Exceptional service in the national interest

Sandia National Laboratories

IOWA STATE UNIVERSITY

UC DAVIS UNIVERSITY OF CALIFORNIA

arpa·e

ALSTOM

ISO new england

On Deployment Barriers and Research Challenges for Stochastic Unit Commitment

Dr. Jean-Paul Watson

Analytics Department

Sandia National Laboratories, Albuquerque, NM

U.S. DEPARTMENT OF ENERGY

NNSA National Nuclear Security Administration

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.

With Thanks To… (But Absolutely No Endorsement Implied!)

- Sandia National Laboratories
  - Cesar Silva-Monroy
- Alstom Grid
  - Kwok Cheung
- Iowa State University
  - Sarah Ryan
- UC Davis
  - Roger Wets
  - David Woodruff

(Some) Historical Barriers to Adoption of Stochastic Unit Commitment

- ***We can't create sufficiently accurate sets of scenarios to capture load and renewables uncertainty***

- ***Even if we could create accurate sets of scenarios, the resulting models are too difficult to solve***

- ***Even if we could solve the resulting models, it would require significant HPC resources – which is a major impediment to industrial adoption***

Topics That Will Not Be Discussed…

- Unless provoked!
- Robust optimization “versus” stochastic programming
- Market design and stochastic programming

Context: An ARPA-e GENI Project

- Execute <u>*stochastic*</u> 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***

The General Structure of a Stochastic Unit Commitment Optimization Model

Objective: Minimize expected cost

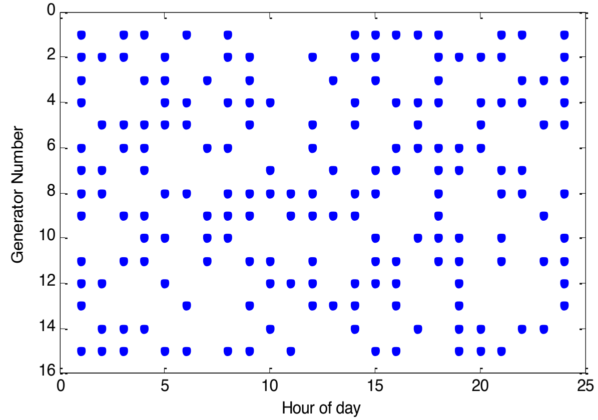

p_1 p_2 … p_N

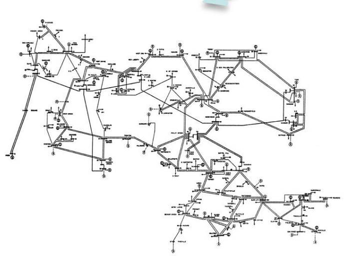

Scenario 1

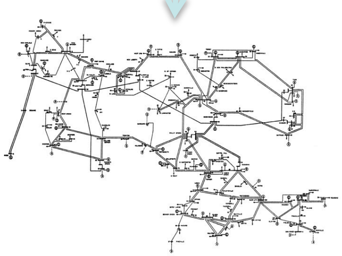

Scenario 2

…

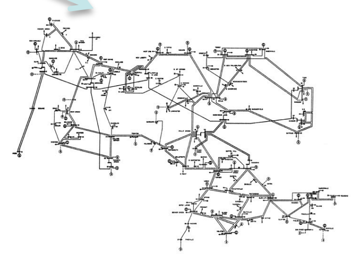

Scenario N

First stage variables:

- Unit On / Off

Nature resolves uncertainty

- Load
- Renewables output
- Forced outages

Second stage variables (*per time period*):

- Generation levels
- Power flows
- Voltage angles
- …

Uncertainty in DAM, RUC, and SCED2 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
  - 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

(Some) Historical Barriers to Adoption of Stochastic Unit Commitment

- ***We can't create sufficiently accurate sets of scenarios to capture load and renewables uncertainty***

- ***Even if we could create accurate sets of scenarios, the resulting models are too difficult to solve***

- ***Even if we could solve the resulting models, it would require significant HPC resources – which is a major impediment to industrial adoption***

On Scenario Generation for Stochastic Unit Commitment…

- Stochastic programming, like all things algorithmic, operates on the “GIGO” principle (Garbage In, Garbage Out)
  - If you don’t get the scenarios right, then the solution will be useless
- Our observation from our ARPA-e project is that scenario generation dominates the time for algorithm development in stochastic unit commitment
  - Roughly an 80%/20% split in “practice”
- There is a huge historical database of forecasted and actual observations, which can be leveraged to create accurate stochastic process models of load, wind, and (maybe) solar
  - This is operations – no need to pretend that distributions don’t exist

