methods. IIT tested TS3ph with UAG data.

Task 6 - Design and implement a restructured multi-thread CAPE core	5/31/15	5/31/15	Delivered multi-threaded CAPE program to IIT	Complete: IIT tested a sample case (SAVNW from Siemens and Electrocon) using CAPE-15 (development version of multi-threaded CAPE program).
Task 7 - Prepare realistic operating scenarios using utility data and test simulator	9/30/14	9/30/14	Cases prepared; results provided in current and previous quarterly reports and during Year 2 site visit.	Complete: Alstom prepared cases using the ESCA60 model. Contingencies were provided by IIT and the following practical scenarios were modeled: Power System Stressed conditions, SVC controls to increase damping, Composite load cases, and Induction Motor cases. ANL tested parallel TS3ph with 2737 and 22,996 bus cases. IIT tested serial TS3ph with utility case.
Task 8 - Demonstrate software enhancements within 30 days after completion of Tasks 5 through 7	9/30/14, 5/31/15	9/30/14 9/30/15	Simulation results presented in Year 2 report and during Year 2 site visit. Phase 2 results presented in Year 3 report and during Year 3 site visit.	Complete: Task 8 completed for 2014 deliverables, as described in Tasks 5 through 7. Task 8 completed for 2015 deliverables. Accomplishments were presented to DOE Program Officer and Program Manager during annual in-person site reviews at IIT.
Task 9 - Perform validation and verification of the developed codes and algorithms (TS3PH vs. DTS/TSAT, CAPE, TS, UAG)	9/30/15	9/30/15	Verification results presented in Year 3 report and during Year 3 site visit.	Complete: Task 9 completed for 2015 deliverables, including verification results and performance evaluation of parallel TS3ph. No-cost-extension allowed team to focus on validation studies with utility partners and parallel performance improvements for TS3ph.
Task 10 - Develop a final project report documenting key findings and operational improvements	5/30/17	5/30/17	Submission of final report.	Complete: Final project report submitted. Faster-than-real-time dynamics simulation achieved for utility case.

Accomplishment details

The research prototype dynamics simulator was created as a research platform for exploring faster-than-real-time dynamics simulation, unbalanced three-phase dynamics simulation, combined dynamics and protection simulation, and single-phase induction motor models with stall capability. Each of these four application areas will be described in this section. Although a prototype simulator was developed, the simulator itself was not a deliverable for this project. Rather, the lessons learned and the technical approaches (some successful, some unsuccessful) are the valuable outcomes. Most of the lessons learned have been disseminated through major conferences and publications. Meanwhile, this report contains the four major success stories that were direct outcomes of the project.

Success Story #1 – Faster than real-time dynamics simulation
One of the main goals of the project was to compute the post-disturbance response of an actual utility dynamics case in less time than the amount of simulated response time. For example, suppose a simulation were configured to determine 10 seconds of dynamic response to a three-phase fault in a large-scale power system model. A typical transmission fault is cleared in 0.1 seconds, but a well-behaved transient response might last for several seconds following the fault. A 10 second simulation would be more than sufficient to determine if the response were stable or unstable for most normally cleared faults under reasonable operating conditions.

However, the amount of wall-clock time, or the amount of time that the user would have to wait, before the simulation results were complete, would be significantly longer than 10 seconds, assuming serial code execution. A user may need to wait 20 or 30 seconds or more, before the simulation would be complete and the results would be available for analysis. This time delay is acceptable today for offline planning studies, and it would be acceptable even for online operational studies that are performed in a control center based on real-time operating conditions.

On the other hand, if it takes 20 or 30 wall-clock seconds, i.e., the amount of time that passes according to the clock on the wall, to determine a 10 second simulated response to a disturbance, then operators could not use the simulation results to help them make a difficult decision during a cascading outage. In other words, the physical power grid could collapse long before the simulation results were complete.

During a cascading outage, the power system dynamics can be unusually fast and furious. Based on today's large-scale energy management systems, there is no way to simulate the response of the physical power grid FASTER than the actual post-disturbance dynamics occur. Hence, the project team sought to simulate the post-disturbance response of a large-scale power system faster-than-real-time.

With a faster-than-real-time simulator, an operator would have a view into the future to see the impact of an ongoing disturbance, BEFORE the post-disturbance response had run its course. Now it's important to understand that "one second faster than real-time" is not very useful. But, if a simulator can run faster-than-real-time, then the longer the simulation, the larger the window into the future becomes. Also, it's important to realize that the speedup of a faster-than-real-time simulator depends on the model size (how many buses), the model complexity (how many control devices), and the severity of the disturbance (how widespread is the impact of the disturbance).

The original project goal was to complete a 10 second simulation of a three-phase unbalanced model for a large-scale utility case in less than 5 seconds. This 2x faster than real time speedup would allow an operator to see 15 seconds into the future after just 15 seconds of wall-clock time. In other words, after 15 seconds passed on the wall-clock, the operator could see the results of a 30 second simulation.

The high performance computing cluster used for the faster-than-real-time simulator had the following major components:

- Three Supermicro SuperStorage Server 6028R-E1CR12L nodes
- InfiniBand EDR networking fabric
- Each node had two Intel Xeon E5-2667V4 / 3.2 GHz 8 core processors
- Each node had 128 GB of 2400MHz DDR4 RAM

The utility dynamics case had the following characteristics:

- 18,946 buses
- 1,263 detailed generators with 6 dynamic states (Eq', Ed', φd, φq, Δω, δ)
- 1,018 IEEE DC1A Exciters with 5 dynamic states (Vc, Vll, Vr, Efd, Vf)
- 12,097 loads

All governors and stabilizers were intentionally removed, such that a 12 cycle duration three-phase fault (from t=0.1 to t=0.3 s) would lead to unstable behavior after approximately 15 seconds of simulated time. This enabled the team to test the simulator on both stable dynamics, which occurred for roughly the first 15 seconds of the simulation, and unstable dynamics, which occurred for roughly the last 15 seconds.

Although the team made significant progress in decreasing the amount of wall-clock time needed to complete the 10 second simulation, the 5 second target was not met for the three-phase model. Figure 1 shows the amount of wall-clock time in seconds on the vertical axis for a series of simulation runs with an increasing number of processing cores, which are shown on the horizontal axis.

Modeling the power system as a three-phase unbalanced network introduces three times as many network variables (Va, Vb, Vc) as compared to a positive-sequence network (V1). In addition, the positive-sequence network current mismatch equation (I1) becomes three current mismatch equations (Ia, Ib, Ic). Furthermore, each complex admittance element of the positive-sequence network becomes nine complex elements in the three-phase bus admittance matrix. Based on empirical studies, this led to a 5x larger computational burden for three-phase unbalanced modeling compared to the standard practice of positive-sequence modeling for electromechanical transient analysis.

Although the team sought to demonstrate faster-than-real-time simulation results for the three-phase representation of the 18,946 bus utility case, it became clear that the model was too large.

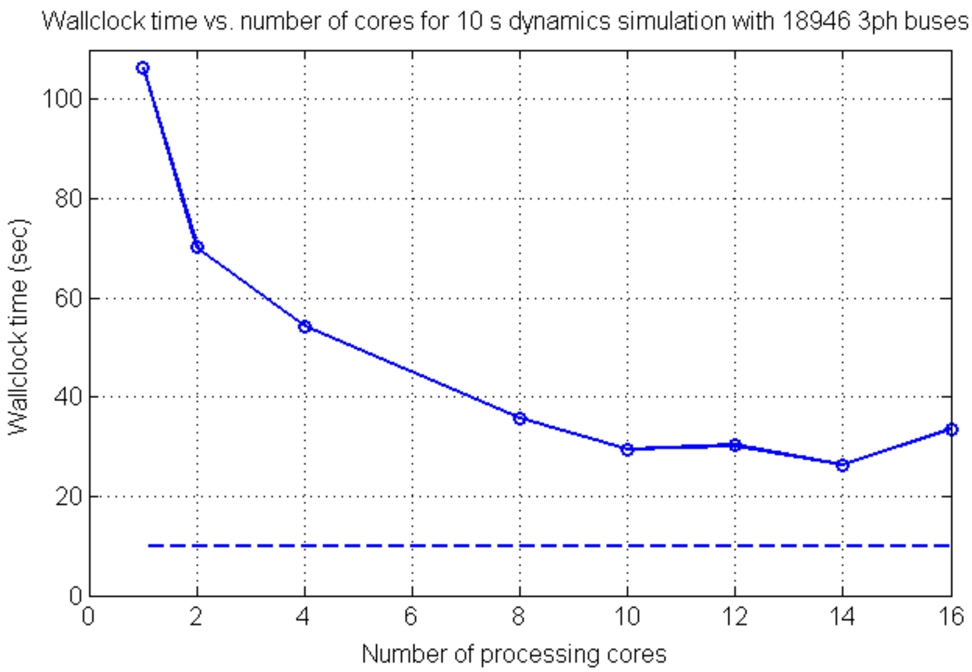

Figure 1: Scalability plot for three-phase network representation of 18,946 bus utility case

At the end of the project, the simulator core modeling functions were rewritten to implement a positive-sequence model for the faster-than-real-time simulator. Figure 2 shows the amount of wall-clock time in seconds on the vertical axis for a series of simulation runs with an increasing number of processing cores, which are shown on the horizontal axis.

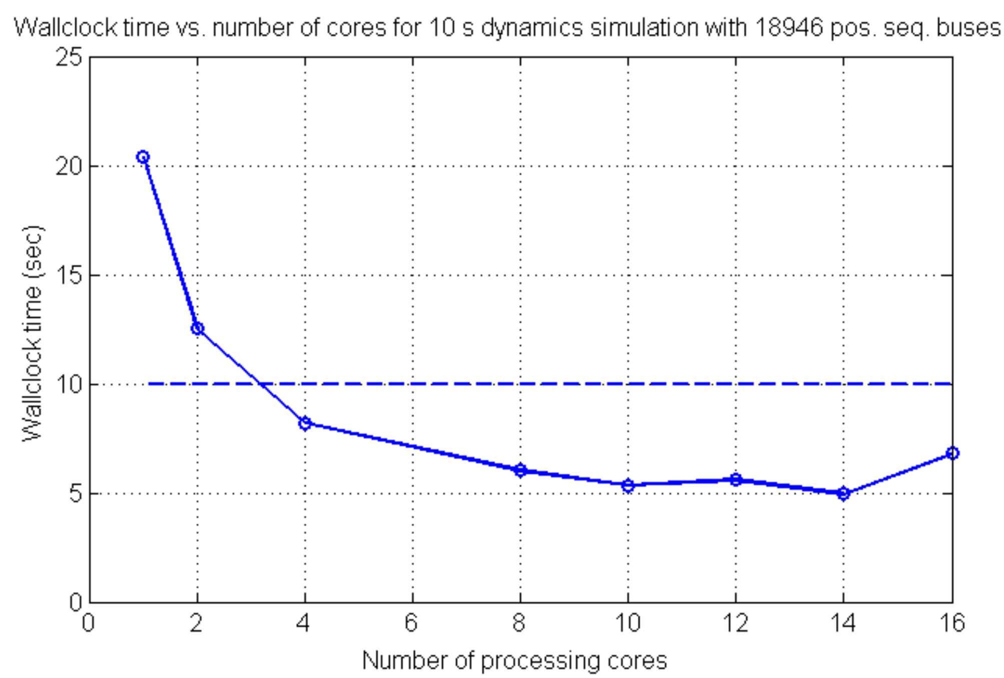

Figure 2: Scalability plot for positive-sequence network representation of 18,946 bus utility case

The success of the faster-than-real-time simulator is largely due to the PETSc toolkit, which is developed by Argonne National Laboratory. The following options were employed via PETSc:

- Integration Scheme: Implicit Trapezoidal (Theta method)
- Time Step: Fixed time step of ½ cycle (0.008333 s)
- Linear solver: GMRES
- Preconditioner: Additive Schwarz Method
- Overlap: 4
- Partitioning routine: ParMETIS

In addition, PETSc allows the user to lag the computation of the preconditioner. The results above were obtained by computing the preconditioner once at the beginning of the simulation (t=0), once at the beginning of the fault-on period (t=0.1s), and once just after the fault was cleared (t=0.3s), i.e., at the beginning of each new network topology period. Since the preconditioner matrix must be factored before being used in the GMRES method, a reduction in the number of factorizations can yield a large improvement in speed. However, if the preconditioner matrix becomes too far away, in an eigenvalue sense, from the Newton update system matrix, then the preconditioner will be less useful and the GMRES method will require more iterations. This tradeoff can be difficult to balance, but the above strategy worked well.

The computational times appear below for the positive-sequence representation of the 18,946 bus case for a 10 second simulation:

# of cores	Wall-clock (sec)
1	20.363
2	12.515
4	8.205
8	6.029
10	5.345
12	5.607
14	4.952
16	6.808

Note that the impressive 2x faster-than-real-time speedup was nearly sustained for a 20 second simulation, in which the same case and the same three-phase fault were simulated. The wall-clock time required was 11.910 seconds. The reason that the 20 second simulation requires more than twice the wall-clock time of the 10 second simulation is due to the unstable dynamics of the case, as discussed below. For unstable dynamics, the 1.68x faster-than-real-time speedup is still impressive.

However, the unstable dynamics (shown in Figure 3) completely erase all of the speedup gains during a 30 second simulation. The wall-clock time for a 30 second simulation of the same case with the same fault was 38.750 seconds. The reason for the slowdown comes directly from the increase in the number of iterations required to find a solution at each time step.

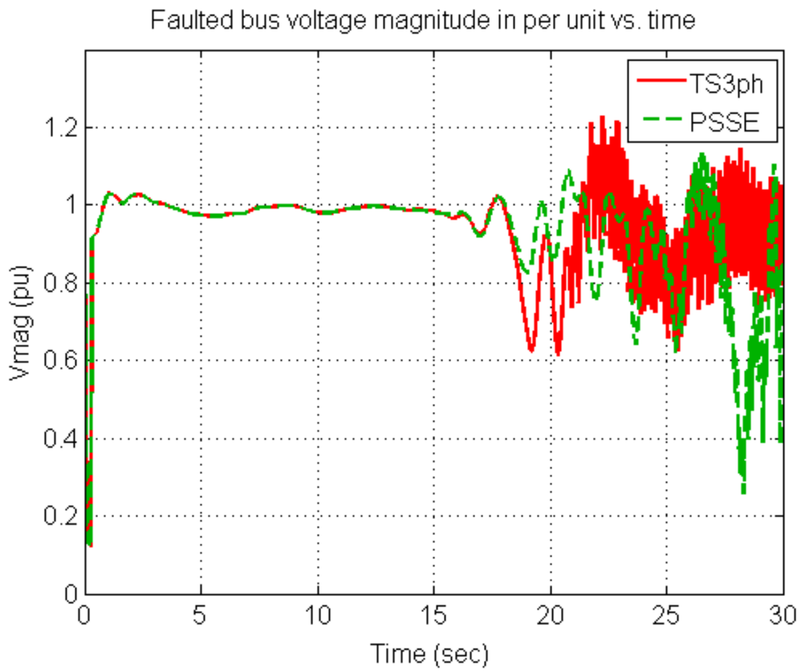

Figure 3: Faulted bus voltage magnitude vs. time in benchmark comparison: 0-30 sec

The benchmark comparison between the research prototype simulator, called TS3ph, and Siemens PTI PSSE illustrates the unstable dynamics at the faulted bus after roughly 15 seconds. The two simulators do not agree on the details of the unstable behavior, but both simulators produce high frequency oscillations in bus voltage magnitude during the second half of the simulation. The reason for this discrepancy will be explored through the next series of plots, but for now it should be clear that the unstable behavior is present in both simulations and that the unstable behavior is concentrated in the second half of the 30 second simulation.

