THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

A decomposition-based warm-start method for stochastic programming

Citation for published version:

Colombo, M & Grothey, A 2013, 'A decomposition-based warm-start method for stochastic programming', *Computational optimization and applications*, vol. 55, no. 2, pp. 311-340. <https://doi.org/10.1007/s10589-012-9530-7>

Digital Object Identifier (DOI):

[10.1007/s10589-012-9530-7](https://doi.org/10.1007/s10589-012-9530-7)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Early version, also known as pre-print

Published In:

Computational optimization and applications

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

A decomposition-based crash-start for stochastic programming

Marco Colombo · Andreas Grothey

the date of receipt and acceptance should be inserted later

Abstract In this paper we propose a crash-start technique for interior point methods applicable to multi-stage stochastic programming problems. The main idea is to generate an initial point for the interior point solver by decomposing the barrier problem associated with the deterministic equivalent at the second stage and using a concatenation of the solutions of the subproblems as a warm-starting point for the complete instance. We analyse this scheme and produce theoretical conditions under which the warm-start iterate is successful. We describe the implementation within the OOPS solver and the results of the numerical tests we performed.

Keywords Stochastic Programming, Interior Point Methods, Warm-starting

1 Introduction

Stochastic Programming Problems [5, 17] are frequently solved by formulating the deterministic equivalent and applying a standard solver to the resulting problem. However for problems formulated over large scenario sets, and in particular for multistage problems, the deterministic equivalent quickly reaches enormous sizes. Due to the special structure of these problems, they can still be efficiently solved by either approaches based on Benders Decomposition [3, 19, 24] or by structure exploiting Interior Point Methods (IPM) [6, 13, 23]. Common to these approaches is that they obtain the deterministic equivalent by discretising the (in general) continuous distribution of the underlying uncertain parameters. An appealing idea is therefore to work initially with a coarser discretisation of the probability distribution (leading to a smaller

This research was supported by the Engineering and Physical Sciences Research Council under grant EP/E036910/1.

School of Mathematics and Maxwell Institute, The University of Edinburgh. E-mail: A.Grothey@ed.ac.uk

problem) and using its solution to crash-start the full problem. Despite seeming an obvious approach, this seems to be a novel idea; probably the reason for this is that both Benders Decomposition and IPM have difficulties in exploiting an advanced crash-start point efficiently. In a previous work [7], the authors have introduced a crash-start approach for stochastic programming based on solving a reduced problem which contains only a trivial fraction of the scenarios. That work provided measurable evidence that a crash-starting scheme based on solving reduced trees within IPMs is workable and leads to measurable efficiency gains.

The approach described in [7] is efficient albeit crude. One of its drawbacks is that the initial point, built from a reduced scenario tree, may be far away from primal-dual feasibility in some of the subproblems. This is particularly the case for trees that are wide and deep, as the representative scenario may not be able to convey all the information of the sub-tree from which it is chosen. On the other hand, a primal or dual feasible point can be easily constructed by performing a single iteration of a decomposition scheme such as Benders Decomposition or Lagrangian Relaxation. This has already been used to construct a crash-start point for network flow problems in [11].

In this paper we combine the two ideas: In a first step the problem is solved on a reduced scenario tree to obtain estimates for the first stage decision. These are then used in a second step to solve recourse problems for each scenario, effectively performing half an iteration of Benders Decomposition. From the resulting solutions we construct a crash-start point for the IPM on the full problem. As constructed the resulting crash-start point is always primal feasible. We show that under suitable conditions on the proximity of the reduced scenario tree to the full tree, the resulting crash-start point is successful as IPM warm-start in the sense of [26, 11], that is a full step in the initial iteration can be taken, absorbing any remaining infeasibility. Numerical results show that significant performance gains can be obtained through this scheme.

At this point we should clarify our terminology: This paper is concerned with the issue of *crash-starting* an interior point method, that is, finding a good starting point that is preferable to Mehrotra's standard starting point heuristic by making use of the structure of the problem. We will analyse this method using results from *warm-starting* IPM in the sense of [26, 11], that is speeding up the solution of the target problem by using an iterate obtained from the solution process of a nearby problem. In that sense we are applying IPM warm-starting technology to the problem of finding a good crash-start point for stochastic programming problems.

The paper is structured as follows. In Section 2 we first review relevant background for Stochastic Programming and warm-starting of Interior Point Methods. In Section 3 we introduce the proposed algorithm, which we analyse in Section 4. In Section 5, we present the numerical results obtained with our implementation. Finally, in Section 6 we draw our conclusions and highlight directions of future research.

2 Stochastic Programming and Interior Point Methods

Stochastic programming [5, 17] models uncertainty through the analysis of possible future scenarios. In stochastic programming, the uncertain environment is described through a stochastic process which is obtained from historical data or conjectured according to some prescribed properties. The continuous process is usually further approximated by a discrete distribution in order to obtain a computationally amenable description. This is done by generating a finite, but usually very large, number of scenarios that represent an approximate description of the possible outcomes. The discrete stochastic process can be represented as an *event tree*: each node of the tree denotes a point in time when a realisation of the random process becomes known and a *recourse* decision is taken.

A linear two-stage stochastic programming problem can be formalised as

$$\begin{aligned} \min_x \quad & c^\top x + \mathbb{E}_\xi[Q(x, \xi)] \quad \text{where} \quad Q(x, \xi) = \min_y c(\xi)^\top y(\xi) \\ \text{s.t.} \quad & W_0 x = h_0 \quad \text{s.t.} \quad T x + W(\xi) y(\xi) = h(\xi) \\ & x \geq 0 \quad \text{s.t.} \quad y(\xi) \geq 0, \end{aligned} \tag{1}$$

where the random variable ξ captures the uncertainty. The optimal solution of problem (1) describes a continuous stochastic process which cannot be solved for directly: as mentioned above, it has to be discretised and rewritten in a form which is viable for computation. For our purposes, we rely on the *deterministic equivalent formulation*. To formulate the deterministic equivalent we adopt the following notation: Let a scenario (π_i, ξ_i) be given by the data $(c_i, W_i, h_i) = (c(\xi_i), W(\xi_i), h(\xi_i))$. A stochastic programming problem is defined by a scenario set (or a *tree*) $\mathcal{T} = \{(\pi_i, \xi_i)_i\}$, so that $P(\mathcal{T})$ denotes the problem

$$\begin{aligned} \min \quad & c^\top x + \sum_{i \in \mathcal{T}} \pi_i c_i^\top y_i, \quad \text{s.t.} \quad W_0 x = h_0 \\ & T x + W_i y_i = h_i, \quad i \in \mathcal{T} \\ & x, y_i \geq 0. \end{aligned} \tag{P(\mathcal{T})}$$

We can introduce a multi-stage decision structure by considering several stages of recourse decisions $(x, y^1, \dots, y^T)$ and a tree of scenarios. For the purposes of this paper we will assume that the stages $2, \dots, T$ are represented as one, that is y_i in $(P(\mathcal{T}))$ denotes the vector $y_i = (y_i^1, \dots, y_i^T)$ and the recourse matrices W_i are structured matrices representing the final $T - 1$ stages of the problem.

Several solution methods for stochastic linear programs have been presented in the literature [3, 19, 24]. These often rely on a variant of Benders' decomposition, such as the L-shaped method. They do not require the explicit generation of the deterministic equivalent problem. An entirely different approach, based on interior point methods, exploits the fact that the augmented system matrix arising from applying an IPM to the deterministic equivalent of a large-scale multi-stage stochastic program displays a nested block structure that can be efficiently exploited in the linear algebra.

In this paper we adopt a decomposition-like scheme to build a crash-start iterate that we then use to solve the deterministic equivalent through an interior point method.

2.1 Interior Point Methods

Consider the linear programming problem in standard form

$$\min c^\top x \quad \text{s.t.} \quad Ax = b, \quad x \geq 0, \quad (2)$$

where $A \in \mathcal{R}^{m \times n}$ is full rank, $x, c \in \mathcal{R}^n$ and $b \in \mathcal{R}^m$. For the purposes of this paper, problem (2) corresponds to the deterministic equivalent generated from a given event tree $\mathcal{T}$, and we will refer to it as the *complete problem*. Interior point methods work with the barrier problem

$$\min c^\top x - \mu \sum \ln x_i \quad \text{s.t.} \quad Ax = b, \quad (3)$$

instead of (2). Problem (3) is a family of strictly convex problems, parameterised by $\mu > 0$, whose unique solution approaches the solution to the original problem as $\mu \rightarrow 0$. The trajectory of solutions to (3) for different values of μ is the *central path*. Interior point methods apply a damped Newton method to the optimality conditions of (3), yielding the direction-finding problem

$$\begin{bmatrix} A & 0 & 0 \\ 0 & A^\top & I \\ Z & 0 & X \end{bmatrix} \begin{bmatrix} \Delta x \\ \Delta y \\ \Delta z \end{bmatrix} = \begin{bmatrix} b - Ax \\ c - A^\top y - z \\ -XZe + \mu e \end{bmatrix} = \begin{bmatrix} \xi_b \\ \xi_c \\ \xi_\mu \end{bmatrix}, \quad (4)$$

which needs to be solved with a specified μ for a search direction $(\Delta x, \Delta y, \Delta z)$ at every iteration.

Path-following methods [25] globalise the Newton iteration by keeping the iterates in a neighbourhood of the central path, thus follow it in approaching the optimal solution. A possible choice for the neighbourhood is

$$\mathcal{N}_2(\theta) = \{(x, y, z) : Ax = b, A^\top y - z = c, (x, z) > 0, \|XZe - \mu e\|_2 \leq \theta\mu, n\mu = x^\top z\}$$

for some $\theta > 0$. For an interior point method to be successful, it is essential that centrality is maintained throughout the iterations, until the optimal partitioning is identified. Approaching a non-optimal vertex too soon can hurt the performance very badly, as the algorithm will strive to recenter the iterate before being able to progress towards optimality. For this reason, the choice of the first iterate is a key issue in an implementation of an interior point method.

