Co-optimization of battery storage over multiple revenue streams and time scales

Federal Enegy Regulatory Commission
Workshop on Increasing Market and Planning Efficiency Through Software

June 21, 2015

Warren B. Powell
Harvey Cheng
PENSA Laboratory
Princeton University
http://energysystems.princeton.edu

The co-optimization challenge

- Revenue streams for battery storage
  - » Frequency regulation
  - » Energy arbitrage
    - In front of the meter (what we have done)
    - Behind the meter
      - – Customer baseline load
  - » Peak shifting
    - Utility – flattening the load curve
    - Customer perspective – demand charge management
      - – Commercial scale
  - » Customer baseline load
    - Behind the meter
  - » Power factor correction
    - Similar to voltage regulation

The co-optimization challenge

- Revenue streams for battery storage
  - » Capacity market
    - Backup; 8 hours out of a year;
    - Historically summer peaks, but can happen any time
    - Likely to be a function of temperature (high/medium/low)
  - » Spinning reserve
    - Hourly market; day-ahead bidding, clears that day; dispatched for 15 minutes; dispatched when a reserve event occurs; zonal-based
  - » Power factor correction
  - » Voltage regulation market (emerging)
    - Signal will specify watts (real) vs. var (reactive) power.
  - » Life cycle management
  - » Emerging markets for variability of solar

The co-optimization challenge

- Pure frequency regulation
  - » Follow the RegD signal from the ISO
  - » Penalties for noncompliance
    - At PJM, these are computed hourly
  - » Noncompliance is due purely to engineering limitations
- Energy arbitrage
  - » Simple control law

$$X^{\pi}(S_t \mid \theta) = \begin{cases} +1 & \text{if } p_t < \theta^{\text{charge}} \\ 0 & \text{if } \theta^{\text{charge}} < p_t < \theta^{\text{discharge}} \\ -1 & \text{if } p_t > \theta^{\text{charge}} \end{cases}$$

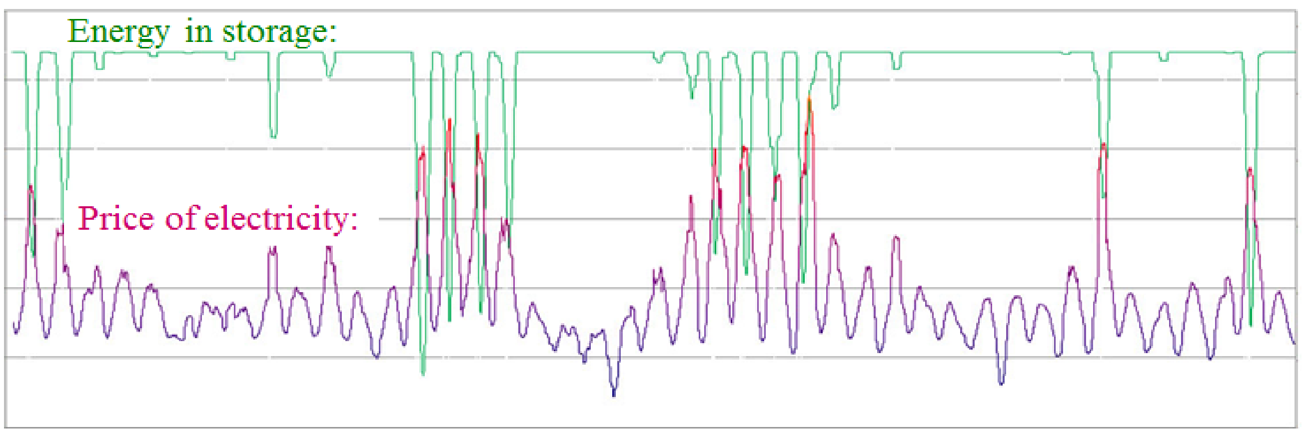

The co-optimization challenge

- Optimizing both at the same time
  - » We may wish to deviate from the PJM signal, trading off loss of FR revenue now to make more money:
    - By selling into a high LMP
    - By holding (or purchasing) energy to use later.
  - » Co-optimization requires making cost/revenue tradeoffs
    - Now…
    - … or in the future
  - » This requires that we formulate the problem as a cost minimization problem over time.
  - » How to solve?

The co-optimization challenge

- Strategy:
  - » Formulate problem as a cost-minimization problem over time
  - » Solve as a Markov decision process
- Challenge
  - » We have to handle multiple time (and spatial) scales:
    - Frequency regulation, with decisions made every 2-4 seconds
    - Hourly planning in 5-minute increments to respond to variations in LMPs, with hourly compliance evaluations
    - Daily planning over the grid, handling both daily cycles and grid congestion.
  - » Problem:
    - 21,600 time periods (4-second increments over a day)
    - 4 dimensional state variable to cover all periods

A multiscale decomposition strategy

- We break the horizon into three time scales

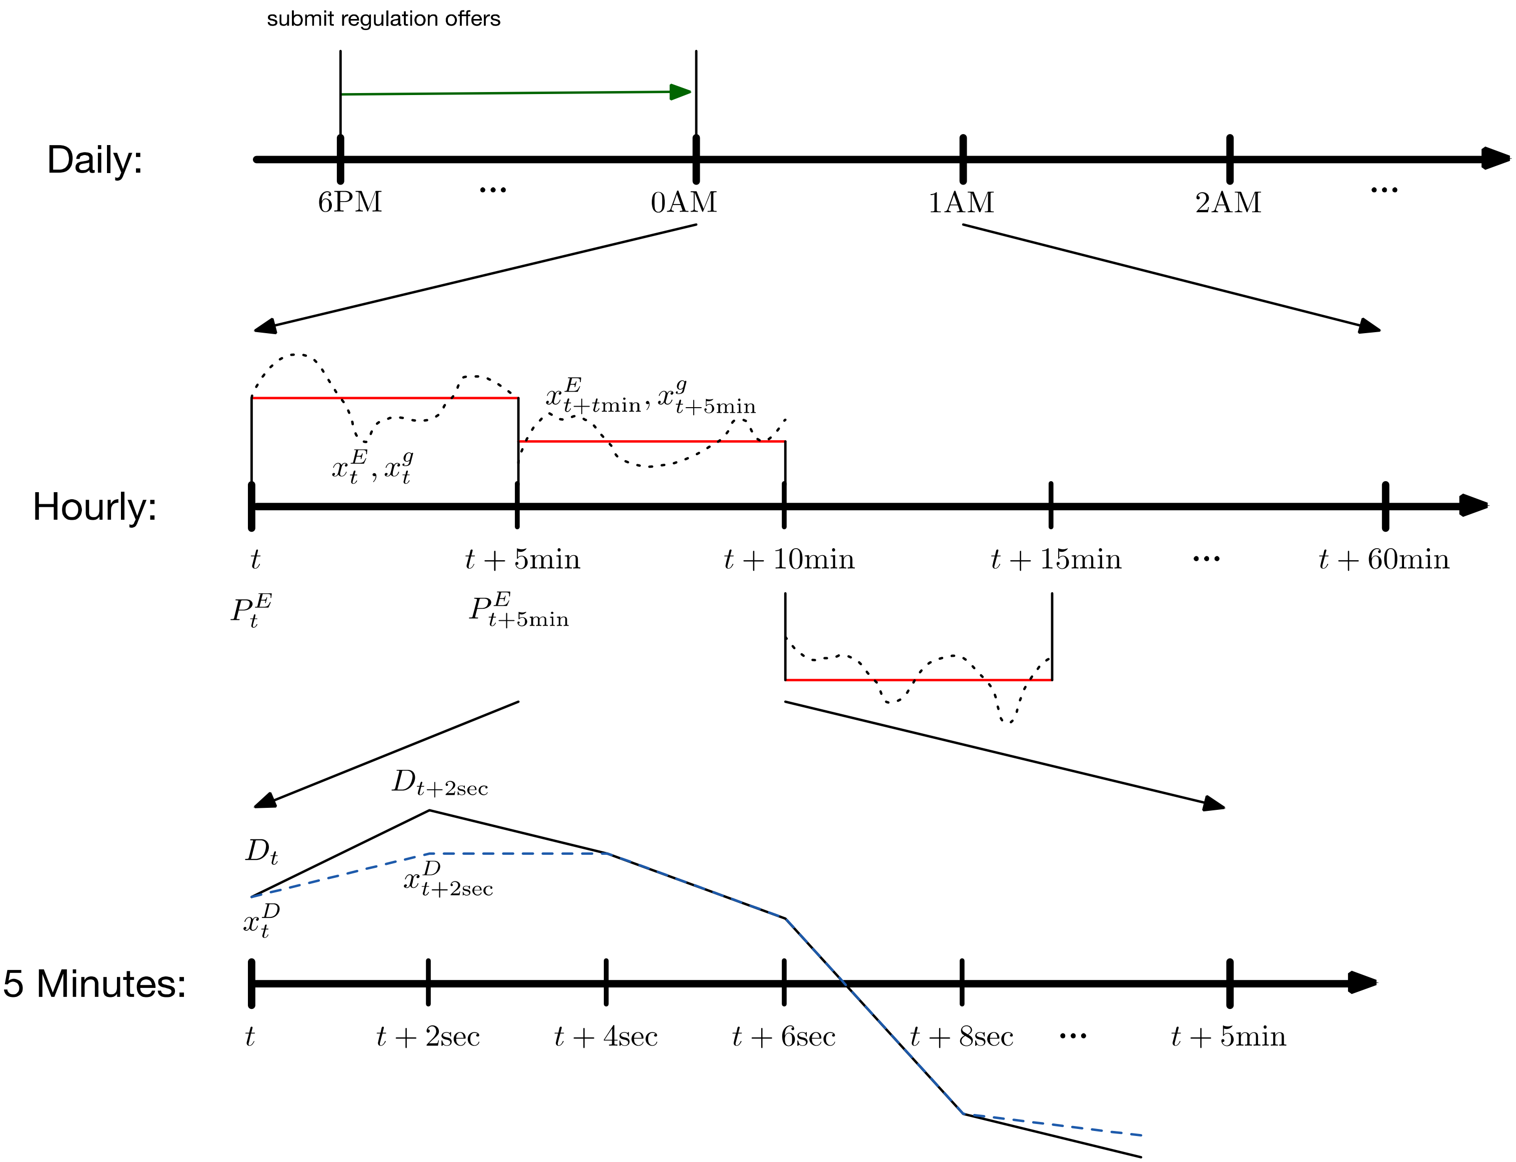

$$S_t = R_t$$

24 time steps

$$S_t = (R_t, G_t, P_t^E)$$

x_t^E = Economic basepoint

12 time steps

$$S_t = (R_t, G_t, D_t)$$

x_t^D = Frequency regulation

75 time steps

The co-optimization challenge

- Algorithmic strategy:
  - » We will use SMART-Storage to solve the grid level storage problem (developed by Daniel Salas)
    - 5-minute increments over an entire day
    - Can handle single storage, or hundreds of batteries
    - Simultaneously models the grid and economic dispatch
  - » We then solve the hourly problem
    - 5-minute increments
    - Use SMART-Storage to capture the value of energy in the battery at the end of each hour
    - Hourly model captures compliance signal in the state variablae
  - » Finally we solve the frequency regulation problem
    - 5-minute horizon in 4-second increments
    - Hourly problem gives the value of energy storage at the end of 5-minutes
    - Captures LMP in the state variable

An energy storage problem

- Consider a basic energy storage problem:

A storage problem

- Solve using Bellman's optimality equation

$$V(S_t) = \min_{x_t \in \mathcal{X}} \left(C(S_t, x_t) + \gamma \mathbb{E}\left\{ V(S_{t+1}(S_t, x_t, W_{t+1})) \mid S_t \right\} \right)$$

$$\begin{bmatrix} R_t \\ G_t \\ D_t \\ P_t^E \\ D_t \end{bmatrix} \qquad x_t \qquad \begin{bmatrix} \delta P_{t+1}^E \\ \delta D_{t+1} \end{bmatrix}$$

Approximate dynamic programming

- Optimizing grid level storage over 24 hours
  - » We approximate the value of energy in storage:

$$V_t(S_t) = \min_{x_t \in \mathcal{X}} \left(C(S_t, x_t) + \gamma \bar{V}_t^x \left(S_t^x(S_t, x_t) \right) \right)$$

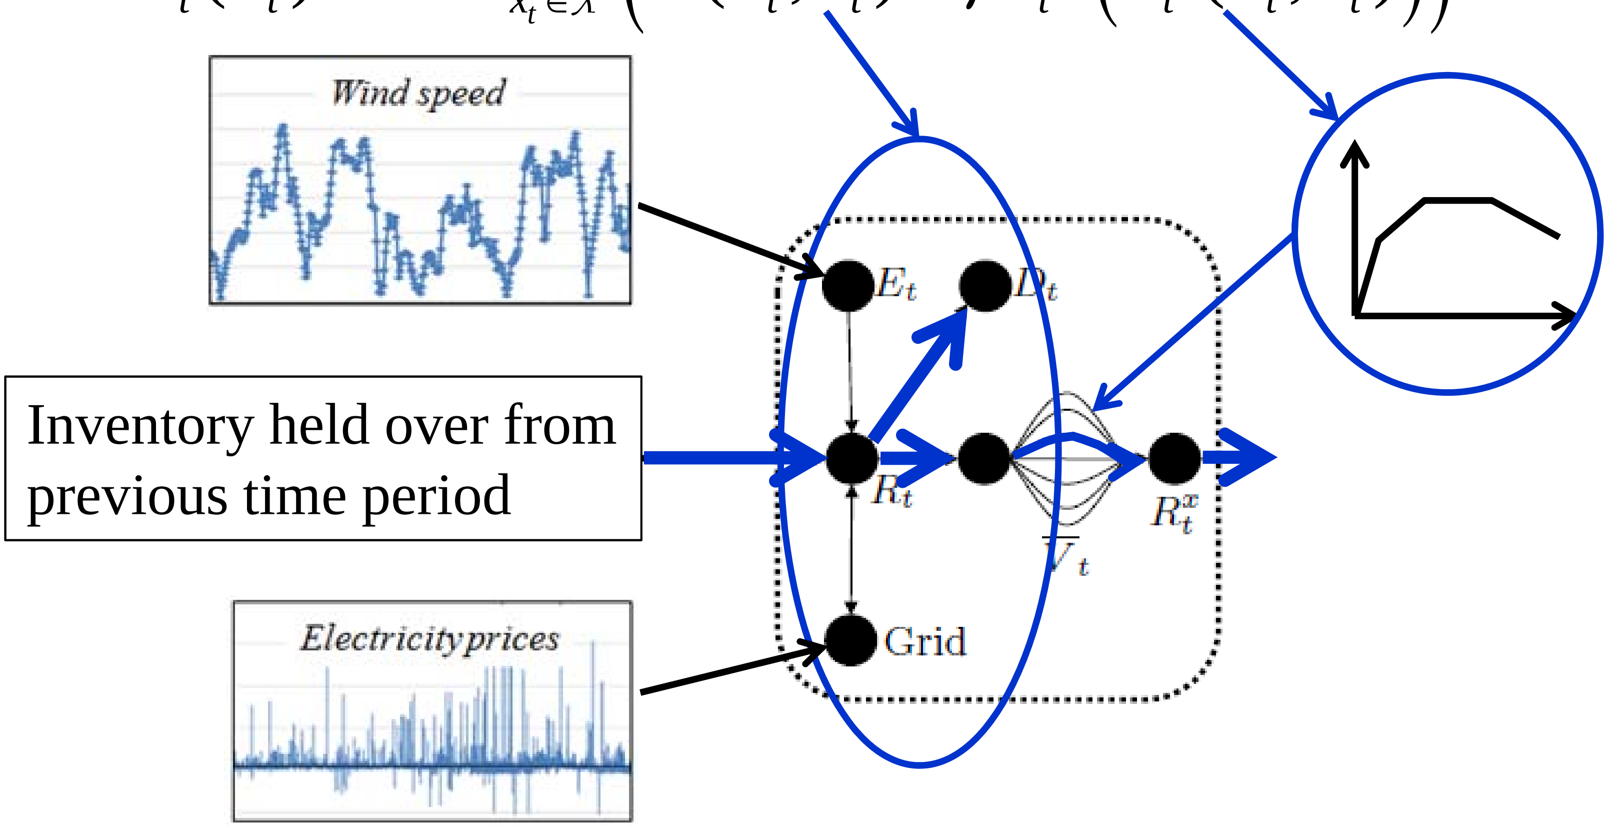

Approximate dynamic programming

- Optimizing grid level storage over 24 hours
  - » We update the piecewise linear value functions by computing estimates of slopes using a backward pass:

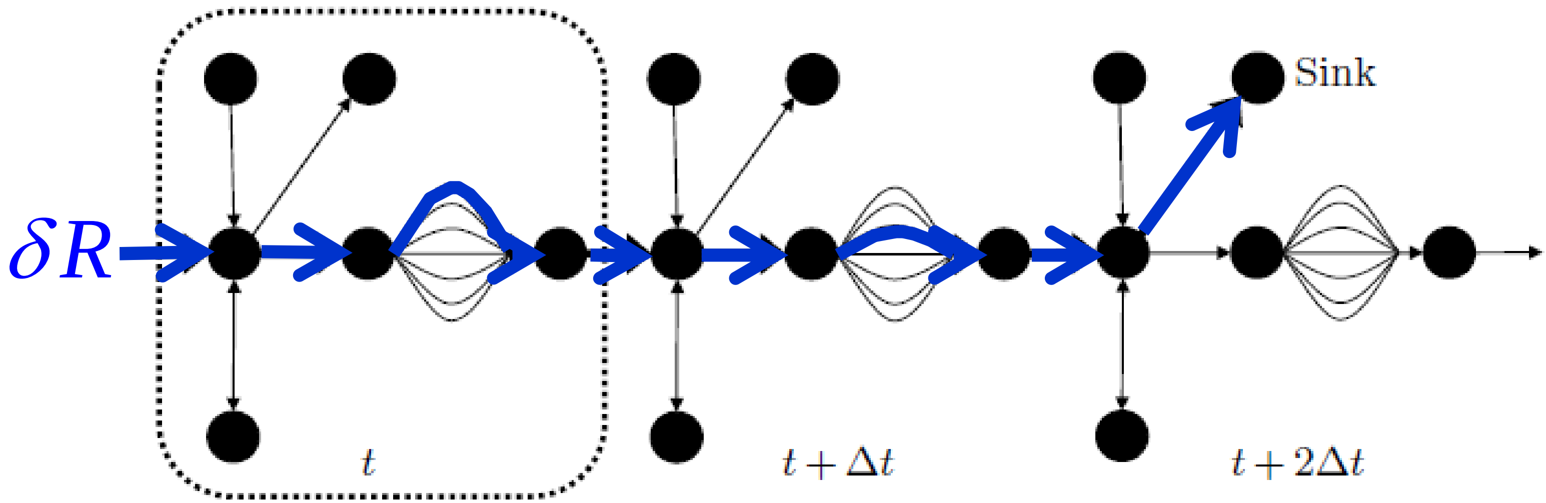

  - » The cost along the marginal path is the derivative of the simulation with respect to the flow perturbation.