Note that the step size in TS3ph is 1/2 cycle, while the step size in PSSE is 1/16 cycle. The smaller step size in PSSE was necessary to capture the unstable dynamics. Since PSSE uses an explicit integration method, it is more sensitive to step size than an implicit integration method.

To achieve faster-than-real-time dynamics simulation, the number of iterations required to find a solution to a large set of nonlinear equations must remain small. If the simulated dynamics become wild, then the number of iterations can easily double, if not triple, which is precisely the reason for the slower-than-real-time dynamic simulation for the 30 second simulation.

Many researchers have investigated variable time-stepping methods to speed up general nonlinear dynamics simulation. Those methods have great promise when the dynamics are stable. In fact, the Argonne team demonstrated 5~6 times faster simulation speeds using a variable step Rosenbrock integration scheme, compared to a fixed time step trapezoidal method. However, for an unstable scenario, all of the variable step schemes were SLOWER than the fixed time step trapezoidal method. Since the goal of the project is to help operators during extreme conditions, such as a cascading outage, the team decided to use a fixed step method.

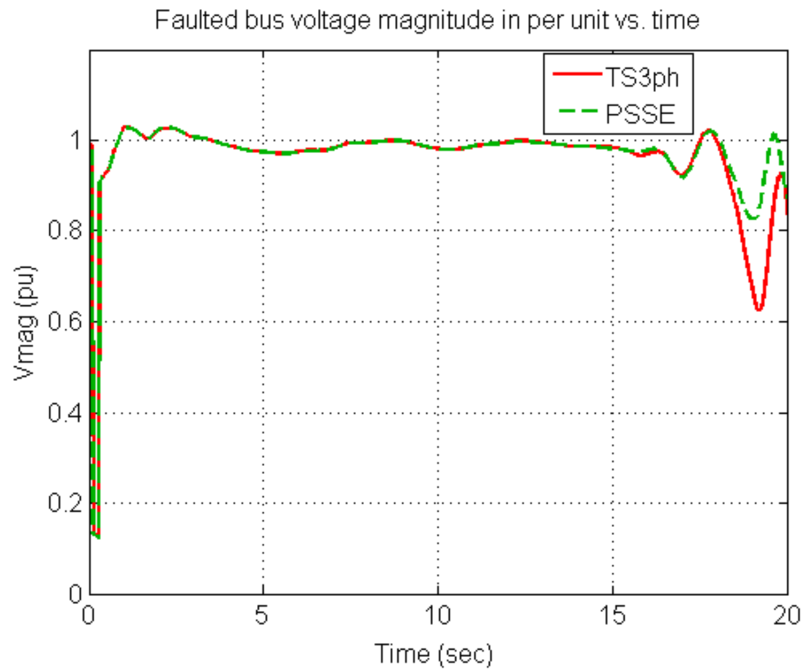

Figure 4: Faulted bus voltage magnitude vs. time in benchmark comparison: 0-20 sec

During the first 20 seconds of the 30 second simulation, the faulted bus voltage magnitude recovers nicely, but then begins to oscillate wildly after roughly 15 seconds. Figure 4 shows the benchmark comparison for the first 20 seconds. The two simulators appear to diverge after 17 seconds for reasons that will become clear later. Figure 5 zooms into the first 4 seconds.

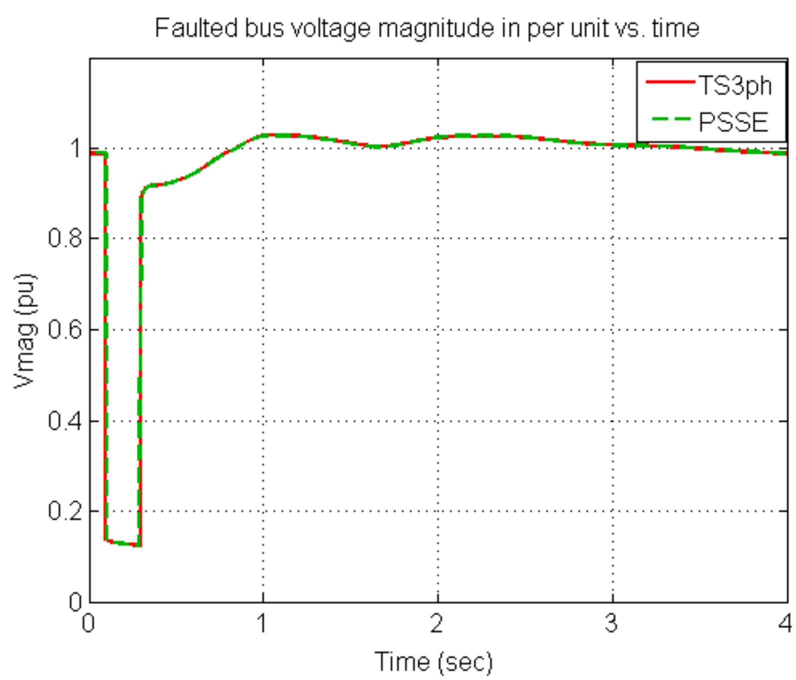

Figure 5: Faulted bus voltage magnitude vs. time in benchmark comparison: 0-4 sec

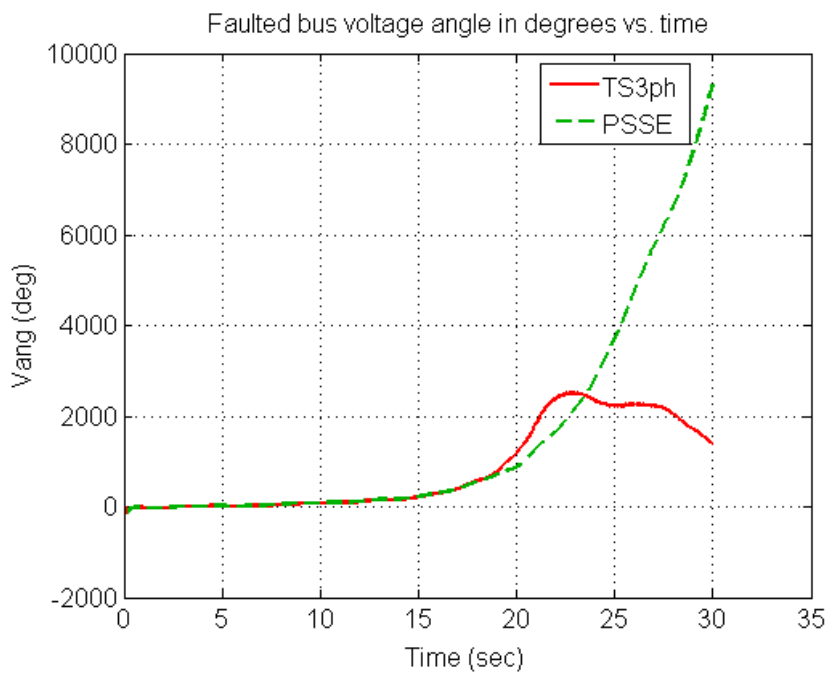

Figure 6: Faulted bus voltage angle vs. time in benchmark comparison: 0-30 sec

The faulted bus angle is shown in Figures 6 and 7 for 30 seconds and 20 seconds respectively. Again, the two simulators diverge after roughly 17 seconds.

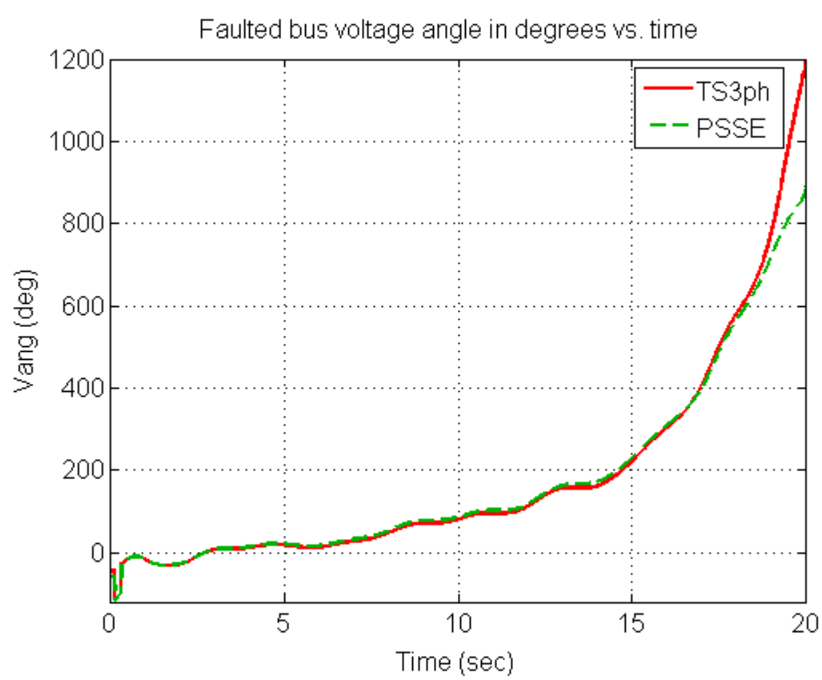

Figure 7: Faulted bus voltage angle vs. time in benchmark comparison: 0-20 sec

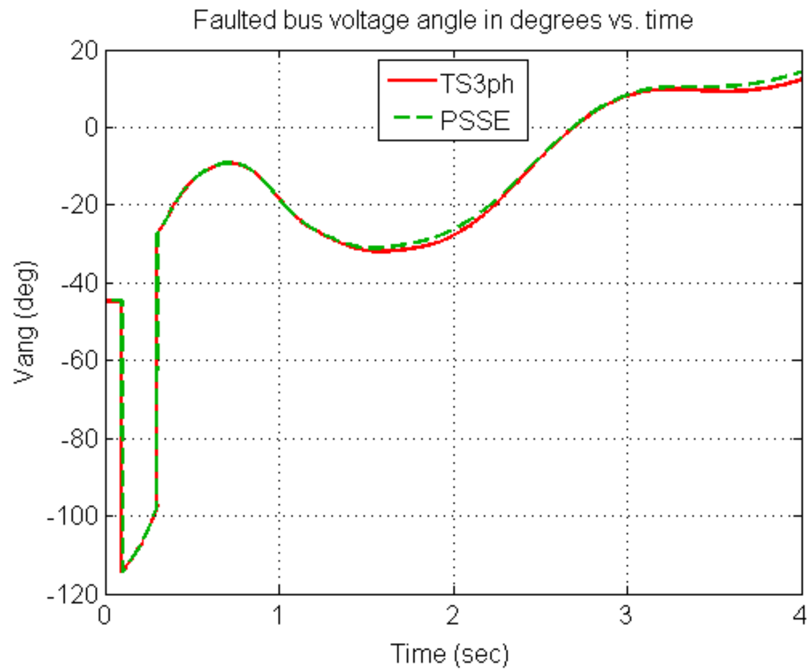

Figure 8: Faulted bus voltage angle vs. time in benchmark comparison: 0-4 sec

The first four seconds of the 30 second simulation yield very similar results from the two simulators, as shown in Figure 8. Going back to the full 30 second simulation in Figure 9, we get another view of the clearly unstable dynamics from a generator near the fault.

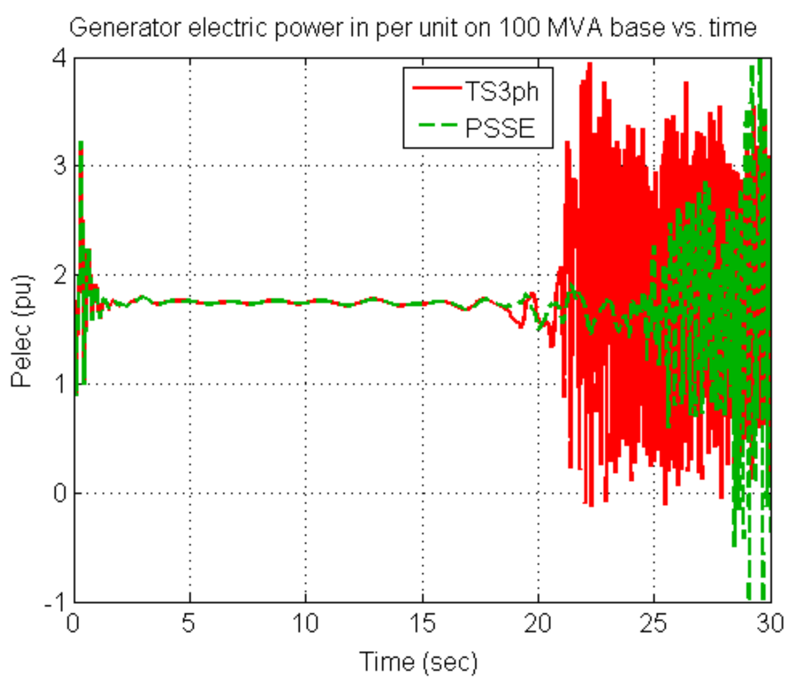

Figure 9: Nearby generator real power vs. time in benchmark comparison: 0-30 sec

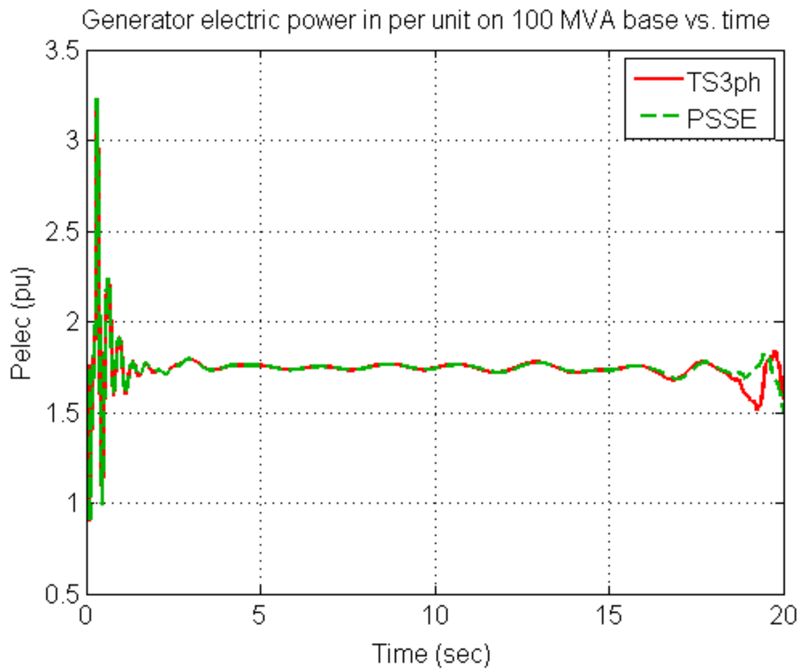

Figure 10: Nearby generator real power vs. time in benchmark comparison: 0-20 sec

Within the first 17 seconds of the 30 second simulation, the two simulators agree on the generator output, as shown in Figures 10 and 11.