Load Scenario Generation: ISO-NE

- We have developed a novel technique for approximating stochastic process models using historical weather data and corresponding actual realizations
  - Based on epi-spline technologies (Wets et al.)
  - Not Monte Carlo, not AR(I)MA
  - Approximates the distributions – no sampling required!
- Accuracy of the expected day-ahead load is consistent with that generated by ISO-NE point forecasts in practice

	Segments (N_S)			
Season	1	3	5	7
Fall	5.45	4.66	4.2	3.99
Spring	3.1	2.88	2.67	2.73
Summer	10.25	4.82	4.14	4.19
Winter	5.25	3.32	3.29	3.47

← MAPEs

Illustrative Load Scenarios: ISO-NE

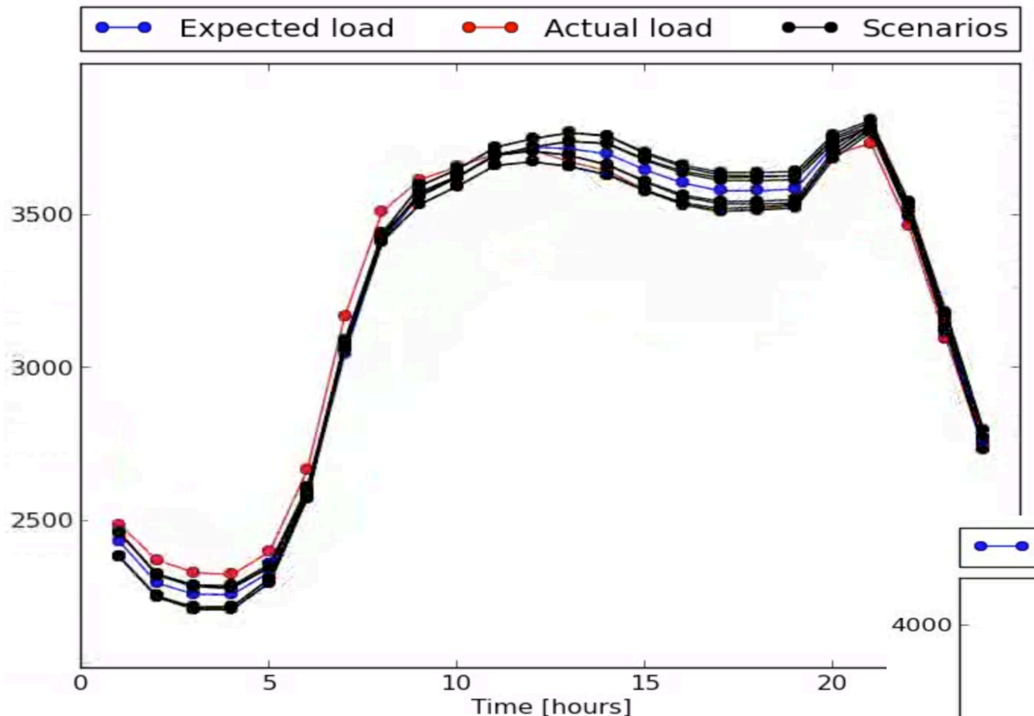

If the historical data indicates no variability, then the scenarios will reflect that consistency

Captures variability in load when present – but predictions are not perfect!

