Advances in DOASA[1]

Andy Philpott
Electric Power Optimization Centre
University of Auckland.
www.epoc.org.nz

Joint work with Lea Kapelevich and Hamish Mellor

(acknowlegdements to Geoff Pritchard, Jerome Clavel, Ziming Guan and Oscar Dowson)

[1] EPOC Winter Workshop, Auckland, September 7, 2016

What is DOASA?

DOASA stands for

Dynamic Outer Approximation Sampling Algorithm.

DOASA is an implementation of SDDP in C++ developed by Geoff Pritchard and Andy Philpott.

EMI-DOASA is the version of DOASA available on the Electricity Authority EMI site for free use in New Zealand. There are CLP and Gurobi versions.

What is JADE?

JADE stands for

Just Another DOASA Environment

JADE is an implementation of DOASA developed by Lea Kapelevich in the scripting language Julia. It uses the Julia packages StochDualDynamicProgram developed by Oscar Dowson, JuMP developed by Iain Dunning and Miles Lubin at MIT, and any LP solver (CLP, Gurobi, CPLEX, XPRESS).

(Everybody named here except Miles Lubin is a current or past student in Engineering Science at Auckland).

Summary

Summary

Water release policy

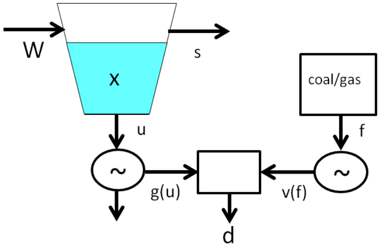

Find a water release policy over next 52 weeks that minimizes the expected thermal fuel cost of meeting demand over this period. Unsatisfied demand is met by shedding load at a shortage cost. A policy is not a fixed plan - it is a decision rule that says at the beginning of each week “given previous observations and actions, do the following”. This means that water release actions are not predetermined but will depend on observations.

Mathematical description

$$
\begin{aligned}
\min \quad & \mathbb{E}\left[\textstyle\sum_{t=1}^{T} c_t^\top v_t + V_{T+1}(x_{T+1})\right] \\
\text{s.t.} \quad & g(u_t) + v_t = d(t), \\
& x_{t+1} = x_t - A(u_t + s_t) + W_t, \\
& x_1 = \bar{x}, \\
& u_t \in \mathcal{U}, \quad v_t \in \mathcal{V}, \; s_t \geq 0, \qquad t = 1, \ldots, T.
\end{aligned}
$$

Expected marginal water value for a single reservoir

- Water release policies for a single-reservoir can be specified by an expected marginal water value measured in $\$/m^3$ This is defined for each possible reservoir level, and is the opportunity cost of releasing one unit of water.
- The optimal policy generates from all thermal plant that are no more expensive (in \$/MWh) than the expected marginal water value multiplied by a conversion factor (cubic metres per MWh).
- The water release policy solves a stage optimization problem: minimize the current cost of thermal generation plus the expected future cost of meeting demand. The slope of this future cost function at level x is (minus) the marginal water value at this level.

Optimal solutions and competitive market outcomes
(Philpott, Ferris and Wets, 2016)

- Water release policies from system optimization give a social planning solution.
- Social planning models correspond to perfectly competitive partial equilibrium when:
  - market participants act as risk-neutral price takers;
  - markets are complete;
  - market participants problems are convex;
  - market participants agree on probability distributions of random outcomes.
- When market participants are risk averse then social planning models correspond to perfectly competitive partial equilibrium when additionally:
  - risk measures are coherent and comparable;
  - markets for risk are complete.

