

A Modified Augmented Lagrangian Method for Problems with Inconsistent Constraints

Martin P. Neuenhofen¹ and Eric C. Kerrigan²

Abstract—We present a numerical method for the minimization of objectives that are augmented with linear inequality constraints and large quadratic penalties of over-determined inconsistent equality constraints. Such objectives arise from quadratic integral penalty methods for the direct transcription of optimal control problems. The Augmented Lagrangian Method (ALM) has a number of advantages over the Quadratic Penalty Method (QPM) for solving this class of problems. However, if the equality constraints are inconsistent, then ALM might not converge to a point that minimizes the bias of the objective and penalty term. Therefore, in this paper we show a modification of ALM that fits our purpose. We prove convergence of the modified method and prove under local uniqueness assumptions that the local rate of convergence of the modified method in general exceeds the one of the unmodified method. Numerical experiments demonstrate that the modified ALM can minimize certain quadratic penalty-augmented functions faster than QPM, whereas the unmodified ALM converges to a minimizer of a significantly different problem.

I. MOTIVATION IN THE OPTIMAL CONTROL CONTEXT

The method of choice for the numerical solution of optimal control problems is direct transcription. Typical direct transcriptions methods use orthogonal collocation [2]. It is known that the latter can struggle with singular arc and high-index differential algebraic equalities (DAEs); the former arising in the example

$$\begin{aligned} \min_{y,u} \quad & J = \int_0^5 (y(t)^2 + t u(t)) dt, \\ \text{s.t.} \quad & y(0) = 0.5, \quad \dot{y}(t) = \frac{1}{2}y(t)^2 + u(t), \quad (\text{OCP}) \\ & y(t), u(t) \in [-1, 1] \quad \forall t \in [0, 5]. \end{aligned}$$

which has the analytic solution $y^*(t) = \frac{1}{2-t}$ at $t \leq 1$ and $J^* = -\frac{51}{16} - \frac{\log(2)}{2} - \sqrt{32} \tanh \sqrt{\frac{25}{32}} \approx -7.5$.

Quadratic integral penalty methods [1], [6], [10] are an alternative to collocation methods, where the squared path equality constraint residual is integrated and added as a penalty into the objective. In [8] the authors present

¹Martin P. Neuenhofen is with the Department of Electrical & Electronic Engineering, Imperial College London, SW7 2AZ London, UK, e-mail: m.neuenhofen19@imperial.ac.uk www.MartinNeuenhofen.de

²Eric C. Kerrigan is with the Department of Electrical & Electronic Engineering and Department of Aeronautics, Imperial College London, SW7 2AZ London, UK, e-mail: e.kerrigan@imperial.ac.uk www.imperial.ac.uk/people/e.kerrigan

such a method with a proof of convergence under mild assumptions, including convergence for singular arcs and high-index DAEs. This is verified in [8] in comparison to collocation methods via numerical experiments.

Before proceeding, we guide the reader through the solution of (OCP) via the quadratic integral penalty method: let y be represented with continuous and u with discontinuous piecewise linear finite element functions y_h, u_h on a uniform mesh of $N \in \mathbb{N}$ elements, $h := \frac{5}{N}$; represented with $\mathbf{x} := [y_h(h), \dots, y_h(Nh), u_h^+(0), u_h^-(h), u_h^+(h), \dots, u_h^-(Nh)]^T \in \mathbb{R}^n$, $n := 3N$. $y_h(0) = 0.5$ is fixed and removed from \mathbf{x} . We minimize

$$\begin{aligned} \min_{\mathbf{x} \in \mathbb{R}^{2N+1}} \quad & \Phi_\omega(\mathbf{x}) := \int_0^{\frac{\pi}{2}} (y_h(t)^2 + t u_h(t)) dt \\ & + \frac{1}{2\omega} \int_0^5 \left\| -\dot{y}_h(t) + \frac{1}{2}y_h(t)^2 + u_h(t) \right\|_2^2 dt \quad (1) \end{aligned}$$

$$\text{s.t.} \quad \begin{bmatrix} -\mathbf{x} \\ \mathbf{x} \end{bmatrix} \geq \begin{bmatrix} -\mathbf{1} \\ -\mathbf{1} \end{bmatrix}.$$

The integrals are evaluated with Gauss-Legendre quadrature of $q = 8$ points per element. Writing τ, α for abscissae and weights, $m := Nq$, and

$$f(\mathbf{x}) := \sum_{j=1}^{Nq} \alpha_j (y_h(\tau_j)^2 + \tau_j u_h(\tau_j)) \quad (2a)$$

$$c(\mathbf{x}) := \begin{bmatrix} \vdots \\ \sqrt{\alpha_j} (-\dot{y}_h(\tau_j)^2 + \frac{1}{2}y_h(\tau_j)^2 + u_h(\tau_j)) \\ \vdots \end{bmatrix} \in \mathbb{R}^m \quad (2b)$$

allows us to express (1) as a *quadratic penalty program*:

$$\begin{aligned} \min_{\mathbf{x} \in \mathbb{R}^n} \quad & \Phi_\omega(\mathbf{x}) = f(\mathbf{x}) + \frac{1}{2\omega} \|c(\mathbf{x})\|_2^2 \quad (\text{QPP}) \\ \text{s.t.} \quad & \mathbf{A}\mathbf{x} \geq \mathbf{b} \end{aligned}$$

