Online Supplement to SDDP vs. ADP: The Effect of Dimensionality in Multistage Stochastic Optimization for Grid Level Energy Storage

Tsvetan Asamov, Daniel F. Salas and Warren B. Powell
Department of Operations Research and Financial Engineering, Princeton University
{tasamov, dsalas, powell}@princeton.edu

1. Supplementary Materials

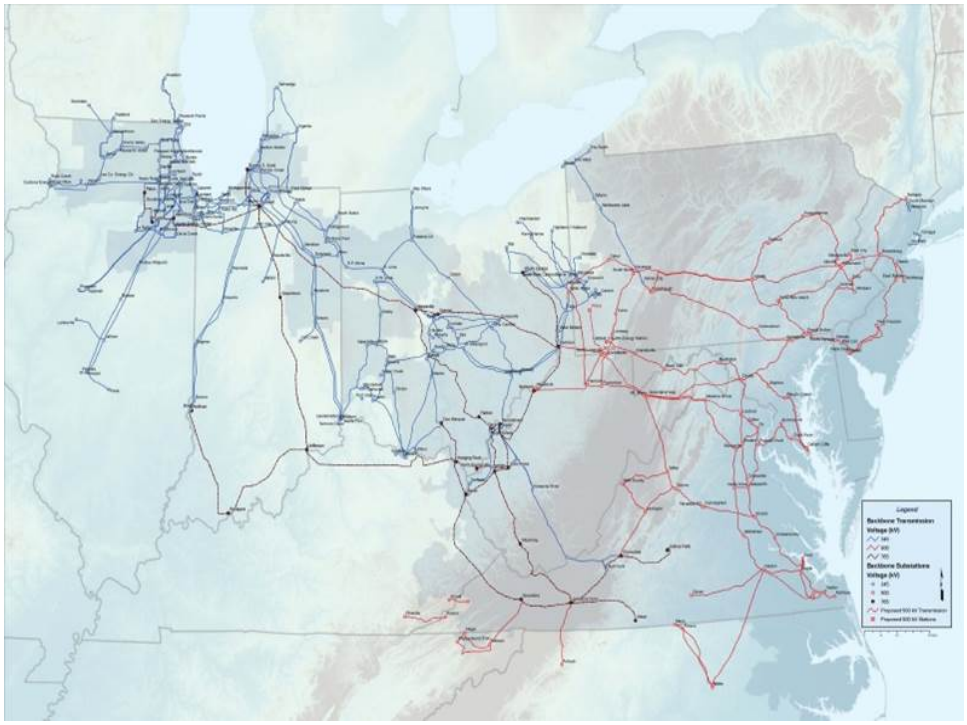

Figure 1 The PJM grid.

Algorithm 1 Markov SDDP Method for Sampled $\widehat{\Omega} \subseteq \Omega$

Choose $Q_t \succeq 0, t = 0, \ldots, T$, and define the sequence $\{\varrho^k\}$.

Define $\widehat{V}_T^k(R_T^x, I_T^x) := V_T^*(R_T^x, I_T^x),\ k = 0, \ldots, K,\ I_T^x \in \mathcal{I}_T^x$.

Define $\widehat{V}_t^0(R_t^x, I_t(\widehat{\omega}_t)) := -\infty,\ \widehat{\omega}_t \in \widehat{\Omega}_t,\ t = 0, \ldots, T-1$.

$(R_{-1}^{x,k}, I_0) \leftarrow S_0,\ k = 0, \ldots, K$

for $k = 0, \ldots, K$ **do**

  Sample $\widehat{\omega} \in \widehat{\Omega}$ using the sampled Markov stochastic process $\{\widehat{W}_t, t = 1, \ldots, T\}$.

  for $t = 0, \ldots, T$ **do**

    if $(k = 0)$ **then**

$$\text{Select } x_t^k \in \underset{x_t \in \mathcal{X}_t(R_{t-1}^{x,k}, I_t(\widehat{\omega}))}{\arg\min} \{C(S_t(\widehat{\omega}), x_t)\}$$

    else

$$x_t^k \in \underset{x_t \in \mathcal{X}_t(R_{t-1}^{x,k}, I_t(\widehat{\omega}))}{\arg\min} \left\{C(S_t(\widehat{\omega}), x_t) + \widehat{V}_t^{k-1}(R_t^x, I_t^x(\widehat{\omega}))\right\}$$

    end if

    Set $R_t^{x,k} \leftarrow B_t^k x_t^k$; $S_{t+1}(\widehat{\omega}) \leftarrow (R_t^{x,k} + \widehat{\omega}_{t+1}, I_{t+1}(\widehat{\omega}))$

  end for

  for $t = T, \ldots, 1$ **do**

$$\text{Define } \underline{V}_t^k(R_{t-1}^x, \widehat{\omega}_t) := \min_{x_t \in \mathcal{X}_t(R_{t-1}^x, I_t(\widehat{\omega}_t))} \left\{C(S_t(\widehat{\omega}_t), x_t) + \widehat{V}_t^k(R_t^x, I_t^x(\widehat{\omega}_t))\right\}$$

    for all $\widehat{\omega}_t \in \widehat{\Omega}_t$ **do**

$$\text{Select } \underline{\beta}_t^k(\widehat{\omega}_t) \in \partial_{R_{t-1}^x} \underline{V}_t^k(R_{t-1}^{x,k}, \widehat{\omega}_t)$$

    end for

    for all $I_{t-1}^x \in \mathcal{I}_{t-1}^x(\widehat{\omega}_{t-1})$ **do**

      $\alpha_{t-1}^k(I_{t-1}^x) \leftarrow \sum_{\widehat{\omega}_t \in \widehat{\Omega}_t} \mathbb{P}(\widehat{\omega}_t | I_{t-1}^x) \underline{V}_t^k(R_t^{x,k}, \widehat{\omega}_t)$; $\beta_{t-1}^k(I_{t-1}^x) \leftarrow \sum_{\widehat{\omega}_t \in \widehat{\Omega}_t} \mathbb{P}(\widehat{\omega}_t | I_{t-1}^x) \underline{\beta}_t^k(\widehat{\omega}_t)$