In practice, the starting point is generally computed by Mehrotra's starting point heuristic [18], which is considered to be computationally effective. In this heuristic, the starting point is found by solving two least squares problems which attempt to satisfy primal and dual constraints; this point is then shifted away from the boundary towards the positive orthant, in order to satisfy $(x, z) > 0$. Alternatively when a closely related problem has been solved

previously, information gained from the solution process of this problem can be used to construct an initial point for the new instance (a process known as warmstarting). Unlike the situation for active set methods (such as the simplex method), when warmstarting IPMs it is not a good idea to use the solution of the previous problem directly as initial iterate. For IPMs this is observed to lead to “jamming” [16,12], that is the search direction points outside the feasible region leading to very small steps. Theoretical insights [26, 11] suggest that the best warmstart point is near the central path for a not-too-small value of μ . Further a *modification* step is necessary that absorbs any primal-dual infeasibility that may be present in the warmstart point. Let $d = (A, b, c)$ be the data describing the first problem in sequence, and likewise $\bar{d} = (\bar{A}, \bar{b}, \bar{c}) = d + (\Delta A, \Delta b, \Delta c)$ for the second problem that is to be warm-started. Let $\mathcal{N}_2$ and $\bar{\mathcal{N}}_2$ be corresponding neighbourhoods of the central paths of the two problems. Further we define a norm $\|d\| := \max(\|A\|_2, \|b\|_2, \|c\|_2)$ on the space of problem instances. With $\mathcal{B}$ be the set of ill-posed (infeasible or unbounded) problem instances, and $\rho(d) = \inf\{\|\Delta d\| : d + \Delta d \in \mathcal{B}\}$ the distance to ill-posedness we can define the Renegar condition number

$$C(d) := \|d\|/\rho(d)$$

which has been suggested as a useful indicator of the difficulty of a problem instance when solved by IPM. Then a typical warmstarting result is

Proposition 1 ([26, Prop. 4.2]) *Let $(x, y, s) \in \mathcal{N}_2(\theta_0)$ be given with $\mu = x^T s/n$ and suppose that $(\Delta x, \Delta y, \Delta s)$ is obtained by a Weighted-Least-Square modification step. Further let $\theta > \theta_0$, and $\xi \in (0, \theta - \theta_0)$. Assuming that*

$$\delta_{bc} \leq \frac{\theta - \theta_0 - \xi}{(2n+1)C(d)} \quad \text{and} \quad \mu \geq \frac{\|d\|}{\xi} 4C(d)^2 \delta_{bc},$$

where $\delta_{bc} = (\|\Delta c\|_2 + 2C(d)\|\Delta b\|_2)/\|d\|$, then $(x + \Delta x, y + \Delta y, s + \Delta s) \in \bar{\mathcal{N}}_2(\theta)$.

The above result assumes that $\Delta A = 0$, although similar results can be obtained in the general case. A common strategy in practice is to take a point on (or near) the central path for the original problem and use this (after a modification step) as the warmstart point in the new problem. Different modification steps have been suggested in [26,11]. The theoretical analysis in Section 4 is based on using the Weighted Least Squares (WLS) modification step of [26], although similar results could be obtained for other choices.

3 Decomposition scheme

In the crash-start approach of [7], we build a reduced tree $\mathcal{T}^R$, from the event tree $\mathcal{T}$ associated with the problem, by picking a small number of available scenarios. An approximate solution to the deterministic equivalent corresponding to $\mathcal{T}^R$ is computed and extended to construct a crash-start iterate for the complete problem. The advantage of the scheme, apart from its simplicity, is

that the reduced problem is much smaller than the complete formulation, and hence much easier to solve. Still, it provides sufficient information to generate an advanced starting point for the complete problem.

Despite its rather crude nature the scheme of [7] can be surprisingly efficient, even when using a very low number of scenarios in the reduced tree. The reason for this seems to be that for many stochastic programming problems the set of active constraints is very similar for most scenarios. Figure 1 visualises this for some standard two-stage test problems. In these plots each horizontal line corresponds to a scenario: for each variable within a scenario, a pale dot was printed if the variable was active at its lower bound at the optimal solution, a black dot otherwise. Thus, looking vertically, we can spot the difference in the active set between scenarios. It is striking to see the great

Fig. 1: Active set for problem `stocfor2`, `pltexpA2-16` and `fmx2-16`.

similarity among all of them for the problems considered. Many of the scenarios lead to the same active set, and where there are differences these are limited to a small fraction of variables. While the crash-start strategy of [7] is thus typically able to find a good starting point for the majority of scenarios, a bad guess for even a small number of scenarios can severely hamper the performance of the scheme. If a crash-start point is close to a constraint that is not active at the solution, the IPM will need many iterations to move away from such a misidentified active constraint, an effect that has been described as “jamming” or “blocking” [16, 12]. In order to avoid this we will investigate a more sophisticated crash-start procedure.

The idea is based on the observation that if we were given the optimal first stage decision x for problem $(P(\mathcal{T}))$, we could obtain the second stage components y_i by solving a separate sub-problem for each scenario. The underlying idea recalls that of Benders’ decomposition.

3.1 A Decomposition-Based Crash-start Scheme

Given a two-stage stochastic programming problem $P(\mathcal{T})$ with the associated scenario set $\mathcal{T}$, we start by considering the corresponding barrier problem:

$$\begin{aligned} \min \quad & c^\top x + \sum_{i \in \mathcal{T}} \pi_i c_i^\top y_i - \mu \sum_j \ln x_j - \mu \sum_i \sum_j \ln y_{i,j} \\ \text{s.t.} \quad & Tx + W_i y_i = h_i, \quad i \in \mathcal{T} \end{aligned} \quad (P_\mu(\mathcal{T}))$$

and its KKT conditions

$$s + \sum_i T^\top \lambda_i = c \quad (5a)$$

$$z_i + W_i^\top \lambda_i = \pi_i c_i \quad (5b)$$

$$Tx + W_i y_i = h_i \quad (5c)$$

$$XSe = \mu e \quad (5d)$$

$$Y_i Z_i e = \mu e \quad (5e)$$

For a given value $\hat{x}$ of the first stage decisions, the barrier problem $P_\mu(\mathcal{T})$ decomposes into the scenario subproblems

$$Q_{i,\mu}(\hat{x}) = \min \{ \pi_i c_i^\top y_i - \mu \sum_j \ln y_{i,j} : W_i y_i = h_i - T_i \hat{x}, y_i \geq 0 \}, \quad (P_{i,\mu}(\hat{x}))$$

with optimality conditions

$$W_i^\top \lambda_i + z_i = \pi_i c_i \quad (6a)$$

$$W_i y_i = h_i - T_i \hat{x} \quad (6b)$$

$$Y_i Z_i e = \mu e \quad (6c)$$

which together cover the conditions (5b),(5c) and (5e). Conditions (5a) and (5d) on the other hand reflect on the centrality and optimality of the chosen $\hat{x}$.

Thus, if we are given a reasonable central and optimal $\hat{x}, \hat{s}$ that are good approximations to the x, s -components of the solution to $P_\mu(\mathcal{T})$, we can use this to solve subproblems $Q_{i,\mu}(\hat{x})$. Combining $\hat{x}, \hat{s}$ with the thus obtained $(y_i^*, z_i^*, \lambda_i^*)$ for all subproblems yields a primal feasible (5a) and central (5d,e) approximation $(\bar{x}, \bar{s}, \bar{y}, \bar{z}, \bar{\lambda}) = (\hat{x}, \hat{s}, y^*, z^*, \lambda^*)$ to the solution of $P_\mu(\mathcal{T})$. Dual feasibility will depend on the quality of the chosen $(\hat{x}, \hat{s})$. In the suggested scheme we will obtain estimates $(\hat{x}, \hat{s})$ by solving $P_\mu(\mathcal{T}^R)$ for a reduced tree $\mathcal{T}^R$. We will show that if $\mathcal{T}^R$ approximates $\mathcal{T}$ sufficiently well, the constructed crash-start point $(\bar{x}, \bar{s}, \bar{y}, \bar{z}, \bar{\lambda})$ will be in an appropriate N_2 -neighbourhood for problem $P_\mu(\mathcal{T})$ and thus a sufficiently good warm-start point by Proposition 1. The resulting decomposition based crash-start algorithm is summarised as Algorithm 1.

Compared to the approach of [7], this scheme provides directly the part of the crash-start vector for every node in the tree, and so it does not require a

Algorithm 1 Decomposition-based warm-start algorithm

Require: The complete event tree $\mathcal{T}$.

- 1: Generate a reduced event tree $\mathcal{T}^R \subset \mathcal{T}$; Fix a target $\bar{\mu} > 0$.
 - 2: Solve the barrier problem on the reduced tree $P_{\bar{\mu}}(\mathcal{T}^R)$ and obtain first stage decisions and duals $(\bar{x}, \bar{s})$ from it.
 - 3: **for** $i \in \{\text{second stage nodes of } \mathcal{T}\}$ **do**
 - 4: Solve scenario sub-problem $Q_{i, \bar{\mu}}(\bar{x})$ to obtain recourse decisions and duals $(y_i^*, z_i^*, \lambda_i^*)$.
 - 5: **end for**
 - 6: Combine these to a crash start iterate $(\bar{x}, \bar{s}, \bar{y}, \bar{z}, \bar{\lambda})$ for the complete problem $P(\mathcal{T})$.
 - 7: Perform a modification step in the full problem and solve to optimality.
-

process of “expanding the solution” by copying parts of solution from extraneous nodes. As such, we expect it to provide an iterate that is closer to the central path of the complete problem: by construction the crash-start point is central and primal feasible, while dual feasibility depends on the quality of the estimate of the first-stage decisions. We will back up these claims in the following section.

4 Theoretical analysis

The aim of this section is to state conditions under which the decomposition-based crash-start procedure is successful. That is we are able to construct a point that is sufficiently primal–dual feasible and central to satisfy the conditions of Proposition 1. We will assume that the problem is a two-stage problem.

We assume that we have full recourse, so that every scenario sub-problem is feasible and bounded for all x . Indeed for the technical results we need to assume more: following the terminology of Section 2.1 we use $d_i(\hat{x}) = (W_i, h_i - T\hat{x}, c_i)$ to denote the problem data of the i -th scenario sub-problem and define the distance to ill-posedness for problem $P_i(x)$ as

$$\rho(d_i(x)) := \inf\{\|\Delta d\| : d_i + \Delta d \in \mathcal{B}\},$$

with $\mathcal{B}$ being the set of “ill-posed” data instances, and we require that

$$\underline{\rho} := \inf_x \min_{i \in \mathcal{T}} \rho(d_i(x)) > 0.$$

Further we assume that the problem data itself is bounded, i.e. $\|d_i\| \leq \overline{\|d\|}$, which gives us the existence of a bound $\overline{C(d)}$ on the Renegar condition number $C(d) = \|d\|/\rho(d)$:

$$C(d_i) \leq \overline{C(d)}, \quad \forall i \in \mathcal{T}.$$

In what follows we deal with three types of problems: The full problem $P(\mathcal{T})$, the reduced problem $P(\mathcal{T}^R)$ and the scenario subproblems $P_i(\hat{x})$. We use the convention that objects associated with the full problem, the reduced problem and the scenario subproblems have superscripts $\mathcal{T}$, R and P_i respectively (as in $\mathcal{N}_2^{\mathcal{T}}, \mathcal{N}_2^R, \mathcal{N}_2^{P_i}$). We will also use the notation such as $x_\mu(\mathcal{T})$ to refer to the x -component of the solution to problem $P_\mu(\mathcal{T})$.