Approximate dynamic programming

- Derivatives are used to estimate a piecewise linear approximation

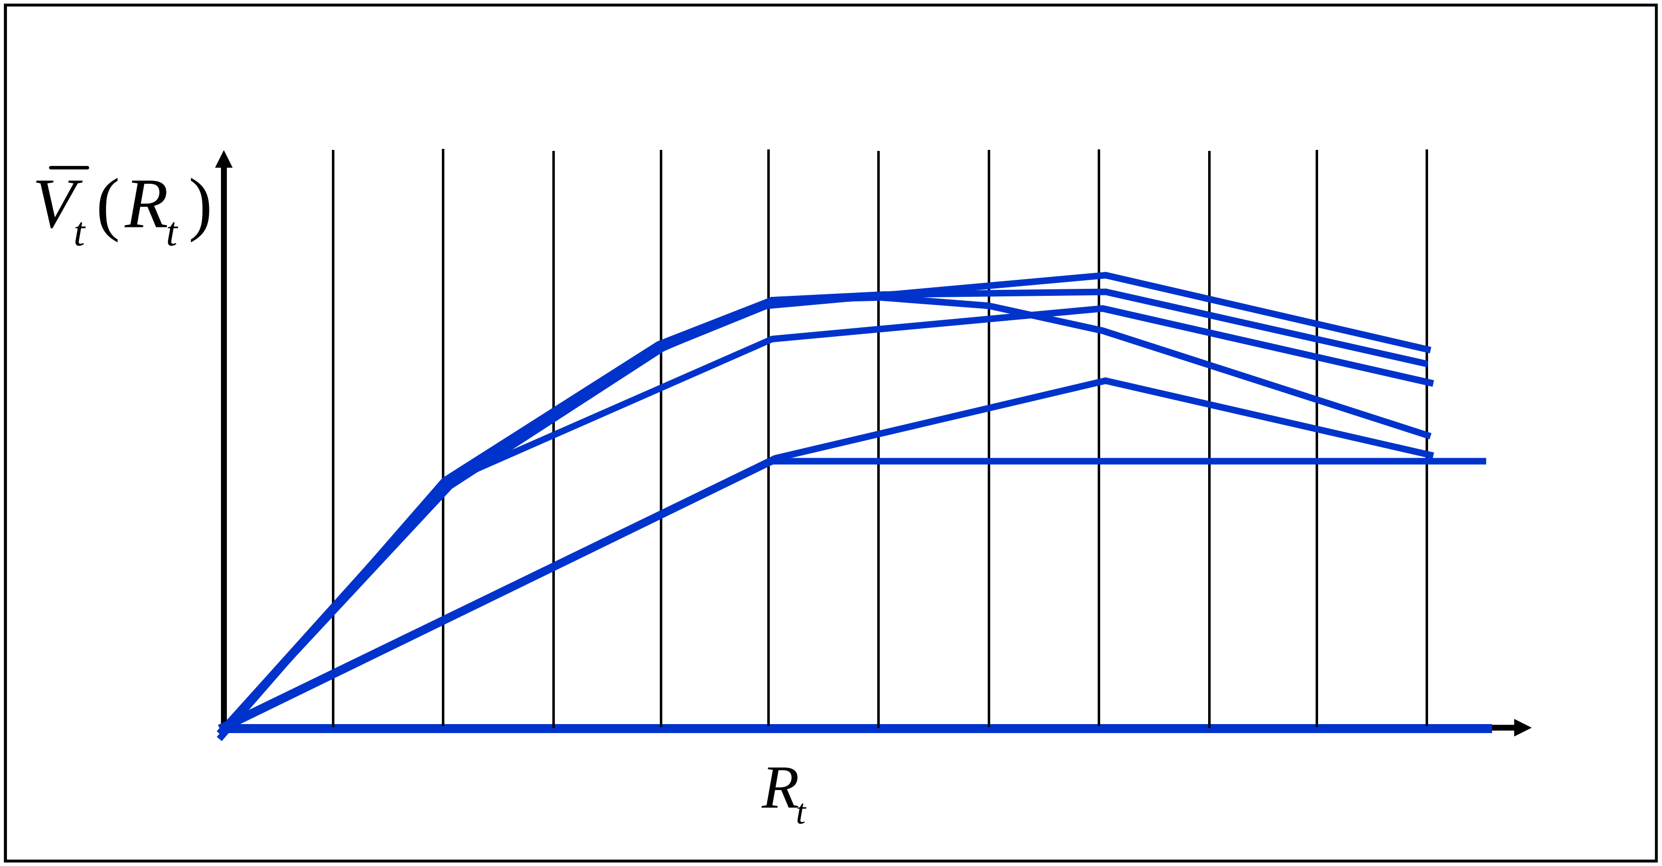

Approximate dynamic programming

- Benchmarking against optimality on a stochastic model

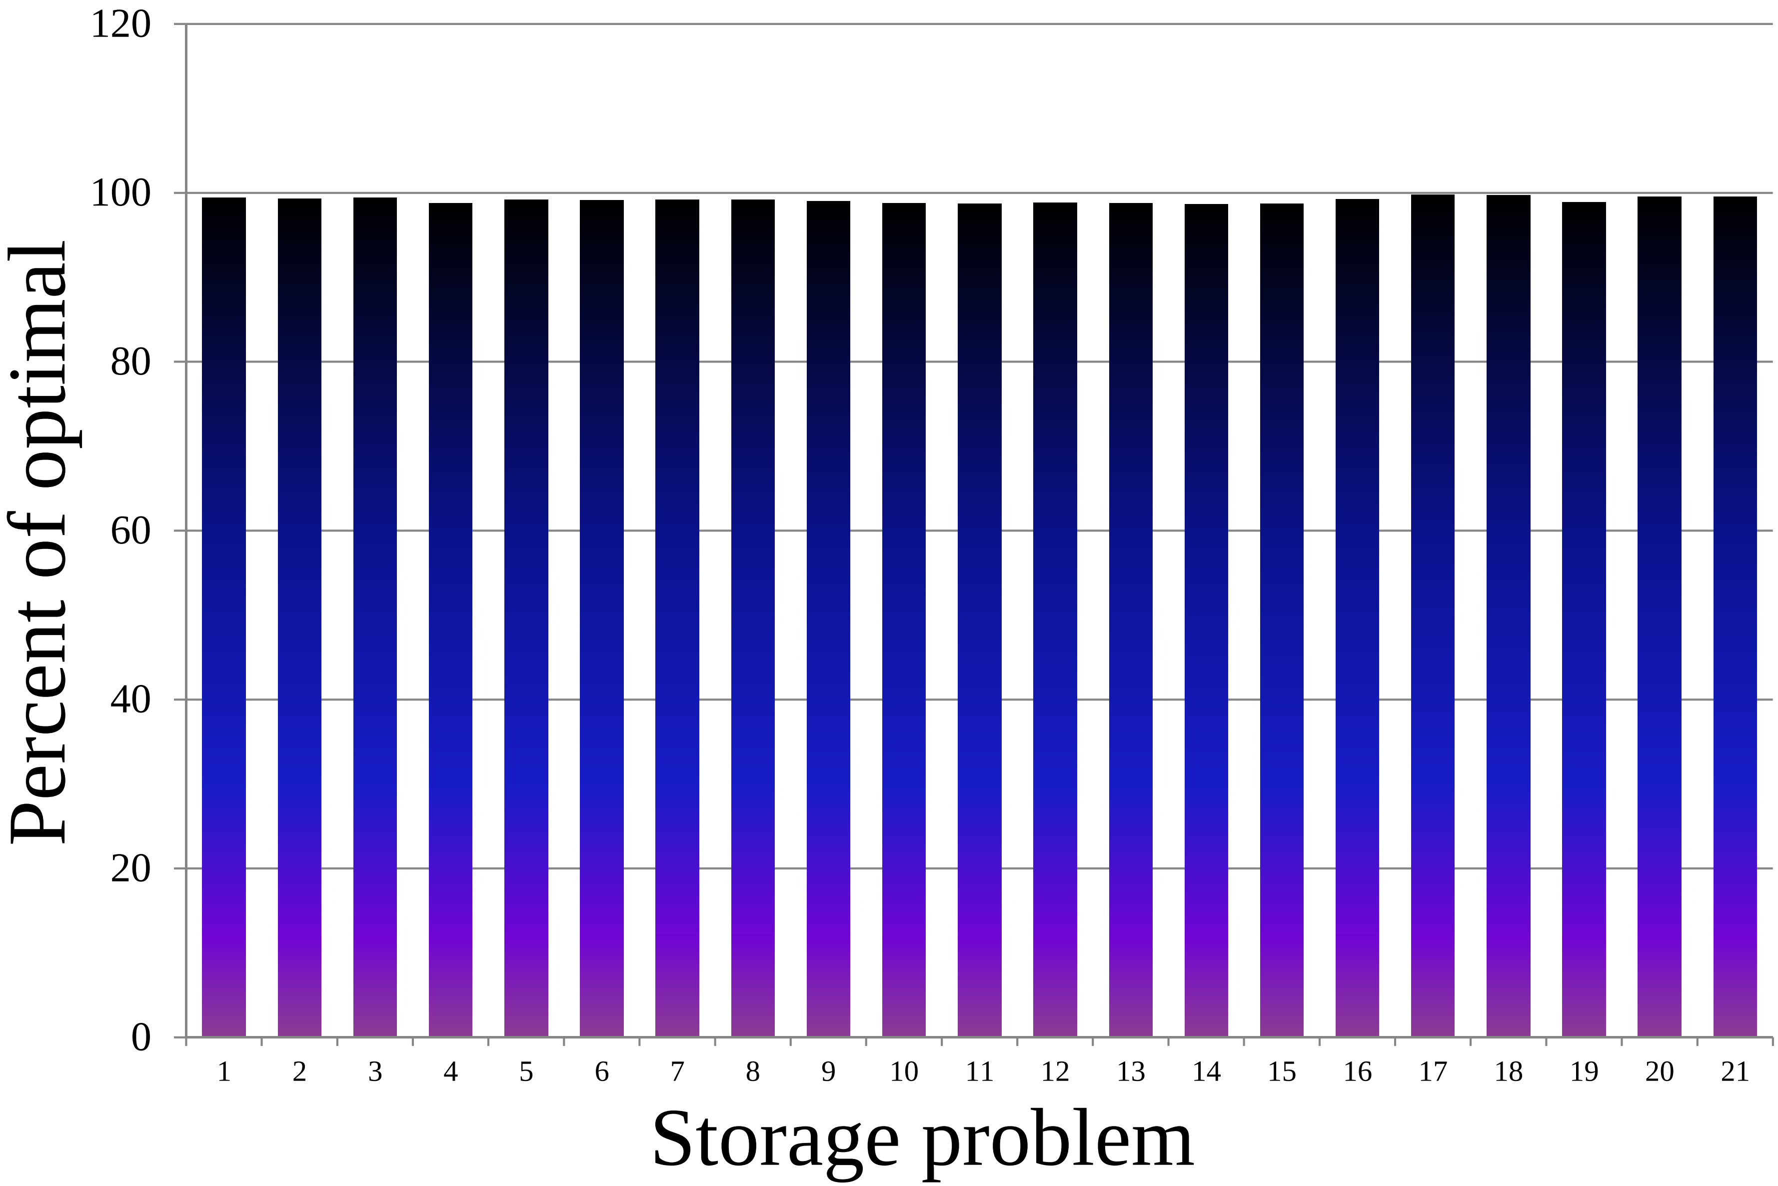

LMPs – Locational marginal prices

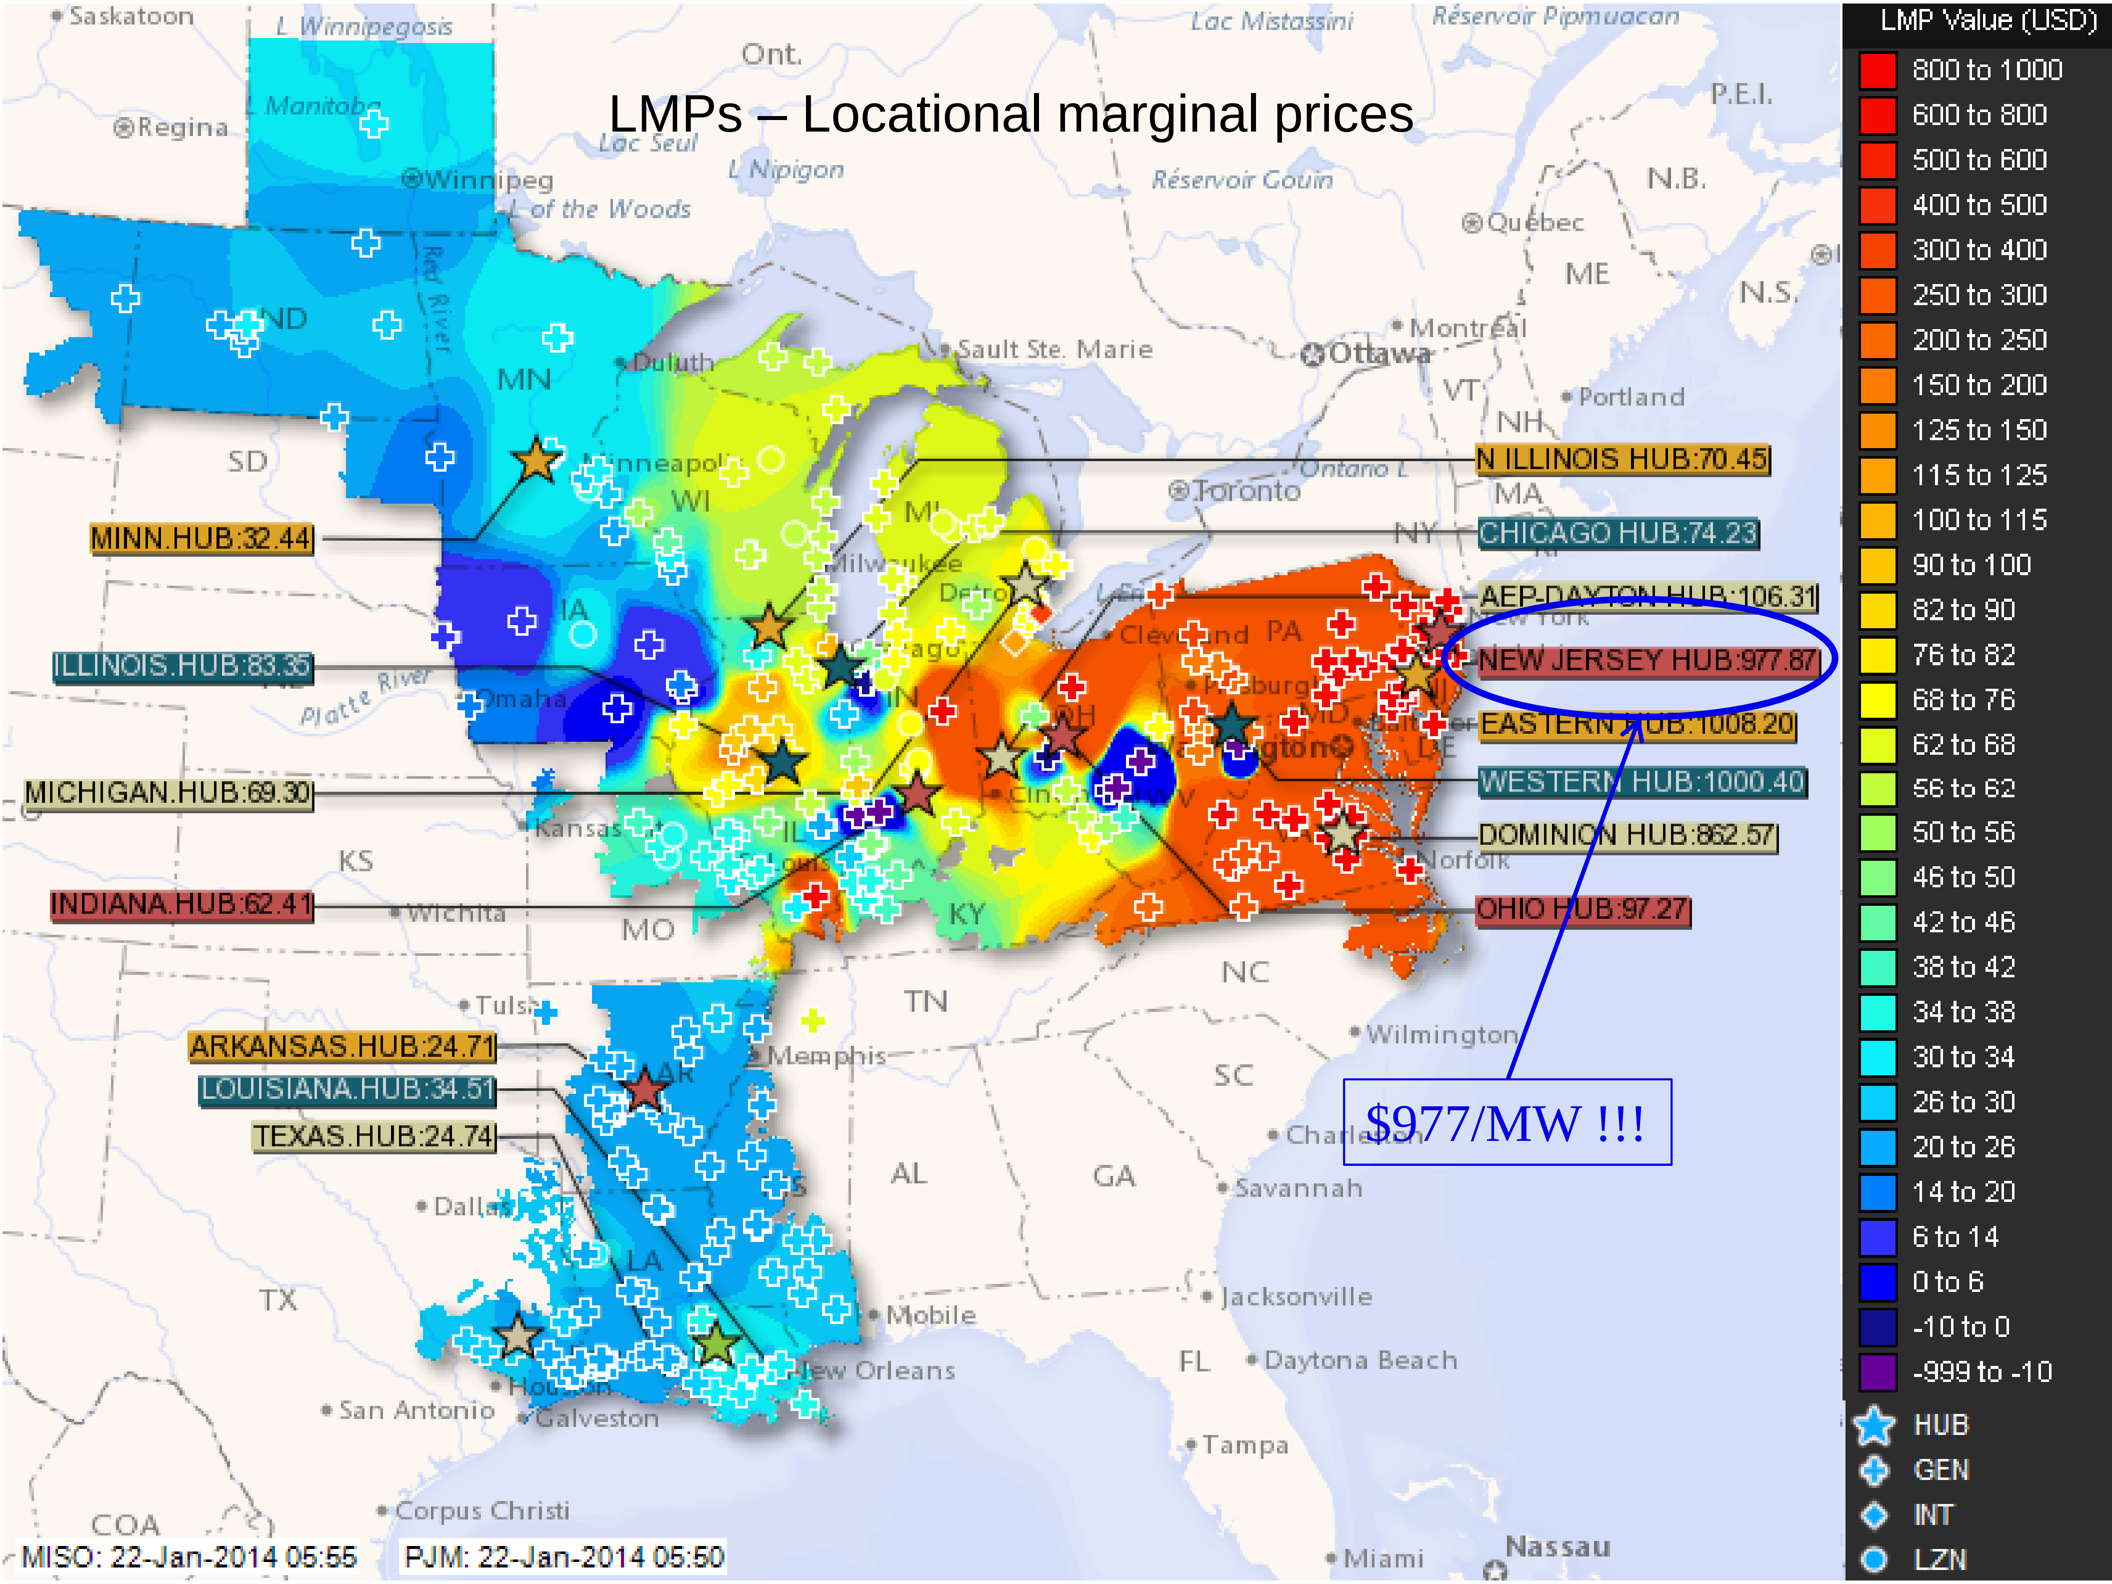

- Imagine 25 large storage devices spread around the PJM grid:

- Imagine 25 large storage devices spread around the PJM grid:

Grid level storage control

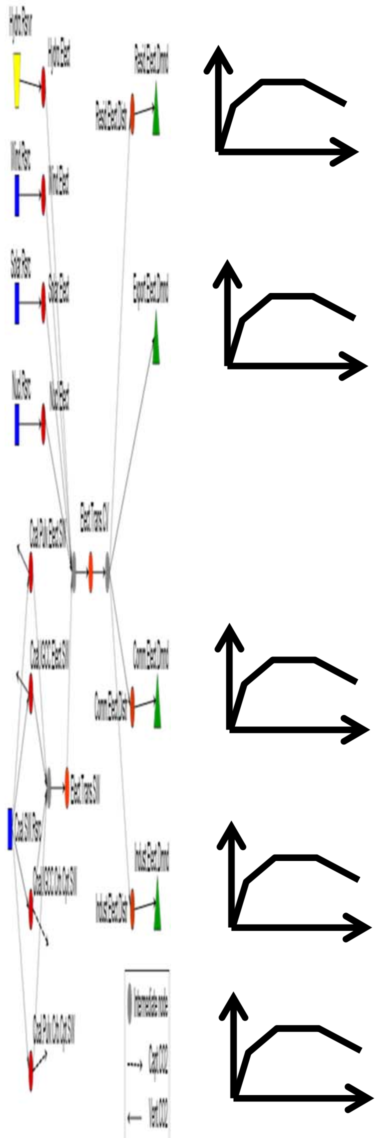

Grid level storage control

Monday

Time **:05** **:10** **:15** **:20**

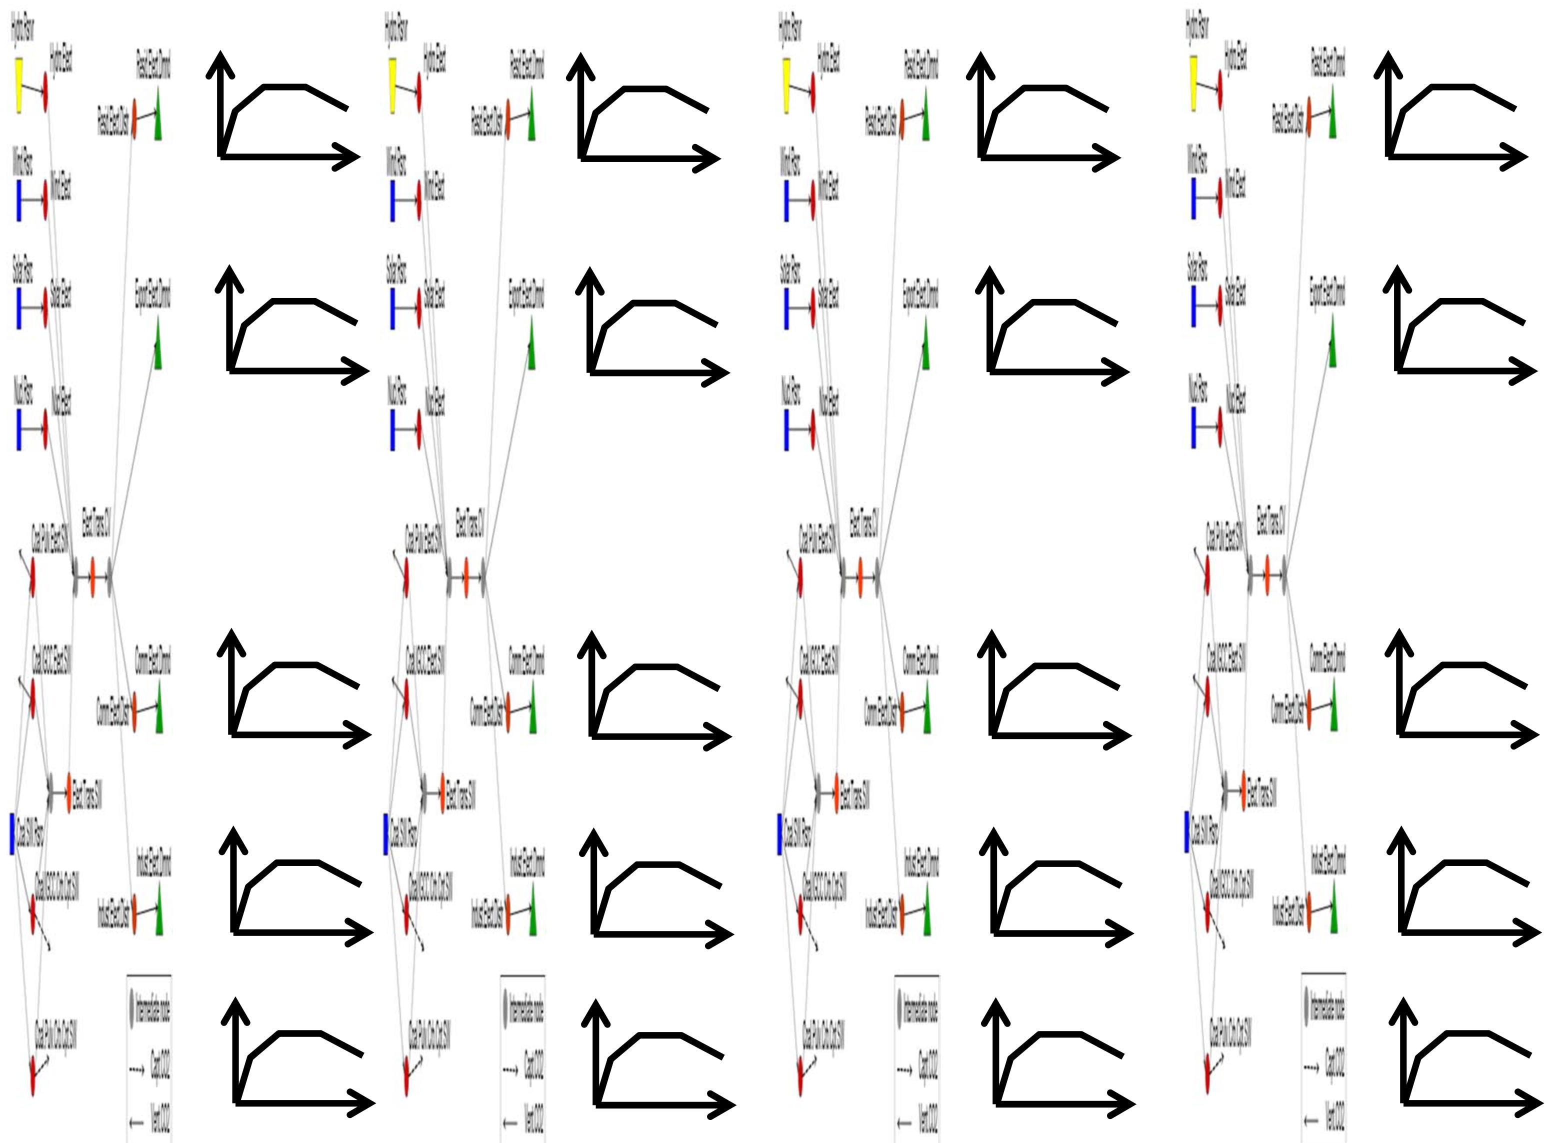

• • •

Grid level storage control

Monday

Time **:05** **:10** **:15** **:20**

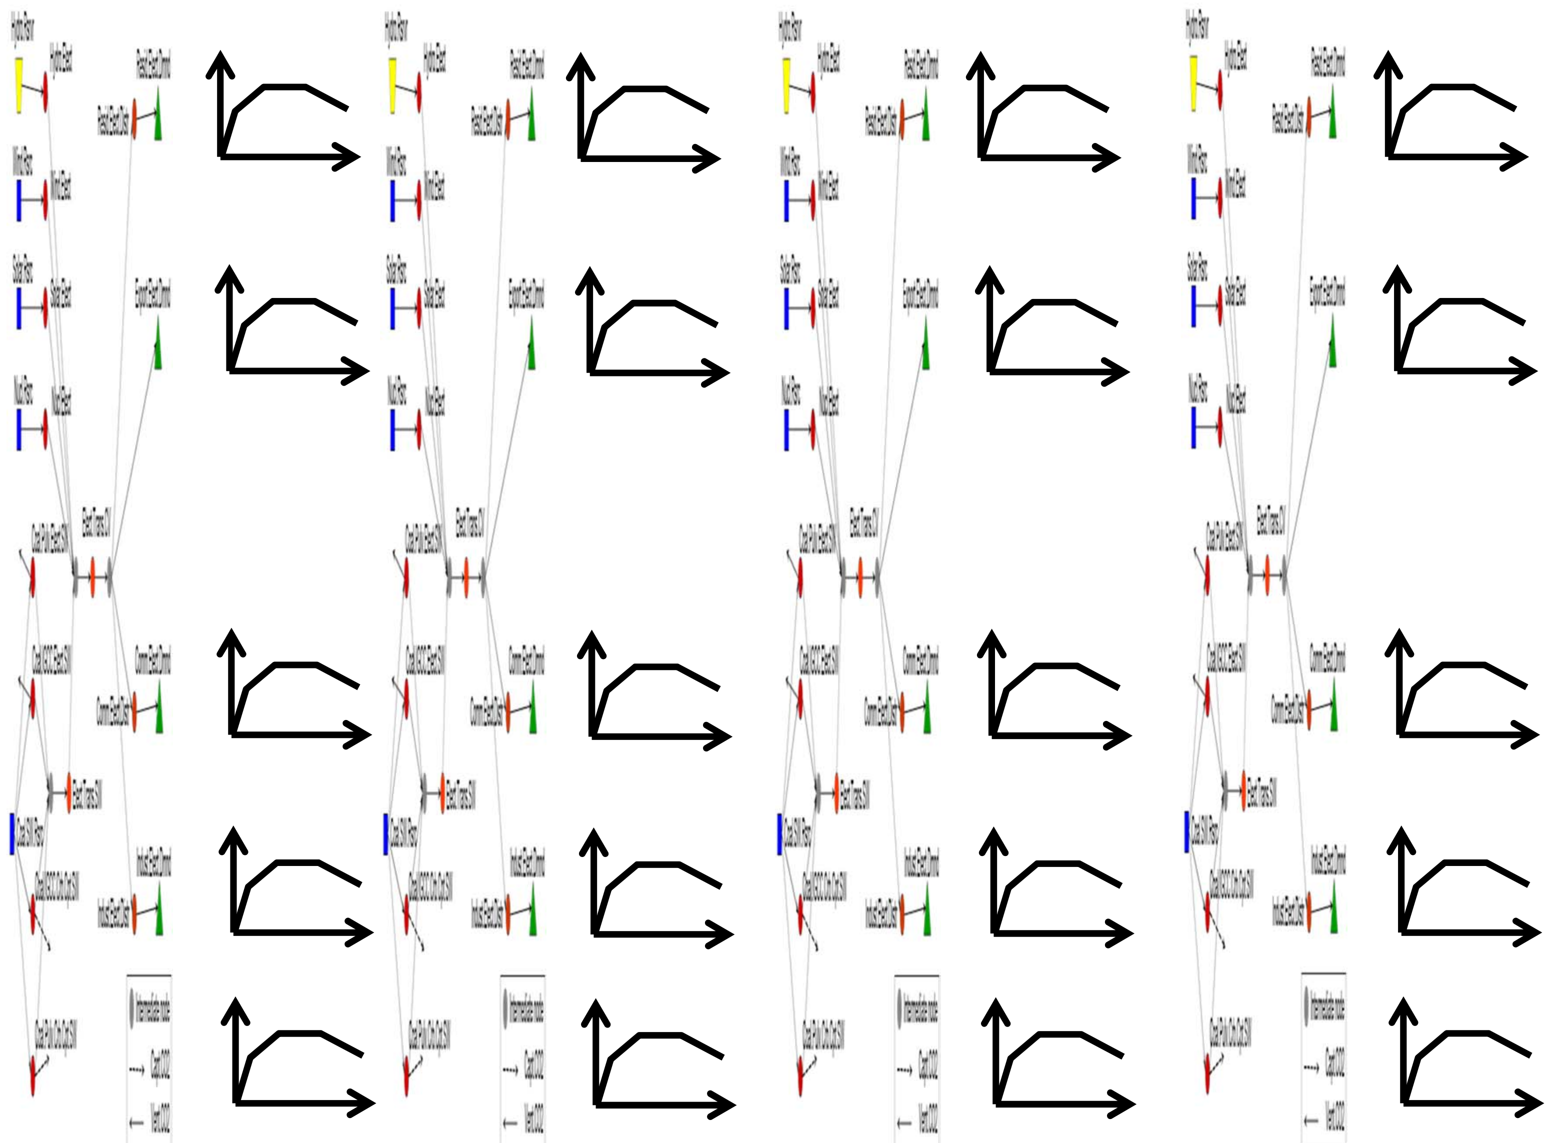

• • •

Grid level storage control

Monday

Time **:05** **:10** **:15** **:20**

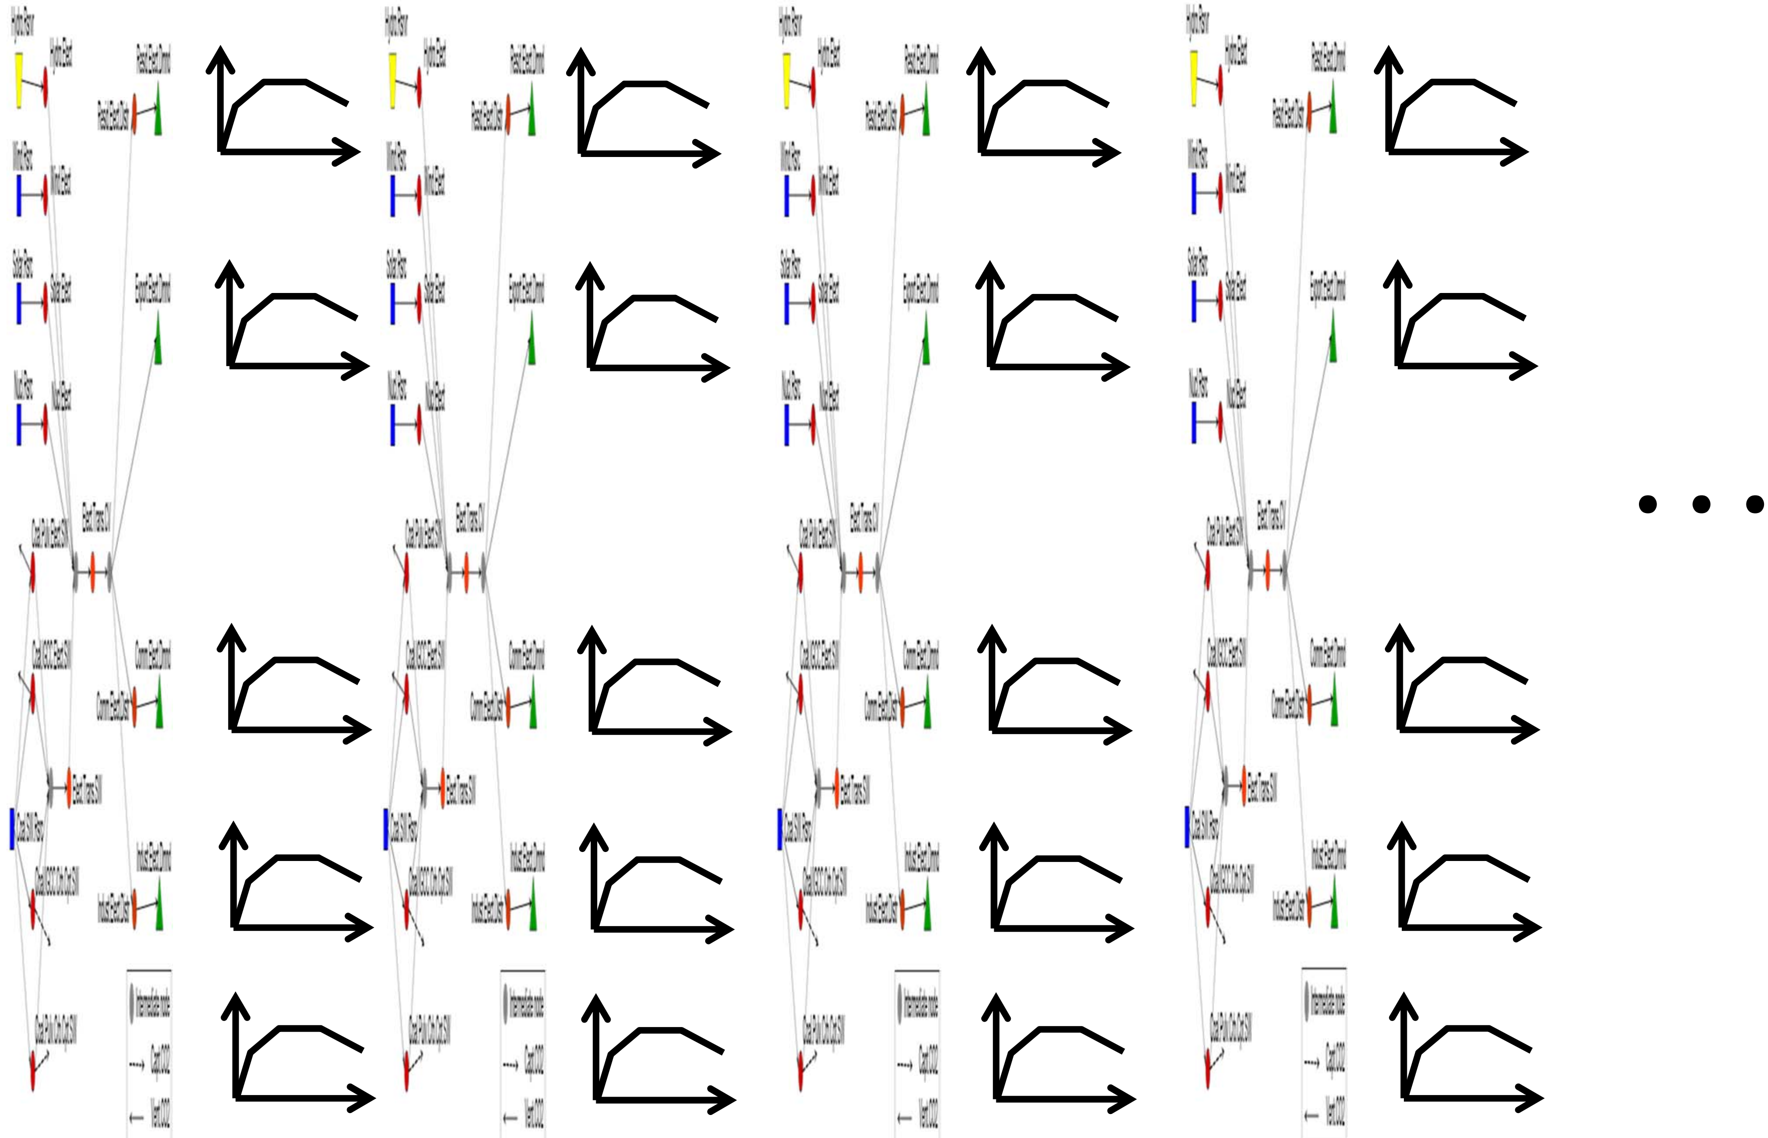

Grid level storage control

- ADP (blue) vs. LP optimal (black)