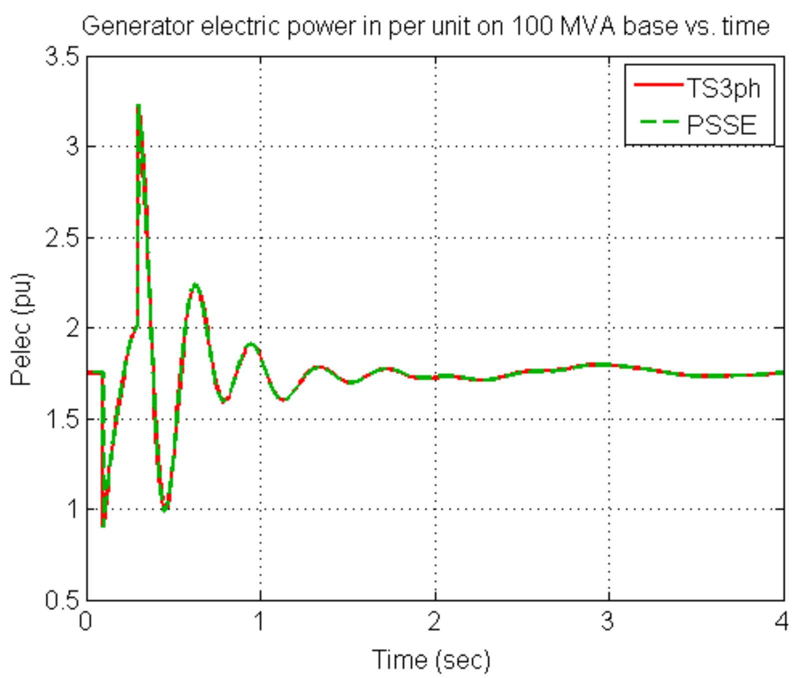

Figure 11: Nearby generator real power vs. time in benchmark comparison: 0-4 sec

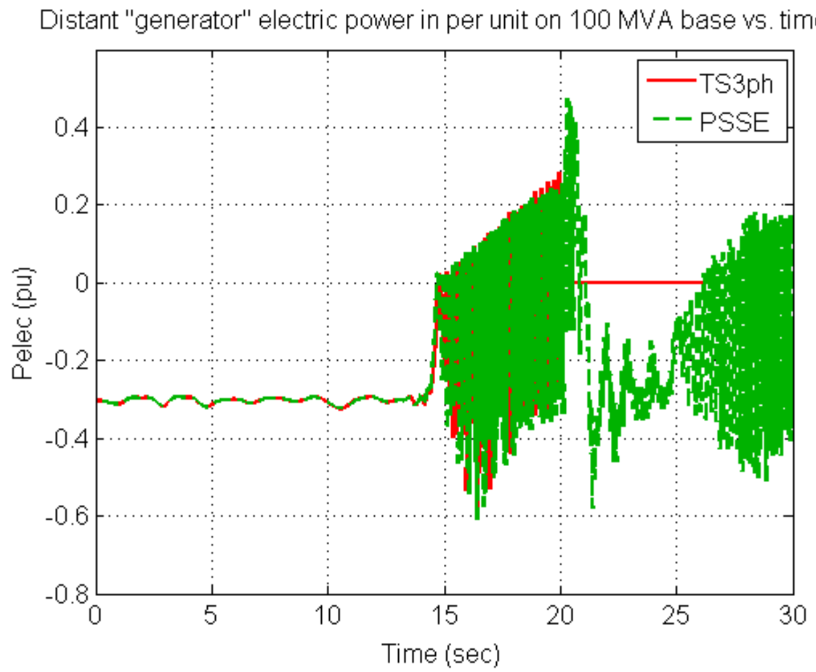

Figure 12: Distant “generator” real power vs. time in benchmark comparison: 0-30 sec

To better understand the discrepancy in the previous plots, a large 30 MW motor load is shown in Figure 12. Notice the small ripple from the very beginning of the simulation, even before the fault. This motor is several hundred miles away from the fault location, yet it experiences wildly unstable dynamics just before t=15 s. Figure 13 zooms in to see more detail.

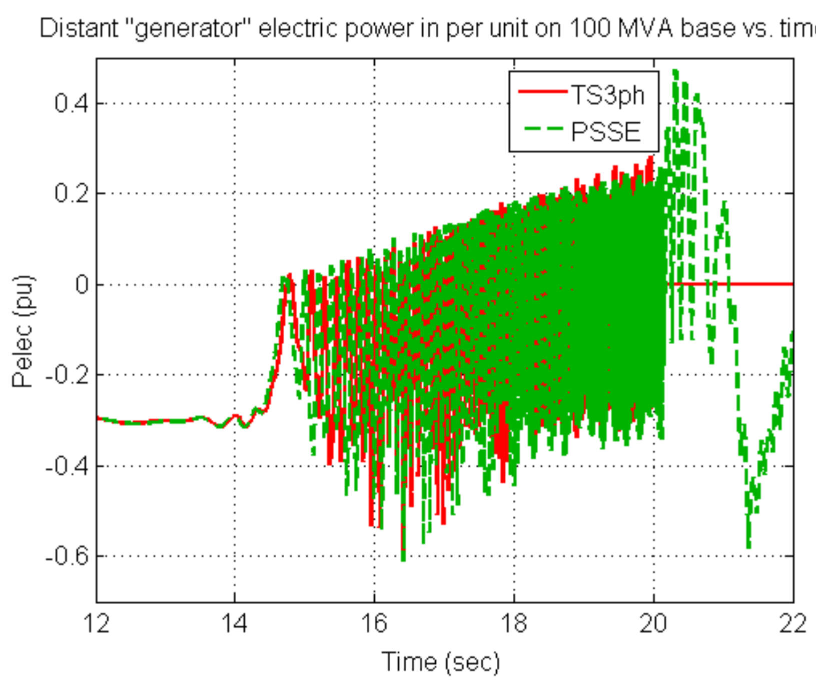

Figure 13: Distant “generator” real power vs. time in benchmark comparison: 12-22 sec

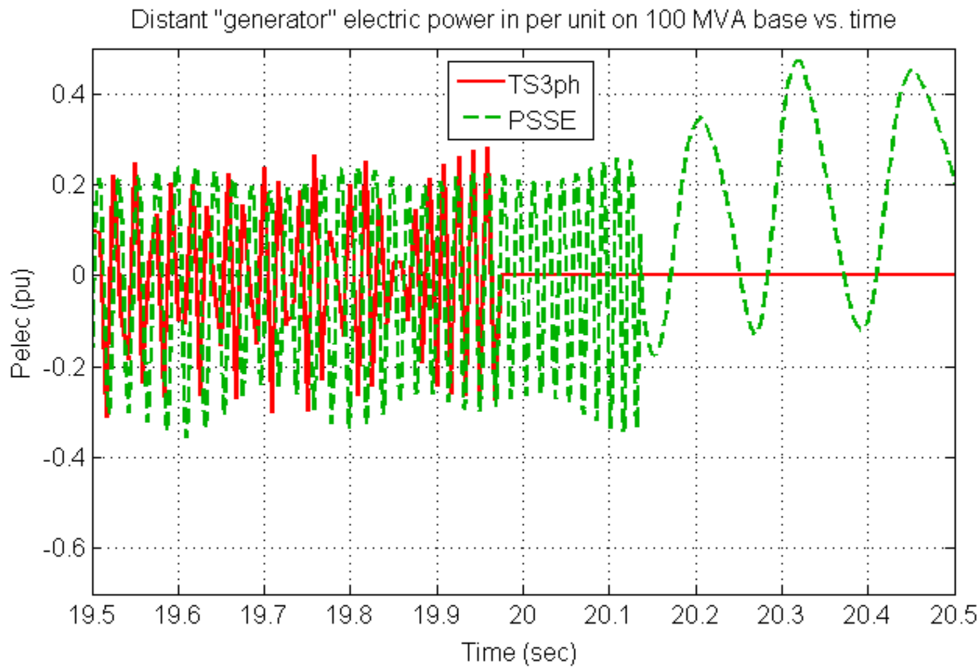

Figure 14: Distant “generator” real power vs. time in benchmark comparison: 19.5-20.5 sec

The large motor load was initially absorbing 0.3 pu, i.e., Pelec was -0.3 pu, before the unstable dynamics. Figure 14 shows the tripping action in the TS3ph simulation just before t=20s, which will be explained by a later plot. Figure 15 shows the two simulators diverging around t=14.5s.

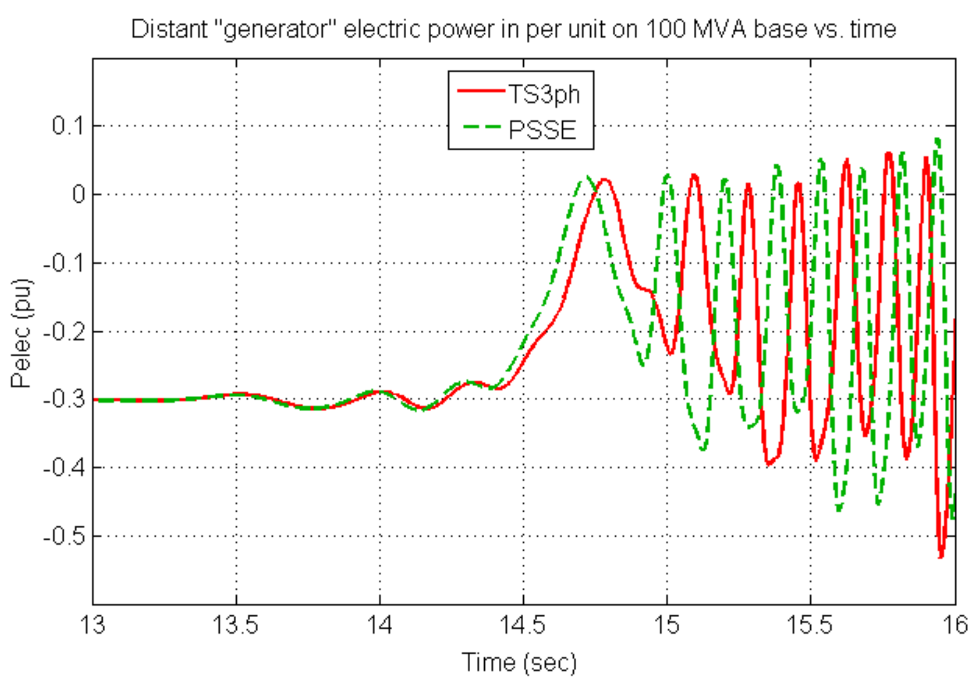

Figure 15: Distant “generator” real power vs. time in benchmark comparison: 13-16 sec

Figure 16: Distant “generator” real power vs. time in benchmark comparison: 0-10 sec

A noticeable ripple can be seen in the large motor power consumption before the fault in Figure 16. Figure 17 shows the large motor’s per unit speed deviation, which explains the tripping action seen in Figure 14. When the speed deviation reaches -1, the motor has stalled.

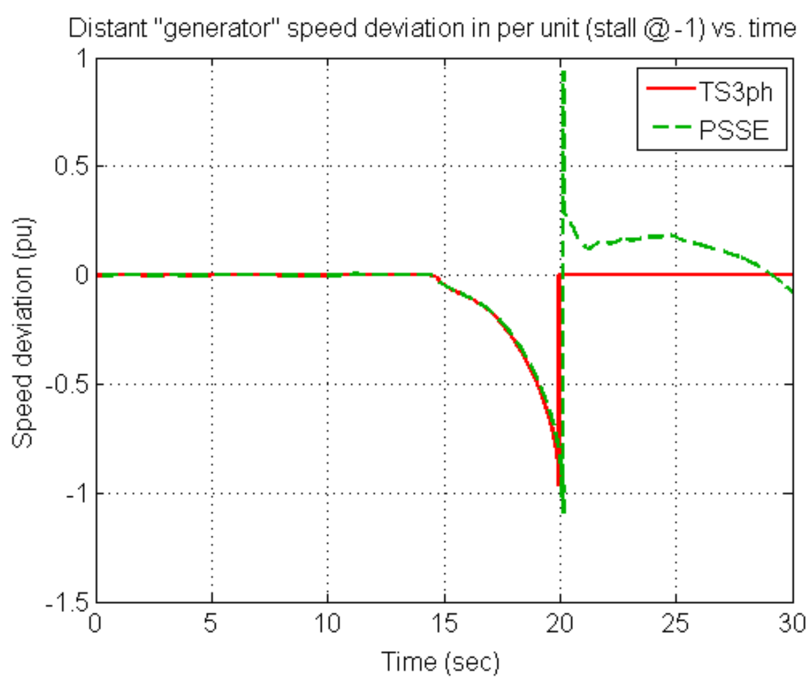

Figure 17: Distant “generator” speed deviation vs. time in benchmark comparison: 0-30 sec

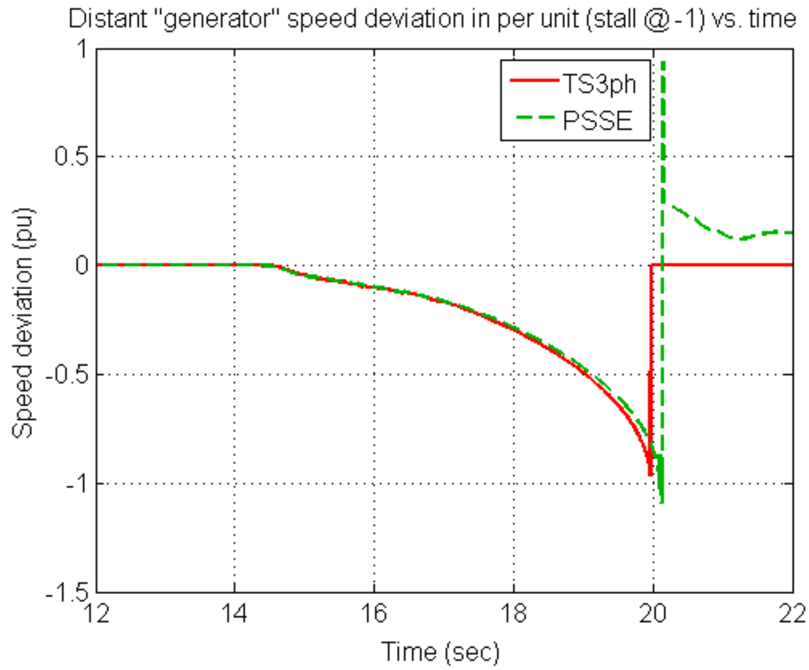

Figure 18: Distant "generator" speed deviation vs. time in benchmark comparison: 12-22 sec

A 30 MW motor would have protection to trip it long before it stalled, but for this case, protection was not included. To prevent numerical instability, TS3ph trips the motor just before t=20s. PSSE does not trip the motor, although protective tripping can be included in PSSE.