In our decomposition-based crash-start method, starting from the problem of interest $P(\mathcal{T})$, we identify a reduced tree $\mathcal{T}^R$, and find a $\bar{\mu}$ -central point $(x^R, y^R, \lambda^R, s^R, z^R) \in \mathcal{N}_2^R(\theta_0)$. We take the x, s -components of this point as our estimate of the optimal first-stage decisions to solve the subproblems $P_i(x), \forall i \in \mathcal{T}$, obtaining $\bar{\mu}$ -central points

$$(y_i^*, \lambda_i^*, z_i^*) \in \mathcal{N}_2^{P_i}(\theta_0), \quad \forall i \in \mathcal{T}.$$

Then (x^R, s^R) and $(y_i^*, \lambda_i^*, z_i^*)_{i \in \mathcal{T}}$ are combined to obtain the *crash-start point* $\bar{w} = (\bar{x}, \bar{s}, \bar{y}, \bar{z}, \bar{\lambda})$. We aim to show that the warmstart from this point in problem $P(\mathcal{T})$ is successful, in that the modification step $\Delta w = (\Delta x, \Delta s, \Delta y, \Delta z, \Delta \lambda)$ from $\bar{w}$ satisfies

$$\bar{w} + \Delta w \in \mathcal{N}_2^{\mathcal{T}}(\theta).$$

In particular we will show that the warm-start is successful if the reduced tree $\mathcal{T}^R$ is close enough to $\mathcal{T}$ in an appropriate measure.

As the measure of closeness of the trees we use the *Wasserstein distance* [10, 21]. In an abstract probability space setting with event space Ω and corresponding Borel σ -field $\mathcal{B}$, the Wasserstein distance of two probability measures $\nu, \tilde{\nu}$ on $(\Omega, \mathcal{B})$ is given by

$$\hat{\mu}_1(\nu, \tilde{\nu}) = \inf \left\{ \int_{\Omega \times \Omega} \|\omega - \tilde{\omega}\| \eta(d(\omega, \tilde{\omega})) : \eta \in \mathcal{P}(\Omega \times \Omega), \eta(B \times \Omega) = \nu(B), \right. \\ \left. \eta(\Omega \times B) = \tilde{\nu}(B), \forall B \in \mathcal{B} \right\}.$$

In our case, where ν and $\tilde{\nu}$ are discrete measures implied by the full and reduced scenario trees as

$$\nu = \sum_{(\pi_i, \xi_i) \in \mathcal{T}} \pi_i \delta_{\xi_i}, \quad \tilde{\nu} = \sum_{(\tilde{\pi}_i, \tilde{\xi}_i) \in \mathcal{T}^R} \tilde{\pi}_i \delta_{\tilde{\xi}_i},$$

where δ_{ξ_i} is the Dirac-measure placed on scenario ξ_i , the corresponding formula for the Wasserstein distance reduces to the transportation distance

$$W_1(\mathcal{T}, \mathcal{T}^R) := \hat{\mu}_1(\nu, \tilde{\nu}) \\ = \min_{\eta \geq 0} \left\{ \sum_{(\pi_i, \xi_i) \in \mathcal{T}} \sum_{(\tilde{\pi}_j, \tilde{\xi}_j) \in \mathcal{T}^R} \|\xi_i - \tilde{\xi}_j\| \eta_{ij} : \sum_i \eta_{ij} = \tilde{\pi}_j, \sum_j \eta_{ij} = \pi_i \right\}.$$

Our analysis first deals with an idealised algorithm in which we choose a target μ value of $\bar{\mu}$ and then obtain the exact μ -center for reduced problem $P(\mathcal{T}^R)$ and scenario subproblems $P_i(x)$. Later we give results for the more realistic algorithm in which we are content with finding points in a neighbourhood of the central path. In what follows we assume that the random parameter ξ only affects the right-hand side h of the problem, and that this influence is bounded by the Lipschitz constant H :

$$\|h(\xi) - h(\tilde{\xi})\| \leq H \|\xi - \tilde{\xi}\|, \quad \forall \xi, \tilde{\xi}.$$

Moreover, we make use of the following relation from Dikin [8]

$$\chi(W) := \sup_{\Sigma \in D^+} \|\Sigma W^\top (W \Sigma W^\top)^{-1}\|_\infty < \infty, \quad (7)$$

where D^+ is the set of all diagonal matrices with strictly positive diagonal elements and W is an arbitrary matrix.

4.1 Results for exact subproblem solutions

We first investigate the dependence of the value functions of the family of barrier problems (for different values of ν)

$$\begin{aligned} \hat{Q}_\nu(x; \xi) = \min & c^\top y - \nu \sum_j \ln y_j \\ \text{s.t. } & Wy = h(\xi) - Tx \end{aligned} \quad (8)$$

on the random parameter ξ . We start with the following Lemma which establishes a Lipschitz result for the function $Q_\nu(x; \xi)$.

Lemma 1 *The functions $\hat{Q}_\nu(x; \xi)$ are Lipschitz in the second argument, that is*

$$|\hat{Q}_\nu(x; \xi) - \hat{Q}_\nu(x; \bar{\xi})| \leq L_Q \|\xi - \bar{\xi}\|_2$$

with Lipschitz constant $L_Q := 3\overline{C(d)}(\overline{C(d)}\|d\| + \nu n)\chi(W)H$.

Proof Optimality conditions for (8) are

$$z + W^\top \lambda = c, \quad Wy + Tx = h(\xi), \quad YZe = \nu e. \quad (9)$$

Under our assumptions, this system is non-singular, so the implicit function theorem assures the existence of y, λ and s as differentiable functions of ξ . Differentiating with respect to ξ gives

$$W \frac{dy}{d\xi} = \frac{dh}{d\xi} \quad (10a)$$

$$\frac{dz}{d\xi} + W^\top \frac{d\lambda}{d\xi} = 0 \quad (10b)$$

$$Y \frac{dz}{d\xi} + Z \frac{dy}{d\xi} = 0. \quad (10c)$$

After rearranging the final equation for $\frac{dz}{d\xi}$ and substituting into the second equation we obtain

$$-Y^{-1}Z \frac{dy}{d\xi} + W^\top \frac{d\lambda}{d\xi} = 0,$$

and hence

$$\frac{dy}{d\xi} = YZ^{-1}W^\top \frac{d\lambda}{d\xi}. \quad (11)$$

Multiplying from the left with W and substituting into (10a) yields

$$\frac{d\lambda}{d\xi} = (WY Z^{-1} W^\top)^{-1} \frac{dh}{d\xi},$$

which together with (11) gives

$$\frac{dy}{d\xi} = Y Z^{-1} W^\top (WY Z^{-1} W^\top)^{-1} \frac{dh}{d\xi}.$$

Now, recalling (7), we get the bound

$$\left\| \frac{dy}{d\xi} \right\|_\infty \leq \chi(W) \left\| \frac{dh}{d\xi} \right\|_\infty.$$

On the other hand, from the definition of Q_ν in (8) we have

$$\frac{dQ_\nu}{d\xi} = c^\top \frac{dy}{d\xi} - \nu Y^{-1} \frac{dy}{d\xi} = (c - z)^\top \frac{dy}{d\xi}.$$

From [20, Theorem 3.1] we have the bound

$$\|z\|_\infty \leq 2C(d)(C(d)\|d\| + \nu n),$$

which together with $\|c\|_\infty \leq \|d\|$ and $\left\| \frac{dh}{d\xi} \right\|_\infty \leq H$ yields

$$\left\| \frac{dQ_\nu}{d\xi} \right\|_\infty \leq [\|d\| + 2C(d)(C(d)\|d\| + \nu n)] \chi(W) H,$$

and the assertion of the Lemma follows since $C(d) \geq 1$.

In what follows we define

$$\eta_T(x) := c^\top x - \mu \sum_j \ln x_j + \rho_T(x) \quad (12)$$

where

$$\rho_T(x) = \sum_{(\pi_i, \xi_i) \in \mathcal{T}} \pi_i Q_{\frac{\mu}{\pi_i}}(x; \xi_i).$$

We are interested in how the optimal first-stage solution $x_\mu(\mathcal{T})$ to the problem

$$\min_x \eta_T(x) \quad (P_\mu(\mathcal{T}))$$

depends on the underlying scenario set $\mathcal{T}$. We need to assume that the set of possible optimal first-stage decisions x for different trees is bounded, that is

$$\|x_\mu(\mathcal{T})\|_\infty \leq \overline{B}, \quad \text{for all considered } \mathcal{T}.$$

Also we assume that the probabilities of all considered scenarios in the full and reduced tree are bounded below by $\bar{\pi}$. As long as we only consider trees that are derived from the original tree by deletion of scenarios and aggregation, both assumptions hold trivially, since there is only a finite number of possible trees to consider.

Lemma 2 *With the notation introduced above we have*

$$\|x_\mu(\mathcal{T}) - x_\mu(\mathcal{T}^R)\|_\infty \leq \frac{2}{\sqrt{\mu}} \bar{B} \sqrt{L_Q(\mu/\bar{\pi})} \sqrt{W_1(\mathcal{T}, \mathcal{T}^R)}$$

and

$$\|s_\mu(\mathcal{T}) - s_\mu(\mathcal{T}^R)\|_\infty \leq \frac{2}{\sqrt{\mu}^3} \bar{B}^3 \sqrt{L_Q(\mu/\bar{\pi})} \sqrt{W_1(\mathcal{T}, \mathcal{T}^R)}.$$

Proof We make use of Remark 2.2 to [21, Proposition 2.1] (see also [22, Theorem 2.10]), which states that given $|h(z) - h(\tilde{z})| \leq L_h \|z - \tilde{z}\|$, we have

$$\left| \int h(z) \nu(dz) - \int h(z) \tilde{\nu}(dz) \right| \leq L_h W_1(\nu, \tilde{\nu}).$$

We apply this result to the function $Q_\mu(x; \xi)$, with $\nu, \tilde{\nu}$, as before, being the probability measures implied by the full and reduced trees. In this case

$$\left| \sum_{i \in \mathcal{T}} \pi_i Q_{\frac{\mu}{\pi_i}}(x, \xi_i) - \sum_{i \in \mathcal{T}^R} \tilde{\pi}_i Q_{\frac{\mu}{\tilde{\pi}_i}}(x, \tilde{\xi}_i) \right| \leq L_Q(\mu/\bar{\pi}) W_1(\mathcal{T}, \mathcal{T}^R)$$

therefore

$$|\rho_{\mathcal{T}}(x) - \rho_{\mathcal{T}^R}(x)| \leq L_Q W_1(\mathcal{T}, \mathcal{T}^R),$$

and hence also

$$|\eta_{\mathcal{T}}(x) - \eta_{\mathcal{T}^R}(x)| \leq L_Q W_1(\mathcal{T}, \mathcal{T}^R). \quad (13)$$

By $x_\mu(\mathcal{T}^R)$ being a minimiser of $\eta_{\mathcal{T}^R}(x)$, it holds that

$$\eta_{\mathcal{T}^R}(x_\mu(\mathcal{T}^R)) \leq \eta_{\mathcal{T}^R}(x_\mu(\mathcal{T})). \quad (14)$$

Bound (13) implies

$$\eta_{\mathcal{T}^R}(x_\mu(\mathcal{T})) \leq \eta_{\mathcal{T}}(x_\mu(\mathcal{T})) + L_Q W_1(\mathcal{T}, \mathcal{T}^R),$$

and hence together with (14)

$$\eta_{\mathcal{T}^R}(x_\mu(\mathcal{T}^R)) \leq \eta_{\mathcal{T}}(x_\mu(\mathcal{T})) + L_Q W_1(\mathcal{T}, \mathcal{T}^R). \quad (15)$$