Summary

Stochastic Dual Dynamic Programming (SDDP)

Computes a water release policy when there are a number of different reservoirs with (possibly) different locations, different inflow processes and different river systems. This can give different results from aggregated reservoirs (see e.g. Philpott, Dallagi & Gallet, 2013). Optimal release is no longer determined by a single marginal water value. The optimal releases u given storages x and next period's inflow W, is computed by the recursive stage problem:

$$\begin{aligned} V_t(x) = \mathbb{E}[\min \quad & c_t^\top v + V_{t+1}(x - A(u+s) + W)] \\ \text{s.t.} \quad & g(u) + v = d(t), \\ & u \in \mathcal{U}, \quad v \in \mathcal{V}, \; s \geq 0. \end{aligned}$$

Cuts

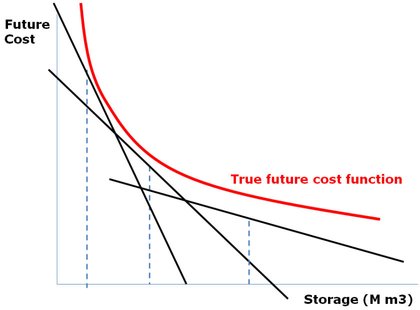

We need to define the function V_{t+1}. This is never known exactly but is approximated by the pointwise maximum of linear functions of x that are called cuts. SDDP algorithms like DOASA update these cuts using sampling procedures. For the New Zealand system about 5000-10000 cuts are needed in each week to define V_t accurately enough. This can take a lot of computation.

Bounds

- At each iteration of the algorithm SDDP gives a lower bound on the expected cost of an optimal policy. This increases as the algorithm proceeds.
- At each iteration of the algorithm SDDP gives a candidate policy (defined by the cuts). This can be simulated to estimate its expected cost and see how close it is to the current lower bound.
- Solutions are judged to be optimal (enough) when a confidence interval of their expected cost is close to the lower bound.

Inflow Models

- Inflow models preserve spatial dependence between catchments
- Stagewise independent inflows sampled from historical record (the default for EMI-DOASA).
- Inflows sampled independently from historical values with variance increased to model correlation (DIA model in EMI-DOASA).
- PARMA models: simplest is AR1 model. (PSR's SDDP and CEPEL's NEWAVE)
- PARMA models for logarithm of inflows (AR model avoiding negative inflows).
- Markov Chain model of climate state and inflows. (Philpott & de Matos, 2012)

Risk-averse SDDP
(Philpott, De Matos and Finardi, 2013)

- SDDP assumes sampled inflow outcomes are equally likely.
- We can bias sampling towards extremes in some forward passes as along as sampling is eventually independent (for convergence). This still gives a policy that minimizes expected cost
- We can bias in the backward pass (with m outcomes in each stage) by inflating the probabilities of low inflows to be $> \frac{1}{m}$.
  - the resulting policy no longer minimizes expected cost.
  - the resulting policy can be interpreted as being optimal while robust to variations in estimates of the inflow distribution.
  - with appropriate probability adjustment, the policy can be interpreted as a risk-averse policy with a dynamic coherent risk measure.

Summary

The EMI-DOASA model

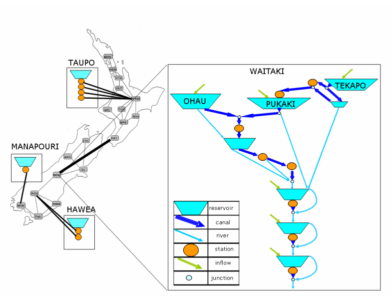

Approximate network representation of New Zealand electricity network showing main hydro-electricity generators.

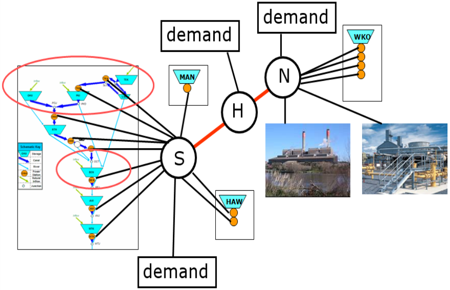

EMI-DOASA model has three demand nodes and up to seven state variables corresponding to seven storage reservoirs.

Modelling demand

- EMI-DOASA has three nodes and three load blocks in each week (peak shoulder and offpeak). These currently are hard coded in the software.
- Demand in each node can be estimated for historical backtesting by running vSPD for every trading period and aggregating the generation in each region (plus the net flow into the region in all vSPD arcs joining regions).
- For runs of DOASA in future years historical data can be used to forecast load blocks (either by scaling of existing load blocks, or by removing baseload e.g. if Rio Tinto smelter were to close).