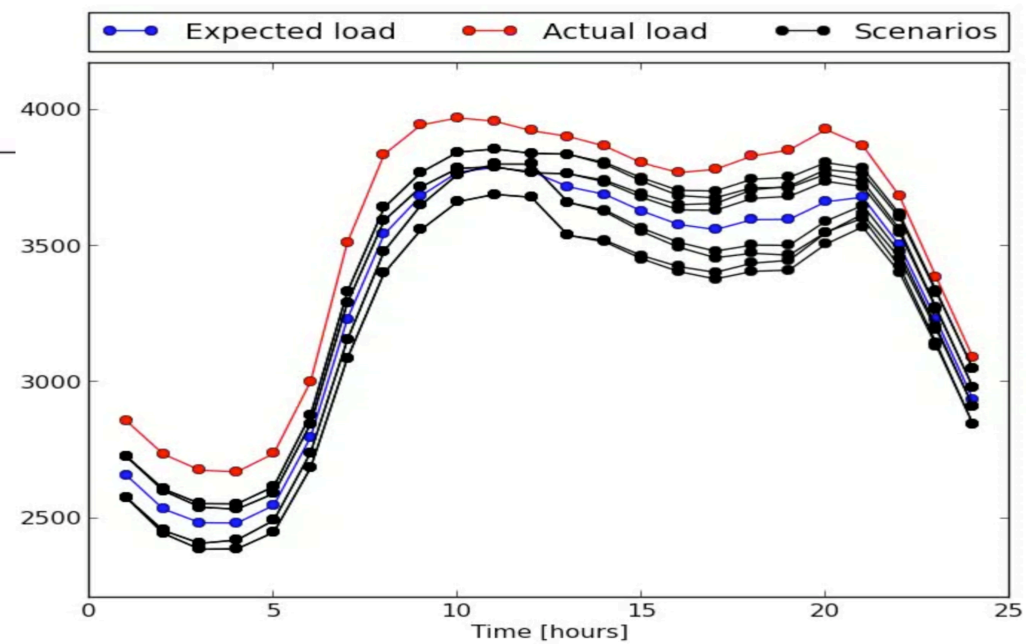

Wind Scenario Generation: BPA

Scenarios generated using Pinson et al. method

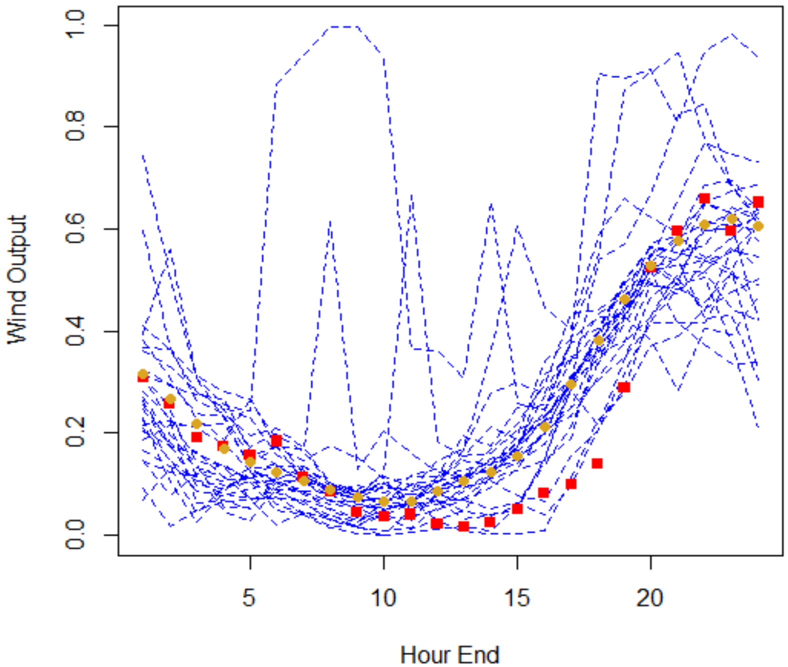

Scenarios generated using our epi-spline approach

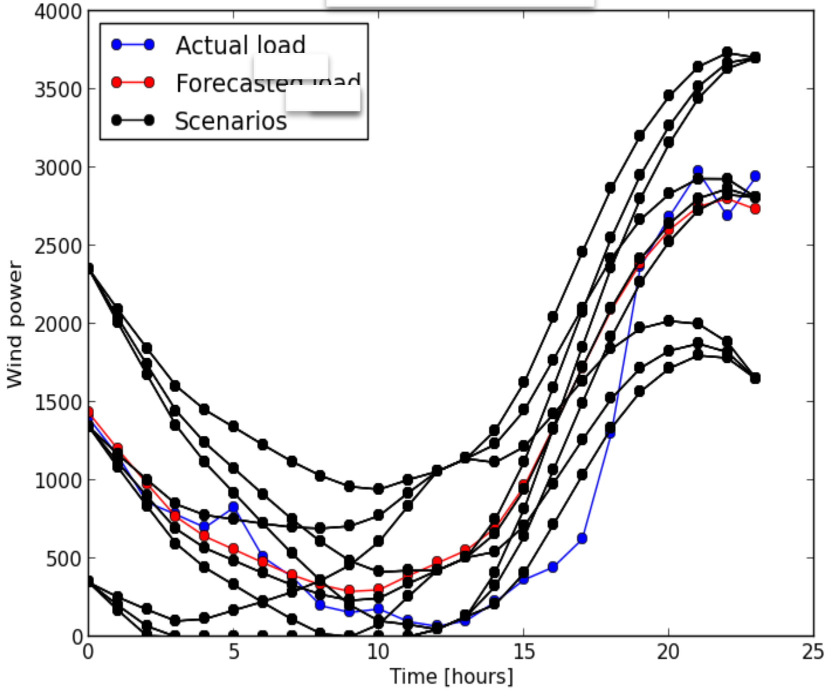

Note: Real wind profiles show significant ramps, but not as extreme as those obtained using (e.g.,) the Pinson et al. method

Scenario Generation: Discussion

- We can and should leverage the significant volume of historical data concerning load and renewables forecast / actuals
  - Arguably do ***not*** need stochastic forecasts from vendors
  - Instead, we can build stochastic models from historical point forecasts

- Stochastic process model accuracy can approach that of state-of-the-art point forecasting techniques
  - But in addition represents variability

- Approximation of stochastic process models, rather than Monte Carlo sampling, can yield significant reductions in the number of scenario required for stochastic unit commitment
  - Enabled by epi-spline-based models of stochastic load and wind

(Some) Historical Barriers to Adoption of Stochastic Unit Commitment

- *We can't create sufficiently accurate sets of scenarios to capture load and renewables uncertainty*

- ***Even if we could create accurate sets of scenarios, the resulting models are too difficult to solve***

- *Even if we could solve the resulting models, it would require significant HPC resources – which is a major impediment to industrial adoption*

On the Difficulty of Stochastic Unit Commitment: Extensive Forms

- 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
- Results

Table 3 Solution quality statistics for the extensive form of the *WECC-240-r1* instance, given 4 hours of run time.

# Scenarios	Objective Value	MIP Lower Bound	Gap %	Run Time (s)
3	64278.20	63797.72	0.75	14491
5	62740.67	62180.86	0.89	14723
10	61563.10	60835.45	1.18	14630
25	61455.55	59963.78	2.36	14960
50	61911.74	59540.87	3.83	15480
100	62388.85	59548.23	4.51	16562

Scenario-Based Decomposition via Progressive Hedging (PH)