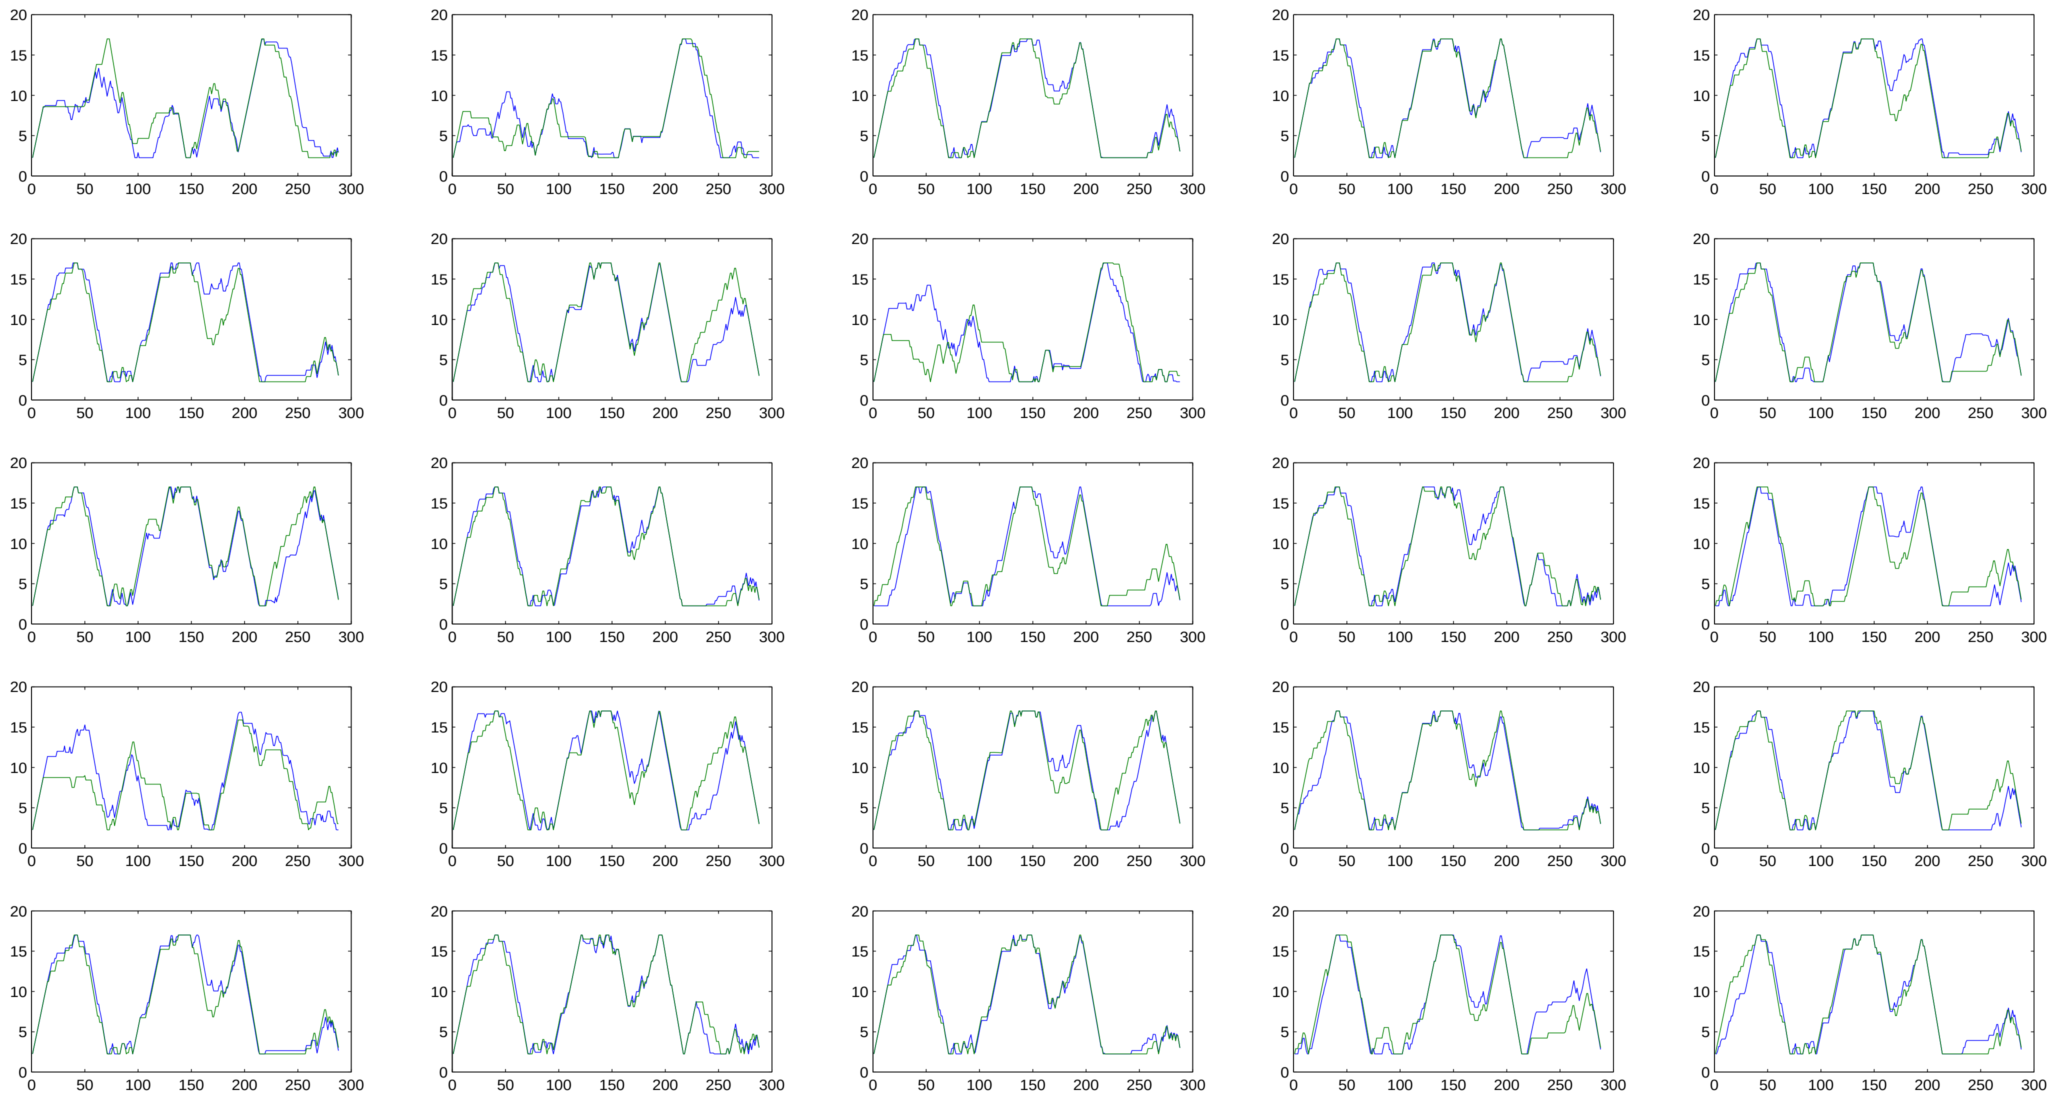

Grid level storage control

- Without storage

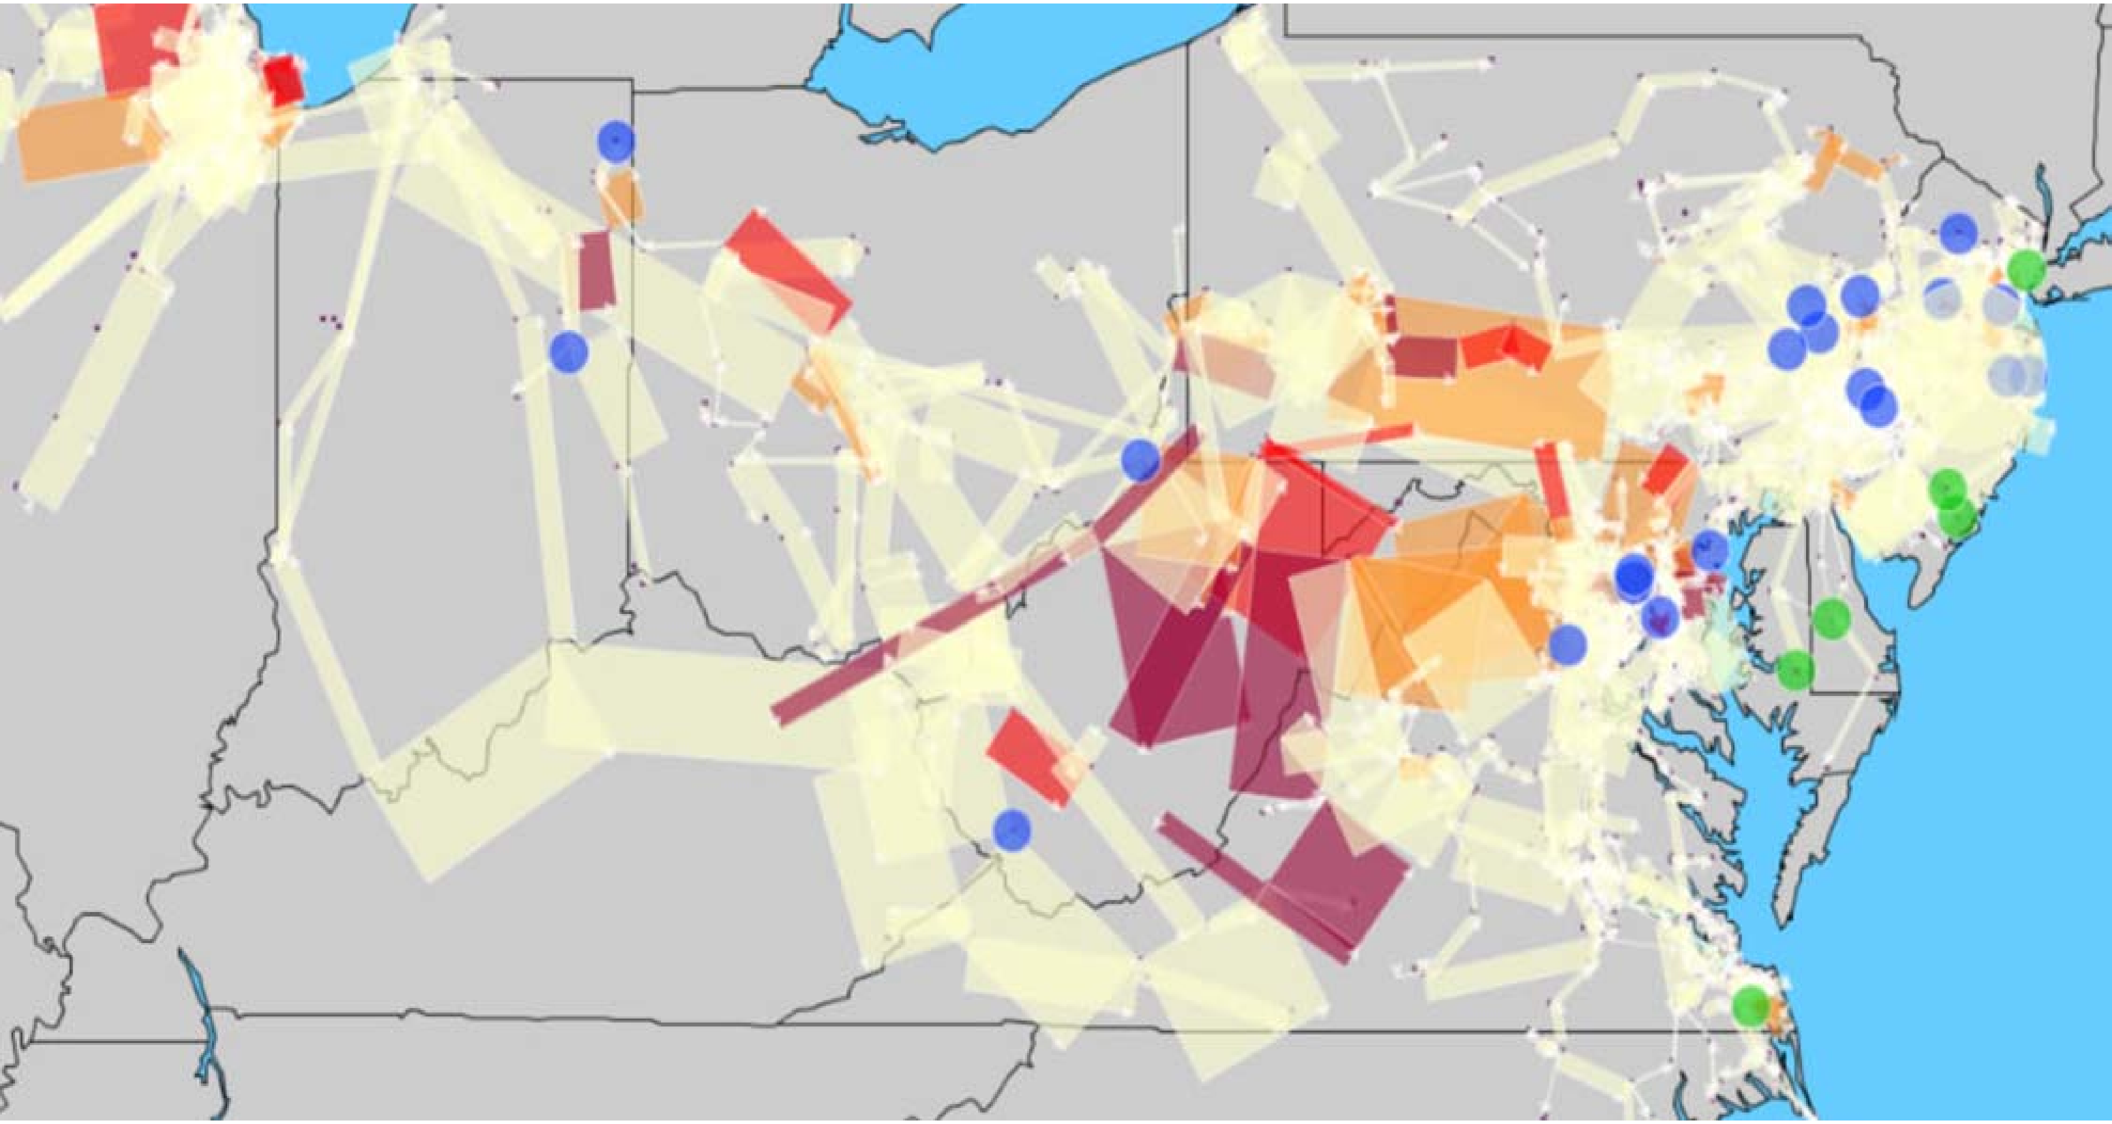

Grid level storage control

- With storage

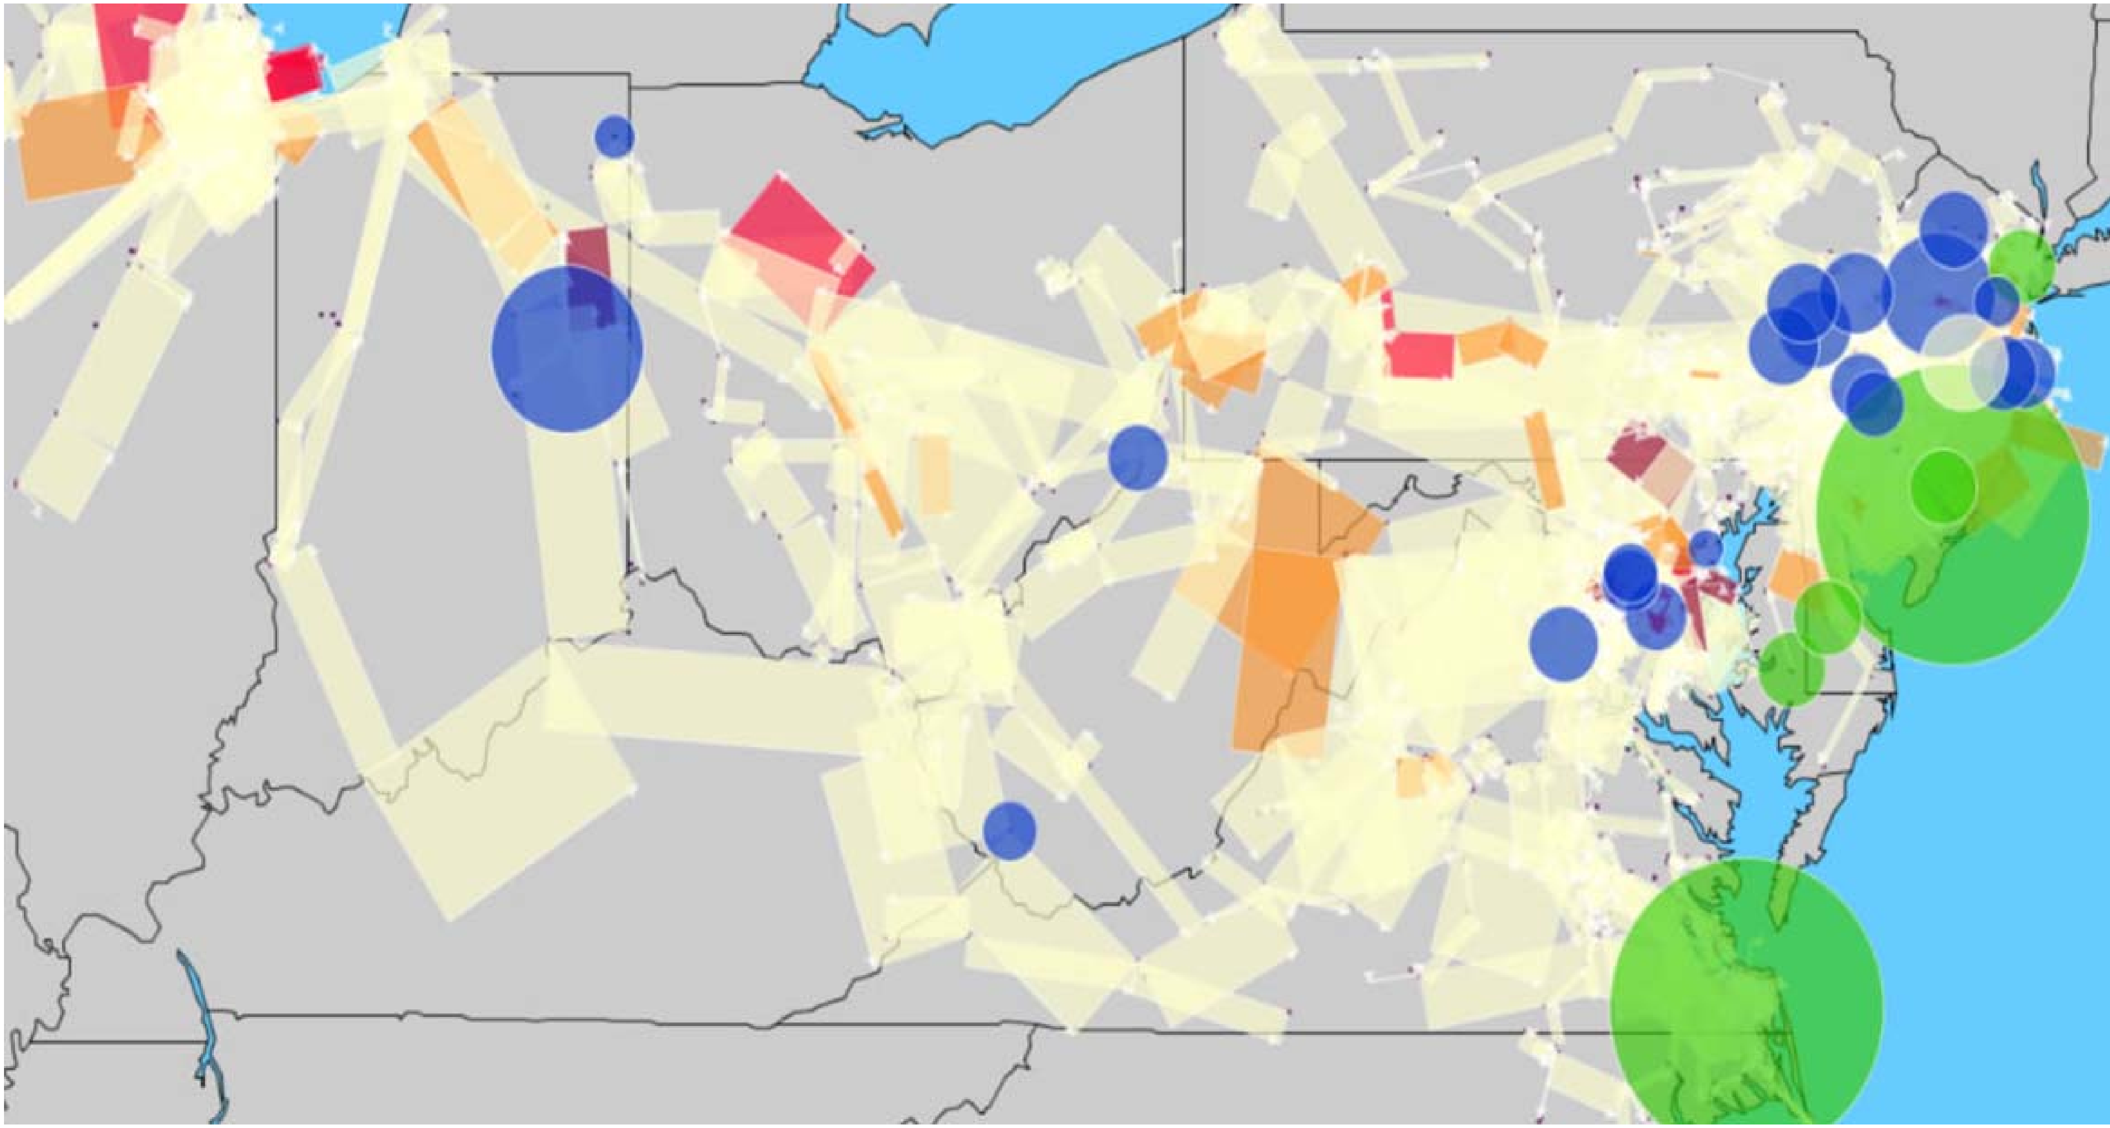

Heterogeneous fleets of batteries

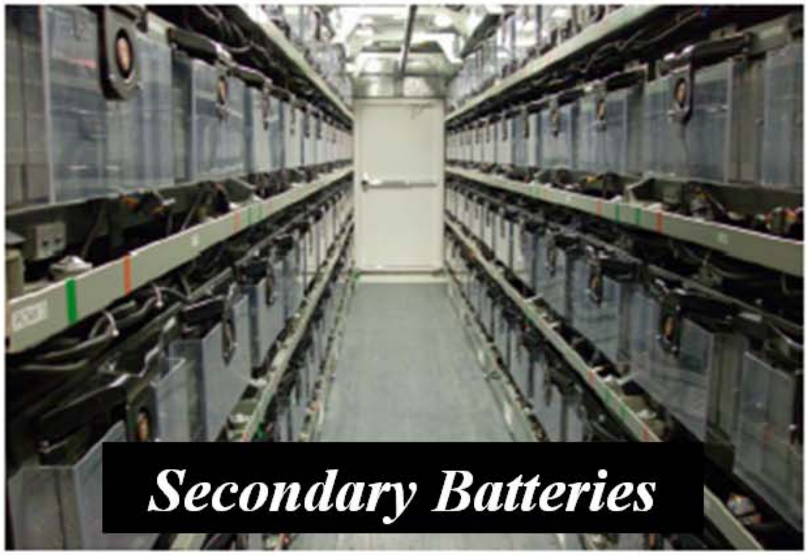

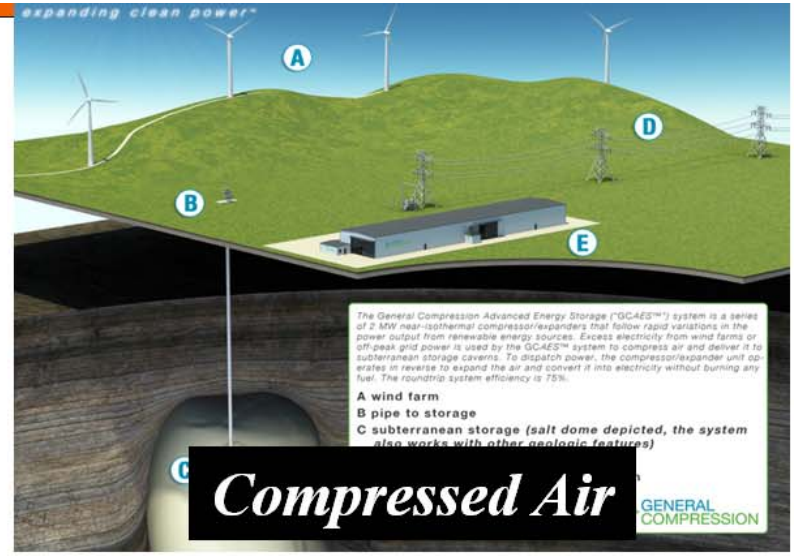

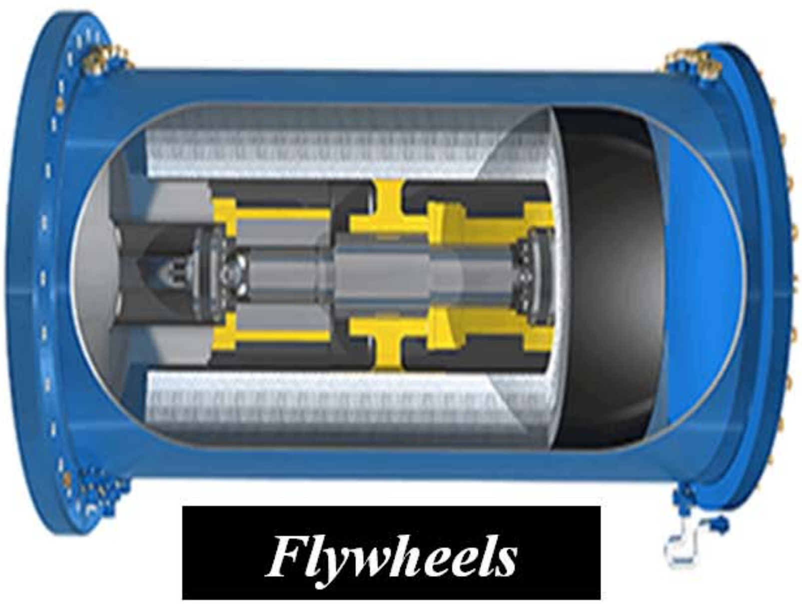

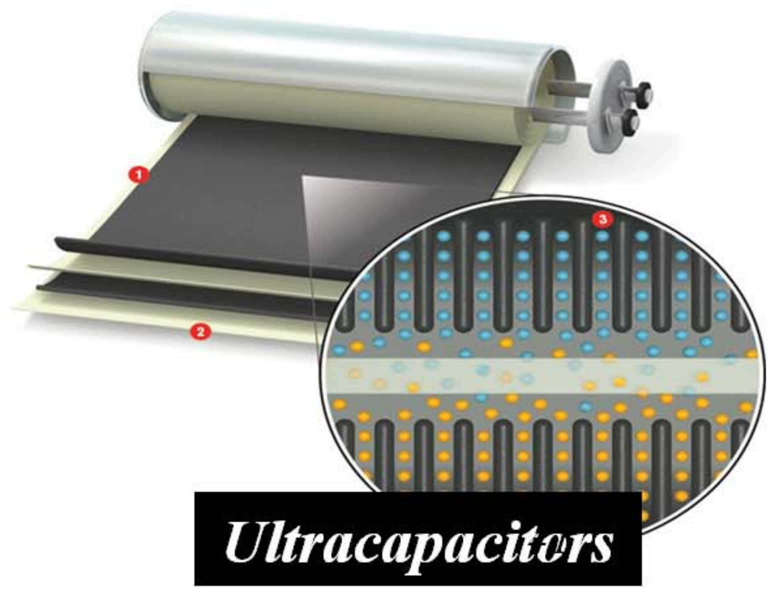

Heterogeneous fleets of batteries

- A tale of two batteries
  - » Ultracapacitor – High power, high efficiency, low capacity
  - » Lead acid – Lower power, lower efficiency, high capacity

