Bilevel Optimization with a Piecewise-Linear Follower

Henri Lefebvre[1] Oleg Prokopyev[2] Martin Schmidt[3] Simon Stevens[3]

[1]LIRMM, University of Montpellier, CNRS [2]University of Zurich [3]Trier University

UNIVERSITÄT TRIER

Mixed-Integer Linear Bilevel Problem with a PWL Follower

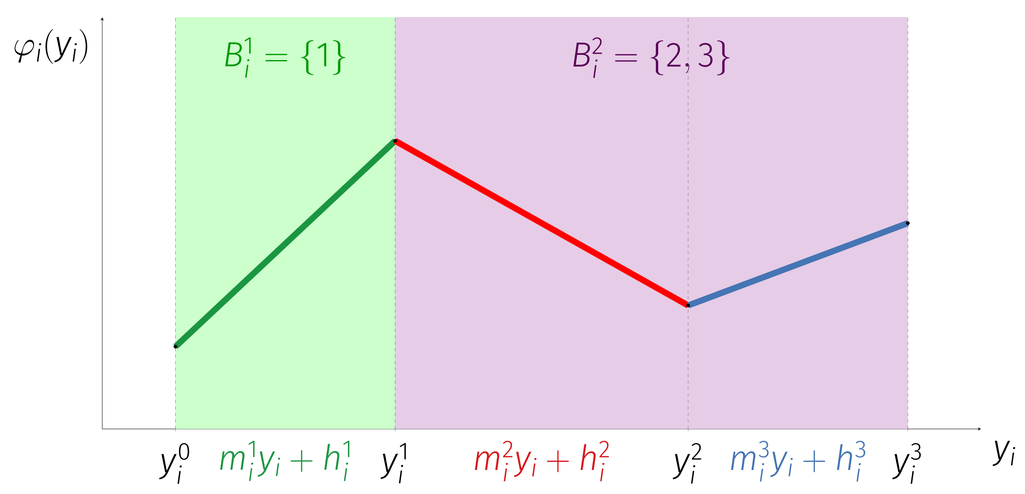
Motivation: piecewise-linear functions can model economies of scale, incremental discounts, or S-shaped returns, and approximate general *nonconvex* separable nonlinear functions. A follower with such costs turns the bilevel problem into one with a *nonconvex* lower level.

$$\begin{aligned} \min_{x\in X,\, y} \quad & c^\top x + d^\top y \\ \text{s.t.} \quad & Ax + By \geq a, \\ & y \in \arg\min_{\bar{y}\in Y(x)} \left\{ f^\top \bar{y} + \sum_{i=1}^{n_y} \varphi_i(\bar{y}_i) \right\}, \end{aligned}$$

with $Y(x) := \{y \in \mathbb{R}^{n_y} : Dy \geq b - Cx,\ y^- \leq y \leq y^+\}$ and where every φ_i is univariate, continuous and piecewise linear (PWL), but need *not* be convex and is defined by

$$\varphi_i(y_i) = m_i^j y_i + h_i^j \quad \text{on } [y_i^{j-1}, y_i^j]$$

for every piece $j \in [p_i] := \{1, \ldots, p_i\}$.



Exact Multi-Follower Reformulation

Let $\mathcal{K} := \prod_i [p_i]$ be the set of all possible piece assignments. Then, each $\boldsymbol{k} = (k_1, \ldots, k_{n_y}) \in \mathcal{K}$ fixes one affine piece in every coordinate. Then, the **follower LP** for this fixed assignment reads

$$\begin{aligned} \min \quad & g_{\boldsymbol{k}}(s, y) := f^\top y + \sum_{i=1}^{n_y} (m_i^{k_i} y_i + h_i^{k_i}) + \kappa s \\ \text{s.t.} \quad & (s, y) \in \Lambda(x, \boldsymbol{k}) := \left\{ (s, y) : Dy + 1s \geq b - Cx,\ y_i^{k_i - 1} \leq y_i \leq y_i^{k_i},\ s \geq 0 \right\}. \end{aligned}$$