      $h_{t-1}^k(R_{t-1}^x, I_{t-1}^x) := \alpha_{t-1}^k(I_{t-1}^x) + \langle \beta_{t-1}^k(I_{t-1}^x), R_{t-1}^x - R_{t-1}^{x,k} \rangle$

      $\widehat{V}_{t-1}^k(R_{t-1}^x, I_{t-1}^x) := \max\left\{\widehat{V}_{t-1}^{k-1}(R_{t-1}^x, I_{t-1}^x),\ h_{t-1}^k(R_{t-1}^x, I_{t-1}^x)\right\}$

    end for

  end for

$$\underline{V}_0^k \leftarrow \left\{ \min_{x_0 \in \mathcal{X}_0(S_0)} C(S_0, x_0) + \widehat{V}_0^k(R_0^x, I_0^x) \right\}$$

  $\widehat{R}_t^{x,k} \leftarrow R_t^{x,k},\ t = 0, \ldots, T-1$

end for

Algorithm 2 ADP with Piecewise Linear Separable Value Function Approximations

Set $\delta > 0$.

Define $\overline{V}_T^k(R_T^x) := V^*(R_T^x), k = 0, \ldots, K$.

$(R_{-1}^{x,k}, I_0) \leftarrow S_0,\ k = 0, \ldots, K$

for $k = 0, \ldots, K$ **do**

  Sample $\omega \in \Omega$ from the full Markov stochastic process $\{W_t, t = 1, \ldots, T\}$.

  for $t = 0, \ldots, T$ **do**

    if $(k = 0)$ **then**

$$(x_t^k, R_t^{x,k}) \leftarrow \left\{ \underset{x_t \in \mathcal{X}_t(R_{t-1}^k, I_t(\omega))}{\arg\min} C(S_t(\omega), x_t) \right\}$$

    else

$$(x_t^k, R_t^{x,k}) \leftarrow \left\{ \underset{x_t \in \mathcal{X}_t(R_{t-1}^{x,k}, I_t(\omega))}{\arg\min} C(S_t(\omega), x_t) + \overline{V}_t^k(R_t^x) \right\}$$

    end if

  end for

  for $t = T, \ldots, 1$ **do**

    for $m = 1, \ldots, |R_t^x|$ **do**

$$\hat{x}_{t,m}^{k+} \leftarrow \left\{ \underset{x_t \in \mathcal{X}_t(R_t^{x,k} + \delta \cdot \mathbf{e}_m, I_t^x(\omega))}{\arg\min} C(S_t, x_t) + \overline{V}_t^{k-1}(R_t^x) \right\}$$

$$\hat{x}_{t,m}^{k-} \leftarrow \left\{ \underset{x_t \in \mathcal{X}_t(R_t^{x,k} - \delta \cdot \mathbf{e}_m, I_t^x(\omega))}{\arg\min} C(S_t, x_t) + \overline{V}_t^{k-1}(R_t^x) \right\}.$$

      $\overline{V}_{t-1,m}^k(R_{t-1,m}^x) \leftarrow CAVE(\hat{x}_{t,m}^{k+}, \hat{x}_{t,m}^{k-}, \overline{V}_{t-1,m}^{k-1}(R_{t-1,m}^x))$

    end for

    $\overline{V}_{t-1}^{k+1}(R_{t-1}^x) := \sum_{m=1}^{|R|} \overline{V}_{t-1,m}^{k+1}(R_{t-1,m}^x)$

  end for

end for

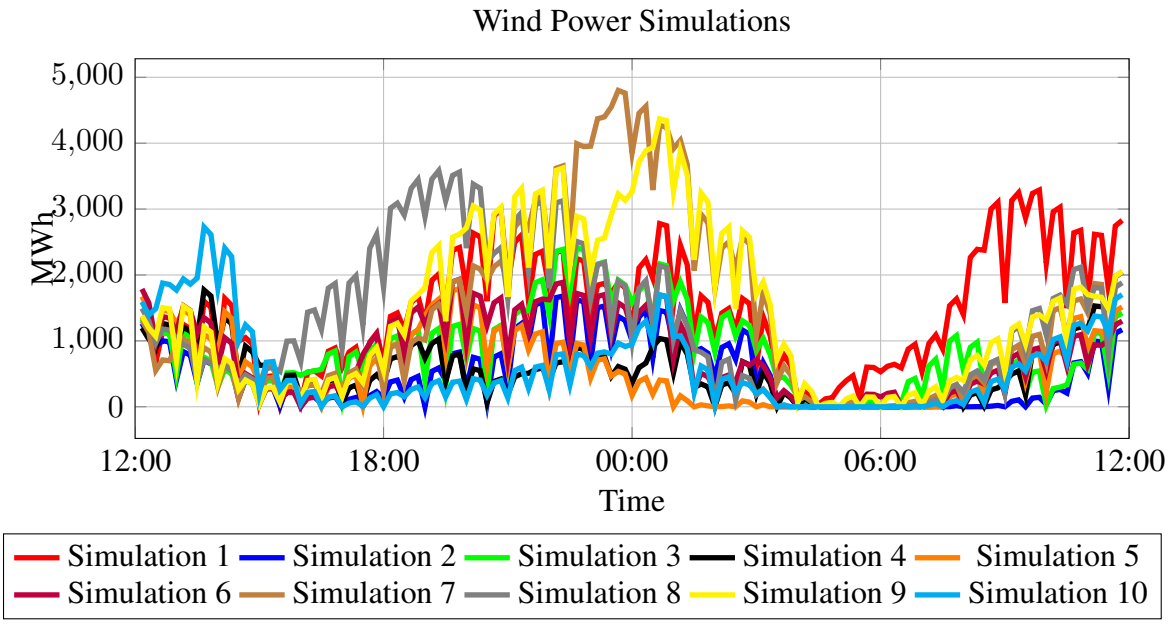

Figure 2 **Simulated daily realizations of wind power for a given wind farm over 24 hour time horizon.**

2. Deterministic Experiments

In this set of experiments, we consider only a single realization of the random process $W_t, t = 1, \ldots, T$, and therefore $|\hat{\Omega}_t| = 1, t = 1, \ldots, T$. In this case, the given problem can be solved to optimality as one large linear programming problem. Thus, we can compare the objective values and solutions that we obtain using the stochastic methods described above to the optimal ones. In the experiments below, we consider three possible cases with $|R_t^x| = 5, 25, 50$.

