

*Exceptional service in the national interest*



IOWA STATE  
UNIVERSITY

UC DAVIS  
UNIVERSITY OF CALIFORNIA



ALSTOM



# Stochastic Unit Commitment: Scalable Computation and Experimental Results

Dr. Jean-Paul Watson

Discrete Math and Complex Systems Department

Sandia National Laboratories, Albuquerque, NM



Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

# Core Project Team

- Sandia National Laboratories
  - Ross Guttromson, MS, MBA, PE
  - Jean-Paul Watson, PhD
  - Cesar Silva Monroy, PhD
- Iowa State University
  - Sarah Ryan, PhD
  - Leigh Tesfatsion, PhD
  - Dionysios Aliprantis, PhD
- Alstom Grid
  - Kwok Cheung, PhD
- UC Davis
  - Roger Wets, PhD
  - David Woodruff, PhD
- ISO New England
  - Eugene Litvinov, PhD

# External Advisors

- Eugene Litvinov, ISO-NE
  - Chairs Advisory Team
- Richard O'Neill, FERC
- Ralph Masiello, KEMA
- David Morton, UT Austin

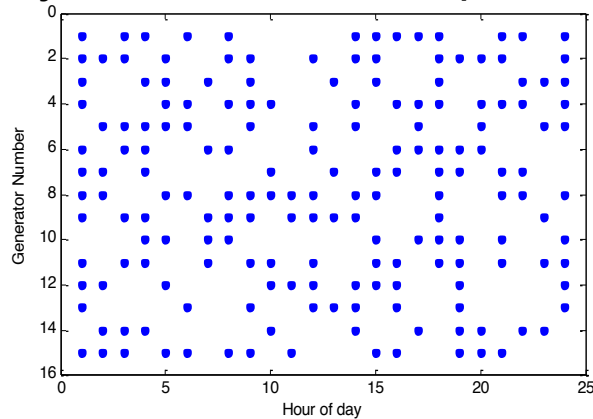
# Project Goals

- Execute stochastic unit commitment (UC) **at scale, on real-world data sets**
  - Stochastic UC state-of-the-art is very limited (tens to low hundreds of units)
  - Our solution must ultimately be useable by an ISO
- Produce solutions **in tractable run-times, with error bounds**
  - Parallel scenario-based decomposition
    - For both upper and lower bounding (Progressive Hedging and Dual Decomp.)
  - Quantification of uncertainty
    - Rigorous confidence intervals on solution cost
- Employ high-accuracy stochastic process models
  - Leveraged to achieve computational tractability while maintaining solution quality and robustness
- Demonstrate potential **cost savings on an ISO-scale system at high renewables penetration levels**

*PART 1:*  
*INTRODUCTION AND CONTEXT*

# The General Structure of a Stochastic Unit Commitment Optimization Model

Objective: Minimize expected cost



First stage variables:

- Unit On / Off



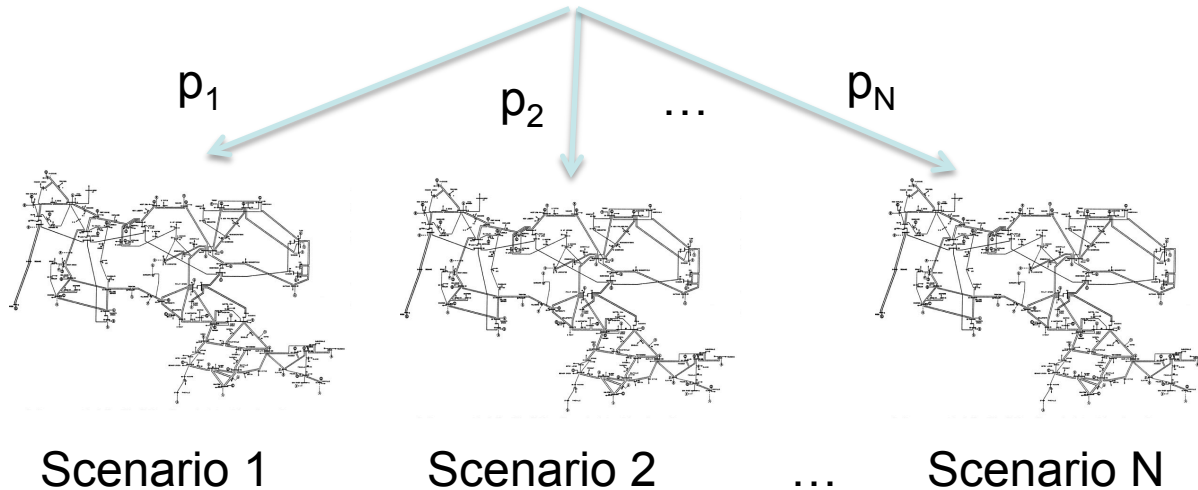
Nature resolves uncertainty

- Renewables output
- Forced outages



Second stage variables  
(*per time period*):

- Generation levels
- Power flows
- Voltage angles
- ...