$\omega \in \mathbb{R}_+ \setminus \{0\}$ controls the quadratic penalty and should be chosen on the order of approximation of the finite element space [6], [8]. $\mathbf{A} \in \mathbb{R}^{p \times n}$, $\mathbf{b} \in \mathbb{R}^p$ are $p \in \mathbb{N}_0$ optional inequality constraints that span a non-empty polyhedral $\mathcal{B} := \{\mathbf{x} \in \mathbb{R}^n | \mathbf{A}\mathbf{x} \geq \mathbf{b}\}$.

A sometimes related problem is the *equality constrained program*:

$$\begin{aligned} \min_{\mathbf{x} \in \mathbb{R}^n} \quad & f(\mathbf{x}) \quad (\text{ECP}) \\ \text{s.t.} \quad & c(\mathbf{x}) = \mathbf{0}, \quad \mathbf{A}\mathbf{x} \geq \mathbf{b} \end{aligned}$$

with Lagrangian $\mathcal{L}(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\eta}) := f(\mathbf{x}) - \boldsymbol{\lambda}^\top c(\mathbf{x}) - \mathbf{A}^\top \boldsymbol{\eta}$ and Lagrange multipliers $\boldsymbol{\lambda} \in \mathbb{R}^m$, $\boldsymbol{\eta} \in \mathbb{R}_{\geq 0}^p$, where $\mathbb{R}_{\geq 0}$ are the non-negative real numbers.

Recall that nonlinear inequality constraints can be incorporated into (QPP) and (ECP) via slack variables.

The Karush-Kuhn Tucker optimality system of (QPP) and (ECP) is

$$\begin{aligned} \nabla_{\mathbf{x}} \mathcal{L}(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\eta}) &= \mathbf{0} \\ c(\mathbf{x}) + \omega \boldsymbol{\lambda} &= \mathbf{0} \end{aligned} \quad (\text{KKT1})$$

and

$$\begin{aligned} \mathbf{A}_i \mathbf{x} &= \mathbf{b}_i \quad \text{and} \quad \boldsymbol{\eta}_i \geq 0 & \forall i \in \mathcal{A} \\ \mathbf{A}_i \mathbf{x} &> \mathbf{b}_i \quad \text{and} \quad \boldsymbol{\eta}_i = 0 & \forall i \notin \mathcal{A}, \end{aligned} \quad (\text{KKT2})$$

where $\mathcal{A} \subseteq \{1, \dots, p\}$ is the *active set*, $\mathbf{A}_i, \mathbf{b}_i, \boldsymbol{\eta}_i$ the i th row of $\mathbf{A}, \mathbf{b}, \boldsymbol{\eta}$, and where $\omega = 0$ for (ECP). The multiplier $\boldsymbol{\lambda}$ in (QPP) is a substitution trick such that $\nabla \mathcal{L}(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\eta}) = \nabla \Phi_\omega(\mathbf{x}) - \mathbf{A}^\top \boldsymbol{\eta}$.

We recommend the use of the above penalty finite element method when numerically solving optimal control problems, because penalty methods have favourable convergence properties over collocation methods [8].

A remaining difficulty is with solving (QPP). This can be more challenging than solving a problem of the form (ECP). Below we describe important details.

A. Problems (QPP) and (ECP) have different solutions

It holds $m = Nq \equiv 8N$ and $n = 3N \ll m$ in our discretization (2), hence (ECP) is *over-determined*. In contrast, problem (QPP) cannot be over-determined.

Considering (ECP) in the context of over-determination poses the risk of *inconsistency*. For instance, suppose (OCP) had the additional constraint $y(5) = -1$. Then $c(\mathbf{x}) \neq \mathbf{0} \forall \mathbf{x} \in \mathbb{R}^n$ would follow, i.e. (ECP) would be infeasible due to inconsistent over-determination, whereas (QPP) remains feasible.

In the example, (ECP) is over-determined but feasible. Namely, regardless of the mesh size $h > 0$, $c(\mathbf{x}) = \mathbf{0}$ is achieved (only) when $y_h(t) = 0.5$, $u_h(t) = -0.125$, because the differential equation has no non-trivial piecewise linear solution. This is far away from y^*, u^* , to which solutions of (QPP) converge as $h, \omega \searrow 0$. In conclusion, solutions to (ECP) and (QPP) can be very different.

The Modified Augmented Lagrangian Method (MALM), discussed in Section II, converges to minimizers of (QPP) instead of (ECP).

