Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Finite-block step-size design for directional discrete-gradient method

Jiabao Yang

g2558001@stu.musashino-u.ac.jp

Musashino University

Research Article

Keywords:

Posted Date: November 26th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11696758/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Finite-block step-size design for directional discrete-gradient methods

Jiabao Yang[1*]

[1]Department of Mathematical Engineering, Graduate School of Engineering, Musashino University, Tokyo, Japan.

Corresponding author(s). E-mail(s): g2558001@stu.musashino-u.ac.jp;

Abstract

Directional updates in the Itoh–Abe method preserve objective decrease for positive step sizes. We study step-size selection over prescribed finite blocks. For positive definite quadratic models, an ordered matrix product characterises the worst-case energy-norm amplification, and strict contraction is equivalent to the block directions spanning the space. Independent-stage design admits an exact finite construction in two dimensions: endpoint projections reduce the optimisation to a scalar multi-affine function, and vertex evaluation with edge interpolation constructs an optimum. The construction extends to higher dimensions under rank-one endpoint compression. Perturbation bounds, grid estimates, and local nonlinear analysis clarify the use of approximate models. Applying existing nonsmooth convergence theory to the chronological stage sequence yields Clarke stationarity of accumulation points under the stated regularity assumptions, cyclic direction density, and uniform positive lower and finite upper step bounds. Two- and three-dimensional benchmarks evaluate the matrix-based designs and their construction costs. Separate nonconvex experiments compare constant, equal, and prescribed Chebyshev-root steps. Equal steps perform best on the tested Rosenbrock cases, while the advantage of the Chebyshev profile on a nonsmooth example depends on the initial point.

Keywords: Discrete gradient, Itoh–Abe method, Step-size selection, Relaxed projections, Numerical optimization

MSC Classification: 65K05 , 65L20 , 90C26

1 Introduction

We consider the unconstrained optimisation problem

$$\min_{\boldsymbol{x}\in\mathbb{R}^n} V(\boldsymbol{x}), \tag{1.1}$$

where $V : \mathbb{R}^n \to \mathbb{R}$ is locally Lipschitz continuous, bounded below, and coercive. Neither differentiability nor convexity is assumed. Additional regularity assumptions are stated only for the results that require them.

Riis et al. [1] developed this approach for locally Lipschitz, nonsmooth, and nonconvex objectives. Their randomised Itoh–Abe method uses objective-value differences to determine updates along search directions. Under suitable assumptions on the directions and step sizes, accumulation points are Clarke stationary. The deterministic directional framework also permits prescribed direction sequences satisfying appropriate coverage conditions. For smooth objectives, discrete-gradient optimization admits solvability results, convergence estimates, and practical step-size prescriptions [2].

These results provide the foundation for the present study. The update formula already preserves objective decrease, so our question concerns the selection of steps within that formula. A single directional update acts only along one direction. Subsequent updates therefore act on errors modified by the preceding stages. A step that is effective for an isolated direction need not be optimal for the complete block.

1.1 Discrete-gradient method

We consider a block formulation of the directional updates underlying the randomised Itoh–Abe method [1].

Let $\mathbb{S}^{n-1} := \{\boldsymbol{x} \in \mathbb{R}^n \mid \|\boldsymbol{x}\|_2 = 1\}$ denote the unit sphere, and let $s \in \mathbb{N}$ be the number of directional updates in each block. For each block k, choose coefficients $a_{k,1}, \ldots, a_{k,s}$ satisfying

$$a_{k,i} > 0 \qquad (i = 1, \ldots, s), \qquad \sum_{i=1}^{s} a_{k,i} = 1. \tag{1.2}$$

For a positive scale h_k, define the stage steps by

$$\tau_{k,i} := h_k a_{k,i}, \qquad k \geq 0, \quad i = 1, \ldots, s. \tag{1.3}$$

The coefficients may be redesigned between blocks. A fixed profile is recovered by taking $a_{k,i} = a_i$. The iteration is defined for positive stage steps. Where required for convergence, we additionally assume that there exist constants $0 < \tau_{\min} \leq \tau_{\max} < \infty$ such that

$$\tau_{\min} \leq \tau_{k,i} \leq \tau_{\max} \qquad (k \geq 0,\ i = 1, \ldots, s). \tag{1.4}$$

These bounds are imposed on the actual stage steps $\tau_{k,i}$.

At the beginning of block k, set $\boldsymbol{g}_0^k = \boldsymbol{x}^k$. Each internal stage is an Itoh–Abe-type update with step size $\tau_{k,i}$ and direction $\boldsymbol{d}_i^k \in \mathbb{S}^{n-1}$. If the current point is Clarke

directionally stationary with respect to this direction (Definition 2.4), no update is performed. Otherwise, a nonzero solution of the scalar discrete-gradient equation is selected. After all stages have been completed, the final internal point becomes the next outer iterate. The procedure is summarised in Algorithm 1.

Algorithm 1 (s-stage Itoh–Abe-type method)

Input: Initial point $\boldsymbol{x}^0$; coefficients $(a_{k,i})$ satisfying (1.2); positive scales (h_k); unit directions $(\boldsymbol{d}_i^k)$; a stopping criterion.

$k \leftarrow 0$
while the stopping criterion is not satisfied **do**
  $\boldsymbol{g}_0^k \leftarrow \boldsymbol{x}^k$
  for $i = 1, \dots, s$ **do**
    $\tau_{k,i} \leftarrow h_k a_{k,i}$
    if V is Clarke directionally stationary at $\boldsymbol{g}_{i-1}^k$ with respect to $\boldsymbol{d}_i^k$ **then**
      $\beta_i^k \leftarrow 0$
    else
      Choose any nonzero solution β_i^k satisfying

$$\beta_i^k = -\frac{V\left(\boldsymbol{g}_{i-1}^k + \tau_{k,i}\beta_i^k \boldsymbol{d}_i^k\right) - V(\boldsymbol{g}_{i-1}^k)}{\tau_{k,i}\beta_i^k}.$$

    end if
    $\boldsymbol{g}_i^k \leftarrow \boldsymbol{g}_{i-1}^k + \tau_{k,i}\beta_i^k \boldsymbol{d}_i^k$
  end for
  $\boldsymbol{x}^{k+1} \leftarrow \boldsymbol{g}_s^k$
  $k \leftarrow k + 1$
end while

Output: Final iterate $\boldsymbol{x}^k$.

1.2 Related work

Discrete-gradient methods provide a framework for constructing optimisation algorithms that preserve objective decrease. Riis et al. [1] developed the randomised Itoh–Abe method for locally Lipschitz, nonconvex, and nonsmooth objectives and established stationarity results under suitable assumptions. For smooth objectives, solvability, convergence estimates, and step-size prescriptions have been studied in [2]. These results form the basis of the directional updates used in the present paper. Our interest is in the choice of positive step sizes for several successive updates and their combined effect over a finite block.

The use of stage-dependent step sizes is also motivated by stabilised methods for smooth optimisation. Eftekhari et al. [3] developed the Runge–Kutta Chebyshev descent method, and Ushiyama et al. [4] investigated related constructions based on stable explicit numerical methods. On quadratic objectives, these methods can be analysed through scalar stability functions associated with Hessian eigenvalues. For successive directional Itoh–Abe updates, however, the amplification generally depends on an ordered product of noncommuting matrices. Accordingly, step-size schedules

motivated by scalar stability functions require a separate analysis when applied to directional updates.

For quadratic objectives and coordinate directions, the Itoh–Abe method is equivalent to successive overrelaxation under a change of parameters [5]. Adaptive relaxation parameters based on line-search conditions have also been studied [6]. In energy coordinates, the directional update matrices are relaxed orthogonal projections, connecting their analysis to alternating projection and subspace-correction methods [7, 8]. Scheduled step selection has also been investigated for randomised Kaczmarz methods using expected-error criteria [9]. In the present study, the direction sequence and block length are prescribed, and the steps are chosen to minimise the worst-case energy-norm amplification of the finite matrix product on a positive definite quadratic model.

Weak discrete gradients offer a related framework for transferring continuous-time dissipation and convergence arguments to discrete optimisation methods [10, 11]. That framework examines gradient approximations and the properties of the resulting schemes. In their preprint *Finite Horizon Optimization: Framework and Applications* (2024, https://arxiv.org/abs/2412.21068), Yushun Zhang, Dmitry Rybin, and Zhi-Quan Luo study finite-horizon step-size design for primal–dual methods in linear programming by controlling singular values of products of nonsymmetric matrices. The present work exploits the structure of directional discrete-gradient factors to obtain endpoint projections and an exact scalar multi-affine reduction.

1.3 Contributions and organization

We study step-size selection within blocks of directional Itoh–Abe updates. For positive definite quadratic models, we characterise the worst-case energy-norm amplification by an ordered matrix product. Strict contraction is equivalent to the block directions spanning the space. Independent-stage design admits an exact finite construction in two dimensions: endpoint projections reduce the problem to a scalar multi-affine function, whose vertex values and interpolation along one edge determine an optimum. The construction extends to higher dimensions under a rank-one endpoint compression condition. Perturbation and grid estimates quantify design errors, and local linearisation connects the quadratic product to nonlinear iteration.

For locally Lipschitz objectives that are bounded below and coercive, we establish dissipation and basic sequence properties. Applying the deterministic Itoh–Abe framework to the chronological stage sequence yields Clarke stationarity of accumulation points under cyclic direction density and uniform positive lower and finite upper bounds on the actual stage steps.

Two- and three-dimensional benchmarks compare matrix-based designs with existing step rules, accounting for schedule-construction costs. Separate smooth and nonsmooth nonconvex experiments compare a prescribed Chebyshev-root profile with constant and equal steps under common counting and stopping rules. These controls distinguish step allocation from a reduction in average step size. The results show problem- and initial-point-dependent benefits; they do not establish optimality of the profile or test the matrix-based design.

The remainder of the paper is organised as follows. Section 2 introduces nonsmooth stationarity and the existing convergence framework. Section 3 defines the

quadratic amplification criterion and admissible step-size designs. Section 4 presents the main design and convergence results. Section 5 reports the numerical experiments, and Section 6 discusses their theoretical and computational implications. Section 7 concludes the paper.

Proofs and supporting results are collected in Appendices A–E. Appendix F provides the step sequences used in the exact-arithmetic certificates, and Appendix G records the nonconvex experiment coefficients and additional diagnostics.

2 Preliminaries on nonsmooth analysis

Because the objective functions considered in this paper need not be differentiable or convex, we first introduce the Clarke subdifferential and Clarke stationary points as notions of stationarity from nonsmooth analysis. We then define Clarke directional stationarity and cyclically dense sequences of directions, which are needed for the convergence analysis, and finally state the deterministic convergence theorem for the Itoh–Abe-type method of Riis et al. used in this paper.

For $\varepsilon > 0$ and $\boldsymbol{x} \in \mathbb{R}^n$, let $B_\varepsilon(\boldsymbol{x}) := \{\boldsymbol{y} \in \mathbb{R}^n \mid \|\boldsymbol{y} - \boldsymbol{x}\|_2 < \varepsilon\}$.

Definition 2.1 (Locally Lipschitz continuous function) A function $V : \mathbb{R}^n \to \mathbb{R}$ is said to be locally Lipschitz continuous if, for every $\boldsymbol{x} \in \mathbb{R}^n$, there exist $\varepsilon > 0$ and $L > 0$ such that

$$|V(\boldsymbol{y}) - V(\boldsymbol{z})| \le L \|\boldsymbol{y} - \boldsymbol{z}\|_2$$

for all $\boldsymbol{y}, \boldsymbol{z} \in B_\varepsilon(\boldsymbol{x})$.

Definition 2.2 (Clarke directional derivative) Let $V : \mathbb{R}^n \to \mathbb{R}$ be locally Lipschitz continuous in a neighborhood of $\boldsymbol{x} \in \mathbb{R}^n$. For $\boldsymbol{d} \in \mathbb{R}^n$, the Clarke directional derivative of V at $\boldsymbol{x}$ in the direction $\boldsymbol{d}$ is defined by

$$V^\circ(\boldsymbol{x}; \boldsymbol{d}) := \limsup_{\boldsymbol{y} \to \boldsymbol{x},\ \lambda \downarrow 0} \frac{V(\boldsymbol{y} + \lambda \boldsymbol{d}) - V(\boldsymbol{y})}{\lambda}.$$

Definition 2.3 (Clarke subdifferential) Let $V : \mathbb{R}^n \to \mathbb{R}$ be locally Lipschitz continuous. The Clarke subdifferential of V at $\boldsymbol{x} \in \mathbb{R}^n$ is defined by

$$\partial V(\boldsymbol{x}) := \left\{ \boldsymbol{\xi} \in \mathbb{R}^n \mid V^\circ(\boldsymbol{x}; \boldsymbol{d}) \ge \langle \boldsymbol{\xi}, \boldsymbol{d} \rangle \quad \text{for all } \boldsymbol{d} \in \mathbb{R}^n \right\}.$$

Definition 2.4 (Clarke directional stationarity) Let $\boldsymbol{d} \in \mathbb{S}^{n-1}$. We say that V is Clarke directionally stationary at $\boldsymbol{x}^* \in \mathbb{R}^n$ with respect to $\boldsymbol{d}$ if

$$\min\{V^\circ(\boldsymbol{x}^*; \boldsymbol{d}), V^\circ(\boldsymbol{x}^*; -\boldsymbol{d})\} \ge 0.$$

Definition 2.5 (Clarke stationary point) A point $\boldsymbol{x}^* \in \mathbb{R}^n$ is called a Clarke stationary point of V if $\boldsymbol{0} \in \partial V(\boldsymbol{x}^*)$.

Remark 2.6 The function V is Clarke directionally stationary at $\boldsymbol{x}^*$ with respect to every $\boldsymbol{d} \in \mathbb{S}^{n-1}$ if and only if $\boldsymbol{x}^*$ is a Clarke stationary point [1, 12].

Definition 2.7 (Coercivity) A function $V : \mathbb{R}^n \to \mathbb{R}$ is coercive if, for every $M \in \mathbb{R}$, the sublevel set

$$\{\boldsymbol{x} \in \mathbb{R}^n \mid V(\boldsymbol{x}) \leq M\}$$

is bounded.

Definition 2.8 (Cyclically dense sequence of directions) A sequence $(\boldsymbol{d}^m)_{m=0}^{\infty} \subset \mathbb{S}^{n-1}$ is said to be cyclically dense on $\mathbb{S}^{n-1}$ if, for every $\varepsilon > 0$, there exists $N \in \mathbb{N}$ such that

$$\mathbb{S}^{n-1} \subset \bigcup_{\ell=m}^{m+N-1} B_\varepsilon(\boldsymbol{d}^\ell)$$

for every $m \geq 0$.

Remark 2.9 Cyclic density means that every sufficiently long consecutive block of the direction sequence approximately covers the entire unit sphere $\mathbb{S}^{n-1}$. This condition is therefore used to detect Clarke directional stationarity not only along a particular direction but in all directions.

Remark 2.10 When $n = 2$, let $\sigma \in (0, 1)$ be irrational and define $\xi_m = \{m\sigma\}$ and

$$\boldsymbol{d}^m = \begin{pmatrix} \cos(2\pi\xi_m) \\ \sin(2\pi\xi_m) \end{pmatrix} \qquad (m = 0, 1, 2, \dots).$$

Then $(\boldsymbol{d}^m)_{m=0}^{\infty}$ is cyclically dense on S^1 [1].

We next state the deterministic convergence theorem for the Itoh–Abe-type method of Riis et al. Later, we show that each internal stage of the proposed multistage method corresponds to an Itoh–Abe update of the following form. By flattening the internal stages into a single iteration sequence, we clarify the connection with the existing convergence theory for Itoh–Abe-type methods.

Theorem 2.11 (Deterministic directional stationarity) [1, Theorem 3.7] *Let $V : \mathbb{R}^n \to \mathbb{R}$ be locally Lipschitz continuous, bounded below, and coercive. Let $(\widehat{\boldsymbol{d}}^m)$ be unit directions and let $0 < \tau_{\min} \leq \widehat{\tau}_m \leq \tau_{\max} < \infty$. At a Clarke directionally stationary point set $\widehat{\beta}_m = 0$; otherwise choose a nonzero solution of*

$$\widehat{\beta}_m = -\frac{V(\boldsymbol{y}^m + \widehat{\tau}_m \widehat{\beta}_m \widehat{\boldsymbol{d}}^m) - V(\boldsymbol{y}^m)}{\widehat{\tau}_m \widehat{\beta}_m}.$$

Define $\boldsymbol{y}^{m+1} = \boldsymbol{y}^m + \widehat{\tau}_m \widehat{\beta}_m \widehat{\boldsymbol{d}}^m$. If the directions are cyclically dense on $\mathbb{S}^{n-1}$, every accumulation point of $(\boldsymbol{y}^m)$ is Clarke stationary.

3 Finite-block step-size design

We now introduce a criterion for selecting the step sizes within a finite block of directional updates. Separately from the convergence analysis for nonsmooth objectives, we consider a positive definite quadratic model and evaluate the largest possible amplification of the initial error. Throughout the design problem, the quadratic model, the directions, and their order are prescribed; the stage step sizes are the design variables.

3.1 Quadratic blocks and energy-norm amplification

Consider the positive definite quadratic objective

$$V_H(\boldsymbol{x}) = V_* + \frac{1}{2}(\boldsymbol{x} - \boldsymbol{x}_*)^\mathsf{T} H(\boldsymbol{x} - \boldsymbol{x}_*), \qquad H = H^\mathsf{T} \succ 0, \tag{3.1}$$

where $\boldsymbol{x}_*$ is its unique minimizer and $V_* = V_H(\boldsymbol{x}_*)$. Define the energy norm and its induced matrix norm by

$$\|\boldsymbol{e}\|_H = (\boldsymbol{e}^\mathsf{T} H \boldsymbol{e})^{1/2}, \qquad \|P\|_H = \sup_{\boldsymbol{e}\neq 0} \frac{\|P\boldsymbol{e}\|_H}{\|\boldsymbol{e}\|_H}.$$

Fix an ordered sequence of unit directions $\mathcal{D} = (\boldsymbol{d}_1, \ldots, \boldsymbol{d}_s)$, with $\boldsymbol{d}_i \in \mathbb{S}^{n-1}$, and finite positive stage steps $\boldsymbol{\tau} = (\tau_1, \ldots, \tau_s)$. Within a single block, we suppress the outer iteration index and denote the internal points by $\boldsymbol{g}_0, \ldots, \boldsymbol{g}_s$. Set

$$\boldsymbol{e}_i = \boldsymbol{g}_i - \boldsymbol{x}_*, \qquad q_i = \boldsymbol{d}_i^\mathsf{T} H \boldsymbol{d}_i > 0.$$

For the quadratic objective (3.1), the scalar stage equation gives

$$\beta_i = -\frac{\boldsymbol{d}_i^\mathsf{T} H \boldsymbol{e}_{i-1}}{1 + \tau_i q_i/2}.$$

This expression also covers a directionally stationary stage, where $\beta_i = 0$. Consequently,

$$\boldsymbol{e}_i = M_i \boldsymbol{e}_{i-1}, \qquad M_i = I - \frac{\tau_i}{1 + \tau_i q_i/2} \boldsymbol{d}_i \boldsymbol{d}_i^\mathsf{T} H. \tag{3.2}$$

The complete block therefore satisfies

$$\boldsymbol{e}_s = P_s(H, \mathcal{D}; \boldsymbol{\tau}) \boldsymbol{e}_0, \qquad P_s(H, \mathcal{D}; \boldsymbol{\tau}) = M_s \cdots M_1.$$

The rightmost factor acts first, so the product retains the prescribed order of the directional updates.

Definition 3.1 (Block amplification) For a fixed model H, direction sequence $\mathcal{D}$, and step vector $\boldsymbol{\tau}$, the worst-case energy-norm amplification is

$$Q_H(\mathcal{D}; \boldsymbol{\tau}) = \|P_s(H, \mathcal{D}; \boldsymbol{\tau})\|_H = \sup_{\boldsymbol{e}_0 \neq 0} \frac{\|P_s \boldsymbol{e}_0\|_H}{\|\boldsymbol{e}_0\|_H}. \tag{3.3}$$

To express this quantity in Euclidean coordinates, introduce

$$r_i = \frac{1 - \tau_i q_i/2}{1 + \tau_i q_i/2}, \quad \boldsymbol{u}_i = \frac{H^{1/2} \boldsymbol{d}_i}{\sqrt{q_i}}, \quad A_i(r_i) = I - (1 - r_i) \boldsymbol{u}_i \boldsymbol{u}_i^\mathsf{T}. \tag{3.4}$$

Then

$$\|\boldsymbol{u}_i\|_2 = 1, \qquad -1 < r_i < 1, \qquad A_i(r_i) = H^{1/2} M_i H^{-1/2}.$$

Thus, in the energy coordinates $\boldsymbol{z}_i = H^{1/2}\boldsymbol{e}_i$, the update becomes

$$\boldsymbol{z}_i = A_i(r_i)\boldsymbol{z}_{i-1}.$$

Each factor multiplies the component along $\boldsymbol{u}_i$ by r_i and leaves its orthogonal complement unchanged. In particular, $r_i = 0$, equivalently $\tau_i = 2/q_i$, gives the orthogonal projection onto $\boldsymbol{u}_i^{\perp}$.

Define

$$S_i = A_i(r_i) \cdots A_1(r_1), \qquad S_0 = I.$$

Since

$$S_s = H^{1/2} P_s H^{-1/2},$$

the block amplification satisfies

$$Q_H(\mathcal{D}; \boldsymbol{\tau}) = \|S_s\|_2 = \sigma_{\max}(S_s). \tag{3.5}$$

Let $\boldsymbol{g}_s(\boldsymbol{x})$ denote the final point obtained by applying this fixed block from $\boldsymbol{g}_0 = \boldsymbol{x}$. The identity

$$V_H(\boldsymbol{x}) - V_* = \frac{1}{2} \|\boldsymbol{x} - \boldsymbol{x}_*\|_H^2$$

then yields

$$\sup_{\boldsymbol{x} \neq \boldsymbol{x}_*} \frac{V_H(\boldsymbol{g}_s(\boldsymbol{x})) - V_*}{V_H(\boldsymbol{x}) - V_*} = Q_H(\mathcal{D}; \boldsymbol{\tau})^2. \tag{3.6}$$

Hence Q_H measures the worst-case amplification of the error norm, whereas Q_H^2 measures the worst-case ratio of objective gaps.