Scenario 1

Scenario 2

...

Scenario N

## Stochastic Programming Models

- Reliability Unit Commitment
  - Renewables generator output, load, forced (unplanned) outages
  - Fewer binaries than DAM, long time horizon, many scenarios
- Look-Ahead Unit Commitment
  - Similar to Reliability Unit Commitment
  - Fewer binaries than RUC, short time horizon, few scenarios
- Day-Ahead Unit Commitment
  - In contrast to RUC and SCED2, an ISO can't really make direct use of a stochastic UC in the DAM without changing DAM procedures
  - With our partners, we are exploring alternative models and experimenting with procedures that incorporate stochastic models
  - We are eager to discuss ideas offline

# Core Unit Commitment Model

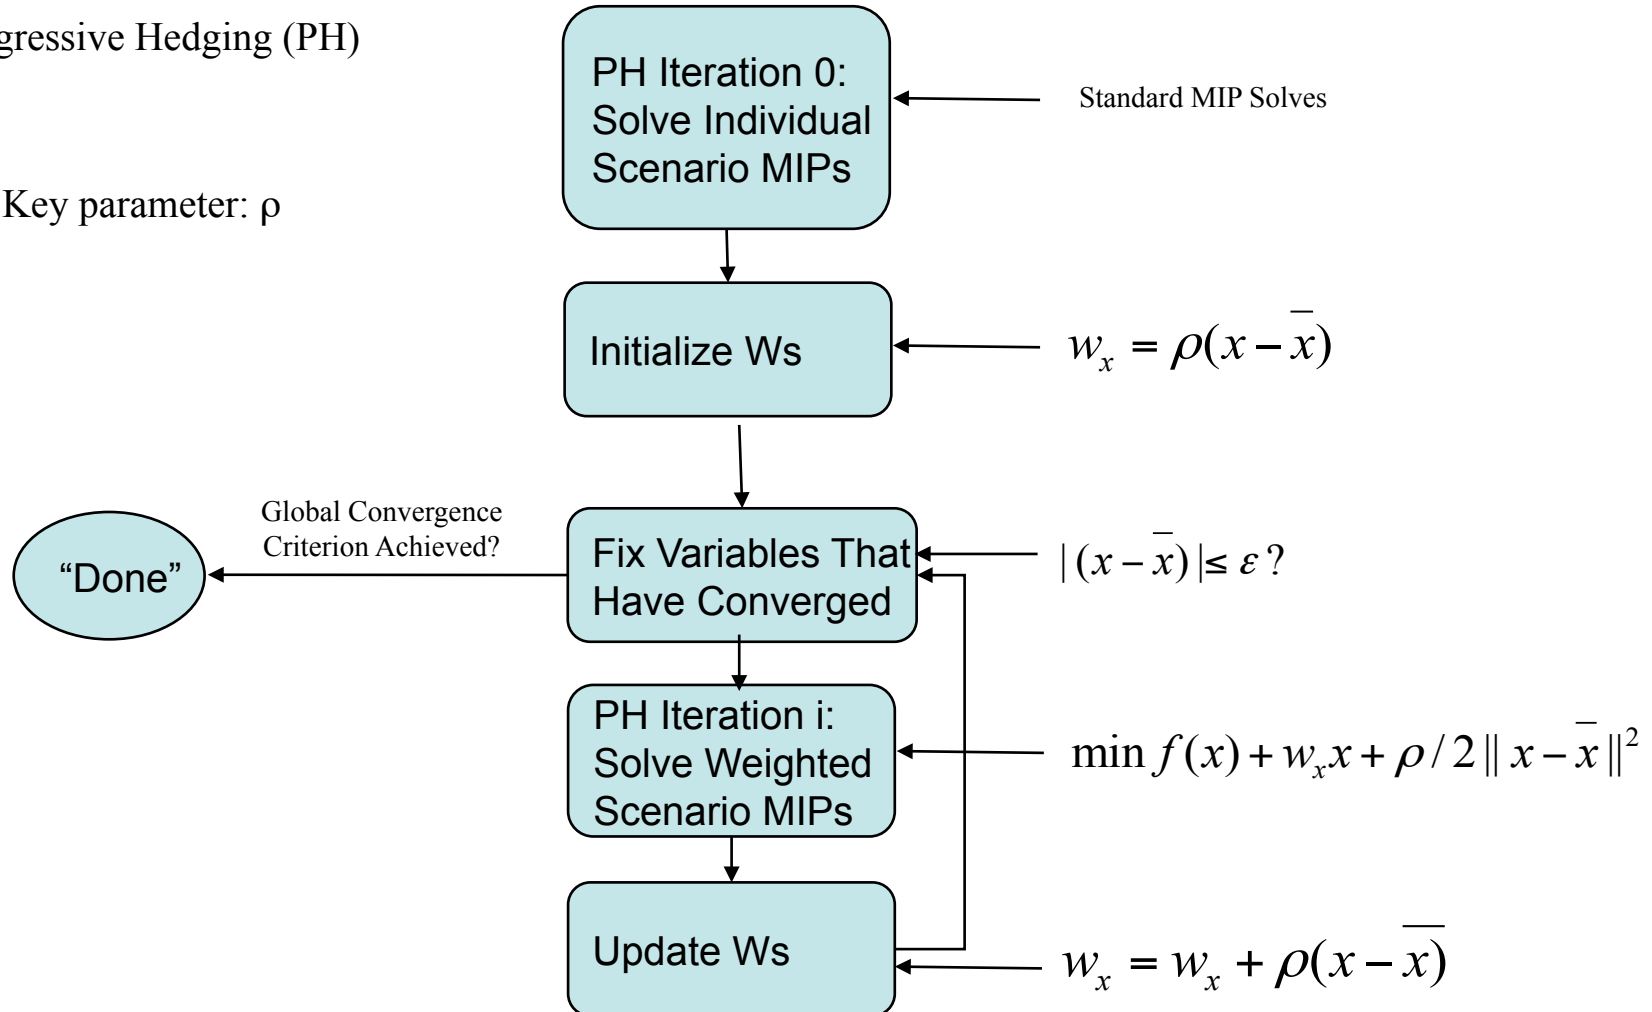
- Basic deterministic / single-scenario unit commitment model
  - Carrion and Arroyo (2006)
  - Alternative to the well-known 3-binary variable formulation
- Based on empirical evidence developed during this and prior projects, we find no serious performance differences between the Carrion and Arroyo formulation and other formulations
  - Highly problem-dependent
- Our UC model deviates from the core Carrion and Arroyo model in two key ways
  - Different startup / shutdown cost modeling components
  - Inclusion of high-fidelity ancillary services modeling components
- Our model is cross-validated with Alstom's UC model
  - Accurate to within solver tolerances

