Efficient LP warmstart for linear modifications of the constraint matrix

Guillaume Derval, Bardhyl Miftari, Damien Ernst, Quentin Louveaux

Montefiore Institute, University of Liège, Liège, Belgium

gderval@uliege.be, bmiftari@uliege.be, dernst@uliege.be, q.louveaux@uliege.be

Abstract

We consider the problem of computing the optimal solution and objective value of a linear program whose constraint matrix changes linearly with a parameter, for a predefined set of values of this parameter. Based on the information given by a precomputed basis, we present three efficient LP warm-starting algorithms, based on the eigenvalue decomposition, the Schur decomposition and the Hessenberg decomposition, to evaluate the optimal solution and optimal objective value of these problems. The three algorithms have an overall complexity of $O(pm^2 + pmn)$ where m (resp. n) is the number of constraints (resp. variables) of the original problem and p the number of parameter values after an initial preprocessing step. We also provide theorems related to the optimality conditions to verify when a basis is still optimal. When the perturbation of the basis matrix has a low rank ρ, as is often the case in practice, we then reduce all the computations to a $\rho \times \rho$ matrix built with ρ solves with the factorization of the basis maintained by the simplex solver; with it, we compute, in each direction, an interval on which the basis is guaranteed to remain optimal, together with a bound on the objective. In practice, these intervals are close to the exact ones, and they remain useful for degenerate bases. We also propose a matrix-free method based on shifted Krylov subspaces. Numerical experiments on 398 instances built from real models, with up to 349 120 constraints, assess the accuracy, the cost and the range of applicability of all the methods. On the interval on which a basis remains optimal, once this interval is known, the low-rank and Krylov methods compute the optimal objective value 117× to 304× faster than a warm start of the solver, in median.

Keywords: Linear Programming, Warmstart, Decomposition, Parametric Programming

Funding: This work was partially funded by the SCK-CEN SMR Chair Program.

1 Introduction

Many real-life linear programs do not have a one-value fit for some of their coefficients. Typically, several values for these coefficients may be valid depending on the setting studied and hypotheses made. Sometimes, these programs evolve over time, which can also lead to changes in the coefficients. Assessing the behavior of both the optimal objective function and the optimal solution in relation to these changes is particularly important. This issue has been addressed by the fields of sensitivity analysis and parametric programming.

An efficient method for dealing with varying coefficients, particularly when they are in the objective function or in the right-hand side, is warmstarting. Indeed, in these cases, either primal or dual feasibility is preserved, which allows us to run a few iterations of either the primal or dual simplex. However, when dealing with varying constraint matrix coefficients, the basis matrix changes, which imposes a heavier computational burden. In this paper, we propose warm-starting algorithms to evaluate the optimal objective function and optimal solution in an exact manner for LPs whose constraint coefficients vary linearly.

We study problems of the type

$$\begin{aligned} \mathcal{P}(\lambda) \equiv \min\ & c^\top x \\ \text{s.t. } & Ax + \lambda Dx = b \\ & x \geq 0 \end{aligned} \tag{1}$$

where λ controls the linear change in coefficients and belongs to a finite discrete set Λ, and the uncertainty matrix D gives the direction of this change. More precisely, we study how to use the information given by a precomputed optimal basis at a given λ to warmstart the computation of optimal solutions for other values of λ.

Given a basis B, let x_B be the basic variables and x_N the nonbasic variables:

$$x_B(\lambda) = (A_B + \lambda D_B)^{-1} b \quad x_N(\lambda) = 0, \tag{2}$$

where, for a matrix M, M_B and M_N denote its basic and nonbasic columns. Like x_B and x_N, they are indexed by the variables $i \in B$ and $j \in N$, with unit vectors e_i and e_j. The basis thus forms a potential solution to $\mathcal{P}(\lambda)$, at least if $A_B + \lambda D_B$ is invertible. A basis is said to be feasible for a given λ if the matrix is indeed invertible and $x_B(\lambda) \geq 0$, and optimal if it provides an optimal solution. We denote by $o^*(\lambda)$ the optimal objective value of $\mathcal{P}(\lambda)$ and, for a basis B such that $A_B + \lambda D_B$ is invertible, by

$$o_B(\lambda) := c_B^\top (A_B + \lambda D_B)^{-1} b = c_B^\top x_B(\lambda) \tag{3}$$

the objective value of its basic solution, whether this solution is feasible or not. If B is optimal for $\mathcal{P}(\lambda)$, then $o^*(\lambda) = o_B(\lambda)$. The dual solution associated with B is $y(\lambda)$, given by $y(\lambda)^\top = c_B^\top (A_B + \lambda D_B)^{-1}$.

The core of the paper revolves around finding methods to reduce the cost of solving the system $(A_B + \lambda D_B) x_B(\lambda) = b$ repeatedly for various λ, in order to compute $x_B(\lambda)$ and quickly check the optimality of the basis for different values of λ.

1.1 Related work

As mentioned above, most articles in the literature tackle the sensitivity analysis of the objective function c or the right-hand side b (Jansen et al. 1997, Gal and Nedoma 1972, Gal 1995, Saaty and Gass 1954, Gass and Saaty 1955). Lee et al. (2019) solve large batches of linear programs that differ only in their right-hand side, by grouping the values that share an optimal basis; in this case, the basis matrix does not change. Only a few papers deal with the problem of assessing the behavior of the optimum value and the optimal solution for a varying constraint matrix. Sherman and Morrison (1949, 1950) update the inverse of a matrix when the entries of one column or one row change, and when one entry changes. Bartlett (1951) extends the formula to any rank-one change $D = uv^\top$. For such a change, this formula turns the conditions for B to remain optimal into linear inequalities in λ (Gal 1995). Greenberg (1999) recalls this result and extends it to the optimal partition of an interior solution. Woodbury (1950) extends the formula of Sherman and Morrison to changes of any rank; Section 3 uses it to reduce all the computations to a small matrix when the perturbation has a low rank. More recently, Derval et al. (2026) provide methods to compute upper and lower bounds on the objective function for our class of problems. Freund (1985) gives the derivatives, and the Taylor series, of the optimal value and of the optimal basic solutions with respect to λ at a nondegenerate optimal basis, and discusses the degenerate case; Section 4 uses the first-order terms of these expansions to certify an interval of optimality, also for degenerate bases.

The algorithm most related to our contribution is presented in Zuidwijk (2005), where Zuidwijk provides an exact algorithm to calculate $o_B(\lambda) = c_B^\top (A_B + \lambda D_B)^{-1} b$. By using realization theory (Bart et al. 1979, Chapter 3), Zuidwijk derives the following formula:

$$o_B(\lambda) = \frac{1}{\lambda} \Big(\prod_{j=1}^{m} \frac{1 + \lambda \beta_j}{1 + \lambda \alpha_j} - 1 \Big) \tag{4}$$

where α_j and β_j are the eigenvalues of $A_B^{-1}D_B$ and $A_B^{-1}(D_B + bc_B^\top)$, respectively. To prove the correctness of this algorithm, the method checks three conditions. It computes the eigenvalues of one matrix for the existence condition, of m matrices for the feasibility condition and of $2(n-m)$ matrices for the optimality condition. This preprocessing step costs $O(nm^\omega)$, assuming matrix-matrix multiplication is $O(m^\omega)$.

Most operations on matrices (matrix-matrix multiplication, inversion, system solve, QR decomposition, eigenvalue problems, ...) can be done in $O(n^{\omega+\eta})$ for any $\eta > 0$ (Demmel et al. 2007); for simplicity, we write $O(n^\omega)$. In practical applications, ω is generally 3, but there exist galactic algorithms which can reduce ω.

1.2 Our contribution

We provide algorithms to compute the exact value of an optimal solution $x_B(\lambda)$ and the optimum value $o^*(\lambda)$ of problem $\mathcal{P}(\lambda)$ around an optimal basis B for a given set of values $\lambda \in \Lambda$. The algorithms are based on three different reformulations of the term $(A_B + \lambda D_B)^{-1}$, which use an eigendecomposition, a Schur decomposition or a Hessenberg decomposition of $A_B^{-1}D_B$; the latter always exists, is computed without eigenvalues in real arithmetic, and its reduction is cheaper than the Schur decomposition. These reformulations limit the computational complexity of finding $x_B(\lambda)$ and $o_B(\lambda)$ for new λ. More precisely, they avoid the need to invert large matrices or solve large systems of unknown structure for every $\lambda \in \Lambda$. The three algorithms have different trade-offs, but all share a global runtime to compute $p = |\Lambda|$ different points $(x_B(\lambda),\ o_B(\lambda))$ of $O(m^\omega + pm^2)$, and of $O(m^\omega + pm^2 + pmn)$ with the optimality test that gives $o^*(\lambda)$, with m (resp. n) the number of constraints (resp. variables) of $\mathcal{P}(\lambda)$. In addition, we present how each of these algorithms ensures that B is feasible and optimal for each $\lambda \in \Lambda$.

We then make these methods applicable to large problems:

- Section 3 shows that, when $A_B^{-1}D_B$ has a low rank ρ, as is often the case in practice, all the quantities above can be computed from a $\rho \times \rho$ matrix and from two $m \times \rho$ factors, built with ρ solves with A_B and ρ with $A_B^\top$. These solves reuse the LU factorization of the simplex solver.
- With these low-rank factors, Section 4 provides, in each direction, an interval around a given λ on which the basis is guaranteed to remain optimal, together with a bound on the objective function on this interval. In the experiments, these intervals are close to the exact ones, and they remain useful when the basis is degenerate, which is common in practice.
- Section 5 presents a matrix-free alternative, based on shifted Krylov subspaces, that needs one solve with the factorized basis matrix per iteration and terminates after at most $\rho + 1$ iterations.

Finally, Section 6 evaluates all the methods on 398 instances built from real models, showing that both the low-rank and the Krylov-based techniques are accurate and that, as long as the basis remains optimal, they compute and certify $o^*(\lambda)$ 3.7× to 5.5× faster than a solver warm start, in median. Once the interval on which the basis is optimal is known, they compute $o^*(\lambda)$ 117× to 304× faster.

2 The matrix $E = A_B^{-1}D_B$, its decompositions and optimality of the basis

This section shows that one decomposition of the matrix $E := A_B^{-1}D_B$, computed once, reduces the cost of computing $x_B(\lambda)$ for a new value of λ from $O(m^\omega)$ to $O(m^2)$ and how it allows us to efficiently check the three conditions for the basis B to remain optimal.

2.1 From $A_B + \lambda D_B$ to decomposing E

We take $\lambda = 0$ as the nominal value and assume that $\mathcal{P}(0)$ has been solved: the solver provides an optimal basis B and maintains an LU decomposition of A_B. When $A_B + \lambda D_B$ is invertible,

$$x_B(\lambda) = (A_B + \lambda D_B)^{-1}b = (A_B(I_m + \lambda E))^{-1}b = (I_m + \lambda E)^{-1}x_B(0), \qquad x_B(0) = A_B^{-1}b. \tag{5}$$

Solving this dense system costs $O(m^\omega)$ (Demmel et al. 2007) for each value of λ. A decomposition of E avoids it.

Proposition 1. *Let $E = QTQ^{-1}$ with Q invertible. For any λ such that $I_m + \lambda E$ is invertible,*

$$x_B(\lambda) = Q(I_m + \lambda T)^{-1}Q^{-1}x_B(0), \qquad o_B(\lambda) = c_B^\top Q(I_m + \lambda T)^{-1}Q^{-1}x_B(0).$$

Proof. $I_m + \lambda E = Q(I_m + \lambda T)Q^{-1}$, and the formulas follow from (5). □

The three following choices of T are computed once, in $O(m^\omega)$. They differ by their requirements and by the cost of the systems with $I_m + \lambda T$ (Table 1):

- **Eigendecomposition**: T is the diagonal matrix of the eigenvalues of E, so that $x_B(\lambda)$ costs $O(m^2)$, and $o_B(\lambda)$ costs $O(m)$ once $c_B^\top Q$ and $Q^{-1}x_B(0)$ are precomputed. The decomposition does not exist when E is defective, which happens in practice.
- **Schur decomposition**: $E = QTQ^H$ with Q unitary and T upper triangular, so that back-substitution gives $x_B(\lambda)$ and $o_B(\lambda)$ in $O(m^2)$. The decomposition always exists, but it requires complex arithmetic.
- **Hessenberg decomposition**: $E = PHP^\top$ with P orthogonal and H upper Hessenberg, computed directly by Householder reflections (Golub and Van Loan 2013), in real arithmetic and without eigenvalues. Gaussian elimination with partial pivoting solves the systems with $I_m + \lambda H$ in $O(m^2)$, as only two rows are candidate pivots at each step and only one row is updated. Its pivots give $\det(I_m + \lambda H)$ up to the sign, so that a (numerically) zero pivot reveals that $I_m + \lambda E$ is singular.