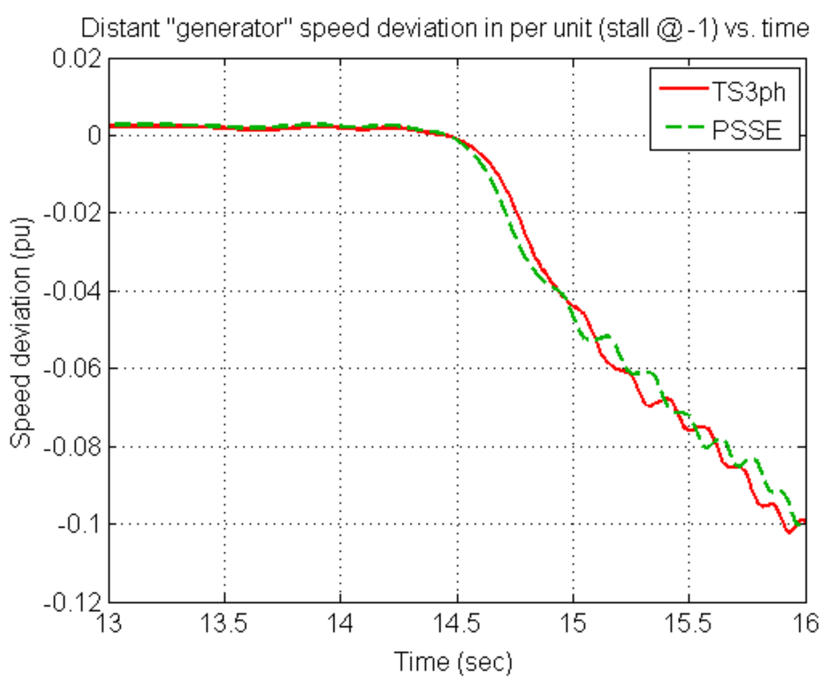

Figure 19: Distant "generator" speed deviation vs. time in benchmark comparison: 13-16 sec

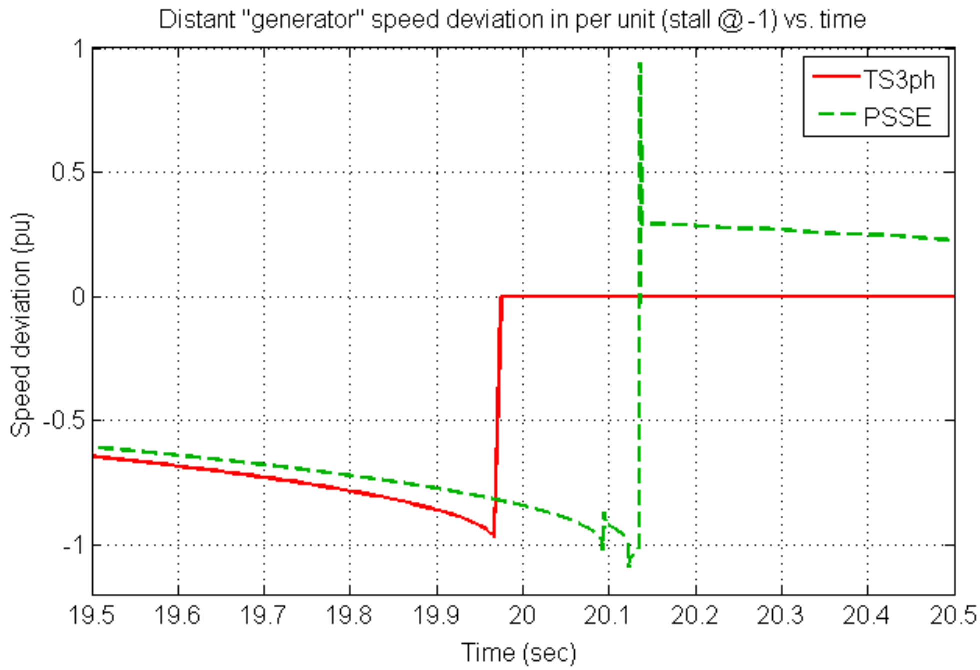

Figure 20: Distant “generator” speed deviation vs. time in benchmark comparison: 19.5-20.5 sec

Figure 19 zooms in to show the two simulators diverging around t=14.5s, although the discrepancy is not large. Figure 20 shows the time (just before t=20s) when the discrepancy becomes large. TS3ph trips the motor, since it has stalled, but the PSSE simulation continues with the motor jumping suddenly from a nearly stalled condition to a new condition in which the motor is now rotating faster than its nominal speed. This sudden jump is a numerical artifact that does not represent actual large motor behavior. Therefore, the PSSE simulation results can no longer be used beyond t=20s. To handle this situation properly, the user would need to trip the motor in PSSE, after determining the speed deviation was near -1.0 per unit.

However, referring back to Figure 13, it is clear that the electric power absorbed by the large motor is no longer stable at t=14s. Therefore, both TS3ph and PSSE are no longer simulating realistic dynamic behavior for the second half of the simulation. A realistic dynamic simulation would require protective devices that would trip generators, motors and other loads that experienced wild oscillations, such as those shown in Figure 13 starting around t=14s.

In conclusion, the prototype dynamics simulator achieved an impressive 2x faster-than-real-time speedup during stable dynamics for an 18,946 bus case with 1,263 detailed generators and 1,018 DC1A exciters. In addition, the faster-than-real-time simulator achieved a 1.68 speedup during a mix of stable and unstable dynamics for a 20 second simulation. Finally, to run faster-than-real-time is not enough. The simulation results need to be accurate. As shown in the above comparisons, the benchmarking process verified the accuracy of the faster-than-real-time simulator.

Success Story #2 – Unbalanced three-phase network modeling
There are many advantages to unbalanced three-phase network models as compared to traditional positive-sequence models:

1. Direct knowledge of system imbalances based on availability of all three complex phase voltages.
2. Direct knowledge of the impact of unbalanced disturbances on all three complex phase voltages.
3. Ability to model effects of untransposed lines.
4. Ability to model impacts of independent phase controls.
5. Ability to interface directly to a detailed substation model and an unbalanced three-phase distribution system.

As we saw in the previous section, unbalanced three-phase network models require more computation time. Therefore, the unbalanced three-phase network models should be used when they provide additional valuable information beyond a traditional positive-sequence network model. Conversely, if the unbalanced three-phase network model does not provide additional valuable information, then the traditional positive-sequence network model should be used.

How can an engineer determine a-priori when to use an unbalanced three-phase network model and when to use a traditional positive-sequence model? The answer appears in the previous paragraph. Use the more detailed model when it provides useful information. However, most transmission planning engineers will not have experience with unbalanced three-phase network modeling, unless they have worked in the areas of electromagnetic transients (EMT), system protection, or distribution. Without the prior experience, it would be good practice to run some simulations with an unbalanced three-phase network model to gain experience with the unbalanced behaviors.

As an illustration of this point, consider a static var compensator (SVC) that has independently controlled phases. For example, some manufacturers are preparing to provide SVC controls that will boost a single-phase-to-ground voltage during a single-phase-to-ground fault, while either leaving the other two phases with zero shunt injection, or providing a negative shunt injection with a reactor to bring down the typically high voltage magnitudes on the un-faulted phases.

The SVC model in this section will be based on the SVSMO2U2 model in PSSE. This model was required by one of the utility partners. First, the research team implemented a similar SVC model in TS3ph and then benchmarked the simulation results with PSSE. Next, the research team modified the SVC model in TS3ph to investigate the impacts of independently controlled shunt phases.

For the SVSMO2U2 SVC model, the team discovered that PSSE was noticeably sensitive to the length of the time step. For a half-cycle time step of 0.00833 seconds, the PSSE simulator gives significantly different results than for a 0.001 second time step. Figure 21 compares TS3ph and PSSE following a three-phase fault with a 0.00833 second time step. In this case, TS3ph is unable to match PSSE. However, in Figure 22, TS3ph and PSSE are shown with a 0.001 second time step for the same fault. In this case, the TS3ph results are the same as the TS3ph results with the previous 0.00833 sec time step, but now TS3ph and PSSE match.

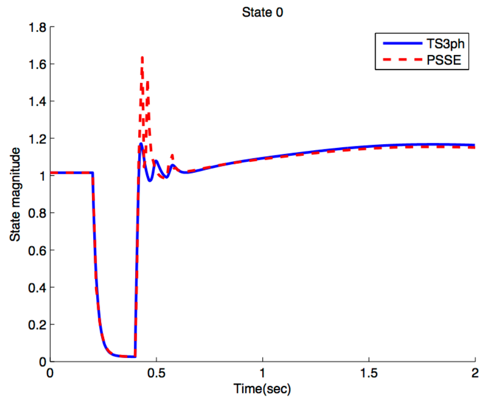

(a) Controlled voltage sensor

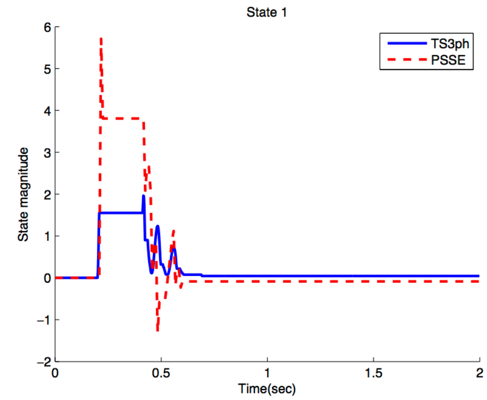

(b) SVC main PI controller integrator

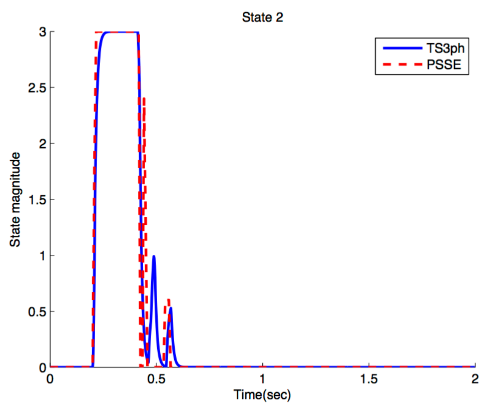

(c) SVC output lag

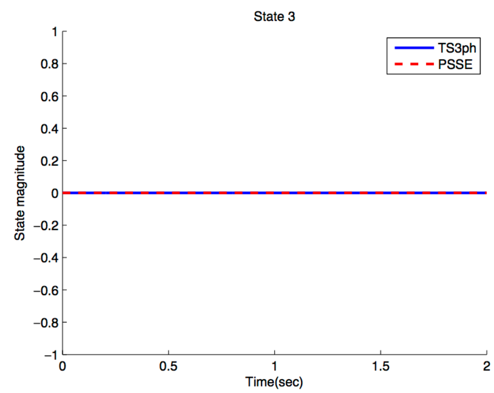

(d) SSC PI controller integrator

Figure 21: SVSMO2U2 states vs. time with 0.00833 second time step (TS3ph: solid blue, PSSE: dashed red)

For a single-line-to-ground fault, TS3ph can simulate the response of the individual phases ABC to the unbalanced fault, but PSSE cannot, due to its positive-sequence-only simulation framework. However, CAPE TS-Link can be used to study unbalanced fault conditions by using CAPE to model the unbalanced fault, as well as the network-equivalent negative sequence and zero sequence impedances. Then, CAPE can use its TS-Link module to combine the network equivalents with the unbalanced fault impedance and provide an equivalent impedance to PSSE for a positive-sequence dynamics simulation. As PSSE simulates the positive-sequence response, CAPE can translate the positive sequence information into phase ABC information to recreate the individual phase voltages and phase currents. In Figures 23-25, a phase A fault to ground is compared in TS3ph-CAPE and CAPE-TSLink. Figure 23 shows the phase A response. Figure 24 shows the phase B response. Figure 25 shows the phase C response.

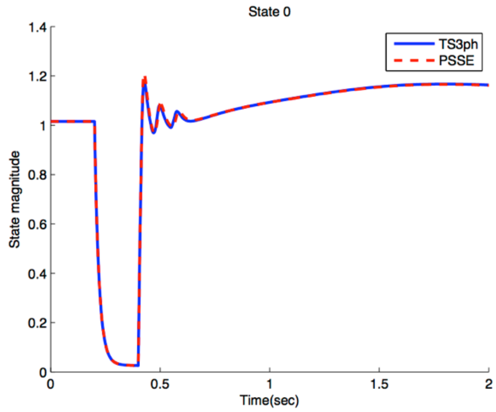

(a) Controlled voltage sensor

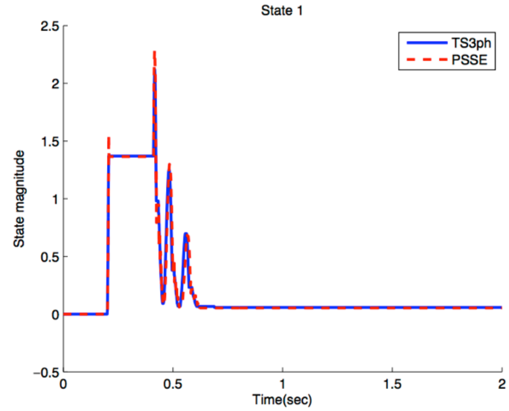

(b) SVC main PI controller integrator

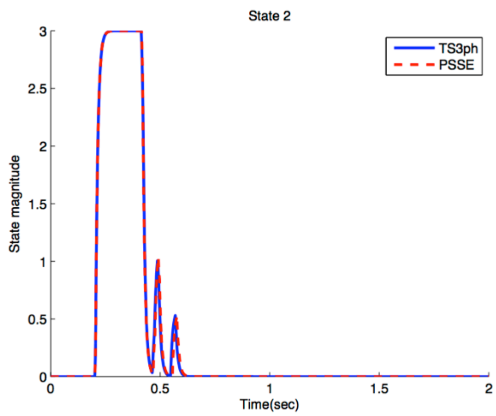

(c) SVC output lag

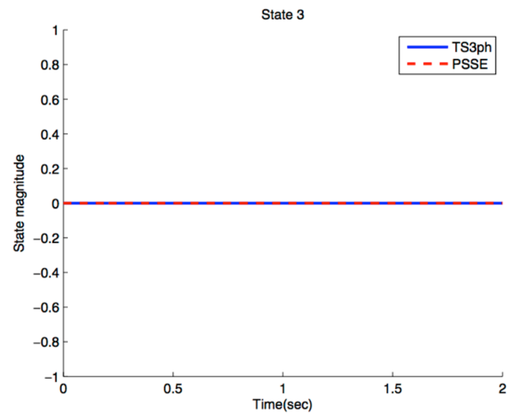

(d) SSC PI controller integrator

Figure 22: SVSMO2U2 states vs. time with 0.001 second time step (TS3ph: solid blue, PSSE: dashed red)