*PART 2:*  
*COMPUTATIONAL RESULTS*

# Scenario-Based Decomposition

Progressive Hedging (PH)

Key parameter:  $\rho$



# Progressive Hedging: Some Algorithmic Issues and their Resolution

- We are dealing with mixed-integer programs
  - So we have to deal with the possibility of cycling and other manifestations of non-convergence
  - See: *Progressive Hedging Innovations for a Class of Stochastic Mixed-Integer Resource Allocation Problems*, J.P. Watson and D.L. Woodruff, Computational Management Science, Vol. 8, No. 4, 2011
- Good values for the  $\rho$  parameter are critical
  - Poor or ad-hoc values of  $\rho$  can lead to atrocious performance
  - The good news in unit commitment
    - We have a lot of information concerning the cost of using a generator
    - Use the LMPs associated with PH iteration 0 scenario solutions
    - Cost-proportional rho is a known, effective strategy in Progressive Hedging
  - Also see Computational Management Science paper indicated above

# Progressive Hedging: Parallelization and Bundling

- Progressive Hedging is, at least conceptually, easily parallelized
  - Scenario sub-problem solves are clearly independent
  - Advantage over Benders, in that “bloat” is distributed
    - Critical in low-memory-per-node cluster environments
  - Parallel efficiency drops rapidly as the number of processors increases
    - But: *Relaxing barrier synchronization does not impact PH convergence*
  
- Why just one scenario per processor?
  - Bundling: Creating miniature “extensive forms” from multiple scenarios
    - Diverse or homogeneous scenario bundles?
  - Empirically results in a large reduction in total number of PH iterations
    - Growth in sub-problem cost *must* be mitigated by drop in iteration count
    - In practice, mitigation is enabled by cross-iteration warm starts