The slack s keeps every follower LP feasible; an exact penalty κ makes $s > 0$ worse than the follower's optimum. With sufficiently large big-Ms, we obtain the exact reformulation

$$\begin{aligned} \min_{x,y,z,\alpha} \quad & c^\top x + d^\top y \\ \text{s.t.} \quad & x \in X, \quad Ax + By \geq a, \\ & z \geq g_{\boldsymbol{k}}(s^{\boldsymbol{k}}, y^{\boldsymbol{k}}) - M_z(1 - \alpha_{\boldsymbol{k}}) && \forall \boldsymbol{k} \in \mathcal{K}, \\ & z \leq g_{\boldsymbol{k}}(s^{\boldsymbol{k}}, y^{\boldsymbol{k}}) && \forall \boldsymbol{k} \in \mathcal{K}, \\ & y = \sum_{\boldsymbol{k}\in\mathcal{K}} \alpha_{\boldsymbol{k}} y^{\boldsymbol{k}}, \qquad \sum_{\boldsymbol{k}\in\mathcal{K}} \alpha_{\boldsymbol{k}} = 1, \\ & s^{\boldsymbol{k}} \leq M_s(1 - \alpha_{\boldsymbol{k}}) && \forall \boldsymbol{k} \in \mathcal{K}, \\ & \alpha_{\boldsymbol{k}} \in \{0, 1\} && \forall \boldsymbol{k} \in \mathcal{K}, \\ & (s^{\boldsymbol{k}}, y^{\boldsymbol{k}}) \in \arg\min_{(\bar{s},\bar{y})\in\Lambda(x,\boldsymbol{k})} g_{\boldsymbol{k}}(\bar{s}, \bar{y}) && \forall \boldsymbol{k} \in \mathcal{K}. \end{aligned}$$

- $\alpha_{\boldsymbol{k}}$ picks the active assignment; $z \leq g_{\boldsymbol{k}}$ for all $\boldsymbol{k}$ and $z \geq g_{\boldsymbol{k}}$ for the selected $\boldsymbol{k}$ make its optimum globally optimal for the follower.
- **Single-level reformulation:** replace every follower LP by strong duality (separate or aggregated) or KKT conditions (SOS1 or binary McCormick).

Convex-Block Aggregation

The size of the reformulation can be reduced by merging adjacent pieces. A **convex block** $B_i \subseteq [p_i]$ is a maximal set of consecutive pieces with nondecreasing slopes, i.e.,

$$m_i^j \leq m_i^{j+1} \quad \text{for all } j, j+1 \in B_i,$$

see the figure at the top. On the interval of each block, φ_i is convex and can thus be modeled by one epigraph variable t_i. Let $\mathcal{B}_i$ collect the convex blocks of φ_i.

Affine pieces	→	Convex blocks
$\mathcal{K} = [p_1] \times \cdots \times [p_{n_y}]$		$\mathcal{B} = \mathcal{B}_1 \times \cdots \times \mathcal{B}_{n_y}$
$\prod_i p_i$ follower LPs		$\prod_i \lvert\mathcal{B}_i\rvert$ follower LPs