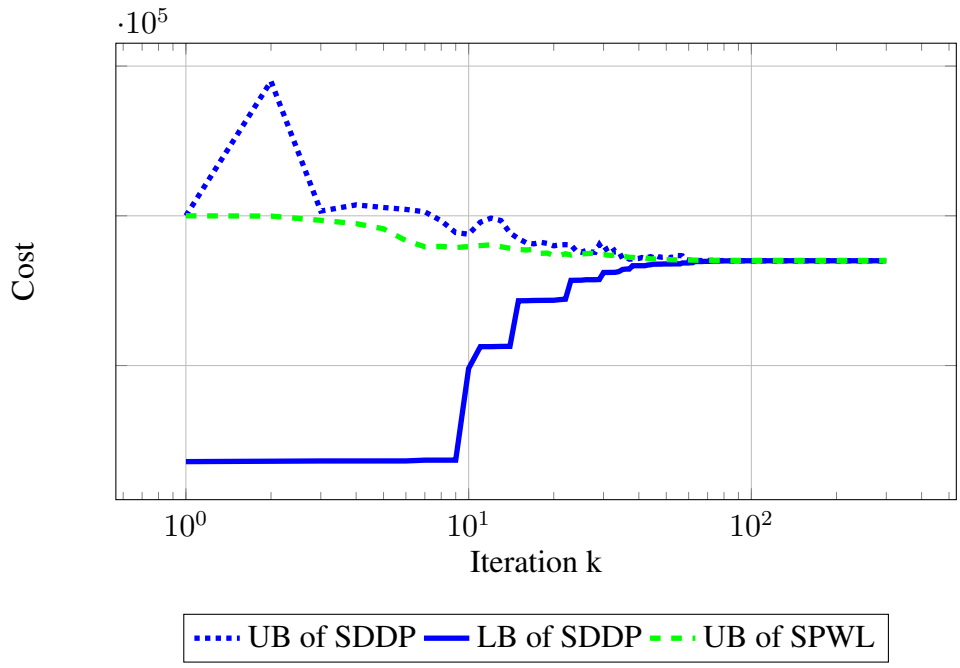

Figure 3 Comparison of multistage stochastic optimization methods for $|R| = 5$.

In this instance, we can see that both methods quickly converge to (near) optimal policies. Still, that does not imply that their solutions are the same, as can be seen from the plots below. Please note that the given plots also include the exact optimal solution (in black).

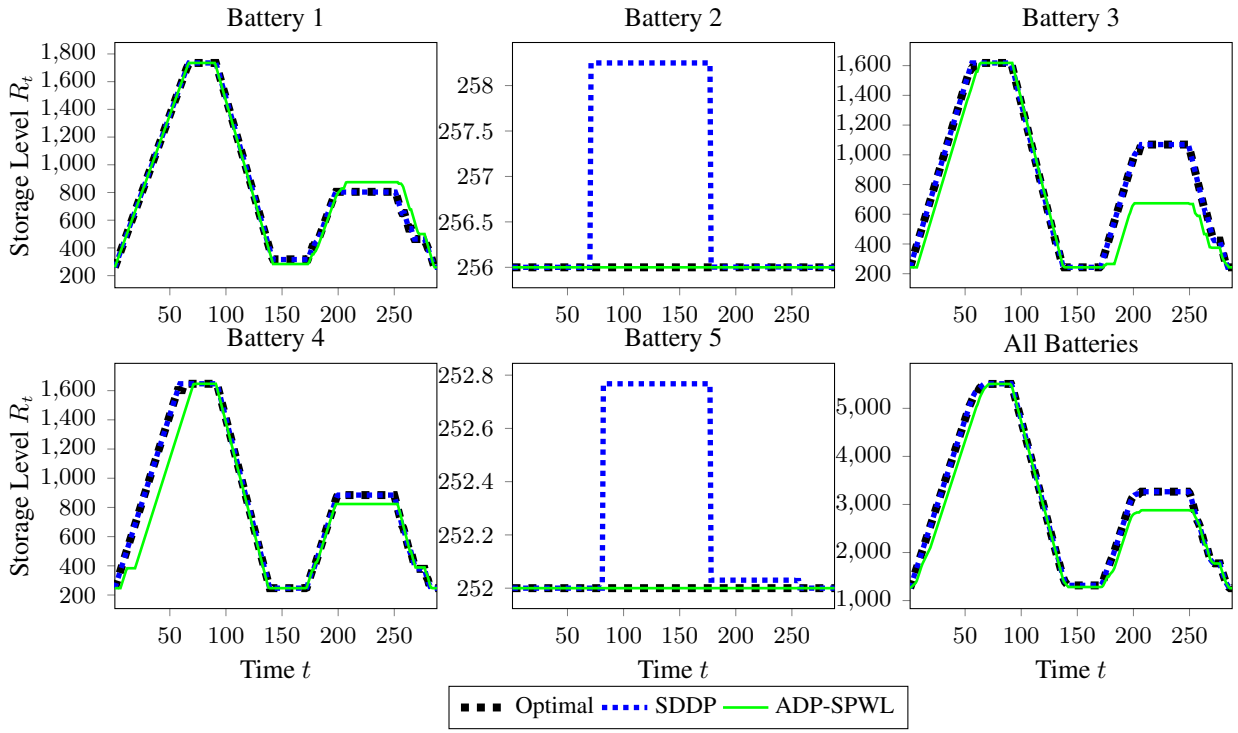

Figure 4 **Numerical comparison of battery storage levels for $|R| = 5$ at iteration 100.**