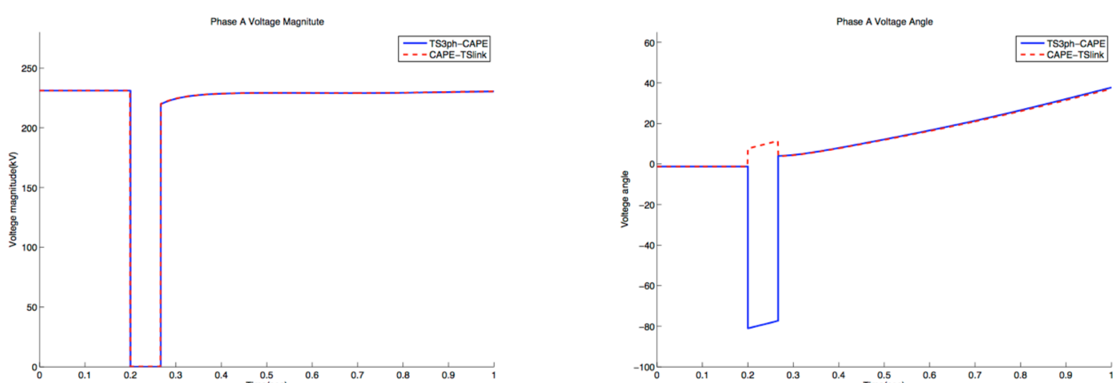

Figure 23: Phase A fault to ground: Phase A voltage magnitude and angle vs. time (TS3ph: solid blue, CAPE-TSLink: dashed red)

Note that CAPE-TSLink shows an insignificant error in the phase A voltage angle during the fault. The error is insignificant since the phase A voltage magnitude is zero, so a vector with zero magnitude can have an arbitrary voltage angle without introducing any error. This issue occurs in PSSE whenever the fault impedance is zero.

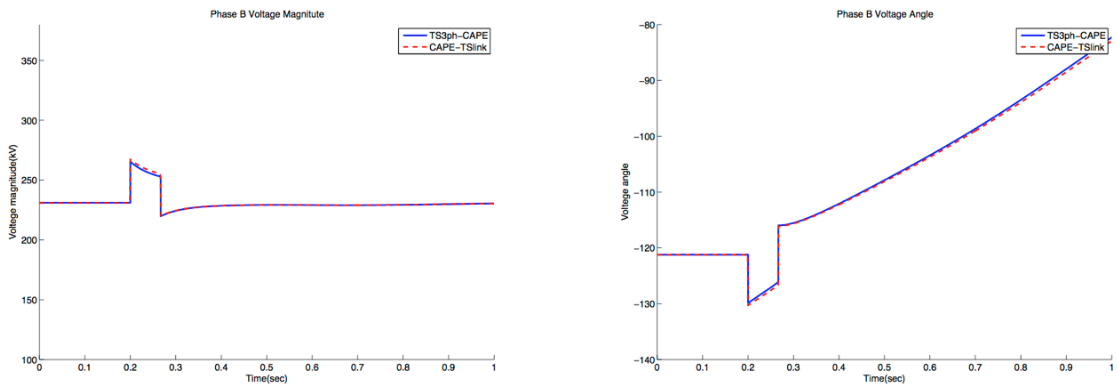

Figure 24: Phase A fault to ground: Phase B voltage magnitude and angle vs. time (TS3ph: solid blue, CAPE-TSLink: dashed red)

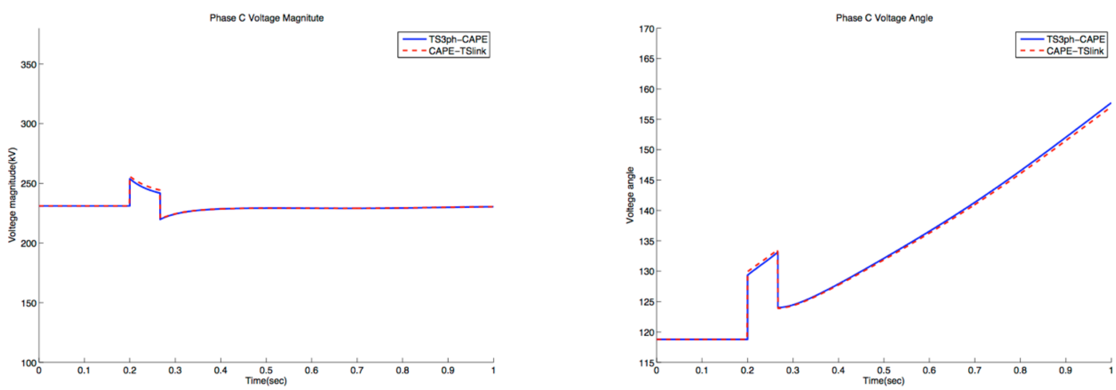

Figure 25: Phase A fault to ground: Phase C voltage magnitude and angle vs. time (TS3ph: solid blue, CAPE-TSLink: dashed red)

Figures 26-28 show the impact of the SVC. For comparison, four curves appear in each figure. There are two TS3ph simulation curves: one for the response without the SVC and one for the response with the SVC. The same is true for the two CAPE-TSLink curves. In Figure 26, without the SVC and with the SVC, the TS3ph results match with the CAPE-TSLink results, except for the insignificant phase A voltage angle error due to the zero phase A voltage magnitude.

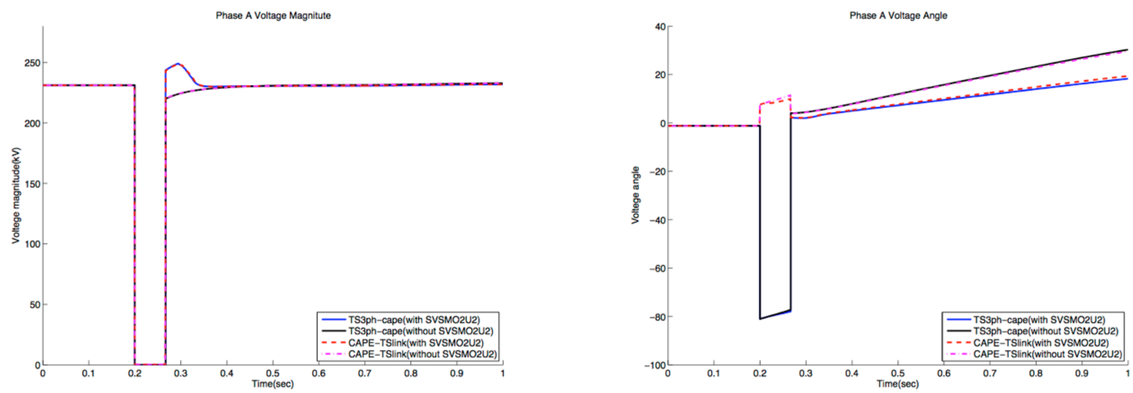

Figure 26: Phase A fault to ground with SVC: Phase A voltage magnitude and angle vs. time (TS3ph: two shades solid blue, CAPE-TSLink: two shades dashed red)

However, in Figures 27 and 28, there is evidence of a slight modeling error in CAPE-TSLink. Note that the phase B and C responses show a slight error in the CAPE-TSLink results due to missing information in the negative and zero sequence networks. In this case, the error is significant since it comes from the assumption that the negative and zero sequence network-equivalents are constant during the entire simulation, which is not true when the SVC inserts a shunt admittance in the network. The shunt admittance is included in the positive-sequence model in PSSE, but it is not represented in the negative and zero sequence network-equivalents in CAPE.

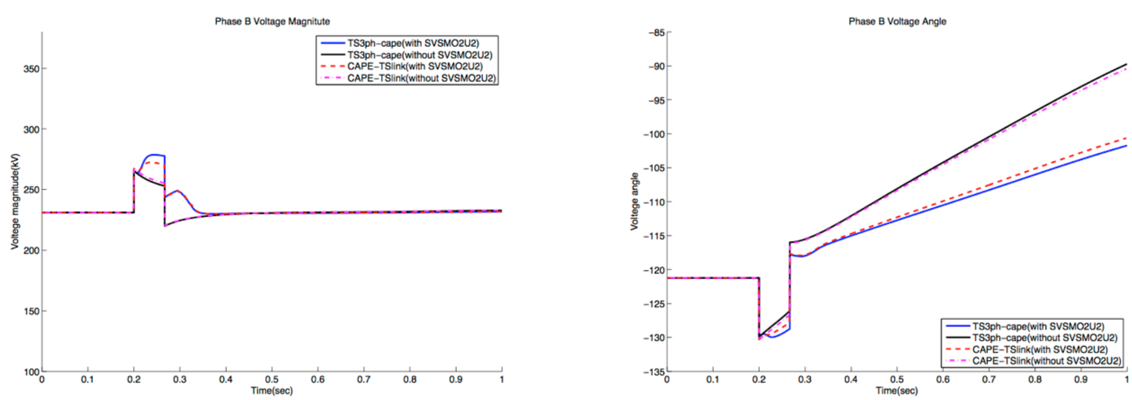

Figure 27: Phase A fault to ground with SVC: Phase B voltage magnitude and angle vs. time (TS3ph: two shades solid blue, CAPE-TSLink: two shades dashed red)

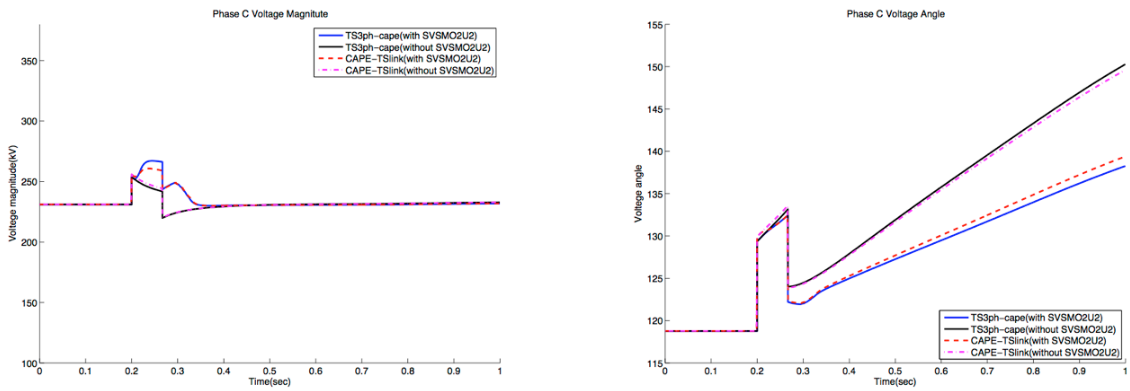

Figure 28: Phase A fault to ground with SVC: Phase C voltage magnitude and angle vs. time (TS3ph: two shades solid blue, CAPE-TSLink: two shades dashed red)

In addition to the three-phase gang-operated SVC controls, the team tested a future control scheme that some vendors are planning to provide in an upcoming SVC product release. During a single-line-to-ground fault, the faulted phase voltage drops close to zero. However, the other two healthy phases experience an increase in their voltage magnitude. With a standard three-phase gang-operated SVC controller, all three phases of the SVC will react to the positive sequence voltage at the bus.

As the positive sequence voltage drops, all three phases will insert shunt capacitance into the network to restore the positive sequence voltage magnitude. However, phases B and C already experience a voltage overshoot during the fault. With the gang-operated controller, the phase B and C voltages will continue to be higher than normal following the fault as the positive sequence voltage recovers to its normal value. The additional overvoltage condition can be avoided with separate controllers operating each phase of the SVC independently.

Figures 29-31 show the results from a modified TS3ph (without the negative and zero sequence SVC components to show how to match CAPE-TSLink) with gang-operated controls (solid blue), CAPE-TSLink with gang-operated controls (dashed red) and TS3ph with independent phase ABC controls (dashed black). Note that the independent phase controls were not tuned for this application, so the phase A overvoltage following the fault could likely be reduced with controller tuning. In Figures 30 and 31, the advantage of the independently controlled SVC phases can be seen in the reduced overvoltage shown by the dashed black line.

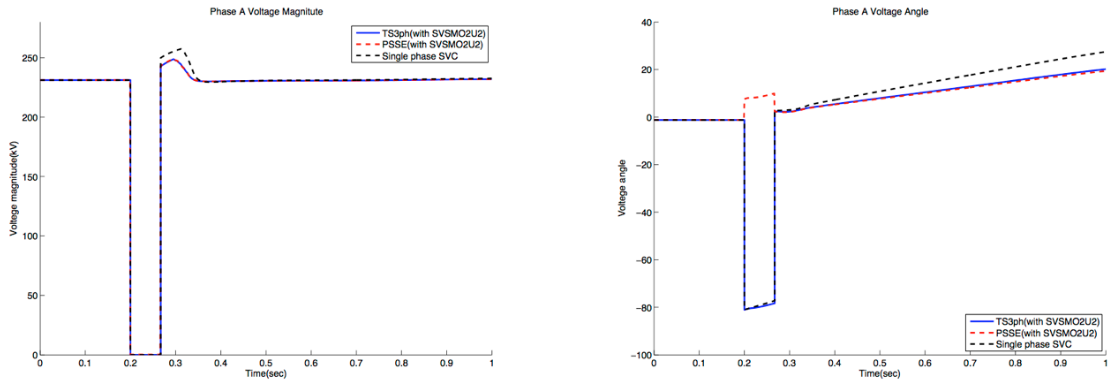

Figure 29: Phase A fault to ground with SVC: Phase A voltage magnitude and angle vs. time (modified TS3ph: solid blue, CAPE-TSLink: dashed red, TS3ph with independent controls: dashed black)