The functions $\rho_{\mathcal{T}}(x), \rho_{\mathcal{T}^R}(x)$ are convex and differentiable, and the term $-\mu \sum_j \ln x_j$ is convex and twice continuously differentiable with Hessian

$$\nabla_x^2 [-\mu \sum_j \ln x_j] = \mu X^{-2}$$

and its lowest eigenvalue satisfies

$$\sigma_1(\mu X^{-2}) \geq \mu / (\max_j x_j)^2 \geq \mu / \bar{B}^2.$$

Hence we get the bound

$$\begin{aligned} \eta_{\mathcal{T}}(x_\mu(\mathcal{T}^R)) &\geq \eta_{\mathcal{T}}(x_\mu(\mathcal{T})) + \nabla_x \eta_{\mathcal{T}}(x_\mu(\mathcal{T}))^\top (x_\mu(\mathcal{T}^R) - x_\mu(\mathcal{T})) \\ &\quad + \frac{\mu}{2\bar{B}^2} \|x_\mu(\mathcal{T}^R) - x_\mu(\mathcal{T})\|_\infty^2. \end{aligned}$$

Since $\eta_{\mathcal{T}}(x)$ is convex and differentiable with minimiser $x_{\mu}(\mathcal{T})$, then $\nabla_x \eta_{\mathcal{T}}(x_{\mu}(\mathcal{T})) = 0$, so that

$$\eta_{\mathcal{T}}(x_{\mu}(\mathcal{T}^R)) \geq \eta_{\mathcal{T}}(x_{\mu}(\mathcal{T})) + \frac{\mu}{2\bar{B}^2} \|x_{\mu}(\mathcal{T}^R) - x_{\mu}(\mathcal{T})\|_{\infty}^2 \quad (16)$$

Therefore, by combining (13) and (16) we get

$$\begin{aligned} \eta_{\mathcal{T}^R}(x_{\mu}(\mathcal{T}^R)) &\geq \eta_{\mathcal{T}}(x_{\mu}(\mathcal{T}^R)) - L_Q W_1(\mathcal{T}, \mathcal{T}^R) \\ &\geq \eta_{\mathcal{T}}(x_{\mu}(\mathcal{T})) + \frac{\mu}{2\bar{B}^2} \|x_{\mu}(\mathcal{T}^R) - x_{\mu}(\mathcal{T})\|_{\infty}^2 - L_Q W_1(\mathcal{T}, \mathcal{T}^R). \end{aligned}$$

On the other hand, the minimum value $\eta_{\mathcal{T}^R}(x_{\mu}(\mathcal{T}^R))$ of $P_{\mu}(\mathcal{T}^R)$ needs to satisfy (15) so that we have

$$\begin{aligned} \eta_{\mathcal{T}}(x_{\mu}(\mathcal{T})) + L_Q W_1(\mathcal{T}, \mathcal{T}^R) \\ \geq \eta_{\mathcal{T}}(x_{\mu}(\mathcal{T})) + \frac{\mu}{2\bar{B}^2} \|x_{\mu}(\mathcal{T}^R) - x_{\mu}(\mathcal{T})\|_{\infty}^2 - L_Q W_1(\mathcal{T}, \mathcal{T}^R). \end{aligned}$$

After rearrangement we are left with

$$\frac{\mu}{2\bar{B}^2} \|x_{\mu}(\mathcal{T}^R) - x_{\mu}(\mathcal{T})\|_{\infty}^2 \leq 2L_Q W_1(\mathcal{T}, \mathcal{T}^R),$$

or equivalently

$$\|x_{\mu}(\mathcal{T}) - x_{\mu}(\mathcal{T}^R)\|_{\infty} \leq \frac{2}{\sqrt{\mu}} \bar{B} \sqrt{L_Q} \sqrt{W_1(\mathcal{T}, \mathcal{T}^R)}.$$

For the bound on $\|s_{\mu}(\mathcal{T}) - s_{\mu}(\mathcal{T}^R)\|_{\infty}$ we note that

$$X_{\mu}(\mathcal{T}) S_{\mu}(\mathcal{T}) = \mu I, \quad X_{\mu}(\mathcal{T}^R) S_{\mu}(\mathcal{T}^R) = \mu I,$$

so that

$$s_{\mu}(\mathcal{T}) = \frac{1}{\mu} S_{\mu}(\mathcal{T}) S_{\mu}(\mathcal{T}^R) x_{\mu}(\mathcal{T}^R), \quad s_{\mu}(\mathcal{T}^R) = \frac{1}{\mu} S_{\mu}(\mathcal{T}) S_{\mu}(\mathcal{T}^R) x_{\mu}(\mathcal{T}),$$

and hence

$$\begin{aligned} \|s_{\mu}(\mathcal{T}) - s_{\mu}(\mathcal{T}^R)\|_{\infty} &= \left\| \frac{1}{\mu} S_{\mu}(\mathcal{T}) S_{\mu}(\mathcal{T}^R) (x_{\mu}(\mathcal{T}^R) - x_{\mu}(\mathcal{T})) \right\|_{\infty} \\ &\leq \frac{\bar{B}^2}{\mu} \|x_{\mu}(\mathcal{T}) - x_{\mu}(\mathcal{T}^R)\|_{\infty}, \end{aligned}$$

yielding the required bound.

Lemma 3 *Let estimates of the first-stage decisions $\bar{x} = x^R, \bar{s} = s^R$ be given and let (y_i, λ_i, z_i) be the (exact) μ -center for the i -th scenario subproblem $P_i(\bar{x})$. Further, let $(x_{\mu}, y_{\mu}, \lambda_{\mu}, s_{\mu}, z_{\mu})$ be the exact μ -center of the full problem $P(\mathcal{T})$. Then*

$$\|\lambda_i - \lambda_{\mu, i}\|_{\infty} \leq C_{\lambda} \|\bar{x} - x_{\mu}\|_{\infty}$$

where

$$C_{\lambda} = C_{\lambda}(\mu) = 4\chi(W) \|T\|_{\infty} \overline{C(d)}^2 [\overline{C(d)} \|d\| + \mu n]^2 / \mu.$$

Proof The optimality conditions for problem $P_i(\bar{x})$ are given by (9). Differentiating with respect to x gives:

$$\frac{dz_i}{dx} + W^\top \frac{d\lambda_i}{dx} = 0, \quad W \frac{dy_i}{dx} = -T, \quad Y_i \frac{dz_i}{dx} + Z_i \frac{dy_i}{dx} = 0.$$

As in the proof to Lemma 1, these equations can be solved for $\frac{d\lambda_i}{dx}$ to obtain

$$\begin{aligned} \frac{d\lambda_i}{dx} &= (WY_iZ_i^{-1}W^\top)^{-1}T \\ &= (WW^\top)^{-1}WY_i^{-1}Z_iY_iZ_i^{-1}W^\top(WY_iZ_i^{-1}W^\top)^{-1}T. \end{aligned}$$

As before we have

$$\|Z_i^{-1}W^\top(WY_iZ_i^{-1}W^\top)^{-1}\|_\infty \leq \chi(W), \quad \|W^\top(WW^\top)^{-1}\|_\infty \leq \chi(W),$$

and, due to (y_i, λ_i, z_i) being the exact μ -center,

$$\|Z_iY_i^{-1}\|_\infty = \|Z_i^2/\mu\|_\infty = \|Z_i\|_\infty^2/\mu.$$

We use again the bound from [20, Theorem 3.1]

$$\|z_i\|_\infty \leq 2C(d_i)[C(d_i)\|d_i\| + \mu n]$$

giving

$$\begin{aligned} \left\| \frac{d\lambda_i}{dx} \right\|_\infty &\leq 4\chi(W)^2\|T\|_\infty C(d_i)^2[C(d_i)\|d_i\| + \mu n]^2/\mu \\ &\leq 4\chi(W)^2\|T\|_\infty \overline{C(d)}^2[\overline{C(d)}\|\overline{d}\| + \mu n]^2/\mu \end{aligned}$$

yielding the assertion of the Lemma.

Theorem 1 *Let $\bar{w} = (\bar{x}, \bar{y}, \bar{\lambda}, \bar{s}, \bar{z})$ be the warm-start point for problem $P(T)$ obtained by following the above algorithm starting from the reduced tree T^R , using a target μ -value of $\bar{\mu}$. If*

$$W_1(T, T^R) \leq \frac{\theta^2}{\overline{C(d)}^2 C_2(\mu)^2} \min \left\{ \frac{\|\overline{d}\|^2}{4(2n+1)^2}, \frac{\mu^2}{64\overline{C(d)}^2} \right\}$$

where $C_2(\mu) = 2\overline{B}\sqrt{\overline{L_Q}}(\|T\|_\infty|T|C_\lambda(\mu) + \overline{B}^2/\mu)/\sqrt{\bar{\mu}}$, then the warmstart is successful, that is the Yildirim - Wright [26] Weighted Least Squares step $(\Delta x, \Delta y, \Delta \lambda, \Delta s, \Delta z)$ from $(\bar{x}, \bar{y}, \bar{\lambda}, \bar{s}, \bar{z})$ is feasible and leads to

$$(\bar{x} + \Delta x, \bar{y} + \Delta y, \bar{\lambda} + \Delta \lambda, \bar{s} + \Delta s, \bar{z} + \Delta z) \in \mathcal{N}_2^T(\theta).$$

Proof Due to the construction of the crash-start point $\bar{w}$ we have

$$T\bar{x} + W\bar{y}_i = h_i, \quad \bar{X}\bar{S}e = \mu e, \quad \bar{Y}_i\bar{Z}_i e = \mu e,$$

that is, the point is primal feasible and central in the complete problem $P(\mathcal{T})$. Let

$$\bar{c} = \sum_i T^\top \bar{\lambda}_i + \bar{s}$$

and consider the problem instance $P(\bar{d})$ obtained from the full problem $P(\mathcal{T})$ by replacing the first-stage cost c with $\bar{c}$. By construction, $\bar{w}$ is also dual feasible for $P(\bar{d})$, hence it is the exact μ -center, or $\bar{w} \in \bar{\mathcal{N}}_2(0)$, where $\bar{\mathcal{N}}_2$ is the $\mathcal{N}_2$ -neighbourhood for problem $P(\bar{d})$. We will treat this point as a warm-start attempt for the (perturbed) problem $P(\mathcal{T})$ starting from a central point for problem $P(\bar{d})$. The change in problem data is

$$\Delta d = (\Delta A, \Delta b, \Delta c) = (0, 0, \Delta c),$$

with

$$\begin{aligned} \Delta c &= c - \bar{c} = \sum_i T^\top \lambda_{\mu,i} + s_\mu - \sum_i T^\top \bar{\lambda}_i - \bar{s} \\ &= \sum_i T^\top (\lambda_{\mu,i} - \bar{\lambda}_i) + (s_\mu - \bar{s}). \end{aligned}$$