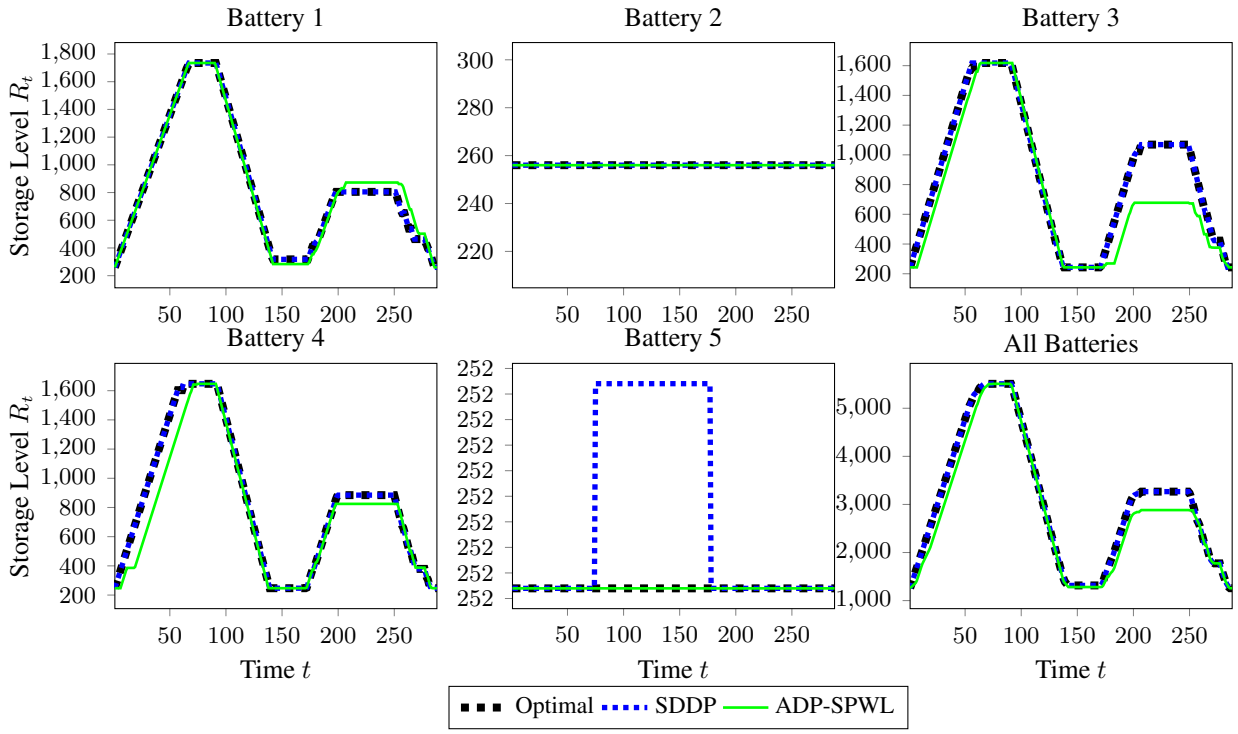

Figure 5 **Numerical comparison of battery storage levels for $|R| = 5$ at iteration 200.**

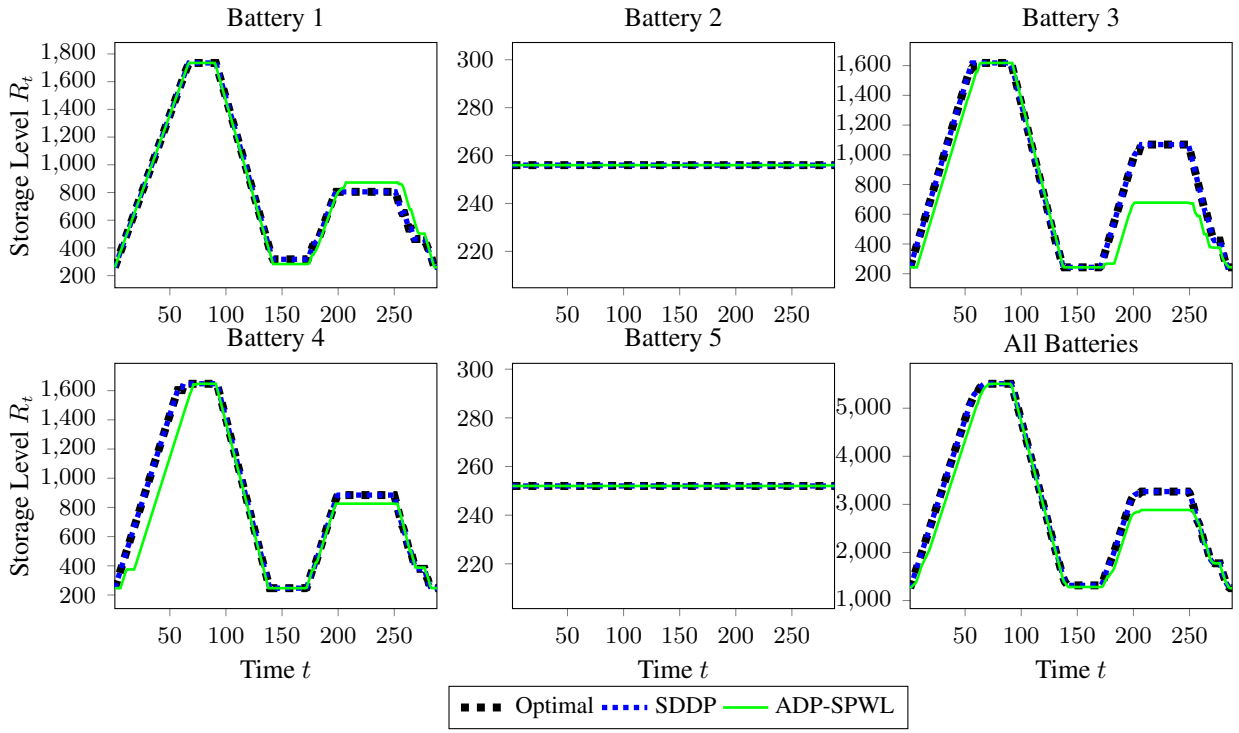

Figure 6 **Numerical comparison of battery storage levels for $|R| = 5$ at iteration 300.**

We can see that for the case of $|R_t^x| = 5$, the Benders value functions generate an exact solution to the problem at iteration 100. On the other hand, the separable value function approximations produce a very close, yet not exact, match of the optimal solution. We attribute this to the fact that the functional form of the separable approximations limits their ability to deal with inter-dimensional dependencies.