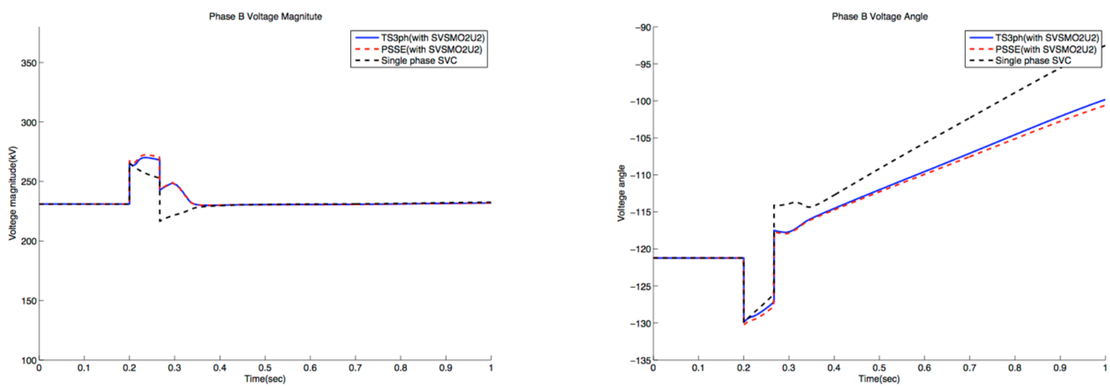

Figure 30: Phase A fault to ground with SVC: Phase B voltage magnitude and angle vs. time (modified TS3ph: solid blue, CAPE-TSLink: dashed red, TS3ph with independent controls: dashed black)

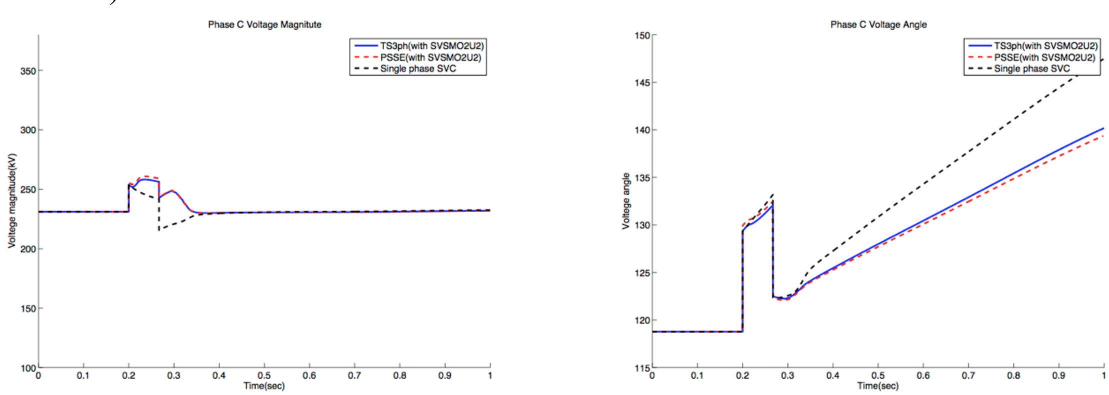

Figure 31: Phase A fault to ground with SVC: Phase C voltage magnitude and angle vs. time (modified TS3ph: solid blue, CAPE-TSLink: dashed red, TS3ph with independent controls: dashed black)

The unbalanced three-phase network modeling is a major breakthrough for transmission dynamics analysis. It enables the investigation of network imbalances, such as those due to untransposed lines or unbalanced faults. It also enables power system engineers to perform model validation on a wider set of recorded disturbances, which are typically unbalanced. Finally, unbalanced three-phase modeling will enable a more accurate simulation of the contributions from and the impacts on future devices with a power electronics interface.

Success Story #3 – Comprehensive dynamics and protection simulation

One of the advantages of using an unbalanced three-phase network model is the ability to interface the transmission model directly with a detailed substation model including the protective devices, such as breakers and their associated relays. Combining the TS3ph dynamics simulator with Electrocon's CAPE protection simulator provided the opportunity to explore the complicated interactions between power system dynamics and power system protection. Examples of dynamics and protection interactions include stable/unstable power swings, load encroachment, under-frequency and under-voltage load shedding, remedial action schemes, and generator protection. Research in the area of dynamics and protection interaction is ongoing.

TS3ph was responsible for determining the voltages, currents and bus frequencies throughout the network. At each time step, the subset of voltages and currents necessary to evaluate the protection system were sent to CAPE through a socket. Then, CAPE would run through the relay logic and decide if any breaker trip signals should be sent back to TS3ph. If CAPE determined that a trip signal should be sent, then TS3ph would receive a breaker operation code and a future operation time. The future operation time accounts for the delay in physically opening the breaker. Figure 32 illustrates the co-simulation architecture.

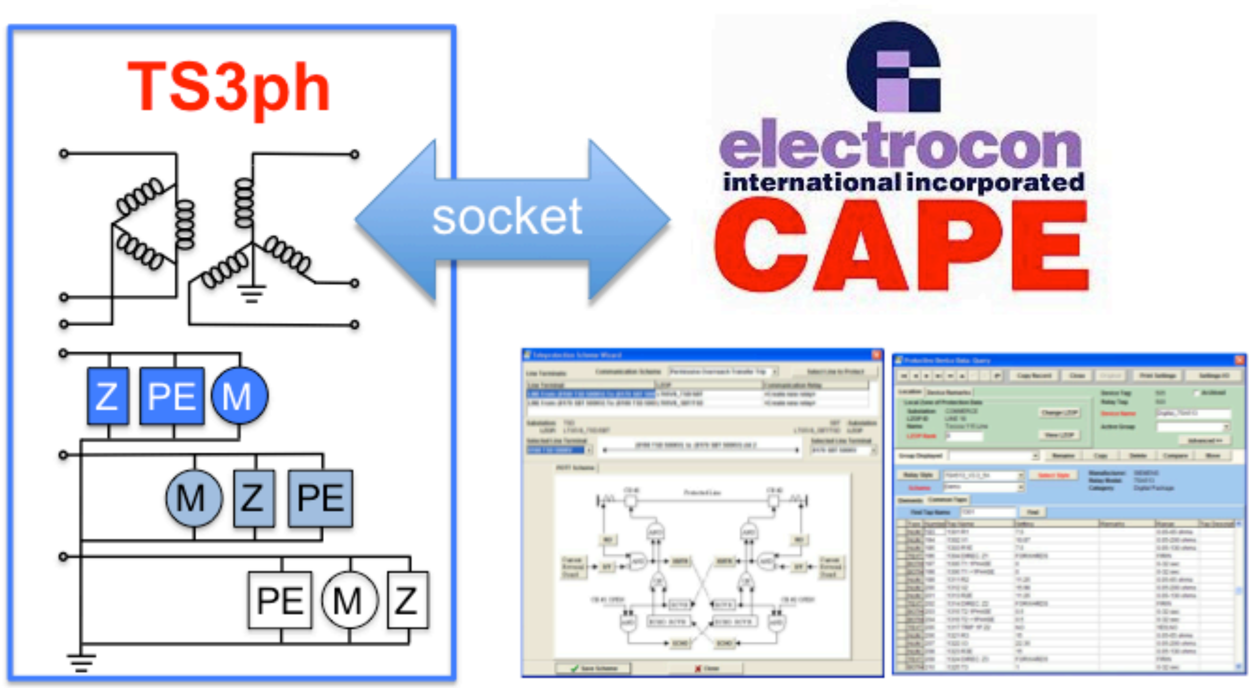

Figure 32: TS3ph-CAPE socket communication for voltages, currents, bus frequencies, and breaker operations

The initial project goals from Electrocon's perspective included the following:

- Link CAPE and TS3ph. CAPE provides the protection model simulation and runs only on Windows OS. TS3ph provides the transient stability simulation and runs only on UNIX platforms.
- Increase the number of relays that can be simulated. CAPE can simulate 5000 relays containing up to 10000 elements, but the goal is to simulate at least 10 times those numbers.
- Decrease the time required to simulate a relay during a time-stepping simulation with voltages and currents supplied by TS3ph. Portions of the simulation that lie outside the time stepping loop are ignored, including building the network model and reading the simulation data from the database.

CAPE and TS3ph continue to execute on their respective operating platforms, and are linked via Berkeley sockets. CAPE runs in server mode, while TS3ph connects as a client. TS3ph tells CAPE what relays to simulate, and in return, CAPE tells TS3ph what voltages and currents are needed at each time step. TS3ph supplies the solution and CAPE returns breaker operations that begin in the current time step, if any. TS3ph processes this information and then continues to the next time step, looping until the disturbance is cleared or the system is determined to be stable.

Previously, the maximum number of objects to simulate in CAPE was determined at compile time. Now, all the simulation storage areas are dynamically allocated and can expand at run-time. The maximum number of relays that can be simulated depends on the size of the simulated network (buses and branches) as well as the type of relay. Simple electro-mechanical relays have only a few elements, but modern digital relays can have tens to over one hundred elements. Studying the complete protection system response for large networks (>50000 buses, with digital relays) will require switching to x64. This task was left for future development.

After implementation of numerous efficiency improvements, the time-stepping portion of the simulation code was improved by a factor of 10 over CAPE_2010 (the starting version for these improvements).

When tested on a 16-core server, the CPU requirements for the time-stepping portion of the simulations were further reduced by a factor of between 5 and 12 to 1, depending on the techniques used. Both multi-threading and multi-processing have been implemented. The ideal solution will use a combination of these two techniques, depending on what system resources are available.

The big advantage to using CAPE for handling the protection system simulation is that CAPE allows protection engineers to store the actual field settings of the relays in a CAPE database. Hence, the protection simulation accuracy improves in two ways: (1) the protection engineers maintain the protection system model, which frees the transmission planning engineers from deciding when to add relays models and when to not add relay models, and (2) the protection engineers maintain the settings, which eliminates the possibility of the transmission planning engineers entering ad-hoc relay models with the wrong characteristics. Of course, the protection engineers also benefit from the ability to study power swings and frequency-related issues.

The TS3ph-CAPE co-simulation framework introduced the world's first high fidelity three-phase unbalanced dynamics and protection simulator. In addition to the comprehensive dynamics and protection system modeling capabilities, the co-simulation framework is ideal for studying cascading failures. Research work in this area is ongoing.

Success Story #4 – Single-phase induction motor modeling with stall capability

In addition to the detailed SVC model, the team implemented and tested a single-phase induction motor model with stall capability, known as IND1. The IND1 model is not available in PSSE, since a single-phase load cannot be completely represented in a positive sequence framework. However, bench tests have been performed on single-phase induction motors that drive residential air-conditioning compressors. Based on work done by Bernie Lesieutre and others, "Load Modeling Transmission Research", *Lawrence Berkeley National Laboratory*, March 2010, the IIT team implemented a single-phase induction motor model to represent the bench-tested air-conditioning loads. Figures 33, 34 and 35 show the single-phase motor response following a three-phase fault. In this case, there are three identical IND1 models connected in a grounded wye configuration, one for each phase, at bus 2.

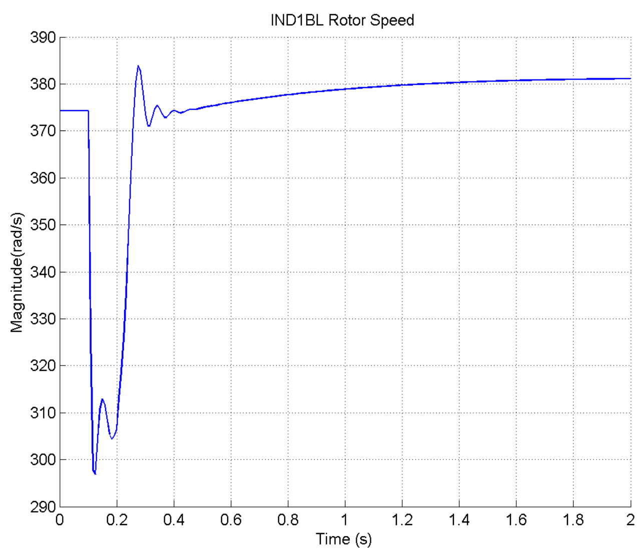

Figure 33: IND1 response to three-phase fault: rotor speed (rad/s) vs. time

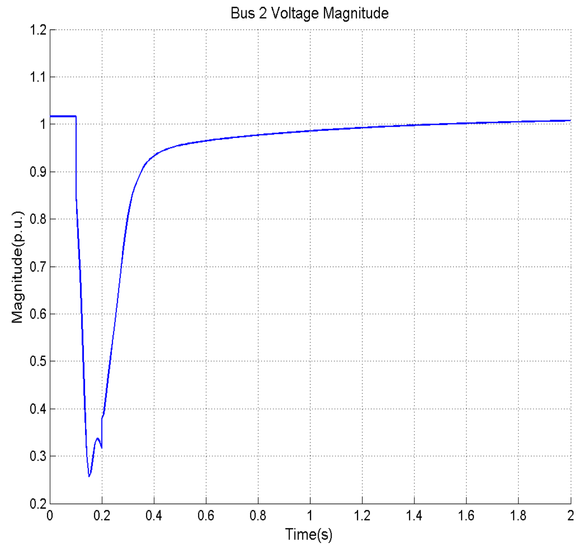

Figure 34: IND1 response to three-phase fault: bus voltage magnitude (pu) vs. time

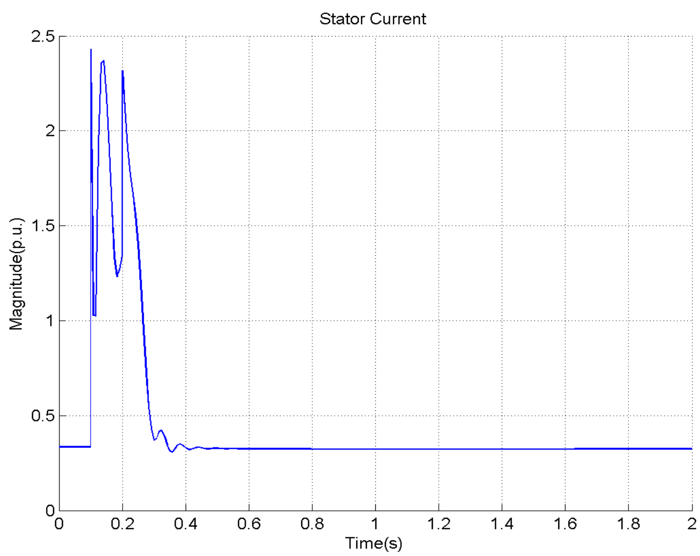

Figure 35: IND1 response to three-phase fault: stator current magnitude (pu) vs. time

In the above figures, it is clear that the single-phase motor does not stall. The speed decreases during the fault, and slowly recovers after the fault is cleared. In addition, the voltage magnitude at the bus decreases sharply during the fault, but then eventually recovers once the fault is cleared. The slow voltage recovery illustrates the point that the loading of the single-phase induction motor is rather high for this system. However, if the loading were increased, or if the fault were more severe, then the motor would likely stall.

Figures 36, 37, and 38 illustrate a stall situation, in which the fault is more severe and causes the single-phase induction motor to stall.