2.2 Optimality of the basis

The methods above compute $x_B(\lambda)$. In the following, we check whether it is optimal.

Let $\nu_1, \dots, \nu_m$ be the eigenvalues of E. The eigenvalues of $I_m + \lambda E$ are the numbers $1 + \lambda\nu_k$, and $A_B + \lambda D_B = A_B(I_m + \lambda E)$ is thus invertible if and only if $\lambda \neq -1/\nu_k$ for every nonzero ν_k. If the three following conditions hold, then B is optimal:

(existence) $\lambda \neq -1/\nu_k$ for every nonzero ν_k, i.e. $A_B + \lambda D_B$ is invertible. These eigenvalues are computed once, or avoided with the pivots of the Hessenberg elimination;

(feasibility) $x_B(\lambda) \geq 0$, checked in $O(m^2)$. Then $o_B(\lambda)$ is an upper bound on $o^*(\lambda)$, even if B is not optimal;

(optimality) the reduced costs are nonnegative, which is sufficient, and also necessary when the basic solution is nondegenerate. With the decomposition, they cost $O(m^2 + mn)$ per value of λ:

$$r(\lambda) = c_N^\top - c_B^\top(A_B + \lambda D_B)^{-1}(A_N + \lambda D_N) = c_N^\top - c_B^\top Q(I_m + \lambda T)^{-1}Q^{-1}A_B^{-1}(A_N + \lambda D_N). \quad (6)$$

2.3 An initial algorithm for warm-starting from the decomposition

The decomposition is computed once. Then, for each $\lambda \in \Lambda$ ($|\Lambda| = p$), the three conditions are checked, and the objective and the solution are computed. If B is not optimal at a value of λ, $\mathcal{P}(\lambda)$ is re-solved with a warm start from B, and the procedure restarts with the new optimal basis for the remaining values of λ.

Table 1 summarizes the asymptotic complexity of all the decomposition methods presented in this paper, and compares it to two methods: the naive method of recomputing the problem in full for each λ and the method presented in Zuidwijk (2005).

As presented, these decompositions are sadly of a quite similar size to the original constraint matrix A_B, which is $m \times m$. When the matrix E has a low rank, the next section reduces the size of the decomposition and thus the cost of computing $x_B(\lambda)$ and $o_B(\lambda)$ for new values of λ.

Table 1: Complexity of various methods to compute $o^*(\lambda)$ and $x_B(\lambda)$ for all $\lambda \in \Lambda$, $|\Lambda| = p$, in big-Oh ($O(\cdot)$) asymptotic complexity.

Method	Preprocessing	Existence	Feasibility	Optimality	Objective	Solution	Total
Naïve solve	/	/	/	/	$\approx n^\omega$	$\approx n^\omega$	$\approx pn^\omega$
Basis system solving	m^ω	m	m^ω	$m^2 + mn$	m^ω	m^ω	$pm^\omega + pmn$
Zuidwijk (2005)	nm^ω	m	m^2	mn	m	N/A	$nm^\omega + pm^2 + pmn$
Eigendecomposition	m^ω	m	m^2	$m^2 + mn$	m	m^2	$m^\omega + pm^2 + pmn$
Schur	m^ω	m	m^2	$m^2 + mn$	m^2	m^2	$m^\omega + pm^2 + pmn$
Hessenberg	m^ω	m^2	m^2	$m^2 + mn$	m^2	m^2	$m^\omega + pm^2 + pmn$

In the table, n is the number of variables in the original problem and m the number of constraints, with $m < n$ in general. $O(n^\omega)$ is the complexity of multiplying two $n \times n$ matrices.

3 Exploiting the low rank of the perturbation

As formulated in Section 2, the decompositions are computed on the $m \times m$ matrix $E = A_B^{-1} D_B$; forming and decomposing it as a dense matrix requires $O(m^\omega)$ operations and $O(m^2)$ memory, which restricts these methods to problems with at most a few thousand constraints. In this section, we assume that D_B has a low rank $\rho \ll m$. This is often the case in practice, as a parameter usually affects a limited part of the model: the nonzero entries of D may form rows, columns, diagonals or combinations of them, and the rank of D_B is small as soon as they span few rows or few columns. This section shows that all the quantities of Section 2 can be obtained from a small $\rho \times \rho$ matrix (with $\rho \ll m$) and two $m \times \rho$ factors, and that building them only requires ρ solves with the basis matrix A_B and ρ with $A_B^\top$. The simplex method already maintains A_B in factored form, so these solves are computed cheaply.

3.1 Low-rank reformulation

The matrix D_B is made of the columns of D of the basic variables: its rows are those of D, and its columns are indexed by the basic variables. Let $\mathcal{J} \subseteq B$ be the set of the basic variables whose column in D is nonzero, $n_D = |\mathcal{J}|$, and let $\mathcal{I}$ be the set of the rows of D_B that contain a nonzero entry (a subset of those of D), $m_D = |\mathcal{I}|$. The matrix D_B is zero outside of the rows $\mathcal{I}$ and of the columns $\mathcal{J}$, so that it admits factorizations

$$D_B = UV^\top, \qquad U, V \in \mathbb{R}^{m \times \rho}, \qquad \rho \geq \operatorname{rank}(D_B). \tag{7}$$

Let $S_\mathcal{J}$ (resp. $S_\mathcal{I}$) be the matrix made of the columns $\mathcal{J}$ (resp. $\mathcal{I}$) of the identity matrix. Two factorizations are available without any computation: $U = D_B S_\mathcal{J}$, $V = S_\mathcal{J}$ (then $\rho = n_D$), and $U = S_\mathcal{I}$, $V = D_B^\top S_\mathcal{I}$ (then $\rho = m_D$), since the columns of D_B outside $\mathcal{J}$ and its rows outside $\mathcal{I}$ are zero, so that $D_B = D_B S_\mathcal{J} S_\mathcal{J}^\top = S_\mathcal{I} S_\mathcal{I}^\top D_B$. A rank-revealing factorization of the $m_D \times n_D$ nonzero block of D_B (e.g. its singular value decomposition) gives the smallest possible value, $\rho = \operatorname{rank}(D_B)$; we call such a factorization *minimal*. The experiments use the smaller of the first two, with $\rho = \min(n_D, m_D)$.

Given a factorization (7), let us define the following matrices and vectors that will be used in the subsequent analysis:

$$\begin{aligned} \hat{U} &= A_B^{-1} U \in \mathbb{R}^{m \times \rho}, & \hat{V} &= (A_B^\top)^{-1} V \in \mathbb{R}^{m \times \rho}, \\ C &= V^\top \hat{U} \in \mathbb{R}^{\rho \times \rho}, & w_c &= \hat{U}^\top c_B \in \mathbb{R}^\rho, \end{aligned} \tag{8}$$

for which $E = A_B^{-1} UV^\top = \hat{U}V^\top$ by construction.

Theorem 1. *Let $D_B = UV^\top$ be a factorization* (7)*, with the notation* (8)*. For any λ, $A_B + \lambda D_B$ is invertible if and only if $I_\rho + \lambda C$ is invertible, and in that case*

$$(A_B + \lambda D_B)^{-1} = A_B^{-1} - \lambda\, \hat{U}\, (I_\rho + \lambda C)^{-1}\, V^\top A_B^{-1}, \tag{9}$$

$$x_B(\lambda) = x_B(0) - \lambda\, \hat{U}(I_\rho + \lambda C)^{-1} V^\top x_B(0), \qquad o_B(\lambda) = o_B(0) - \lambda\, w_c^\top (I_\rho + \lambda C)^{-1} V^\top x_B(0). \tag{10}$$

Proof. Since $D_B = UV^\top = A_B \hat{U} V^\top$, $A_B + \lambda D_B = A_B(I_m + \lambda \hat{U} V^\top)$ and, by Sylvester's determinant identity, $\det(I_m + \lambda \hat{U} V^\top) = \det(I_\rho + \lambda V^\top \hat{U})$: the two matrices are simultaneously invertible. Equation (9) is the Sherman–Morrison–Woodbury formula (Woodbury 1950, Hager 1989). Multiplying (9) by b on the right gives $x_B(\lambda)$ in (10), since $A_B^{-1} b = x_B(0)$, and multiplying it by $c_B^\top$ on the left gives $o_B(\lambda)$, since $o_B(\lambda) = c_B^\top x_B(\lambda)$ and $c_B^\top \hat{U} = w_c^\top$. □

Only the $\rho \times \rho$ matrix $I_\rho + \lambda C$ depends on λ, and so any of the decompositions of Section 2 can be applied to C instead of E, with ρ instead of m: with $C = PHP^\top$ or $C = QTQ^H$, the systems with $I_\rho + \lambda C$ cost $O(\rho^2)$ per value of λ after a decomposition in $O(\rho^\omega)$, instead of $O(m^2)$ and $O(m^\omega)$, and an eigendecomposition of C, when it exists, gives the objective in $O(\rho)$.

3.2 Spectrum, existence and diagonalizability

The spectra of E and C are closely related:

Lemma 1. *The nonzero eigenvalues of $E = A_B^{-1} D_B = \hat{U} V^\top$ are the nonzero eigenvalues of $C = V^\top \hat{U}$, with the same algebraic and geometric multiplicities. In particular, the existence condition of Section 2.2 reads: $A_B + \lambda D_B$ is singular if and only if $\lambda = -1/\mu$ for a nonzero eigenvalue μ of C.*

Lemma 2. *If the factorization* (7) *is minimal, then E is diagonalizable if and only if C is nonsingular and diagonalizable.*

The proofs of these lemmas are in Appendix D. Lemma 2 explains why the eigendecomposition of E often fails to exist in practice (as we will see in Section 6): the eigenvalue 0 of E has an algebraic multiplicity of at least $m - \rho$, and it is defective as soon as the (minimal) matrix C is singular, that is, as soon as $V^\top A_B^{-1} U$ loses rank. For instance, if $C = 0$ but $\hat{U} V^\top \neq 0$, then $E^2 = \hat{U} C V^\top = 0$ and E is a nonzero nilpotent matrix. The low-rank formulation is less demanding: Theorem 1 does not require C to be nonsingular, and an eigendecomposition of C only requires C to be diagonalizable ($C = 0$, for instance, is diagonalizable).

3.3 Optimality conditions

Inserting (9) in the first expression of (6) gives, for every λ such that $I_\rho + \lambda C$ is invertible,

$$r(\lambda) = c_N^\top - \big(y(0)^\top - \lambda\, w_c^\top (I_\rho + \lambda C)^{-1} \hat{V}^\top\big)(A_N + \lambda D_N). \tag{11}$$

The vector in parentheses requires $O(\rho^2 + m\rho)$ operations per value of λ, and the product with $A_N + \lambda D_N$ is sparse, i.e. $O(\mathrm{nnz}(A_N) + \mathrm{nnz}(D_N))$ operations instead of the $O(m^2 + mn)$ of Section 2.2. The feasibility of B for a value of λ is thus checked in $O(\rho^2 + m\rho)$ and its optimality in $O(\rho^2 + m\rho + \mathrm{nnz}(A_N) + \mathrm{nnz}(D_N))$ operations.

3.4 Complexity and implementation with a simplex solver

To form $C = V^\top \hat{U}$, we need $\hat{U} = A_B^{-1} U$, which requires ρ solves with A_B. The reduced costs (11) also need $\hat{V} = (A_B^\top)^{-1} V$, which requires ρ solves with $A_B^\top$. Their right-hand sides are sparse (the columns of U are supported on the rows $\mathcal{I}$ and those of V on the rows $\mathcal{J}$). LP solvers expose these solves with A_B and $A_B^\top$, called FTRAN and BTRAN, so no factorization of A_B has to be computed. Let τ be the cost of one such solve. This cost is proportional to the number of nonzeros of the LU factors of A_B, and much smaller than m^2 in practice. Table 2 compares the low-rank variants with the methods of Section 2.