Shortage Costs

Energy deficit in any stage is met by load shedding at an increasing shortage cost in three tranches. This is equivalent to having three dummy thermal plant at each location with capacities equal to 5% of load, 5% of load and 90% of load, for each load sector, and costs as follows

	Up to 5%	Up to 10%	VOLL	North Is	South Is
Industrial	$1,000	$2,000	$10,000	0.34	0.58
Commercial	$2,000	$4,000	$10,000	0.27	0.15
Residential	$2,000	$4,000	$10,000	0.39	0.27

Load reduction costs (NZD/MWh) and proportions of load that is industrial, commercial, and residential in each island.

Summary

JADE
(Based on Year 4 project work of Lea Kapelevich)

JADE stands for

Just Another DOASA Environment

JADE is an instance of the SDDP package in the scripting language Julia. JADE was developed by Lea Kapelevich . It was developed completely independently from EMI-DOASA, but adopts the same modelling assumptions and reads the same input files. Because of its use of the JuMP modelling language, JADE can be easily configured to verify EMI-DOASA code changes and detect bugs.

JADE and EMI-DOASA lower bounds

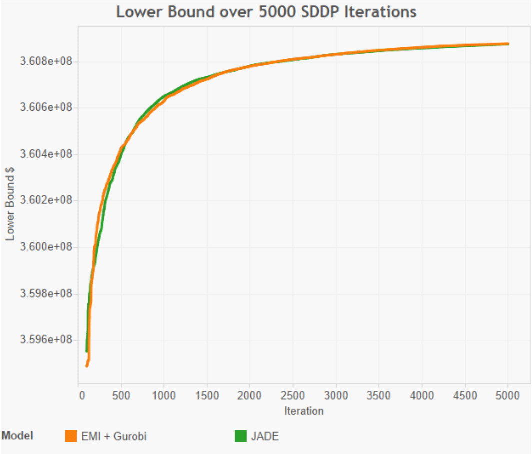

JADE and DOASA follow the same trajectories in computing lower bounds for the expected cost of the optimal policy.

Simulated value of policies

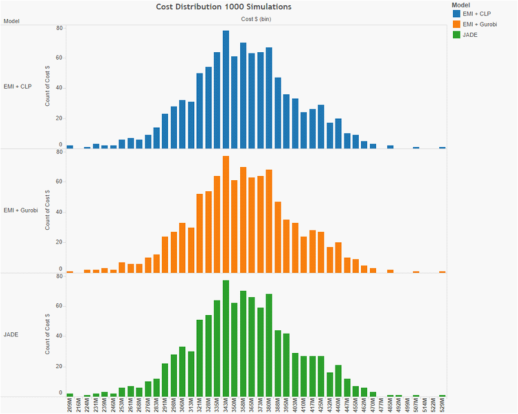

Policies give essentially the same distributions when simulated with Monte Carlo inflows.