Each factor $A_i(r_i)$ is nonexpansive, so $Q_H(\mathcal{D}; \boldsymbol{\tau}) \leq 1$. The necessary and sufficient condition for strict contraction is established in Theorem 4.1.

3.2 Admissible steps and the design problem

For the fixed quadratic model H and direction sequence $\mathcal{D}$, we select the stage steps by minimizing the block amplification. Given a nonempty compact admissible set $\Theta \subset (0, \infty)^s$, the design problem is

$$\min_{\boldsymbol{\tau} \in \Theta} Q_H(\mathcal{D}; \boldsymbol{\tau}). \tag{3.7}$$

The objective is continuous in $\boldsymbol{\tau}$, so a minimizer exists. We consider two admissible families, distinguished by whether the stage-step ratios are prescribed or optimized.

Definition 3.2 (Common-scale and independent-stage designs) A *common-scale design* uses a prescribed positive profile $\boldsymbol{a} = (a_1, \ldots, a_s)$, normalized by $\sum_{i=1}^s a_i = 1$. Its admissible set is

$$\Theta_{\mathrm{com}} = \{h\boldsymbol{a} \mid h \in [h_-, h_+]\}, \qquad 0 < h_- \leq h_+ < \infty.$$

Thus $\tau_i = ha_i$, and the optimization reduces to

$$\min_{h\in[h_-,h_+]} Q_H(\mathcal{D}; h\boldsymbol{a}).$$

The profile is fixed only within this design family; h determines the common scale.

An *independent-stage design* fixes $0 < \eta < 1$ and allows the parameters r_i in (3.4) to vary independently over $[-\eta, \eta]$. The inverse change of variables is

$$\tau_i(r_i) = \frac{2}{q_i}\frac{1-r_i}{1+r_i}, \qquad i = 1, \dots, s. \tag{3.8}$$

Its admissible step set is

$$\Theta_{\text{ind}} = \prod_{i=1}^{s} \left[\frac{2(1-\eta)}{q_i(1+\eta)}, \frac{2(1+\eta)}{q_i(1-\eta)}\right].$$

Equivalently, the design problem is

$$Q_s^*(H, \mathcal{D}; \eta) = \min_{\boldsymbol{r}\in[-\eta,\eta]^s} \mathcal{Q}(\boldsymbol{r}), \quad \mathcal{Q}(\boldsymbol{r}) \quad = \|A_s(r_s)\cdots A_1(r_1)\|_2\,. \tag{3.9}$$

Since the transformation (3.8) is strictly decreasing, the independent parameter bounds are equivalent to

$$\frac{2(1-\eta)}{q_i(1+\eta)} \le \tau_i \le \frac{2(1+\eta)}{q_i(1-\eta)}. \tag{3.10}$$

Both admissible families are nonempty compact subsets of $(0,\infty)^s$. In the independent-stage family, the bounds depend on H and $\mathcal{D}$ through q_i, and neither the sum nor the ratios of the stage steps are prescribed.

The independent design is consistent with the representation $\tau_i = ha_i$. Indeed, any positive step vector determines

$$h = \sum_{j=1}^{s} \tau_j, \qquad a_i = \frac{\tau_i}{h},$$

so that $a_i > 0$ and $\sum_{i=1}^{s} a_i = 1$. Here the scale and allocation coefficients are determined by the selected steps, rather than fixed in advance.

The independence of the stage parameters is essential to the exact construction developed below. For $s \ge 2$, setting $r_1 = r_s = 0$ while keeping all interior parameters unchanged remains feasible in $[-\eta, \eta]^s$. The endpoint factors then become orthogonal projections. Such independent endpoint choices need not be feasible in a common-scale family.

3.3 Why the ordered matrix product is needed

The parameter r_i describes the action of a single update along the transformed direction $\boldsymbol{u}_i$. The orthogonal complement of this direction is unchanged. Consequently, multiplying the scalar factors r_i does not, in general, give the amplification of the complete error vector, even when the update matrices commute. Moreover, the factors associated with different directions generally do not commute, so their order must also be retained.

Example 3.3 (Scalar factors and block amplification) Let

$$H = I_2, \qquad (\boldsymbol{d}_1, \boldsymbol{d}_2) = (\mathbf{e}_1, \mathbf{e}_2), \qquad \tau_1 = \tau_2 = \frac{2}{3},$$

where $\mathbf{e}_1, \mathbf{e}_2$ are the standard coordinate vectors. Then $r_1 = r_2 = 1/2$ and

$$M_1 = \begin{pmatrix} 1/2 & 0 \\ 0 & 1 \end{pmatrix}, \qquad M_2 = \begin{pmatrix} 1 & 0 \\ 0 & 1/2 \end{pmatrix}.$$

Hence

$$P_2 = M_2 M_1 = \frac{1}{2} I_2, \qquad Q_H = \frac{1}{2},$$

whereas $r_1 r_2 = 1/4$. Thus the scalar product does not represent the block amplification, even in this commuting example.

More generally, $A_i(r_i)$ has eigenvalue r_i along $\boldsymbol{u}_i$ and eigenvalue one on $\boldsymbol{u}_i^{\perp}$. Since M_i is similar to $A_i(r_i)$,

$$\det P_s = \prod_{i=1}^{s} \det M_i = \prod_{i=1}^{s} r_i.$$

The scalar product therefore determines the determinant, rather than the worst-case error amplification.

The spectral radius of the block matrix also need not equal its induced energy norm.

Example 3.4 (Spectral radius and finite-block amplification) Let

$$H = \begin{pmatrix} 2 & 1 \\ 1 & 2 \end{pmatrix}, \qquad (\boldsymbol{d}_1, \boldsymbol{d}_2) = (\mathbf{e}_1, \mathbf{e}_2), \qquad \tau_1 = \tau_2 = 1.$$

The stage matrices are

$$M_1 = \begin{pmatrix} 0 & -1/2 \\ 0 & 1 \end{pmatrix}, \qquad M_2 = \begin{pmatrix} 1 & 0 \\ -1/2 & 0 \end{pmatrix},$$

and therefore

$$P_2 = M_2 M_1 = \begin{pmatrix} 0 & -1/2 \\ 0 & 1/4 \end{pmatrix}.$$

Its spectral radius is $1/4$, whereas

$$Q_H = \|P_2\|_H = \frac{1}{2}.$$

The spectral radius describes the asymptotic behavior of powers of a fixed block matrix, but does not determine its worst-case amplification over a single block. For the finite-block criterion considered here, the relevant quantity is therefore

$$Q_H(\mathcal{D}; \boldsymbol{\tau}) = \sigma_{\max}(A_s(r_s) \cdots A_1(r_1)).$$

Table 1 summarizes the settings addressed by the subsequent results. In the independent-stage designs, $0 < \eta < 1$ is fixed and each parameter ranges independently over $[-\eta, \eta]$.

Table 1 Finite-block analysis and step-size design. The quadratic model, directions, and direction order are prescribed

Setting	Assumptions or design family	Result
Any dimension	Fixed positive finite steps	Strict contraction if and only if the directions span $\mathbb{R}^n$
Two dimensions, $s \geq 2$	Independent-stage design	Exact construction from 2^{s-2} signed vertex values and, when needed, interpolation along one edge
Structured $n \geq 3$, $s \geq 2(n-1)$	Independent-stage design; each endpoint group contains $n-1$ mutually H-conjugate directions	Rank-one endpoint compression and an exact scalar multi-affine reduction
General $n \geq 3$, $s \geq 2$	Independent-stage design with endpoint projections	A reduced matrix-norm optimization problem
Common scale	Prescribed positive profile $\tau_i = ha_i$	An error bound for logarithmic-grid search

The exact scalar reduction relies on rank-one endpoint compression. This structure is available in two dimensions and under the stated endpoint conditions in higher dimensions. For general higher-dimensional direction sequences, the reduced objective remains a matrix norm.

All design statements in this section concern the prescribed quadratic model. Their implications for smooth nonlinear iteration are treated separately. For nonsmooth objectives, Clarke stationarity of accumulation points requires the step-size and direction conditions of the convergence analysis.

4 Main results

We first present the quadratic amplification and step-size design results. The central result is an exact finite construction for independently bounded stage parameters in two dimensions, with an extension under rank-one endpoint compression. We also give perturbation estimates and a connection to local nonlinear convergence.

We then establish Clarke stationarity of accumulation points for the proposed block iteration using the deterministic convergence framework of Riis et al. The proofs are collected in Appendices A–E.

4.1 Quadratic amplification and finite-block step-size design

For the quadratic model and admissible sets introduced in Section 3, we first characterize strict contraction of a finite block. We then give an exact independent-stage design in two dimensions and extend it to higher dimensions under a rank-one endpoint compression condition. Finally, we quantify the effect of an approximate Hessian and connect the quadratic criterion to local nonlinear convergence.

4.1.1 Contraction of a quadratic block

Each energy-coordinate factor satisfies

$$\|A_i(r_i)\boldsymbol{z}\|_2^2 = \|\boldsymbol{z}\|_2^2 - (1-r_i^2)(\boldsymbol{u}_i^{\mathsf{T}}\boldsymbol{z})^2. \tag{4.1}$$

Summing these identities along the block gives the following characterization.

Theorem 4.1 (Exact contraction criterion) *Let $H = H^{\mathsf{T}} \succ 0$, let $\mathcal{D} = (\boldsymbol{d}_1, \dots, \boldsymbol{d}_s)$ be prescribed unit directions, and let $\boldsymbol{\tau} \in (0,\infty)^s$. Then*

$$\begin{aligned} &Q_H(\mathcal{D};\boldsymbol{\tau}) \le 1, \\ &Q_H(\mathcal{D};\boldsymbol{\tau}) < 1 \quad\Longleftrightarrow\quad \operatorname{span}\{\boldsymbol{d}_1,\dots,\boldsymbol{d}_s\} = \mathbb{R}^n. \end{aligned} \tag{4.2}$$

Moreover, define

$$G_s = \sum_{i=1}^{s} (1-r_i^2) S_{i-1}^{\mathsf{T}} \boldsymbol{u}_i \boldsymbol{u}_i^{\mathsf{T}} S_{i-1}. \tag{4.3}$$

Then

$$\begin{aligned} S_s^{\mathsf{T}} S_s &= I - G_s, \qquad 0 \preceq G_s \preceq I, \\ Q_H(\mathcal{D};\boldsymbol{\tau})^2 &= 1 - \lambda_{\min}(G_s). \end{aligned} \tag{4.4}$$

The proof and a triangular representation of G_s are given in Appendix B.2. The theorem expresses the contraction mechanism of relaxed projection products in the variables used for step-size design [8]. Spanning determines whether strict contraction occurs, whereas G_s also records the effects of the step sizes, direction correlations, and update order. In particular, a block with $s < n$ leaves some nonzero error unchanged and therefore has amplification equal to one.

4.1.2 Exact independent-stage design in two dimensions

Setting $r_i = 0$ gives the orthogonal projection

$$\Pi_i = I - \boldsymbol{u}_i \boldsymbol{u}_i^{\mathsf{T}}.$$

For $s \ge 2$, replacing the first and last factors by Π_1 and Π_s cannot increase the block norm. This replacement is feasible in the independent-stage design because zero belongs to each parameter interval.

In two dimensions, these projections have rank one. Define

$$J = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}, \qquad \boldsymbol{v}_i = J\boldsymbol{u}_i.$$

Then $\Pi_i = \boldsymbol{v}_i \boldsymbol{v}_i^{\mathsf{T}}$. With $r_1 = r_s = 0$, the block product becomes

$$\begin{aligned} S_s &= g(\boldsymbol{r}_{\mathrm{mid}}) \boldsymbol{v}_s \boldsymbol{v}_1^{\mathsf{T}}, \\ g(\boldsymbol{r}_{\mathrm{mid}}) &= \boldsymbol{v}_s^{\mathsf{T}} A_{s-1}(r_{s-1}) \cdots A_2(r_2) \boldsymbol{v}_1, \end{aligned} \tag{4.5}$$

where $\boldsymbol{r}_{\mathrm{mid}} = (r_2, \ldots, r_{s-1})$. Since the endpoint vectors are unit vectors, $\|S_s\|_2 = |g(\boldsymbol{r}_{\mathrm{mid}})|$. Each middle parameter occurs in exactly one affine factor, so g is multi-affine.

Theorem 4.2 (Exact two-dimensional design) *Let $H = H^{\mathsf{T}} \succ 0$, let $\boldsymbol{d}_1, \ldots, \boldsymbol{d}_s$ be prescribed unit directions in $\mathbb{R}^2$, with $s \geq 2$, and let $0 < \eta < 1$. Define*

$$g_- = \min_{\boldsymbol{r}_{\mathrm{mid}} \in \{-\eta, \eta\}^{s-2}} g(\boldsymbol{r}_{\mathrm{mid}}), \qquad g_+ = \max_{\boldsymbol{r}_{\mathrm{mid}} \in \{-\eta, \eta\}^{s-2}} g(\boldsymbol{r}_{\mathrm{mid}}).$$

Then

$$Q_s^*(H, \mathcal{D}; \eta) = \begin{cases} 0, & g_- \leq 0 \leq g_+, \\ g_-, & g_- > 0, \\ -g_+, & g_+ < 0. \end{cases} \tag{4.6}$$

A minimizer can be chosen with $r_1 = r_s = 0$ and at most one middle parameter strictly inside its admissible interval. For $s = 2$, the middle parameter set consists of one point, and the empty product in (4.5) is the identity.

The proof is given in Appendix C. If all vertex values have the same strict sign, a vertex with the smallest absolute value is optimal. If a vertex value is zero, that vertex is optimal. Otherwise, a path of cube edges between vertices of opposite signs contains an edge whose endpoint values have opposite signs. Affine interpolation along that edge gives a zero.

Algorithm 2 (Independent-stage design in two dimensions)

Input: $H = H^{\mathsf{T}} \succ 0$; unit directions $\boldsymbol{d}_1, \ldots, \boldsymbol{d}_s$; $s \geq 2$; $0 < \eta < 1$.

Compute q_i, $\boldsymbol{u}_i$, and $\boldsymbol{v}_i$ for $i = 1, \ldots, s$.
Set $r_1 = r_s = 0$.
Evaluate g at all vertices of $\{-\eta, \eta\}^{s-2}$.
if a vertex value is zero **then**
  Set $\boldsymbol{r}_{\mathrm{mid}}$ to a vertex with value zero.
else if all vertex values have the same strict sign **then**
  Set $\boldsymbol{r}_{\mathrm{mid}}$ to a vertex with the smallest absolute value.
else
  Connect a negative vertex to a positive vertex by a path of cube edges.
  Select adjacent vertices $\boldsymbol{r}^-$ and $\boldsymbol{r}^+$ satisfying $g(\boldsymbol{r}^-) < 0 < g(\boldsymbol{r}^+)$.
  Set

$$t_* = -\frac{g(\boldsymbol{r}^-)}{g(\boldsymbol{r}^+) - g(\boldsymbol{r}^-)}.$$

  Set

$$\boldsymbol{r}_{\mathrm{mid}} = (1 - t_*)\boldsymbol{r}^- + t_* \boldsymbol{r}^+.$$

end if
Compute $\tau_i = \dfrac{2}{q_i} \dfrac{1 - r_i}{1 + r_i}$ for $i = 1, \ldots, s$.
Evaluate $Q_H(\mathcal{D}; \boldsymbol{\tau})$.

Output: The step vector $\boldsymbol{\tau}$ and its block amplification.

Interpolation is performed between adjacent vertices because g is affine along an edge. Direct enumeration requires $O(s2^{s-2})$ small matrix-vector operations. The construction is therefore practical for short blocks, although its cost grows exponentially with the number of stages. The theorem concerns exact arithmetic; numerical implementation requires particular care when determining the signs of vertex values close to zero.

Example 4.3 (Two and three stages) For two stages, define

$$c = \boldsymbol{u}_1^{\mathsf{T}} \boldsymbol{u}_2 = \frac{\boldsymbol{d}_1^{\mathsf{T}} H \boldsymbol{d}_2}{\sqrt{q_1 q_2}}.$$

Then

$$Q_2^*(H, \mathcal{D}; \eta) = |c|,$$

attained at $r_1 = r_2 = 0$. For three stages, write $c_{ij} = \boldsymbol{u}_i^{\mathsf{T}} \boldsymbol{u}_j$. Then

$$\begin{aligned} g(r_2) &= c_{32}c_{21} + (c_{31} - c_{32}c_{21})r_2, \\ Q_3^*(H, \mathcal{D}; \eta) &= \max\left\{0,\ |c_{32}c_{21}| - \eta|c_{31} - c_{32}c_{21}|\right\}. \end{aligned} \tag{4.7}$$

For

$$H = \begin{pmatrix} 2 & 1 \\ 1 & 2 \end{pmatrix}, \qquad (\boldsymbol{d}_1, \boldsymbol{d}_2, \boldsymbol{d}_3) = (\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_1),$$

we obtain

$$g(r_2) = \frac{1}{4} + \frac{3}{4} r_2.$$

If $\eta \geq 1/3$, the choice $(r_1, r_2, r_3) = (0, -1/3, 0)$ is feasible. The corresponding steps are $(1, 2, 1)$, and the block matrix satisfies $P_3 = 0$.

The middle step in this example is selected for its effect on the complete block. Thus blockwise design can yield exact error elimination on the quadratic model, even though each individual stage leaves an orthogonal complement unchanged.

4.1.3 Rank-one endpoint compression in higher dimensions

The same scalar reduction applies when groups of endpoint directions compress the surviving error to a one-dimensional subspace. Two directions are called H-conjugate if $\boldsymbol{d}_i^{\mathsf{T}} H \boldsymbol{d}_j = 0$.

Theorem 4.4 (Rank-one endpoint reduction) *Let $n \geq 2$, $s \geq 2(n-1)$, $H = H^{\mathsf{T}} \succ 0$, and $0 < \eta < 1$. Suppose that the first $n-1$ directions are mutually H-conjugate and that the last $n-1$ directions are mutually H-conjugate within their own group. Then (3.9) has a minimizer with all parameters in these endpoint groups equal to zero.*

Put $m = n - 1$ and $p = s - 2m$. Choose unit vectors $\boldsymbol{a}, \boldsymbol{b}$ such that

$$I - \sum_{i=1}^{m} \boldsymbol{u}_i \boldsymbol{u}_i^{\mathsf{T}} = \boldsymbol{a}\boldsymbol{a}^{\mathsf{T}}, \qquad I - \sum_{i=s-m+1}^{s} \boldsymbol{u}_i \boldsymbol{u}_i^{\mathsf{T}} = \boldsymbol{b}\boldsymbol{b}^{\mathsf{T}}.$$

Define

$$g(\boldsymbol{r}_{\mathrm{mid}}) = \boldsymbol{b}^{\mathsf{T}} A_{s-m}(r_{s-m}) \cdots A_{m+1}(r_{m+1}) \boldsymbol{a}. \tag{4.8}$$

Let g_- and g_+ denote its minimum and maximum over the 2^p vertices of $[-\eta,\eta]^p$. Then

$$Q_s^*(H,\mathcal{D};\eta) = \operatorname{dist}(0,[g_-,g_+]). \tag{4.9}$$

A minimizer can be constructed from these vertex values and, when needed, interpolation along one edge. At most one middle parameter needs to lie strictly inside its interval. For $p=0$, the middle product is the identity and the parameter set consists of one point.

The proof in Appendix C uses endpoint compression followed by the scalar multi-affine argument. Conjugacy is required within each endpoint group, but not between the two groups. For $n = 2$, each group contains one direction, so the condition is automatic. For higher dimensions, it imposes a specific structure. For general direction sequences, endpoint projection leaves a matrix-valued reduced problem; see Section 6.2.

4.1.4 Design with an approximate Hessian

In a nonlinear implementation, the steps may be designed using an available positive definite model B. Let $H = H^{\mathsf{T}} \succ 0$ be the target matrix and $B = B^{\mathsf{T}} \succ 0$ the model. We compare their amplifications for the same direction sequence and the same actual step vector. Define

$$T = B^{-1/2} H B^{-1/2}, \qquad \delta = \|T - I\|_2 < 1. \tag{4.10}$$

No commutativity between H and B is assumed.

Theorem 4.5 (Uniform Hessian perturbation bound) *Under (4.10), for every prescribed unit direction sequence $\mathcal{D}$ and every $\boldsymbol{\tau} \in (0,\infty)^s$,*

$$\begin{aligned} |Q_H(\mathcal{D};\boldsymbol{\tau}) - Q_B(\mathcal{D};\boldsymbol{\tau})| &\le E_s(\delta), \\ E_s(\delta) &:= \min\left\{1,\ \frac{3s\delta}{2(1-\delta)}\right\}. \end{aligned} \tag{4.11}$$

The bound is uniform over all finite positive stage steps. It therefore applies both to a prescribed schedule and to one obtained by approximate optimization.

Corollary 4.6 (Transfer of an approximate design) *Under the assumptions of Theorem 4.5, let $\Theta \subset (0,\infty)^s$ be the same nonempty compact admissible step set for both models. Suppose that $\widehat{\boldsymbol{\tau}} \in \Theta$ satisfies*

$$Q_B(\mathcal{D};\widehat{\boldsymbol{\tau}}) \le \min_{\boldsymbol{\tau}\in\Theta} Q_B(\mathcal{D};\boldsymbol{\tau}) + \varepsilon_{\mathrm{opt}}, \qquad \varepsilon_{\mathrm{opt}} \ge 0.$$

Then

$$\begin{aligned} Q_H(\mathcal{D};\widehat{\boldsymbol{\tau}}) \le{}& \min_{\boldsymbol{\tau}\in\Theta} Q_H(\mathcal{D};\boldsymbol{\tau}) \\ &+ 2E_s(\delta) + \varepsilon_{\mathrm{opt}}. \end{aligned} \tag{4.12}$$

Furthermore, if a baseline $\boldsymbol{\tau}^{\mathrm{base}} \in \Theta$ and $\gamma > 0$ satisfy

$$Q_B(\mathcal{D};\widehat{\boldsymbol{\tau}}) \le Q_B(\mathcal{D};\boldsymbol{\tau}^{\mathrm{base}}) - \gamma,$$

then

$$\begin{aligned} Q_H(\mathcal{D};\widehat{\boldsymbol{\tau}}) \le Q_H(\mathcal{D};\boldsymbol{\tau}^{\mathrm{base}}) \\ -\gamma + 2E_s(\delta). \end{aligned} \tag{4.13}$$

The proofs are given in Appendix D. In particular, an improvement $\gamma > 2E_s(\delta)$ on the model B guarantees an improvement on H for the same step vectors.