Table 2: Asymptotic complexity ($O(\cdot)$) of the low-rank variants, where $\rho = \min(n_D, m_D)$ or $\rho = \operatorname{rank}(D_B)$, τ is the cost of a solve with the factorized basis matrix and $\mathrm{nz} = \mathrm{nnz}(A_N) + \mathrm{nnz}(D_N)$.

	Preprocessing	Per value of λ		
		Objective	Solution	Optimality
Dense E (Sec. 2)	m^ω	m or m^2	m^2	$m^2 + mn$
Low-rank, eig. of C	$2\rho\tau + \rho^\omega + m\rho$	ρ	$m\rho$	$m\rho + \mathrm{nz}$
Low-rank, Schur/Hess. of C	$2\rho\tau + \rho^\omega + m\rho$	ρ^2	$\rho^2 + m\rho$	$\rho^2 + m\rho + \mathrm{nz}$

Forming $C = V^\top \hat{U}$ costs $O(\mathrm{nnz}(V)\,\rho)$ more (nothing for $V = S_{\mathcal{J}}$), and the optional rank-revealing factorization $O(n_D m_D \min(n_D, m_D))$. The eigendecomposition variant requires C to be diagonalizable.

4 Local radius of optimality

The methods above compute $o_B(\lambda)$ efficiently, and allow us to check whether the computed values are optimal, i.e. if locally $o_B(\lambda) = o^*(\lambda)$. A piecewise-linear plot of the computed values $o_B(\lambda_i)$, $\lambda_i \in \Lambda$, implicitly assumes that $o_B(\lambda)$ does not deviate much from this linear interpolation and that the basis remains optimal between two sampled points. In general, $o_B(\lambda)$ has no such property (Derval et al. 2026). In this section, we compute, with the low-rank factors of Section 3 and for a given optimal basis B at $\lambda = 0$, an interval $(-\Delta^-, \Delta^+)$ on which B is guaranteed to remain optimal; we call Δ^- and Δ^+ the *local radii of optimality*. We also compute lower and upper bounds on $o^*(\lambda)$ on this interval, which bound in particular

$$\sup_{-\Delta^- < \lambda < \Delta^+} |o^*(\lambda) - o^*(0)|.$$

4.1 First-order expansion

The functions $x_B(\lambda)$, $r(\lambda)$ and $o_B(\lambda)$ are rational functions of λ. Their denominator, $\det(I_m + \lambda E)$, has degree up to m, so their extrema on an interval are hard to locate when m is large. We therefore write each function as its first-order expansion at $\lambda = 0$, which is explicit, plus a remainder of order λ^2. Section 4.2 bounds these remainders.

For the basic solution, let $W = (I_m + \lambda E)^{-1}$, so that $x_B(\lambda) = W x_B(0)$ by (5). The identity $(I_m + \lambda E)W = W(I_m + \lambda E) = I_m$ gives $W = I_m - \lambda WE = I_m - \lambda EW$. Using the first form, then the second one, gives the exact equality

$$x_B(\lambda) = x_B(0) - \lambda W E x_B(0) = x_B(0) + \lambda \dot{x}_B - \lambda^2 EW \dot{x}_B, \qquad \dot{x}_B = -E x_B(0).$$

Its first two terms form the first-order expansion of x_B at $\lambda = 0$, with the derivative $\dot{x}_B$, and the last term is the remainder. With the low-rank notation (8), $E = \hat{U} V^\top$ and $EW = \hat{U}(I_\rho + \lambda C)^{-1} V^\top$, so the remainder only involves the $\rho \times \rho$ matrix C. The same steps apply to $r(\lambda)$ and $o_B(\lambda)$.

Lemma 3. *Let B be an optimal basis for $\mathcal{P}(0)$, with the notation* (8)*, and let*

$$\dot{x}_B = -E x_B(0), \quad G = E A_B^{-1} A_N - A_B^{-1} D_N, \qquad \dot{r} = c_B^\top G, \quad \dot{o}_B = c_B^\top \dot{x}_B,$$
$$w_x = V^\top \dot{x}_B, \qquad \theta_j = V^\top G e_j = C \hat{V}^\top A_N e_j - \hat{V}^\top D_N e_j \quad (j \in N), \quad w_c = \hat{U}^\top c_B.$$

For every λ such that $I_m + \lambda E$ is invertible, each basic variable $i \in B$ and each nonbasic variable $j \in N$,

$$(x_B(\lambda))_i = (x_B(0))_i + \lambda\,(\dot{x}_B)_i \quad - \lambda^2\,(\hat{U}^\top e_i)^\top (I_\rho + \lambda C)^{-1} w_x, \tag{12}$$

$$(r(\lambda))_j = (r(0))_j + \lambda\,(\dot{r})_j \qquad - \lambda^2\,\theta_j^\top (I_\rho + \lambda C^\top)^{-1} w_c, \tag{13}$$

$$o_B(\lambda) = o_B(0) + \lambda\,\dot{o}_B \qquad - \lambda^2\, w_c^\top (I_\rho + \lambda C)^{-1} w_x. \tag{14}$$

Table 3: Terms of the general form (15) for each function of Lemma 3.

Function	Eq.	δ	$\dot{\delta}$	K	ξ	ζ
Basic variable i	(12)	$(x_B(0))_i$	$(\dot{x}_B)_i$	C	$\hat{U}^\top e_i$	w_x
Reduced cost $j \in N$	(13)	$(r(0))_j$	$(\dot{r})_j$	$C^\top$	θ_j	w_c
Objective	(14)	$o_B(0)$	$\dot{o}_B$	C	w_c	w_x

The proof is in Appendix D. Equations (12)–(14) share the general form

$$\underbrace{\delta + \lambda\dot{\delta}}_{\text{first-order expansion}} \quad \underbrace{-\lambda^2\, \xi^\top (I_\rho + \lambda K)^{-1} \zeta}_{\text{remainder}}, \tag{15}$$

with the values of Table 3.

4.2 Bounding the remainders

To certify a condition on a whole interval of values of λ, we bound its remainder in (15) on this interval with quantities computed at $\lambda = 0$ only. This amounts to bounding the scalar function $\xi^\top (I_\rho + \lambda K)^{-1}\zeta$, with $\xi, \zeta \in \mathbb{R}^\rho$ and $K = C$ or $C^\top$. The Neumann series (Golub and Van Loan 2013, Section 2.3) and Hölder's inequality give $|\xi^\top (I_\rho + \lambda K)^{-1}\zeta| \le |\xi|_1 |\zeta|_\infty / (1 - |\lambda| \|K\|_\infty)$ when $|\lambda|\|K\|_\infty < 1$, where $\|\cdot\|_\infty$ is the matrix norm induced by the vector norm $|\cdot|_\infty$. However, $\|K\|_\infty$ can be orders of magnitude larger than the spectral radius of K. Since $(I_\rho + \lambda K)^{-1}\zeta$ only depends on the action of K on the vectors $\zeta, K\zeta, K^2\zeta, \ldots$, the matrix K can be replaced by its restriction to the Krylov subspace spanned by these vectors. This subspace often has a dimension much smaller than ρ. We call this replacement the *reduction to the Krylov subspace*. Together with a diagonal scaling, it gives the following, sharper, bound. Section 4.4 gives the scalings of our implementation.

Lemma 4. *Let $K \in \mathbb{R}^{\rho\times\rho}$ and $\xi, \zeta \in \mathbb{R}^\rho$, let the columns of $Q \in \mathbb{R}^{\rho\times d}$ be an orthonormal basis of the Krylov subspace $\mathcal{K}_\rho(K, \zeta) = \operatorname{span}\{\zeta, K\zeta, \ldots, K^{\rho-1}\zeta\}$, of dimension d, let $H = Q^\top K Q \in \mathbb{R}^{d\times d}$ be the matrix of the restriction of K to this subspace in the basis Q, and let Ω be a positive diagonal matrix of order d and $\gamma = \|\Omega^{-1} H \Omega\|_\infty$. For every λ such that $I_\rho + \lambda K$ is invertible and $|\lambda|\gamma < 1$,*

$$|\xi^\top (I_\rho + \lambda K)^{-1}\zeta| \le \frac{|\Omega Q^\top \xi|_1 \, |\Omega^{-1} Q^\top \zeta|_\infty}{1 - |\lambda|\gamma},$$

where $|\cdot|_1$ is the vector 1-norm.

The proof is in Appendix D.

4.3 Radius of the basis

Each condition of optimality of B is a scalar inequality: a basic variable $(x_B(\lambda))_i$ or a reduced cost $(r(\lambda))_j$ must remain nonnegative. Without loss of generality, we consider $\lambda \ge 0$ (the case $\lambda \le 0$ follows Theorem 2). The value of a condition (i.e. either (12) or (13)) is then given by (15) and Table 3, and its *margin* δ, i.e. its value at $\lambda = 0$, is nonnegative since B is optimal for $\mathcal{P}(0)$. Lemma 4 bounds the remainder: with $a = |\Omega Q^\top \xi|_1 |\Omega^{-1} Q^\top \zeta|_\infty$, it is at least $-a\lambda^2/(1 - \gamma\lambda)$ as long as $I_\rho + \lambda K$ is invertible and $\gamma\lambda < 1$. A condition ((12) or (13)) is therefore guaranteed to hold at λ as soon as

$$f(\lambda) = \delta + \dot{\delta}\,\lambda - \frac{a\,\lambda^2}{1 - \gamma\lambda} \ge 0,$$

and we look for the largest T such that $f \ge 0$ on the whole interval $[0, T)$: the condition then holds for every $\lambda \in [0, T)$. The following lemma gives this T in closed form.

Lemma 5. *Let $\delta \geq 0$, $\dot{\delta} \in \mathbb{R}$, $a \geq 0$ and $\gamma \geq 0$, and let $f(t) = \delta + \dot{\delta}\, t - at^2/(1-\gamma t)$ for $t \in [0, 1/\gamma)$, with $1/0 = +\infty$. Let t_q be the smallest root in $(0, 1/\gamma)$ of the quadratic polynomial $q(t) = \delta + (\dot{\delta} - \gamma\delta)t - (a + \gamma\dot{\delta})t^2$, or $1/\gamma$ if q has no root in $(0, 1/\gamma)$, and*

$$\mathrm{rad}(\delta, \dot{\delta}, a; \gamma) = \begin{cases} 0 & \text{if } \delta = 0 \text{ and } \big(\dot{\delta} < 0 \text{ or } (\dot{\delta} = 0 \text{ and } a > 0)\big), \\ 1/\gamma & \text{if } \delta = \dot{\delta} = a = 0, \\ t_q & \text{otherwise.} \end{cases} \tag{16}$$

Then $\mathrm{rad}(\delta, \dot{\delta}, a; \gamma)$ *is the largest $T \leq 1/\gamma$ such that $f \geq 0$ on $[0, T)$.*

The proof is in Appendix D.

Lemma 5 gives each basic variable $i \in B$ and each reduced cost $j \in N$ its own radius, R_i^x and R_j^r, equal to $\mathrm{rad}(\delta, \dot{\delta}, a; \gamma)$ with the quantities of the corresponding row of Table 3 and with a scaling potentially chosen independently for each of them. The basis matrix limits the radius as well: by Lemma 1, $A_B + \lambda D_B$ is singular if and only if $\lambda = -1/\mu$ for a real eigenvalue $\mu < 0$ of C, and we denote by η^+ the smallest of these values ($\eta^+ = +\infty$ if there is none).

Theorem 2. *Let B be an optimal basis for $\mathcal{P}(0)$ and, with the notation above,*

$$\Delta^+ = \min\Big(\eta^+,\ \min_{i \in B} R_i^x,\ \min_{j \in N} R_j^r\Big).$$

Then B is optimal for $\mathcal{P}(\lambda)$ for every $\lambda \in [0, \Delta^+)$.