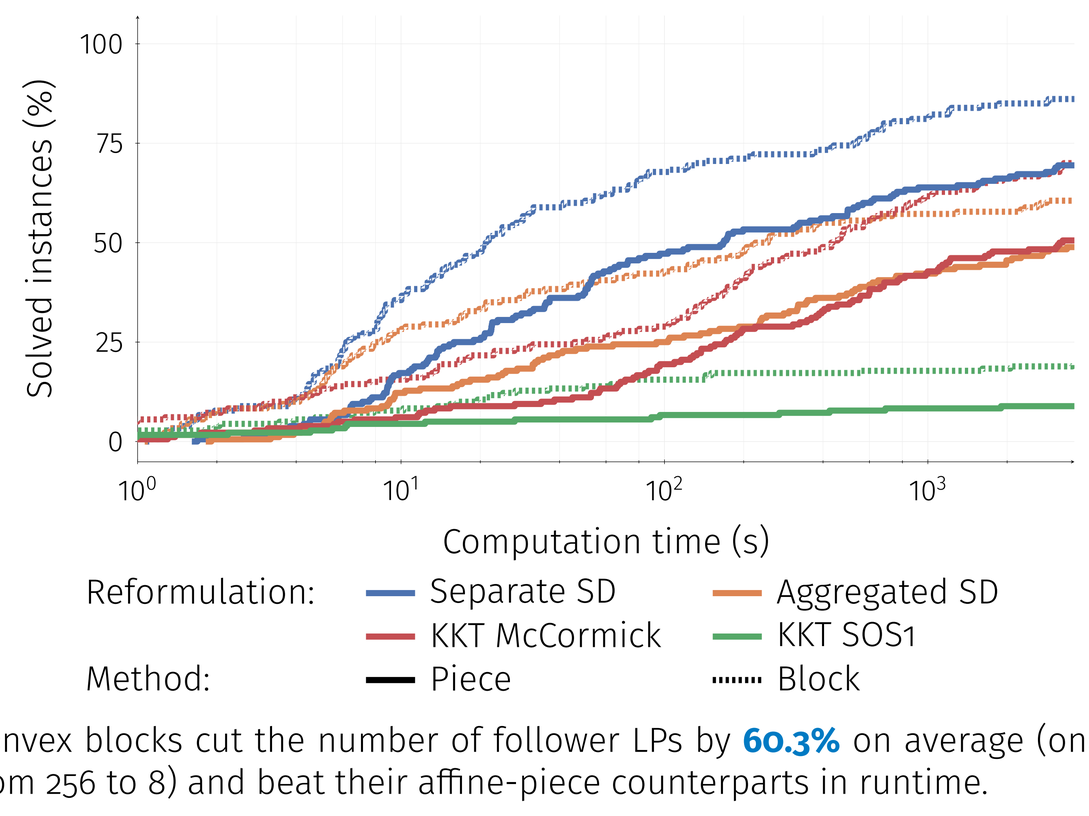
Each $\boldsymbol{B} = (B_1, \ldots, B_{n_y}) \in \mathcal{B}$ fixes one convex block in every coordinate, and we can equivalently rewrite the reformulation for all $\boldsymbol{B} \in \mathcal{B}$ instead of $\boldsymbol{k} \in \mathcal{K}$.

Complexity

General case	Σ_2^p**-hard**, even with only two pieces per PWL term (via DeRi bilevel knapsack [1]).
Convex PWL terms	**NP-complete**; an epigraph formulation gives a polynomial-size MILP–LP bilevel model.
Fixed dimension	With continuous leader variables, fixed n_y, and fixed piece counts p_i, the problem is **polynomial-time solvable**.

Numerical Results

Setup: 180 BOBILib [2] knapsack-interdiction instances, PWL follower, 1 hour.



Convex blocks cut the number of follower LPs by **60.3%** on average (once from 256 to 8) and beat their affine-piece counterparts in runtime.

Contributions and Outlook

- We classify the **complexity**: from Σ_2^p-hard in general to polynomial-time solvable cases in fixed dimensions.
- We introduce an **exact reformulation** for a nonconvex PWL follower: one follower LP per piece assignment.
- We reduce the size of the reformulation and improve the computation times by using **convex-block aggregation**.
- **Outlook:** an **iterative method** that starts from simple overestimators and reveals convex blocks only where needed.

[1] A. Caprara, M. Carvalho, A. Lodi, G. J. Woeginger. A complexity and approximability study of the bilevel knapsack problem. *IPCO*, 2013.

[2] J. Thürauf, T. Kleinert, I. Ljubić, T. Ralphs, M. Schmidt. BOBILib: Bilevel Optimization (Benchmark) Instance Library. *Math. Prog. Comp.*, 2026.

Numerical Study