# But PH is just a Heuristic!

- So is any complete optimization algorithm that is *not* allowed to run to – completion!
  - Key point is that we don't believe it will be possible to obtain optimal solutions to stochastic unit commitment problems at scale, in tractable wall clock times
- But PH doesn't provide bounds!
  - No longer true
  - Now comes with (rather tight) lower bounds
  - For details, see Sarah Ryan's talk at the upcoming Stochastic Programming conference in Bergamo, Italy, in July
- More seriously
  - We have a lot of work going on in the realm of lower bounding, which we are happy to discuss off-line

# Illustrative Results: WECC-240

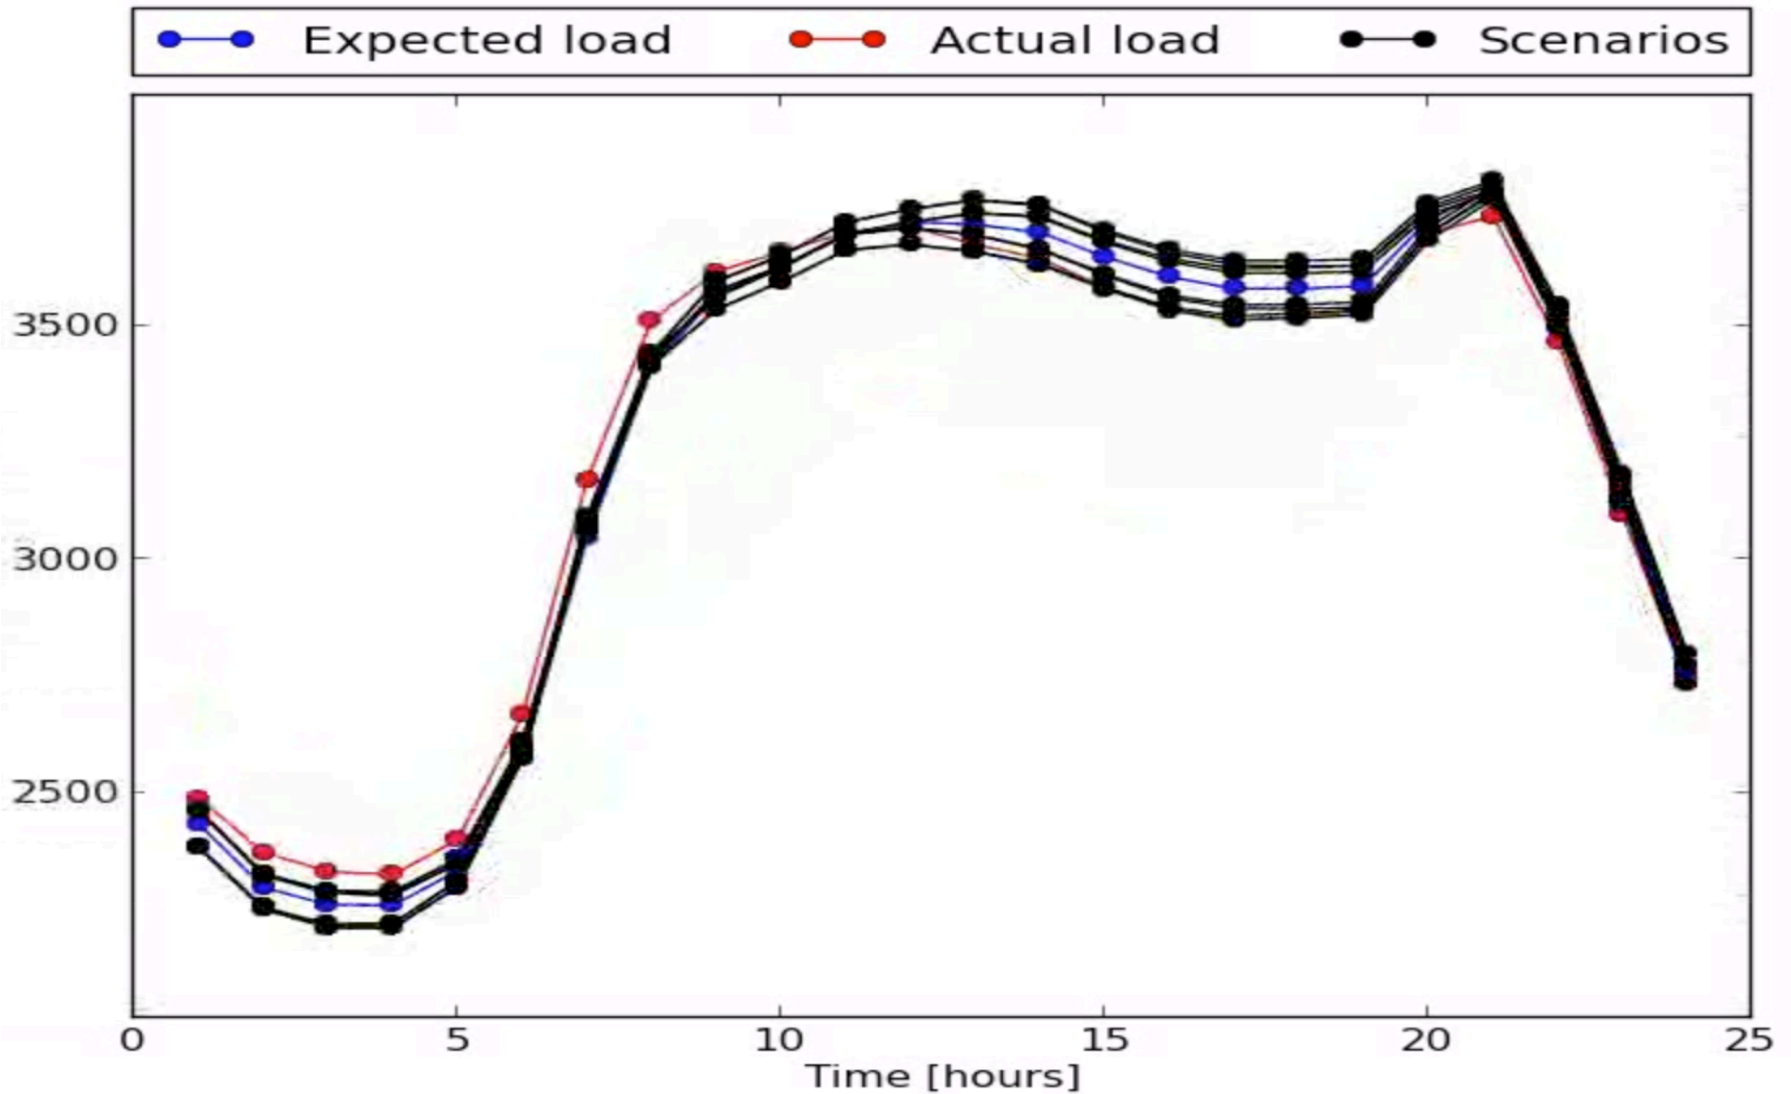
- Test instance
  - Modified WECC-240 instance for reliability unit commitment
  - Stochastic demand, 100 scenarios
- Extensive form
  - CPLEX, after 1 day of CPU on a 16-core workstation
    - No feasible incumbent solution
- PH, 20 iterations, post-EF solve - serial
  - ~13 hours, 2.5% optimality gap
- PH, 20 iterations, post-EF solve – parallel
  - ~11 minutes, 2.5% optimality gap
- PH, 20 iterations, post-EF solve – parallel with bundling
  - ~12 minutes, 1.5% optimality gap