Proof. Let $0 \leq \lambda < \Delta^+$. Since $\lambda < \eta^+$, $A_B + \lambda D_B = A_B(I_m + \lambda E)$ is invertible, and so are $I_m + \lambda E$ and, by Theorem 1, $I_\rho + \lambda C$. For each basic variable and each reduced cost, λ is smaller than its radius, which is at most $1/\gamma$, so that Lemma 4 applies to its remainder and its value at λ is at least $f(\lambda)$, which is nonnegative by Lemma 5. The basis B is thus feasible and optimal for $\mathcal{P}(\lambda)$. □

For $\lambda \leq 0$, $\mathcal{P}(\lambda)$ is $\mathcal{P}(-\lambda)$ with D replaced by $-D$. This changes the signs of E, C, $\dot{\delta}$, ξ and ζ, but not a and γ. Theorem 2 thus gives the radius Δ^- with $-\dot{\delta}$ instead of $\dot{\delta}$ and the real eigenvalues $\mu > 0$ of C.

Every condition is handled separately, which makes the radius resistant to degeneracy. A degenerate condition, i.e. $\delta = 0$, gives $\mathrm{rad}(0, \dot{\delta}, a; \gamma) = \dot{\delta}/(a + \gamma\dot{\delta}) > 0$ if $\dot{\delta} > 0$, and a zero radius if $\dot{\delta} < 0$. The zero radius is the exact answer: the condition is then violated at first order, for any small $\lambda > 0$.

Corollary 1. *With the notation of Theorem 2, let (a_o, γ_o) be the coefficients of Lemma 4 for the objective (Table 3). For every $\lambda \in (-\Delta^-, \Delta^+)$ such that $\gamma_o|\lambda| < 1$,*

$$|o^*(\lambda) - o^*(0) - \lambda\, \dot{o}_B| \leq \frac{a_o\, \lambda^2}{1 - \gamma_o|\lambda|}. \tag{17}$$

Proof. By Theorem 2, $o^*(\lambda) = o_B(\lambda)$, and the bound follows from (14) and Lemma 4. □

This bound gives the guarantee on the objective that motivated this section: for a tolerance $\epsilon > 0$, $|o^*(\lambda) - o^*(0)| \leq \epsilon$ for all $\lambda \in (-\Delta^-, \Delta^+)$ such that $|\lambda| < \mathrm{rad}(\epsilon, -|\dot{o}_B|, a_o; \gamma_o)$, by Lemma 5.

4.4 Choice of the scaling

Lemma 4 holds for every positive diagonal scaling Ω. The resulting radius depends on Ω, and the best scaling can differ from one condition to another. As every scaling gives a valid radius, we compute, for each condition, the radius with two classical scalings and keep the larger one:

- the *balancing* scaling (Osborne 1960, Parlett and Reinsch 1969), computed by the LAPACK routine `xGEBAL`, a general-purpose choice that reduces the norm of $\Omega^{-1} H \Omega$;
- the *Perron* scaling (Safonov 1982) $\Omega = \mathrm{diag}(\psi)$, where the Perron vector ψ of $|H|$ is its eigenvector with positive entries, associated with its largest eigenvalue: when $|H|$ is irreducible, this vector exists (Horn and Johnson 2012, Chapter 8) and gives the smallest possible value of γ.

Appendix B gives the details of these scalings, and Section 6.4 measures their effect.

4.5 Computation

The radius only uses the low-rank factors of Section 3: no solve with A_B is needed beyond those that form $\hat{U}$ and $\hat{V}$ (and, with bounded variables, the vector h of Appendix A.1). Moreover, it does not need one Krylov subspace per condition: the Arnoldi process (which builds an orthonormal basis of a Krylov subspace one vector at a time, and is recalled in Section 5) on C started at w_x serves all the basic variables and the objective, and the one on $C^\top$ started at w_c serves all the reduced costs. With d and d_c the dimensions of these two subspaces, usually much smaller than ρ (in the experiments of Section 6, the median of d/ρ is 0.12 and that of d_c/ρ is 0.29), and $\mathrm{nz} = \mathrm{nnz}(A_N) + \mathrm{nnz}(D_N)$, the radius requires $O(m\rho(d + d_c) + \mathrm{nz}\, d_c + d_c^2|N|)$ operations in total. Appendix B gives the details.

5 Matrix-free evaluation with shifted Krylov subspaces

This section computes $x_B(\lambda)$ without forming E or the low-rank factors: the systems $(I_m + \lambda E)\, x = g$ for all the values of λ are solved simultaneously in one Krylov subspace of small dimension, built with one solve with A_B per dimension. This is useful when ρ is large, since the low-rank method of Section 3 needs ρ solves with A_B and ρ with $A_B^\top$ before the first value of λ can be evaluated, and the dense methods need the full matrix E.

5.1 Shifted Arnoldi process

For every λ, $x_B(\lambda)$ is the solution of $(I_m + \lambda E)\, x = x_B(0)$. For a vector g, the Krylov subspace $\mathcal{K}_j(E, g) = \mathrm{span}\{g, Eg, \ldots, E^{j-1}g\}$ is the same for all the matrices $I_m + \lambda E$: $\mathcal{K}_j(I_m + \lambda E, g) = \mathcal{K}_j(E, g)$ for $\lambda \neq 0$. A single basis of $\mathcal{K}_j(E, x_B(0))$ therefore serves all the values of λ, which is the principle of the Krylov methods for shifted systems (Datta and Saad 1991, Frommer and Glässner 1998, Simoncini and Szyld 2007). This principle is classical. This section applies it to $E = A_B^{-1} D_B$ with the factorization of the solver, and shows that the low rank of D_B bounds the number of steps (Theorem 3). The Arnoldi process (Saad 2003) builds an orthonormal basis $Q_j = [v_1, \ldots, v_j]$ of this subspace one vector at a time, starting from $v_1 = x_B(0)/\|x_B(0)\|_2$. Step j computes $w = Ev_j = A_B^{-1}(D_B v_j)$ with a sparse product with D_B and one solve with A_B (an FTRAN), so that E is never formed. It then orthogonalizes w against $v_1, \ldots, v_j$ by Gram–Schmidt: for $i = 1, \ldots, j$, $h_{i,j} = v_i^\top w$ and $w \leftarrow w - h_{i,j} v_i$. Finally, $h_{j+1,j} = \|w\|_2$ and $v_{j+1} = w/h_{j+1,j}$, and the process stops when $h_{j+1,j} = 0$. The coefficients $h_{i,j}$ form an upper Hessenberg matrix $\underline{H}_j \in \mathbb{R}^{(j+1)\times j}$ such that

$$EQ_j = Q_{j+1}\underline{H}_j = Q_j H_j + h_{j+1,j}\, v_{j+1} e_j^\top, \tag{18}$$

where $H_j = Q_j^\top E Q_j$ is made of the first j rows of $\underline{H}_j$ and $h_{j+1,j}$ is its last entry.

We approximate $x_B(\lambda)$ by a vector of the subspace, $x_j(\lambda) = Q_j z_j(\lambda)$, whose coordinates $z_j(\lambda) \in \mathbb{R}^j$ in the basis Q_j are chosen so that the residual $x_B(0) - (I_m + \lambda E)\, x_j(\lambda)$ is orthogonal to the subspace (the full orthogonalization method (Saad 2003)). Since $Q_j^\top x_B(0) = \|x_B(0)\|_2\, e_1$ and, by (18), $Q_j^\top (I_m + \lambda E) Q_j = I_j + \lambda H_j$, this condition is the $j \times j$ system

$$(I_j + \lambda H_j)\, z_j(\lambda) = \|x_B(0)\|_2\, e_1, \tag{19}$$

a Hessenberg system that is solved in $O(j^2)$ for each value of λ by Gaussian elimination with partial pivoting (Golub and Van Loan 2013).

As in the full orthogonalization method (Saad 2003, Section 6.4, Proposition 6.7), (18) gives, for every λ such that $I_j + \lambda H_j$ is invertible, the residual

$$x_B(0) - (I_m + \lambda E)\, x_j(\lambda) = -\lambda\, h_{j+1,j}\, \left(e_j^\top z_j(\lambda)\right) v_{j+1}. \tag{20}$$

Its norm $|\lambda\, h_{j+1,j}\, (z_j(\lambda))_j|$ is thus available for all the values of λ at no extra cost. It does not tell when the approximation becomes exact. While for a general matrix the Arnoldi process may need up to m steps, in our case the low rank of D_B bounds the dimension of the Krylov subspace:

Theorem 3. $\dim \mathcal{K}_j(E, g) \leq \operatorname{rank}(D_B) + 1$ *for all* j *and all* g. *Consequently, in exact arithmetic, the Arnoldi process on* E *stops (with* $h_{j+1,j} = 0$*) after at most* $\operatorname{rank}(D_B) + 1 \leq \rho + 1$ *steps, and then* $x_j(\lambda) = x_B(\lambda)$ *for every* λ *such that* $I_m + \lambda E$ *is invertible (the eigenvalues of* H_j *are then eigenvalues of* E*, so* $I_j + \lambda H_j$ *is invertible too).*

Proof. For $i \geq 1$, $E^i g = A_B^{-1} D_B (E^{i-1} g)$ lies in the range of $A_B^{-1} D_B$, of dimension $\operatorname{rank}(D_B)$; hence $\mathcal{K}_j(E, g) \subseteq \operatorname{span}\{g\} + \operatorname{range}(A_B^{-1} D_B)$. The process stops when $\mathcal{K}_j$ becomes invariant under E, i.e. when $h_{j+1,j} = 0$; H_j is then the restriction of E to this invariant subspace, and the residual (20) vanishes. □

The dual solution satisfies $(I_m + \lambda \widetilde{E})\, y(\lambda) = y(0)$ with $\widetilde{E} = (A_B^\top)^{-1} D_B^\top$, so that the same process, applied to $\widetilde{E}$ with one solve with $A_B^\top$ (a BTRAN) per step, gives $y(\lambda)$ and the reduced costs. The objective only needs the vector $c_B^\top Q_j$, and requires $O(j^2)$ operations per value of λ.

The eigenvalues of H_j (the Ritz values) only include the eigenvalues of E reached from $x_B(0)$, so the Krylov process does not provide a complete existence test. For this, it is easier to use the eigenvalues of C as computed in Section 3.

5.2 Adaptive evaluation

The residual (20) gives an adaptive algorithm: the subspace is extended, one solve with A_B at a time, until the relative residual of every requested value of λ is below a tolerance (10^{-12} in our experiments), and it can be extended later if new values of λ are requested; a value of λ whose primal or dual residual has not converged is never declared optimal. The cost is j solves with A_B, j sparse products with D_B and $O(mj^2)$ operations for the orthogonalization, plus $O(j^2 + mj)$ per value of λ; Theorem 3 guarantees $j \leq \rho + 1$ in exact arithmetic. In practice, j is much smaller: a median of 8 steps in the experiments of Section 6, against a median of 29 for ρ.

6 Numerical experiments

This section evaluates the methods of Section 2, together with the low-rank and Krylov variants of Sections 3 and 5, on 398 parametric linear programs built from real models. We address five questions: (i) do the decompositions exist numerically for the optimal bases of real problems, and how accurate are their results; (ii) what do the preprocessing and the evaluation of one value of λ cost, compared with re-solving $\mathcal{P}(\lambda)$; (iii) on which range of λ does an optimal basis remain optimal, and how much of it does the local radius of Section 4 certify; (iv) is reusing the optimal bases faster than re-solving when computing $o^*(\lambda)$ on a grid of values of λ; and (v) how much faster are the methods on the stability interval of an optimal basis, with and without a proof of optimality at each value. Implementation details, including the baselines (cold solves, solver warm starts, sparse LU) and the hardware, are given in Appendix B. The code and the scripts that reproduce all the experiments are publicly available together with the dataset (Derval 2025).

6.1 Instances

The instances come from the dataset of Derval (2025), with five families: the LPs of Netlib and of MIPLIB 2017 with randomly perturbed coefficients, capacitated facility location, unit commitment, and real models where D describes a physically meaningful change. They have up to 349 120 constraints, and the mixed-integer models are taken as their LP relaxations (Appendix C). The nominal value λ_0 is the middle of the range of λ (the instances are shifted so that λ_0 becomes 0, as in the rest of the paper), and Λ is a grid of p equally spaced values of this range, with $p = 101$ unless stated otherwise.