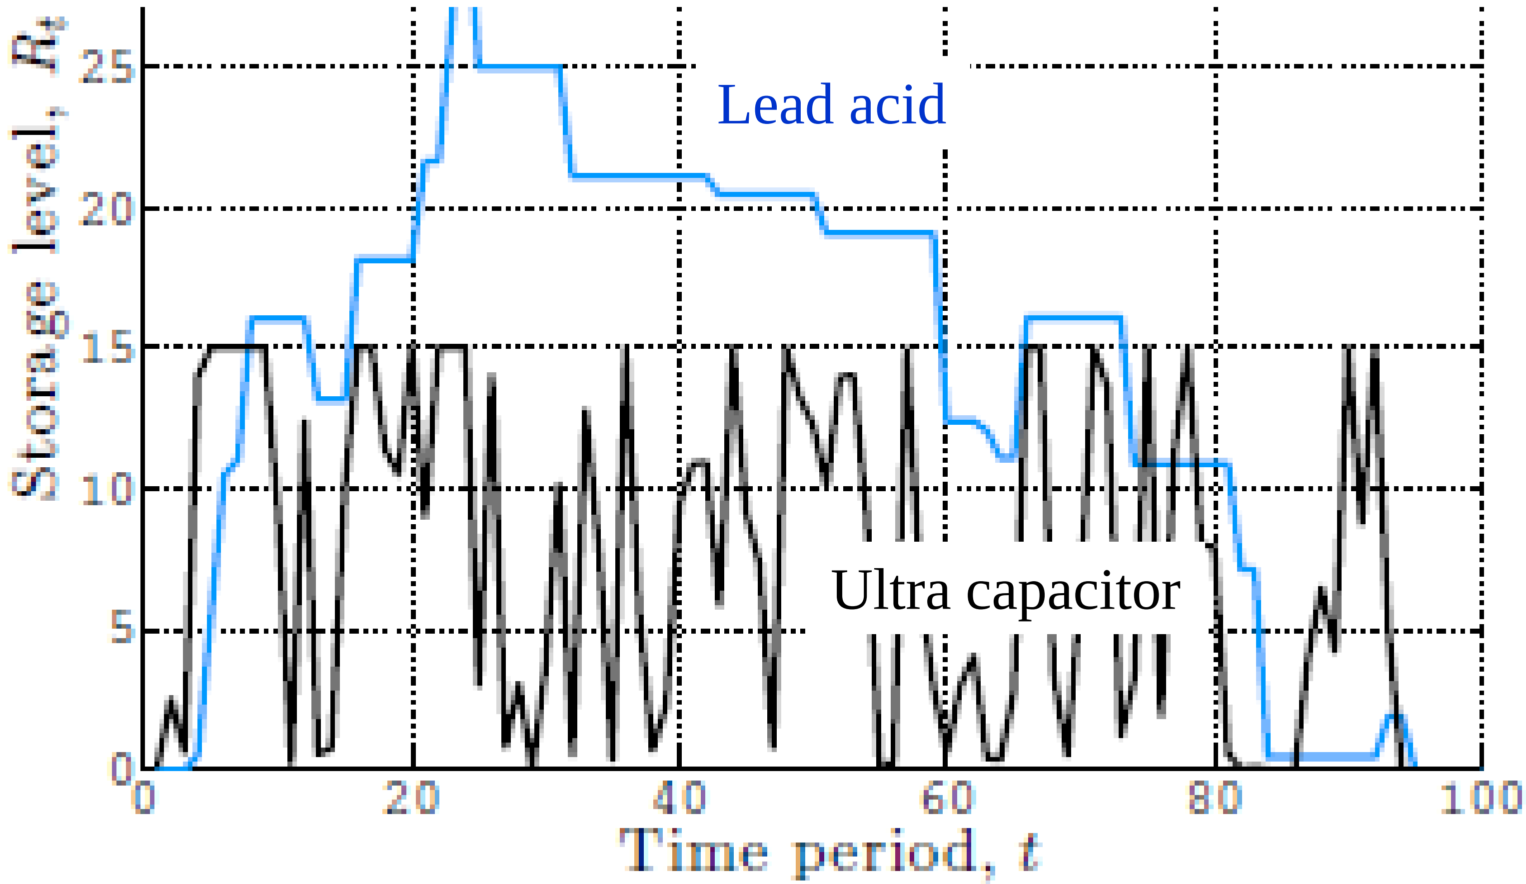

Co-optimizing multiple revenue streams

- Each time scale is addressing different issues, with different state variables
  » Daily model may need to recognize that solar energy during low load periods (in winter) are needed during higher usage periods in the evening.
  » Hourly model needs to monitor PJM compliance for frequency regulation, while optimizing buying and selling energy.
  » Every 2-seconds, need to balance frequency regulation compliance against short-term revenue opportunities and longer term energy shifting or battery arbitrage.
- How do we make decisions at the 2-second and 5-minute scales "see" hours into the future?

Co-optimizing multiple revenue streams

- Solution strategy I – Solve as a single optimization model
  - » 4-second increments over a day = 21,600 time periods
  - » State variables required:
    - Energy in the battery
    - Regulation price
    - Compliance metric
    - Electricity price
    - Load
  - » 21,600 time periods and 4-dimensional state variable!
- Solution strategy II – Three models, three scales
  - » Daily, hourly increments
    - Tracks hour-of-day patterns,
  - » Hourly model, 5-minute increments
    - Tracks PJM frequency regulation compliance
  - » 5-minute model, 4-second increments
    - Balances frequency regulation signal against buying/selling electricity now and holding for the future.

A multiscale model

- The challenge of time scales:

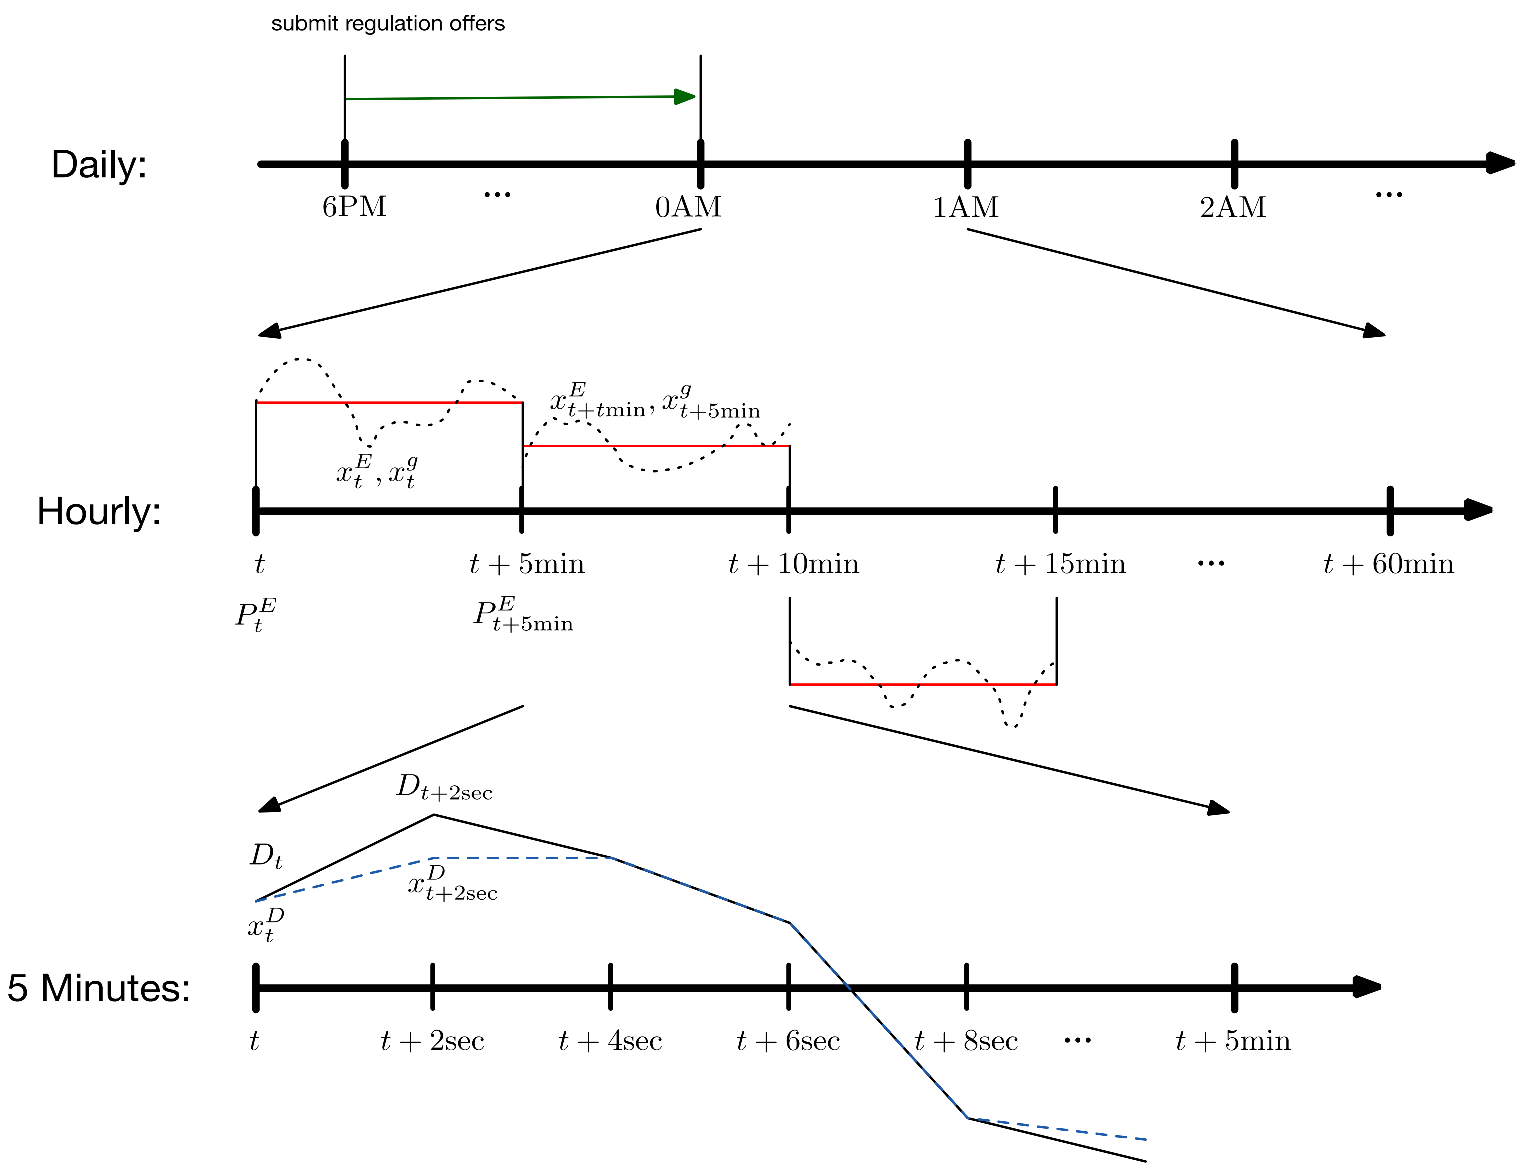

$S_t = R_t$

24 time steps

$S_t = (R_t, G_t, P^E_t)$

x^E_t = Economic basepoint

12 time steps

$S_t = (R_t, G_t, D_t)$

x^D_t = Frequency regulation

75 time steps

Handling multiple time scales

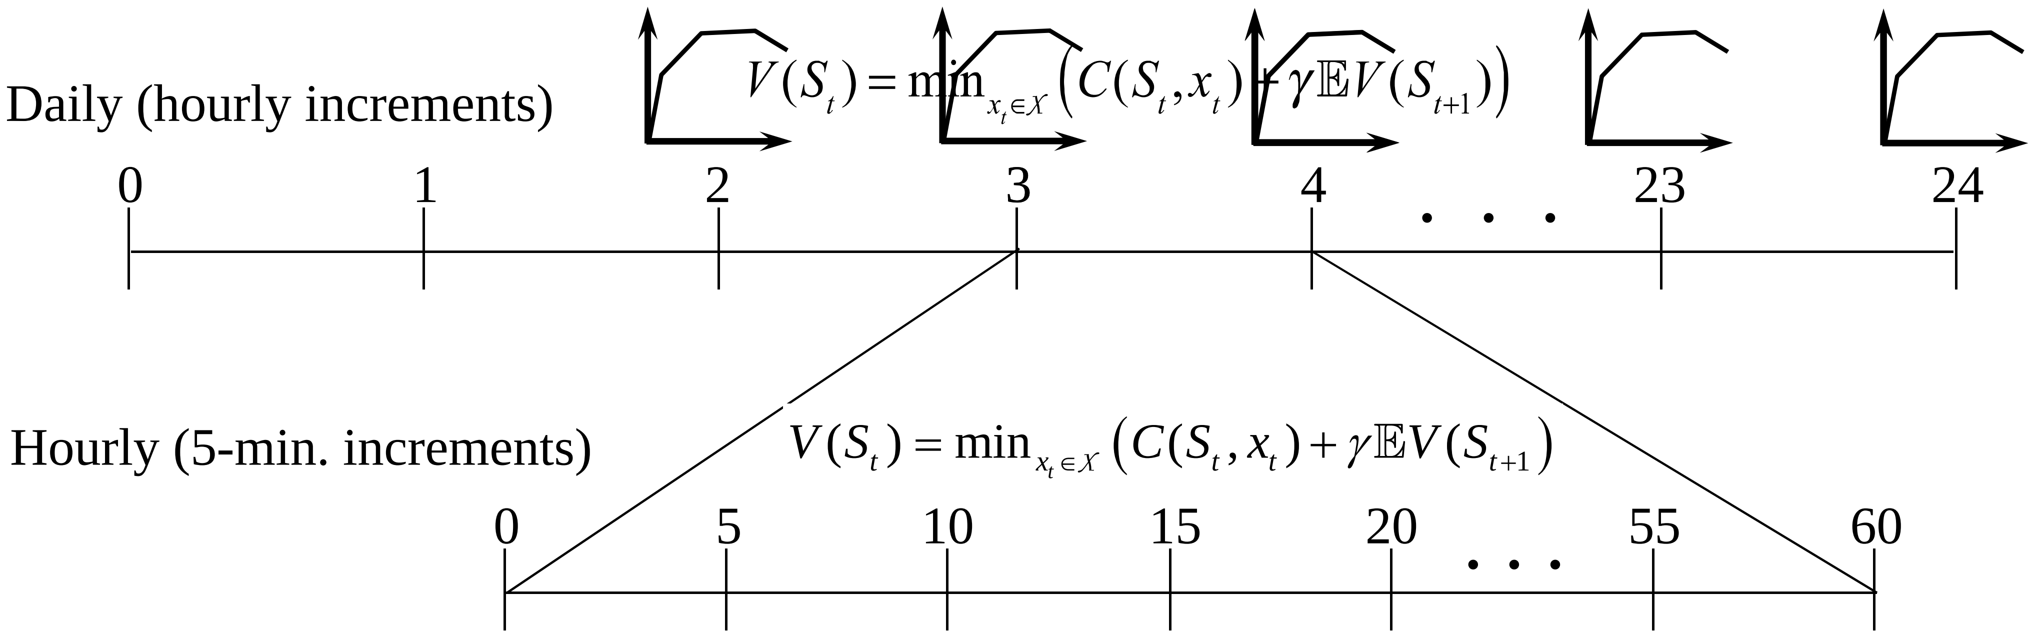

Handling multiple time scales

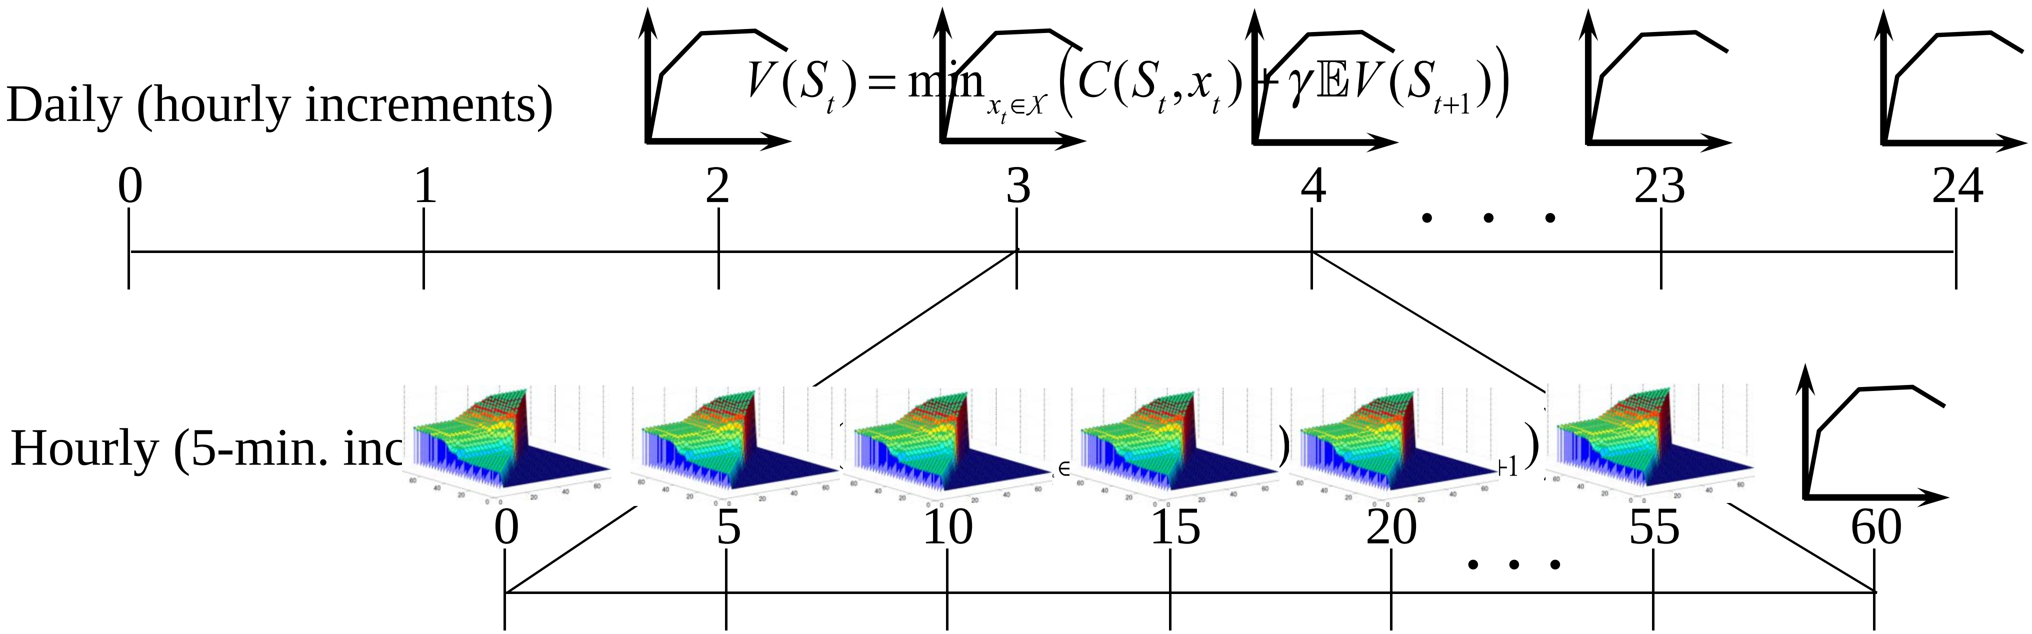

Handling multiple time scales

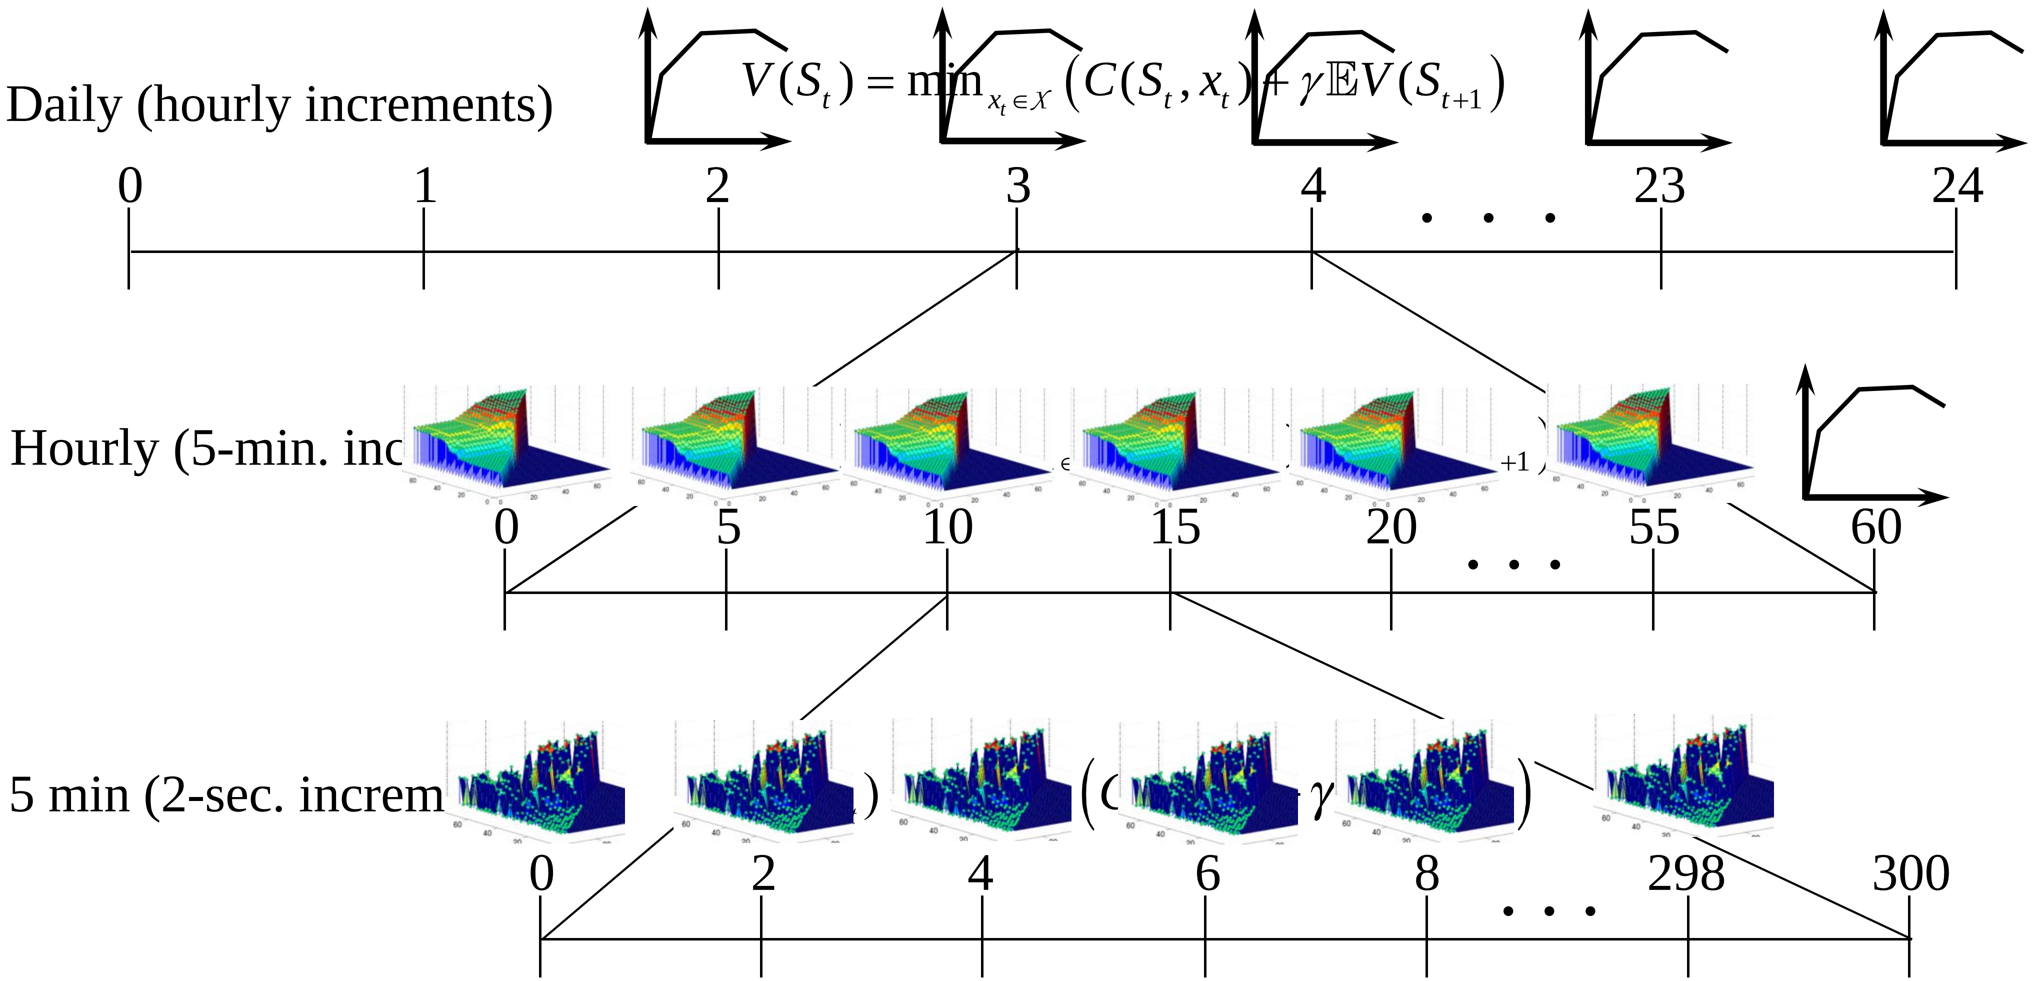

Frequency Regulation

- PJM sends charge/discharge signals to the generators every 2 seconds to smooth out fluctuations in supply/demand balance.

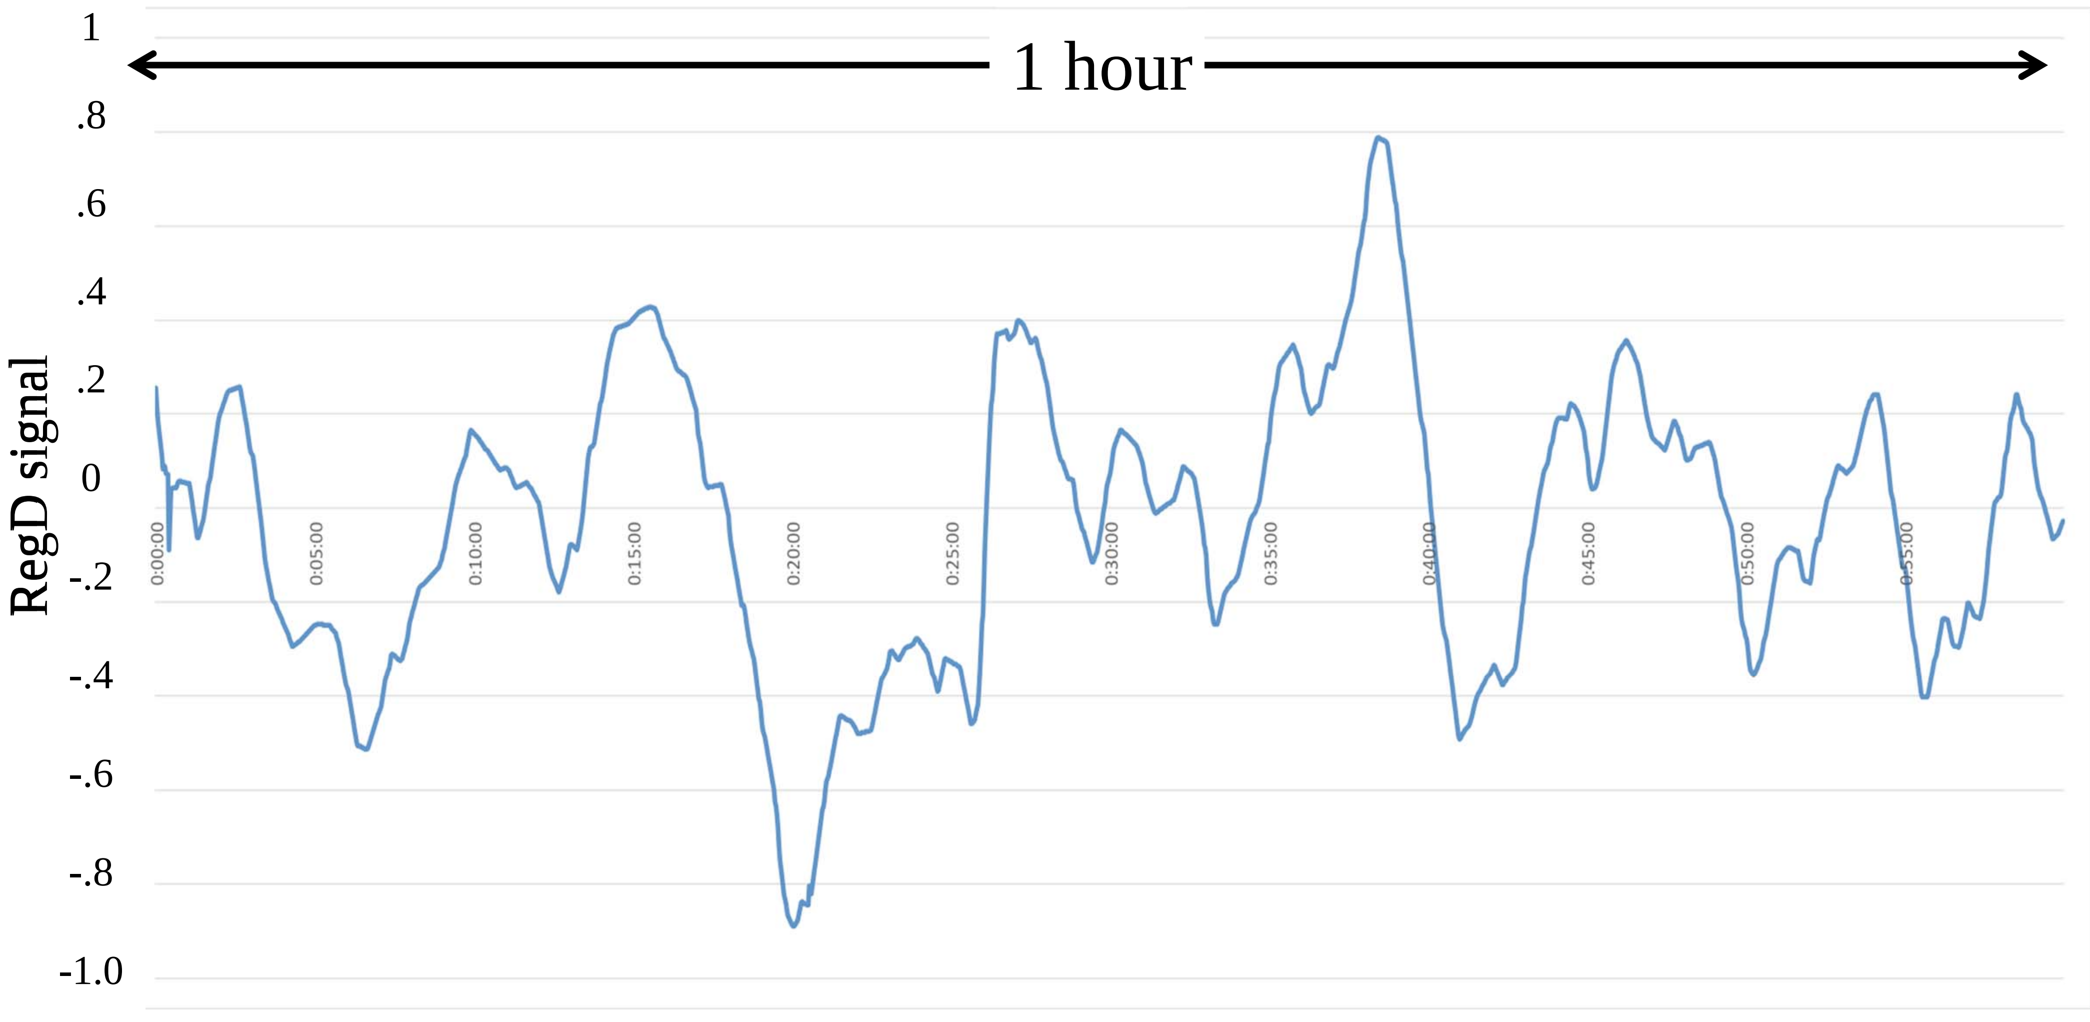

Optimizing a multi-dimensional problem

- The frequency regulation DP

Step 0: Initialize $V_{76}^{FR}(S_{76}^{FR}) = V_{76}^{EB}(S_{76}^{EB})$ for all states.