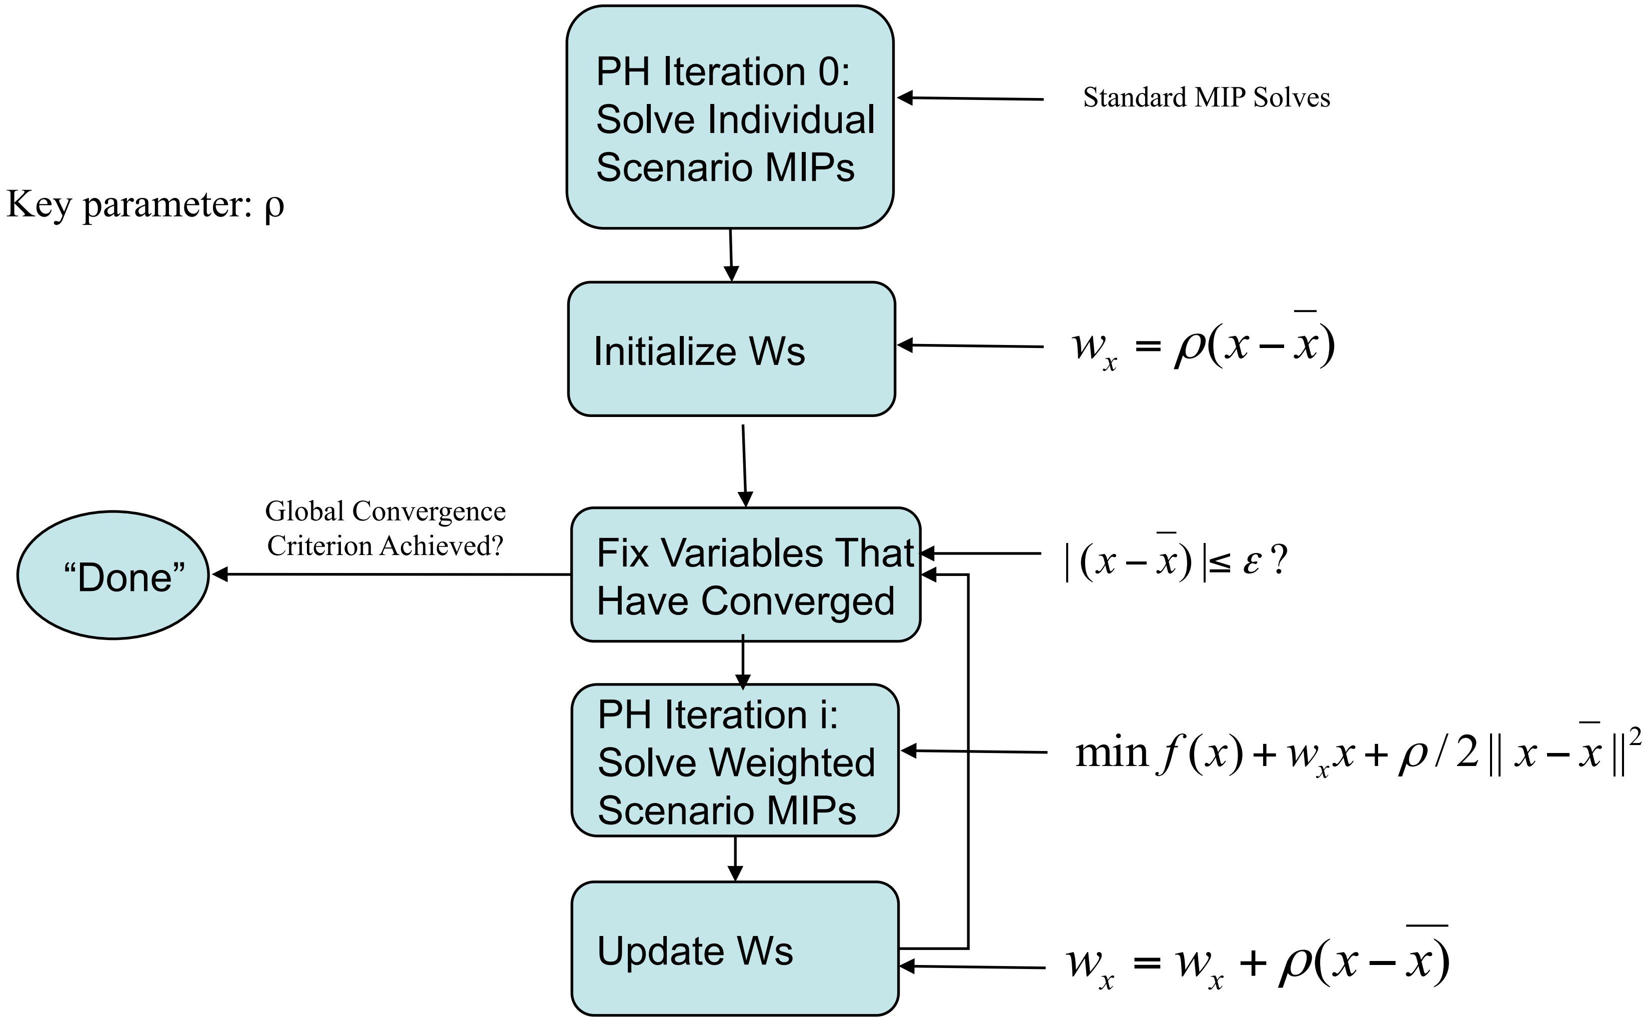

Progressive Hedging: Parallelization and Bundling

- 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

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) – 39th 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….

Progressive Hedging Results: WECC-240++

Table 7 Solve time (in seconds) and solution quality statistics for PH executing on the *WECC-240-r1* instance, with $\alpha = 0.5$, $\mu = 6$, and $\gamma = 0.025$

# Scenarios	Convergence Metric	Obj. Value	PH L.B.	# Vars Fx.	Time
		64-Core Workstation Results			
3	0.0 (20 iters)	64213.397	63235.381	4080	508
5	0.0 (in 18 iters)	62642.531	61767.253	4079	674
10	0.0 (in 35 iters)	61396.553	60476.604	4066	648
25	0.0 (in 22 iters)	60935.040	59992.622	4066	761
50	0.0 (in 15 iters)	60625.149	59631.839	4034	1076
100	0.0 (in 25 iters)	61155.387	60014.571	4080	1735
		Red Sky Results			
50	0.0 (in 16 iters)	60623.343	59779.813	4007	404
100	0.0 (in 25 iters)	61120.943	60275.744	4080	549