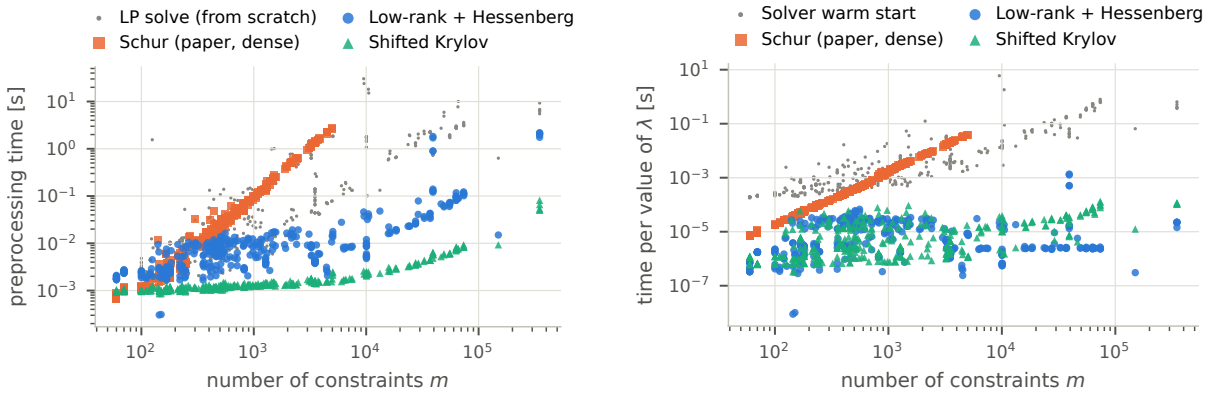
6.2 Accuracy and robustness of the decompositions

We compare, for the optimal basis B of $\mathcal{P}(\lambda_0)$ and the 101 values of Λ, the results of each method with those of a sparse LU factorization of $A_B + \lambda D_B$, on the values of λ for which B is optimal (the dense

Table 4: Median times (one thread).

	310 instances with $m \leq 4985$				All the 398 instances			
Method	Prep.	Obj. / λ	Cert. / λ	Speed-up	Prep.	Obj. / λ	Cert. / λ	Speed-up
Eigendecomposition	3.5 ms	3.2 μs	74 μs	3.2×	–	–	–	–
Schur	18 ms	297 μs	1.1 ms	0.7×	–	–	–	–
Hessenberg (dense)	12 ms	447 μs	1.8 ms	0.4×	–	–	–	–
Zuidwijk	20 ms	64 μs	–	–	–	–	–	–
Low-rank + eig.	2.9 ms	0.38 μs	74 μs	7.0×	4.8 ms	0.41 μs	159 μs	9.9×
Low-rank + Schur	3.9 ms	2.7 μs	136 μs	6.5×	5.7 ms	2.3 μs	190 μs	7.9×
Low-rank + Hessenberg	4.1 ms	3.4 μs	133 μs	6.6×	6.2 ms	3 μs	180 μs	8.9×
Shifted Krylov	1.2 ms	1.9 μs	128 μs	7.4×	1.3 ms	2.7 μs	172 μs	9.1×
LP solve from scratch	–		10 ms		–		25 ms	
Solver warm start	–		602 μs		–		1.1 ms	

Preprocessing, objective and certified $o^*(\lambda)$ per value of λ; the dense methods only run on the left part. The speed-up is the median, over the instances, of the ratio between the time of a solver warm start and the time to compute and certify $o_B(\lambda)$ (the eigendecomposition methods are only timed on the instances where they exist). The times of the Krylov method per value of λ include its Arnoldi steps, amortized over the values evaluated; its preprocessing does not.



(a) Preprocessing time of the Schur method (dense), of the low-rank method with a Hessenberg decomposition of C and of the Krylov method, compared with the time to solve $\mathcal{P}(\lambda)$ from scratch.

(b) Time to compute $o_B(\lambda)$ for one value of λ, without the optimality test, compared with the time of a solver warm start of $\mathcal{P}(\lambda)$.

Figure 1: Times of the methods against the number of constraints m.

methods are run on the 310 instances with $m \leq 4985$). All the methods behave correctly: the error is at most $1.9 \cdot 10^{-10}$ on the objective and $1.5 \cdot 10^{-9}$ on $x_B(\lambda)$, and the optimality tests always agree with the reference. There are two exceptions. The eigendecomposition of E does not exist (numerically) for 214 of the 310 instances (69%), confirming that defective matrices E are the rule rather than the exception on real problems (Lemmas 1 and 2 explain why), and that of C does not exist for 128 of the 398 instances. Zuidwijk's formula, which only gives the objective, reaches an error of $3.8 \cdot 10^{-9}$.

6.3 Computational cost

Table 4 reports the median preprocessing time of each method, the median time to compute $o_B(\lambda)$ for one value of λ (amortized over 1000 values), and the median time to compute $o_B(\lambda)$ *and* certify that B is optimal for λ (existence, feasibility and optimality tests), which is the quantity to compare with one re-solve of $\mathcal{P}(\lambda)$. Figure 1 shows how these times grow with m.

The dense methods follow their $O(m^2)$ per-λ complexity: the Schur method is faster than a solver warm start on all the instances with $m \leq 200$, but slower on 97% of those with $m > 1000$.

The preprocessing of the low-rank method amounts to its 2ρ solves and the decomposition of C; the

Table 5: Certified radius Δ^σ of the optimal basis B_0 of $\mathcal{P}(\lambda_0)$ against the corresponding side $\Delta^\sigma_{\text{ex}}$ of its exact stability interval, on the sides longer than 10^{-6}.

Family	#	Sides	$\Delta^\sigma/\Delta^\sigma_{\text{ex}}$			$\Delta^\sigma \geq \Delta^\sigma_{\text{ex}}/2$		
			median	$\geq 1/2$	$\geq 1/10$	identity	no red.	norms E
Netlib	115	101	0.79	58%	78%	46%	16%	6%
MIPLIB	31	28	0.87	64%	79%	61%	39%	30%
Facility	72	143	1.00	100%	100%	100%	62%	37%
Unit comm.	54	92	1.00	100%	100%	100%	100%	96%
Real models	126	231	1.00	96%	99%	89%	61%	25%
All	398	595	1.00	90%	95%	85%	59%	29%

Number of instances and of sides; median ratio between the radius computed as in Section 4 and the exact side, and share of the sides where it reaches $1/2$ and $1/10$ of the exact side; same share ($1/2$) with the identity scaling only, without the reduction to the Krylov subspaces, and with the norms of E instead of the low-rank coefficients (on the 458 sides where they could be computed).

Krylov method only extracts D_B in its preprocessing, and its evaluation times include its solves. Both methods evaluate $o_B(\lambda)$ in a few microseconds in median (at most 1.4 ms for the low-rank method and 133 μs for the Krylov method over all the instances), two to three orders of magnitude faster than a solver warm start. Certifying the optimality of B is more expensive, because the reduced costs require a product with A_N (Equation (11)); it remains 8.9× (low-rank) and 9.1× (Krylov) faster than a solver warm start, in median.

These costs are low because the dimensions involved are small. The low-rank matrix C has a median dimension of 29 (at most 669), and the Krylov method reaches a relative residual of 10^{-12} for the 101 values of λ of the whole range after a median of 8 and at most 38 steps (largest over the primal, dual and, when $h \neq 0$, h subspaces of Appendix A.1), usually far below the bound $\rho + 1$ of Theorem 3 (the median ratio to ρ is 0.24, and the bound is only reached on 8 instances, all with $\rho \leq 10$).

6.4 Stability of the optimal basis and local radius

The methods compute $o^*(\lambda)$ exactly only where B remains optimal. For the optimal basis B_0 of $\mathcal{P}(\lambda_0)$, we compute the exact stability interval, i.e. the largest interval containing λ_0 on which B_0 is optimal (with the tolerances of Appendix B, and up to the resolution of the scan), by scanning the range and bisecting on the first value where B_0 is not optimal (the fast evaluation of Section 3 makes this cheap: 119 ms in median), and the certified radii Δ^- and Δ^+ of Theorem 2, denoted Δ^σ with $\sigma \in \{-, +\}$. We call *sides* of the interval the distances Δ^-_{ex} and Δ^+_{ex} from λ_0 to its two ends; sides and radii are limited to the range of λ.

The stability intervals are small: their median width is $8.6 \cdot 10^{-3}$ for ranges of width 2 (0.3 for the facility location instances), and B_0 is optimal on a median of 1 of the 101 values of Λ. Exactly one of their sides is shorter than 10^{-6} for 10% of the instances, and both of them for 20%: the basis is degenerate, and a basic variable or a reduced cost leaves its bound as soon as λ moves in one (or either) direction, so that the median of the smaller side, $\min(\Delta^-_{\text{ex}}, \Delta^+_{\text{ex}})$, is only $7.3 \cdot 10^{-4}$. Degeneracy is the rule: 85% of the optimal bases have a basic variable at one of its bounds.

Table 5 compares the certified radius Δ^σ with the corresponding side $\Delta^\sigma_{\text{ex}}$ of the exact interval. The radius detects all the 201 sides shorter than 10^{-6}, on which B_0 does not satisfy the optimality conditions in that direction. On the 595 longer ones, it is a median of 0.9996 times the exact side, at least half of it on 90% of them and at least a tenth on 95% (80 of them end at the limit of the range). It costs 14 ms in median on top of the low-rank preprocessing (13× less than the exact interval). We check the radii numerically, by evaluating the optimality conditions at 40 points of each certified interval, including its ends and its midpoint: they hold up to rounding errors.

The reduction to the Krylov subspaces and the scalings both matter (Table 5): without the reduction, the radius reaches half of the exact side on 59% of the sides instead of 90%. With a simpler bound that

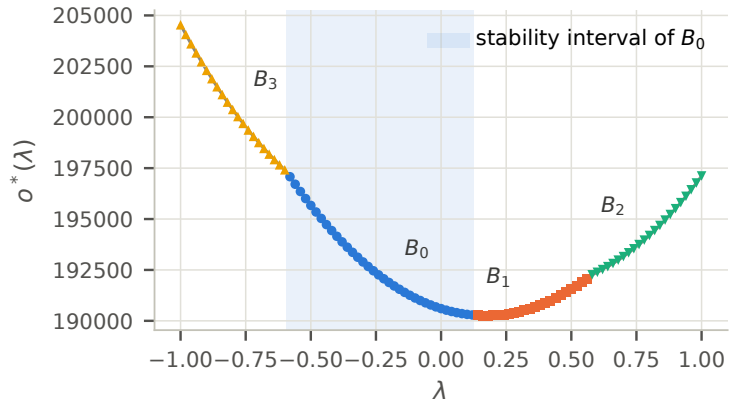


Figure 2: $o^*(\lambda)$ for the instance `real/sndlib_abilene__east_west_shift` on the grid of 101 values: each color and marker shape corresponds to one optimal basis; the bases are numbered in the order of their discovery, and the low-rank method computes their values; the shaded area is the exact stability interval of B_0.

Table 6: Computation of $o^*(\lambda)$ on the grid of $p = 101$ values.

Family	#	Bases / point	Coverage of B_0 # points	Coverage of B_0 radius	Total time / cold re-solves solver warm start	LR+Hess.	Krylov	LU
Netlib	115	0.48	1	$1.0 \cdot 10^{-9}$	0.08	0.30	0.19	0.59
MIPLIB	31	0.12	1	$2.2 \cdot 10^{-9}$	0.05	0.07	0.05	0.35
Facility	72	0.07	26	$1.3 \cdot 10^{-2}$	0.08	0.09	0.09	0.42
Unit comm.	54	1.00	1	$4.1 \cdot 10^{-5}$	0.09	0.16	0.16	0.31
Real models	126	0.80	1	$2.2 \cdot 10^{-3}$	0.06	0.13	0.12	0.48
All	398	0.52	1	$7.3 \cdot 10^{-4}$	0.07	0.13	0.12	0.44

Median number of optimal bases per value of λ, median number of values covered by B_0 and median radius of its stability interval, and median total times relative to p cold re-solves, for the solver warm starts and for the basis-reuse procedure with the low-rank (Hessenberg), Krylov and LU methods.

uses the norms of the $m \times m$ matrix E, i.e. $\gamma = \|E\|_\infty$, without reduction or scaling, it reaches half of the exact side on only 29% of the 458 sides where they could be computed.