The common admissible step set in Corollary 4.6 is essential. Transforming the same r-box separately using B and H generally produces different step sets, to which the stated optimal-value comparison does not directly apply.

4.1.5 Connection to local nonlinear convergence

We finally relate the quadratic block product to the derivative of the nonlinear block map. The assumptions in this subsection are local and do not impose global smoothness or convexity on the objective.

Suppose that V is twice continuously differentiable near $\boldsymbol{x}_*$ and that

$$\nabla V(\boldsymbol{x}_*) = \mathbf{0}, \qquad H_* := \nabla^2 V(\boldsymbol{x}_*) \succ 0.$$

Fix unit directions $\mathcal{D} = (\boldsymbol{d}_1, \dots, \boldsymbol{d}_s)$ and finite positive steps $\boldsymbol{\tau} = (\tau_1, \dots, \tau_s)$.

For stage i, write the smooth scalar equation as

$$F_i(\boldsymbol{x}, \beta) := \beta + \int_0^1 \boldsymbol{d}_i^{\mathsf{T}} \nabla V(\boldsymbol{x} + t\tau_i \beta \boldsymbol{d}_i)\, \mathrm{d}t = 0. \tag{4.14}$$

At each stage, use the local solution branch through $(\boldsymbol{x}_*, 0)$. Let T_{blk} denote the composition of the resulting stage maps.

Theorem 4.7 (Local block contraction) *The stage maps are continuously differentiable near $\boldsymbol{x}_*$, and*

$$DT_{\mathrm{blk}}(\boldsymbol{x}_*) = P_s(H_*, \mathcal{D}; \boldsymbol{\tau}).$$

Consequently,

$$\begin{aligned} T_{\mathrm{blk}}(\boldsymbol{x}_* + \boldsymbol{e}) - \boldsymbol{x}_* = P_s(H_*, \mathcal{D}; \boldsymbol{\tau})\boldsymbol{e} \\ + o(\|\boldsymbol{e}\|_2). \end{aligned} \tag{4.15}$$

If the directions span $\mathbb{R}^n$, then $Q_{H_*}(\mathcal{D};\boldsymbol{\tau}) < 1$. For every

$$q \in \big(Q_{H_*}(\mathcal{D};\boldsymbol{\tau}), 1\big),$$

a sufficiently small H_*-norm ball centered at $\boldsymbol{x}_*$ is invariant under T_{blk}, and

$$\|T_{\mathrm{blk}}(\boldsymbol{x}) - \boldsymbol{x}_*\|_{H_*} \le q\, \|\boldsymbol{x} - \boldsymbol{x}_*\|_{H_*} \tag{4.16}$$

holds throughout that ball. Repeated application of this fixed block therefore converges locally linearly to $\boldsymbol{x}_*$.

If $P_s(H_*, \mathcal{D}; \boldsymbol{\tau}) = 0$, the local error satisfies

$$\|T_{\mathrm{blk}}(\boldsymbol{x}) - \boldsymbol{x}_*\|_{H_*} = o\Big(\|\boldsymbol{x} - \boldsymbol{x}_*\|_{H_*}\Big).$$

Thus the repeated block iteration either terminates at $\boldsymbol{x}_*$ or converges superlinearly.

The proof is given in Appendix E. The theorem identifies the quadratic block product with the derivative of the nonlinear iteration. A zero quadratic block matrix gives a superlinear local estimate, but does not by itself imply finite termination for a nonlinear objective.

A supporting smooth-convex convergence argument is also recorded in Appendix E. Its uniform direction-coverage condition is distinct from the cyclic-density assumption used for nonsmooth stationarity.

4.2 Stationarity of accumulation points for nonconvex and nonsmooth problems

We now consider Algorithm 1 without assuming differentiability or convexity of V. The convergence conditions are imposed directly on the actual stage steps $\tau_{k,i}$. In particular, when $\tau_{k,i} = h_k a_{k,i}$, the allocation coefficients may vary between blocks, provided that the resulting stage steps satisfy the bounds stated below.

The scalar stage equation and the update formula give

$$V(\boldsymbol{g}_{i-1}^k) - V(\boldsymbol{g}_i^k) = \tau_{k,i}(\beta_i^k)^2 = \frac{\left\|\boldsymbol{g}_i^k - \boldsymbol{g}_{i-1}^k\right\|_2^2}{\tau_{k,i}}. \tag{4.17}$$

At a directionally stationary stage, $\beta_i^k = 0$, so the same identity holds. Summing over the stages yields

$$V(\boldsymbol{x}^k) - V(\boldsymbol{x}^{k+1}) = \sum_{i=1}^{s} \frac{\left\|\boldsymbol{g}_i^k - \boldsymbol{g}_{i-1}^k\right\|_2^2}{\tau_{k,i}} \geq 0. \tag{4.18}$$

The following theorem establishes stationarity for the proposed block iteration by applying the deterministic Itoh–Abe convergence framework of Riis et al. [1] to the chronological stage sequence. The proof is given in Appendix A.

Theorem 4.8 (Clarke stationarity of accumulation points) *Let $V : \mathbb{R}^n \to \mathbb{R}$ be locally Lipschitz continuous, bounded below, and coercive. Consider Algorithm 1 with a fixed number of stages $s \geq 1$.*

At each internal stage, set $\beta_i^k = 0$ if V is Clarke directionally stationary at $\boldsymbol{g}_{i-1}^k$ with respect to $\boldsymbol{d}_i^k$. Otherwise, choose any nonzero solution β_i^k of the scalar stage equation.

Suppose that there exist constants $0 < \tau_{\min} \leq \tau_{\max} < \infty$ such that

$$\tau_{\min} \leq \tau_{k,i} \leq \tau_{\max} \qquad (k \geq 0,\ i = 1, \ldots, s). \tag{4.19}$$

Define the chronological direction sequence by

$$\widehat{\boldsymbol{d}}^{\,ks+i-1} := \boldsymbol{d}_i^k \qquad (k \geq 0,\ i = 1, \ldots, s).$$

If $(\widehat{\boldsymbol{d}}^{\,m})_{m=0}^{\infty}$ is cyclically dense on $\mathbb{S}^{n-1}$, then the iterate sequence $(\boldsymbol{x}^k)_{k=0}^{\infty}$ is bounded and has at least one accumulation point. Every accumulation point $\boldsymbol{x}^*$ satisfies

$$\mathbf{0} \in \partial V(\boldsymbol{x}^*).$$

The cyclic density assumption is distinct from the finite-block spanning condition used in the quadratic analysis. For $n \geq 2$, periodically repeating a finite set of directions does not produce a cyclically dense sequence on the sphere. Thus, finite-block spanning alone does not imply the nonsmooth convergence conclusion above.

The theorem establishes stationarity of all accumulation points, rather than convergence of the entire sequence to a single point. For a nonconvex objective, a Clarke stationary point need not be a minimizer.

5 Numerical experiments

The experiments address two questions. Quadratic and smooth strongly convex benchmarks compare the matrix-based designs with existing step rules, including design costs. Separate smooth and nonsmooth nonconvex examples compare constant steps, equal steps, and a prescribed Chebyshev-root profile within the same directional update. The second comparison examines the effect of a prescribed allocation; it does not test optimality of the matrix-based design.

5.1 Quadratic and smooth nonquadratic benchmarks

We consider

$$\begin{aligned} H_\rho &= (\rho^{|i-j|})_{i,j=1}^n, \\ V_{\rho,\varepsilon}(\boldsymbol{x}) &= \frac{1}{2}\boldsymbol{x}^\mathsf{T} H_\rho \boldsymbol{x} + \varepsilon \sum_{j=1}^n \log\cosh(x_j), \end{aligned} \tag{5.1}$$

where $0 \leq \rho < 1$, $\varepsilon \geq 0$, and $n = 2$ or 3. The leading principal determinants of H_ρ are 1, $1-\rho^2$, and, in three dimensions, $(1-\rho^2)^2$. Thus H_ρ is positive definite. Moreover,

$$\begin{aligned} \nabla^2 V_{\rho,\varepsilon}(\boldsymbol{x}) &= H_\rho + \varepsilon \operatorname{diag}\big(\operatorname{sech}^2(x_1), \ldots, \operatorname{sech}^2(x_n)\big), \\ H_\rho &\preceq \nabla^2 V_{\rho,\varepsilon}(\boldsymbol{x}) \preceq H_\rho + \varepsilon I. \end{aligned} \tag{5.2}$$

Consequently, these objectives are strongly convex and have globally Lipschitz continuous gradients. Their unique minimizer is $\boldsymbol{x}_* = \mathbf{0}$, with $V_* = 0$. For $\rho > 0$, the quadratic term couples all coordinates.

The five benchmark cases and their model-refresh policies are listed in Table 2.

Table 2 Benchmark cases and model-refresh policies for the designed schedules

Case	n	ρ	ε	Design model
Q2	2	0.95	0	Known H_ρ, fixed
N2	2	0.99	0.01	Current Hessian every block
Q3	3	0.8	0	Known H_ρ, fixed
N3	3	0.8	0.01	Initial Hessian, fixed
H3	3	0.999	0.001	Initial Hessian, fixed

For these benchmarks, let $\boldsymbol{y}^m$ denote the point after m individual direction updates, so that $\boldsymbol{y}^{ks+i} = \boldsymbol{g}_i^k$ for $i = 0, \dots, s$, with the shared block endpoints identified. We set

$$y_j^0 = (-1)^{j-1}, \qquad j = 1, \dots, n,$$

and use the same coordinate sequence for every method,

$$\widehat{\boldsymbol{d}}^{\,m} = \mathbf{e}_{1+(m \bmod n)}, \qquad m \geq 0.$$

Each design block contains nine direction updates. The coordinate sequence continues across block boundaries. In two dimensions, consecutive blocks therefore start with different coordinates.

The stopping condition is checked after every direction update. A run terminates when

$$\frac{V_{\rho,\varepsilon}(\boldsymbol{y}^m)}{V_{\rho,\varepsilon}(\boldsymbol{y}^0)} \leq 10^{-12}, \tag{5.3}$$

or after 200000 direction updates. For these benchmarks, iteration counts refer to individual direction updates, rather than coordinate sweeps or nine-stage blocks.

The coordinate sequence satisfies

$$\sum_{\ell=m}^{m+n-1} \widehat{\boldsymbol{d}}^{\,\ell} (\widehat{\boldsymbol{d}}^{\,\ell})^{\mathsf{T}} = I \qquad (m \geq 0).$$

It therefore provides uniform finite-window coverage, but is not cyclically dense on the sphere for $n \geq 2$. The nonsmooth cyclic-density theorem is not used to justify these coordinate-based benchmarks.

5.2 Comparison step-size prescriptions

For the benchmark family (5.1), define

$$\begin{aligned} L_i &= 1 + \varepsilon, \\ \overline{L}_i^2 &= (1+\varepsilon)^2 + \sum_{j \neq i} (H_\rho)_{ij}^2, \\ \overline{L}_\Sigma &= \left(\sum_{i=1}^n \overline{L}_i^2 \right)^{1/2}. \end{aligned} \tag{5.4}$$

Here L_i bounds the curvature along coordinate i, whereas $\overline{L}_i$ is a Lipschitz constant for the scalar partial derivative $\partial_i V$ with respect to the full variable.

We use three prescriptions from Ehrhardt et al. [2]:

$$\begin{aligned} \text{R–T:} \quad & \tau_i = \frac{1}{\overline{L}_\Sigma}, \\ \text{R–S:} \quad & \tau_i = \frac{2}{\sqrt{n}\, L_i}, \\ \text{R–C:} \quad & \tau_i = \frac{2}{L_i}. \end{aligned} \tag{5.5}$$

R–T uses the cyclic Itoh–Abe step in Table 1 of that paper. R–S follows the formula in the main text of Section 8.2.4. R–C is the practical coordinate prescription used in Section 8.1. On Q2 and Q3, R–C gives exact coordinate minimization. On the nonquadratic benchmarks, it gives a fixed coordinate step.

We compare these prescriptions with two designs based on the nine-update amplification criterion.

- **P–U: common-step design.** For a given model B and direction block $\mathcal{D}$, approximately minimize

$$Q_B(\mathcal{D}; \tau \mathbf{1}) \qquad \text{over } \tau \in [0.01, 100].$$

  This is the common-scale family with equal stage steps.
- **P–I: independent-stage design.** Optimize the middle parameters subject to $|r_i| \le 0.9$, with $r_1 = r_9 = 0$. In two dimensions, use Algorithm 2, which evaluates the 128 signed middle-vertex values and, when needed, performs edge interpolation. Theorem 4.2 guarantees global optimality for this box problem in exact arithmetic. In three dimensions, use a multistart numerical search for the remaining seven parameters. No global optimality is asserted for that search.

For Q2 and Q3, the known matrix H_ρ is used and the designed schedules are retained. In Q2, interchange of the two coordinates leaves H_ρ invariant, so the alternating block phase preserves the block norm. For N2, the Hessian and the schedule are recomputed at the beginning of every block. For N3 and H3, the initial Hessian is computed once and the resulting schedule is reused.

These policies distinguish repeated local design from reuse of a fixed schedule. Implementation details are given below; design and iteration costs are reported separately in Section 5.6.

5.3 Implementation of the matrix-based benchmarks

The iteration for the five main benchmarks is summarized in Algorithm 3. For a positive definite model B, we compute an upper Cholesky factor F satisfying $F^\mathsf{T} F = B$ and use

$$\boldsymbol{u}_i = \frac{F \boldsymbol{d}_i}{\sqrt{\boldsymbol{d}_i^\mathsf{T} B \boldsymbol{d}_i}}.$$

This gives the same block singular values as the principal-square-root construction and avoids explicitly forming $B^{1/2}$.

Algorithm 3 (Coordinate iteration with finite-block design)

Input: Initial point $\boldsymbol{y}^0$; coordinate sequence $(\widehat{\boldsymbol{d}}^{\,m})$; stage count s; model-refresh policy; stopping conditions.

Set $m \leftarrow 0$.
while neither stopping condition is satisfied **do**
  if this is the first block or the policy requests a new design **then**
    Form the positive definite model B and collect the next s directions.
    Compute and store a common-step or independent-stage schedule $(\tau_1, \ldots, \tau_s)$.
  end if
  for $i = 1, \ldots, s$ **do**
    Compute the directional update with step τ_i along $\widehat{\boldsymbol{d}}^{\,m}$, using (3.2) or (5.6).
    Store the new point as $\boldsymbol{y}^{m+1}$ and set $m \leftarrow m + 1$.
    if either stopping condition is satisfied **then**
      return $\boldsymbol{y}^m$.
    end if
  end for
end while

Output: Final point $\boldsymbol{y}^m$.

For P–U, we evaluate 41 logarithmically spaced grid points in two dimensions and 61 in three dimensions. The search variable is the natural logarithm of the common step. Bounded scalar minimization then refines the intervals around grid local minima, with at most the three best intervals refined in two dimensions. All grid values and endpoints remain candidates. The refinement tolerances in the logarithmic variable are 10^{-9} and 10^{-10}, respectively. We use the resulting steps as approximate common-step designs.

For two-dimensional P–I, vertices are ordered lexicographically with -0.9 preceding 0.9. The first vertices attaining the minimum and maximum signed values are selected. If a vertex value is exactly zero, the first such vertex is used. Otherwise, when the signed extrema bracket zero, a path flips differing coordinates in increasing index order and its first sign-changing edge is interpolated. This specifies the selection when the optimal step vector is not unique.

For three-dimensional P–I, we apply L-BFGS-B to the squared block norm from eight starting points: the zero vector, the clipped common-step design, the constant vectors -0.7 and -0.9, and four vectors drawn uniformly from $[-0.9, 0.9]^7$ with random seed 20261015. We set `ftol`$= 10^{-16}$, `gtol`$= 10^{-12}$, `maxiter`$= 1000$, and `maxls`$=$ 50. The analytic derivative obtained from a largest singular-vector pair is used. It is the ordinary derivative when the largest singular value is simple; at a repeated largest singular value the objective may be nondifferentiable. We therefore treat the computation as a heuristic search and select the smallest actual block norm among all starting and returned points. The common-step search supplying the initial guess is included in the cost of P–I.

For quadratic objectives, every method uses the closed-form update (3.2). For the nonquadratic benchmark family, write a coordinate update as $\boldsymbol{x}^+ = \boldsymbol{x} + \alpha \mathbf{e}_j$, where

$\alpha = \tau\beta$ is the actual displacement. The scalar equation is

$$F_j(\alpha) = \frac{\alpha}{\tau} + (H_\rho \boldsymbol{x})_j + \frac{(H_\rho)_{jj}}{2}\alpha + \varepsilon D(x_j, \alpha) = 0, \tag{5.6}$$

where, for scalar z,

$$D(z,\alpha) = \begin{cases} \dfrac{\log\cosh(z+\alpha) - \log\cosh(z)}{\alpha}, & \alpha \neq 0, \\ \tanh z, & \alpha = 0. \end{cases}$$

All methods use Brent's method, with absolute and relative tolerances 2×10^{-14} and at most 100 iterations. Set

$$g = (H_\rho \boldsymbol{x})_j + \varepsilon \tanh(x_j), \qquad c = \frac{1}{\tau} + \frac{(H_\rho)_{jj}}{2}.$$

Since

$$F_j'(\alpha) = c + \varepsilon \int_0^1 t \operatorname{sech}^2(x_j + t\alpha)\, \mathrm{d}t \geq c > 0,$$

the interval with endpoints 0 and $-(1 + 10^{-12})g/c$ brackets the unique root when $g \neq 0$. If $g = 0$, we set $\alpha = 0$. Stable evaluation of D uses a cubic small-displacement expansion for $|\alpha| < 10^{-5}$. For $10^{-5} \leq |\alpha| \leq 0.5$, the log-cosh difference is evaluated as $\log(1 + \tanh(z)\sinh(\alpha) + 2\sinh^2(\alpha/2))$ using `log1p`. For larger displacements, the identity $\log\cosh z = |z| + \log(1 + e^{-2|z|}) - \log 2$ is used. Post-run scalar residual checks are excluded from the solve counts.

The root solver uses analytic directional derivatives for bracketing, and the proposed designs use analytic Hessians. The implementation therefore measures performance with this information available. Scalar residual counts refer to evaluations of the scalar residual in (5.6), rather than calls to a general objective-value oracle.

5.4 Quadratic examples

We first consider Q2 and Q3. Here the finite-block criterion gives the exact worst-case energy-norm amplification, so these examples directly test the effect of accounting for the interaction among directions. Table 3 collects the stopping counts for all five cases. On Q2, the practical coordinate rule R–C requires 271 updates, while P–U and P–I require 63 and 9. Figure 1 shows the corresponding decrease in the objective value.

The same effect occurs in the fully coupled three-dimensional example Q3: R–C, P–U, and P–I require 109, 86, and 9 updates, respectively. For each schedule, let $\mathcal{P}_m$ denote the ordered product of the first m actual direction-update matrices, including any changes in block phase, and set $\mathcal{P}_0 = I$. Define

$$m_{\mathrm{wc}} = \min\left\{ m \geq 0 : \|\mathcal{P}_m\|_H^2 \leq 10^{-12} \right\}.$$

Table 4 compares this worst-case stopping count with the initial block factors. In this section, Q_9 denotes the amplification of the relevant nine-update model block.

Table 3 Direction updates required to reach the relative objective-gap tolerance 10^{-12}

Case	R–T	R–S	R–C	P–U	P–I
Q2	1077	387	271	63	9
N2	2784	991	698	75	81
Q3	695	227	109	86	9
N3	655	215	101	83	36
H3	60116	15572	9892	8255	5769

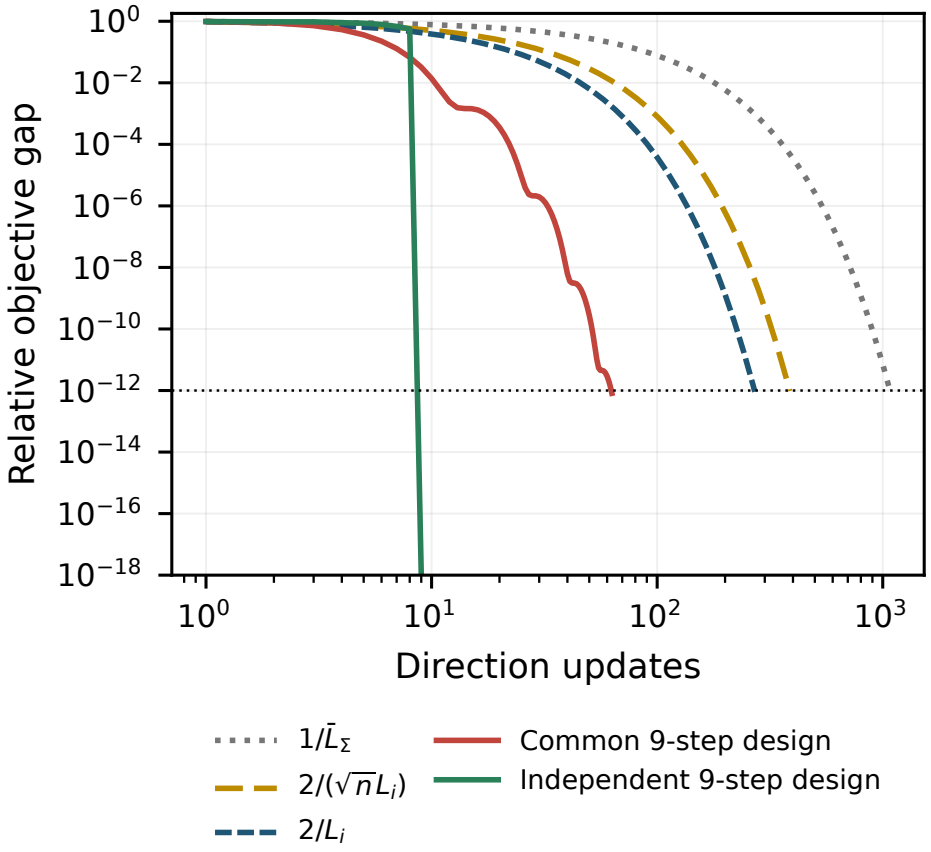

Fig. 1 Q2: relative objective gap against direction updates for all five step rules. The horizontal axis starts at one update on a logarithmic scale. The dotted horizontal line is the stopping tolerance 10^{-12}; values below 10^{-18} are displayed at 10^{-18}

This calculation accounts for all initial errors. For example, the selected initial vector reaches the tolerance after 109 R–C updates on Q3, while the full-product criterion requires 113. Both tests give nine updates for P–I.