Using the bounds from Lemma 2 and Lemma 3 we have

$$\begin{aligned} \|\Delta c\|_\infty &\leq \|T\|_\infty |\mathcal{T}| \max_i \|\lambda_{\mu,i} - \bar{\lambda}_i\|_\infty + \|s_\mu - \bar{s}\|_\infty \\ &\leq \|T\|_\infty |\mathcal{T}| C_\lambda(\mu) \|x_\mu - \bar{x}\|_\infty + \|s_\mu - \bar{s}\|_\infty \\ &\leq \left(\|T\|_\infty |\mathcal{T}| C_\lambda(\mu) + \frac{\bar{B}^2}{\mu} \right) \frac{2}{\sqrt{\mu}} \bar{B} \sqrt{L_Q} \sqrt{W_1(\mathcal{T}, \mathcal{T}^R)} \\ &= C_2(\mu) \sqrt{W_1(\mathcal{T}, \mathcal{T}^R)} \end{aligned}$$

with

$$C_2(\mu) = \left(\|T\|_\infty |\mathcal{T}| C_\lambda(\mu) + \frac{\bar{B}^2}{\mu} \right) \frac{2}{\sqrt{\mu}} \bar{B} \sqrt{L_Q}.$$

According to [26, Proposition 4.2] sufficient conditions for a successful warm-start are (using $\theta_0 = 0$, $\xi = \theta/2$):

$$\|\Delta c\|_\infty \leq \frac{\theta}{2(2n+1)C(d)} \|d\|, \quad \text{and} \quad \mu \geq 8 \frac{C(d)^2}{\theta} \|\Delta c\|_\infty,$$

which can be combined to produce

$$\|\Delta c\|_\infty \leq \frac{\theta}{C(d)} \min \left\{ \frac{\|d\|}{2(2n+1)}, \frac{\mu}{8C(d)} \right\}.$$

Together with the above bound on $\|\Delta c\|_\infty$, we obtain the condition

$$C_2(\mu)\sqrt{W_1(\mathcal{T}, \mathcal{T}^R)} \leq \frac{\theta}{C(d)} \min \left\{ \frac{\|d\|}{2(2n+1)}, \frac{\mu}{8C(d)} \right\},$$

and, after rearranging, the condition given in the Theorem.

The question remains how the condition of the Theorem can be satisfied in practice. Clearly the bound on $W_1(\mathcal{T}, \mathcal{T}^R)$ is largest when the two expressions in the min are equal, that is if

$$\frac{\|d\|}{2(2n+1)} = \frac{\mu}{8C(d)} \quad \text{or} \quad \mu = \frac{4\|d\|C(d)}{2n+1}.$$

For an (optimally) chosen μ -value, Theorem 1 gives conditions on the closeness of the trees $\mathcal{T}$ and $\mathcal{T}^R$. It is not practically possible to determine this value exactly, so the μ value would in practice be chosen by an appropriate heuristic.

Closer to the practical application is probably to choose the two trees first and then the corresponding μ value to be used in the solution of the reduced problem $P(\mathcal{T}^R)$ and the subproblems $P_i(\bar{x})$. In this view of things, Theorem 1 gives a minimum closeness of the approximating tree $\mathcal{T}^R$ that needs to be achieved and also gives conditions on the selection of the corresponding μ value. In practice we expect the conditions to be much too tight and will proceed with the warm-start attempt regardless.

4.2 Results for approximate subproblem solutions

The requirement to obtain an exact μ -center for the reduced problem $P(\mathcal{T}^R)$ or the decomposed subproblems $P_i(\bar{x})$ is too demanding. Rather we will be satisfied with obtaining primal-dual feasible and reasonably central points, i.e.

$$(\tilde{x}_\mu^R, \tilde{y}_\mu^R, \tilde{\lambda}_\mu^R, \tilde{s}_\mu^R, \tilde{z}_\mu^R) \in \mathcal{N}_2^R(\tilde{\theta}) \quad (17)$$

$$(\tilde{y}_i, \tilde{\lambda}_i, \tilde{z}_i) \in \mathcal{N}_2^{(i)}(\tilde{\theta}) \quad (18)$$

for some $\tilde{\theta} \in (0, 1)$. We can then establish bounds on the resulting error in the relevant components of these points compared to the exact μ -centers (Lemma 5, Appendix). Following this strand it is possible to derive a counterpart of Theorem 1 in

Theorem 2 *Let $\tilde{w}$ be the crash-start point for problem $P(\mathcal{T})$ obtained by following the above algorithm starting from the reduced tree $\mathcal{T}^R$, using a target μ -value of $\bar{\mu}$ and approximate solutions satisfying (17) and (18). If we choose $W_1(\mathcal{T}, \mathcal{T}^R)$ and $\tilde{\theta}$ small enough – specifically so that they satisfy condition (26) – then the warm-start is successful, that is the Yildırım - Wright [26] Weighted Least Squares step Δw from $\tilde{w}$ is feasible and leads to $\tilde{w} + \Delta w \in \mathcal{N}_2^T(\theta)$.*

The detailed analysis of this situation together with the proof of Theorem 2 can be found in the appendix.

Problem	scenarios	rows	columns	nonzeros
AIRL1	25	152	306	706
AIRL2	25	152	306	706
AIRL3	676	4,058	8,118	18,934
cargo-4node32	32	2,382	6,396	15,656
cargo-4node64	64	4,750	12,732	31,016
cargo-4node128	128	9,486	25,404	61,736
cargo-4node256	256	18,958	50,748	123,176
cargo-4node512	512	37,902	101,436	246,056
cargo-4node1024	1024	75,790	202,812	491,816
cargo-4node2048	2048	151,566	405,564	983,336
cargo-4node4096	4096	303,118	811,068	1,966,376
cargo-4node8192	8196	606,222	1,622,076	3,932,456
cargo-4node16384	16384	1,212,430	3,244,092	7,864,616
asset1	100	505	1,313	2,621
asset2	37500	187,505	487,513	975,021
env.1200	1200	57,648	102,085	220,972
env.1875	1875	90,048	159,460	345,172
env.3780	3780	181,488	321,385	695,692
env.5292	5292	254,064	449,905	973,900
env.lрге	8232	395,184	699,805	1,514,860
stocfor2	64	6,543	9,237	29,985
sslp_10_50_100	100	6,001	52,011	101,911
sslp_10_50_500	500	30,001	260,011	509,511
sslp_10_50_1000	1000	60,001	520,011	1,019,011
dcap233_500	500	7,506	16,518	31,518
dcap243_500	500	9,006	21,018	39,018
storm27	27	14,441	37,485	94,274
storm125	125	66,185	172,431	433,256
storm1000	1000	528,185	1,377,306	3,459,881

Table 1: Characteristics of 2-stage problems

5 Implementation and numerical results

In this section we present the numerical results of the implementation of the proposed decomposition crash-start scheme within the interior point code OOPS.

The setup is the following. Given a 2-stage stochastic programming problem with a given scenario set $\mathcal{T}$, we first generate a reduced tree $\mathcal{T}^R$ with a specified number of scenarios by eliminating scenarios using the scenario reduction technique of [15]. The reduced problem is solved until the first iterate for which $\mu^{(k)} < \bar{\mu}$. If this point is not primal-dual feasible, additional pure centering iterations with a target of $\mu = \bar{\mu}$ are performed until we obtain a primal-dual feasible point. The x -components of this point are used to set-up scenario subproblems $P_{i,\mu}(x)$ which again are solved to obtain a primal-dual feasible point corresponding to $\mu \approx \bar{\mu}$ following the procedure above.

We concentrate our tests on the two-stage problems summarised in Table 1. The first block of problems are from the collection of Ariyawansa and Felt [2], 'sslp' and 'dcap' are from the SIPLIB collection of Ahmed [1] and 'storm' is from the POSTS collection [4]. We compare the number of itera-

tions until converged from the crash-start point and total running time (for both crash-start generation and solution phase) for the suggested algorithm against the reduced tree crash-start of [7] and a cold-start using Mehrotra's starting point[18]. Results are summarised in Table 2. All computations were performed on a Linux PC with a 3.0GHz Intel Core 2 processor and 3GB of RAM. Final convergence tolerance was to reduce the relative primal-dual gap below 10^{-7} .

Problem	$ T $	cold		dec				red			
		iter	time	iter	time	$ T^R $	$\bar{\mu}$	iter	time	$ T^R $	$\bar{\mu}$
AIRL1	25	12	0.05	3	0.05	12	1.0	14	0.07	5	10
AIRL2	25	10	0.07	4	0.05	12	1.0	8	0.04	2	10
AIRL3	676	11	1.45	7	1.66	60	1.0	31	4.23	100	10
4node32	32	21	1.6	7	0.9	6	0.1	7	0.6	2	0.01
4node64	64	25	4.4	9	2.1	10	0.1	7	1.14	2	0.1
4node128	128	24	8.0	7	3.1	10	0.01	7	2.17	2	0.1
4node256	256	22	14.8	7	6.2	10	0.01	7	4.3	2	0.1
4node512	512	30	41.5	9	9.7	64	0.01	7	4.8	2	0.01
4node1024	1024	30	80.7	9	17.9	64	0.01	8	10.2	2	0.01
4node2048	2048	33	189	10	41.3	128	0.01	8	22.6	2	0.01
4node4096	4096	43	491	12	90.1	128	0.01	10	54.9	2	0.1
4node8192	8192	37	856	12	202	512	0.01	10	106	2	0.01
4node16384	16384	42	2083	12	402	256	0.01	10	222	2	0.01
asset1	100	11	0.18	5	0.23	10	0.01	4	0.10	20	0.01
asset2	37500	15	174	7	717	50	0.001	6	85.6	2	0.01
env.1200	1200	31	37.8	6	14.7	10	0.001	6	8.1	10	0.1
env.1875	1875	31	61.4	7	24.2	10	0.001	22	45.1	100	0.001
env.3780	3780	34	133	9	59.3	50	0.001	100	-	-	-
env.5292	5292	33	191	7	78.2	100	0.001	25	148	50	0.1
env.lрге	8232	38	368	8	141	100	0.001	43	404	100	0.1
stocfor2	64	21	3.9	9	1.6	2	10	8	0.9	1	10
sslp10_50_100	100	26	8.0	15	6.8	10	0.001	27	9.2	2	0.1
sslp10_50_500	500	49	80.0	21	44.1	50	0.001	49	84.7	20	0.1
sslp10_50_1000	1000	55	183	21	87.4	50	0.001	59	204	50	10
dcap233_500	500	22	3.9	7	1.8	50	0.01	13	2.3	10	10
dcap243_500	500	23	4.7	7	1.4	50	0.01	13	2.7	10	10
storm27	27	81	37.3	6	9.3	5	10	29	14.1	1	10
storm125	125	98	310	8	39.0	10	1.0	35	65.8	1	10
storm1000	1000	108	1868	21	420	20	10	22	378	50	10

Table 2: Results for 2-stage problems.

As can be seen the decomposition-based crash start offers significant savings both in terms of iterations and time when compared to the coldstart. Savings average about 60% in terms of iterations and 30% in terms of time and are more pronounced for larger problems, in many cases achieving a halving of total solution time. It also offers savings compared to the reduced tree crash-start. This advantage is especially pronounced for problems where the latter algorithm struggles.

The results in Table 2 correspond to the best $\bar{\mu}$ value found for either approach (using trial values of $\bar{\mu} = 10^k, k = -4, \dots, 1$). Indications are that the best $\bar{\mu}$ value is fairly constant within a problem class. However, this is slightly unsatisfactory.