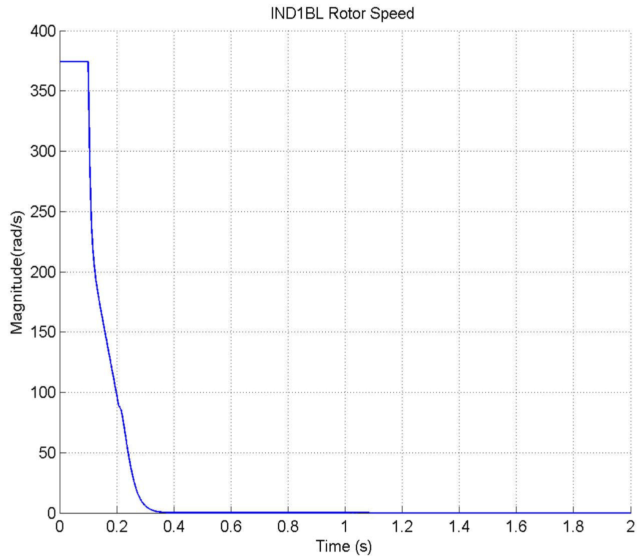

Figure 36: IND1 stall response to severe three-phase fault: rotor speed (rad/s) vs. time

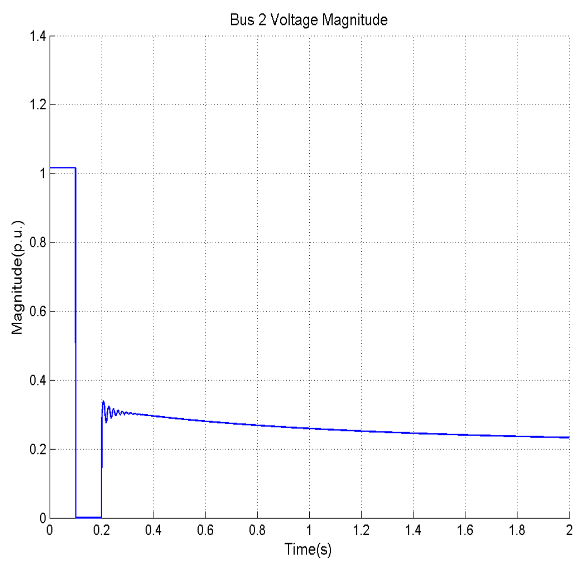

Figure 37: IND1 stall response to severe three-phase fault: bus voltage magnitude (pu) vs. time

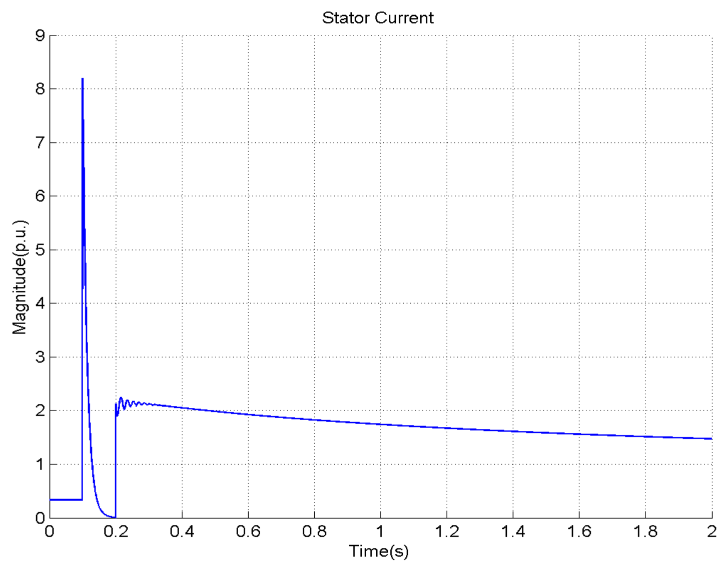

Figure 38: IND1 stall response to severe three-phase fault: stator current magnitude (pu) vs. time

Note that the bus voltage drops lower during the fault in Figure 37 than it did during the weaker fault. The lower voltage excursion initiates the stalling process. Also note the stator current magnitude for the stronger fault is initially higher at the beginning of the fault, but then decays due to the zero voltage condition on the terminals of the machine.

Lastly, note the stator current magnitude after the fault has cleared at t=0.2 s. In this situation, the stalled induction motor appears to the grid as a constant impedance, hence it draws a current in proportion to its impedance. Furthermore, the equivalent rotor impedance DECREASES as the motor slip increases to 1 per unit (i.e., stalled rotor) as compared to its steady-state slip of roughly 0.01 pu. This causes the single-phase induction motor to draw a higher current while stalled as compared to the current drawn before the fault.

In conclusion, the successful implementation and demonstration of the single-phase induction motor model with stall capability is a major breakthrough. The stall dynamics are critical to modeling the phenomenon known as fault induced delayed voltage recovery. The single-phase induction motor dynamics also play a critical role in slower time scale voltage collapse. Therefore, the first-of-its-kind implementation of a single-phase induction motor model with stall capability will allow engineers to simulate and design mitigation strategies against voltage collapse and cascading outages.

Summary of project activities

Due to two no-cost-extensions, the three-year project was extended to four years and five months. The extra time was allotted to parallel performance improvements, necessary to achieve faster-than-real-time dynamics simulation, as well as modeling improvements, necessary to demonstrate detailed simulation capabilities of three-phase unbalanced devices, single-phase induction motors with stall capability, and comprehensive dynamics plus protection simulation.

Over the span of the project, the team's activities included the following:

- Create Utility Advisory Group (UAG)
    - After initial discussions with 6 companies, we settled on two: Commonwealth Edison ("ComEd") of Chicago and AltaLink of Alberta. ComEd represents the class of eastern utilities with dense urban systems and AltaLink represents the typical Western or Southwestern utility with load centers and generation separated by substantial distances.
- Execute UAG Non-Disclosure Agreements
- Gather UAG data
- Solicit areas of interest from the UAG. For AltaLink, it was the performance of their connected wind generation plants to system disturbances (including interactions between the wind generation, their HVDC line and SVCs). For ComEd, it included detailed end-user induction motor modeling and the reaction of their SVC network to system disturbances.
- Conduct bi-monthly UAG conference calls (22 meetings over 4 years).
- Conduct monthly research team calls with DOE
- Host annual in-person project reviews at IIT
- Advanced Grid Modeling workshop in Knoxville, TN (Feb 2013)
- Peer Review Meeting in Washington, DC (Jun 2014)
- SIAM Annual Meeting Paper Session in Chicago, IL (Jul 2014)
- IEEE PES GM Panel Session in Washington, DC (Jul 2014)
- IEEE PES GM Panel Sessions in Boston, MA (Jul 2016)
- TS3ph simulator development based on Argonne's PETSc toolkit
- TS3ph-CAPE co-simulation development
- TS3ph generator model development (e.g., round rotor 6 state differential equation model)
- TS3ph generator controller model development (e.g., exciters, governors, stabilizers)
- TS3ph SVC model development
- TS3ph single-phase induction motor model development
- TS3ph performance tuning via Argonne's PETSc high performance scientific computing toolkit, including Time Stepping, Scalable Nonlinear Equations Solvers, Krylov Subspace Methods, Preconditioning, and Partitioning.
- TS3ph verification/validation with existing commercial tools (e.g., Siemens PTI PSSE, PowerTech TSAT, GE PSLF, PowerWorld)
- Linux-based high performance computing cluster specification, installation, configuration and tuning, including several Dell and SuperMicro 2U rack servers with Intel Xeon v3

and v4 processors, high density/high bandwidth memory modules, gigabit ethernet and InfiniBand communication fabric.

- Design and implement CAPE Server, with a sample client and API documentation. Communication is via Berkeley sockets. CAPE Server responds to a small set of new client/server specific commands, as well as any command that standard CAPE understands.
- Focus on two key project objectives
  - How to simulate more relays
  - How to reduce the CPU burden during the time stepping simulations
- How to simulate more relays
  - Minimize the amount of storage required for each relay. Any duplicate/unused storage could be removed to allow more relays to be stored. A thorough review of the data structures revealed that there was not enough duplicate/unused storage that could be removed to appreciably effect the overall requirements. We were then left with two methods to increase the available storage.
  - Convert CAPE to x64. The best, long-term, solution is to convert CAPE to x64. This would allow a virtually unlimited size for both network and protection model. However, this would have required extensive work to the CAPE interface, which would not benefit this project, so this approach was considered out of scope.
  - Use memory-mapped files. Memory-mapped files are a method to utilize available memory/disk space as program storage that is outside of the 32-bit address space. This essentially extends the available memory to near 64-bit allowances. A complete memory-mapped file implementation was developed for all of the relay/element data. The data for each relay is gathered into an object and that object is written to the mapped file. When the relay is needed, only that portion of the file is mapped into the 32-bit address space and the relay data are available for use/updating. When finished, the memory is freed so that a different relay can be mapped. By using only a portion of the mapped file, the 32-bit address limits are never exceeded. This is very effective in increasing the number of relays that can be simulated. However, there is a performance penalty for continually mapping/remapping a portion of the file. The program was approximately 5% slower using memory mapped files. In addition, while the mapped file can be shared among different processes/threads on the same computer, it cannot be shared across a network of computers.
  - The conclusion is that eventually, CAPE should be converted to x64.
- How to reduce the CPU burden during the time stepping simulations. There are three approaches to achieve this.
  - Look for inefficient code/processes and improve them. Using tools from both Microsoft and Intel, the existing code was analyzed for bottlenecks. This cycle of analyze/fix was continued until the resulting code was approximately 10 times faster.
  - Implement multiple threads. After the efficiency improvements became harder to realize (or fixing those found was decided to be out of scope), using multiple threads was considered. OpenMP was selected as the standard to use for this phase. OpenMP allows FORTRAN cycle loops (Do loops) to be easily enclosed

with high-level statements that effectively break the loop and allow it to be processed by several threads simultaneously, with each thread doing a share of the total work. While easy to implement, the underlying code must be thread safe, and CAPE was not. With great difficulty, and much time, the code was made thread-safe. The resulting speedup was as high as 12:1 when run on a 16-core computer.

    - Implement multiple processes. After multiple threads were implemented and working properly, we decided to implement multiple processing. OpenMPI was the standard chosen for this work. This standard essentially allows one program to be run n-times simultaneously. One of the n-programs acts as the parent, and doles out work to the remaining (n-1) programs, which act as children. The programs can be running on the same or different computers, and exchange information in order to stay synchronized. The observed speedup was 5:1 with 10 child processes all running on the 16-core server. While not as effective as using multiple threads, using multiple processes is incredibly easier to implement.
- Alstom Grid (now GE Energy) contributed their expertise to several aspects of the project including the following:
    - GE used Powertech's Transient Stability Assessment Tool (TSAT) as a validation tool for IIT's TS3ph Three-Phase Simulator. For this purpose, a reduced model was utilized that comprised of 5 buses, 4 lines, and 2 composite loads, all connected to an infinite bus representing the external generation in the system (Figure 39).

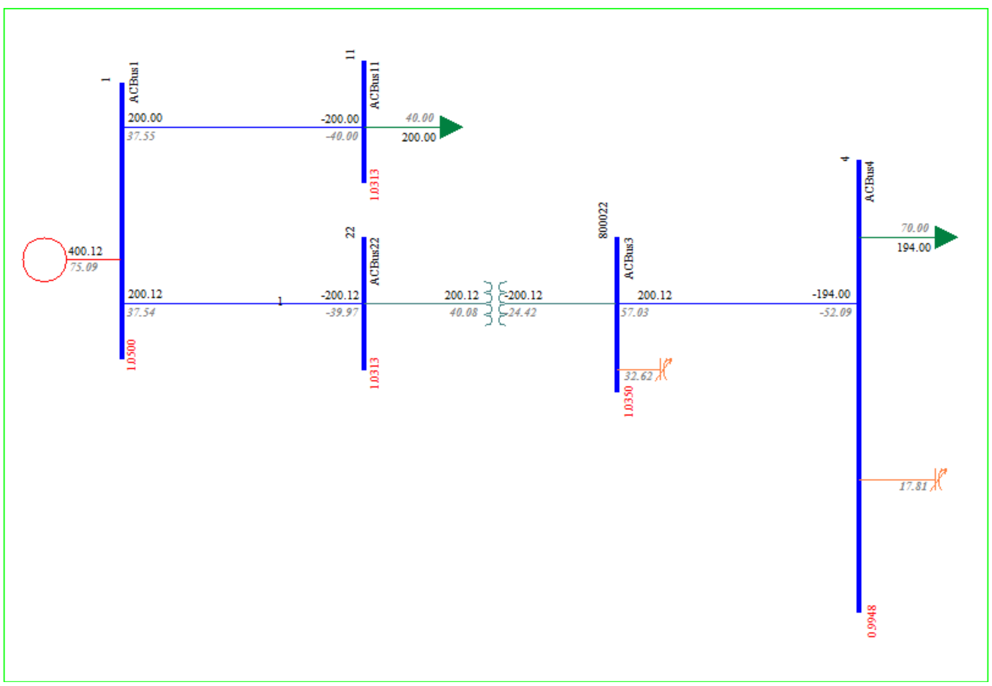

Figure 39: Reduced 5 bus system

- For composite load model verification purposes, various scenarios were studied. These included composite models with single-phase induction motor models as well as three-phase induction load model (Figure 40).

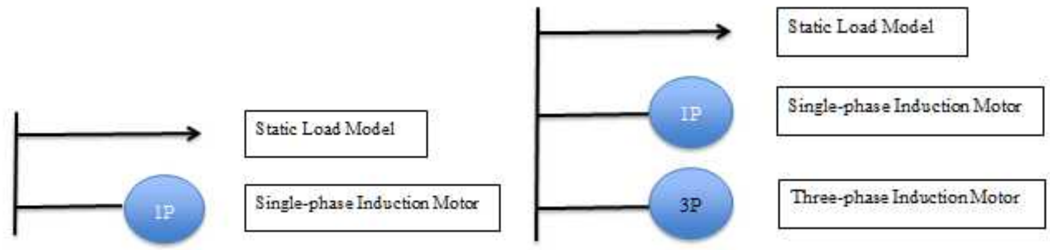

Figure 40: Composite load models with single-phase and three-phase induction motors

- For each of the scenarios, different contingences were analyzed including a gradual step-down of generator voltage until the motor stalls and imitating a fault induced delayed voltage recovery (FIDVR) event. The scenarios illustrated the impact of the dynamic nature of the three-phase motor model including motor slip reflected in the motor speed of the three-phase motor (Figure 41).