*PART 3:*  
*EMPIRICAL RESULTS*

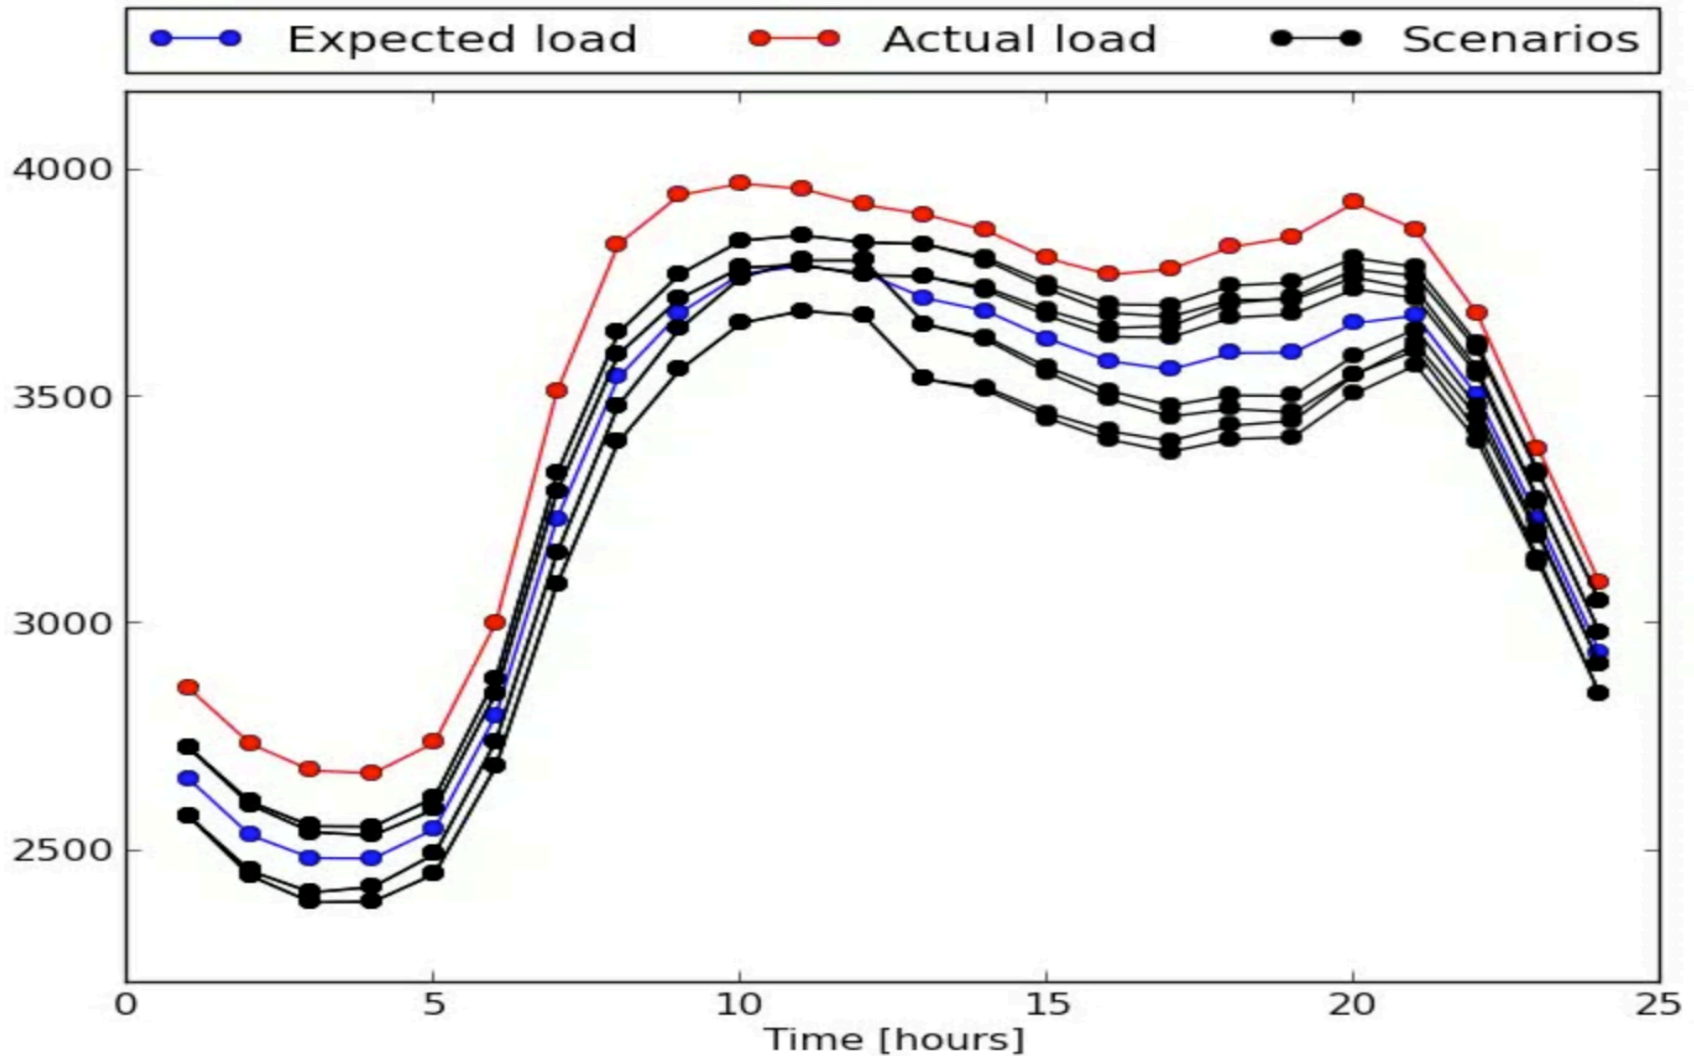
# RUC Test Instance: WECC-240++

- J.E. Price, Reduced Network Modeling of WECC as a Market Design Prototype, 2011 IEEE PES General Meeting
- Changes necessary to create viable RUC test case
  - Addition of realistic ramping rates and min up/down time constraints
  - Culling renewables generators present in the original instance
- Scale ISO-NE load forecast models for 2011 to WECC-240