Historical simulations

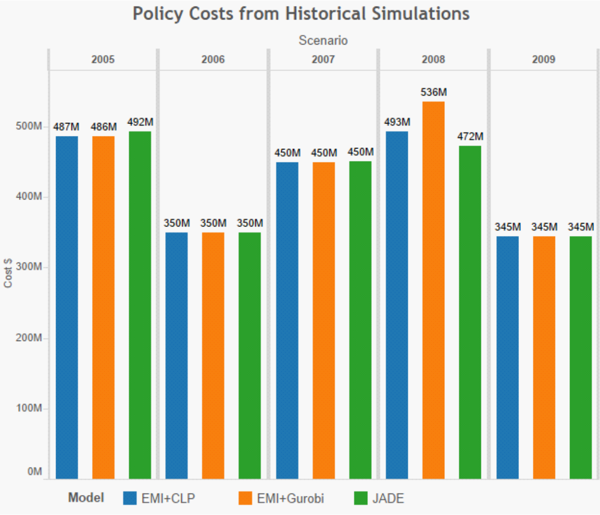

Solution times

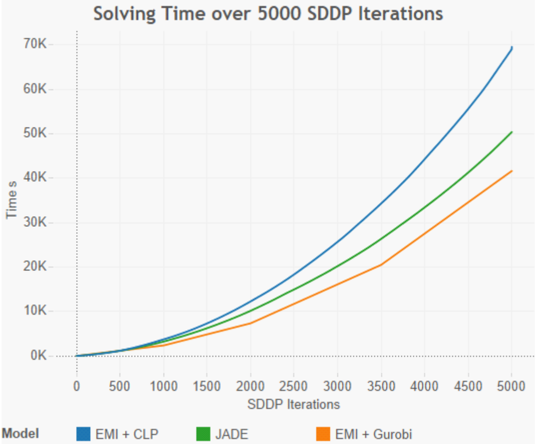

JADE using the Gurobi solver is faster than EMI-DOASA using CLP, but a bit slower than EMI-DOASA using Gurobi.

Cut selection
(De Matos, Philpott and Finardi, 2015)

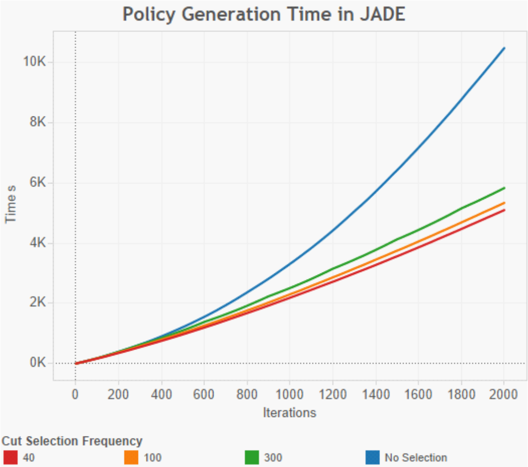

In each stage optimization, Level 1 dominance cut selection ignores dominated cuts at visited states (but restores them at the end).

Cut selection policy evaluation

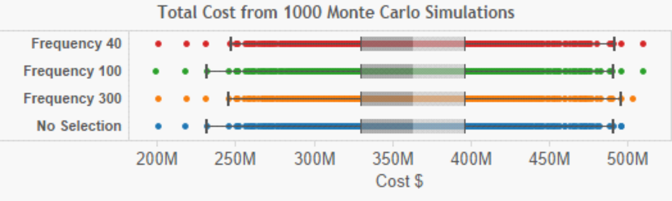

Effect of cut selection on simulated policy value.

Historical simulations with cut selection

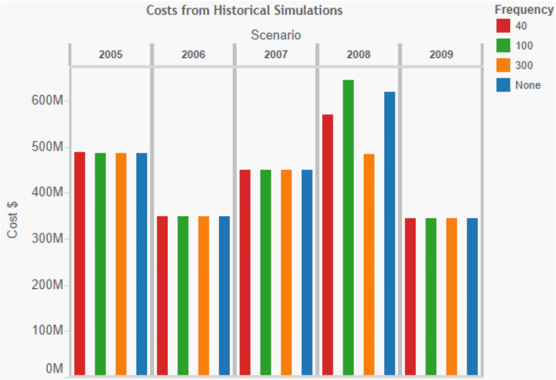

Cut selection affects policy in dry year.

Tests with different inflow models

Inflow model	Scenario cost (NZ)$M				
	2005	2006	2007	2008	2009
Independent	491.90	349.62	450.57	553.69	345.51
AR1	486.22	351.12	452.66	454.45	348.43
Log AR1	489.97	349.71	450.87	600.79	343.64
DIA 2	485.40	349.77	450.54	453.41	346.14
DIA 5	484.00	350.00	451.00	516.00	349.00
Risk Averse	484.82	349.68	451.47	447.09	345.20

JADE policy computed with 5000 cuts and simulated in historical years.

Test with more load blocks

Load Blocks	3	6	9
Lower Bound ($M)	356.17	359.21	359.27
CI Lower	353.54	356.57	356.65
CI Upper	359.59	362.60	362.70
Mean	356.57	359.58	359.67
Hours to solve	1.92	3.58	15.20

Results of JADE with 4000 cuts, cut selection and different numbers of load blocks in each week.