Next, we perform similar tests with 25 and 50 batteries. For both of those, we consider plots of the objective values, and the corresponding battery policies.

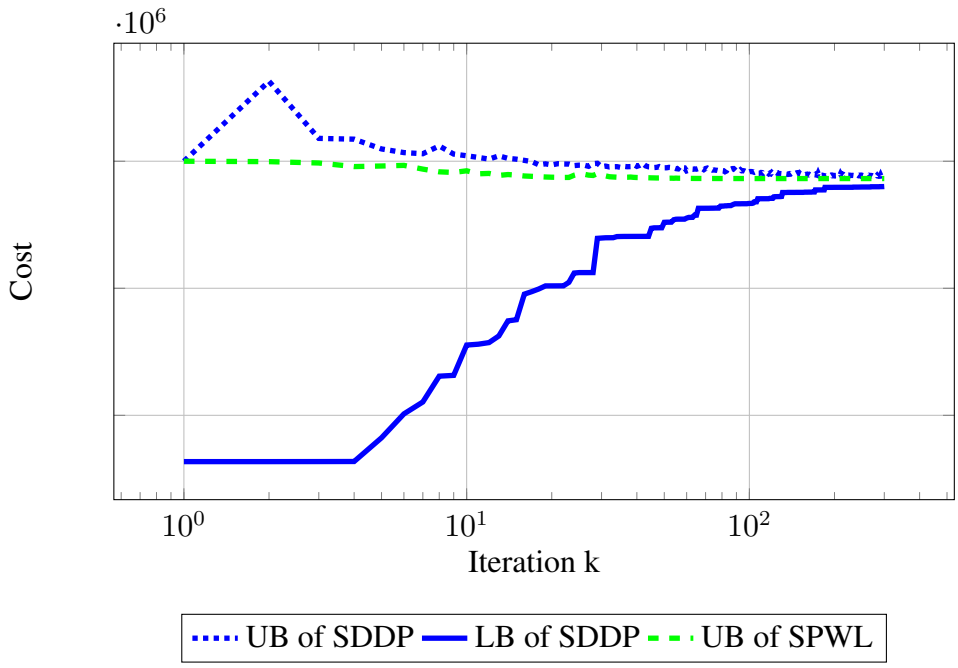

Figure 7 **Comparison of multistage stochastic optimization methods for $|R| = 25$.**

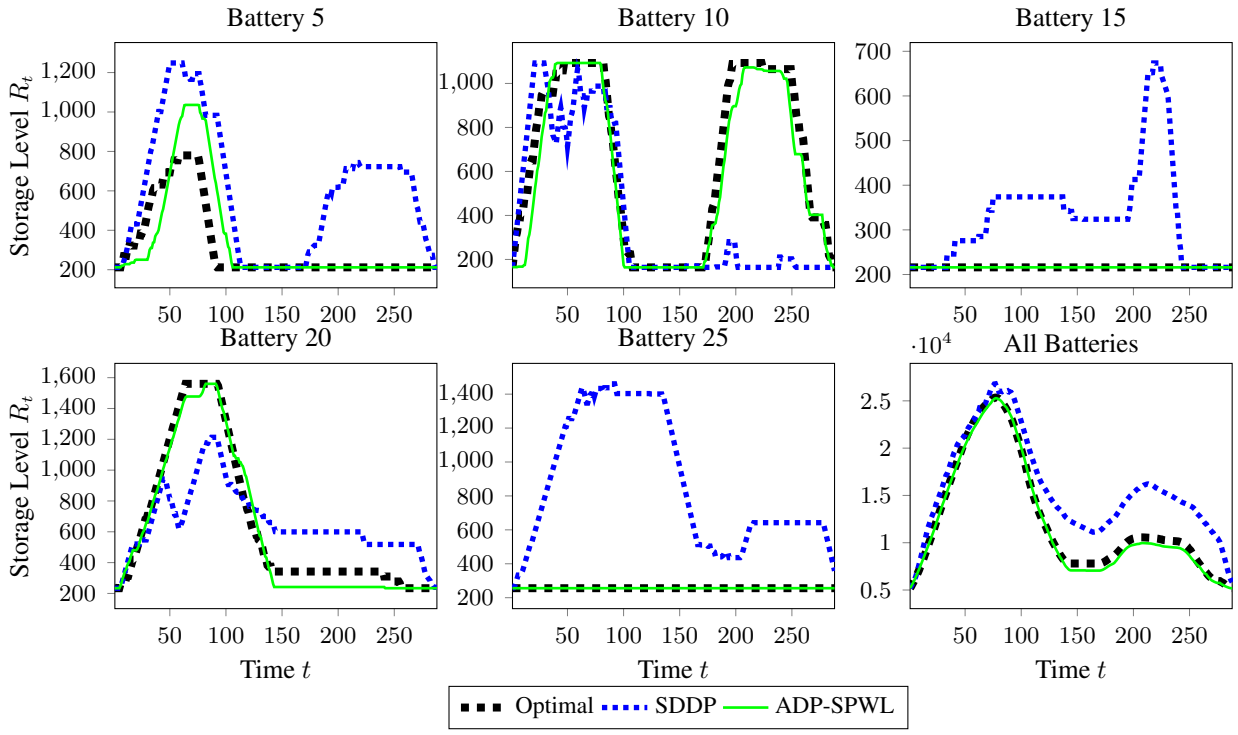

Figure 8 **Numerical comparison of battery storage levels for $|R| = 25$ at iteration 100.**