6.5 Computing $o^*(\lambda)$ on a grid

We finally compute $o^*(\lambda)$ for all $\lambda \in \Lambda$ with the procedure of Section 2.3: starting from B_0, the values of Λ are swept from λ_0 towards each end of the range; while the current basis is optimal, $o^*(\lambda)$ is given by the method, and on the first value where it is not, $\mathcal{P}(\lambda)$ is re-solved with a warm start from the current basis and the method restarts from the new basis. The values are evaluated in blocks of increasing size, so that a basis that is quickly abandoned costs little. All the computed values agree with the solver: they differ by at most $7.4 \cdot 10^{-8}$ from its warm-started or its cold answer. The largest difference with the cold answer alone, $1.4 \cdot 10^{-4}$, is that of the solver itself, whose warm-started and cold answers differ by up to $1.4 \cdot 10^{-4}$ (on instance `miplib/sp150x300d_eq`). $\mathcal{P}(\lambda)$ is infeasible on 10 values of each of 4 facility instances, and the solver warm start does not return an optimal solution on 4 values of 2 other instances, where a cold solve does: these values have no result. On the finer grids below, the computed values differ by at most $3.9 \cdot 10^{-7}$ from the solver warm starts. Figure 2 illustrates the procedure on one instance, and Table 6 summarizes the results for $p = 101$; Table 7 repeats the experiment with up to $p = 10001$ values, on the 392 instances with $m \leq 10^5$.

The time of the solver warm starts grows linearly with p (52× from $p = 101$ to $p = 10001$, in median), since each value needs a re-solve, whereas that of basis reuse grows much more slowly (7.3× with the Krylov method), since one basis serves many values. The Krylov method takes a median of 1.70, 1.00 and 0.41 times the time of the solver warm starts for $p = 101$, 1001 and 10001, and it is faster on 18%, 50% and 80% of the instances, by up to 34× (Table 7).

Table 7: Computation of $o^*(\lambda)$ on grids of p values, on the 392 instances with $m \leq 10^5$.

Family	#	p	Basis changes	Distinct bases	Total time (median) Solver warm start	LR+Hess.	Krylov	Reuse / solver warm start LR+Hess.	Krylov	Faster (Krylov)
Netlib	115	101	48	47	99 ms	553 ms	392 ms	3.75	2.53	13%
		1001	85	72	653 ms	1.4 s	787 ms	1.69	0.99	51%
		10001	106	86	6 s	3.4 s	2.1 s	0.51	0.28	91%
MIPLIB	30	101	13	13	350 ms	704 ms	479 ms	1.08	1.00	50%
		1001	20	14	2.2 s	2 s	1.3 s	0.71	0.52	73%
		10001	25	14	19 s	9.6 s	7.6 s	0.43	0.35	93%
Facility	72	101	7	7	36 ms	53 ms	57 ms	1.29	1.25	40%
		1001	8	8	283 ms	121 ms	132 ms	0.39	0.43	99%
		10001	8	8	2.8 s	710 ms	716 ms	0.26	0.27	100%
Unit comm.	54	101	101	101	12 s	17 s	17 s	1.44	1.36	0%
		1001	1000	1000	37 s	101 s	82 s	2.68	2.20	0%
		10001	5669	5669	193 s	463 s	352 s	2.62	1.96	4%
Real models	121	101	78	78	274 ms	823 ms	810 ms	2.29	2.05	9%
		1001	230	228	1.3 s	2.6 s	2.3 s	1.57	1.51	36%
		10001	326	302	11 s	6.9 s	6.3 s	0.62	0.49	87%
All	392	101	52	52	126 ms	563 ms	486 ms	1.81	1.70	18%
		1001	112	99	763 ms	1.6 s	1.2 s	1.45	1.00	50%
		10001	155	130	6.6 s	4.9 s	3.5 s	0.50	0.41	80%

Median number of changes of basis and of distinct optimal bases along the sweep of the Krylov method, median total times, and median ratio between the time of the basis-reuse procedure and that of the solver warm starts (below 1: basis reuse is faster), with the fraction of the instances where basis reuse with the Krylov method is faster.

Table 8: Local analysis, with and without the optimality test.

			With the optimality test				Without the optimality test			
Family	#	Solver warm start / λ	LR+Hess.	LR+eig.	Krylov	LU	LR+Hess.	LR+eig.	Krylov	LU
Netlib	115	579 μs	3.8×	5.2×	7.6×	0.1×	24.7×	125.0×	256.9×	0.3×
MIPLIB	31	1.8 ms	3.0×	3.8×	4.5×	0.1×	242.3×	506.6×	819.2×	0.5×
Facility	72	287 μs	3.5×	4.0×	4.1×	0.2×	114.9×	211.4×	186.7×	0.7×
Unit comm.	54	9.7 ms	3.6×	3.2×	4.9×	0.1×	508.5×	561.6×	961.3×	0.2×
Real models	126	1.1 ms	3.9×	3.8×	6.5×	0.1×	319.4×	200.4×	395.0×	0.3×
All	398	650 μs	3.7×	3.9×	5.5×	0.1×	117.3×	238.1×	303.9×	0.3×

Median time of a solver warm start per value of λ in the stability interval of B_0, and median speed-up of each method for computing $o^*(\lambda)$ at 1001 values of the interval, preprocessing included, with the optimality test at each value and without it (then without the eigenvalues of the existence test). Without the test, the LR+eig. column only covers the 270 instances where the decomposition exists.

The number of distinct bases is the key quantity. The optimal bases met along the range of λ are essentially a property of the instance, but a grid only visits those whose stability interval contains one of its values: a coarse grid misses the short intervals, so that the number of bases grows with p until the grid resolves them, whereas the cost of the re-solves keeps growing linearly with p. The median number of distinct bases is 52, 99 and 130 for $p = 101$, 1001 and 10001. It stops growing on the Facility instances and, in median, on the MIPLIB instances: a coarse grid already visits most of their bases. It still grows slowly on the Netlib and real-model instances. The unit commitment instances have very short stability intervals: with $p = 10001$, they still need a new basis at more than one value out of two, and basis reuse remains slower on almost all of them for all p. With $p = 101$, a median of 0.52 bases per value of λ is needed: the curves of many instances have a breakpoint every few values (a median of 0.48 bases per value for the Netlib instances, whose perturbation touches 100 constraints), and the solver warm start, which needs only a few pivots per value, is then hard to beat. With $p = 10001$, the outcome is governed by the number of distinct bases per value of λ: basis reuse with the Krylov method is faster on 95% of the instances that need at most one new basis every 20 values, and on 0% of those that need more than one every 5 values. With $p = 101$, the LU baseline is slower than the Krylov method on all the instances, which confirms that the gain comes from avoiding a factorization per value of λ.

6.6 Local analysis

The methods of this paper do not replace the re-solve at each change of basis. Their natural use is the analysis of o^* in the neighborhood of a solved problem, on the stability interval of its optimal basis. Table 8 reports the speed-up, with respect to a solver warm start, of computing and certifying $o^*(\lambda)$ for 1001 values of λ inside the exact stability interval of B_0 (the solver warm starts start from B_0 and need no pivot there, except on 2 instances).

Within the stability interval, $o^*(\lambda)$ is obtained 5.5× (Krylov) and 3.7× (low-rank) faster than with a solver warm start, in median, with a certificate of optimality. Almost all of this time is the optimality test (Table 4), which is not needed at each value: the radius of Section 4 or the exact stability interval is computed once and certifies all the values inside it. The right part of Table 8 reports the speed-up without this test: when the interval is known, it reaches 117× (low-rank) and 304× (Krylov) in median, preprocessing included. With the one-time cost of the interval and of the low-rank factors, amortized over the 1001 values, it is 34× (low-rank) and 39× (Krylov) with the radius (14 ms in median on top of the factors), and 6.1× and 6.4× with the exact interval (119 ms). These speed-ups hold inside the interval only, and they depend on the number of values over which the preprocessing is amortized. The dense methods of Section 2 are only competitive on the smallest instances (the instances on which the Schur method is faster than a solver warm start have a median of 162 constraints, against 504 for those on which it is slower).

6.7 Discussion

The low-rank and Krylov methods handle instances with hundreds of thousands of constraints, with the factorization already maintained by the solver. The Krylov method is the method of choice for evaluating $o^*(\lambda)$ at many values of λ: it only builds the necessary dimensions (a median of 8 steps, against a median dimension of 29 for C) and never forms the low-rank factors. The low-rank method is preferable when a certificate of optimality on an interval or a complete existence test is required, since its factors give the local radius of Section 4 and the eigenvalues of C give all the values of λ at which $A_B + \lambda D_B$ is singular.

7 Conclusion

For linear programs $\mathcal{P}(\lambda)$ whose constraint matrix varies linearly with a parameter λ, this paper reuses an optimal basis B of $\mathcal{P}(0)$ to compute $x_B(\lambda)$ and $o_B(\lambda)$ for many values of λ, together with a proof that B remains optimal. When the perturbation of the basis matrix has a low rank ρ, all the computations reduce to a $\rho \times \rho$ matrix built with a few solves with the factorization maintained by the simplex solver; shifted Krylov subspaces give a matrix-free alternative, which is cheaper when only $o_B(\lambda)$ is needed. Finally, the low-rank factors give, in each direction, a radius up to which B is certified to remain optimal, with a bound on the objective; this radius handles degenerate bases.

On 398 instances built from real models, with up to 349 120 constraints, the low-rank and Krylov methods compute $o^*(\lambda)$ on the stability interval of a basis 3.7× to 5.5× faster than a solver warm start with a certificate of optimality at each value, and 117× to 304× faster once the interval is known, in median and preprocessing included. Over the whole range of λ, reusing the bases becomes faster than solver warm starts as the grid gets finer: with 10001 values, the Krylov method takes a median of 0.41 times their time, unless the basis changes at almost every value. The certified radius reaches at least half of the exact stability interval on 90% of the sides longer than 10^{-6}. As optimal bases of real problems typically remain optimal on short ranges of λ, the methods are best used for local analyses and dense samplings of λ, combined with solver warm starts at the breakpoints of $o^*(\lambda)$.

Agentic AI usage disclosure

This article was written with the help of multiple AI tools, including agentic AI systems, after an initial draft by the authors. The authors supervised all interactions with these tools and verified the content produced, agree with its content and take responsibility for its correctness.

References

Bart H, Gohberg I, Kaashoek MA (1979) *Minimal Factorization of Matrix and Operator Functions*, volume 1 of *Operator Theory: Advances and Applications* (Basel: Birkhäuser), URL `http://dx.doi.org/10.1007/978-3-0348-6293-6`.

Bartlett MS (1951) An inverse matrix adjustment arising in discriminant analysis. *The Annals of Mathematical Statistics* 22(1):107–111, URL `http://dx.doi.org/10.1214/aoms/1177729698`.

Beasley JE, Meade N, Chang TJ (2003) An evolutionary heuristic for the index tracking problem. *European Journal of Operational Research* 148(3):621–643, URL `http://dx.doi.org/10.1016/S0377-2217(02)00425-3`.

Brown T, Hörsch J, Schlachtberger D (2018) PyPSA: Python for power system analysis. *Journal of Open Research Software* 6(1):4, URL `http://dx.doi.org/10.5334/jors.188`.

Datta BN, Saad Y (1991) Arnoldi methods for large Sylvester-like observer matrix equations, and an associated algorithm for partial spectrum assignment. *Linear Algebra and its Applications* 154–156:225–244, URL `http://dx.doi.org/10.1016/0024-3795(91)90378-A`.

Demmel J, Dumitriu I, Holtz O (2007) Fast linear algebra is stable. *Numerische Mathematik* 108(1):59–91, URL `http://dx.doi.org/10.1007/s00211-007-0114-x`.