Table 4 Initial quadratic block factors and the first update count m_{wc} at which the full accumulated product satisfies $\|\mathcal{P}_m\|_H^2 \le 10^{-12}$. The operator norms evaluate worst-case amplification over all initial errors in floating-point arithmetic

Rule	Q_9 (Q2)	m_{wc}	Q_9 (Q3)	m_{wc}
R–T	0.893823	1077	0.843994	722
R–S	0.742881	387	0.631463	236
R–C	0.663420	271	0.506785	113
P–U	0.294122	64	0.457336	89
P–I	3.419427×10^{-17}	9	2.379199×10^{-9}	9

For Q2, the rational schedule in Appendix F gives $P_9 = 0$ exactly when $\rho = 19/20$. This rational-arithmetic identity explains why a nine-stage block can eliminate every initial error in the exact quadratic problem. For Q3, the P–I schedule obtained by the present numerical implementation has $Q_9 \approx 2.3792 \times 10^{-9}$.

The exact-arithmetic guarantee stated here applies to the reference Q3 schedule printed in Table F1. Interpreting those displayed decimals as exact rational numbers and taking $\rho = 4/5$, the three leading principal minors of

$$10^{-16} H_\rho - P_9^{\mathsf{T}} H_\rho P_9 \tag{5.7}$$

are strictly positive. Hence $Q_9^2 < 10^{-16}$ for that reference schedule. Its last digits differ from the runtime schedule used for the floating-point result above. Appendix F states the rational inputs, and the accompanying source code contains executable exact-arithmetic checks. This certifies the displayed reference schedule without asserting global optimality of the three-dimensional numerical search.

5.5 Smooth nonquadratic examples

Next, we consider the effect of using a quadratic model to choose steps for the original nonlinear iteration. In N2, the Hessian and the schedule are refreshed every nine updates. The results are shown in Figure 2. R–C requires 698 updates, compared with 75 for P–U and 81 for P–I. Thus both designs retain a substantial reduction in updates on this nonquadratic problem. The slightly smaller count for P–U also shows that minimizing the terminal norm of each local model does not determine the ordering of the complete nonlinear trajectories.

In the three-dimensional case N3, the schedule is designed from the initial Hessian and then reused. Even with this fixed schedule, the counts decrease from 101 for R–C to 83 for P–U and 36 for P–I. The model still needs to be interpreted locally: the initial P–I model has $Q_9^2 \approx 6.685 \times 10^{-18}$, whereas the actual nonlinear relative gap after the first nine updates is approximately 8.335×10^{-5}. The matrix norm describes the frozen quadratic model; Theorem 4.7 relates such a product to the derivative at the true minimizer, and Theorem 4.5 compares fixed quadratic models. Neither result identifies the quadratic factor with a finite-displacement nonlinear error ratio.

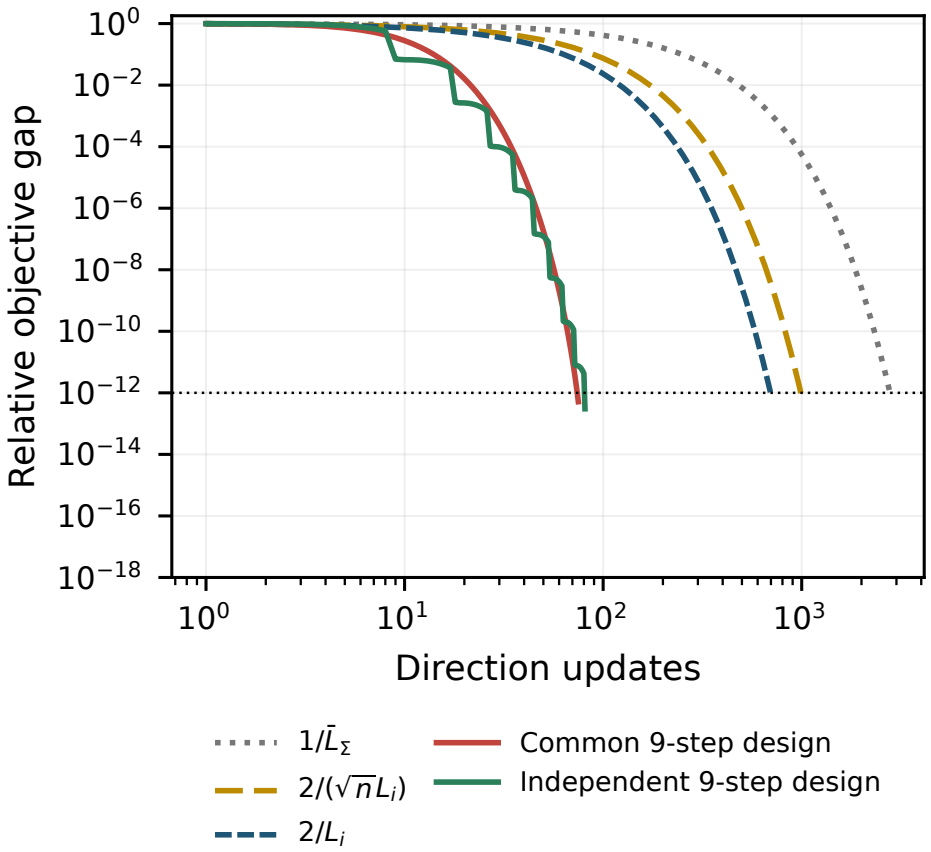

Fig. 2 N2: relative objective gap with Hessian and schedule refresh every nine direction updates for P–U and P–I. The direction cycle continues across block boundaries. The logarithmic axes and display floor are as in Fig. 1

Finally, H3 examines a more ill-conditioned nonquadratic problem with the same one-time design policy. Figure 3 shows the convergence curves. The stopping counts are 9892 for R–C, 8255 for P–U, and 5769 for P–I. The advantage in update counts persists, but the additional effort used to obtain the independent schedule becomes important, as the timing results below show.

5.6 Computational cost and interpretation

We now compare the cost of obtaining the steps with the savings during the iteration. The reconstructed computations used Python 3.13.5, NumPy 2.1.3, and SciPy 1.15.3 on macOS arm64. These new measurements replace the earlier Linux results. These timing and residual measurements concern the five main benchmarks, not the supplementary Chebyshev-root experiments. OpenBLAS and OpenMP were restricted to one thread. Each method and case was timed five times after one warm-up, with randomized method order using seed 20261031. The reported milliseconds measure

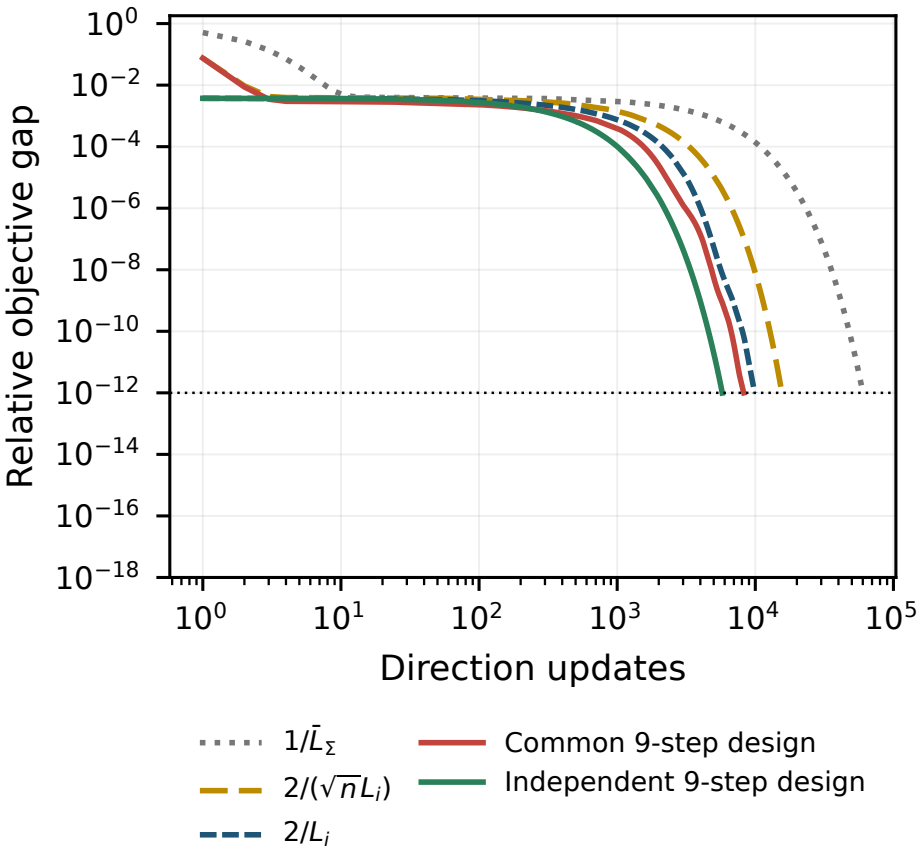

Fig. 3 H3: convergence on the fully coupled ill-conditioned nonquadratic problem. Both designed schedules use the initial Hessian without subsequent refresh. Total costs, including the design, are given in Tables 5 and 6

this implementation and execution environment. Total time includes step preparation, model construction, design, scalar solves, and objective stopping checks. Fixed benchmark construction, trajectory recording, validation, and plotting are excluded. Design and iteration times are recorded within each run; their separate medians need not sum to the median total.

Table 5 gives the total times, and Table 6 separates design and iteration for P–U and P–I. For N2, the finite endpoint design is inexpensive enough that P–I reduces the total time from 7.589 ms for R–C to 1.774 ms. P–U takes 8.020 ms, since repeated design offsets its shorter iteration. In H3, the one-time common design reduces the total from 104.697 to 89.198 ms, while the independent design takes 116.394 ms. On Q3 and N3, the independent search dominates an otherwise short iteration. These results distinguish improved error damping from reduced total running time.

In Table 6, N_H counts full Hessian constructions, with zero charged for a known quadratic matrix, and N_Q counts evaluations of the model criterion. The 128 vertices

Table 5 Total running time in milliseconds, including design: medians of five runs

Case	R–T	R–S	R–C	P–U	P–I
Q2	4.804	1.816	1.218	1.124	0.251
N2	29.485	10.596	7.589	8.020	1.774
Q3	3.163	1.052	0.509	1.607	54.559
N3	7.125	2.435	1.142	2.104	50.768
H3	615.195	163.143	104.697	89.198	116.394

Table 6 Design work and its measured cost for the proposed schedules. N_H counts Hessian constructions and N_Q counts model-criterion evaluations. Times are medians in milliseconds

Case	Rule	N_H	N_Q	Design	Iteration
Q2	P–U	0	50	0.835	0.287
Q2	P–I	0	129	0.202	0.046
N2	P–U	9	441	7.104	0.895
N2	P–I	9	1161	0.816	0.925
Q3	P–U	0	72	1.211	0.396
Q3	P–I	0	1310	54.508	0.048
N3	P–U	1	70	1.187	0.925
N3	P–I	1	1221	50.334	0.433
H3	P–U	1	76	1.316	87.881
H3	P–I	1	1322	54.543	61.926

of the two-dimensional design count as 128 model evaluations even though they are evaluated together. Saved vertex values are reused along the selected cube path; the final operator-norm check is one additional evaluation. A joint value-and-gradient evaluation in the three-dimensional search counts once, and its common-step initialization is included. Table 7 reports the scalar residual work in the nonquadratic examples. The different costs of these operations account for the fact that their counts alone do not predict the total time.

Table 7 Scalar residual evaluations in the nonquadratic runs. The count is not a black-box objective-evaluation count

Case	R–T	R–S	R–C	P–U	P–I
N2	11420	4185	2969	357	368
N3	2818	950	461	385	167
H3	230555	61533	39741	33540	23395

Two further examples clarify the role of the block objective. For $n = 2$, $\rho = 0$, and $\varepsilon = 0$, the stopping counts for R–C, P–U, and P–I are 2, 2, and 8. The independent

design attains the optimal terminal norm, but its schedule need not terminate as early as the coordinate rule. For $n = 2$, $\rho = 0.98$, and $\varepsilon = 0$, the corresponding counts are 685, 70, and 81, while the P–U and P–I block norms are approximately 0.6366 and 0.1919. Thus the smaller worst-case block factor can coexist with a larger stopping count from a specified initial vector.

The numerical solves are consistent with the dissipation identity. Across the five main cases, the largest absolute scalar residual is 1.518452×10^{-14}, and the maximum normalized residual

$$\frac{|V(\boldsymbol{x}^+) - V(\boldsymbol{x}) + \alpha^2/\tau|}{\max\{V(\boldsymbol{y}^0), |V(\boldsymbol{x})|, |V(\boldsymbol{x}^+)|\}} \tag{5.8}$$

is 7.986837×10^{-15}. The largest objective increase is 5.650305×10^{-20}. These measurements assess the accuracy of the computed iterations; the rational checks above concern exact arithmetic for the stated coefficients.

Taken together, the examples show that the matrix criterion can produce substantially shorter iterations, while the benefit in running time depends on the cost and reuse of the design. The comparison concerns step selection within the same directional scheme and with the stated curvature information. Assessing it against optimized SOR, conjugate gradients, direct solvers, or algorithms with other information requirements would require a broader comparison under matched cost assumptions.

5.7 Setup for the nonconvex applications

We compare three prescribed step rules within the same directional Itoh–Abe iteration. Let $\boldsymbol{y}^m$ denote the point after m individual direction updates. The rules are

$$\begin{aligned} \mathrm{C}: &\quad \tau_m = h, \\ \mathrm{E}: &\quad \tau_m = h/9, \\ \mathrm{Ch}: &\quad \tau_{9k+i-1} = h a_i^{\mathrm{ch}}, \qquad i = 1, \dots, 9. \end{aligned} \tag{5.9}$$

Rules E and Ch have the same total step size h over nine updates; rule C has total $9h$. With chronological directions and stopping checks after every update, rule E is also identical to a one-stage method with step $h/9$. Grouping these updates into blocks changes neither its trajectory nor its evaluation count.

For rule Ch, put $s = 9$ and $\varepsilon_{\mathrm{ch}} = 0.05$ and set

$$\begin{aligned} w_0 &= 1 + \varepsilon_{\mathrm{ch}}/s^2, & w_1 &= T_s(w_0)/T_s'(w_0), \\ z_i &= \frac{\cos((2i-1)\pi/(2s)) - w_0}{w_1}, & a_i^{\mathrm{ch}} &= -1/z_i. \end{aligned} \tag{5.10}$$

Here T_s is the Chebyshev polynomial of the first kind. The coefficients are positive and sum to one in exact arithmetic. They are applied in the order $i = 1, \dots, 9$, without reordering. Table G1 records the actual floating-point coefficients; the implementation uses $a_i^{\mathrm{ch}} = -1/z_i$ without an additional normalisation.

All rules use the initial points in Table 8 and the same chronological direction sequence

$$\sigma = \frac{\sqrt{5}-1}{2}, \qquad \xi_m = \{m\sigma\}, \qquad \widehat{\boldsymbol{d}}^{\,m} = \begin{pmatrix} \cos(2\pi\xi_m) \\ \sin(2\pi\xi_m) \end{pmatrix}. \tag{5.11}$$

Here $\{t\}$ denotes the fractional part. The phase is zero and the direction index continues across block boundaries. Each computed direction is normalised to unit length in floating-point arithmetic.

Table 8 Initial points for the nonconvex applications

Label	(A)	(B)	(C)	(D)
Initial point	$(-2,-2)^\mathsf{T}$	$(-2,2)^\mathsf{T}$	$(2,-2)^\mathsf{T}$	$(2,2)^\mathsf{T}$

For every rule, the scalar equation

$$\phi(\beta) = V(\boldsymbol{x} + \tau\beta\boldsymbol{d}) - V(\boldsymbol{x}) + \tau\beta^2 = 0$$

is solved on a nonzero sign-changing bracket using Brent's method. The absolute and relative tolerances are both 10^{-12}, with at most 256 iterations. Starting at $p = 10^{-5}/\max\{1,\tau\}$, we evaluate $\phi(p)$ and $\phi(-p)$. If neither is negative, p is halved while $p \geq 10^{-14}$. If both are negative, the smaller value selects the sign, with the positive sign preferred in a tie. The selected signed probe is then doubled until a positive value is obtained, allowing at most 128 further expansions after the first doubling. Brent's method uses the last sign-changing interval. This rule does not require the root of smallest absolute value.

If probing detects no negative value, a numerical zero update is recorded. This is an implementation convention, not an exact test of Clarke directional stationarity. Bracketing or solver failure terminates the run and is reported separately; a non-root is never substituted. Neither event occurred in the 24 primary runs.

The principal cost measure is the number N_f of objective-function evaluations. It includes the initial value, all probing and bracketing calls, Brent iterations, and the final candidate evaluation. The current objective value is cached. Targets are checked after each direction update using unrounded values. Evaluation budgets are enforced before every objective call. If a budget is reached during a solve, the incomplete update is discarded and the last accepted point is returned; evaluations spent on that solve remain included in N_f. A common limit of 200000 direction updates was inactive in every run. We report individual direction updates m, rather than outer block counts.

5.8 Application to the Rosenbrock function

We consider the smooth nonconvex Rosenbrock objective [13],

$$V(x,y) = (1-x)^2 + 100(y-x^2)^2. \tag{5.12}$$

Its unique minimiser is $(1,1)^{\mathsf{T}}$, with minimum value zero. We set $h = 1$, use target $V \leq 10^{-8}$, and impose an evaluation budget of 10^5 for every rule.

Table 9 and Figure 4 report the results from the same saved trajectories.

Table 9 Rosenbrock results with $h = 1$ and target 10^{-8}. Rules C, E and Ch are defined in (5.9). N_f counts objective evaluations and m counts individual direction updates. T: target reached; B: evaluation budget reached. Final values are rounded for display; stopping uses unrounded values

Initial	Rule	Final V	N_f	m	Stop
(A)	C	9.9696363×10^{-9}	31372	1431	T
(A)	E	9.9827696×10^{-9}	20215	942	T
(A)	Ch	9.9574751×10^{-9}	43970	2286	T
(B)	C	7.4250958×10^{-9}	35893	1625	T
(B)	E	9.4433398×10^{-9}	14460	727	T
(B)	Ch	9.9484919×10^{-9}	43918	2286	T
(C)	C	3.5125909×10^{1}	100000	4393	B
(C)	E	9.7006596×10^{-9}	61078	2608	T
(C)	Ch	9.9490600×10^{-9}	95800	4171	T
(D)	C	9.9969274×10^{-9}	52295	2399	T
(D)	E	9.0058917×10^{-9}	21578	994	T
(D)	Ch	9.5106290×10^{-9}	55127	2756	T

Rule E reaches the target with the fewest evaluations from all four initial points. Rule Ch requires more evaluations than C from (A), (B), and (D). From (C), rule C exhausts the budget at $V \approx 35.13$, while E and Ch reach the target after 61078 and 95800 evaluations, respectively. Thus the prescribed Chebyshev allocation provides no improvement over equal steps in these Rosenbrock tests. In particular, an advantage over rule C alone would not identify a benefit from the nonuniform allocation.

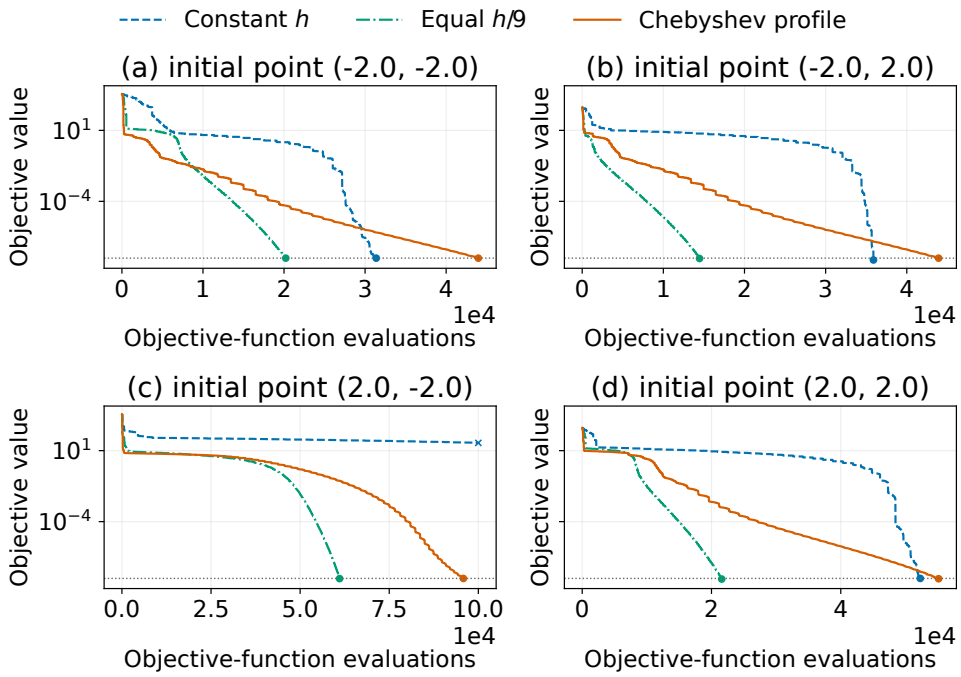

Fig. 4 Rosenbrock objective histories against objective-function evaluations for rules C, E and Ch. The dotted horizontal line marks 10^{-8}. Circles mark target attainment and a cross marks budget exhaustion. Endpoints use the same records as Table 9

5.9 Application to a nonsmooth Chebyshev–Rosenbrock function

We next consider Nesterov's second nonsmooth Chebyshev–Rosenbrock function in two dimensions [14],

$$V(x, y) = \frac{1}{4}|x - 1| + |y - 2|x|| + 1|\,. \tag{5.13}$$

Its unique minimiser is $(1, 1)^{\mathsf{T}}$, with minimum value zero, and $y = 2|x| - 1$ is a piecewise-linear valley. We set $h = 3$, use target $V \leq 10^{-5}$, and impose a common evaluation budget of 3×10^5.