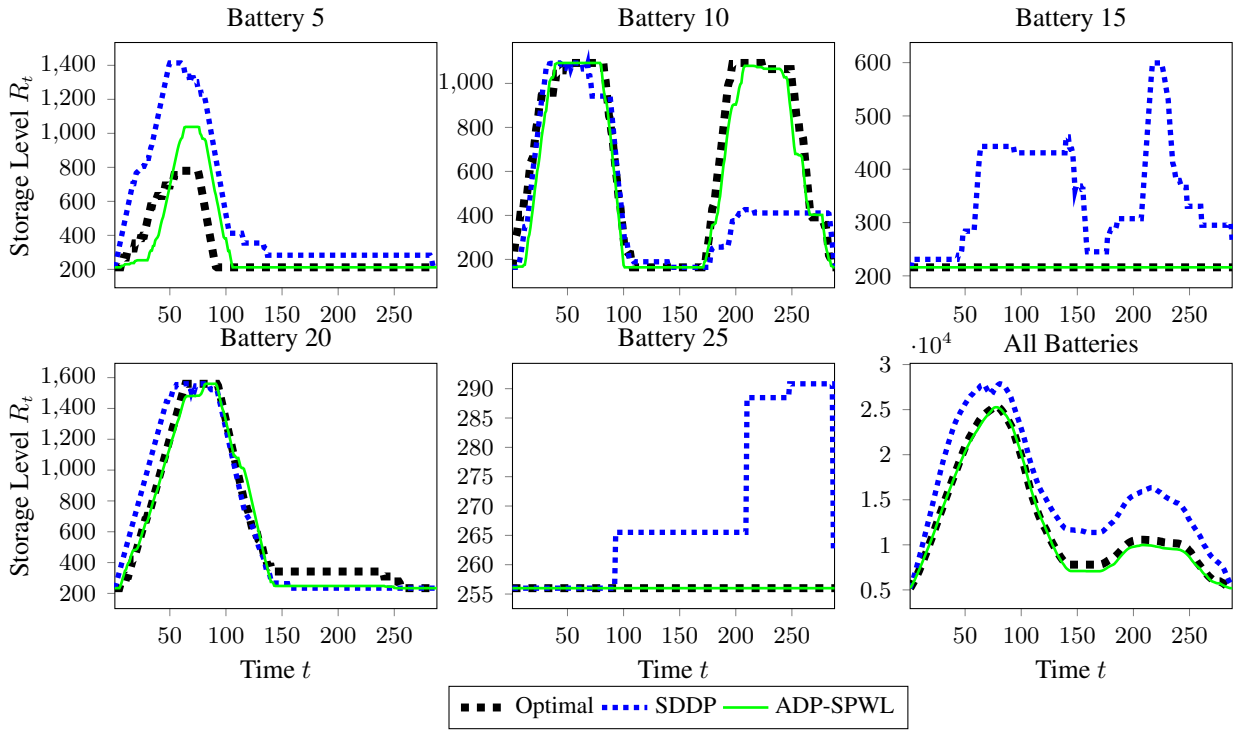

Figure 9 **Numerical comparison of battery storage levels for $|R| = 25$ at iteration 200.**

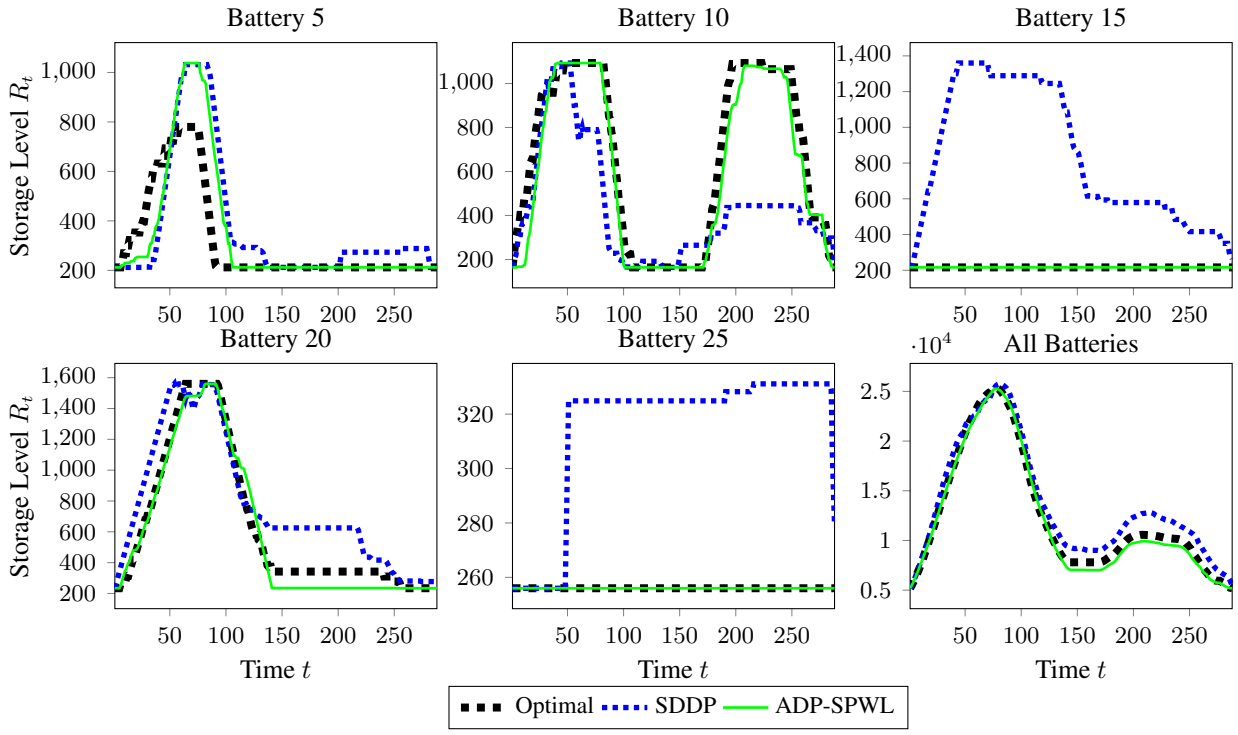

Figure 10 **Numerical comparison of battery storage levels for $|R| = 25$ at iteration 300.**