To evaluate the sensitivity of the proposed crash-start algorithm with respect to the target barrier parameter $\bar{\mu}$ and the reduced tree size $|\mathcal{T}^R|$ and to give a comparison with the reduced tree crash-start of [7] we have set up the following experiment: for 24 of the 29 problems reported in Tables 1 and 2 (choosing 5 problems of the **cargo** series to limit the influence of one particular problem class) we have chosen 5 values for $\bar{\mu}$, namely $\bar{\mu} = 10^k$ for $k = -3, \dots, 1$ and 4 different reduced tree sizes (between 2 and $|\mathcal{T}|/4$). For these $24 \times 20 = 480$ variations we have again compared the performance of the decomposition based crash-start to that of the reduced tree crash-start and a cold start. We compare both IPM iterations in the full problem and total solution time. Instances that took more than 100 iterations for the warmstarted problem are deemed unsuccessful and are regarded as infinite iterations and solution time. We wish to note that we generally expect the decomposition based crash-start to be successful as long as the conditions indicated in Theorem 1 are met at least approximately: that is a small target $\bar{\mu}$ should not be combined with a too roughly approximating $\mathcal{T}^R$. As long as the warmstart is successful however, we would generally expect that a small target $\bar{\mu}$ leads to fewer full problem iterations.

Complete results of these trials are available from [14]. Here we give a summary. Figure 2 presents results of the trials in form of performance profiles[9]. Curves labelled 'dec', 'red' and 'cold' correspond to the decomposition based crash-start, the reduced tree crash-start and the cold start respectively.

Fig. 2: Performance profile: number of IPM iterations (left) and total solution time (right). The performance profile plots the percentage of all instances for which the performance of a given method is within a factor (corresponding to the reading of the horizontal axis) of the best method.

We can see that the decomposition based crash-start requires the (joint) least number of iterations in just under 60% of cases, with the reduced tree crash-start being optimal in just over 30% of cases. As expected these savings do not quite manifest themselves with respect to total solution time, but again the decomposition based crash-start is the fastest method in just under 40% of instances, outperforming both the cold start and the reduced tree crash-start. Further more than 80% of instances can be solved within twice the time needed by the best approach (again outperforming the other two alternatives in this metric). This shows that the decomposition based crash-start is not only preferable to both other methods in most situations but is also significantly less sensitive to the correct choice of $\bar{\mu}$ and $\mathcal{T}^R$ than the reduced tree crash-start.

A more detailed analysis of the results reveals the following: Of the 480 instances, in 155 (that is just under 1/3) the problem cannot be solved with the reduced tree crash-start within the allowable number of iterations. Of those 325 cases where it can be solved, in just over half of the cases (187, 58%) the reduced tree crash-start needs fewer iterations, and in 180 cases (55%) also less time than the cold-start. On the other hand if we take the best combination of tree size and target- $\bar{\mu}$, then in 75% of problems the reduced tree crash-start needs less iterations (and in 71% also less time) than the cold start. In other words, while the reduced tree crash-start is able to produce significant time-savings it does suffer from a certain sensitivity with regards to the optimal choice of tree size and target- $\bar{\mu}$. It should be noted however that good choices are relatively stable for each problem class, so with experience good guesses can be made.

For the decomposition based crash-start on the other hand, 455 of the 480 trials (95%) are successful, of these 373 (82%) yield an improvement in terms of number of iterations and still 285 (63%) an improvement in terms of time when compared to the cold start. The best combination of $\bar{\mu}$ and $\mathcal{T}^R$ for each problem is in all cases better than the cold-start in terms of iterations (although worse in terms of solution time for AIRL3, asset1, asset2) and leads to an average improvement of 70.1% of iterations and a saving of 33.2% in terms of total solution time. The unsuccessful trials occur for the smallest reduced tree size ($|\mathcal{T}^R| = 2$) in the **env.5292**, **env.lrgc** and **sslp_10_50_100** problems and for half of the cases (corresponding to combinations of small $\bar{\mu}$ and small $|\mathcal{T}^R|$) in the **storm125** and **storm1000** problems.

Figure 3 shows the relative performance of both the decomposition based crash-start (blue solid line) and the reduced tree crash-start (green dashed line) with respect to the cold start for all 480 trials. Trials in each case are sorted in decreasing order of crash-start effectiveness, with the instances of failure removed, and plotted on a logarithmic scale on the y -axis. As can be seen, in most cases the decomposition based crash-start results in better performance compared to cold start (corresponding to the points below 1), although in some cases the performance is drastically worse. Most of the bad cases (where the decomposition based crash-start is successful but takes more than $2\times$ the number of iterations than the cold start) are from a single problem (AIRL3) where a large tree size combined with small μ results in uniformly bad

performance (although a small tree size with large $\bar{\mu}$ results in a successful warmstart).

Figure 4, which is typical of how the crash-start effectiveness depends on $\bar{\mu}$ and $\mathcal{T}^R$ gives some further insight: In terms of IPM iterations (left plot) there is a pronounced “sweet valley” for a combinations of small $\bar{\mu}$ and large $\mathcal{T}^R$ through to large $\bar{\mu}$ and small $\mathcal{T}^R$ for which the crash-start works very well. Outside this valley small reduced tree sizes lead to more iterations in the full problem, even if the crash-start itself is successful. On the other side a too small $\bar{\mu}$ combined with a small reduced tree (i.e. a far converged but not accurate approximation) results in a failed warmstart and hence a drastically increased iteration number. For the total solution time the situation is much the same except for a slightly larger relative penalty for large reduced tree sizes. The observed problems of the method with a combination of small $\bar{\mu}$ and small $|\mathcal{T}^R|$ is to be expected from the theoretical analysis of Section 4. It is encouraging however that outside this identified troublesome region the method seems robust with regard to the parameter choices. Note that a similar “sweet valley” can be observed for the reduced tree crash-start, at least for the problems on which this approach works well. Generally, however, the performance of the reduced tree crash-start is much more erratic (see Figure 3 and the full results [14]).

Fig. 3: Performance ratios (left: iterations, right: time) of reduced tree and decomposition tree crash-starts relative to cold-started IPM across all 480 problem instances. Warm-starts provide an improvement over cold start for ratios below 1, and a deterioration for ratios above 1 (note the logarithmic scale on the y -axis). The reduced tree curve is much shorter because of the higher number of failures.

For the **cargo**-series of problems we have also compared how the effectiveness of the crash-start for various $\bar{\mu}/\mathcal{T}^R$ combinations changes as the problem size increases. Results are presented in Figure 5. As can be seen both in terms of numbers of iterations (left hand plot) and total solution time (right hand plot) there is almost no discernible difference, apart from the fact that sensitivity to values in $\bar{\mu}$ and $\mathcal{T}^R$ become less pronounced for larger problems.

Fig. 4: Difference in iterations (left) and time (right) caused by decomposition based crash-start for problem **cargo-4node512** against different $\bar{\mu}$ and T^R -values.

Fig. 5: Percentage of saved iterations (left) and time (right) for all combinations of $\bar{\mu}$ and T^R against problem number for the **cargo** problems. Each line corresponds to one $\bar{\mu}/T^R$ combination, while the logarithmic horizontal axis corresponds to the size (in number of scenarios) of the problem in the series.

Finally we have evaluated the algorithm on some multistage problems from the literature (Table 3). For multistage problems the proposed decomposition based crash-start will decompose the problem at the second stage and then solve each sub-tree originating from a second stage node separately. This approach is in effect treating a multistage problem as a two stage problem with a large core matrix. The reduced scenario tree $\mathcal{T}^R$ is obtained by choosing a single scenario originating from every second stage node. We choose the one resulting in the minimal Wasserstein distance from the remaining nodes in this sub-tree. We compare the number of iterations in the full problem and the total solution time for the same crash start approaches as before namely the decomposition based crash-start presented in this paper, the reduced tree

Problem	stages	scenarios	rows	columns	nonzeros
fxm3-6	3	36	6,200	12,628	57,722
fxm3-16	3	256	41,340	85,575	392,252
fxm4-6	4	216	22,400	47,185	265,442
stocfor3	7	512	16,675	23,541	76,473
pltxpA3-16	3	256	28,350	74,172	150,801
pltxpA4-6	4	216	26,894	70,364	143,059
pltxpA5-6	5	1,296	161,678	422,876	859,747
swing8-4	8	65,536	262,142	349,522	786,422
watson10-128	8	128	26,237	49,664	129,159
watson10-512	10	512	67,069	129,024	351,751
watson10-1024	10	1024	134,127	258,032	703,473

Table 3: Characteristics of multistage test problems.

crash-start and a cold start. As before we have run tests with different selection of values for $\bar{\mu}$ and $|\mathcal{T}^R|$ where we have performed at least 20 tests for each problem. Full results can be obtained from [14], a summary is given in Table 4. For each problem we report the performance for the best choice of parameters and the median performance. We also give the percentage of successful warmstart attempts.

The decomposition based crash starts leads to a significant reduction in the number of iterations compared to the cold start, both when looking at the best and the median performance. While not quite as apparent in terms of total solution time, we are still able to beat the cold start on almost all problems (`stocfor3` being the only exception). When compared with the reduced tree crash-start the advantages are not quite as clear as for the two stage problems. While the median number of iterations still tends to be less, the scenario subproblems (which are here partial trees rooted at a first stage node) do now contribute to the overall solution time, leading generally to a worse performance than for the reduced tree crash-start. Nevertheless it can be seen that the decomposition based crash-start is more robust, with significantly fewer failures.

6 Conclusions and future research

In this paper we proposed a technique to generate a crash-start point for interior point methods applied to two-stage and multistage stochastic programming problems. The method solves a problem defined on a reduced tree in order to get the first-stage decisions, then a sequence of completely independent scenario subproblems.

We presented a thorough theoretical analysis of our procedure and derived sufficient conditions for a successful crash-start. The results show how the maximum allowable distance between full and reduced trees depends on the Renegar condition number of the problem instance and the target value $\bar{\mu}$ adopted in the solution of the reduced problem.

Problem	cold		red						dec					
			succ	best		median		succ	best		median			
	it	tm		it	tm	it	tm		it	tm	it	tm		
fxm3-6	26	2.5	100	4	0.7	8	1.1	100	4	1.8	5	2.4		
fxm3-16	30	18.9	100	6	4.4	8	5.4	100	5	13.4	9	15.2		
fxm4-6	25	9.7	100	6	2.8	10	4.3	100	5	8.3	7	9.5		
stocfor3	33	2.5	92	4	0.6	11	1.1	100	14	2.5	19	3.2		
pltxpA3-16	35	5.7	58	17	2.9	25	4.2	63	11	4.6	13	4.8		
pltxpA4-6	42	15.2	50	13	5.4	26	9.4	58	8	7.6	10	8.9		
pltxpA5-6	77	234	38	37	117	49	159	44	39	189	56	209		
swing8-4	27	44.5	100	7	16.7	41	98.0	100	5	34.2	8	39.5		
watson10-128	54	15.3	57	19	9.1	22	10.5	100	5	12.1	26	14.1		
watson10-512	79	61.0	54	23	27.7	25	38.7	89	4	47.3	38	55.9		
watson10-1024	91	143	45	22	65.0	26	82.0	96	6	102	44	123		

Table 4: Number of iterations and solution time in seconds for best and median performance of different crash start algorithms. Column **succ** gives the percentage of successful warmstarts.