Table 10 and Figure 5 report the results.

Rules E and Ch reach the target from every initial point, whereas rule C exhausts the common budget in all four cases. Rule Ch requires fewer objective evaluations than E from (A), (C), and (D), but more from (B): 113862 versus 98960. From (B), Ch nevertheless uses fewer direction updates, 6615 versus 7462. Objective-evaluation counts and direction-update counts therefore give different orderings in this case. The benefit of the prescribed allocation depends on the initial point.

Repeating all 12 nonsmooth cases with root tolerances 10^{-8} and 10^{-10} preserves the target-attainment outcomes and the ordering of E and Ch by evaluation count; see Appendix G. In a separate diagnostic, extending the budget of rule C from (A)

Table 10 Nonsmooth Chebyshev–Rosenbrock results with $h = 3$ and target 10^{-5}. Rules, counts, and stopping codes are as in Table 9. Budget-limited counts are not costs of target attainment

Initial	Rule	Final V	N_f	m	Stop
(A)	C	1.8940518×10^{-5}	300000	26116	B
(A)	E	9.9920164×10^{-6}	76995	6166	T
(A)	Ch	9.7844296×10^{-6}	21997	1303	T
(B)	C	1.9258084×10^{-5}	300000	25967	B
(B)	E	9.9966326×10^{-6}	98960	7462	T
(B)	Ch	9.4411458×10^{-6}	113862	6615	T
(C)	C	1.9064430×10^{-5}	300000	26050	B
(C)	E	9.9873592×10^{-6}	76978	6166	T
(C)	Ch	9.6040361×10^{-6}	23316	1358	T
(D)	C	1.9108253×10^{-5}	300000	26018	B
(D)	E	9.9329651×10^{-6}	73291	6005	T
(D)	Ch	9.7085391×10^{-6}	17598	1070	T

to 10^6 yields target attainment after 583077 evaluations and 49340 updates, with $V = 9.997085102964043 \times 10^{-6}$. This extended run is excluded from the common-budget comparison. The budget-limited results therefore indicate slow progress, rather than failure to approach the target.

Across the 24 primary runs, the maximum normalised dissipation residual, defined as in (5.8) with $\alpha = \tau\beta$, is 3.904520×10^{-12}. The largest objective increase is 3.330670×10^{-16}. These diagnostics quantify numerical accuracy; they do not replace the assumptions of the exact-iteration convergence results.

Taken together, the nonconvex examples show that the value of a prescribed allocation is problem- and initial-point-dependent. They do not establish general superiority of Chebyshev-root steps and do not test optimality of the matrix-based independent-stage design. Figure 5 shows the nonsmooth objective-evaluation histories. Direction-update histories for the two nonconvex problems are shown in Figures 6 and 7. Outer block counts are not used as a performance comparison.

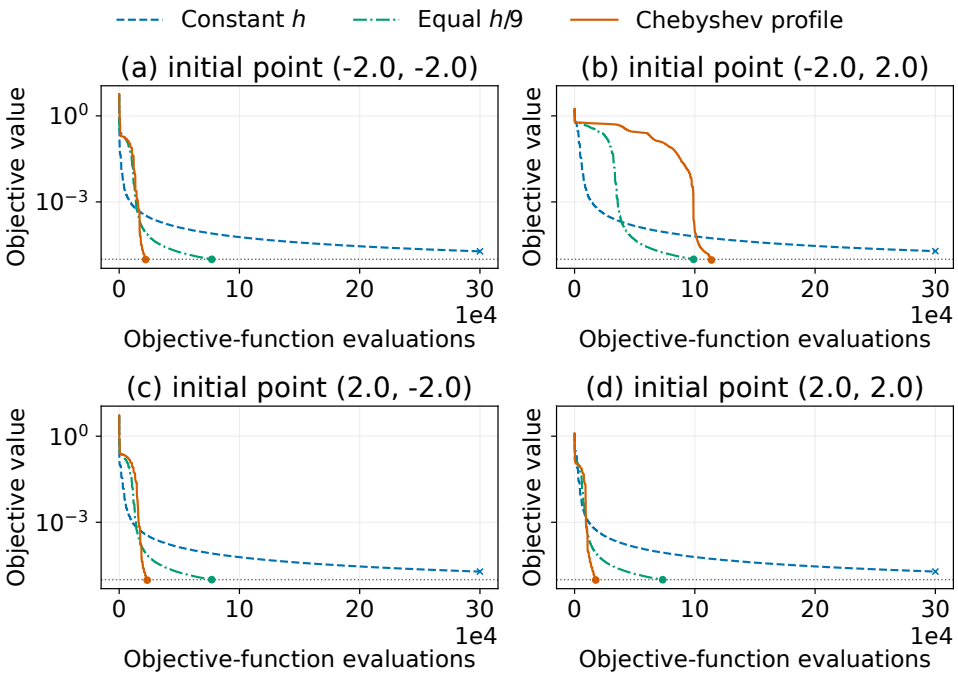

Fig. 5 Nonsmooth Chebyshev–Rosenbrock objective histories against objective-function evaluations. Rules and markers are as in Figure 4; the target is 10^{-5}. Endpoints use the same records as Table 10

Histories against individual direction updates

Figures 6 and 7 use individual direction updates as the horizontal axis. They use the same trajectories as the evaluation-count figures. A final repeated state may include evaluations spent on an interrupted solve; it does not represent an additional accepted direction update.

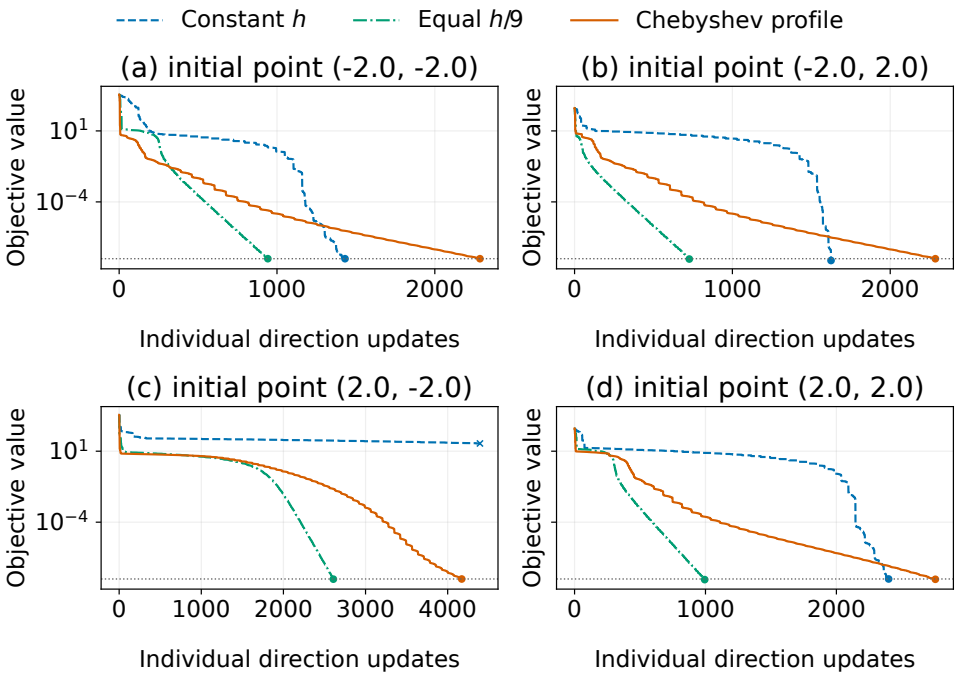

Fig. 6 Rosenbrock histories against individual direction updates. Rules C, E and Ch and the target are as in Figure 4. Each horizontal unit represents one direction stage for every rule, including any numerical zero stage

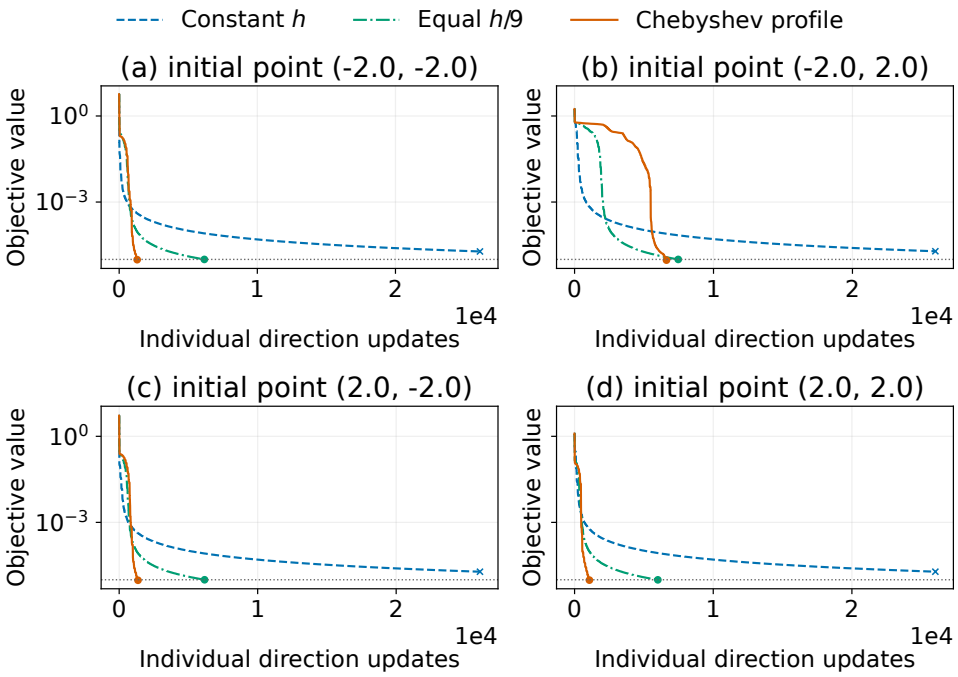

Fig. 7 Nonsmooth Chebyshev–Rosenbrock histories against individual direction updates. Rules, markers, and the target are as in Figure 5

Reproducibility

The single-file Python implementation accompanying this manuscript, `numerical_experiments.py`, contains both experiment suites, the prescribed coefficient generation, rational-arithmetic checks, CSV export, figure generation, and validation routines. It requires NumPy, SciPy, and Matplotlib, with no other local source files or pre-existing data archives. The default execution generates the matrix-based benchmarks and the nonconvex experiments in separate output directories. The `--validate-only` option checks saved histories and recomputes residual maxima; the `--audit-counters` option replays the matrix-based cases with external counters for criterion evaluations and scalar residual callbacks. Instrumented replay times are excluded from the reported timing comparisons. Execution commands and output file names are documented in the accompanying README. The source consolidates the numerical kernels used for the reported runs. Rerunning the code generates new timing measurements; the reported timing values refer to the recorded reference runs in the stated execution environment.

The Q2–H3 benchmarks were rerun using a standalone implementation reconstructed from the stated definitions, Algorithm 2, and implementation settings. The original executable files were unavailable. The reconstructed implementation reproduces the 25 direction-update counts in Table 3 and the 10 worst-case counts in Table 4. The timings, criterion and residual evaluation counts, and accuracy diagnostics are obtained from the new runs.

The nonconvex experiments use a common implementation. All computations were performed using Python 3.13.5, NumPy 2.1.3, and SciPy 1.15.3 on macOS arm64, with OpenBLAS and OpenMP configured for one thread. Floating-point search results and the numbers of criterion, scalar-residual, and objective-function evaluations may depend on the numerical-library versions and execution platform. The tabulated values correspond to the reference environment specified above. The nonconvex experiments use deterministic directions and require no random seeds. Tables and figures are generated from the same saved iteration histories.

For the nonconvex experiments, single-run elapsed times are retained as diagnostics and are not used to claim a runtime speedup. The matrix-based benchmarks instead use the five-run timing protocol described in Section 5.6.

For the nonconvex experiments, the primary absolute and relative root tolerances are both 10^{-12}. The nonquadratic matrix-based benchmarks use absolute and relative tolerances of 2×10^{-14}, as stated in Section 5.3. Targets are checked after every direction update, and evaluation budgets are enforced before every objective-function call. The same counting and stopping rules are used for all methods in each comparison.

6 Discussion

The results provide a step-size design principle within a prescribed directional iteration. We discuss its relation to relaxation methods, computation beyond the exact scalar reductions, and the scope of the convergence guarantees.

6.1 Relaxation, direction geometry, and the stage count

Writing $\omega_i = 1 - r_i$, the quadratic update becomes

$$M_i = I - \frac{\omega_i}{q_i} \boldsymbol{d}_i \boldsymbol{d}_i^{\mathsf{T}} H, \qquad \omega_i = \frac{\tau_i q_i}{1 + \tau_i q_i / 2} \in (0, 2). \tag{6.1}$$

For coordinate directions, this is the SOR coordinate update with a stage-dependent relaxation parameter [5]. The choice $r_i = 0$, equivalently $\tau_i = 2/q_i$, gives exact minimization along $\boldsymbol{d}_i$ for the quadratic objective.

If the prescribed directions form an H-conjugate basis, the transformed directions $\boldsymbol{u}_i$ are orthonormal, and

$$S_n = \sum_{i=1}^{n} r_i \boldsymbol{u}_i \boldsymbol{u}_i^{\mathsf{T}}, \qquad Q_H = \max_{1 \le i \le n} |r_i|.$$

Setting all $r_i = 0$ therefore eliminates the quadratic error in n direction updates. The present design problem addresses prescribed direction sequences that need not have this conjugacy structure.

The stage count specifies the number of directional updates included in the design objective. It does not specify the order of accuracy of a time integrator. For fixed directions and a fixed positive profile $\boldsymbol{a}$, let $\tau_i = h a_i$. If V is C^2 near the point under consideration and the local solution branches are used, then

$$T_h(\boldsymbol{x}) = \boldsymbol{x} - h \left(\sum_{i=1}^{s} a_i \boldsymbol{d}_i \boldsymbol{d}_i^{\mathsf{T}} \right) \nabla V(\boldsymbol{x}) + O(h^2) \qquad (h \downarrow 0).$$

The leading continuous system therefore depends on the weighted directions; see Appendix B. In particular, the nine-stage designs used in Section 5 optimize a product of nine directional factors and do not constitute ninth-order time integrators.

6.2 General higher-dimensional design and its cost

Endpoint projection remains valid in every dimension $n \ge 2$; see Lemma C.1. Let $V_1, V_s \in \mathbb{R}^{n \times (n-1)}$ have orthonormal columns spanning $\boldsymbol{u}_1^{\perp}$ and $\boldsymbol{u}_s^{\perp}$, respectively. For $s \ge 2$, the independent-stage problem reduces to

$$\min_{\boldsymbol{r}_{\mathrm{mid}} \in [-\eta, \eta]^{s-2}} \left\| V_s^{\mathsf{T}} A_{s-1}(r_{s-1}) \cdots A_2(r_2) V_1 \right\|_2 . \tag{6.2}$$

For $s = 2$, the middle product is the identity.

When $n \ge 3$, the reduced matrix need not be scalar or have rank one. For example, take $H = I_3$ and

$$(\boldsymbol{d}_1, \boldsymbol{d}_2, \boldsymbol{d}_3, \boldsymbol{d}_4) = (\mathbf{e}_3, \mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3).$$

With $r_1 = r_4 = 0$,

$$S_4 = \mathrm{diag}(r_2, r_3, 0), \qquad \mathcal{Q}(r_2, r_3) = \max\{|r_2|, |r_3|\}. \tag{6.3}$$

The minimum over $[-\eta,\eta]^2$ is zero at $(0,0)$, whereas every boundary edge has norm at least η. Thus a minimizer cannot always be obtained from a vertex or an edge. The exact construction in Theorem 4.2 depends on the scalar rank-one reduction.

For fixed values of all other parameters, the product is affine in r_i. Indeed, define

$$P_{>i} = A_s \cdots A_{i+1}, \qquad P_{<i} = A_{i-1} \cdots A_1,$$

with empty products interpreted as the identity. Then

$$S_s = P_{>i} \left(\Pi_i + r_i \boldsymbol{u}_i \boldsymbol{u}_i^{\mathsf{T}}\right) P_{<i} = C_i + r_i D_i,$$

where

$$C_i = P_{>i} \Pi_i P_{<i}, \qquad D_i = P_{>i} \boldsymbol{u}_i \boldsymbol{u}_i^{\mathsf{T}} P_{<i}.$$

Each coordinate subproblem is therefore convex. Its spectral-norm epigraph admits the semidefinite representation [15]

$$\begin{aligned} \min_{r_i,t} \quad & t \\ \text{subject to} \quad & \begin{pmatrix} tI & C_i + r_i D_i \\ (C_i + r_i D_i)^{\mathsf{T}} & tI \end{pmatrix} \succeq 0, \\ & -\eta \le r_i \le \eta. \end{aligned} \tag{6.4}$$

This coordinatewise convexity supports coordinate-based optimization, but does not imply joint convexity or global convergence of coordinate minimization. The three-dimensional experiments instead use a multistart search and evaluate the selected product directly.

Once the vectors $\boldsymbol{u}_i$ are available, applying S_s or S_s^{T} to a vector costs $O(ns)$ operations using rank-one updates. Forming the dense product costs $O(sn^2)$, and a dense singular-value decomposition costs $O(n^3)$. Curvature estimation and repeated schedule optimization can nevertheless dominate the savings from fewer updates. Schedule reuse and structured models are therefore important for computational efficiency. If an approximate singular-value calculation is used to certify an upper bound for Q_H, its approximation error must also be controlled.

6.3 Common-scale design and search accuracy

Common-scale design remains useful when the stage steps are tied to a prescribed profile. The following estimate bounds the accuracy of a logarithmic-grid search without assuming unimodality.

Proposition 6.1 (Lipschitz and grid bounds) *For fixed unit vectors $\boldsymbol{u}_i$ and $\boldsymbol{r}, \widetilde{\boldsymbol{r}} \in [-1,1]^s$,*

$$|\mathcal{Q}(\boldsymbol{r}) - \mathcal{Q}(\widetilde{\boldsymbol{r}})| \le \sum_{i=1}^{s} |r_i - \widetilde{r}_i|. \tag{6.5}$$

For a fixed positive profile $\boldsymbol{a}$, define

$$f(t) = Q_H(\mathcal{D}; \exp(t)\boldsymbol{a}).$$

Then

$$|f(t) - f(\widetilde{t})| \le \frac{s}{2}|t - \widetilde{t}|. \tag{6.6}$$

Let $0 < h_- \le h_+ < \infty$. On $[\log h_-, \log h_+]$, consider a grid containing both endpoints and having maximum spacing Δ. Suppose that every computed grid value has absolute error at most $\varepsilon_{\mathrm{ev}}$. If $\widehat{t}$ is a grid point with the smallest computed value, then

$$f(\widehat{t}) \le \min_{t \in [\log h_-, \log h_+]} f(t) + \frac{s\Delta}{4} + 2\varepsilon_{\mathrm{ev}}. \tag{6.7}$$

For a singleton interval, take $\Delta = 0$.

The proof is given in Appendix D.1. The same bound holds after local refinement if the final candidate is selected by comparing computed values over a set containing the entire grid, and all candidate values satisfy the same absolute error bound.

The searches in Section 5 produce approximate common-scale designs. A certified application of (6.7) would additionally require a verified value of $\varepsilon_{\mathrm{ev}}$.

6.4 Interpretation beyond the quadratic model

Three sources of approximation should be distinguished. First, the Hessian model may differ from the target curvature. Theorem 4.5 compares two fixed quadratic models for the same directions and actual steps. Its bound may become uninformative when $s\delta$ is large, and it does not by itself control curvature variation along a nonlinear trajectory.

Second, the design problem may be solved approximately. Corollary 4.6 accounts for this error when a verified optimization tolerance $\varepsilon_{\mathrm{opt}}$ is available. An optimizer's stopping tolerance alone need not provide such a global optimality bound.

Third, the scalar update equations may be solved inexactly. For a quadratic iteration of the form

$$\boldsymbol{e}_i = M_i \boldsymbol{e}_{i-1} + \boldsymbol{\zeta}_i,$$

nonexpansiveness of the exact stage matrices gives

$$\|\boldsymbol{e}_s\|_H \le Q_H(\mathcal{D}; \boldsymbol{\tau}) \|\boldsymbol{e}_0\|_H + \sum_{i=1}^{s} \|\boldsymbol{\zeta}_i\|_H . \tag{6.8}$$

Appendix E.3 relates $\boldsymbol{\zeta}_i$ to the scalar residual. A fixed absolute root tolerance alone does not guarantee exact asymptotic convergence. The numerical experiments therefore report finite-tolerance stopping together with residual and dissipation diagnostics.

The local convergence result concerns a fixed block map. It also extends to continuously differentiable direction and step-selection rules, provided that the local scalar solution branches are used. At $\boldsymbol{x}_*$, derivatives of these rules multiply $\nabla V(\boldsymbol{x}_*)$ or the vanishing stage coefficients $\beta_i(\boldsymbol{x}_*)$. The derivative therefore has the same factorization, with directions and steps evaluated at $\boldsymbol{x}_*$. A search-based selection rule may be discontinuous, so this differentiable-map argument does not automatically apply to every adaptive implementation.