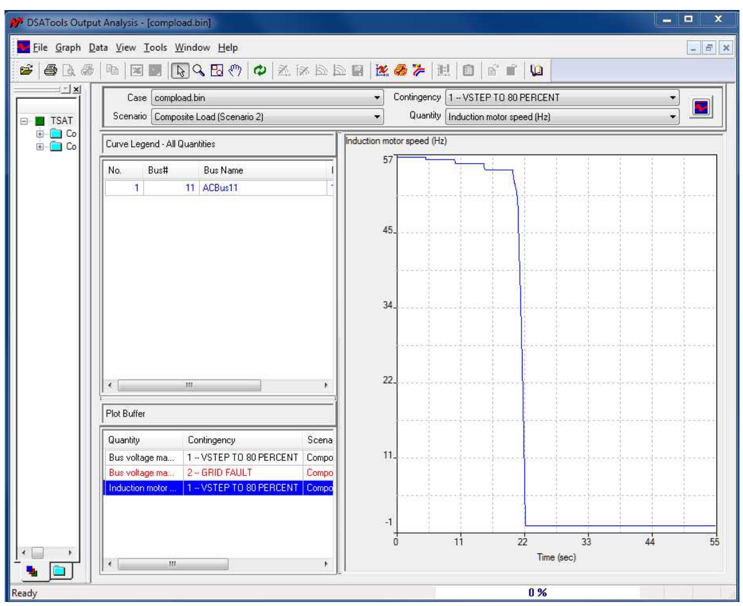

Figure 41: Motor slip from gradual step-down of generator voltage

- Additionally, a larger 23-bus model cases with 25% and 50% motor load at each load bus was also simulated for further verification of theTS3ph simulator. For both scenarios, a fault was introduced with a fault duration of 0.05 seconds (3 cycles), and replicated using TSAT. Both, an explicit integration scheme (RK2) and trapezoidal integration method were used for the numerical simulation.
- GE provided IIT with the e-terraphasoranalytics application as a tool to facilitate a more detailed comparison of simulation results with actual field data captured by

Phasor Measurement Units (PMUs) or Disturbance Fault Recorders (DFRs). The application natively can import data from multiple sources including directly from TSAT in its native 'bin' format as well as other simulation tools in a .csv format. Within the same application, synchrophasor data can simultaneously be imported in a COMTRADE format, thereby making it easier to conduct a side-by-side comparison of the TS3ph simulations against PMU output data.

- The e-terraphasoranalytics application that was provided to IIT offers four ways for comparing the simulation results and measurement data:
    - **Time Domain Plots**: This is a graphical overlay of both simulation and measurement trends as a function of time on the same chart.
    - **Frequency Domain Analysis**: Here one can overlay the frequency domain power system spectral plots (i.e. function of frequency) to compare the simulation results with the measurement data in the frequency domain. Peaks in the spectral plots correspond to resonant modes, thereby allowing the end use to compare the spectral features across the two datasets.
    - **Ringdown Analysis**: This analytic uses system identification techniques such as Prony or Hankel-Total Least Squares (HTLS) approach to decompose the signals into its dominant modes within a user-defined frequency range. The analysis results include mode frequencies, mode damping, and mode shape (i.e. magnitude and phase). The analysis therefore allows the end-user to conduct a side-by-side comparison of the mode decomposition between the measured signals and the simulation results (Figure 42).

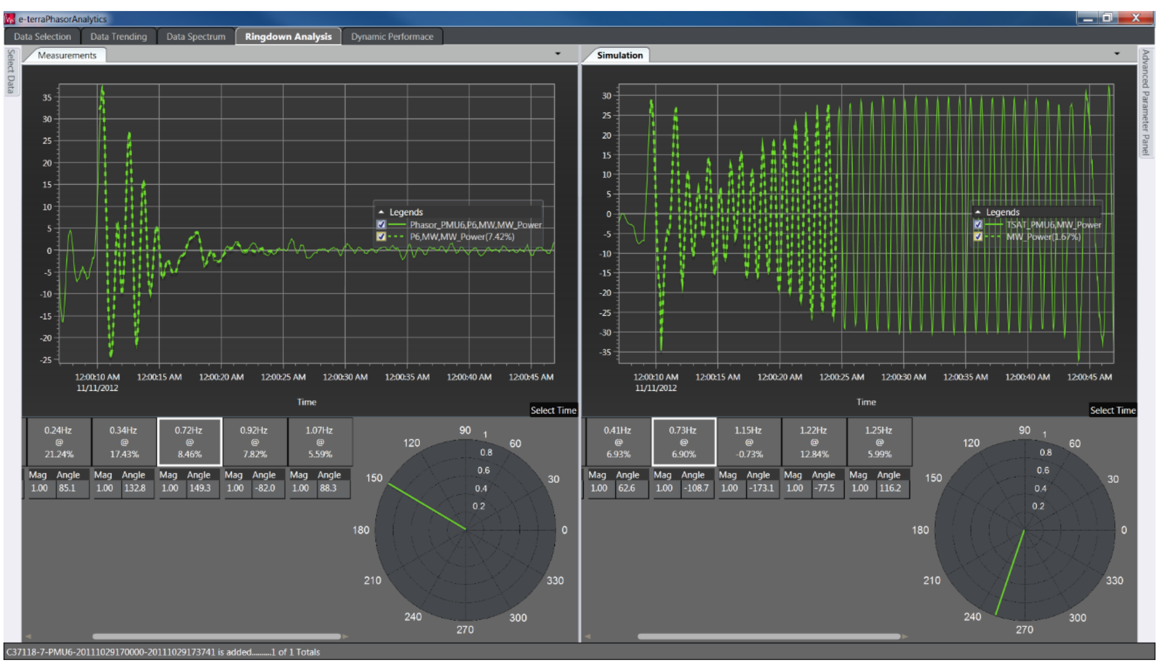

Figure 42: Illustration of side-by-side ringdown analysis between measured signals (left) and simulation results (right)

- **Primary Frequency Response Characterization**: For generation-load mismatch events (such as generation trip), the analysis estimates key parameters associated with the primary frequency response that is observed following the contingency. These parameters include inertia, damping, droop, frequency response coefficient, etc.

- The application's COMTRADE export capability was enhanced to support DFR COMTRADE data imports into the application (Figure 43), since ComEd has DFRs in the field and the tool would allow IIT to compare the field data with TS3ph simulations.

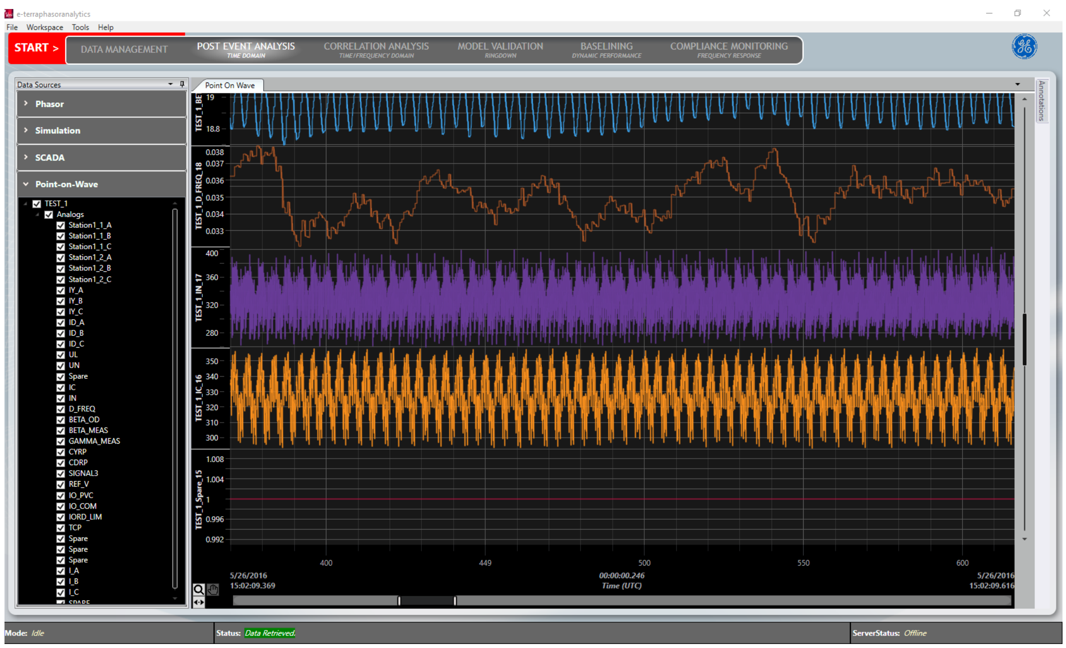

Figure 43: DFR COMTRADE import into e-terraphasoranalytics

Hypotheses:

- Faster than real-time dynamics simulation is possible for large-scale electric power grid models.
- Three-phase unbalanced network models can adequately describe operating conditions on interconnected power grids.
- Detailed dynamics models can be combined with detailed protection models to yield a comprehensive dynamics plus protection simulation.
- Single-phase induction motors, such as those used to drive residential air-conditioning compressors, and their undesirable stalling behavior can be incorporated into unbalanced three-phase dynamic grid simulators.

Assumptions:

- Dynamic device models represent the true dynamic structure of the individual grid components.
- Dynamic device model data and transmission network model data represent the true parameters of the grid infrastructure.

- Complex power grid dynamic behavior can be predicted accurately if given sufficiently accurate models.
- Model inaccuracies can be corrected via model validation procedures given synchrophasor data from phasor measurement units (PMUs).
- Initial operating conditions can be determined prior to a grid disturbance.
- Post-disturbance dynamic behavior can be predicted if the initiating event is known, such as a fault being detected by a relay, or a change in operating conditions being detected by a PMU.
- Faster than real-time simulations can help grid operators avoid blackouts by supplying look-ahead predictions of future grid dynamics and the associated "future state of health". Based on the look-ahead information, operators will have the support they need to take decisive action even while experiencing the uncertainty of not knowing exactly what is happening to the grid during a severe disturbance, such as a cascading outage.

Approaches:

- Leverage recent mathematical and computational advances via high performance parallel computing to produce comprehensive simulation results faster than real-time.
- Create three-phase unbalanced network models to accurately capture dynamic behaviors, including unbalanced dynamics.
- Implement detailed controller models to accurately capture emergent dynamics of interconnected grids.
- Develop new models for grid controls and dynamic loads that better represent existing and future grid components.
- Combine detailed dynamics models with detailed protection models to create a comprehensive simulation framework for predicting dynamic grid behavior.
- Implement multi-threading and message-passing techniques in a subset of the protection simulator code, which will lead to the biggest benefit-to-cost ratio.

Problems:

- Computational complexity: Detailed generators, detailed controllers and detailed induction motor models contain nonlinear elements due to nonlinear physical phenomena and controller limits, including both windup and non-windup limits.
- Weak scalability due to tightly coupled model equations: Power grid models do not lend themselves to naturally parallel decompositions as the network structure is highly meshed, similar to computational fluid dynamics, yet at the same time, the network graphs are unstructured, which is unlike computational fluid dynamics problems.
- Memory bandwidth limits for large sparse matrices: Modern computer architectures exploit data locality, which is a desirable feature in most computational problems, with a hierarchy of caches between the processing core and main memory. Unfortunately, power grid simulations depend on sparse matrix representations. Hence, indirect addressing is required, which significantly weakens the data locality properties of large-scale grid models. Therefore, the improvements in modern computer architectures and their associated subsystem components only weakly benefit large-scale grid simulations. The trend toward higher core counts of modern processors will continue to benefit naturally parallel computational problems, but electric power grid simulations will remain outside the mainstream market focus.

- Detailed protection models require extensive communication (i.e., collect voltages, currents and breaker status signals) and extensive computation (i.e., determine whether or not a relay element will assert that a specific condition has been met, and evaluate breaker operation logic) to determine protection actions. While much of this could be done in parallel, that was not the framework used to develop the protection simulator over the past few decades.

Departures from planned methodologies (and assessment of their impact on project):

- Three-phase unbalanced models of large-scale power systems may not yield faster than real time simulation results, so a tradeoff between modeling detail and simulation performance must be considered.

Conclusion:

- The project was a success with several major breakthroughs in the area of large-scale power system dynamics simulation, including
    - A validated faster than real-time simulation of both stable and unstable transient dynamics in a large-scale positive sequence transmission grid model,
    - A three-phase unbalanced simulation platform for modeling new grid devices, such as independent phase controllers for single-phase static var compensators,
    - The world's first high fidelity three-phase unbalanced dynamics and protection simulator based on Electrocon's CAPE program, and
    - A first-of-its-kind implementation of a single-phase induction motor model with stall capability.
- Beyond the current project, there are ongoing research and development initiatives that were launched from this project:
    - CAPE is currently being redesigned with all of the knowledge gained during this project, to implement all of the efficiency improvements that were considered to be out of scope. This implementation will include x64, multiple threads, and multiple processes.
    - The research prototype TS3ph-CAPE simulator is being used to explore cascading outages.

Products developed under the award

Publications (journal, volume, issue)

- "Parallel Dynamics Simulation using Krylov-Schwarz Linear Solution Scheme", *IEEE Transactions on Smart Grid*, Volume 8, Issue 3, 2017

Conference papers

- "Variable-Step Multi-Stage Integration Methods for Fast and Accurate Power System Dynamics Simulation", *X Bulk Power Systems Dynamics and Control Symposium*, IREP 2017, Portugal

Conference presentations

- "Predicting the Future with Faster Than Real-Time Power System Dynamics Simulation", *Proceedings of the SIAM Annual Meeting*, 2014
- "High-Fidelity, Faster than Real-Time Dynamics Simulation", *Proceedings of the IEEE Power & Energy Society General Meeting*, 2014
- "What's next? High-fidelity power system dynamics and protection simulation", *Proceedings of the IEEE Power & Energy Society General Meeting*, 2016
- "High Performance Computing for Grid Operations - What Can Be Learned from Other Fields?" *Proceedings of the IEEE Power & Energy Society General Meeting*, 2016

Conference panel sessions organized and chaired

- Panel Session CAMS-4: "Faster than Real-time Dynamics Simulation", *IEEE PES General Meeting*, 2014
- Panel Session PSACE CAMS: "Cascading Outages - Dynamics, Protection, Validation and Data", *IEEE PES General Meeting*, 2016

Networks/collaborations fostered

- Combined dynamics and protection simulation is vital for understanding cascading outages. This project enabled IIT and Electrocon to develop a co-simulation framework that ultimately will help transmission operations and planning engineers work with system protection engineers to plan and operate large-scale power grids in a more efficient and reliable manner.
- Integrating new mathematical and computational advances into electromechanical transients analysis will lead to enhanced modeling and simulation capabilities for the nation's electric grid. This project enabled IIT and Argonne National Laboratory to evaluate new computational methodologies in a detailed dynamics simulator prototype. The interaction improved the simulator prototype performance, while also providing valuable insight for applied mathematicians and computer scientists into the domain of large-scale grid simulation.
- Analyzing dynamic device model structure and model parameters within large-scale grid datasets continues to be an important aspect of power grid dynamic analysis. This project enabled IIT and the Utility Advisory Group to explore the fidelity of several classes of dynamic device models in the areas of generator controllers, non-rotating dynamic voltage support, and induction motor loads.

Technologies/Techniques

- High performance computing applied to large-scale power system dynamics simulation
- Three-phase modeling of transmission infrastructure
- Combined dynamics and protection simulation

Data/Databases/Software/Models

- Research prototype for three-phase transmission dynamics simulation
- Research prototype for combined dynamics and protection simulation
- Research prototype for faster-than-real-time dynamics simulation
- Independent single-phase control of static var compensators
- Single-phase induction motor model with stall capability