  - Very realistic and accurate load scenarios
- Use EWITS wind power scenarios developed for ISO-NE
  - Obtained by 3TIER, derived using *analog* methodology
  - Scaled to WECC-240 system to represent varying penetration levels
- Bottom line
  - Our WECC-240 test case is a realistic ISO system, evaluated using accurate stochastic process models developed for ISO-NE

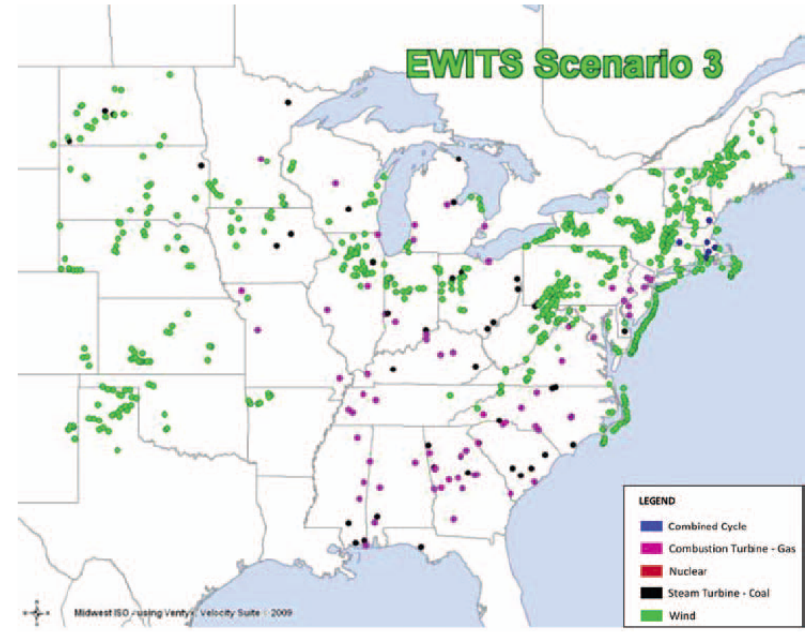
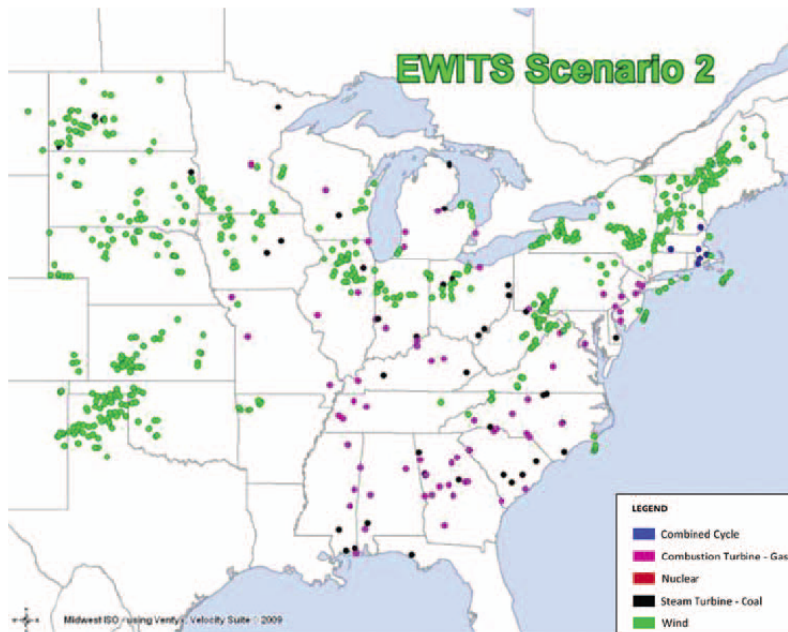
# Near-Ideal Load Scenarios