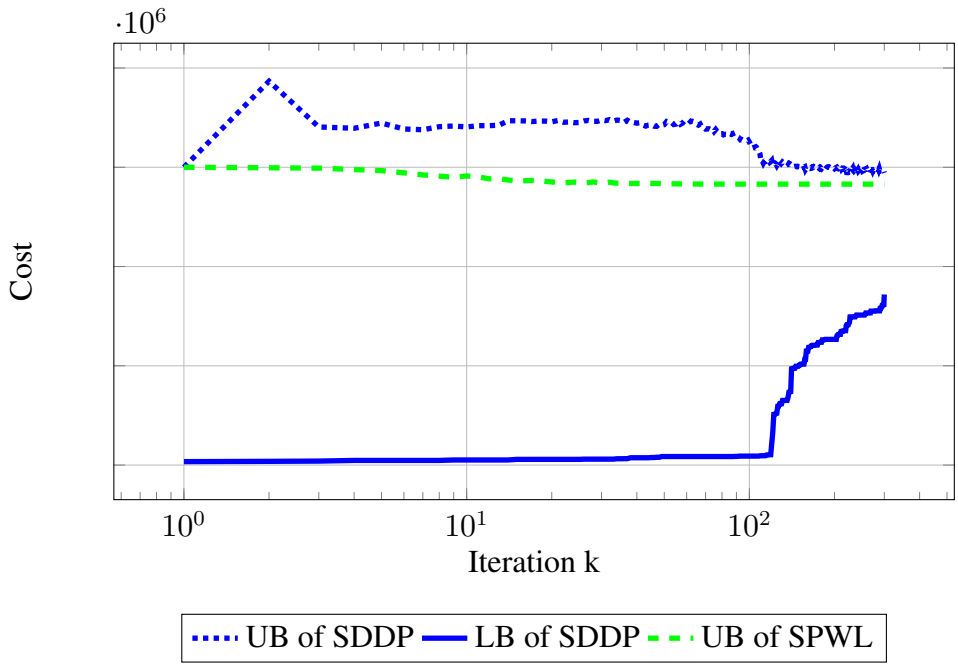

Figure 11 **Comparison of multistage stochastic optimization methods for $|R| = 50$.**

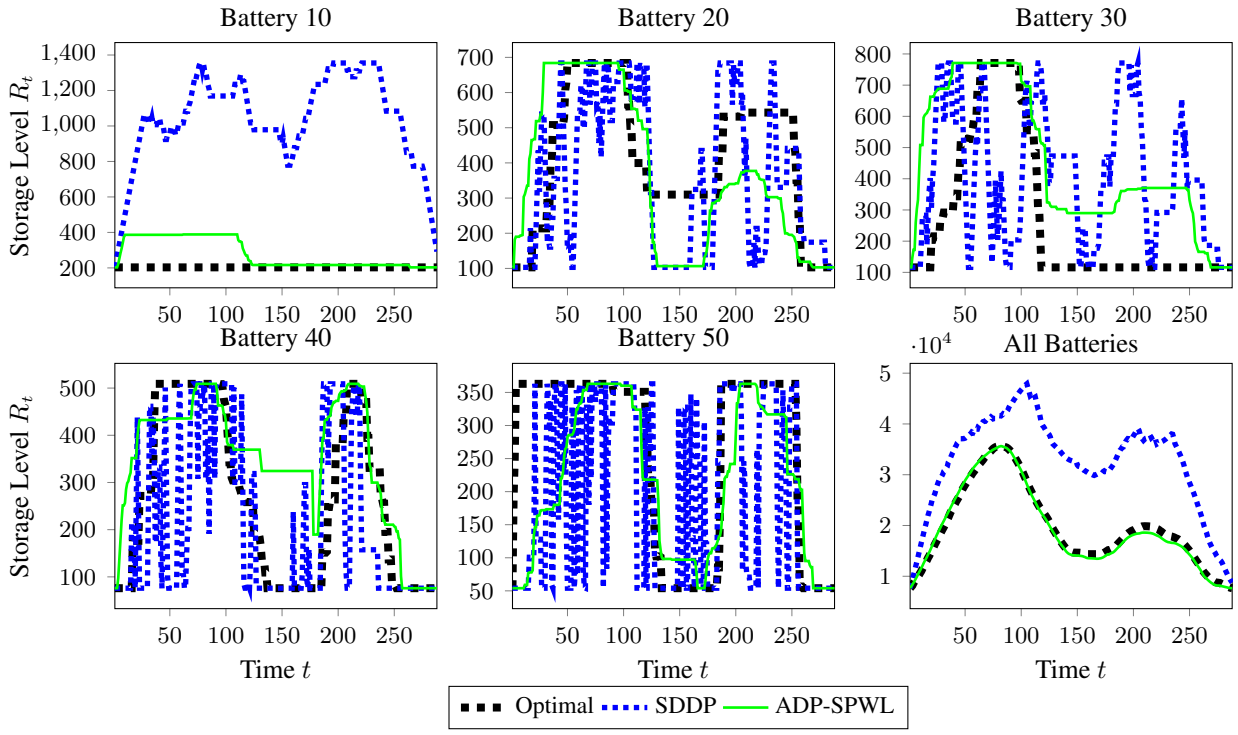

Figure 12 **Numerical comparison of battery storage levels for $|R| = 50$ at iteration 100.**