Our numerical results show that, provided the reduced tree is a good approximation to the full tree, this approach leads to an advanced iterate from which an interior point method can be warm-started successfully, which often translates in significant time savings.

For problems with a large number of second stage nodes and for multi-stage problems, solving all scenario subproblems can be a significant cost. In our implementation the scenario subproblems are solved independently, using Mehrotra’s starting point. Clearly since the scenario subproblems are usually closely related, they could be warmstarted as well, leading to further efficiency gains. For multistage problems, the presented approach could also be applied in a nested manner whereby the sub-trees rooted at the second stage nodes are not solved exactly but rather the presented method is applied to these subproblems recursively to obtain a crash-start point. We leave this for future research.

Acknowledgements We would like to thank the two anonymous referees for their helpful comments which have improved the quality of the manuscript.

References

1. Ahmed, S.: SIPLIB: A stochastic integer programming test problem library. <http://www2.isye.gatech.edu/~sahmed/siplib/> (2004)
2. Ariyawansa, K.A., Felt, A.J.: On a new collection of stochastic linear programming test problems. *INFORMS J Comput* **16**(3), 291–299 (2004)
3. Birge, J.R.: Decomposition and partitioning methods for multistage stochastic linear programs. *Oper Res* **33**, 989–1007 (1985)
4. Birge, J.R., Holmes, D.: A portable stochastic programming test set (POSTS). <http://users.iems.northwestern.edu/~jrbirge/html/dholmes/post.html>
5. Birge, J.R., Louveaux, F.: *Introduction to Stochastic Programming*. Springer-Verlag, New York (1997)
6. Blomvall, J., Lindberg, P.O.: A Riccati-based primal interior point solver for multistage stochastic programming. *Eur J Oper Res* **143**, 452–461 (2002)

7. Colombo, M., Gondzio, J., Grothey, A.: A warm-start approach for large-scale stochastic linear programs. *Math Program* **127**(2), 371–397 (2011)
8. Dikin, I.I.: On the speed of an iterative process. *Upravlaemye Sistemy* **12**, 54–60 (1974)
9. Dolan, E., Moré, J.: Benchmarking optimization software with performance profiles. *Math Program* **91**, 201–213 (2002)
10. Dupačová, J., Gröwe-Kuska, N., Römisch, W.: Scenario reduction in stochastic programming. *Math Program* **95**(3), 493–511 (2003)
11. Gondzio, J., Grothey, A.: Reoptimization with the primal-dual interior point method. *SIAM J Optimiz* **13**(3), 842–864 (2003)
12. Gondzio, J., Grothey, A.: A new unblocking technique to warmstart interior point methods based on sensitivity analysis. *SIAM J Optimiz* **19**(3), 1184–1210 (2008)
13. Gondzio, J., Sarkissian, R.: Parallel interior point solver for structured linear programs. *Math Program* **96**(3), 561–584 (2003)
14. Grothey, A.: Additional results for "a decomposition-based crash-start for stochastic programming" (2012). URL <http://www.maths.ed.ac.uk/~agr/data/SPDecWSResults.pdf>
15. H. Heitsch, W.R.: Scenario reduction algorithms in stochastic programming. *Comput Optim Appl* **24**, 187–206 (2003)
16. Hipolito, A.L.: A weighted least squares study of robustness in interior point linear programming. *Comput Optim Appl* **2**, 29–46 (1993)
17. Kall, P., Wallace, S.W.: *Stochastic Programming*. John Wiley & Sons, Chichester (1994)
18. Mehrotra, S.: On the implementation of a primal-dual interior point method. *SIAM J Optimiz* **2**, 575–601 (1992)
19. Mulvey, J., Ruszczyński, A.: A new scenario decomposition method for large scale stochastic optimization. *Oper Res* **43**, 477–490 (1995)
20. Nunez, M.A., Freund, R.M.: Condition measures and properties of the central trajectory. *Math Program* **83**, 1–28 (1998)
21. Römisch, W., Schultz, R.: Stability analysis for stochastic programming. *Ann Oper Res* **30**, 241–266 (1991)
22. Schultz, R.: Some aspects of stability in stochastic programming. *Ann Oper Res* **100**, 55–84 (2000)
23. Steinbach, M.: Hierarchical sparsity in multistage convex stochastic programs. In: S. Uryasev, P.M. Pardalos (eds.) *Stochastic Optimization: Algorithms and Applications*, pp. 363–388. Kluwer Academic Publishers (2000)
24. Sun, J., Liu, X.: Scenario formulation of stochastic linear programs and the homogeneous self-dual interior point method. *INFORMS J Comput* **18**(4), 444–454 (2006)
25. Wright, S.J.: *Primal-Dual Interior-Point Methods*. SIAM, Philadelphia (1997)
26. Yildirim, E.A., Wright, S.J.: Warm-start strategies in interior-point methods for linear programming. *SIAM J Optimiz* **12**(3), 782–810 (2002)

Appendix

We now give the postponed analysis of the situation of Section 4.2 where the crash start point is not constructed from exact μ -centers, but approximate subproblem solutions satisfying conditions (17) and (18) are used.

The following Lemma 5 gives bounds on the resulting error in the relevant components of these points compared to the exact μ -centers. Before we proceed, we need a general result stating how far the components of a point in the $\mathcal{N}_2$ -neighbourhood can deviate from the exact μ -center.

Lemma 4 *Let (x_μ, y_μ, s_μ) be the exact μ -center for the linear programming problem*

$$\min_x c^\top x, \text{ s.t. } Ax = b, x \geq 0,$$

and let $(\tilde{x}_\mu, \tilde{y}_\mu, \tilde{s}_\mu) \in \mathcal{N}_2(\theta)$ with average complementarity product $\tilde{x}_\mu^\top \tilde{s}_\mu/n = \mu$. Then there are constants $C_x, C_s > 0$, only dependent on the problem data and μ , but not on θ , such that

$$\|\tilde{x}_\mu - x_\mu\|_\infty \leq C_x(\mu) \frac{\theta}{1-\theta}, \quad \|\tilde{s}_\mu - s_\mu\|_\infty \leq C_s(\mu) \frac{\theta}{1-\theta}, \quad \|\tilde{y}_\mu - y_\mu\|_\infty \leq C_x(\mu) \frac{\theta}{1-\theta}.$$

Proof Let $\bar{\mu} \in \mathcal{R}_+^n$ and $(x(\bar{\mu}), y(\bar{\mu}), s(\bar{\mu}))$ be the unique solution to

$$\begin{aligned} A^\top y(\bar{\mu}) + s(\bar{\mu}) &= c \\ Ax(\bar{\mu}) &= b \\ S(\bar{\mu})X(\bar{\mu})e &= \bar{\mu} \\ s(\bar{\mu}), x(\bar{\mu}) &> 0 \end{aligned} \quad (19)$$

then we have $(x_\mu, y_\mu, s_\mu) = (x(\mu e), y(\mu e), s(\mu e))$ and there is a $\tilde{\mu} \in \mathcal{R}_+^n$, such that

$$\|\mu e - \tilde{\mu}\|_2 \leq \theta \mu, \quad e^\top \tilde{\mu} / n = \mu \quad (20)$$

with

$$(\tilde{x}_\mu, \tilde{y}_\mu, \tilde{s}_\mu) = (x(\tilde{\mu}), y(\tilde{\mu}), s(\tilde{\mu})).$$

Differentiating (19) with respect to a component $\bar{\mu}_j$ of $\bar{\mu}$ gives

$$\begin{aligned} A^\top \frac{dy(\bar{\mu})}{d\bar{\mu}_j} + \frac{ds(\bar{\mu})}{d\bar{\mu}_j} &= 0 \\ A \frac{dx(\bar{\mu})}{d\bar{\mu}_j} &= 0 \\ S(\bar{\mu}) \frac{dx(\bar{\mu})}{d\bar{\mu}_j} X(\bar{\mu}) \frac{ds(\bar{\mu})}{d\bar{\mu}_j} &= e_j, \end{aligned}$$

which we can solve for $\frac{dx(\bar{\mu})}{d\bar{\mu}_j}, \frac{dy(\bar{\mu})}{d\bar{\mu}_j}, \frac{ds(\bar{\mu})}{d\bar{\mu}_j}$ to get

$$\frac{dy(\bar{\mu})}{d\bar{\mu}_j} = -(AXS^{-1}A^\top)^{-1}AS^{-1}e_j, \quad (21a)$$

$$\frac{ds(\bar{\mu})}{d\bar{\mu}_j} = A^\top(AXS^{-1}A^\top)^{-1}AS^{-1}e_j, \quad (21b)$$

$$\frac{dx(\bar{\mu})}{d\bar{\mu}_j} = (I - S^{-1}XA^\top(AXS^{-1}A^\top)^{-1}A)S^{-1}e_j. \quad (21c)$$

For any $(x, y, s) \in \mathcal{N}_2(\theta)$ we have

$$(1 - \theta)\mu \leq x_j s_j \leq (1 + \theta)\mu$$

and hence $s_j^{-1} \leq \frac{1}{(1-\theta)\mu} x_j$ which yields

$$\|S^{-1}\|_\infty \leq \frac{1}{(1-\theta)\mu} \|X\|_\infty.$$

Further, from [20, Theorem 3.1] we have the relations

$$|x_j| \leq 2C(d)[C(d)\|d\| + \mu n], \quad |s_j| \leq 2C(d)[C(d)\|d\| + \mu n],$$

which together with (21c) give the bound

$$\left\| \frac{dx(\bar{\mu})}{d\bar{\mu}_j} \right\|_\infty \leq \frac{2}{(1-\theta)\mu} (1 + \chi(A)) \|A\|_\infty C(d)[C(d)\|d\| + \mu n].$$

For a bound on $\left\| \frac{ds(\bar{\mu})}{d\bar{\mu}_j} \right\|_\infty$ we can rewrite (21b) as

$$\frac{ds(\bar{\mu})}{d\bar{\mu}_j} = (X^{-1}S)XS^{-1}A^\top(AXS^{-1}A^\top)^{-1}AS^{-1}e_j,$$

and using $s_j/x_j = s_j^2/(x_j s_j) \leq \frac{1}{\mu(1-\theta)} |s_j|^2$ we obtain

$$\begin{aligned} \left\| \frac{ds(\bar{\mu})}{d\bar{\mu}_j} \right\|_\infty &\leq \frac{1}{\mu^2(1-\theta)^2} \chi(A) \|A\|_\infty \|X\|_\infty \|S\|_\infty^2 \\ &\leq \frac{8}{\mu^2(1-\theta)^2} \chi(A) \|A\|_\infty C(d)^3 [C(d)\|d\| + \mu n]^3. \end{aligned} \quad (22)$$

Finally, for a bound on $\|\frac{dy(\bar{\mu})}{d\bar{\mu}_j}\|_\infty$ we can rewrite (21a) as

$$\frac{dy(\bar{\mu})}{d\bar{\mu}_j} = -(AXS^{-1}A^\top)^{-1}AS^{-1}X(X^{-1}e_j),$$

which yields

$$\|\frac{dy(\bar{\mu})}{d\bar{\mu}_j}\|_\infty \leq \chi(A)\|X^{-1}\| \leq \frac{2}{\mu(1-\theta)}\chi(A)C(d)[C(d)\|d\| + \mu n].$$