# Less-Than Ideal Load Scenarios



# EWITS Scenario 2 and 3 Maps



# Example Wind Power Analogs



*Line width is correlated with scenario probability*

# Simulation Methodology

- General idea
  - Evaluate *energy production cost* observed while executing deterministic and stochastic UC over an extended time horizon
    - RUC-only, and reserve pricing is ignored
  - Costs quantified using fixed costs obtained from unit commitment and variable costs obtained from hourly economic dispatch
- Out-of-sample scenarios are used to simulate what would “really have happened” in our alternative ISO universe
  - Contrasts with Monte Carlo simulations over a focused set of days
- Deterministic UC executed using expected scenarios
  - For both load and wind power
- Static reserves employed in both UC variants
  - Reserves necessary by definition in deterministic UC
  - Limits on forecasting and sampling precision necessitates the use of reserves in an operational stochastic UC context


# Illustrative Results: Load-Only Scenarios


- Simulated for April – June 2011
- 50 load scenarios per day
  - More do not have an appreciable impact on system performance
- Deterministic base case
  - Static reserves employed, 7% required to avoid load shedding
- Stochastic comparative case
  - Static reserves also employed, but at 2% level to avoid load shedding
- Stochastic solution is ~1.5% less costly in terms of the cost of energy production than the deterministic solution
- NOTES:
  - Stochastic solve times vary between 7 and 16 minutes (wall clock)
  - *Not* presently factoring in cost of reserves
  - Inclusion of generator outages increases savings to 3-5%

# Illustrative Results: Load and Wind Power Scenarios

- Results from initial use of wind power scenarios
- Limit to 100 scenarios in aggregate (which is not enough)
  - Using simple cross technique to form load + wind power scenarios
  - 10 load scenarios, 10 wind power scenarios
- Scaled to represent approximately 10% wind penetration in WECC-240 RUC test case
- Adjustment of static reserve levels
  - 3% for stochastic UC (due to limited number of wind scenarios)
  - 9% for deterministic UC
- Notes
  - Again, not factoring in costs of reserves
  - Stochastic solve times rise to between 9 and 18 minutes (wall clock)

# Our Software Environment: Coop

- Project homepage
  - <http://software.sandia.gov/coopr>
- “The Book” 
- Mathematical Programming Computation papers
  - Pyomo: Modeling and Solving Mathematical Programs in Python (Vol. 3, No. 3, 2011)
  - PySP: Modeling and Solving Stochastic Programs in Python (Vol. 4, No. 2, 2012)



# Our Hardware Environments

- Our objective is to run on commodity clusters
  - Utilities don't have, and don't want, supercomputers
  - But they do or might have multi-hundred node clusters
- Sandia Red Sky (Unclassified Segment) – 39<sup>th</sup> fastest on TOP500
  - Sun X6275 blades
  - 2816 dual socket / quad core nodes (22,528 cores)
    - 2.93 GHz Nehalem X5570 processors
    - 12 GB RAM per compute node (1.5 GB per core) << IMPORTANT!
  - For us, the interconnection is largely irrelevant
  - Red Hat Linux (RHEL 5)
- Multi-Core SMP Workstation
  - 64-core AMD, 512GB of RAM
  - For only \$17K from Dell....

# Conclusions

- Stochastic unit commitment has been studied in the literature
  - Indications are that it holds promise
  - Computational challenges have prevented industrial adoption
  - Far easier on paper and in academia than in practice...
- We have developed a promising approach to scalable stochastic unit commitment
  - Initial results are very promising in terms of necessary response times
  - Using reasonable, high-accuracy stochastic process models
- Our technology is starting to allow us to rigorously quantify, at scale, the potential benefits of stochastic unit commitment
  - Critically dependent on high-accuracy stochastic process models
- We are happy to talk to:
  - ISOs, vendors, and academics working toward related goals

# Next (Immediate) Steps

- Improve methodology for accurately accounting for costs associated with reserves
  - Failure to do so limits ultimate impact of stochastic unit commitment
- Extension of WECC-240 simulation results to higher wind penetration levels
- Analysis on internal ISO-NE test case
- Release of stochastic load scenarios for FERC / PJM case

# References

- Y. Feng, D. Gade, S.M. Ryan, J.P. Watson, R.J.B. Wets, and D.L. Woodruff (2013). A New Approximation Method for Generating Day-Ahead Load Scenarios. In *Proceedings of the IEEE 2013 IEEE Power and Energy Society General Meeting*.
- S.M. Ryan, C. Silva-Monroy, J.P. Watson, R.J.B. Wets, and D.L. Woodruff (2013). Toward Scalable, Parallel Progressive Hedging for Stochastic Unit Commitment. In *Proceedings of the IEEE 2013 IEEE Power and Energy Society General Meeting*.

**QUESTIONS**