ISO-NE results are obtained on Red Sky on average in 15 minutes, 25 minutes in the worst case (with 100 scenarios)

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 (< 5 minutes)
- But PH doesn't provide bounds!
  - No longer true
  - Now comes with (rather tight) lower bounds
  - See "Obtaining Lower Bounds from the Progressive Hedging Algorithm for Stochastic Mixed-Integer Programs" (Under review)
- More seriously
  - We have a lot of work going on in the realm of lower bounding, which we are happy to discuss off-line

(Some) Historical Barriers to Adoption of Stochastic Unit Commitment

- *We can't create sufficiently accurate sets of scenarios to capture load and renewables uncertainty*
- *Even if we could create accurate sets of scenarios, the resulting models are too difficult to solve*
- ***Even if we could solve the resulting models, it would require significant HPC resources – which is a major impediment to industrial adoption***

On HPC and Stochastic Unit Commitment…

- We observe that stochastic unit commitment solvers ***do not*** require HPC for execution on industrial scale problems
  - Commodity clusters are sufficient for many analyses
  - Execution on the cloud (e.g., Gurobi with Amazon EC2) is feasible

- There is little evidence that hundreds of thousands to millions of scenarios are required for stochastic unit commitment
  - Approximation of stochastic process models can avoid scalability issues associated with Monte Carlo approaches

On Bounding, Incumbents, and Stochastic Unit Commitment

- Assertion
  - We are never going to solve large-scale stochastic unit commitment problems to optimality in operational contexts
  - Small gaps (1-2%) are sufficient
  - If you're comfortable with reserve margins…

- There is much discussion in the research literature of bounding versus finding high-quality incumbents
  - We don't think the emphasis is where it needs to be

- Our findings
  - Finding high-quality lower bounds is relatively easy
  - Locating high-quality incumbents (solutions) is very difficult

Our Next Steps for Stochastic Unit Commitment (1)…

- We feel that a "continual commitment" process is the ultimate conclusion of stochastic unit commitment research
  - But what does this entail?
- More emphasis on multi-stage stochastic unit commitment
  - Enables more flexibility, required for minimal discontinuities between stages
  - Algorithmic extensions are conceptually straightforward, but in reality…
  - Non-trivial research questions relating to evaluation, comparison
- Strong focus on incremental forecast update technology
  - Present models exhibit discontinuities at (artificial) decision boundaries
  - Allows commitment processes to take advantage of information as it becomes available

Our Next Steps for Stochastic Unit Commitment (2)…

- How do we effectively set endogenous and dynamic reserve levels in stochastic unit commitment?
  - Reserves will always exist, but at a significantly reduced level
- Analysis of scenarios should provide some indication of the degree of reserves required
  - E.g., lower variability => lower reserves
- Scenarios capture aspects of uncertainty, but not everything
  - Understanding the remaining degrees of freedom (relating to modeling assumptions) will drive new methods for reserve quantification

Our Next Steps for Stochastic Unit Commitment (3)…

- Operations is a unique situation for stochastic programming
  - We are repeating the game, with no risk of ruin
  - Can we leverage any information from prior to solves to accelerate time-to-solution?

- Even in deterministic unit commitment…
  - Can we mine historical information to accelerate branch and cut?

- In stochastic unit commitment…
  - Can we mine similar days from the past to provide good initial conditions for algorithms (e.g., weights in PH)?
  - Pilot experiments indicate order-of-magnitude run-times are possible

Final Thoughts

- Even after demonstrating operational feasibility, there remain significant barriers to adoption of stochastic unit commitment
- The biggest open issues involve human and social dimensions
  - How to effectively communicate a stochastic solution to an operator?
  - How to analyze / interpret operator perceptions of risk in an algorithmic context?
  - …
- Doesn't mean there isn't research – but it's different research
  - Not necessarily algorithmic, but equally important for T2M

(Some) References

- Y. Feng, I. Rios, S.M. Ryan, K. Spurkel, J.P. Watson, R.J.B. Wets, and D.L. Woodruff. ***Toward Scalable Stochastic Unit Commitment – Part 1: Load Scenario Generation***. Under review, available from Optimization Online.

- K. Cheung, D. Gade, C. Silva-Monroy, S.M. Ryan, J.P. Watson, R.J.B. Wets, and D.L. Woodruff. ***Toward Scalable Stochastic Unit Commitment – Part 2: Solver Configuration and Performance Assessment***. Under review, available from Optimization Online.

QUESTIONS