Step 1: Step backward $t = 75$ (5 mins), 74, 73, ..., 1

Step 2: Loop over $S_t = \left(R_t, G_t, D_t\right)$ (three loops, 10 million states)

Step 3: Loop over storage decisions x_t (21 actions - increments of $\pm .05$MW)

Step 4: Take the expectation over the RegD signal $d = \left(D_t\right)$ (21)

Compute $Q_t^{FR}(S_t, x_t) = C(S_t, x_t) + \sum_{d=-1}^{1} V_{t+1}\left(S^M\left(S_t, x_t, D_{t+1} = d\right)\right) P^{\mathrm{Re}gD}(d)$

End step 4;

End Step 3;

Find $V_t^{FR}(S_t) = \max_{x_t} Q_t^{FR}(S_t, x_t)$

Store $X_t^{\pi^*}(S_t) = \arg\max_{x_t} Q_t^{FR}(S_t, x_t)$. (This is our policy)

End Step 2;

End Step 1;

Computational challenges

- Brute force solution of full MDP is intractable
  - » At the lowest (FR) level:
    - Value functions have 10 million states
    - These have to be computed for each LMP (we clustered observed LMPs into 7 levels)
    - Have to be computed and stored for 21,000 time periods.
    - ~12.8 terabytes of data
  - » Singular value decomposition (SVD)
    - We can represent a matrix M as:

$$\underbrace{M}_{900\times(600\times 20)} = \underbrace{U}_{900\times 10}\,\underbrace{\Sigma}_{10\times 10}\,\underbrace{V^T}_{10\times(600\times 20)}$$

10 million elements

130,000 elements

Computational challenges

- Error in low-rank approximation

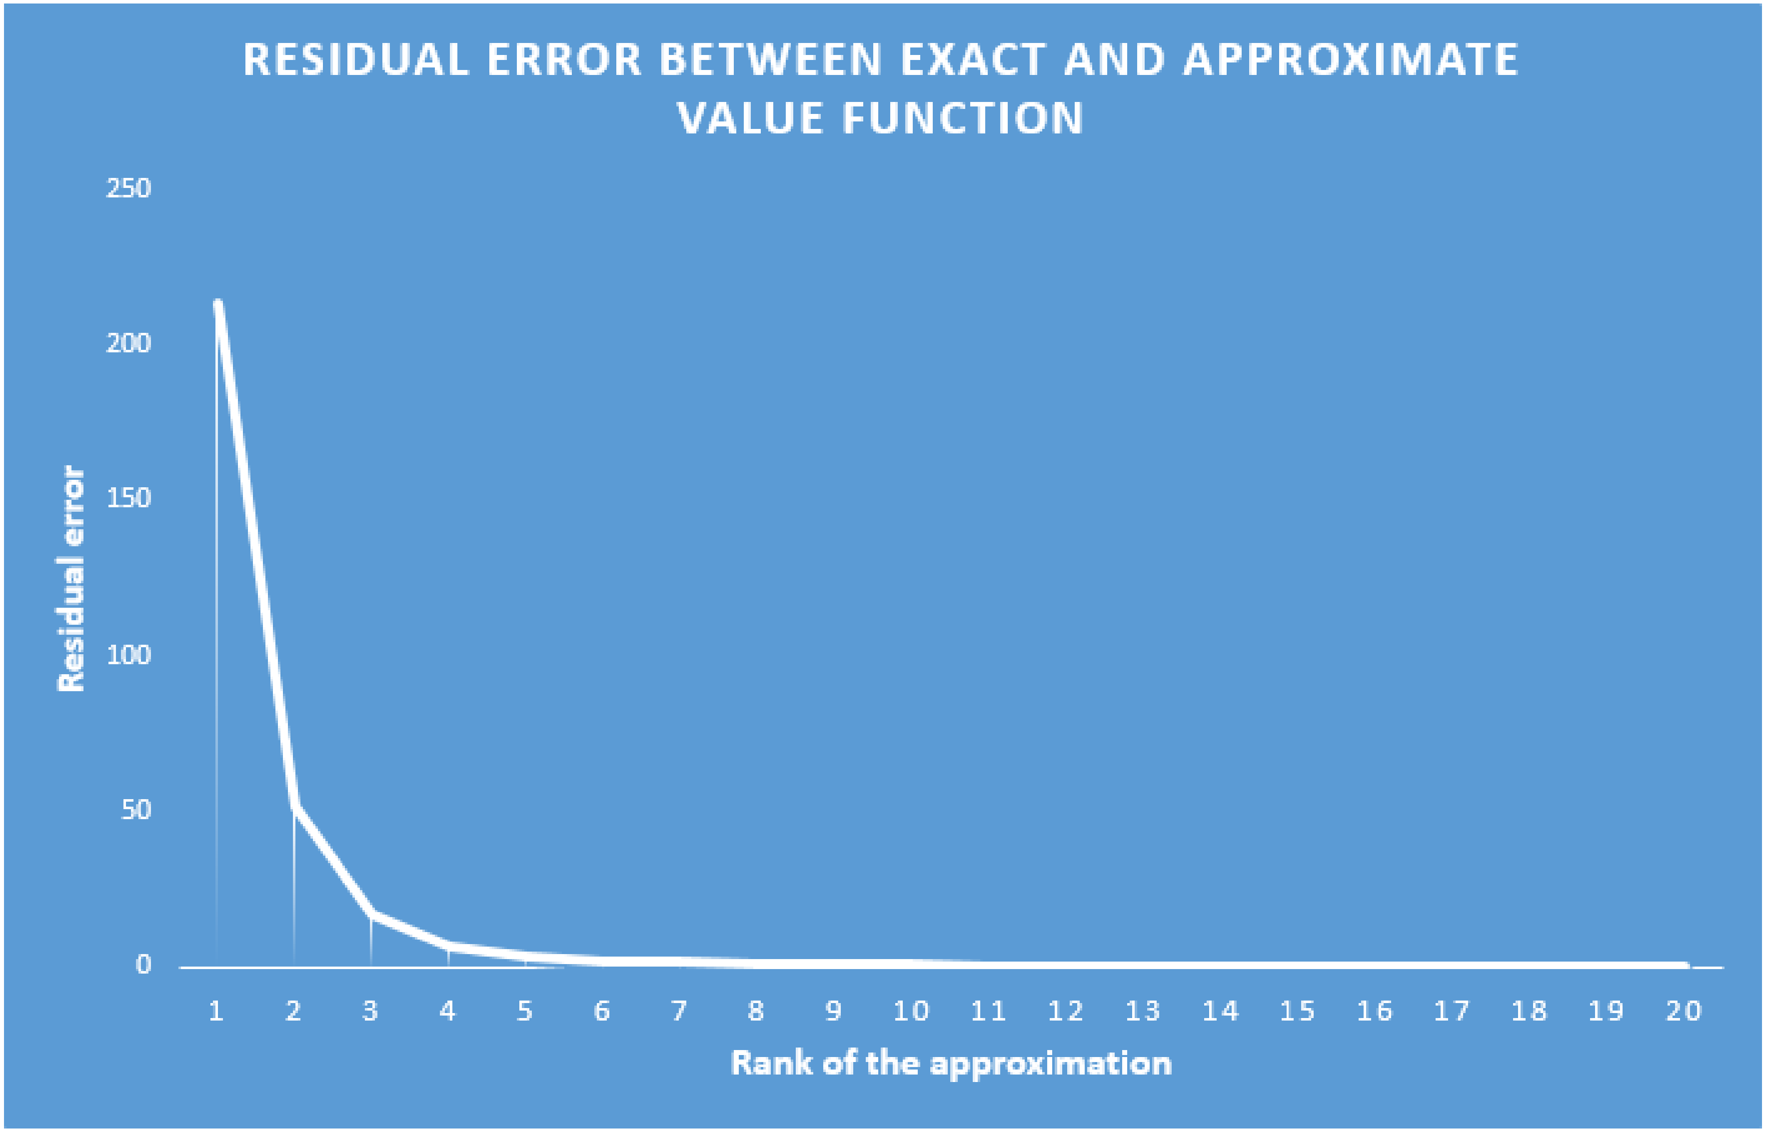

Simulations

- 24 hour simulation
  - » Training:
    - 60 sample paths for LMPs
    - 20 sample paths for RegD signal
  - » Testing
    - 10 sample paths for LMPs
    - Fixed RegD price over a range

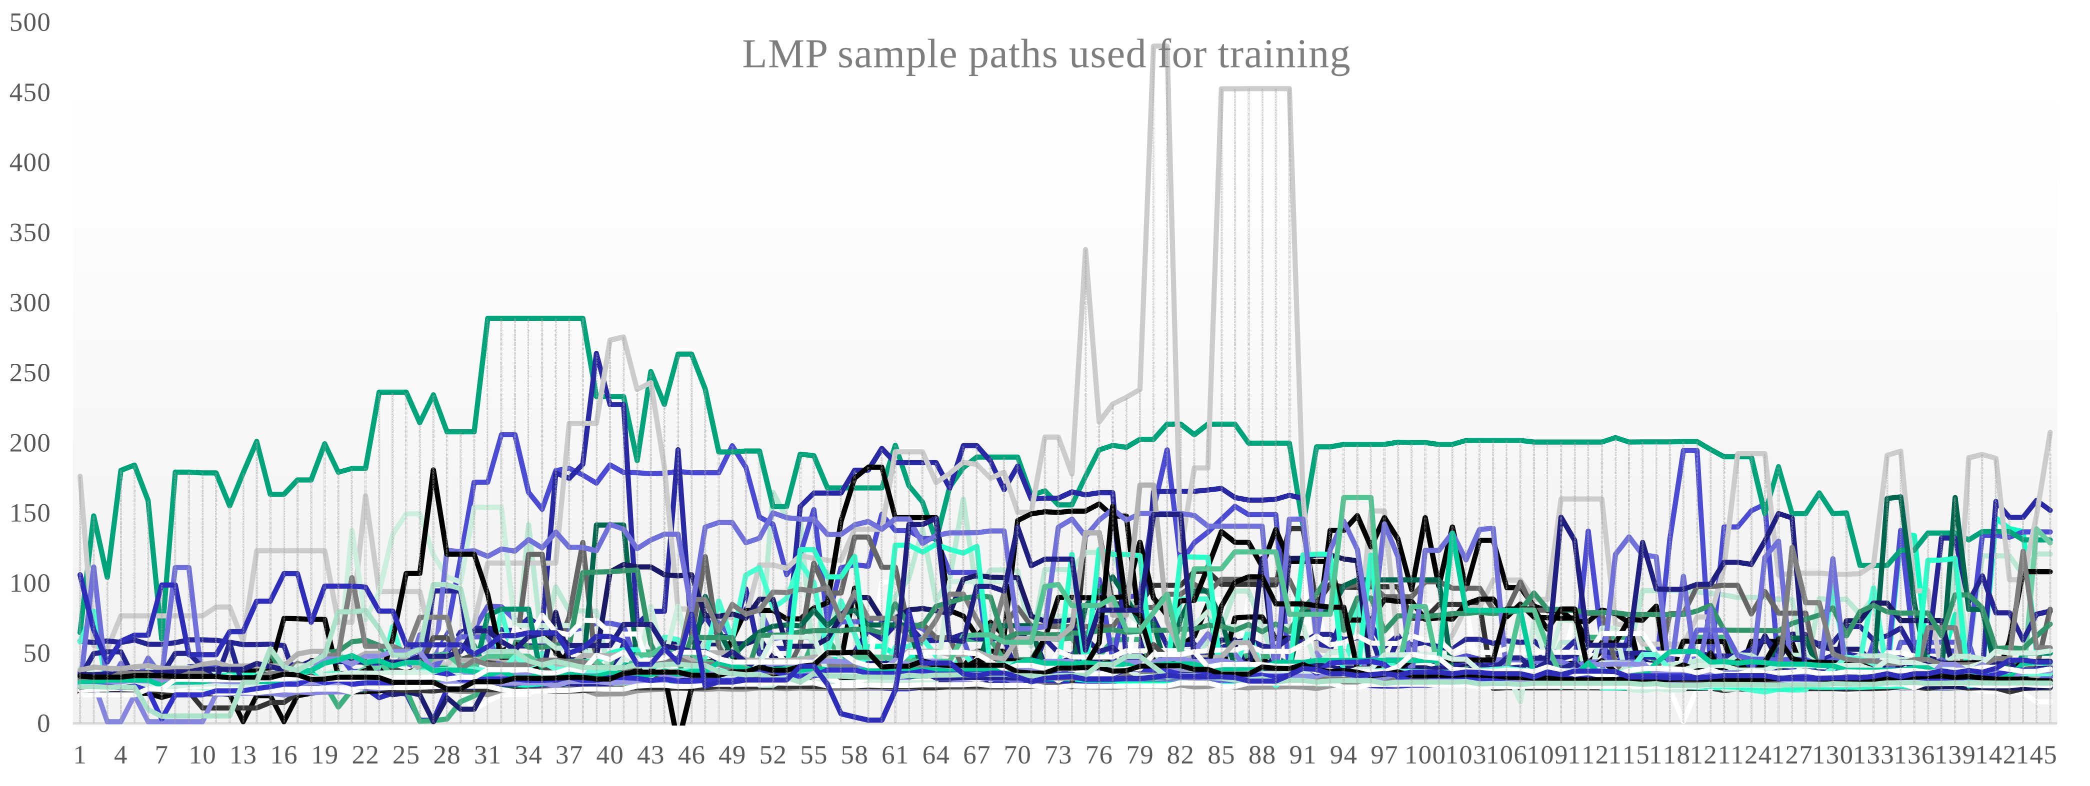

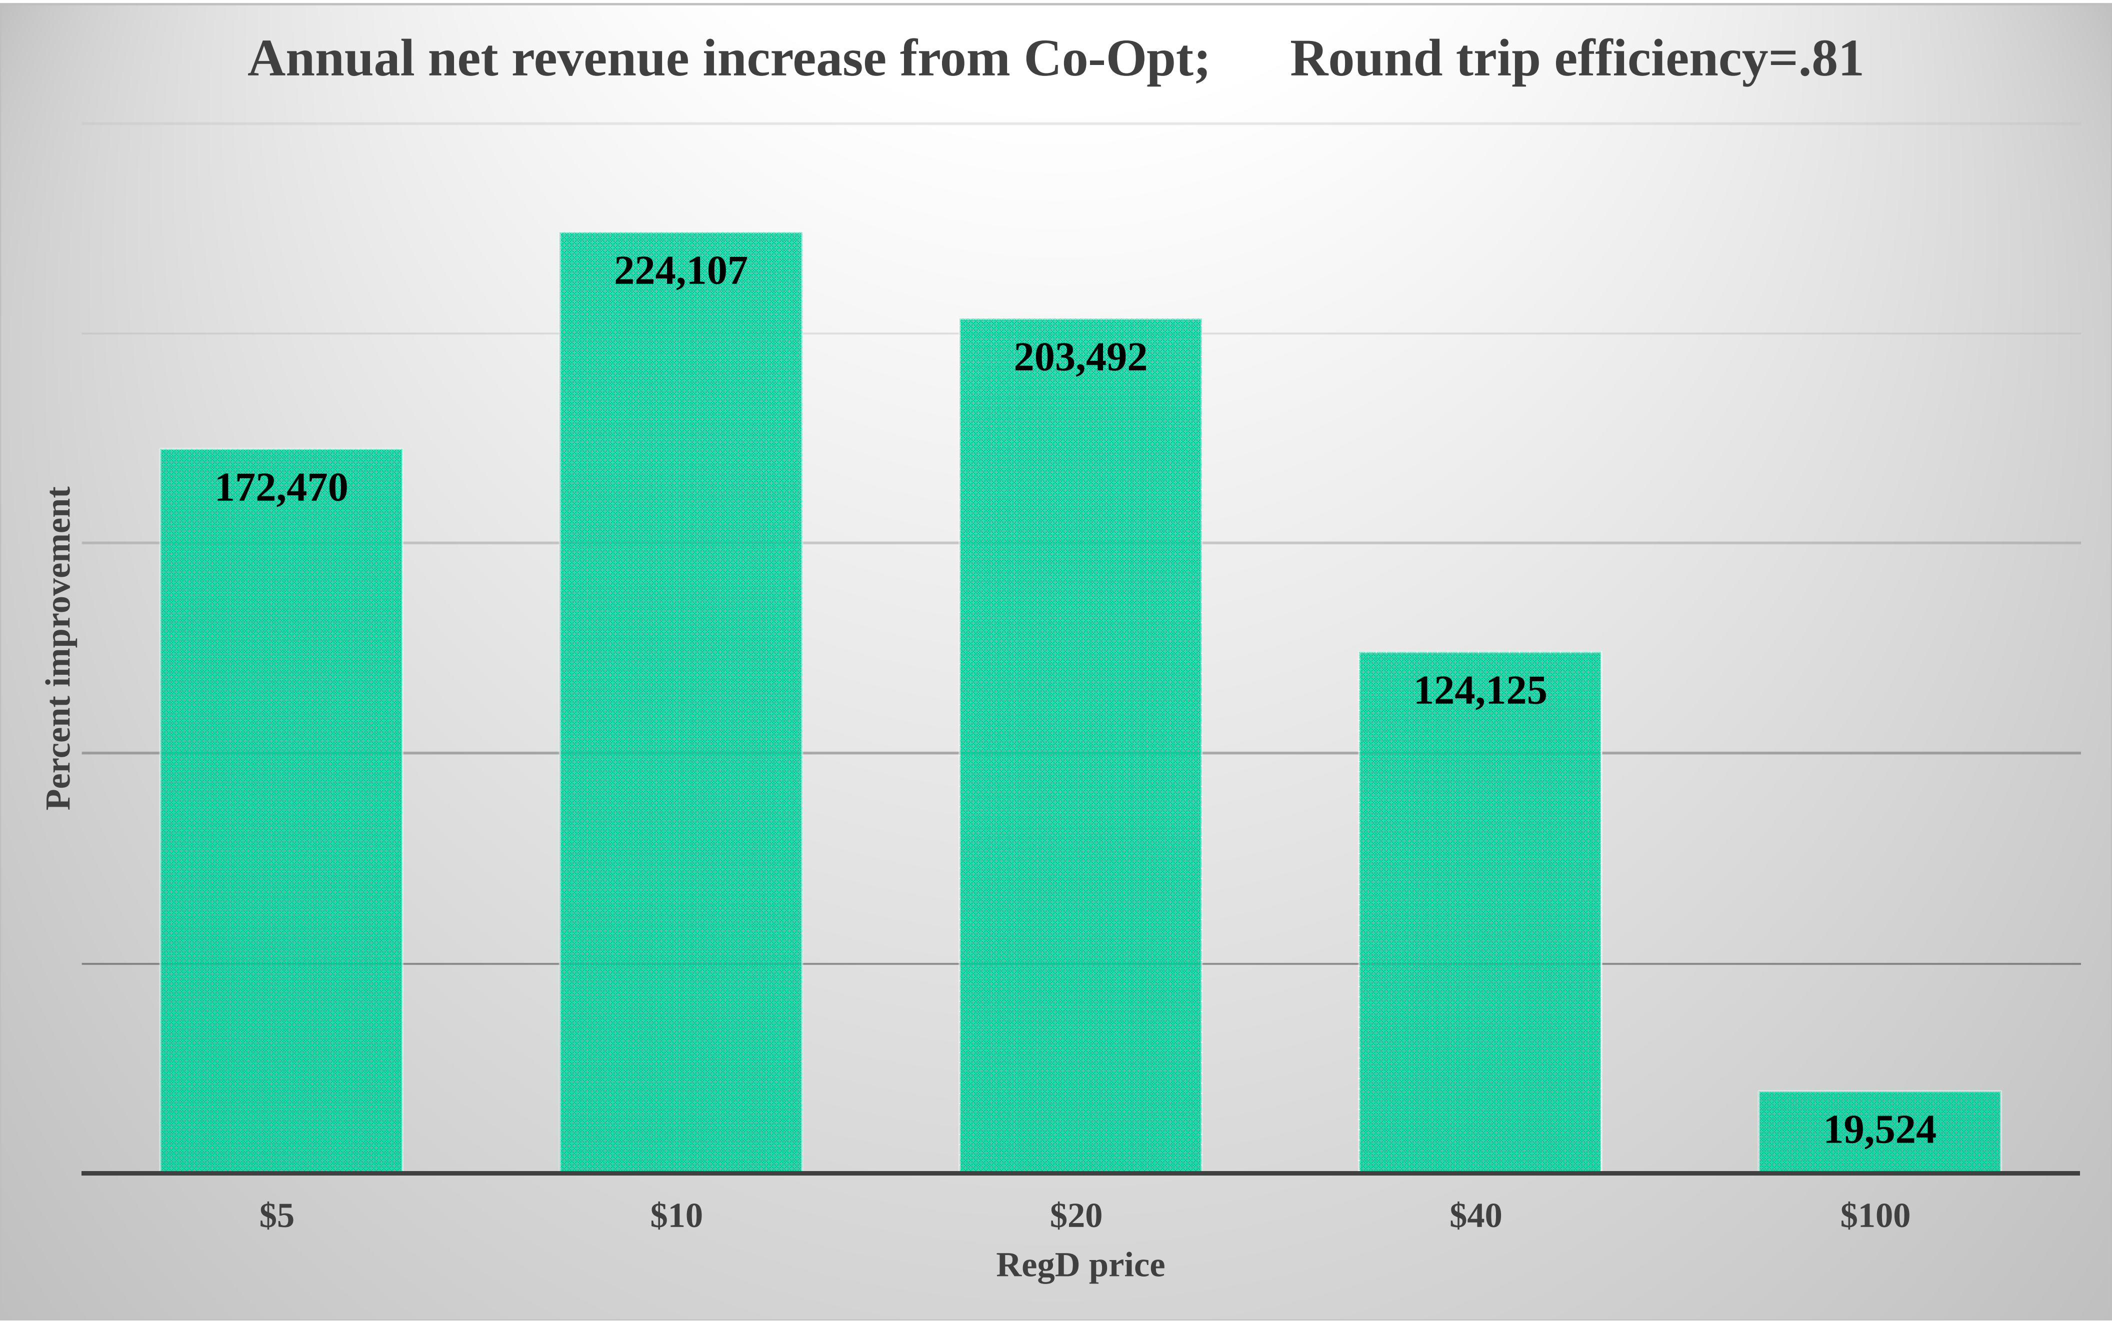
Annual net revenue increase from Co-Opt;   Round trip efficiency=.81
Percent improvement
172,470
224,107
203,492
124,125
19,524
$5
$10
$20
$40
$100
RegD price