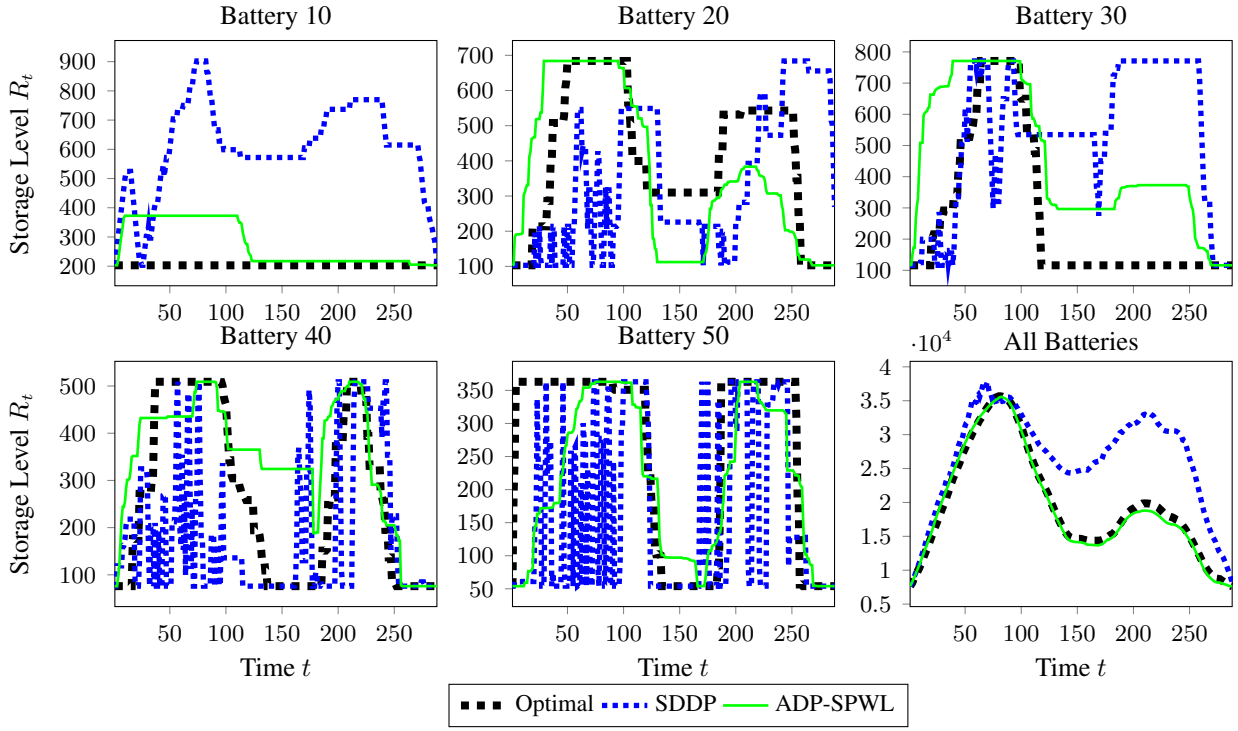

Figure 13 **Numerical comparison of battery storage levels for $|R| = 50$ at iteration 200.**

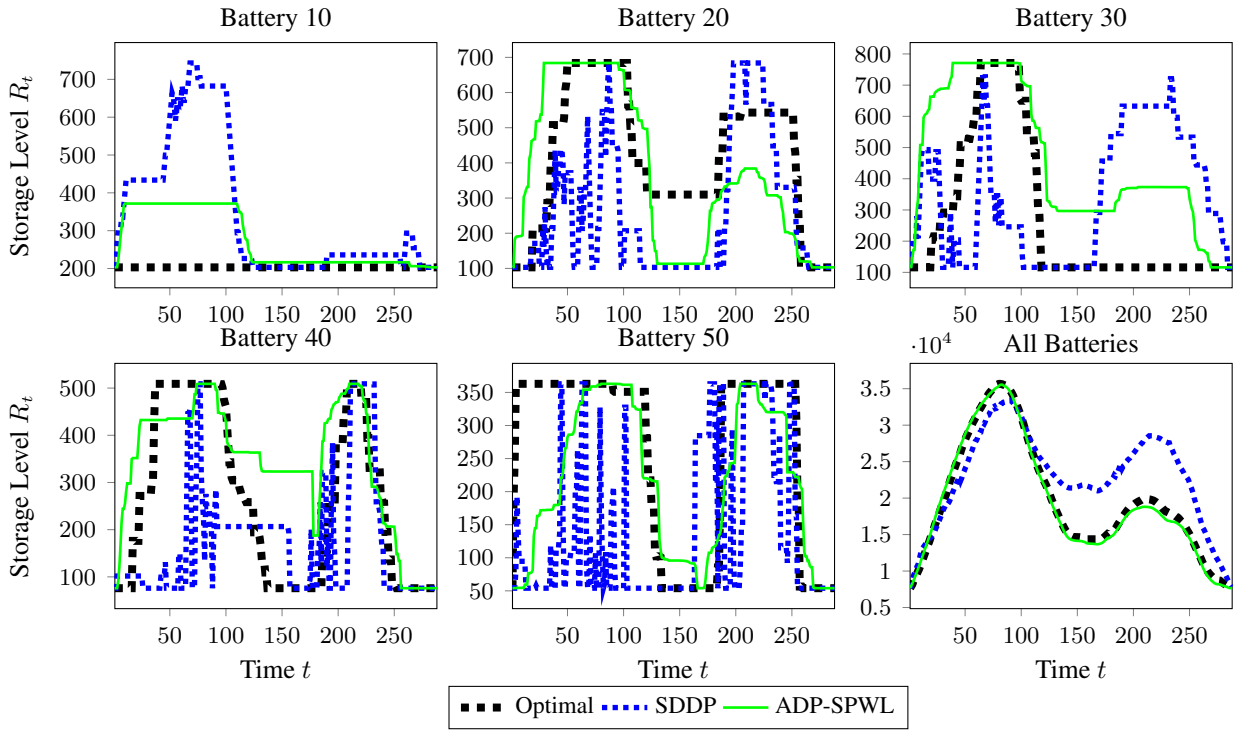

Figure 14 **Numerical comparison of battery storage levels for $|R| = 50$ at iteration 300.**

We can see that as the dimension of the post-decision value functions increases, the separable value function approximations produce policies of lower cost compared to their Benders counterparts. In addition, when we examine the resulting solutions, we see that the solutions obtained via separable value function approximations are closer to the optimal solutions. Moreover, the total amount of energy stored in all batteries by the policy with separable value functions is a very close match to the optimal.

Finally, we should mention that the SDDP solution performed surprisingly well in the case with 25 batteries, as it produced a very small optimality gap. This is unexpected, since SDDP is typically used for smaller dimensional formulations. We attribute this to the use of regularization (for more details on regularization of SDDP, please see the references in the main text). Furthermore, the total amount of energy stored by the Benders policy in Figure 10 matches very closely with the optimal solution.