Derval G (2025) Parametric linear optimization: real problems dataset. `https://github.com/GuillaumeDerval/parametric-linear-optimization`, branch `real-problems-dataset`.

Derval G, Ernst D, Louveaux Q, Miftari B (2026) Sensitivity analysis for linear changes of the constraint matrix of a (mixed-integer) linear program, arXiv:2410.14443.

Ernst AT, Krishnamoorthy M (1999) Solution algorithms for the capacitated single allocation hub location problem. *Annals of Operations Research* 86:141–159, URL http://dx.doi.org/10.1023/A:1018994432663.

Feng X, Barcelos G, Gaboardi JD, Knaap E, Wei R, Wolf LJ, Zhao Q, Rey SJ (2022) spopt: a Python package for solving spatial optimization problems in PySAL. *Journal of Open Source Software* 7(74):3330, URL http://dx.doi.org/10.21105/joss.03330.

Flanders H (1951) Elementary divisors of AB and BA. *Proceedings of the American Mathematical Society* 2(6):871–874, URL http://dx.doi.org/10.1090/S0002-9939-1951-0044493-X.

Frangioni A, Gentile C, Lacalandra F (2009) Tighter approximated MILP formulations for unit commitment problems. *IEEE Transactions on Power Systems* 24(1):105–113, URL http://dx.doi.org/10.1109/TPWRS.2008.2004744.

Freund RM (1985) Postoptimal analysis of a linear program under simultaneous changes in matrix coefficients. *Mathematical Programming Study* 24:1–13, URL http://dx.doi.org/10.1007/BFb0121039.

Frommer A, Glässner U (1998) Restarted GMRES for shifted linear systems. *SIAM Journal on Scientific Computing* 19(1):15–26, URL http://dx.doi.org/10.1137/S1064827596304563.

Gal T (1995) *Postoptimal Analyses, Parametric Programming, and Related Topics: Degeneracy, Multicriteria Decision Making, Redundancy* (Berlin, New York: Walter de Gruyter), second edition, URL http://dx.doi.org/10.1515/9783110871203.

Gal T, Nedoma J (1972) Multiparametric linear programming. *Management Science* 18(7):406–422, URL http://dx.doi.org/10.1287/mnsc.18.7.406.

Gass SI, Saaty TL (1955) Parametric objective function (Part 2)—generalization. *Journal of the Operations Research Society of America* 3(4):395–401, URL http://dx.doi.org/10.1287/opre.3.4.395.

Gay DM (1985) Electronic mail distribution of linear programming test problems. *Mathematical Programming Society COAL Newsletter* 13:10–12, the collection is available at https://netlib.org/lp/.

Gleixner A, Hendel G, Gamrath G, Achterberg T, Bastubbe M, Berthold T, Christophel P, Jarck K, Koch T, Linderoth J, Lübbecke M, Mittelmann HD, Ozyurt D, Ralphs TK, Salvagnin D, Shinano Y (2021) MIPLIB 2017: Data-driven compilation of the 6th mixed-integer programming library. *Mathematical Programming Computation* 13(3):443–490, URL http://dx.doi.org/10.1007/s12532-020-00194-3.

Golub GH, Van Loan CF (2013) *Matrix Computations* (Baltimore: Johns Hopkins University Press), fourth edition.

Greenberg HJ (1999) Matrix sensitivity analysis from an interior solution of a linear program. *INFORMS Journal on Computing* 11(3):316–327, URL http://dx.doi.org/10.1287/ijoc.11.3.316.

Hager WW (1989) Updating the inverse of a matrix. *SIAM Review* 31(2):221–239, URL http://dx.doi.org/10.1137/1031049.

Holmberg K, Rönnqvist M, Yuan D (1999) An exact algorithm for the capacitated facility location problems with single sourcing. *European Journal of Operational Research* 113(3):544–559, URL http://dx.doi.org/10.1016/S0377-2217(98)00008-3.

Horn RA, Johnson CR (2012) *Matrix Analysis* (Cambridge: Cambridge University Press), second edition.

IBM (2022) *CPLEX Callable Library Reference Manual, version 22.1.1*. URL https://www.ibm.com/docs/en/icos/22.1.1.

Jansen B, de Jong J, Roos C, Terlaky T (1997) Sensitivity analysis in linear programming: just be careful! *European Journal of Operational Research* 101(1):15–28, URL http://dx.doi.org/10.1016/S0377-2217(96)00172-5.

Knueven B, Ostrowski J, Watson JP (2020) On mixed-integer programming formulations for the unit commitment problem. *INFORMS Journal on Computing* 32(4):857–876, URL http://dx.doi.org/10.1287/ijoc.2019.0944.

Kondili E, Pantelides CC, Sargent RWH (1993) A general algorithm for short-term scheduling of batch operations—I. MILP formulation. *Computers & Chemical Engineering* 17(2):211–227, URL http://dx.doi.org/10.1016/0098-1354(93)80015-F.

Lee I, Curry S, Serban N (2019) Solving large batches of linear programs. *INFORMS Journal on Computing* 31(2):302–317, URL http://dx.doi.org/10.1287/ijoc.2018.0838.

Orlowski S, Wessäly R, Pióro M, Tomaszewski A (2010) SNDlib 1.0—Survivable network design library. *Networks* 55(3):276–286, URL http://dx.doi.org/10.1002/net.20371.

Osborne EE (1960) On pre-conditioning of matrices. *Journal of the ACM* 7(4):338–345, URL http://dx.doi.org/10.1145/321043.321048.

Parlett BN, Reinsch C (1969) Balancing a matrix for calculation of eigenvalues and eigenvectors. *Numerische Mathematik* 13(4):293–304, URL http://dx.doi.org/10.1007/BF02165404.

Saad Y (2003) *Iterative Methods for Sparse Linear Systems* (Philadelphia: SIAM), second edition, URL http://dx.doi.org/10.1137/1.9780898718003.

Saaty TL, Gass SI (1954) Parametric objective function (Part 1). *Journal of the Operations Research Society of America* 2(3):316–319, URL http://dx.doi.org/10.1287/opre.2.3.316.

Safonov MG (1982) Stability margins of diagonally perturbed multivariable feedback systems. *IEE Proceedings D (Control Theory and Applications)* 129(6):251–256, URL http://dx.doi.org/10.1049/ip-d.1982.0054.

Sherman J, Morrison WJ (1949) Adjustment of an inverse matrix corresponding to changes in the elements of a given column or a given row of the original matrix. *The Annals of Mathematical Statistics* 20(4):621, URL http://dx.doi.org/10.1214/aoms/1177729959, abstract.

Sherman J, Morrison WJ (1950) Adjustment of an inverse matrix corresponding to a change in one element of a given matrix. *The Annals of Mathematical Statistics* 21(1):124–127, URL http://dx.doi.org/10.1214/aoms/1177729893.

Simoncini V, Szyld DB (2007) Recent computational developments in Krylov subspace methods for linear systems. *Numerical Linear Algebra with Applications* 14(1):1–59, URL http://dx.doi.org/10.1002/nla.499.

Woodbury MA (1950) Inverting modified matrices. Memorandum Report 42, Statistical Research Group, Princeton University, Princeton, NJ.

Zuidwijk RA (2005) Linear parametric sensitivity analysis of the constraint coefficient matrix in linear programs. ERIM Report Series Research in Management ERS-2005-055-LIS, Erasmus Research Institute of Management.

A Bounded variables

The paper considers the standard form of Section 1, where the nonbasic variables are zero. This appendix extends the low-rank formulas of Section 3 and the local radius of Section 4 to bounded variables. Solvers handle linear programs in this form.

A.1 Low-rank formulas

With bounds $\ell \le x \le u$, the nonbasic variables of an optimal basis sit at one of their bounds rather than at zero. With x_N the values of the nonbasic variables (slack variables included, with columns of the identity), we redefine $x_B(0) = A_B^{-1}(b - A_N x_N)$, and the basic solution is

$$\begin{aligned} x_B(\lambda) &= (A_B + \lambda D_B)^{-1}\,(b - (A_N + \lambda D_N)\,x_N) \\ &= (I_m + \lambda E)^{-1} g(\lambda), \qquad g(\lambda) = x_B(0) + \lambda h, \qquad h = -A_B^{-1} D_N x_N . \end{aligned}$$

The results of Section 3 hold with $x_B(0)$ replaced by $g(\lambda)$: $x_B(\lambda) = g(\lambda) - \lambda\hat{U}(I_\rho + \lambda C)^{-1}V^\top g(\lambda)$, and the objective is $c^\top x(\lambda) = c_N^\top x_N + c_B^\top g(\lambda) - \lambda w_c^\top (I_\rho + \lambda C)^{-1} V^\top g(\lambda)$. The vector h costs one more solve with A_B (it is zero in the setting of the paper, where $x_N = 0$). Feasibility becomes $\ell_B \le x_B(\lambda) \le u_B$, and optimality requires $r_j(\lambda) \ge 0$ for the nonbasic variables at their lower bound, $r_j(\lambda) \le 0$ for those at their upper bound and $r_j(\lambda) = 0$ for free nonbasic variables. The fixed variables ($\ell_j = u_j$, e.g. the slack variables of equality constraints) impose no condition. In the Krylov method of Section 5, the term λh of $g(\lambda)$ is handled by a second Krylov subspace started at h. All the experiments of Section 6 use this general form.

A.2 Local radius of optimality

With bounded variables, the nonbasic variables are fixed at one of their bounds, x_N, and $x_B(\lambda) = (I_m + \lambda E)^{-1}(x_B(0) + \lambda h)$ with $x_B(0) = A_B^{-1}(b - A_N x_N)$ and $h = -A_B^{-1} D_N x_N$ (Appendix A.1). Lemma 3, Theorem 2 and Corollary 1 then hold with $\dot{x}_B = h - E x_B(0)$, since $x_B(\lambda) - x_B(0) = \lambda(I_m + \lambda E)^{-1}(h - E x_B(0))$, and each finite bound of a basic variable and each sign condition on a reduced cost is a condition with its own margin (the distance to the bound, or $|r_j(0)|$) and its own derivative ($\pm(\dot{x}_B)_i$ or $\pm(\dot{r})_j$); the fixed nonbasic variables impose no condition.

B Numerical and implementation details

All the methods are implemented in Python, with NumPy and SciPy (LAPACK) for the dense linear algebra and Numba for the shifted triangular and Hessenberg solves. The LP $\mathcal{P}(\lambda_0)$, where λ_0 is the nominal value defined in Section 6.1, is solved by the dual simplex method of the solver, and the basis information and the solves with the optimal basis matrix are obtained from the solver itself, through its C interface (`CPXgetbhead`, `CPXftran` and `CPXbtran` of CPLEX 22.1.1 (IBM 2022)). The LU factorization of A_B maintained by the simplex method is thus reused, and no factorization is computed by our code (except for the reference solutions below). The matrix E of the dense methods is formed with n_D solves with A_B, one per nonzero column of D_B.

The methods are compared with three baselines: *cold*, which solves $\mathcal{P}(\lambda)$ from scratch for each λ (dual simplex, new model); *solver warm start*, which updates the coefficients $A + \lambda D$ in the model of the solver and re-optimizes from the previous optimal basis; and *LU*, which computes a sparse LU factorization (SuperLU) of $A_B + \lambda D_B$ for each λ, the "basis system solving" method of Table 1. The LU solutions also serve as references for the accuracy of the methods. All the codes run on a single thread (solver, BLAS and Numba), on a server with two Intel Xeon E5-2650 v4 processors (2.2 GHz, 24 cores in total) and 128 GB of memory. The server processes 12 instances in parallel.

The instances are general LPs, with bounded variables and both equality and inequality constraints, handled in the form of Appendix A.1: a slack variable with a column of the identity is added to each

constraint, the nonbasic variables are fixed at the bound given by the solver, and the conditions of Section 2.2 become $\ell_B \le x_B(\lambda) \le u_B$ and sign conditions on the reduced costs that depend on the bound of each nonbasic variable. A basis is declared optimal for λ when it satisfies the conditions of Section 2.2 with the accuracy of the certificate given by the solver at λ_0. The paragraphs below give the tolerances of these tests, and how the formula of Zuidwijk and the radius of Section 4 are evaluated.