Economic analysis

- RegD = $100

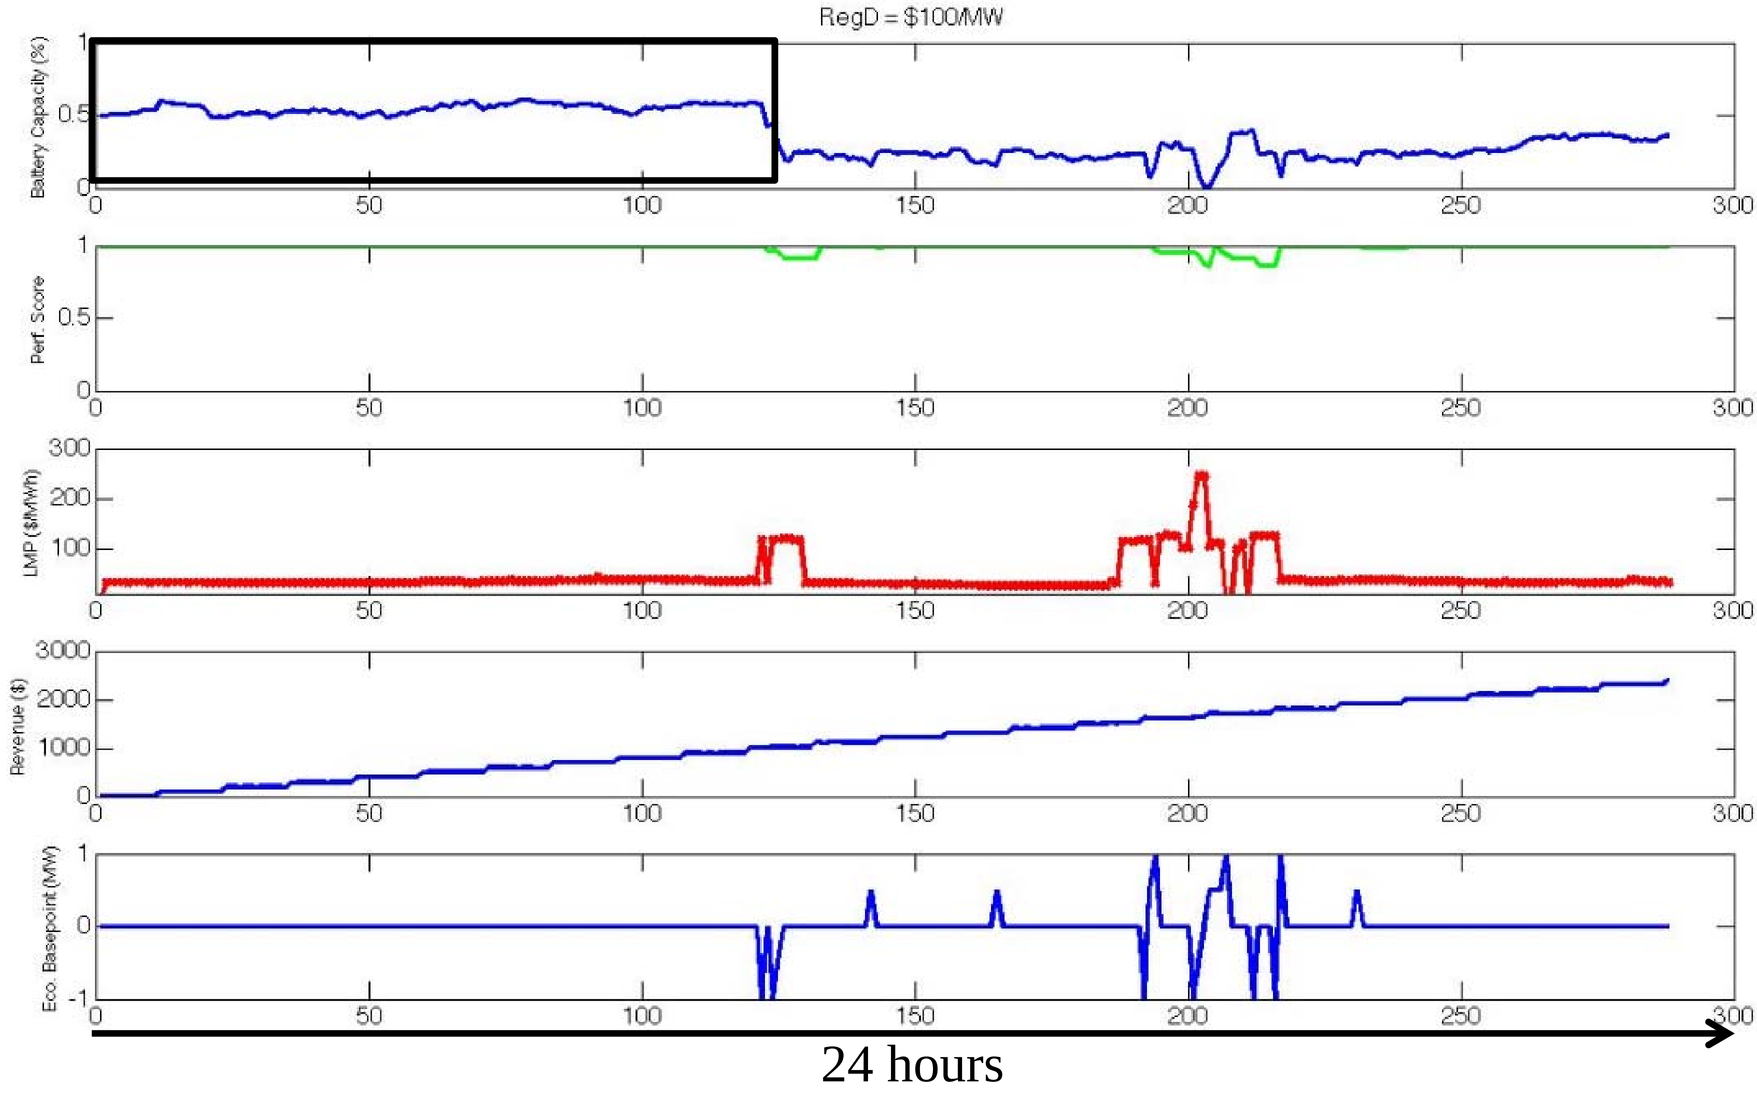

Economic analysis

- RegD = $20

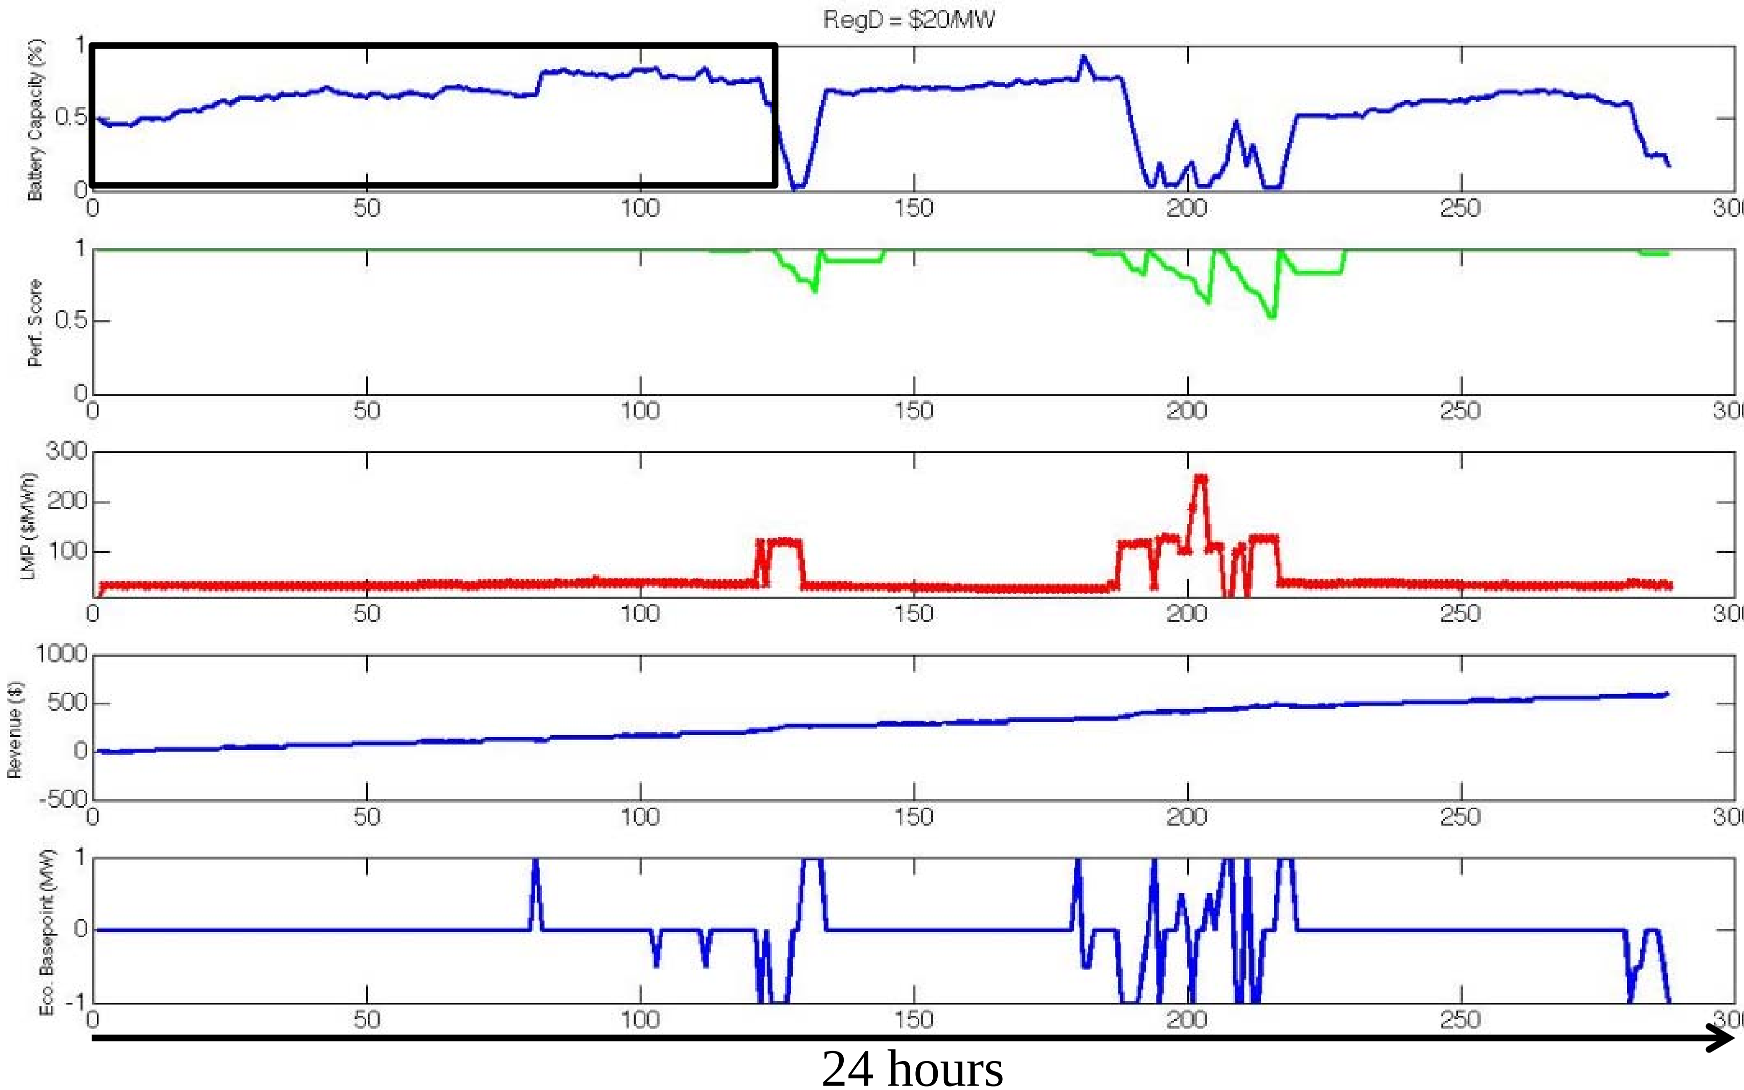

Economic analysis

- RegD = $100

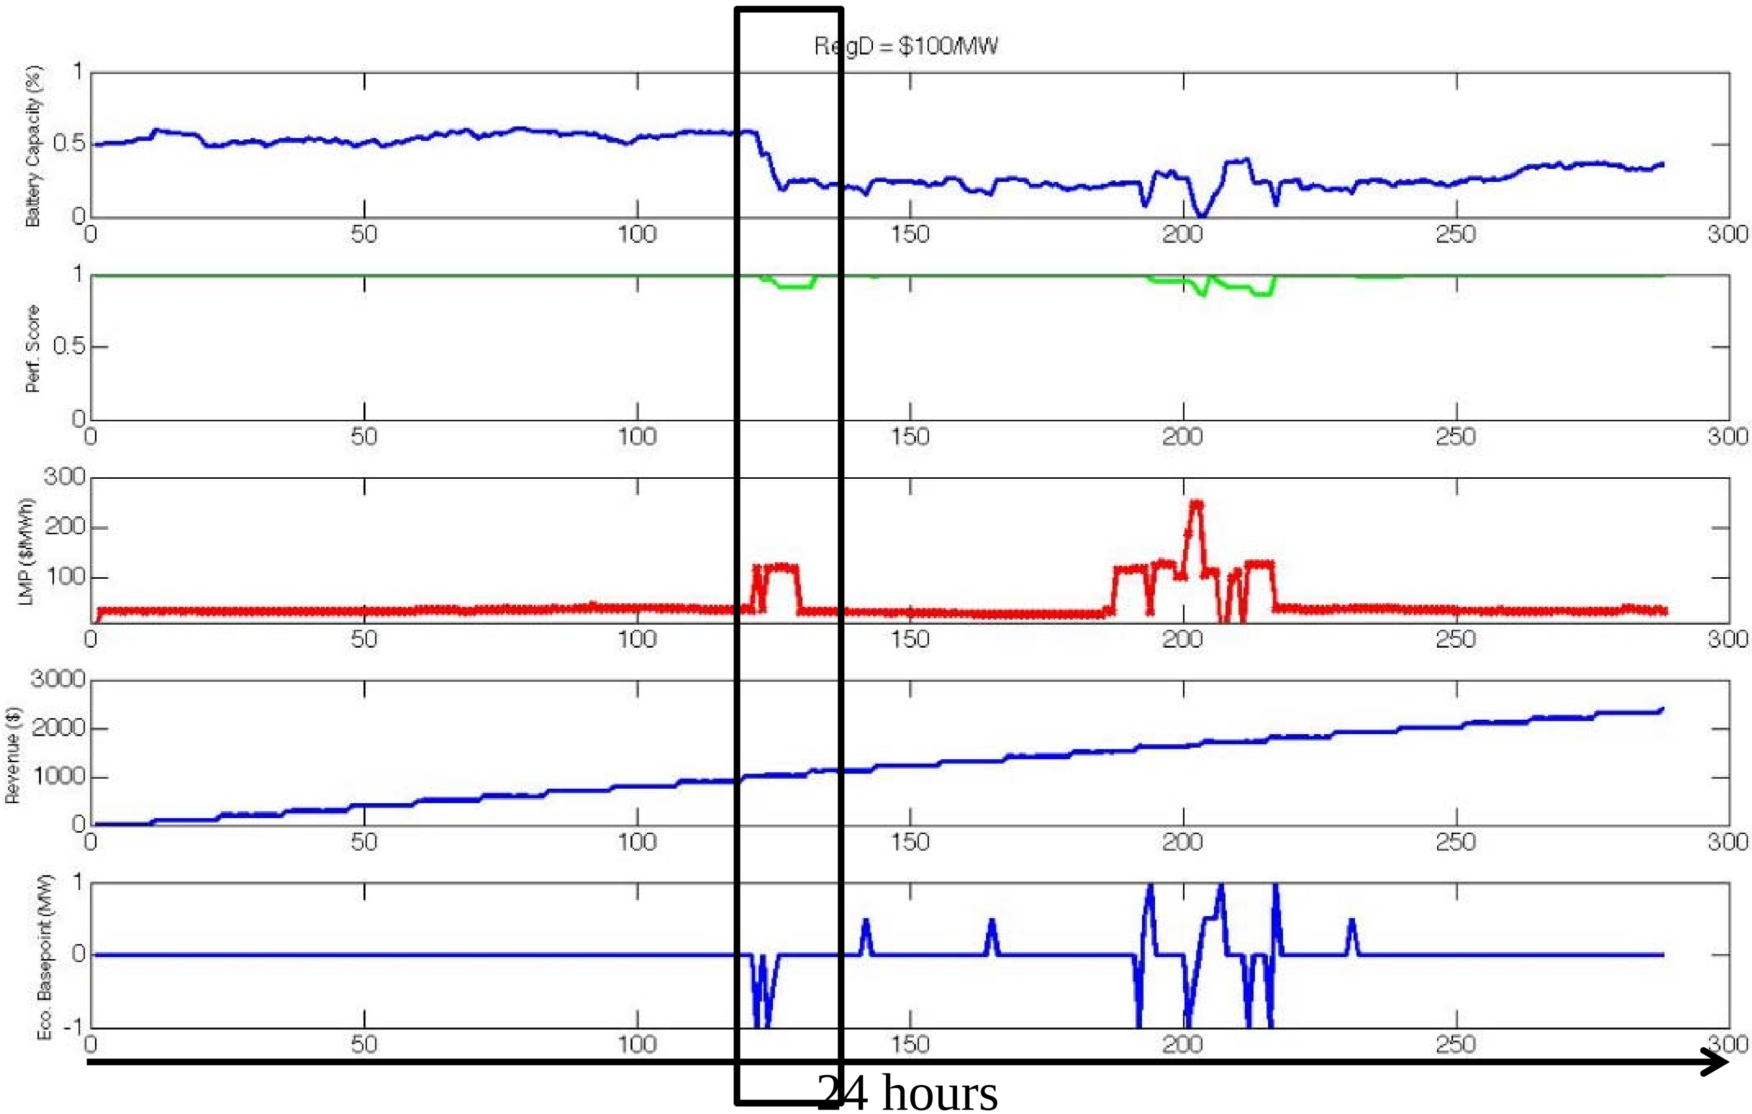

Economic analysis

- RegD = $20

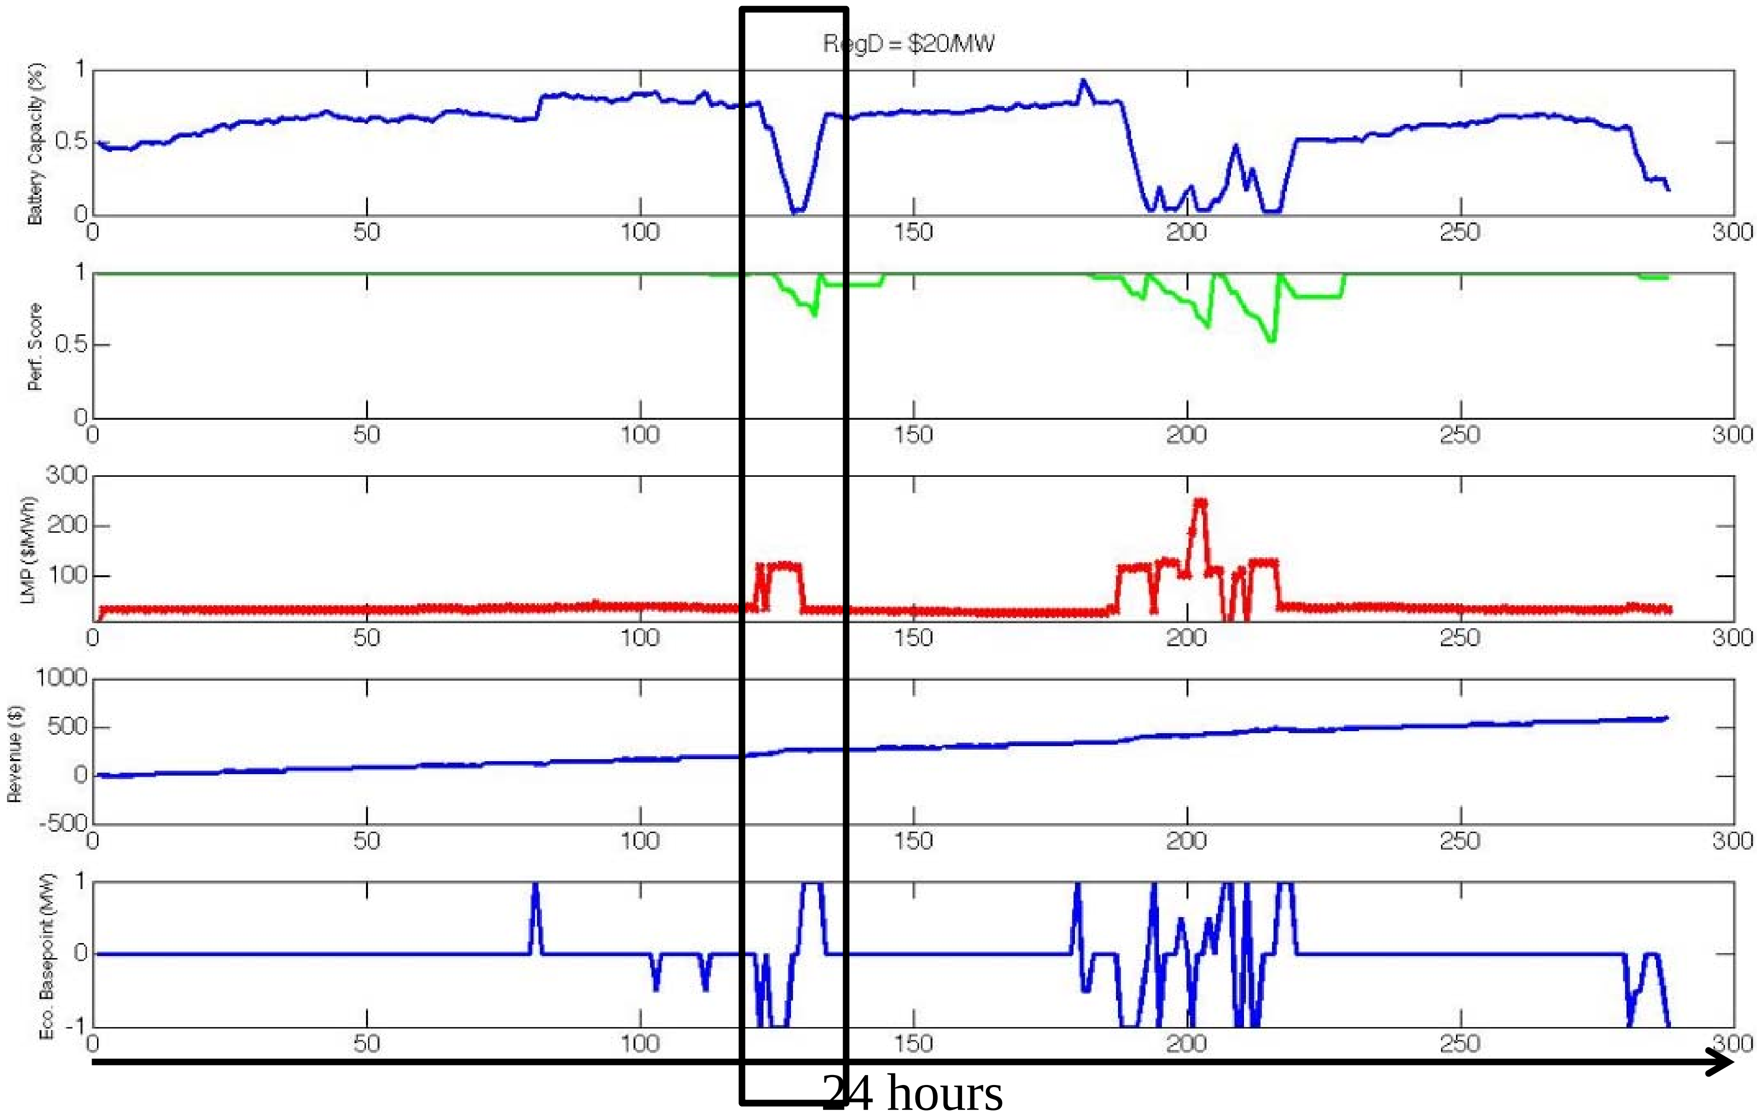

Economic analysis

- RegD = $100

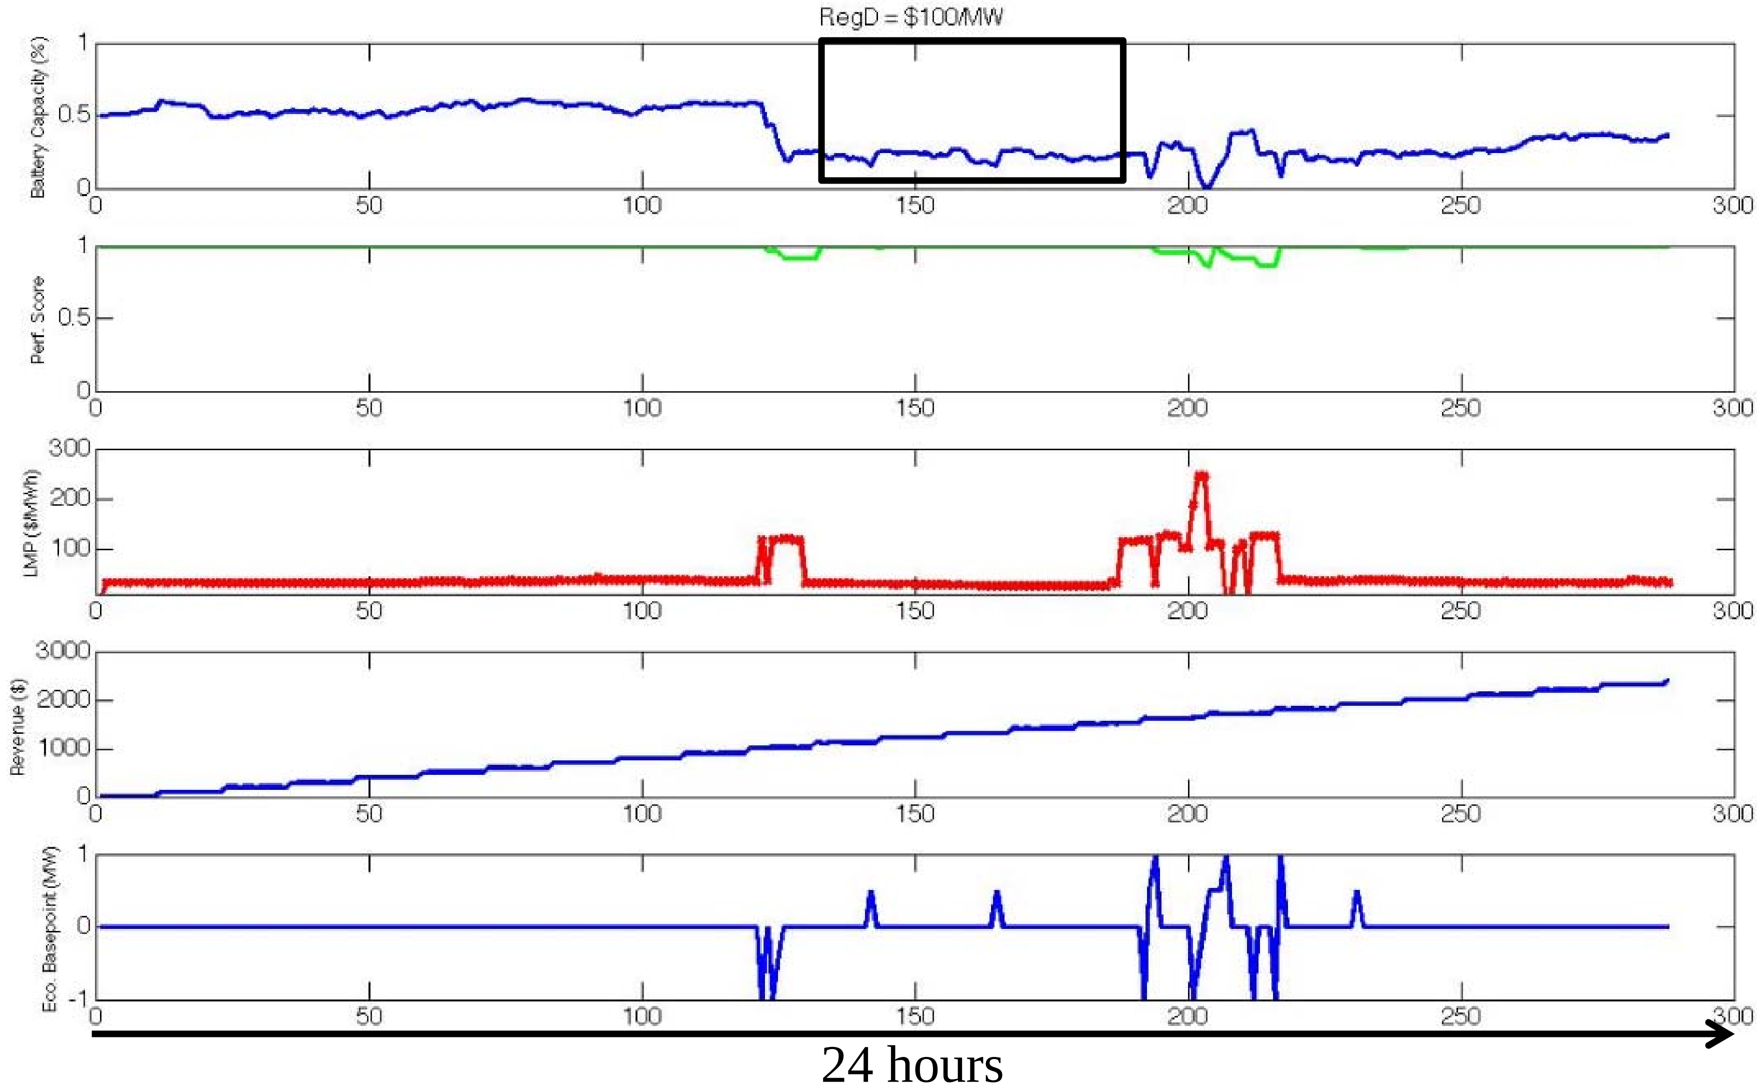

Economic analysis

- RegD = $20

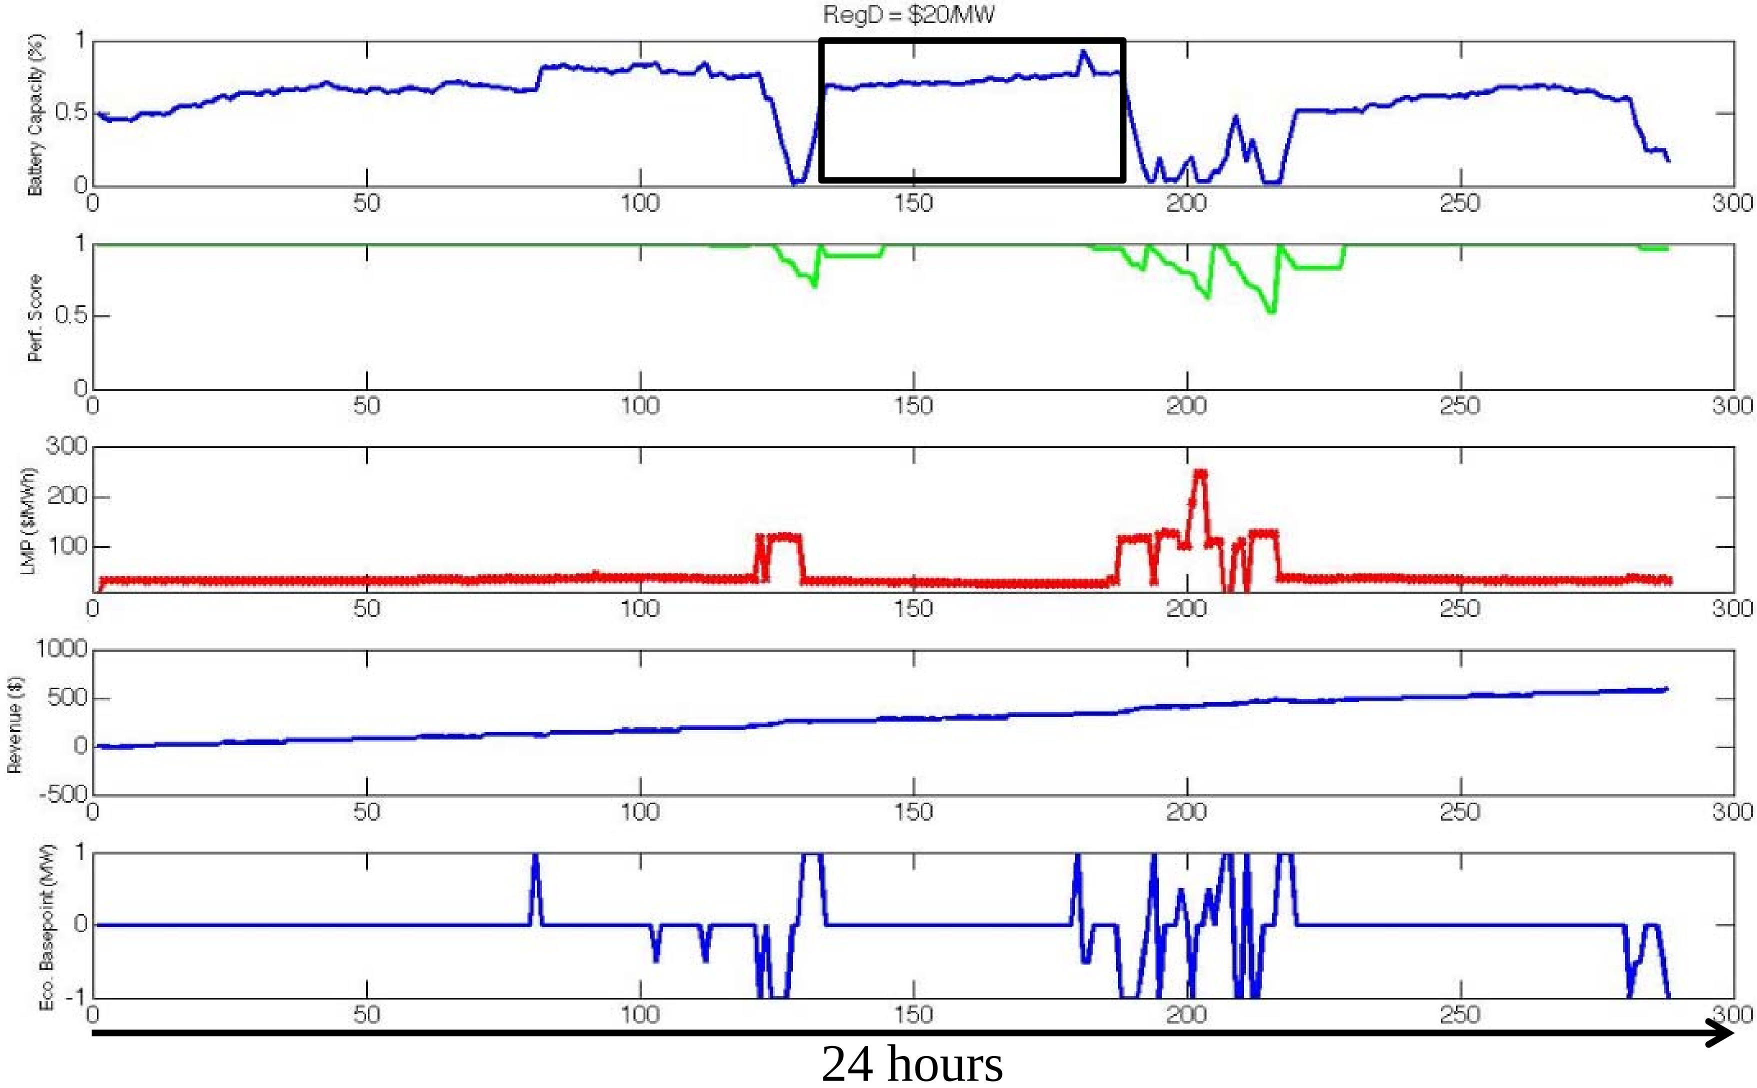

Behaviors

- RegD = $100

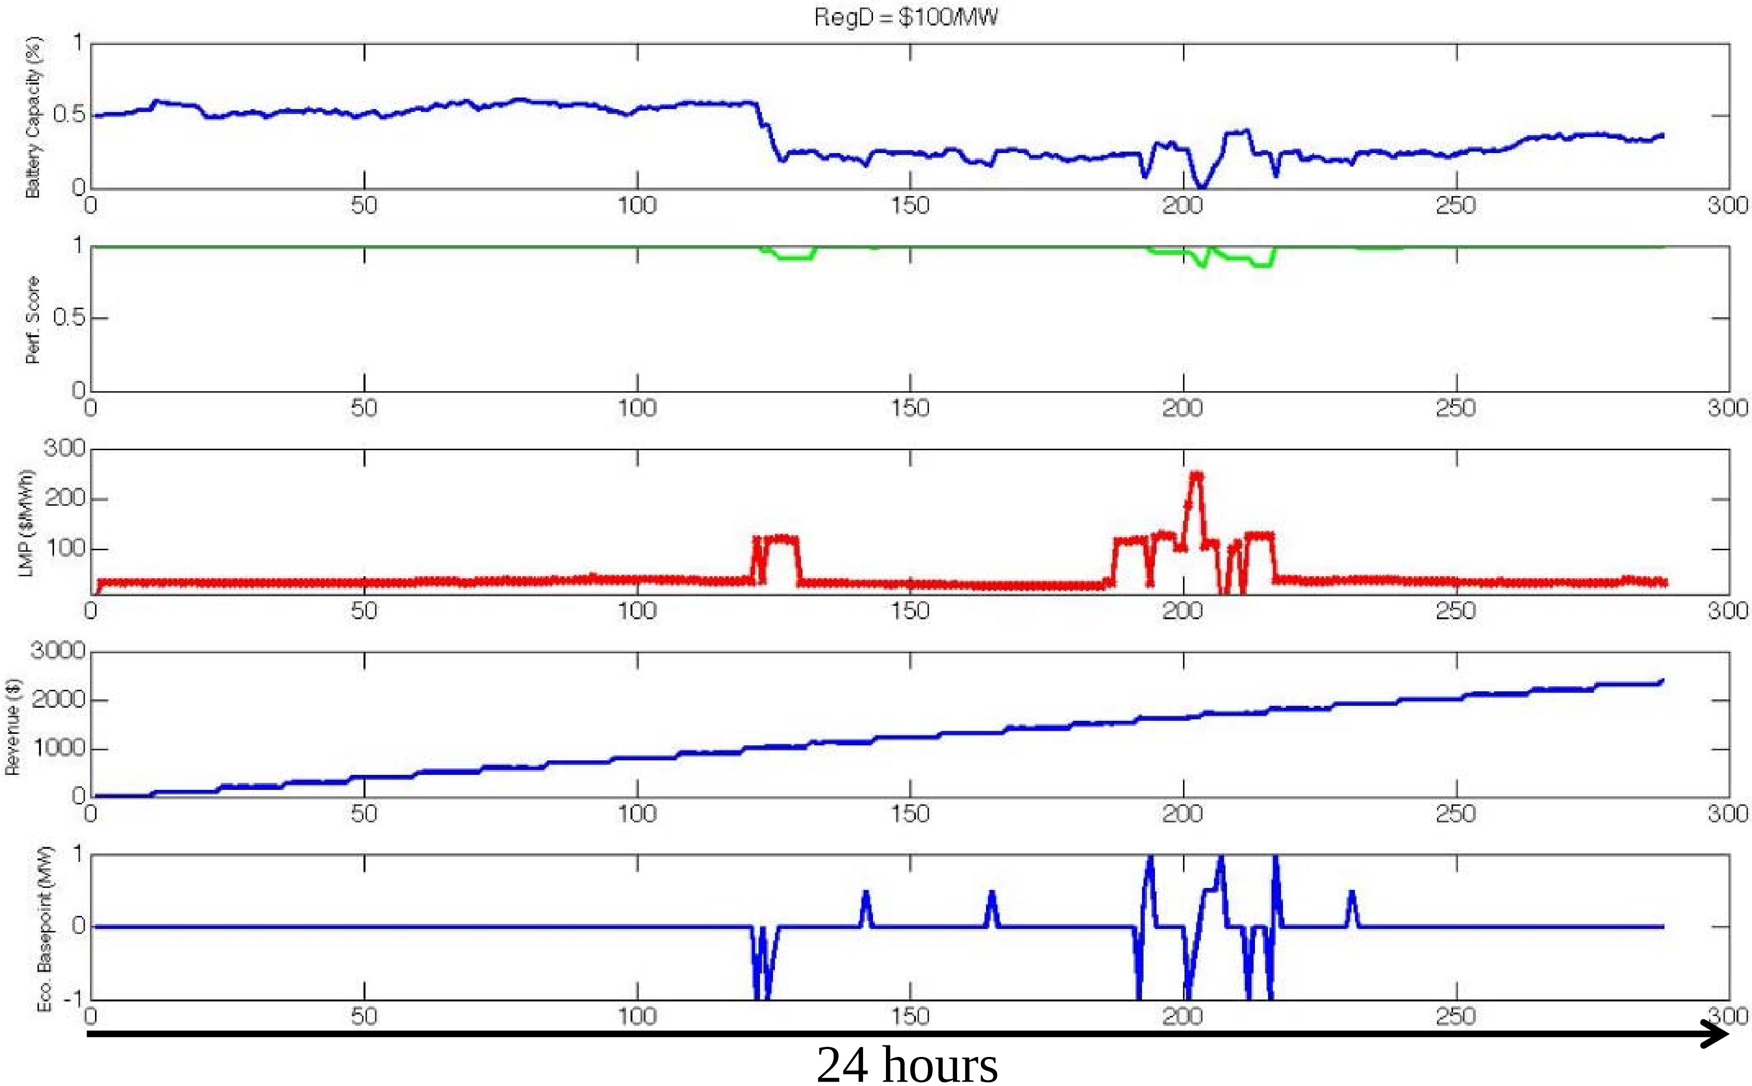