The nonsmooth stationarity theorem permits adaptive stage steps under uniform positive bounds and cyclic density of the chronological direction sequence. For an

independent-stage design, uniform spectral bounds on the models provide one way to ensure the required step bounds. Indeed, suppose

$$\underline{b}I \preceq B_k \preceq \overline{b}I, \qquad 0 < \underline{b} \leq \overline{b} < \infty,$$

and $|r_{k,i}| \leq \eta < 1$. Since the directions are unit vectors,

$$\frac{2(1-\eta)}{\overline{b}(1+\eta)} \leq \tau_{k,i} \leq \frac{2(1+\eta)}{\underline{b}(1-\eta)}. \tag{6.9}$$

The direction-density condition remains a separate requirement.

Both bounds on the actual steps have a role. Consider the one-dimensional objective $V(x) = x^2/2$ with $x^0 = 1$. If

$$\tau_k = \frac{2}{2(k+2)^2 - 1},$$

the update multiplier is $1 - (k+2)^{-2}$, and

$$x^N = \frac{N+2}{2(N+1)} \longrightarrow \frac{1}{2}. \tag{6.10}$$

Here the steps approach zero. If instead

$$\tau_k = 4(k+2)^2 - 2,$$

the multiplier is $-[1 - (k+2)^{-2}]$, giving

$$x^N = (-1)^N \frac{N+2}{2(N+1)}.$$

Here the steps are unbounded. In both cases, the objective decreases but does not approach its minimum. Thus positivity and dissipation alone do not replace the step-size assumptions in the convergence results.

Finally, Hessian-based design and local linearization describe smooth quadratic models and smooth neighborhoods. For general nonsmooth objectives, accumulation-point stationarity follows from the separate nonsmooth analysis and its step-size and direction assumptions [1]. A fixed finite coordinate set does not satisfy cyclic density on $\mathbb{S}^{n-1}$ when $n \geq 2$.

7 Conclusions and perspectives

We studied step-size selection for finite blocks of directional Itoh–Abe updates. For locally Lipschitz objectives that are bounded below and coercive, uniform positive lower and finite upper bounds on the stage steps, together with cyclic direction density, ensure Clarke stationarity of accumulation points.

For positive definite quadratic models, the ordered matrix product gives an exact worst-case energy-norm amplification criterion. Strict contraction holds if and only if the block directions span the space. Independent-stage design admits an exact finite construction in two dimensions through endpoint projections and scalar multi-affine interpolation. This construction extends to higher dimensions under rank-one endpoint compression. Perturbation estimates, grid bounds, and local nonlinear analysis clarify the scope of the design.

The numerical benchmarks show that suitable matrix-based schedules can reduce direction updates, although design costs may offset these savings. In separate nonconvex tests, equal steps attain the Rosenbrock target with fewer objective evaluations than the prescribed Chebyshev-root profile from every tested initial point. On the nonsmooth example, that profile improves upon equal steps from three of the four initial points. These results demonstrate that its benefit is problem- and initial-point-dependent; they do not establish general superiority of the profile or test optimality of the matrix-based design. Future work includes reusable schedules, inexpensive curvature models, and step-size designs tailored to nonconvex and nonsmooth objectives.

Statements and Declarations

Funding. No funding was received for conducting this study.

Competing interests. The author has no relevant financial or non-financial interests to disclose.

Author contributions. Jiabao Yang is the sole author and is responsible for the research, software, numerical experiments, and manuscript preparation.

Data availability. The numerical data supporting the findings of this study, including saved iteration histories and timing records, are available from the corresponding author upon reasonable request.

Code availability. A single-file version of the Python implementation for the numerical experiments is supplied with this manuscript as `numerical_experiments.py`. It generates the experiment data, coefficients, table inputs, and convergence figures and includes validation and counter-audit routines. The code is also available from the corresponding author upon reasonable request.

Use of generative AI. During preparation of this manuscript, ChatGPT (OpenAI) was used to assist with manuscript organization, English drafting, mathematical derivations and proof checking, and code preparation. The author remains responsible for the scientific content and final verification of the manuscript.

References

[1] Riis, E.S., Ehrhardt, M.J., Quispel, G.R.W., Schönlieb, C.-B.: A geometric integration approach to nonsmooth, nonconvex optimisation. Foundations

of Computational Mathematics **22**, 1351–1394 (2022) https://doi.org/10.1007/s10208-020-09489-2

[2] Ehrhardt, M.J., Riis, E.S., Ringholm, T., Schönlieb, C.-B.: A geometric integration approach to smooth optimization: foundations of the discrete gradient method. IMA Journal of Numerical Analysis **45**(3), 1269–1299 (2025) https://doi.org/10.1093/imanum/drae037

[3] Eftekhari, A., Vandereycken, B., Vilmart, G., Zygalakis, K.C.: Explicit stabilised gradient descent for faster strongly convex optimisation. BIT Numerical Mathematics **61**, 119–139 (2021) https://doi.org/10.1007/s10543-020-00819-y

[4] Ushiyama, K., Sato, S., Matsuo, T.: Deriving efficient optimization methods based on stable explicit numerical methods. JSIAM Letters **14**, 29–32 (2022) https://doi.org/10.14495/jsiaml.14.29

[5] Miyatake, Y., Sogabe, T., Zhang, S.-L.: On the equivalence between SOR-type methods for linear systems and the discrete gradient methods for gradient systems. Journal of Computational and Applied Mathematics **342**, 58–69 (2018) https://doi.org/10.1016/j.cam.2018.04.013

[6] Miyatake, Y., Sogabe, T., Zhang, S.-L.: Adaptive SOR methods based on the Wolfe conditions. Numerical Algorithms **84**(1), 117–132 (2020) https://doi.org/10.1007/s11075-019-00748-0

[7] Fält, M., Giselsson, P.: Optimal convergence rates for generalized alternating projections. In: 2017 IEEE 56th Conference on Decision and Control (CDC), pp. 2268–2274 (2017). https://doi.org/10.1109/CDC.2017.8263980

[8] Xu, J., Zikatanov, L.: The method of alternating projections and the method of subspace corrections in Hilbert space. Journal of the American Mathematical Society **15**(3), 573–597 (2002) https://doi.org/10.1090/S0894-0347-02-00398-3

[9] Marshall, N.F., Mickelin, O.: An optimal scheduled learning rate for a randomized Kaczmarz algorithm. SIAM Journal on Matrix Analysis and Applications **44**(1), 312–330 (2023) https://doi.org/10.1137/22M148803X

[10] Ushiyama, K., Sato, S., Matsuo, T.: A unified discretization framework for differential equation approach with Lyapunov arguments for convex optimization. In: Advances in Neural Information Processing Systems, vol. 36, pp. 26092–26120 (2023). https://doi.org/10.52202/075280-1136

[11] Ushiyama, K., Sato, S., Matsuo, T.: Properties and practicability of convergence-guaranteed optimization methods derived from weak discrete gradients. Numerical Algorithms **96**, 1331–1362 (2024) https://doi.org/10.1007/s11075-024-01790-3

[12] Clarke, F.H.: Optimization and Nonsmooth Analysis. SIAM, Philadelphia (1990). https://doi.org/10.1137/1.9781611971309

[13] Rosenbrock, H.H.: An automatic method for finding the greatest or least value of a function. The Computer Journal **3**, 175–184 (1960) https://doi.org/10.1093/comjnl/3.3.175

[14] Gürbüzbalaban, M., Overton, M.L.: On Nesterov's nonsmooth Chebyshev–Rosenbrock functions. Nonlinear Analysis: Theory, Methods & Applications **75**, 1282–1289 (2012) https://doi.org/10.1016/j.na.2011.07.062

[15] Boyd, S., Vandenberghe, L.: Convex Optimization. Cambridge University Press, Cambridge (2004). https://doi.org/10.1017/CBO9780511804441

Appendix A Nonsmooth solvability and stationarity

This appendix establishes the auxiliary results used in Theorem 4.8. The one-stage existence argument follows [1, Lemma 3.1]. After establishing the basic dissipation properties, we obtain Theorem 4.8 by applying Theorem 2.11 to the chronological sequence of internal stages, using uniform bounds on the actual stage steps.

A.1 Solvability

We first establish the existence of a nonzero scalar solution whenever the current point is not Clarke directionally stationary along the selected direction.

Lemma A.1 (Solvability of a one-stage update) *Let $V : \mathbb{R}^n \to \mathbb{R}$ be locally Lipschitz continuous and bounded below. For every $\boldsymbol{x} \in \mathbb{R}^n$, $\boldsymbol{d} \in \mathbb{S}^{n-1}$, and $\tau > 0$, at least one of the following statements holds:*

1. *There exists $\beta \neq 0$ satisfying*

$$\beta = -\frac{V(\boldsymbol{x} + \tau\beta\boldsymbol{d}) - V(\boldsymbol{x})}{\tau\beta}. \tag{A1}$$

2. *V is Clarke directionally stationary at $\boldsymbol{x}$ with respect to $\boldsymbol{d}$.*

The two statements need not be mutually exclusive.

Proof Suppose that the second statement does not hold. There exist $\sigma \in \{-1, 1\}$ and $\varepsilon > 0$ such that

$$V^\circ(\boldsymbol{x}; \sigma\boldsymbol{d}) < -\varepsilon.$$

Since $\tau > 0$ is fixed, the definition of the Clarke directional derivative gives $\delta > 0$ such that

$$\frac{V(\boldsymbol{x} + \tau b\sigma\boldsymbol{d}) - V(\boldsymbol{x})}{\tau b} < -\frac{\varepsilon}{2} \qquad (0 < b < \delta). \tag{A2}$$

Define the continuous function

$$\phi(b) = V(\boldsymbol{x} + \tau b\sigma\boldsymbol{d}) - V(\boldsymbol{x}) + \tau b^2.$$

Choose

$$b_1 = \min\left\{\frac{\delta}{2}, \frac{\varepsilon}{4}\right\} > 0.$$

Then (A2) implies

$$\phi(b_1) < \tau b_1 \left(b_1 - \frac{\varepsilon}{2}\right) < 0.$$

Let $\underline{V} \in \mathbb{R}$ be a lower bound for V. Choose $b_2 > b_1$ sufficiently large that

$$\tau b_2^2 > V(\boldsymbol{x}) - \underline{V}.$$

It follows that

$$\phi(b_2) \geq \underline{V} - V(\boldsymbol{x}) + \tau b_2^2 > 0.$$

By the intermediate value theorem, there exists $b_* \in (b_1, b_2)$ such that $\phi(b_*) = 0$. Setting $\beta = \sigma b_* \neq 0$ gives

$$V(\boldsymbol{x} + \tau\beta\boldsymbol{d}) - V(\boldsymbol{x}) = -\tau\beta^2,$$

which is equivalent to (A1). □

A.2 Stagewise and blockwise dissipation

We first verify the identities used in the main text. If a stage is directionally stationary, the prescribed update has $\beta_i^k = 0$ and $\boldsymbol{g}_i^k = \boldsymbol{g}_{i-1}^k$, so both sides of (4.17) vanish. Otherwise, multiplying the scalar stage equation by $\tau_{k,i}\beta_i^k$ gives

$$V(\boldsymbol{g}_i^k) - V(\boldsymbol{g}_{i-1}^k) = -\tau_{k,i}(\beta_i^k)^2.$$

The update formula and the unit length of $\boldsymbol{d}_i^k$ also give

$$\left\|\boldsymbol{g}_i^k - \boldsymbol{g}_{i-1}^k\right\|_2^2 = \tau_{k,i}^2(\beta_i^k)^2.$$

These two equalities prove (4.17) for every stage. Summing over $i = 1, \dots, s$ and using $\boldsymbol{g}_0^k = \boldsymbol{x}^k$ and $\boldsymbol{g}_s^k = \boldsymbol{x}^{k+1}$ gives

$$\begin{aligned} V(\boldsymbol{x}^k) - V(\boldsymbol{x}^{k+1}) &= \sum_{i=1}^{s} \tau_{k,i}(\beta_i^k)^2 \\ &= \sum_{i=1}^{s} \frac{\left\|\boldsymbol{g}_i^k - \boldsymbol{g}_{i-1}^k\right\|_2^2}{\tau_{k,i}}, \end{aligned}$$

which is (4.18). No differentiability or convexity is used in this argument.

A.3 Properties of the generated sequence

We next collect consequences of the stagewise dissipation identity. Throughout this appendix, the number of stages s is fixed and finite.

Lemma A.2 (Basic properties of the generated sequence) *Suppose that $V : \mathbb{R}^n \to \mathbb{R}$ is locally Lipschitz continuous, bounded below, and coercive. Let the internal stages be generated by exact scalar solves, with zero updates at directionally stationary stages. Assume*

$$0 < \tau_{\min} \le \tau_{k,i} \le \tau_{\max} < \infty \qquad (k \ge 0,\ i = 1, \dots, s).$$

Then every internal point belongs to the compact set

$$K = \left\{\boldsymbol{x} \in \mathbb{R}^n : V(\boldsymbol{x}) \le V(\boldsymbol{x}^0)\right\}.$$

Moreover:

1. *The objective values satisfy*

$$V(\boldsymbol{x}^{k+1}) \le V(\boldsymbol{x}^k), \qquad V(\boldsymbol{x}^k) \to V_\infty$$

   for some finite V_∞.
2. *The scalar stage coefficients satisfy*

$$\sum_{k=0}^{\infty}\sum_{i=1}^{s}(\beta_i^k)^2 < \infty, \qquad \max_{1\le i\le s}|\beta_i^k| \to 0.$$

3. The internal increments satisfy

$$\sum_{k=0}^{\infty}\sum_{i=1}^{s}\left\|\boldsymbol{g}_i^k-\boldsymbol{g}_{i-1}^k\right\|_2^2<\infty.$$

In particular,

$$\max_{1\le i\le s}\left\|\boldsymbol{g}_i^k-\boldsymbol{g}_{i-1}^k\right\|_2\to 0,\qquad \left\|\boldsymbol{x}^{k+1}-\boldsymbol{x}^k\right\|_2\to 0.$$

4. The outer iterate sequence is bounded and has at least one accumulation point.

Proof The stagewise identity (4.17) gives

$$V(\boldsymbol{g}_{i-1}^k)-V(\boldsymbol{g}_i^k)=\tau_{k,i}(\beta_i^k)^2=\frac{\left\|\boldsymbol{g}_i^k-\boldsymbol{g}_{i-1}^k\right\|_2^2}{\tau_{k,i}}\ge 0.$$

Since $\boldsymbol{g}_0^k=\boldsymbol{x}^k$ and $\boldsymbol{g}_s^k=\boldsymbol{x}^{k+1}$, all internal points lie in K. Continuity makes K closed, and coercivity makes it bounded. Hence K is compact.

Summing the stage identities yields

$$V(\boldsymbol{x}^k)-V(\boldsymbol{x}^{k+1})=\sum_{i=1}^{s}\tau_{k,i}(\beta_i^k)^2\ge 0.$$

Thus the objective values are nonincreasing. Since V is bounded below, they converge to a finite value V_∞.

Telescoping over the outer iterations gives

$$\begin{aligned}V(\boldsymbol{x}^0)-V_\infty&=\sum_{k=0}^{\infty}\sum_{i=1}^{s}\tau_{k,i}(\beta_i^k)^2\\&=\sum_{k=0}^{\infty}\sum_{i=1}^{s}\frac{\left\|\boldsymbol{g}_i^k-\boldsymbol{g}_{i-1}^k\right\|_2^2}{\tau_{k,i}}.\end{aligned}\tag{A3}$$

The lower step bound implies

$$\sum_{k=0}^{\infty}\sum_{i=1}^{s}(\beta_i^k)^2\le\frac{V(\boldsymbol{x}^0)-V_\infty}{\tau_{\min}}<\infty.$$

Similarly, the upper step bound gives

$$\sum_{k=0}^{\infty}\sum_{i=1}^{s}\left\|\boldsymbol{g}_i^k-\boldsymbol{g}_{i-1}^k\right\|_2^2\le\tau_{\max}\big(V(\boldsymbol{x}^0)-V_\infty\big)<\infty.$$

These estimates imply that every scalar stage coefficient and every internal increment tends to zero. Since s is fixed and finite, their maxima over the stages also tend to zero.

Furthermore,

$$\left\|\boldsymbol{x}^{k+1}-\boldsymbol{x}^k\right\|_2\le\sum_{i=1}^{s}\left\|\boldsymbol{g}_i^k-\boldsymbol{g}_{i-1}^k\right\|_2\to 0.$$

Finally, the outer iterates belong to the compact set K. They are therefore bounded and have an accumulation point. □

A.4 Stationarity of accumulation points

Proof of Theorem 4.8 For $k \geq 0$ and $i = 0, \ldots, s-1$, define

$$\begin{aligned} \boldsymbol{y}^{ks+i} &:= \boldsymbol{g}_i^k, & \widehat{\boldsymbol{d}}^{\,ks+i} &:= \boldsymbol{d}_{i+1}^k, \\ \widehat{\tau}_{ks+i} &:= \tau_{k,i+1}, & \widehat{\beta}_{ks+i} &:= \beta_{i+1}^k. \end{aligned} \tag{A4}$$

The identity $\boldsymbol{g}_s^k = \boldsymbol{g}_0^{k+1}$ makes these variables a chronological sequence satisfying

$$\boldsymbol{y}^{m+1} = \boldsymbol{y}^m + \widehat{\tau}_m \widehat{\beta}_m \widehat{\boldsymbol{d}}^{\,m}.$$

Its scalar updates, step bounds, and direction sequence satisfy the assumptions of Theorem 2.11. Thus every accumulation point of $(\boldsymbol{y}^m)$ is Clarke stationary. Since $\boldsymbol{x}^k = \boldsymbol{y}^{ks}$, the same conclusion holds for the block iterates. Boundedness follows from stagewise decrease and compactness of the initial sublevel set. □

Appendix B Directional updates and quadratic block identities

B.1 Smooth solvability, dissipation, and consistency

For $V \in C^1(\mathbb{R}^n)$, a unit direction $\boldsymbol{d}$, and $\tau > 0$, define the continuous scalar equation

$$\begin{aligned} F_{\boldsymbol{x},\boldsymbol{d},\tau}(\beta) &:= \beta + \int_0^1 \boldsymbol{d}^{\mathsf{T}} \nabla V(\boldsymbol{x} + t\tau\beta\boldsymbol{d})\, \mathrm{d}t = 0, \\ \boldsymbol{x}^+ &:= \boldsymbol{x} + \tau\beta\boldsymbol{d}. \end{aligned} \tag{B1}$$

For $\beta \neq 0$, this is equivalent to the quotient equation used in the directional iteration. At $\beta = 0$, it holds precisely when $\boldsymbol{d}^{\mathsf{T}} \nabla V(\boldsymbol{x}) = 0$.

Proposition B.1 (Smooth directional solvability and dissipation) *If $V \in C^1(\mathbb{R}^n)$ is convex, (B1) has a unique scalar solution for every $\boldsymbol{x} \in \mathbb{R}^n$, $\boldsymbol{d} \in \mathbb{S}^{n-1}$, and $\tau > 0$. Every solution of (B1), whether or not V is convex, satisfies*

$$V(\boldsymbol{x}^+) - V(\boldsymbol{x}) = -\tau\beta^2 = -\frac{\|\boldsymbol{x}^+ - \boldsymbol{x}\|_2^2}{\tau}. \tag{B2}$$

Proof Set $\varphi(a) = V(\boldsymbol{x} + a\boldsymbol{d})$. If V is convex, then φ' is nondecreasing. For $b_2 > b_1$,

$$\begin{aligned} &F_{\boldsymbol{x},\boldsymbol{d},\tau}(b_2) - F_{\boldsymbol{x},\boldsymbol{d},\tau}(b_1) \\ &\quad = b_2 - b_1 + \int_0^1 \big(\varphi'(t\tau b_2) - \varphi'(t\tau b_1)\big)\, \mathrm{d}t \\ &\quad \geq b_2 - b_1. \end{aligned}$$

Thus $F_{\boldsymbol{x},\boldsymbol{d},\tau}$ is continuous and strictly increasing. Moreover,

$$F_{\boldsymbol{x},\boldsymbol{d},\tau}(b) \geq F_{\boldsymbol{x},\boldsymbol{d},\tau}(0) + b \qquad (b > 0),$$

and

$$F_{\boldsymbol{x},\boldsymbol{d},\tau}(b) \leq F_{\boldsymbol{x},\boldsymbol{d},\tau}(0) + b \qquad (b < 0).$$

Its limits as $b \to \pm\infty$ are therefore $\pm\infty$, proving existence and uniqueness of a zero.

For any solution, the fundamental theorem of calculus gives

$$\begin{aligned} V(\boldsymbol{x}^+) - V(\boldsymbol{x}) &= \tau\beta \int_0^1 \boldsymbol{d}^\mathsf{T} \nabla V(\boldsymbol{x} + t\tau\beta\boldsymbol{d})\, \mathrm{d}t \\ &= -\tau\beta^2. \end{aligned}$$

Since $\|\boldsymbol{d}\|_2 = 1$, this proves (B2). The identity is also immediate when $\beta = 0$. □

Suppose now that V is C^2 near $\boldsymbol{x}$. At $\tau = 0$, the continuous scalar equation becomes

$$\beta + \boldsymbol{d}^\mathsf{T} \nabla V(\boldsymbol{x}) = 0,$$

and its derivative with respect to β is one. The implicit function theorem therefore gives a local solution branch satisfying

$$\begin{aligned} \beta(\boldsymbol{x}, \boldsymbol{d}, \tau) &= -\boldsymbol{d}^\mathsf{T} \nabla V(\boldsymbol{x}) + O(\tau), \\ \boldsymbol{x}^+ &= \boldsymbol{x} - \tau \boldsymbol{d}\boldsymbol{d}^\mathsf{T} \nabla V(\boldsymbol{x}) + O(\tau^2). \end{aligned} \tag{B3}$$

Thus one stage is consistent with the projected flow

$$\dot{\boldsymbol{x}} = -\boldsymbol{d}\boldsymbol{d}^\mathsf{T} \nabla V(\boldsymbol{x}).$$

For fixed unit directions $\mathcal{D} = (\boldsymbol{d}_1, \ldots, \boldsymbol{d}_s)$, a fixed positive profile $\boldsymbol{a}$, and steps $\tau_i = h a_i$, composition of these local branches gives

$$\begin{aligned} T_h(\boldsymbol{x}) &= \boldsymbol{x} - h C_{\mathcal{D},\boldsymbol{a}} \nabla V(\boldsymbol{x}) + O(h^2), \\ C_{\mathcal{D},\boldsymbol{a}} &:= \sum_{i=1}^{s} a_i \boldsymbol{d}_i \boldsymbol{d}_i^\mathsf{T}. \end{aligned} \tag{B4}$$

Indeed, each intermediate point is $\boldsymbol{x} + O(h)$, and the gradient is locally Lipschitz continuous. Summing the stage expansions therefore proves (B4).