Optimality tolerances. A basis is declared feasible (resp. optimal) for λ when no bound (resp. no bound and no sign condition on the reduced costs) is violated, in relative terms, by more than the largest of 10^{-9} and its violation at λ_0, increased by 10^{-11} for the rounding errors of the evaluation: the basis is thus certified at λ with the accuracy of the certificate given by the solver at λ_0, whose tolerances are 10^{-6}. (A plain relative tolerance of 10^{-6} accepts, on a few degenerate instances, bases whose reduced costs have become slightly negative and whose objective is then 10^{-5} above the optimum.)

Singularity and failed decompositions. $A_B + \lambda D_B$ is considered singular when $|1 + \lambda\mu| \le 10^{-10}$ for an eigenvalue μ of E, C or H_j, depending on the method. For this test, our implementation of the Hessenberg variants computes the eigenvalues of H once, in the preprocessing, rather than using the pivots of the elimination. The eigendecomposition of a matrix is declared failed (numerically defective) when the condition number $\|Q\|_1\|Q^{-1}\|_1$ of its eigenvector matrix exceeds 10^{15}.

Zuidwijk's formula. Zuidwijk's formula is evaluated as $o_B(\lambda) = \lambda^{-1}\left(\exp\left(\sum_j \log(1+\lambda\beta_j) - \sum_j \log(1+\lambda\alpha_j)\right) - 1\right)$ to avoid overflows, with the eigenvalues α_j of E and β_j of $E + x_B(0)c_B^\top$.

Margins of the radius. The radius of Section 4 uses the two scalings of Section 4.4 and keeps the larger radius for each condition. The margins of the conditions are increased by the tolerances of the certificate above: rounding errors give derivatives of the order of 10^{-15} to conditions that do not move, and with exact margins the radius vanishes in the direction in which these derivatives happen to be negative (on 32% of the sides of the stability intervals longer than 10^{-6}).

Errors and violations. Errors are measured relative to $\max(1, |\cdot|)$ for the objective and to $\max(1, \|\cdot\|_\infty)$ for x_B; the violations of a bound and of a sign condition are relative to $1 + |\text{bound}|$ and $1 + |c_j|$, respectively.

Breakdown of the Arnoldi process. In the Krylov method of Section 5, the breakdown test of the Arnoldi process is relative to $\|Ev_j\|_2$, not to $\|x_B(0)\|_2$: on problems with large basic solutions, the latter declares spurious invariant subspaces.

Scalings. The scalings of Section 4.4 rely on the following facts. A scaling changes $\gamma = \|\Omega^{-1}H\Omega\|_\infty$, which is at least the spectral radius π of $|H|$, since a matrix norm is at least the spectral radius and $\Omega^{-1}|H|\Omega$ has the same eigenvalues as $|H|$; it also changes the vectors $\Omega Q^\top\xi$ and $\Omega^{-1}Q^\top\zeta$ that define a. When $|H|$ is irreducible, the Perron–Frobenius theorem (Horn and Johnson 2012, Chapter 8) gives a vector ψ with positive entries such that $|H|\psi = \pi\psi$, and $\Omega = \operatorname{diag}(\psi)$ attains $\gamma = \pi$, since every row of $\Omega^{-1}|H|\Omega$ then sums to π. We approximate ψ by the power method.

Cost of the radius. In the computation of Section 4.5, the coefficients of the basic variables are given by the product $\hat{U}Q$ ($O(m\rho d)$ operations), and those of the reduced costs by $Q_c^\top\theta_j = H_c^\top(\hat{V}Q_c)^\top A_N e_j - (\hat{V}Q_c)^\top D_N e_j$, where Q_c is the basis of the Krylov subspace of $C^\top$ and $C^\top Q_c = Q_c H_c$ ($O(m\rho d_c + \operatorname{nz} d_c + d_c^2|N|)$ operations). The Arnoldi processes require $O(\rho^2(d + d_c))$ operations and the radii, once the coefficients are known, $O(m + |N|)$.

Table 9: Instances.

Family	#	m med (min–max)	n med (min–max)	nnz(D) med	n_D med (max)	ρ med (max)
Netlib	115	551 (123–3100)	1371 (83–13525)	148	91 (153)	83 (100)
MIPLIB	31	4435 (450–149855)	10100 (600–111827)	141	60 (135)	58 (98)
Facility	72	115 (60–230)	2075 (510–6030)	20	20 (30)	20 (30)
Unit comm.	54	32792 (3247–73792)	7500 (750–22500)	2400	2391 (4793)	24 (24)
Real models	126	1275 (144–349120)	3296 (270–305712)	43	21 (669)	21 (669)
All	398	712 (60–349120)	2480 (83–305712)	118	48 (4793)	29 (669)

Number of instances, number of constraints m and of variables n (after the variable bounds are restored as bounds), nonzeros of D, number n_D of perturbed basic columns and dimension ρ of the low-rank matrix C for the optimal basis at λ_0.

Conditions that do not move. A condition with a zero margin and a zero derivative gets a zero radius if its coefficient a is nonzero. This is a limitation of a first-order bound, since such a condition may remain satisfied at second order; when the condition does not move at all, however, all the coefficients $\xi^\top K^k \zeta$ of its remainder vanish, so that ξ is orthogonal to the Krylov subspace of Lemma 4 and a is zero.

C Instances

The instances come from the dataset of Derval (2025) and are summarized in Table 9. All of them define a matrix D and a range $[\underline{\lambda}, \overline{\lambda}]$ for λ:

- **Netlib**: the LPs of the Netlib collection (Gay 1985) that are solved within a few seconds; 100 randomly chosen equality constraints (instances `_eq`) or inequality constraints (`_ineq`) are perturbed, by changing up to three of their coefficients a_{ij} into $a_{ij}(1 \pm \lambda u_{ij})$ with u_{ij} uniform in $[0.3, 0.5]$, for $\lambda \in [-1, 1]$;
- **MIPLIB**: instances of MIPLIB 2017 (Gleixner et al. 2021), perturbed in the same way;
- **Facility**: the capacitated facility location instances of Holmberg et al. (1999), with all capacities multiplied by $\lambda \in [0.7, 1]$;
- **Unit commitment**: the unit commitment instances used by Frangioni et al. (2009), in which the coefficients of the power outputs of the units in the hourly load-balance constraints are perturbed by random relative amounts of at most ±30%, for $\lambda \in [-1, 1]$;
- **Real models**: 126 (model, sensitivity) pairs built from public models of hub location (Ernst and Krishnamoorthy 1999), index tracking (Beasley et al. 2003), unit commitment (Knueven et al. 2020), network-constrained unit commitment (Brown et al. 2018), telecommunication network design (Orlowski et al. 2010), retail facility location (Feng et al. 2022) and batch process scheduling (Kondili et al. 1993), where D describes a physically meaningful change of the model (a capacity transferred between hubs or reactors, a fuel price, a reaction selectivity, a line reactance, ...), for $\lambda \in [-1, 1]$ (or $[-0.9, 0.9]$).

Except for Netlib, the models are mixed-integer, and we consider their LP relaxations. The methods of this paper apply to these relaxations, and branch and bound solves such a relaxation at each node. Table 9 also reports n_D, the number of nonzero columns of D_B for the optimal basis B of $\mathcal{P}(\lambda_0)$, and the dimension ρ of the low-rank matrix C of Section 3.

D Proofs

Proof of Lemma 1. The first claim is a classical result on the products $\hat{U}V^\top$ and $V^\top\hat{U}$ of two matrices $\hat{U}, V \in \mathbb{R}^{m\times\rho}$ (Flanders 1951). The second claim follows from Theorem 1 and $\det(I_\rho + \lambda C) = \prod_\mu (1+\lambda\mu)$,

the product running over the eigenvalues of C counted with multiplicities. □

Proof of Lemma 2. A matrix is diagonalizable if and only if each of its eigenvalues has equal algebraic and geometric multiplicities. By Lemma 1, this holds for the nonzero eigenvalues of E if and only if it holds for those of C. With a minimal factorization, $\hat{U}, V \in \mathbb{R}^{m\times\rho}$ have full column rank ρ. As $Ex = \hat{U}(V^\top x)$ and the columns of $\hat{U}$ are linearly independent, $Ex = 0$ if and only if $V^\top x = 0$, so that the geometric multiplicity of the eigenvalue 0 of E is $\dim\ker(E) = \dim\ker(V^\top) = m - \rho$. Its algebraic multiplicity is $m - \rho$ plus that of 0 in C (the algebraic multiplicities add up to m in E and to ρ in C, and coincide for the nonzero eigenvalues by Lemma 1). The two multiplicities of 0 in E are thus equal if and only if C is nonsingular, and then all the eigenvalues of C are nonzero. □

Proof of Lemma 3. Section 4.1 derives $x_B(\lambda) = x_B(0) + \lambda\dot{x}_B - \lambda^2 EW\dot{x}_B$, with $W = (I_m + \lambda E)^{-1}$. The identity $(I_m+\lambda\hat{U}V^\top)^{-1}\hat{U} = \hat{U}(I_\rho+\lambda C)^{-1}$ gives $EW = \hat{U}(I_\rho+\lambda C)^{-1}V^\top$, hence (12), and multiplying it by $c_B^\top$ gives (14). For the reduced costs, $A_B + \lambda D_B = A_B(I_m + \lambda E)$ gives $r(\lambda) = c_N^\top - c_B^\top W A_B^{-1}(A_N + \lambda D_N)$. Replacing W by $I_m - \lambda WE$ where it multiplies $A_B^{-1}A_N$ gives $r(\lambda) = r(0) + \lambda c_B^\top WG$, and replacing W by $I_m - \lambda EW$ then gives $r(\lambda) = r(0) + \lambda\dot{r} - \lambda^2 c_B^\top EWG$. With $c_B^\top\hat{U} = w_c^\top$, the last term is $-\lambda^2 w_c^\top(I_\rho + \lambda C)^{-1}V^\top G e_j = -\lambda^2\theta_j^\top(I_\rho + \lambda C^\top)^{-1}w_c$ for each $j \in N$. □

Proof of Lemma 4. The Krylov subspace $\mathcal{K} = \mathcal{K}_\rho(K, \zeta)$ contains ζ and is invariant under K: by the Cayley–Hamilton theorem, $K^\rho\zeta$ is a linear combination of $\zeta, K\zeta, \ldots, K^{\rho-1}\zeta$. The matrix $QQ^\top$ is the orthogonal projector onto $\mathcal{K}$: it leaves the vectors of $\mathcal{K}$ unchanged, but it is not I_ρ when $d < \rho$. As $\mathcal{K}$ is invariant under K, the columns of KQ belong to $\mathcal{K}$, so that $KQ = QQ^\top KQ = QH$ and $(I_\rho + \lambda K)Q = Q(I_d + \lambda H)$, where $I_d + \lambda H$ is invertible since $|\lambda|\|\Omega^{-1}H\Omega\|_\infty < 1$. As $\zeta \in \mathcal{K}$, $\zeta = QQ^\top\zeta$, so that $(I_\rho + \lambda K)^{-1}\zeta = Q(I_d + \lambda H)^{-1}Q^\top\zeta$, and

$$\xi^\top(I_\rho + \lambda K)^{-1}\zeta = (\Omega Q^\top\xi)^\top\,(I_d + \lambda\,\Omega^{-1}H\Omega)^{-1}\,(\Omega^{-1}Q^\top\zeta).$$

The bound follows from Hölder's inequality and the Neumann bound of Section 4.2. □

Proof of Lemma 5. The function $t \mapsto t^2/(1 - \gamma t)$ is convex on $[0, 1/\gamma)$, so that f is concave there and, as $f(0) = \delta \geq 0$, the set of the t such that $f(t) \geq 0$ is an interval that contains 0. In the first case, $f(t) < 0$ for all $t > 0$, and in the second one f is identically zero. Otherwise, f is positive on some interval $(0, \varepsilon)$, and the interval ends at the first root of f in $(0, 1/\gamma)$, or at $1/\gamma$ if there is none. As $1 - \gamma t > 0$ on $[0, 1/\gamma)$, these roots are those of $q(t) = (1 - \gamma t)f(t)$. □