Behaviors

- RegD = $20

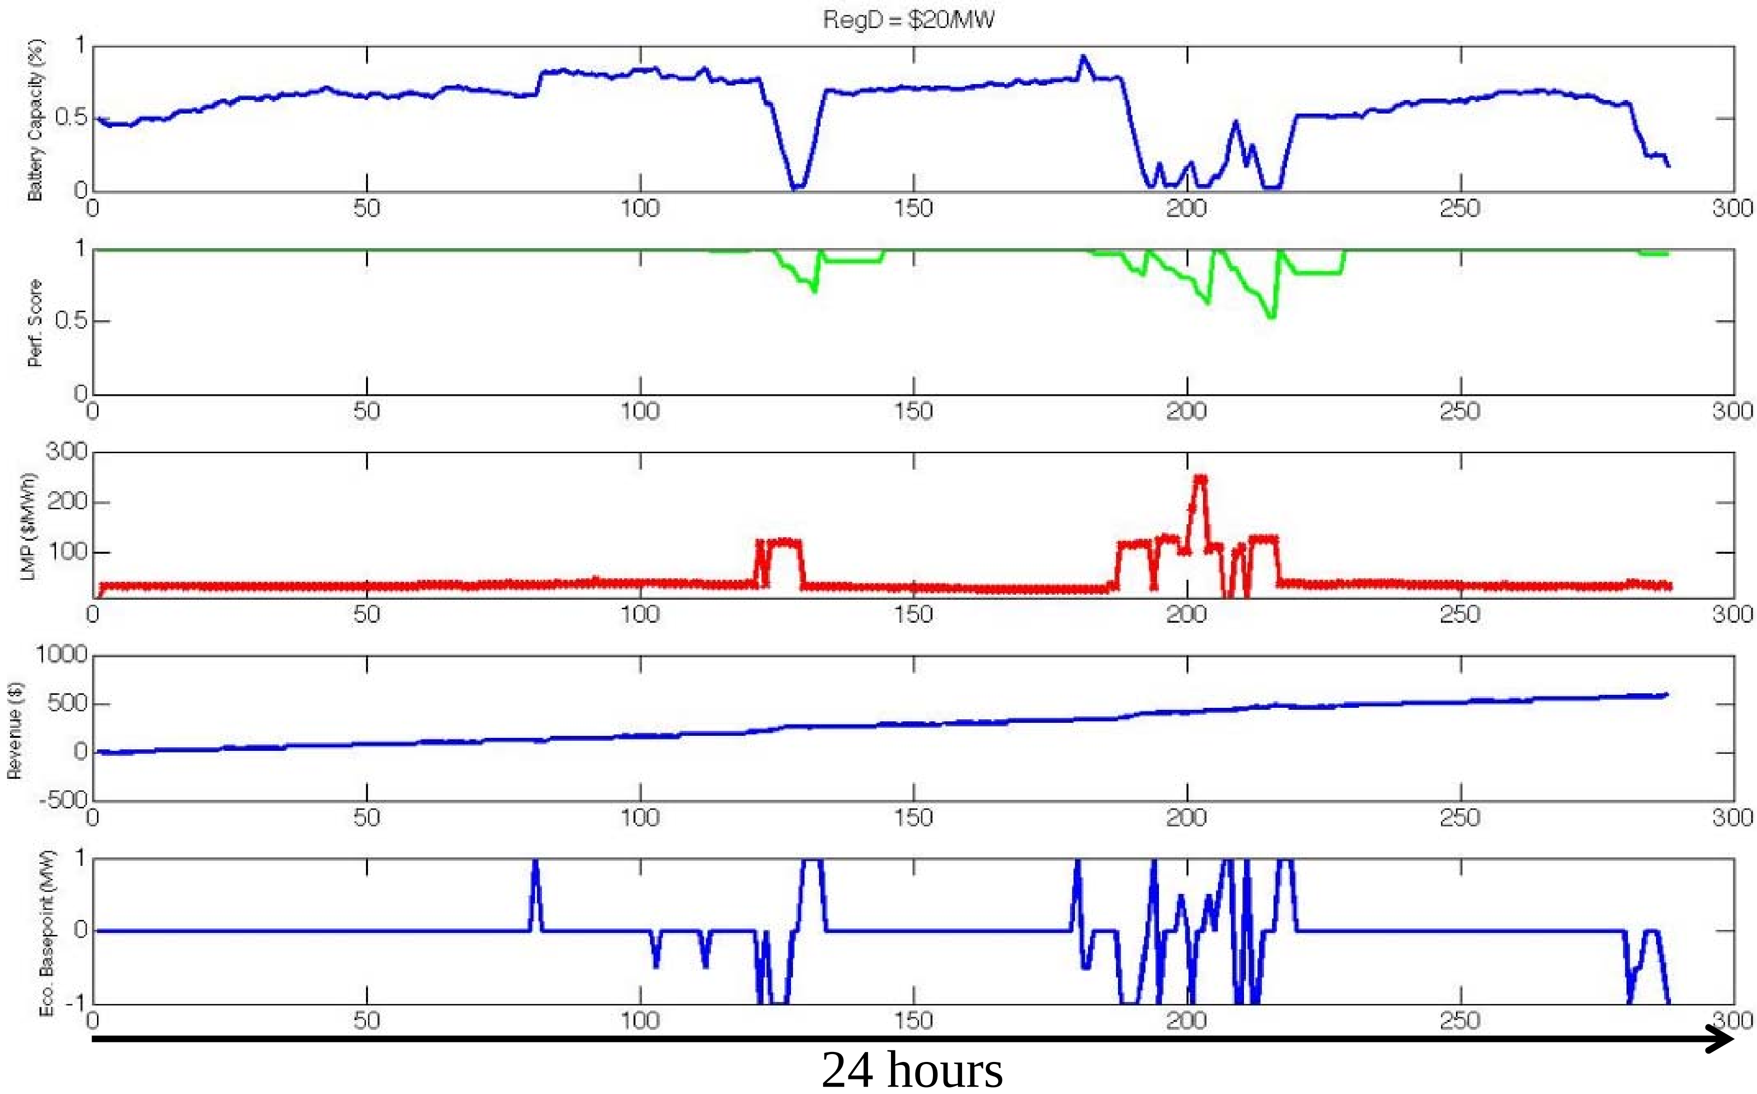

Behaviors

- RegD = $10

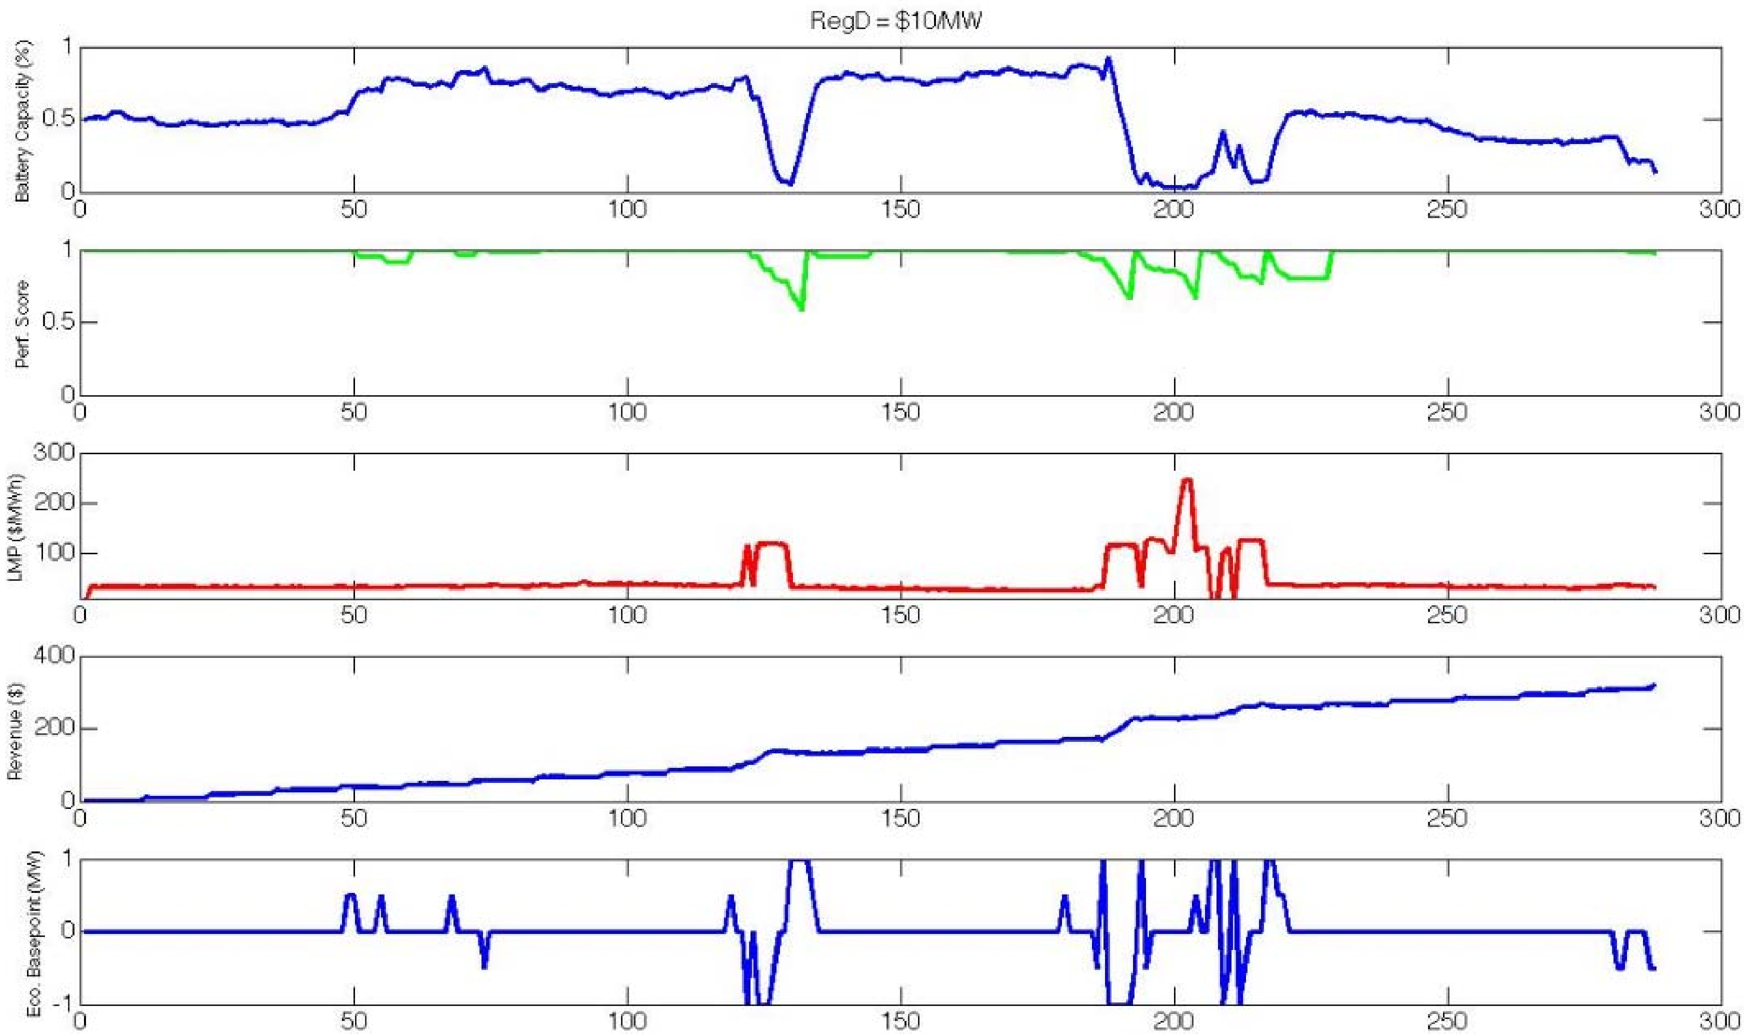

Behaviors

- RegD = $5

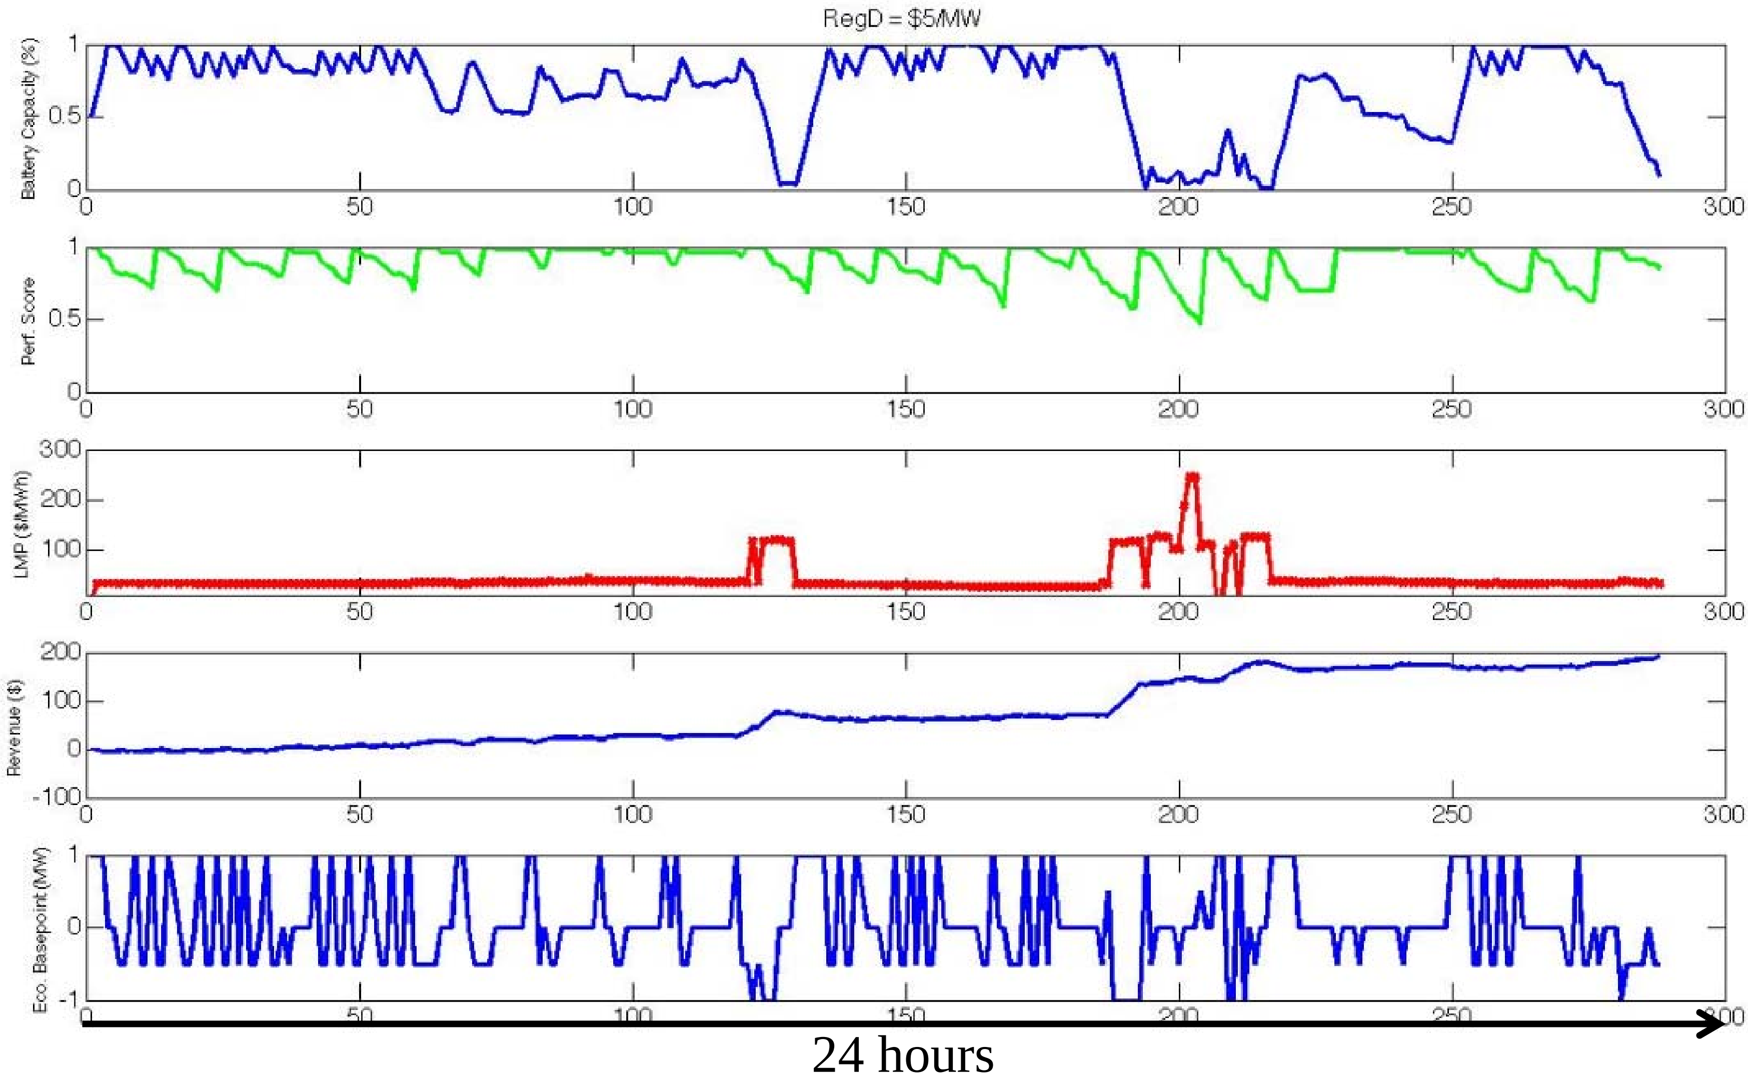

Behaviors

- RegD = $0

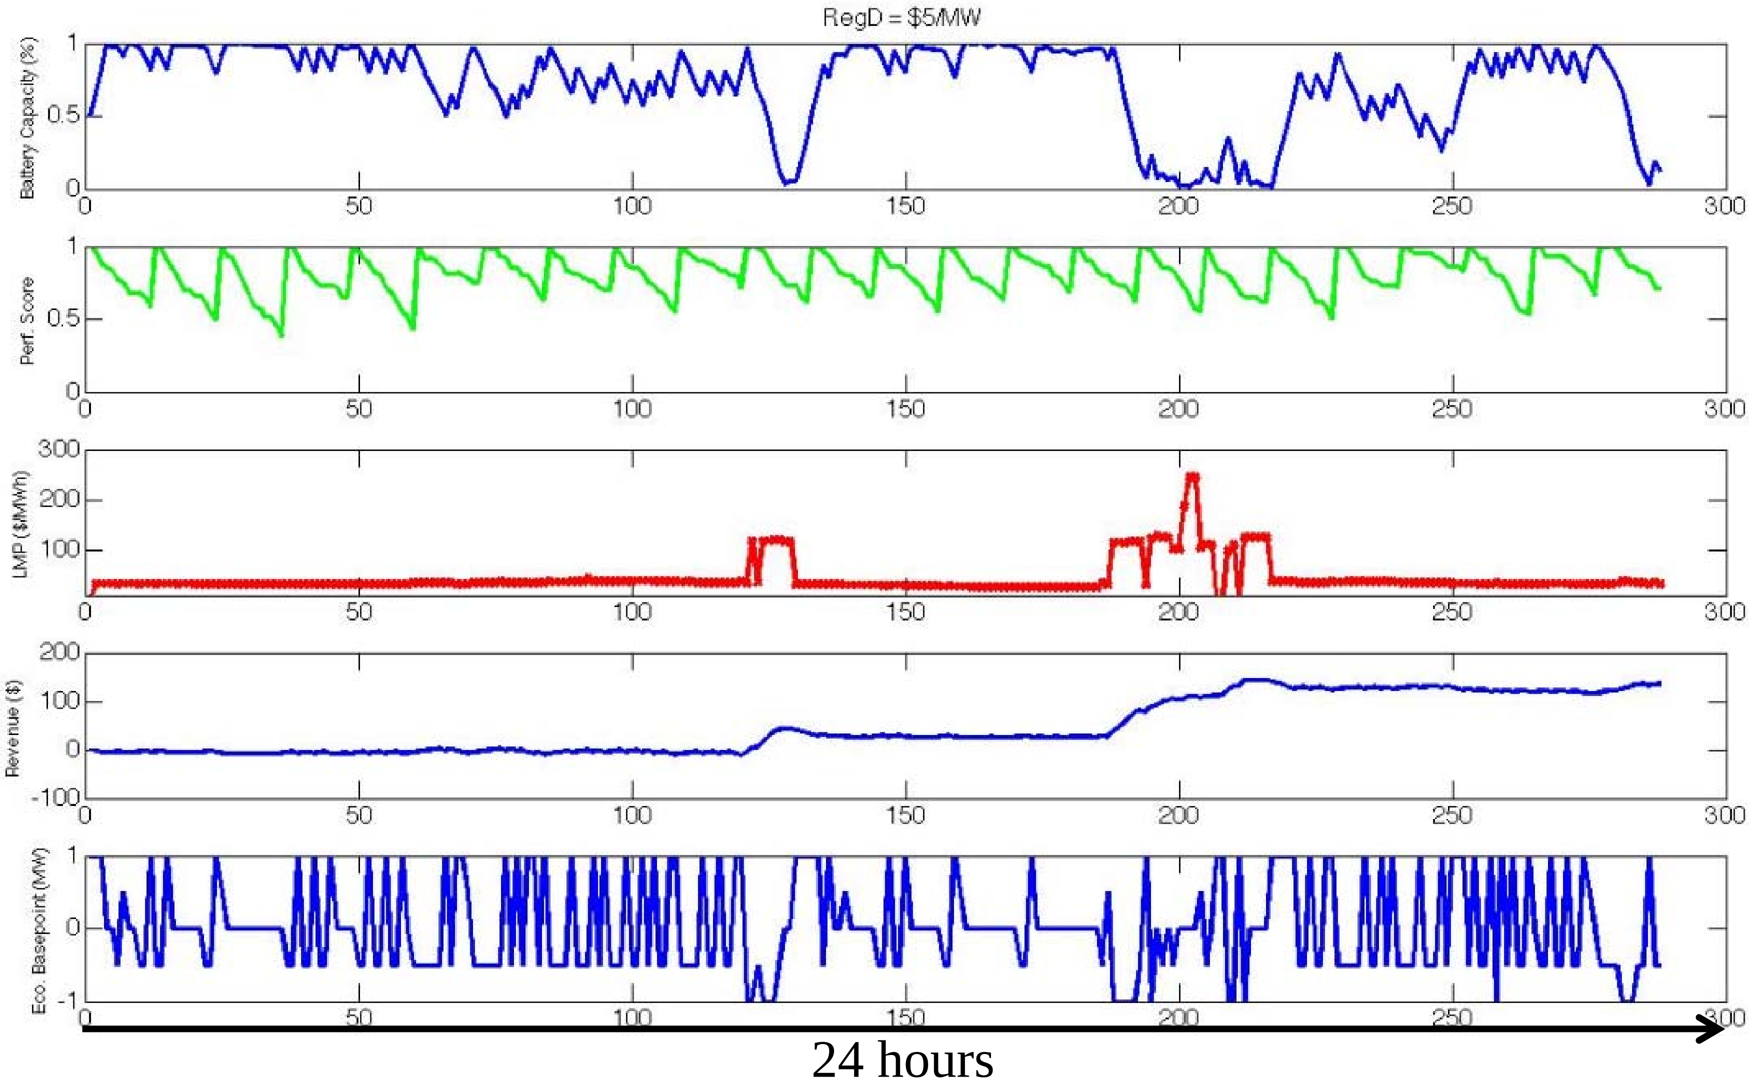

Economic analysis

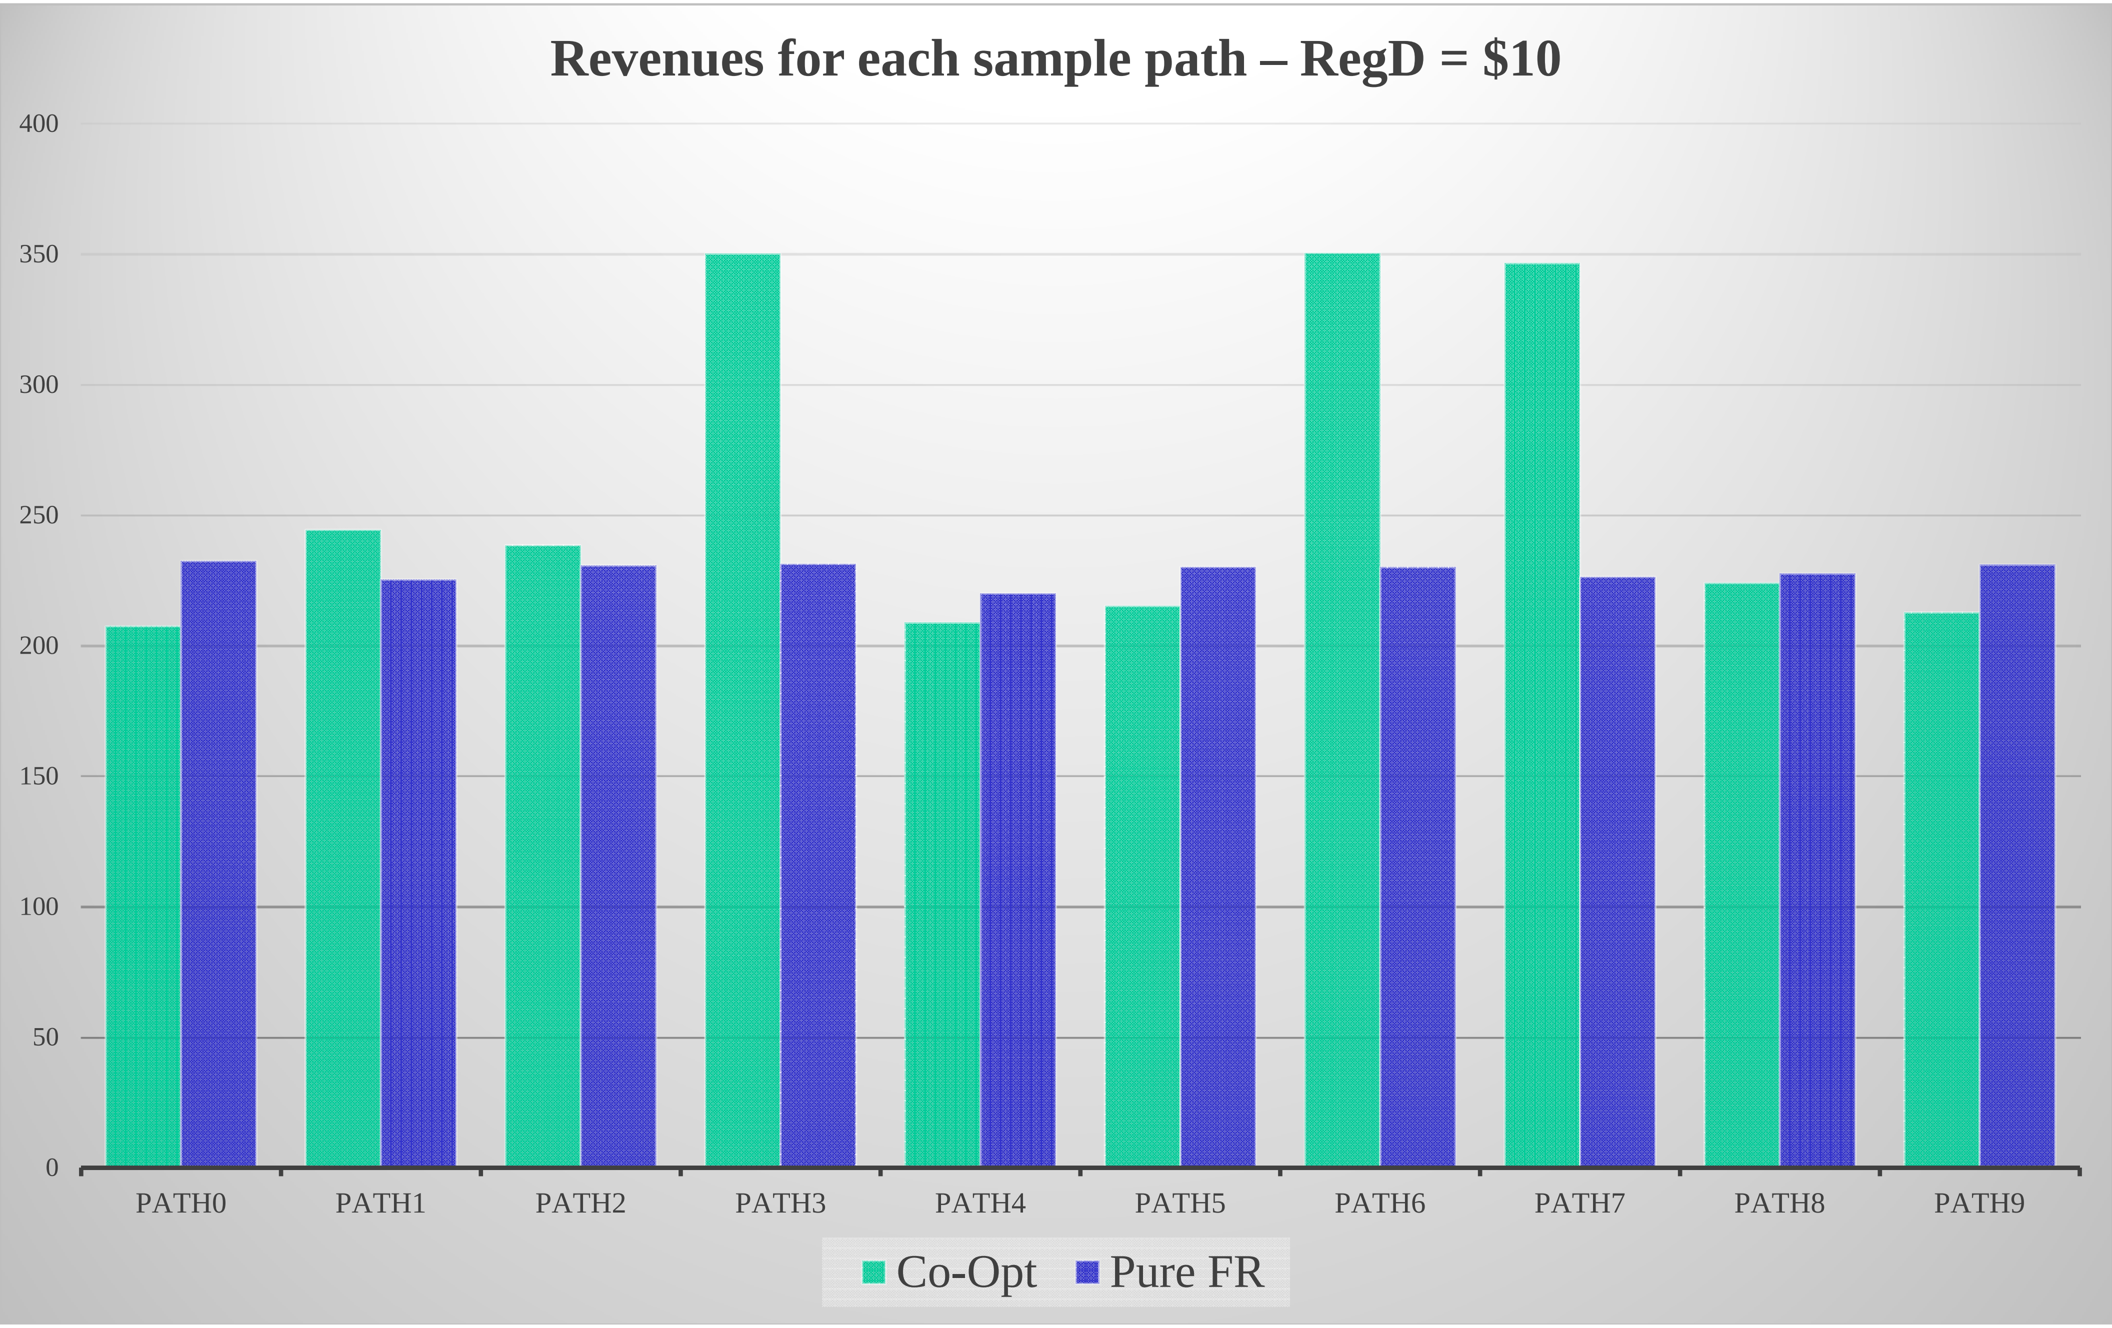

Economic analysis

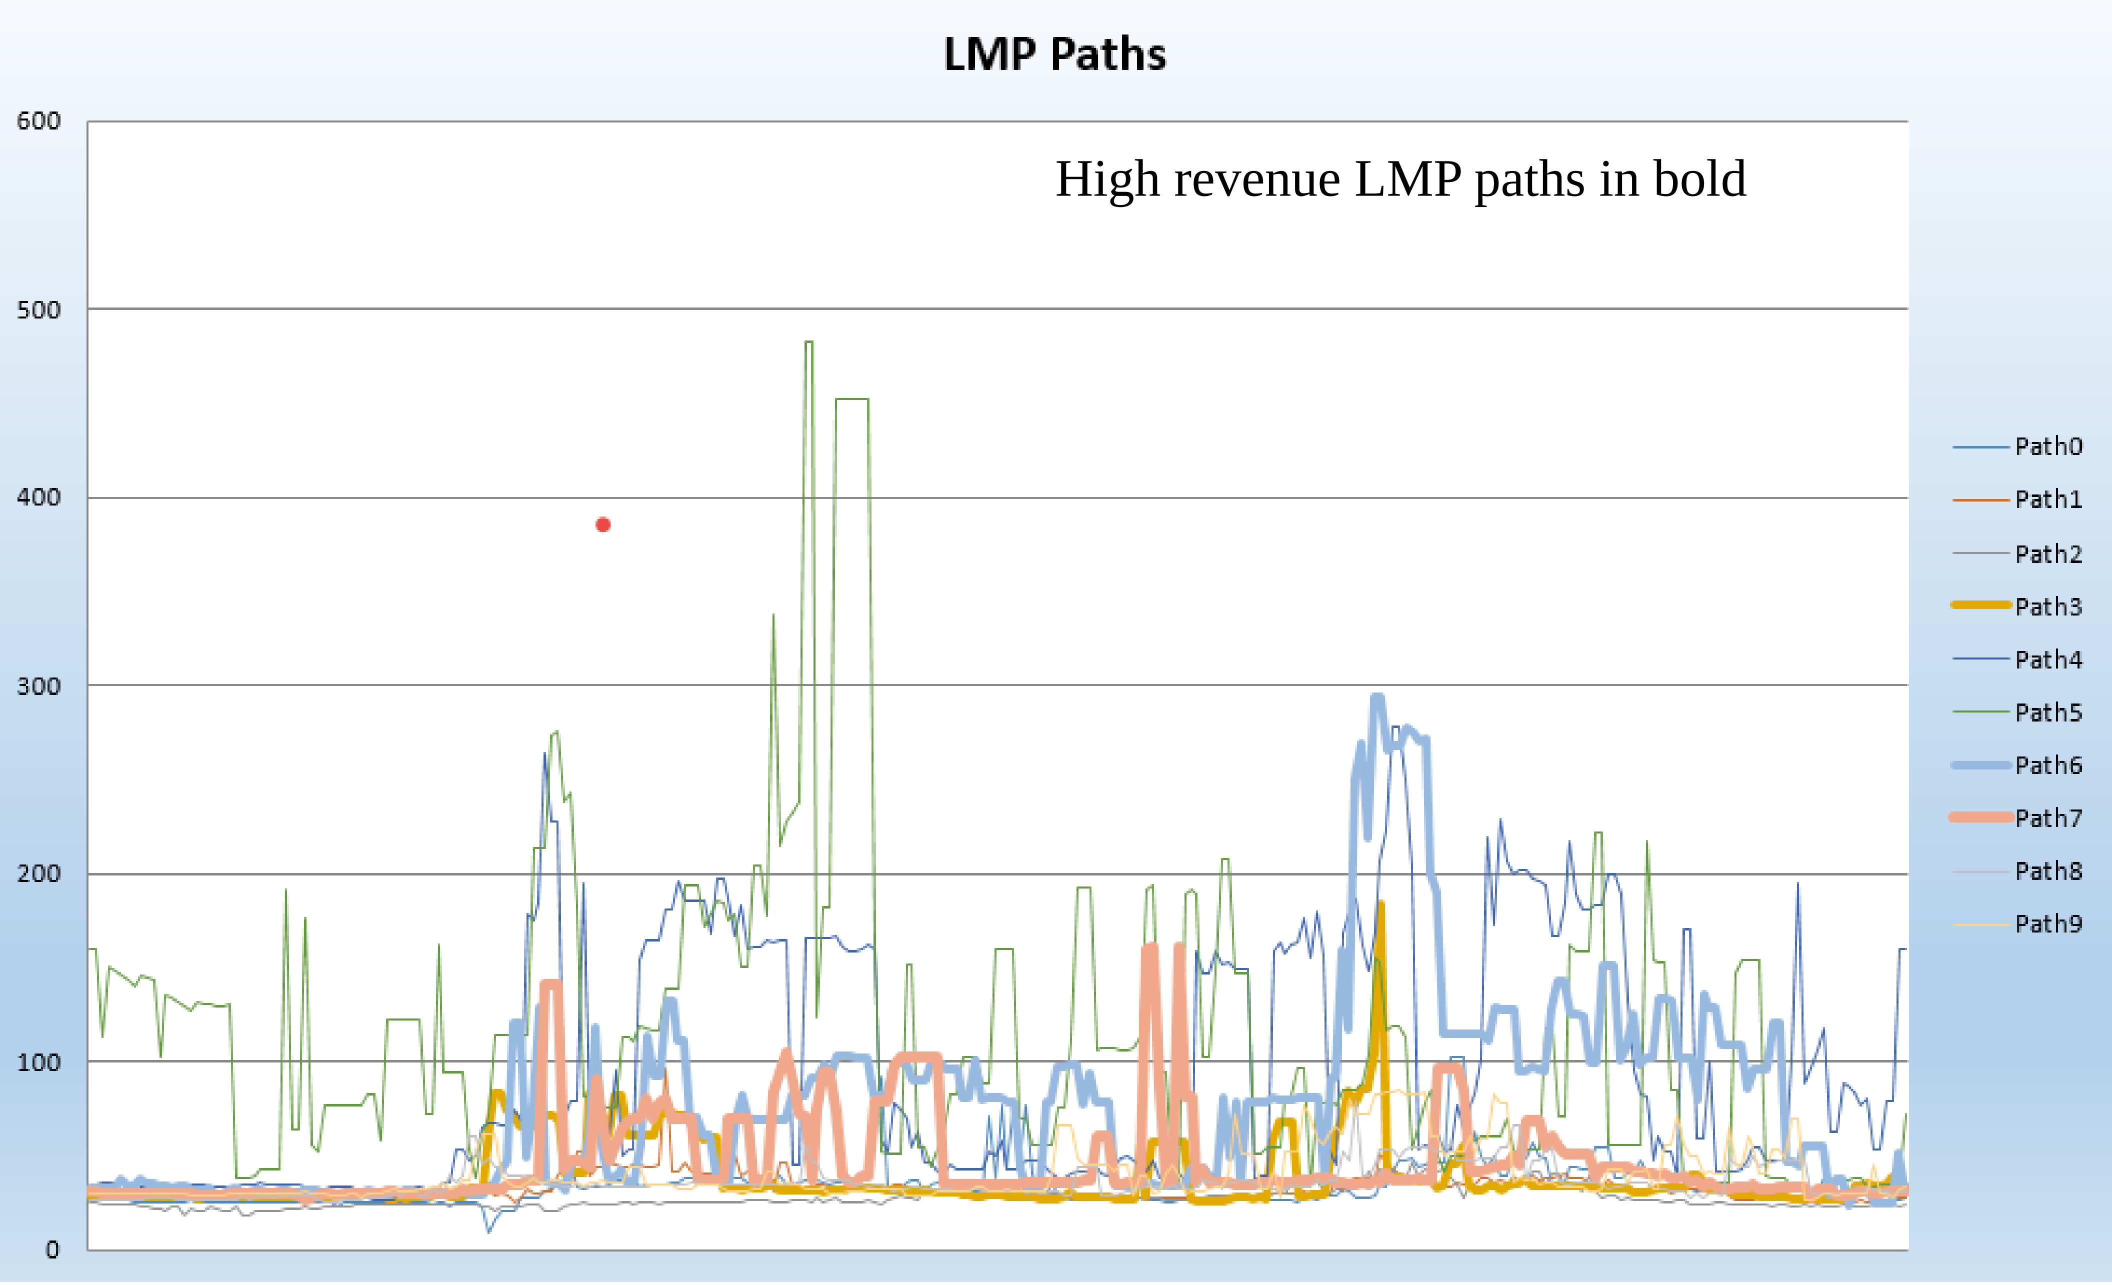

Economic analysis

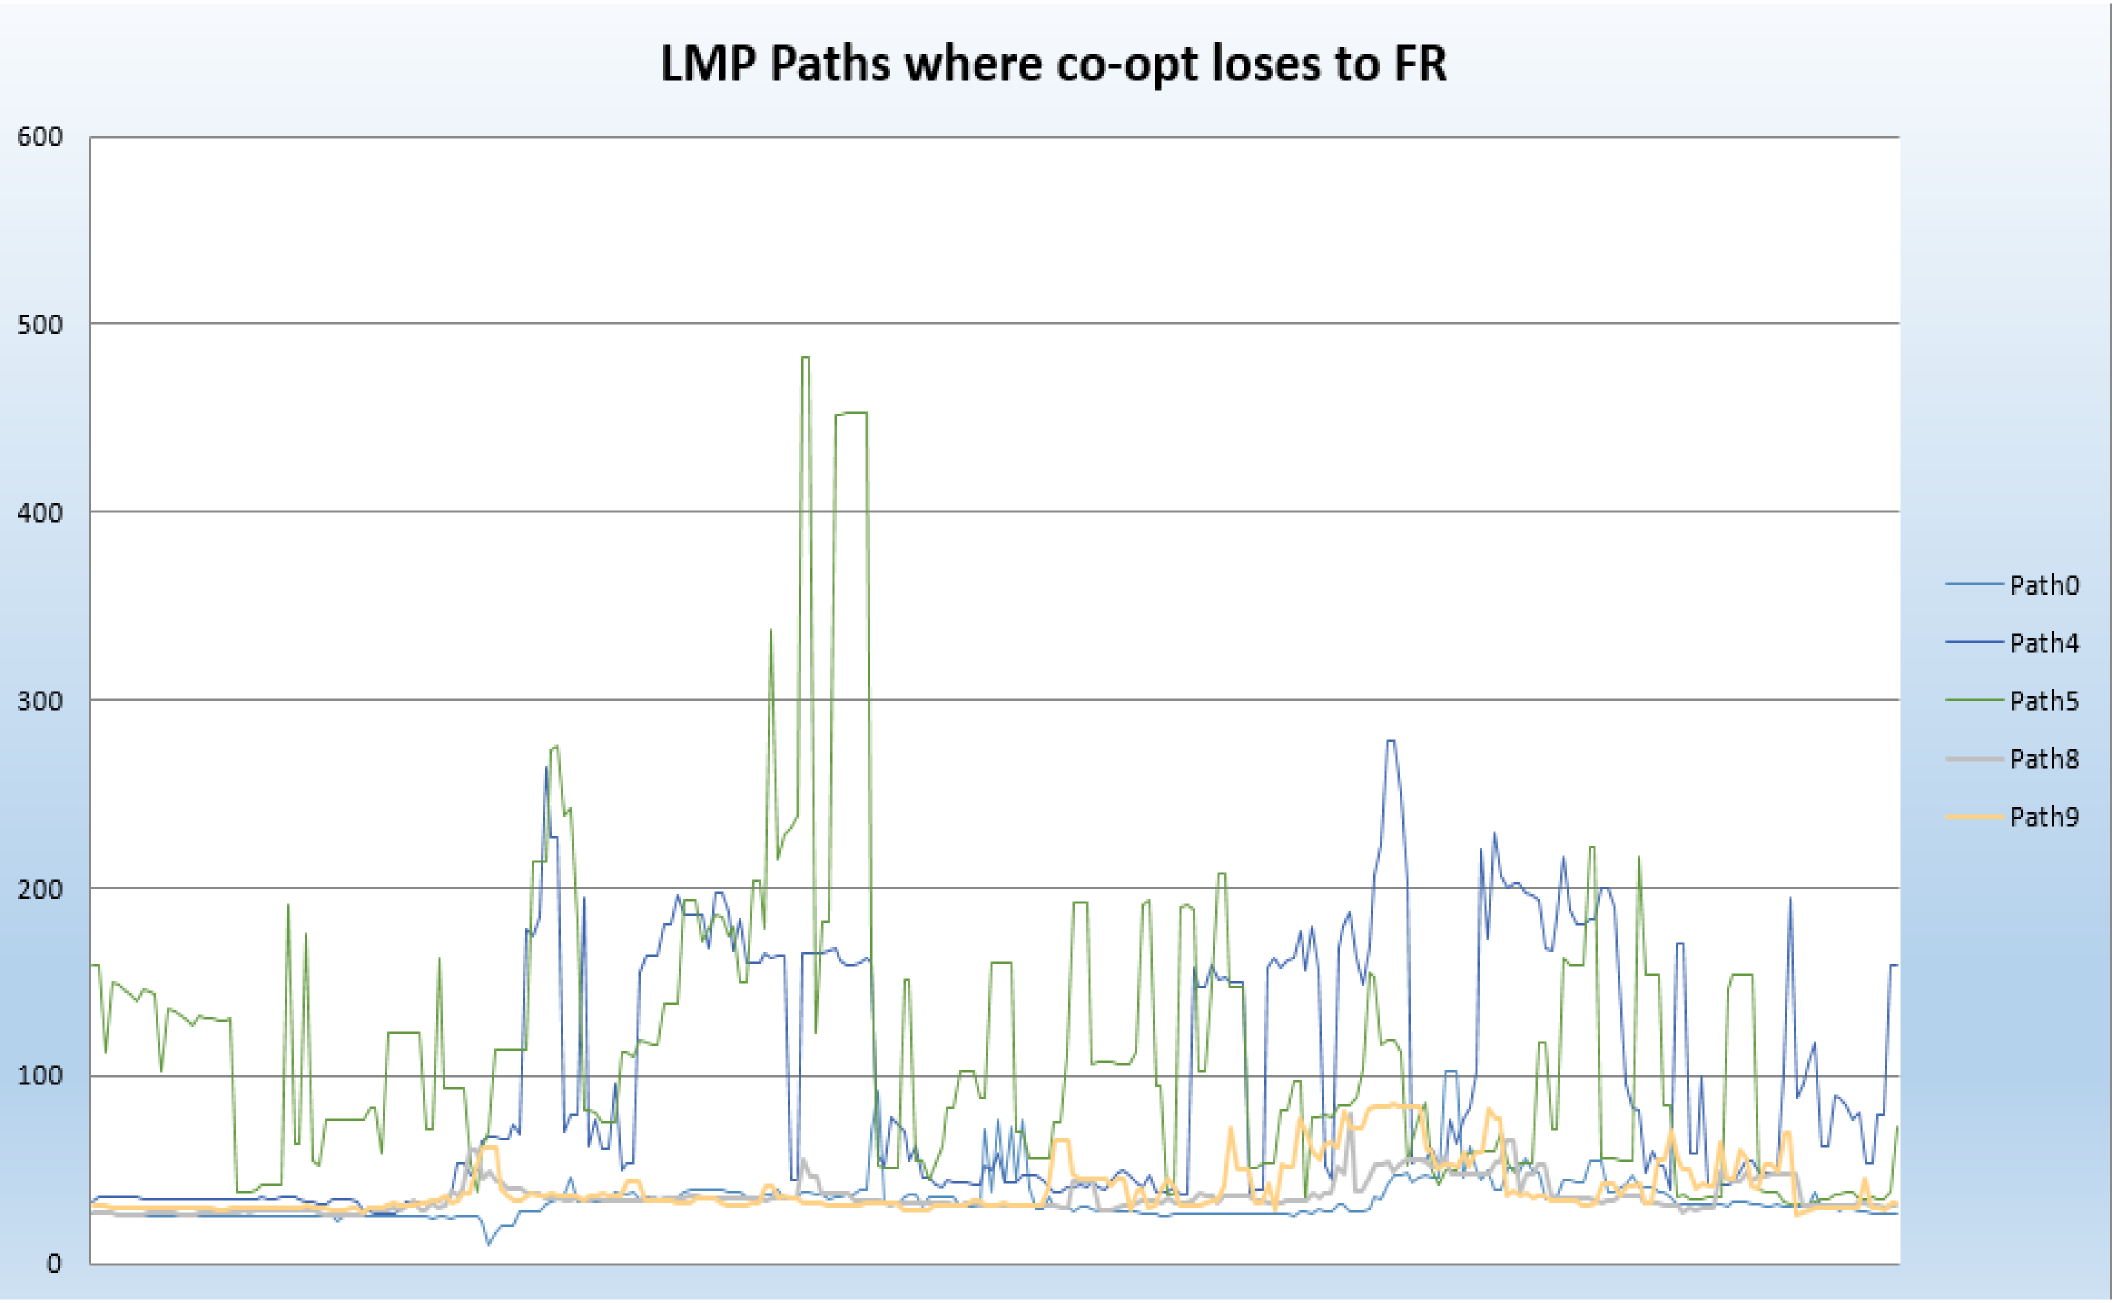

Conclusions

- The algorithm:
  - » Produces very near-optimal optimal policy over all revenue streams and time scales
  - » Balances RegD penalties against other revenue streams
  - » The algorithm can run in real-time (e.g. making RegD decisions every 2 seconds)
  - » But there is considerable offline computation (which we are working to reduce)
- The benefits
  - » In initial studies, value of co-optimization is small
    - Frequency regulation revenues still dominate
  - » We are starting the process of optimizing over a wide range of revenue streams