The matrix $C_{\mathcal{D},\boldsymbol{a}}$ is positive definite if and only if the directions span $\mathbb{R}^n$, because

$$\boldsymbol{z}^\mathsf{T} C_{\mathcal{D},\boldsymbol{a}} \boldsymbol{z} = \sum_{i=1}^{s} a_i (\boldsymbol{d}_i^\mathsf{T} \boldsymbol{z})^2.$$

It is generally not the identity. If $\sum_i a_i = 1$ and $C_{\mathcal{D},\boldsymbol{a}}$ is a scalar multiple of I, then

$$\operatorname{tr}(C_{\mathcal{D},\boldsymbol{a}}) = 1$$

implies $C_{\mathcal{D},\boldsymbol{a}} = I/n$. In that case, the leading update corresponds to a full-gradient-flow step of size h/n. The number of stages does not determine an order of accuracy.

B.2 Energy coordinates and the contraction criterion

For the quadratic objective, substitution of

$$\nabla V_H(\boldsymbol{x}) = H(\boldsymbol{x} - \boldsymbol{x}_*)$$

into (B1) gives (3.2). Writing

$$B_i = \boldsymbol{d}_i \boldsymbol{d}_i^{\mathsf{T}} H,$$

we have $B_i^2 = q_i B_i$. The same stage matrix therefore admits the representation

$$M_i = \left(I + \frac{\tau_i}{2} B_i\right)^{-1} \left(I - \frac{\tau_i}{2} B_i\right). \tag{B5}$$

The inverse exists because $q_i > 0$ and $\tau_i > 0$.

Conjugation by $H^{1/2}$ gives (3.4). Since $\boldsymbol{u}_i$ is a unit vector,

$$A_i(r_i)^{\mathsf{T}} A_i(r_i) = I - (1 - r_i^2) \boldsymbol{u}_i \boldsymbol{u}_i^{\mathsf{T}},$$

which proves (4.1). The identity $S_s = H^{1/2} P_s H^{-1/2}$ then gives (3.5).

Proof of Theorem 4.1 For an initial vector $\boldsymbol{z}_0$, set

$$\boldsymbol{z}_i = S_i \boldsymbol{z}_0, \qquad i = 0, \dots, s.$$

Equation (4.1) shows that every stage is nonexpansive, and hence $Q_H \le 1$.

Suppose that a nonzero $\boldsymbol{z}_0$ satisfies

$$\|S_s \boldsymbol{z}_0\|_2 = \|\boldsymbol{z}_0\|_2 .$$

Every nonnegative stage decrease must then vanish. Since $1 - r_i^2 > 0$,

$$\boldsymbol{u}_i^{\mathsf{T}} \boldsymbol{z}_{i-1} = 0$$

for every i. Consequently, $\boldsymbol{z}_i = \boldsymbol{z}_{i-1}$, so $\boldsymbol{z}_0$ is orthogonal to every $\boldsymbol{u}_i$.

Conversely, any vector orthogonal to all $\boldsymbol{u}_i$ is unchanged by every stage. Such a nonzero vector exists precisely when the $\boldsymbol{u}_i$ fail to span $\mathbb{R}^n$. Because $H^{1/2}$ is invertible and each $q_i > 0$, this is equivalent to failure of the original directions to span $\mathbb{R}^n$.

If the directions span, every nonzero vector is strictly contracted by the block. Continuity and compactness of the unit sphere then imply $\|S_s\|_2 < 1$. This proves (4.2).

Finally,

$$\begin{aligned} S_{i-1}^{\mathsf{T}} S_{i-1} - S_i^{\mathsf{T}} S_i &= S_{i-1}^{\mathsf{T}} \big(I - A_i^{\mathsf{T}} A_i\big) S_{i-1} \\ &= (1 - r_i^2) S_{i-1}^{\mathsf{T}} \boldsymbol{u}_i \boldsymbol{u}_i^{\mathsf{T}} S_{i-1}. \end{aligned}$$

Summing over the stages yields $S_s^{\mathsf{T}} S_s = I - G_s$. Every summand defining G_s is positive semidefinite, and $I - G_s = S_s^{\mathsf{T}} S_s \succeq 0$. Thus $0 \preceq G_s \preceq I$, and

$$Q_H^2 = \lambda_{\max}(S_s^{\mathsf{T}} S_s) = 1 - \lambda_{\min}(G_s).$$

□

The active subspace can also be isolated exactly. Let

$$\mathcal{U} = \operatorname{span}\{\boldsymbol{u}_1, \dots, \boldsymbol{u}_s\}, \qquad \ell = \dim \mathcal{U},$$

and let $W \in \mathbb{R}^{n\times \ell}$ have orthonormal columns spanning $\mathcal{U}$. Set

$$\widehat{\boldsymbol{u}}_i = W^{\mathsf{T}}\boldsymbol{u}_i, \qquad \widehat{A}_i = I_\ell - (1-r_i)\widehat{\boldsymbol{u}}_i\widehat{\boldsymbol{u}}_i^{\mathsf{T}}.$$

Each stage preserves $\mathcal{U}$ and is the identity on $\mathcal{U}^\perp$. Therefore,

$$S_s = W\big(\widehat{A}_s \cdots \widehat{A}_1\big)W^{\mathsf{T}} + I - WW^{\mathsf{T}}. \tag{B6}$$

The singular values of the compressed product are thus supplemented by $n-\ell$ copies of one.

Proposition B.2 (Triangular representation) *Let $U \in \mathbb{R}^{s\times n}$ have rows $\boldsymbol{u}_i^{\mathsf{T}}$, and define*

$$\Lambda = \operatorname{diag}(1-r_1^2, \dots, 1-r_s^2).$$

Let $\mathcal{L}$ be the unit lower triangular matrix

$$\mathcal{L}_{ij} = \begin{cases} 1, & i = j, \\ (1-r_j)\boldsymbol{u}_i^{\mathsf{T}}\boldsymbol{u}_j, & j < i, \\ 0, & j > i. \end{cases} \tag{B7}$$

Then

$$\begin{aligned} G_s &= U^{\mathsf{T}}\mathcal{L}^{-\mathsf{T}}\Lambda\mathcal{L}^{-1}U, \\ Q_H^2 &\le 1 - \frac{\min_i(1-r_i^2)\lambda_{\min}(U^{\mathsf{T}}U)}{\|\mathcal{L}\|_2^2}. \end{aligned} \tag{B8}$$

Proof For any $\boldsymbol{z}_0 \in \mathbb{R}^n$, set

$$\boldsymbol{z}_i = S_i\boldsymbol{z}_0, \qquad b_i = \boldsymbol{u}_i^{\mathsf{T}}\boldsymbol{z}_{i-1}.$$

Expanding the preceding updates gives

$$\boldsymbol{u}_i^{\mathsf{T}}\boldsymbol{z}_0 = b_i + \sum_{j<i}(1-r_j)(\boldsymbol{u}_i^{\mathsf{T}}\boldsymbol{u}_j)b_j.$$

Thus

$$U\boldsymbol{z}_0 = \mathcal{L}\boldsymbol{b}, \qquad \boldsymbol{b} = (b_1, \dots, b_s)^{\mathsf{T}}.$$

The total squared-norm decrease is

$$\boldsymbol{b}^{\mathsf{T}}\Lambda\boldsymbol{b} = \boldsymbol{z}_0^{\mathsf{T}}U^{\mathsf{T}}\mathcal{L}^{-\mathsf{T}}\Lambda\mathcal{L}^{-1}U\boldsymbol{z}_0.$$

Comparison with the definition of G_s proves the matrix identity.

Moreover,

$$\left\|\mathcal{L}^{-1}U\boldsymbol{z}_0\right\|_2 \ge \frac{\|U\boldsymbol{z}_0\|_2}{\|\mathcal{L}\|_2}.$$

Consequently,

$$\lambda_{\min}(G_s) \ge \frac{\min_i(1-r_i^2)\lambda_{\min}(U^{\mathsf{T}}U)}{\|\mathcal{L}\|_2^2}.$$

Equation (4.4) gives the stated bound. □

As an explicit consequence, suppose that

$$mI \preceq H \preceq LI, \qquad \sum_{i=1}^{s} \boldsymbol{d}_i \boldsymbol{d}_i^{\mathsf{T}} \succeq \alpha I, \qquad |r_i| \le \eta,$$

where $0 < m \le L$, $\alpha > 0$, and $0 \le \eta < 1$. Then

$$Q_H^2 \le 1 - \frac{(1-\eta^2)\alpha m}{L\big(1+\sqrt{2s(s-1)}\big)^2} < 1. \tag{B9}$$

Indeed,

$$\begin{aligned} U^{\mathsf{T}}U &= H^{1/2}\left(\sum_{i=1}^{s} \frac{\boldsymbol{d}_i \boldsymbol{d}_i^{\mathsf{T}}}{q_i}\right) H^{1/2} \\ &\succeq \frac{\alpha}{L} H \succeq \frac{\alpha m}{L} I, \end{aligned}$$

and

$$\|\mathcal{L}\|_2 \le 1 + \|\mathcal{L} - I\|_{\mathrm{F}} \le 1 + \sqrt{2s(s-1)}.$$

This conservative estimate separates direction coverage, conditioning of the quadratic model, and the avoidance of nearly undamped stage factors.

Appendix C Exact independent-stage designs

We establish the two ingredients used in the exact-design theorems: endpoint projection and the range of a scalar multi-affine function.

Lemma C.1 (Endpoint projection) *Let $n \ge 2$ and $s \ge 2$. The independent-stage problem (3.9) has a minimizer with $r_1 = r_s = 0$.*

Proof Let

$$B = A_{s-1}(r_{s-1}) \cdots A_2(r_2),$$

with $B = I$ when $s = 2$. The identities

$$\Pi_i A_i(r_i) = A_i(r_i)\Pi_i = \Pi_i$$

give

$$\begin{aligned} A_s(0)BA_1(0) &= \Pi_s B \Pi_1 \\ &= \Pi_s\big(A_s(r_s)BA_1(r_1)\big)\Pi_1. \end{aligned} \tag{C1}$$

Since orthogonal projections have operator norm at most one, replacing the endpoint parameters by zero cannot increase the block norm. The replacement preserves feasibility because $0 \in [-\eta, \eta]$. Applying it to an attained minimizer proves the result. □

The same argument applies to independent compact parameter intervals whose first and last intervals contain zero. Coupled constraints, such as fixed step ratios or a prescribed total step, need not be preserved by the endpoint replacement.

Lemma C.2 (Range of a scalar multi-affine function) *Let $\eta > 0$ and let p be a nonnegative integer. Suppose that $g : [-\eta, \eta]^p \to \mathbb{R}$ is affine in each variable separately. Let g_- and g_+ denote its minimum and maximum over the 2^p vertices. Then*

$$\begin{aligned} g([-\eta, \eta]^p) &= [g_-, g_+], \\ \min_{\boldsymbol{r} \in [-\eta,\eta]^p} |g(\boldsymbol{r})| &= \operatorname{dist}(0, [g_-, g_+]). \end{aligned} \tag{C2}$$

A minimizer can be chosen with at most one coordinate strictly inside its interval. For $p = 0$, the domain consists of one point.

Proof For $p = 0$, the assertions are immediate. Suppose $p \geq 1$ and write

$$r_i = (1 - t_i)(-\eta) + t_i \eta, \qquad 0 \leq t_i \leq 1.$$

Repeated affine interpolation gives

$$g(\boldsymbol{r}) = \sum_{\boldsymbol{\epsilon} \in \{-1,1\}^p} w_{\boldsymbol{\epsilon}}(\boldsymbol{t}) g(\eta \boldsymbol{\epsilon}),$$

where

$$w_{\boldsymbol{\epsilon}}(\boldsymbol{t}) = \prod_{i:\epsilon_i = 1} t_i \prod_{i:\epsilon_i = -1} (1 - t_i).$$

These weights are nonnegative and sum to one. Every function value therefore lies in $[g_-, g_+]$, and both extrema are attained at vertices. Continuity and connectedness of the domain imply that the image is the entire interval $[g_-, g_+]$. This proves (C2).

If $g_- > 0$ or $g_+ < 0$, a vertex whose value is closest to zero minimizes $|g|$. If any vertex value is zero, that vertex is a minimizer. Otherwise, when $g_- < 0 < g_+$, connect vertices with opposite signs by a path of cube edges. Some adjacent pair on the path has opposite signs. Along their common edge, g is affine and therefore has a zero. All coordinates except the one varying along that edge remain at interval endpoints. □

Proof of Theorem 4.2 By Lemma C.1, an optimal independent-stage design can be chosen with $r_1 = r_s = 0$. In two dimensions,

$$\Pi_i = \boldsymbol{v}_i \boldsymbol{v}_i^{\mathsf{T}}, \qquad \|\boldsymbol{v}_i\|_2 = 1.$$

The endpoint-reduced product is therefore

$$S_s = g(\boldsymbol{r}_{\mathrm{mid}}) \boldsymbol{v}_s \boldsymbol{v}_1^{\mathsf{T}},$$

with g given by (4.5). Since $\left\| \boldsymbol{v}_s \boldsymbol{v}_1^{\mathsf{T}} \right\|_2 = 1$, its norm is $|g(\boldsymbol{r}_{\mathrm{mid}})|$.

Each middle parameter occurs in exactly one affine matrix factor, so g is multi-affine. Lemma C.2 gives

$$\min_{\boldsymbol{r}_{\mathrm{mid}} \in [-\eta,\eta]^{s-2}} |g(\boldsymbol{r}_{\mathrm{mid}})| = \operatorname{dist}(0, [g_-, g_+]),$$

which is the piecewise expression (4.6). The same lemma shows that a minimizer can be chosen with at most one middle parameter strictly inside its interval.

Finally, (3.8) gives finite positive stage steps because $q_i > 0$ and $|r_i| \leq \eta < 1$. When $s = 2$, the middle product is the identity, and the same argument applies to the singleton parameter set. □

Proof of Theorem 4.4 Put $m = n - 1$ and $p = s - 2m$. The H-conjugacy assumptions imply that the transformed directions within each endpoint group are orthonormal. Consequently,

$$\Pi_- := A_m(0) \cdots A_1(0) = I - \sum_{i=1}^{m} \boldsymbol{u}_i \boldsymbol{u}_i^{\mathsf{T}} = \boldsymbol{a}\boldsymbol{a}^{\mathsf{T}},$$

$$\Pi_+ := A_s(0) \cdots A_{s-m+1}(0) = I - \sum_{i=s-m+1}^{s} \boldsymbol{u}_i \boldsymbol{u}_i^{\mathsf{T}} = \boldsymbol{b}\boldsymbol{b}^{\mathsf{T}}.$$

Here $\boldsymbol{a}$ and $\boldsymbol{b}$ are unit vectors.

For arbitrary admissible endpoint parameters, every factor in the first group fixes the range of Π_-, and every factor in the last group fixes the range of Π_+. Since the factors are symmetric,

$$A_m(r_m) \cdots A_1(r_1)\Pi_- = \Pi_-,$$

and

$$\Pi_+ A_s(r_s) \cdots A_{s-m+1}(r_{s-m+1}) = \Pi_+.$$

Writing the middle product as

$$B_{\mathrm{mid}} = A_{s-m}(r_{s-m}) \cdots A_{m+1}(r_{m+1}),$$

with $B_{\mathrm{mid}} = I$ when $p = 0$, we obtain

$$\Pi_+ S_s \Pi_- = \Pi_+ B_{\mathrm{mid}} \Pi_-.$$

The right-hand side is precisely the product obtained by setting all endpoint-group parameters to zero.

Multiplication by orthogonal projections cannot increase the operator norm. The replacement also preserves feasibility. Applying it to an attained minimizer proves that an optimal design exists with both endpoint groups set to zero.

The resulting product is

$$\Pi_+ B_{\mathrm{mid}} \Pi_- = g(\boldsymbol{r}_{\mathrm{mid}}) \boldsymbol{b}\boldsymbol{a}^{\mathsf{T}},$$

where g is defined by (4.8). Its norm is $|g|$, and g is multi-affine in the p middle parameters. Lemma C.2 therefore proves (4.9) and the vertex-and-edge construction, including the singleton case $p = 0$. □

For completeness, the two-stage amplification in two dimensions can be computed directly. Set

$$c = \boldsymbol{u}_1^{\mathsf{T}} \boldsymbol{u}_2, \qquad \Gamma = r_1^2 + r_2^2 + c^2(1 - r_1^2)(1 - r_2^2).$$

Using

$$A_i(r_i)^2 = I - (1 - r_i^2)\boldsymbol{u}_i \boldsymbol{u}_i^{\mathsf{T}}$$

and cyclic invariance of the trace gives

$$\operatorname{tr}(S_2^{\mathsf{T}} S_2) = \Gamma.$$

Moreover,

$$\det(S_2^{\mathsf{T}} S_2) = r_1^2 r_2^2.$$

The larger eigenvalue of $S_2^{\mathsf{T}} S_2$ therefore yields

$$Q_2 = \sqrt{\frac{\Gamma + \sqrt{\Gamma^2 - 4r_1^2 r_2^2}}{2}}. \tag{C3}$$

At $r_1 = r_2 = 0$, this gives $Q_2 = |c|$, as stated in Example 4.3.

For $n > 2$ and nonparallel directions, (C3) gives the norm restricted to $\operatorname{span}\{\boldsymbol{u}_1, \boldsymbol{u}_2\}$. The full operator norm is one because the orthogonal complement is nontrivial and unchanged.

For three stages in two dimensions,

$$A_2(r_2) = \boldsymbol{v}_2\boldsymbol{v}_2^{\mathsf{T}} + r_2\boldsymbol{u}_2\boldsymbol{u}_2^{\mathsf{T}}.$$

Since $\boldsymbol{v}_i^{\mathsf{T}}\boldsymbol{v}_j = c_{ij}$, we obtain

$$\begin{aligned} g(r_2) &= c_{32}c_{21} + r_2(\boldsymbol{v}_3^{\mathsf{T}}\boldsymbol{u}_2)(\boldsymbol{u}_2^{\mathsf{T}}\boldsymbol{v}_1) \\ &= c_{32}c_{21} + r_2(c_{31} - c_{32}c_{21}). \end{aligned}$$

The second equality follows by inserting $I = \boldsymbol{v}_2\boldsymbol{v}_2^{\mathsf{T}} + \boldsymbol{u}_2\boldsymbol{u}_2^{\mathsf{T}}$ into $\boldsymbol{v}_3^{\mathsf{T}}\boldsymbol{v}_1$. Minimizing the absolute value of this affine function over $[-\eta, \eta]$ gives (4.7).

Appendix D Perturbation and search estimates

The perturbation argument compares the two quadratic models in compatible energy coordinates. Under (4.10), define

$$F_B = B^{1/2}, \qquad F_H = T^{1/2}B^{1/2}.$$

Then

$$F_B^{\mathsf{T}}F_B = B, \qquad F_H^{\mathsf{T}}F_H = H.$$

Thus these matrices map the corresponding energy norms isometrically to the Euclidean norm.

Although F_H need not be the principal square root of H, the matrix $F_H H^{-1/2}$ is orthogonal. Consequently, conjugation using F_H gives the same block singular values as conjugation using $H^{1/2}$.

Proof of Theorem 4.5 For one prescribed unit direction $\boldsymbol{d}$, define

$$q_B = \boldsymbol{d}^{\mathsf{T}}B\boldsymbol{d}, \qquad \boldsymbol{u} = \frac{B^{1/2}\boldsymbol{d}}{\sqrt{q_B}}, \qquad \theta = \boldsymbol{u}^{\mathsf{T}}T\boldsymbol{u},$$

and

$$\boldsymbol{v} = \frac{T^{1/2}\boldsymbol{u}}{\sqrt{\theta}}, \qquad t = \frac{\tau q_B}{2}.$$

Both $\boldsymbol{u}$ and $\boldsymbol{v}$ are unit vectors. In the coordinates induced by F_B and F_H, the stage factors are

$$\begin{aligned} A_B &= I - (1 - r_B)\boldsymbol{u}\boldsymbol{u}^{\mathsf{T}}, \quad r_B = \frac{1-t}{1+t}, \\ A_H &= I - (1 - r_H)\boldsymbol{v}\boldsymbol{v}^{\mathsf{T}}, \quad r_H = \frac{1-\theta t}{1+\theta t}. \end{aligned}$$

The assumption $\|T - I\|_2 = \delta < 1$ implies

$$1 - \delta \le \theta \le 1 + \delta.$$

Therefore,

$$
\begin{aligned}
|r_H - r_B| &= \frac{2t|\theta - 1|}{(1+t)(1+\theta t)} \\
&\le \frac{2t\delta}{(1-\delta)(1+t)^2} \le \frac{\delta}{2(1-\delta)}.
\end{aligned}
\tag{D1}
$$

The last inequality uses $2t/(1+t)^2 \le 1/2$.