Test on 18-node transmission system

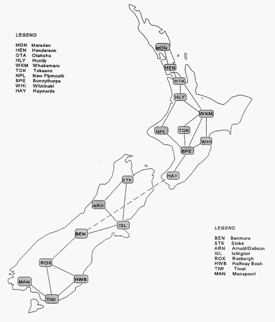

Model	Lower	CI lower	CI upper	Mean	Time(h)
3 nodes	339.98	337.79	343.79	340.79	0.77
3 nodes + ∞	339.98	337.87	343.88	340.87	1.60
18 nodes	341.23	338.81	344.78	341.80	1.18

JADE run for 1000 iterations with cut selection every 200. Results show lower bound and simulated value of policy.

Summary

Backtesting dry years
(Year 4 project work of Hamish Mellor)

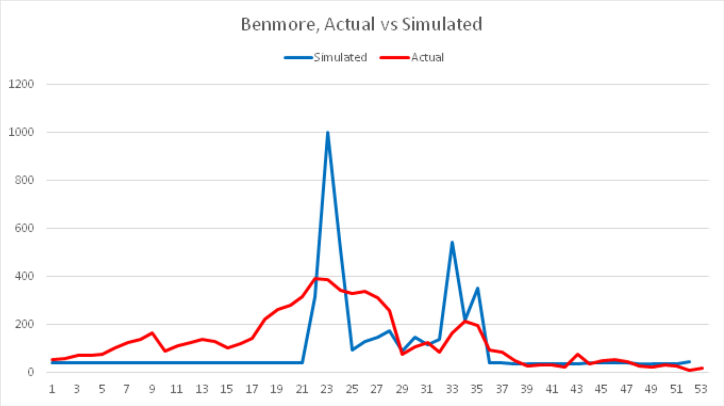

Historical weekly average prices at Benmore in 2008 compared with DOASA predictions at start of year.

Marking up thermal offers

GENERATOR	HEAT_RATE	CAPACITY	$/MWh (coal)
Huntly_main_g1_a	1	95	4
Huntly_main_g1_b	30	44	120
Huntly_main_g1_c	50	41	200
Huntly_main_g1_d	75	40	300
Huntly_main_g1_e	100	40	400

Each large thermal unit can be represented in DOASA as several with increasing heat rates.

2008 price comparison

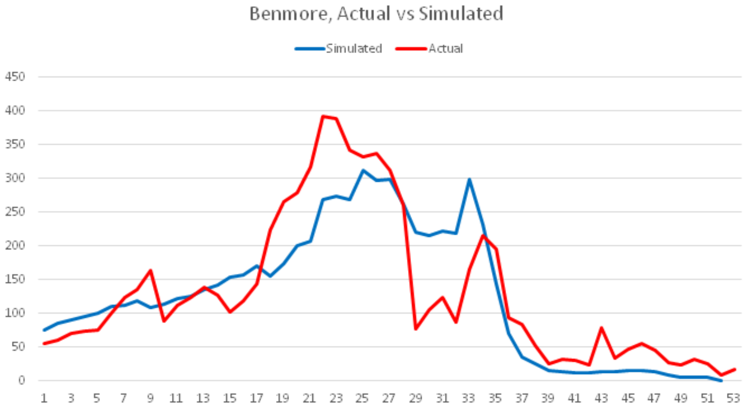

2012 comparison

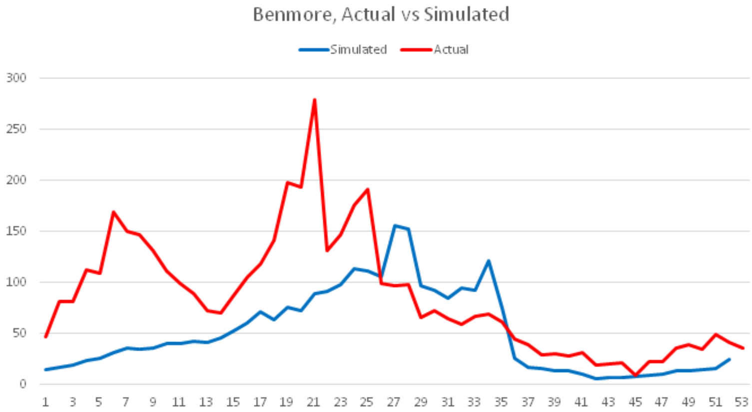

Simulated time average prices in 2012
(Year 4 project work of Hamish Mellor)