By defining

$$\begin{aligned} C_x = C_x(\mu) &= \frac{2}{\mu}(1 + \chi(A))\|A\|_\infty C(d)[C(d)\|d\| + \mu n] \\ C_s = C_s(\mu) &= \frac{8}{\mu^2}\chi(A)\|A\|_\infty C(d)^3[C(d)\|d\| + \mu n]^3 \end{aligned}$$

and since $1 \leq 1/(1-\theta)$ we have

$$\|\frac{dx(\bar{\mu})}{d\bar{\mu}_j}\|_\infty \leq \frac{1}{1-\theta}C_x, \quad \|\frac{ds(\bar{\mu})}{d\bar{\mu}_j}\|_\infty \leq \frac{1}{(1-\theta)^2}C_s, \quad \|\frac{dy(\bar{\mu})}{d\bar{\mu}_j}\|_\infty \leq \frac{1}{1-\theta}C_x.$$

Together with (20) we get

$$\begin{aligned} \|\tilde{x}_\mu - x_\mu\|_\infty &\leq \frac{1}{1-\theta}C_x\|\mu e - \tilde{\mu}\|_\infty \leq \frac{\theta}{1-\theta}C_x(\mu)\mu \\ \|\tilde{s}_\mu - s_\mu\|_\infty &\leq \frac{1}{(1-\theta)^2}C_s\|\mu e - \tilde{\mu}\|_\infty \leq \frac{\theta}{(1-\theta)^2}C_s(\mu)\mu \\ \|\tilde{y}_\mu - y_\mu\|_\infty &\leq \frac{1}{1-\theta}C_x\|\mu e - \tilde{\mu}\|_\infty \leq \frac{\theta}{1-\theta}C_x(\mu)\mu \end{aligned}$$

Lemma 5 For $\theta \in (0, 1)$, let $(\tilde{x}_\mu^R, \tilde{y}_\mu^R, \tilde{\lambda}_\mu^R, \tilde{s}_\mu^R, \tilde{z}_\mu^R) \in \mathcal{N}_2^R(\theta)$. Then there is a $C_3 > 0$ independent of θ such that

$$\|\tilde{x}_\mu^R - x_\mu(\mathcal{T}^R)\|_\infty \leq C_3 \frac{\theta}{1-\theta}, \quad \|\tilde{\lambda}_\mu^R - \lambda_\mu(\mathcal{T}^R)\|_\infty \leq C_3 \frac{\theta}{1-\theta}, \quad \|\tilde{s}_\mu^R - s_\mu(\mathcal{T}^R)\|_\infty \leq C_3 \frac{\theta}{(1-\theta)^2}.$$

Proof This is an immediate consequence of Lemma 4.

From the previous lemma we get that we can bound the difference in the primal-dual first-stage decisions (x, s) of the true μ -center of the full problem $(x_\mu(\mathcal{T}), s_\mu(\mathcal{T}))$ to the calculated approximate μ -center for the reduced problem $(\tilde{x}_\mu^R, \tilde{s}_\mu^R)$ by

$$\begin{aligned} \|x_\mu(\mathcal{T}) - \tilde{x}_\mu^R\|_\infty &\leq \frac{2}{\sqrt{\mu}}\bar{B}\sqrt{L_Q}\sqrt{W_1(\mathcal{T}, \mathcal{T}^R)} + C_x(\mu)\mu\frac{\theta}{1-\theta} \\ \|s_\mu(\mathcal{T}) - \tilde{s}_\mu^R\|_\infty &\leq \frac{2}{\sqrt{\mu^3}}\bar{B}^3\sqrt{L_Q}\sqrt{W_1(\mathcal{T}, \mathcal{T}^R)} + C_s(\mu)\mu\frac{\theta}{(1-\theta)^2} \end{aligned}$$

In the second step of the algorithm we will not find the exact μ -center for all subproblems $P_i(\tilde{x}_\mu^R)$, but rather find points

$$(\tilde{y}_i, \tilde{\lambda}_i, \tilde{s}_i) \in \mathcal{N}_2^{(i)}(\theta).$$

Again we need a bound on the implied error in the dual components λ_i . According to Lemma 5 there is a $C_3 > 0$ such that

$$\|\tilde{\lambda}_i - \lambda_i\|_\infty \leq C_3 \frac{\theta}{1-\theta}.$$

This affects the bound on $\|\Delta c\|$ in the proof of Theorem 1. Finally we are in a position to prove Theorem 2.

Proof (Theorem 2) From (17) and (18) we have that

$$\|\tilde{X}\tilde{S}e - \bar{\mu}e\|_2 \leq \tilde{\theta}\bar{\mu}, \quad \|\tilde{Y}_i\tilde{Z}_ie - \bar{\mu}e\|_2 \leq \tilde{\theta}\bar{\mu}$$

and therefore

$$\begin{aligned} \|(\tilde{X}, \tilde{Y}_1, \dots, \tilde{Y}_n)(\tilde{S}, \tilde{Z}_1, \dots, \tilde{Z}_n)e - \bar{\mu}e\|_2^2 &= \|\tilde{X}\tilde{S}e - \bar{\mu}e\|_2^2 + \sum_i \|\tilde{Y}_i\tilde{Z}_ie - \bar{\mu}e\|_2^2 \\ &\leq (|T| + 1)\tilde{\theta}^2\bar{\mu}^2. \end{aligned} \quad (23)$$

As in the proof to Theorem 1 let

$$\tilde{c} = \sum_i T^\top \tilde{\lambda}_i + \tilde{s}$$

and consider the problem instance $P(\tilde{d})$ obtained from $\mathcal{P}(T)$ by replacing the first-stage cost c with $\tilde{c}$. Then by construction the crash-start point $\tilde{w}$ is primal-dual feasible for $P(\tilde{d})$ and due to (23) satisfies

$$\tilde{w} \in \mathcal{N}_2^{\tilde{d}}(\sqrt{|T|+1}\tilde{\theta}).$$

We will analyse the crash-start as a warm-start attempt for the (perturbed) problem $P(T)$ starting from a point in the $\mathcal{N}_2$ -neighbourhood for problem $P(\tilde{d})$. The change in problem data is

$$\Delta d = (\Delta A, \Delta b, \Delta c) = (0, 0, \Delta c)$$

with

$$\Delta c = c - \tilde{c} = \sum_i T^\top \lambda_{\mu,i} + s_\mu - \left(\sum_i T^\top \tilde{\lambda}_i + \tilde{s} \right) = \sum_i T^\top (\lambda_{\mu,i} - \tilde{\lambda}_i) + (s_\mu - \tilde{s}).$$

Using the bounds from Lemma 3 and Lemma 5 we have

$$\begin{aligned} \|\Delta c\| &= \left\| \sum_i T^\top (\lambda_{\mu,i} - \tilde{\lambda}_i + \tilde{\lambda}_i - \tilde{\lambda}_i) + (s_\mu - \tilde{s} + \tilde{s} - \tilde{s}) \right\| \\ &\leq \|T\| \left(\|\lambda_{\mu,i} - \tilde{\lambda}_i\| + \|\tilde{\lambda}_i - \tilde{\lambda}_i\| \right) + \|s_\mu - \tilde{s}\| + \|\tilde{s} - \tilde{s}\| \\ &\leq \|T\| \left(C_\lambda \|x_\mu - \tilde{x}\| + C_3 \frac{\tilde{\theta}}{1 - \tilde{\theta}} \right) + \|s_\mu - \tilde{s}\| + C_3 \frac{\tilde{\theta}}{1 - \tilde{\theta}}. \end{aligned}$$

Moreover, using the bounds from Lemma 2 we get

$$\begin{aligned} \|\Delta c\| &\leq \|T\| \left(C_\lambda \frac{2}{\sqrt{\mu}} \sqrt{L_Q \bar{B}} \sqrt{W_1(T, T^R)} + C_3 \frac{\tilde{\theta}}{1 - \tilde{\theta}} \right) \\ &\quad + \frac{2}{\sqrt{\mu^3}} \sqrt{L_Q \bar{B}^3} \sqrt{W_1(T, T^R)} + C_3 \frac{\tilde{\theta}}{1 - \tilde{\theta}} \\ &= \left(\|T\| C_\lambda(\mu) + \frac{\bar{B}^2}{\mu} \right) \frac{2}{\sqrt{\mu}} \sqrt{L_Q \bar{B}} \sqrt{W_1(T, T^R)} \\ &\quad + (\|T\| C_\lambda(\mu) + 1) C_3 \frac{\tilde{\theta}}{1 - \tilde{\theta}} \\ &\leq C_4(\mu) \sqrt{W_1(T, T^R)} + C_5 \frac{\tilde{\theta}}{1 - \tilde{\theta}} \end{aligned} \quad (24)$$

where

$$C_4(\mu) = \left(\|T\| C_\lambda(\mu) + \frac{\bar{B}^2}{\mu} \right) \frac{2}{\sqrt{\mu}} \bar{B} \sqrt{L_Q}, \quad C_5 = (\|T\| C_\lambda(\mu) + 1) C_3.$$

Proposition 4.2 of [26] can now be applied with $\theta_0 = \tilde{\theta}\sqrt{|\mathcal{T}|+1}$ and $\xi = \frac{1}{2}(\theta - \tilde{\theta}\sqrt{|\mathcal{T}|+1})$ from which we get

$$\theta - \theta_0 - \xi = \frac{1}{2}(\theta - \tilde{\theta}\sqrt{|\mathcal{T}|+1})$$

and therefore the conditions for a successful warmstart are

$$\|\Delta c\|_\infty \leq \frac{\theta - \tilde{\theta}\sqrt{|\mathcal{T}|+1}}{2(2n+1)\overline{C(d)}}\|\overline{d}\|, \quad \mu \geq \frac{8\overline{C(d)}^2}{\theta - \tilde{\theta}\sqrt{|\mathcal{T}|+1}}\|\Delta c\|_\infty$$

which can be combined to obtain

$$\|\Delta c\|_\infty \leq \frac{\theta - \tilde{\theta}\sqrt{|\mathcal{T}|+1}}{2\overline{C(d)}} \min \left\{ \frac{\|\overline{d}\|}{2n+1}, \frac{\mu}{4\overline{C(d)}} \right\}. \quad (25)$$

Combining (24) and (25) we get as the condition for a successful warmstart

$$C_4(\mu)\sqrt{W_1(\mathcal{T}, \mathcal{T}^R)} + C_5\frac{\tilde{\theta}}{1-\tilde{\theta}} \leq \frac{\theta - \tilde{\theta}\sqrt{|\mathcal{T}|+1}}{2\overline{C(d)}} \min \left\{ \frac{\|\overline{d}\|}{2n+1}, \frac{\mu}{4\overline{C(d)}} \right\} \quad (26)$$

which can be satisfied by keeping $W_1(\mathcal{T}, \mathcal{T}^R)$ and $\tilde{\theta}$ small enough.