To compare the rank-one projections, put

$$
C = T^{1/2}, \qquad c_- = \sqrt{1-\delta}, \qquad c_+ = \sqrt{1+\delta}.
$$

The eigenvalues of C lie in $[c_-, c_+]$. The norm of the difference of the projections onto two unit vectors equals the sine of the angle between them. Hence

$$
\begin{aligned}
&\left\| \boldsymbol{v}\boldsymbol{v}^{\mathsf{T}} - \boldsymbol{u}\boldsymbol{u}^{\mathsf{T}} \right\|_2 \\
&\quad = \frac{\left\| (I - \boldsymbol{u}\boldsymbol{u}^{\mathsf{T}}) C \boldsymbol{u} \right\|_2}{\|C\boldsymbol{u}\|_2} \\
&\quad \le \frac{\left\| C - \frac{1}{2}(c_- + c_+) I \right\|_2}{c_-} \\
&\quad \le \frac{c_+ - c_-}{2c_-} = \frac{\delta}{c_-(c_+ + c_-)} \le \frac{\delta}{2(1-\delta)}.
\end{aligned}
\tag{D2}
$$

Using

$$
\begin{aligned}
A_H - A_B = (r_H - r_B)\boldsymbol{v}\boldsymbol{v}^{\mathsf{T}} \\
+ (1 - r_B)\left(\boldsymbol{u}\boldsymbol{u}^{\mathsf{T}} - \boldsymbol{v}\boldsymbol{v}^{\mathsf{T}} \right)
\end{aligned}
$$

and $|1 - r_B| \le 2$, we conclude that

$$
\|A_H - A_B\|_2 \le \frac{3\delta}{2(1-\delta)}.
$$

Apply this estimate to every stage. Writing $S_H = A_{H,s} \cdots A_{H,1}$ and $S_B = A_{B,s} \cdots A_{B,1}$, the telescoping identity is

$$
S_H - S_B = \sum_{i=1}^{s} A_{H,s} \cdots A_{H,i+1} (A_{H,i} - A_{B,i}) A_{B,i-1} \cdots A_{B,1},
$$

where empty products are identities. Both families are nonexpansive, so taking norms gives

$$
\left\| A_{H,s} \cdots A_{H,1} - A_{B,s} \cdots A_{B,1} \right\|_2 \le \frac{3s\delta}{2(1-\delta)}.
$$

The reverse triangle inequality then bounds the difference of their operator norms by the same quantity. These norms are Q_H and Q_B, respectively. Since both belong to $[0,1]$, their difference is also bounded by one. Combining the two bounds proves (4.11). □

Proof of Corollary 4.6 For brevity, write $E = E_s(\delta)$ and suppress the fixed direction sequence $\mathcal{D}$. Continuity and compactness give a minimizer $\boldsymbol{\tau}_H^* \in \Theta$ for the H-model. Theorem 4.5 and the assumed optimization accuracy imply

$$
\begin{aligned}
Q_H(\widehat{\boldsymbol{\tau}}) &\le Q_B(\widehat{\boldsymbol{\tau}}) + E \\
&\le \min_{\boldsymbol{\tau} \in \Theta} Q_B(\boldsymbol{\tau}) + \varepsilon_{\mathrm{opt}} + E \\
&\le Q_B(\boldsymbol{\tau}_H^*) + \varepsilon_{\mathrm{opt}} + E \\
&\le Q_H(\boldsymbol{\tau}_H^*) + \varepsilon_{\mathrm{opt}} + 2E.
\end{aligned}
$$

This proves (4.12).

For the baseline comparison,

$$\begin{aligned} Q_H(\widehat{\boldsymbol{\tau}}) &\le Q_B(\widehat{\boldsymbol{\tau}}) + E \\ &\le Q_B(\boldsymbol{\tau}^{\text{base}}) - \gamma + E \\ &\le Q_H(\boldsymbol{\tau}^{\text{base}}) - \gamma + 2E, \end{aligned}$$

which is (4.13). □

D.1 Logarithmic-grid accuracy

Proof of Proposition 6.1 For $r_i, \widetilde{r}_i \in [-1, 1]$, both $A_i(r_i)$ and $A_i(\widetilde{r}_i)$ have norm at most one. Moreover,

$$A_i(r_i) - A_i(\widetilde{r}_i) = (r_i - \widetilde{r}_i)\boldsymbol{u}_i\boldsymbol{u}_i^{\mathsf{T}},$$

so

$$\|A_i(r_i) - A_i(\widetilde{r}_i)\|_2 = |r_i - \widetilde{r}_i|.$$

Telescoping the two ordered products therefore gives

$$\|S_s(\boldsymbol{r}) - S_s(\widetilde{\boldsymbol{r}})\|_2 \le \sum_{i=1}^{s} |r_i - \widetilde{r}_i|.$$

The reverse triangle inequality proves (6.5).

For the common-scale family, put

$$b_i(t) = \frac{\exp(t)a_iq_i}{2}, \qquad r_i(t) = \frac{1 - b_i(t)}{1 + b_i(t)}.$$

Since $b_i'(t) = b_i(t) > 0$,

$$|r_i'(t)| = \frac{2b_i(t)}{(1 + b_i(t))^2} \le \frac{1}{2}.$$

Combining this bound with (6.5) yields

$$|f(t) - f(\widetilde{t})| \le \sum_{i=1}^{s} |r_i(t) - r_i(\widetilde{t})| \le \frac{s}{2}|t - \widetilde{t}|,$$

proving (6.6).

Let t_* minimize f on $[\log h_-, \log h_+]$. Because the grid contains both endpoints and has maximum spacing Δ, some grid point t_g satisfies

$$|t_g - t_*| \le \frac{\Delta}{2}.$$

Thus

$$f(t_g) \le f(t_*) + \frac{s\Delta}{4}.$$

Let $\widetilde{f}$ denote the computed values, with $|\widetilde{f}(t) - f(t)| \le \varepsilon_{\text{ev}}$ at each grid point. By the definition of $\widehat{t}$,

$$\begin{aligned} f(\widehat{t}) &\le \widetilde{f}(\widehat{t}) + \varepsilon_{\text{ev}} \\ &\le \widetilde{f}(t_g) + \varepsilon_{\text{ev}} \\ &\le f(t_g) + 2\varepsilon_{\text{ev}} \\ &\le f(t_*) + \frac{s\Delta}{4} + 2\varepsilon_{\text{ev}}. \end{aligned}$$

This proves (6.7). The singleton-interval case follows with $\Delta = 0$. □

Appendix E Nonlinear convergence and inexact updates

E.1 Local convergence

Proof of Theorem 4.7 Step 1: existence and differentiation of the stage maps. For stage i, let F_i be defined by (4.14). Since V is C^2 near $\boldsymbol{x}_*$, the function F_i is continuously differentiable near $(\boldsymbol{x}_*, 0)$. Moreover,

$$F_i(\boldsymbol{x}_*, 0) = 0,$$

and

$$\partial_\beta F_i(\boldsymbol{x}_*, 0) = 1 + \frac{\tau_i}{2} \boldsymbol{d}_i^\mathsf{T} H_* \boldsymbol{d}_i > 0.$$

The implicit function theorem therefore gives a unique local C^1 solution branch $\beta_i = \beta_i(\boldsymbol{x})$ satisfying $\beta_i(\boldsymbol{x}_*) = 0$.

Since

$$D_{\boldsymbol{x}} F_i(\boldsymbol{x}_*, 0) = \boldsymbol{d}_i^\mathsf{T} H_*,$$

implicit differentiation yields

$$D\beta_i(\boldsymbol{x}_*) = -\frac{\boldsymbol{d}_i^\mathsf{T} H_*}{1 + \tau_i \boldsymbol{d}_i^\mathsf{T} H_* \boldsymbol{d}_i / 2}.$$

Thus the local stage map

$$T_i(\boldsymbol{x}) = \boldsymbol{x} + \tau_i \beta_i(\boldsymbol{x}) \boldsymbol{d}_i$$

satisfies

$$DT_i(\boldsymbol{x}_*) = I - \frac{\tau_i}{1 + \tau_i \boldsymbol{d}_i^\mathsf{T} H_* \boldsymbol{d}_i / 2} \boldsymbol{d}_i \boldsymbol{d}_i^\mathsf{T} H_*.$$

This is precisely the quadratic stage matrix (3.2) evaluated at $H = H_*$.

Step 2: differentiation of the block map. Every stage map fixes $\boldsymbol{x}_*$. The chain rule therefore gives

$$DT_{\mathrm{blk}}(\boldsymbol{x}_*) = P_s(H_*, \mathcal{D}; \boldsymbol{\tau}),$$

and differentiability proves (4.15).

Step 3: an invariant neighborhood and a linear rate. Suppose that the directions span $\mathbb{R}^n$. By Theorem 4.1,

$$Q_* := Q_{H_*}(\mathcal{D}; \boldsymbol{\tau}) < 1.$$

Fix $q \in (Q_*, 1)$. Norm equivalence and the remainder in (4.15) imply that, for sufficiently small $\delta > 0$,

$$\begin{aligned} &\|T_{\mathrm{blk}}(\boldsymbol{x}_* + \boldsymbol{e}) - \boldsymbol{x}_*\|_{H_*} \\ &\quad \le Q_* \|\boldsymbol{e}\|_{H_*} + (q - Q_*) \|\boldsymbol{e}\|_{H_*} = q \|\boldsymbol{e}\|_{H_*} \end{aligned}$$

whenever $\|\boldsymbol{e}\|_{H_*} < \delta$. Choose this ball within the domain of the composed local map; such a neighborhood exists because the finitely many stage maps are continuous and fix $\boldsymbol{x}_*$. The estimate makes the ball invariant and proves (4.16). Iteration yields

$$\left\| \boldsymbol{x}^k - \boldsymbol{x}_* \right\|_{H_*} \le q^k \left\| \boldsymbol{x}^0 - \boldsymbol{x}_* \right\|_{H_*}.$$

Step 4: the case of a vanishing block derivative. If $P_s(H_*, \mathcal{D}; \boldsymbol{\tau}) = 0$, the linear term vanishes and

$$\|T_{\mathrm{blk}}(\boldsymbol{x}) - \boldsymbol{x}_*\|_{H_*} = o\Big(\|\boldsymbol{x} - \boldsymbol{x}_*\|_{H_*} \Big).$$

Hence a sufficiently close iteration either reaches $\boldsymbol{x}_*$ in finitely many blocks or has a superlinear error ratio. □

E.2 A supporting smooth-convex convergence argument

The following result records sufficient conditions for convergence of the prescribed-direction iteration on smooth convex objectives. It clarifies the role of uniform direction coverage and is not presented as a principal contribution. General smooth discrete-gradient convergence estimates are discussed in [2].

Theorem E.1 (Convergence under uniform direction coverage) *Let $V \in C^1(\mathbb{R}^n)$ be convex, and suppose that*

$$K = \left\{ \boldsymbol{x} \in \mathbb{R}^n : V(\boldsymbol{x}) \le V(\boldsymbol{y}^0) \right\}$$

is compact. Let $(\widehat{\boldsymbol{d}}^{\,m})_{m=0}^{\infty}$ be a sequence of unit directions, and solve (B1) exactly at every update.

Suppose that there exist constants $0 < \tau_- \le \tau_+ < \infty$, an integer $N \ge 1$, and $\alpha > 0$ such that

$$\begin{aligned} \tau_- \le \widehat{\tau}_m \le \tau_+, \\ \sum_{\ell=0}^{N-1} \widehat{\boldsymbol{d}}^{\,m+\ell} (\widehat{\boldsymbol{d}}^{\,m+\ell})^{\mathsf{T}} \succeq \alpha I \qquad (m \ge 0). \end{aligned} \tag{E1}$$

Then $X_* := \operatorname{argmin}_{\mathbb{R}^n} V$ is nonempty and

$$\begin{aligned} \nabla V(\boldsymbol{y}^m) &\to \boldsymbol{0}, \\ V(\boldsymbol{y}^m) &\to V_* := \min_{\mathbb{R}^n} V, \\ \operatorname{dist}(\boldsymbol{y}^m, X_*) &\to 0. \end{aligned} \tag{E2}$$

If the minimizer is unique, then $\boldsymbol{y}^m \to \boldsymbol{x}_*$.

Proof Step 1: solvability, boundedness, and existence of minimizers. By Proposition B.1, every update exists uniquely and decreases the objective. Thus all iterates belong to K.

Continuity and compactness imply that V attains its minimum on K. This minimum is global, because every point outside K has objective value greater than $V(\boldsymbol{y}^0)$. Hence X_* is nonempty.

Step 2: vanishing increments and directional derivatives. Summing (B2) gives

$$\sum_{m=0}^{\infty} \left\| \boldsymbol{y}^{m+1} - \boldsymbol{y}^m \right\|_2^2 \le \tau_+ \big(V(\boldsymbol{y}^0) - V_* \big) < \infty.$$

Therefore,

$$\left\| \boldsymbol{y}^{m+1} - \boldsymbol{y}^m \right\|_2 \to 0.$$

Since

$$|\widehat{\beta}_m| = \frac{\left\| \boldsymbol{y}^{m+1} - \boldsymbol{y}^m \right\|_2}{\widehat{\tau}_m},$$

the lower step bound implies $\widehat{\beta}_m \to 0$.

The set K is convex because V is convex. Every segment joining consecutive iterates therefore lies in K, where ∇V is uniformly continuous. The integral equation gives

$$\widehat{\beta}_m = -\int_0^1 (\widehat{\boldsymbol{d}}^{\,m})^{\mathsf{T}} \nabla V\Big(\boldsymbol{y}^m + t(\boldsymbol{y}^{m+1} - \boldsymbol{y}^m) \Big) \, \mathrm{d}t.$$

Let

$$\omega_K(r) = \sup_{\substack{\boldsymbol{u},\boldsymbol{v}\in K \\ \|\boldsymbol{u}-\boldsymbol{v}\|_2\le r}} \|\nabla V(\boldsymbol{u}) - \nabla V(\boldsymbol{v})\|_2 .$$

Uniform continuity gives $\omega_K(r) \to 0$ as $r \downarrow 0$. The integral equation implies

$$\left|(\widehat{\boldsymbol{d}}^{\,m})^{\mathsf{T}}\nabla V(\boldsymbol{y}^m) + \widehat{\beta}_m\right| \le \omega_K\left(\left\|\boldsymbol{y}^{m+1} - \boldsymbol{y}^m\right\|_2\right) \to 0.$$

Since $\widehat{\beta}_m \to 0$, it follows that $(\widehat{\boldsymbol{d}}^{\,m})^{\mathsf{T}}\nabla V(\boldsymbol{y}^m) \to 0$.

Step 3: use of uniform direction coverage. For each fixed $\ell \ge 0$,

$$\left\|\boldsymbol{y}^{m+\ell} - \boldsymbol{y}^m\right\|_2 \to 0.$$

Applying the preceding directional estimate at index $m+\ell$ and using uniform continuity once more gives

$$(\widehat{\boldsymbol{d}}^{\,m+\ell})^{\mathsf{T}}\nabla V(\boldsymbol{y}^m) \to 0.$$

Consequently, (E1) yields

$$\alpha\left\|\nabla V(\boldsymbol{y}^m)\right\|_2^2 \le \sum_{\ell=0}^{N-1}\left|(\widehat{\boldsymbol{d}}^{\,m+\ell})^{\mathsf{T}}\nabla V(\boldsymbol{y}^m)\right|^2 \to 0.$$

Thus the gradient converges to zero.

Step 4: convergence of values and distance to the solution set. Every accumulation point is stationary by continuity of the gradient and is therefore globally minimizing by convexity. The nonincreasing objective values have a subsequence converging to V_*, so the entire sequence of values converges to V_*.

If the distance to X_* did not tend to zero, there would be a subsequence whose distance is bounded away from zero. Compactness would give a further convergent subsequence with a limit outside X_*, contradicting the preceding characterization of accumulation points. Finally, if $X_* = \{\boldsymbol{x}_*\}$, distance convergence is convergence to $\boldsymbol{x}_*$. □

E.3 Inexact scalar solves

Consider a quadratic block with inexact updates

$$\boldsymbol{e}_i = M_i\boldsymbol{e}_{i-1} + \boldsymbol{\zeta}_i.$$

Expansion of the recurrence gives

$$\boldsymbol{e}_s = P_s\boldsymbol{e}_0 + \sum_{i=1}^{s} M_s\cdots M_{i+1}\boldsymbol{\zeta}_i,$$

where the trailing product is the identity when $i = s$. Every trailing product is nonexpansive in the H-norm. Hence

$$\|\boldsymbol{e}_s\|_H \le Q_H(\mathcal{D};\boldsymbol{\tau})\,\|\boldsymbol{e}_0\|_H + \sum_{i=1}^{s}\|\boldsymbol{\zeta}_i\|_H ,$$

which proves (6.8).

Suppose the computed scalar coefficient is $\widehat{\beta}_i$, with residual

$$\rho_i = \left(1 + \frac{\tau_i q_i}{2}\right) \widehat{\beta}_i + \boldsymbol{d}_i^{\mathsf{T}} H \boldsymbol{e}_{i-1}.$$

Let β_i^{ex} denote the exact scalar coefficient for the current input error $\boldsymbol{e}_{i-1}$:

$$\beta_i^{\mathrm{ex}} = -\frac{\boldsymbol{d}_i^{\mathsf{T}} H \boldsymbol{e}_{i-1}}{1 + \tau_i q_i/2}.$$

The residual definition gives

$$\widehat{\beta}_i - \beta_i^{\mathrm{ex}} = \frac{\rho_i}{1 + \tau_i q_i/2}.$$

Thus the implemented update can be written as

$$\boldsymbol{e}_i = M_i \boldsymbol{e}_{i-1} + \frac{\tau_i \rho_i}{1 + \tau_i q_i/2} \boldsymbol{d}_i.$$

Therefore,

$$\boldsymbol{\zeta}_i = \frac{\tau_i \rho_i}{1 + \tau_i q_i/2} \boldsymbol{d}_i, \qquad \|\boldsymbol{\zeta}_i\|_H = \frac{\tau_i \sqrt{q_i}}{1 + \tau_i q_i/2} |\rho_i|.$$

This identity separates the error in solving the scalar equation from the error in selecting the step sequence.

Appendix F Coefficient sequences used for exact-arithmetic checks

Table F1 gives the reference step sequences for the exact-arithmetic checks in Section 5.4. The Q2 sequence is a rational witness of zero amplification. For Q3, every displayed decimal is interpreted as an exact rational number. The guarantee $Q_9^2 < 10^{-16}$ stated here concerns this displayed reference Q3 sequence. Its last digits differ from the runtime sequence used for the floating-point results in Section 5.4. Thus the certificate can be checked from the coefficients printed in this table and the exact-arithmetic routines in the accompanying source file.

For each sequence, form the rational stage matrices (3.2) and multiply them in the prescribed order. For Q2 with $\rho = 19/20$, all entries of P_9 vanish, so the block eliminates every initial quadratic error in exact arithmetic.

For Q3 with $\rho = 4/5$, define

$$C = 10^{-16} H_\rho - P_9^{\mathsf{T}} H_\rho P_9.$$

All three leading principal determinants of C are strictly positive. Since C is symmetric, Sylvester's criterion gives

$$P_9^{\mathsf{T}} H_\rho P_9 \prec 10^{-16} H_\rho.$$

Table F1 Reference step sequences for the exact-arithmetic checks. Directions follow the coordinate cycle in the corresponding dimension

Stage	Q2: rational τ_i	Q3: exact decimal τ_i
1	2	2
2	$\frac{2}{19}$	0.46275238905830002
3	$\frac{2}{19}$	37.999579832241281
4	$\frac{9075101946559600}{1425322381566439}$	1.9999999975760792
5	38	10.554377170650737
6	38	0.27828779723725983
7	38	37.999117955280809
8	38	18.895411948011084
9	2	2

Consequently,

$$Q_9^2 < 10^{-16}.$$

These finite algebraic checks establish bounds over all initial errors without sampling error directions.

Appendix G Nonconvex coefficients and numerical diagnostics

Table G1 lists the computed profile in its application order. The generation formula is (5.10). In floating-point arithmetic $w_0 = 1.0006172839506173$ and $w_1 = 0.012757066520140505$.

Table G1 Computed Chebyshev-root coefficients in application order. The effective steps are a_i^{ch} for Rosenbrock and $3a_i^{\mathrm{ch}}$ for the nonsmooth experiment

i	a_i^{ch}
1	0.80692251843773344
2	0.094783329457836088
3	0.035651225814005018
4	0.019370060592785471
5	0.012749196645667869
6	0.009501497769222678
7	0.007762582775268957 2
8	0.006342305701899796
9	0.006425357937288551 1

Table G2 compares the nonsmooth results at three root tolerances. Rules E and Ch reach the 10^{-5} target from every initial point at every tolerance, and their ordering by N_f is unchanged. Rule C exhausts the 3×10^5 budget in all these runs. Tightening the root tolerance also increases the cost per update, so its effect cannot be inferred from budget-limited endpoints alone. At tolerance 10^{-8}, rule C from (A) records 19 numerical zero updates. No such events occur in the other runs. No solver failures occur in any of the 49 runs. The separate extended-budget diagnostic from (A), reported in Section 5.9, reaches the original target at tolerance 10^{-12}.

Table G2 Root-tolerance diagnostics for the nonsmooth problem. N_f is the target-attainment count for E and Ch. Rule C reaches the common budget of 300000 in every row; its final objective is listed instead. Absolute and relative root tolerances are both equal to the value in the second column

Initial	Tolerance	Final V, C	N_f, E	N_f, Ch
(A)	10^{-8}	1.75498×10^{-5}	71954	21047
(A)	10^{-10}	1.81990×10^{-5}	74222	21566
(A)	10^{-12}	1.89405×10^{-5}	76995	21997
(B)	10^{-8}	1.79583×10^{-5}	93135	109048
(B)	10^{-10}	1.84909×10^{-5}	95824	111855
(B)	10^{-12}	1.92581×10^{-5}	98960	113862
(C)	10^{-8}	1.78428×10^{-5}	71973	22305
(C)	10^{-10}	1.83125×10^{-5}	74225	22869
(C)	10^{-12}	1.90644×10^{-5}	76978	23316
(D)	10^{-8}	1.78382×10^{-5}	68355	16798
(D)	10^{-10}	1.83527×10^{-5}	70567	17248
(D)	10^{-12}	1.91083×10^{-5}	73291	17598