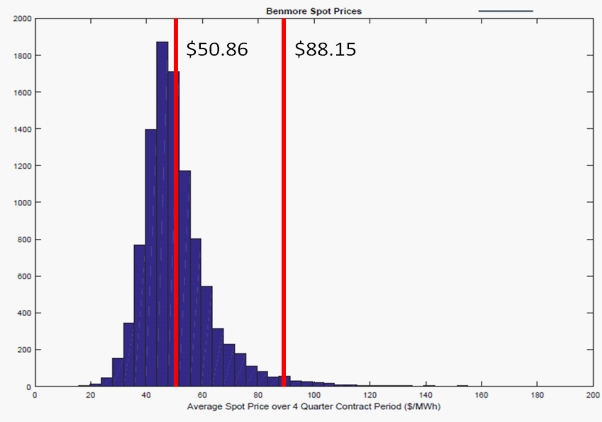

Benmore time average price simulations using DOASA policy for 2012. Highest time average price is $450/MWh, but expected price is $50.86/MWh. CFDs traded 1/1/2012 at $88.15/MWh.

Rooftop solar evaluation
(Joint work with Jerome Clavel, EPFL)

- Run EMI-DOASA with different assumptions on level of rooftop solar installations.
- Use NIWA irradiation figures and panel efficiencies to reduce demand in trading periods and times of year.
- Estimate panel costs at (NZ)$500-600 per square metre (accounting for inverters).
- This amounts to $15/m2/year over 35-40 year life.

EMI-DOASA results

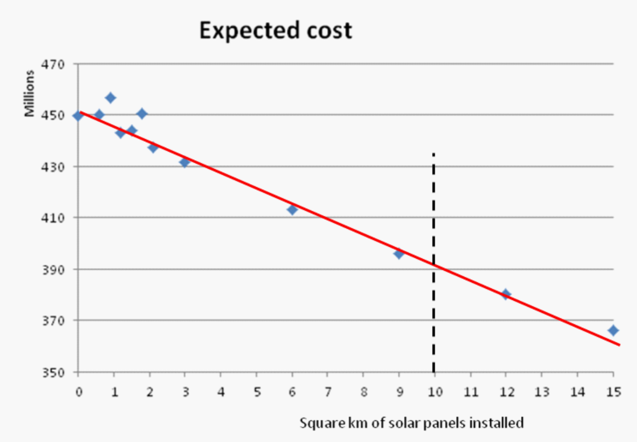

Assuming 2008 demand, 1 km^2 of panels decreases expected cost by $6 M. Current capital cost of panels and inverters = $15M/year ($\$15/m^2/year$ assuming 35 years life).

The End

THE END

References (downloadable from www.epoc.org.nz)

Philpott, A.B., Ferris, M.C. and Wets, R.J-B., Equilibrium, uncertainty and risk in hydrothermal electricity systems, *Mathematical Programming B*, 157,2, 483-513, 2016.

De Matos, V., Philpott, A.B. and Finardi, E.C., Improving the performance of Stochastic Dual Dynamic Programming, *Journal of Computational and Applied Mathematics*, 290, 196-208, 2015.

Philpott, A. B., Matos, V. L. D., and Finardi, E. C., On Solving Multistage Stochastic Programs with Coherent Risk Measures. *Operations Research*, 61, 957-970, 2013.

Philpott, A.B. and de Matos, V.L., Dynamic sampling algorithms for multi-stage stochastic programs with risk aversion, *European Journal of Operational Research*, 218, 470-483, 2012.

Philpott, A.B., Dallagi, A. and Gallet, E., On cutting-plane algorithms and dynamic programming for hydro-electricity generation, in *Risk Management and Energy Trading*, R.M. Kovacevic., G.Ch. Pflug, and M.T. Vespucci (ed.) Springer, 105-127, 2013.