B. Solutions of (QPP) depend on the value of ω

As experimentally verified in [6] and outlined in the analysis in [8], the discretization (1)–(2) converges when both $h, \omega \searrow 0$. That is, for fixed h , too large values of ω result in bad feasibility of the numerical optimal control solution, whereas too small values of ω result in feasible, yet far-from-optimal solutions.

Fig. 1 demonstrates this. Our discretization of (OCP) with $N = 50$ is solved $\forall \omega \in \{10^1, 10^{-3}, 10^{-9}\}$. The value ω determines the bias between minimization of $J := f(\mathbf{x})$ and $r := \|c(\mathbf{x})\|_2 \equiv \|\dot{y}_h + \frac{1}{2}y_h + u_h\|_{L^2(0, \pi/2)}$;




Fig. 1. Numerical solution to (OCP) for $N = 50$ and different ω .

note that the latter \equiv is accurate due to the order of quadrature. For $\omega = 10^{-1}$ the solution achieves a good trade-off between feasibility and optimality on that coarse mesh.

In conclusion, the value of ω has a significant influence on the solution of (QPP). Hence, it is important that (QPP) be minimized for the specified value of ω . The value ω appears in the dual update formula of MALM. This is important and ensures that MALM converges to minimizers of (QPP) for the specific value of ω .

C. Direct minimization of (QPP) is inefficient

To the unprejudiced it appears natural to minimize the objective Φ_ω directly. However, (unless c is an affine function) this will result in very many iterations. This is so because the *nonlinear* penalties result in *curved* valleys through which iterative minimization methods make slow progress.

To demonstrate this inefficiency, consider the instance

$$f(\mathbf{x}) := -x_1 - x_2 \quad (3a)$$

$$c(\mathbf{x}) := \begin{bmatrix} (x_1 + \varepsilon)^2 + x_2^2 - 2 \\ (x_1 - \varepsilon)^2 + x_2^2 - 2 \end{bmatrix} \in \mathbb{R}^m \quad (3b)$$

$$\mathbf{A} := \begin{bmatrix} 0 & 1 \\ 1 & -1 \end{bmatrix}, \mathbf{b} := \begin{bmatrix} 0 \\ 0 \end{bmatrix} \in \mathbb{R}^p \quad (3c)$$

with primal and dual initial guesses $\mathbf{x}_0 := [2 \ 1]^\top$ and $\boldsymbol{\lambda}_0 := \mathbf{0}$, for $\varepsilon = 0$. We discuss later with Table II that minimization of (QPP) of (3) with a direct minimization method takes 134 iterations when $\omega = 10^{-6}$. This is inefficient when compared to MALM, which solves the same instance in only 31 iterations.

The Augmented Lagrangian Method (ALM) uses a local minimization method (typically a quasi-Newton variant) for the primal variables, then updates the duals. Since ALM eventually solves (ECP), the quasi-Newton system must solve (KKT1) with $\omega = 0$. In contrast, MALM converges to minimizers of (QPP), i.e. solves (KKT1) with $\omega > 0$. This yields a dual regularization, which keeps the magnitude of $\|\boldsymbol{\lambda}\|_2$ bounded and locally unique, which improves the convergence of the Newton iteration.

D. Contributions

A modified augmented Lagrangian method for convex quadratic objectives with linear equality constraints is presented in [14]. In [9] we generalize this method from equality-constrained convex linear-quadratic to non-convex nonlinear optimization problems with nonlinear equality constraints. In the present paper, we generalize it further by adding linear inequality constraints.

The authors in [14] show convergence of their method for equality-constrained convex linear-quadratic programs. In the present paper, we show convergence of our generalized method for non-convex nonlinear programs with additional linear inequality constraints under the following assumptions: twice continuous differentiability of f, c , boundedness of \mathcal{B} , and boundedness of Φ_ω over \mathcal{B} . In addition, we derive an improved local convergence result that is valid even for our generalized method, under these additional assumptions: local primal uniqueness, local Lipschitz continuity of first and second derivatives, and strict complementarity. Our analysis draws connections between the rate of convergence for a modified and non-modified augmented Lagrangian method.

E. Structure of the Paper

Section II derives the proposed generalized modified augmented Lagrangian method. Section III presents numerical experiments. This section also elaborates on the numerical difficulties of solving either (QPP) or (ECP) and suitable values of ω for a given instance $\mathbf{x}_0, \boldsymbol{\lambda}_0, f, c$.

II. DERIVATION OF MALM

MALM is a solution method for (QPP). MALM has been presented in [14] for the special case when f is quadratic, c is linear, and Φ_ω is convex. Here, we derive MALM for general nonlinear non-convex programs, and in a stronger relation to its origins in ALM [7], [12]. This will later allow us to show global convergence not only for convex linear-quadratic problems, but also for non-convex nonlinear problems with additional inequality constraints. Furthermore, we will give a rate-of-convergence result in the appendix, which shows that in the limit MALM converges generally faster than ALM for the same problem instance.

We derive MALM for (QPP) from ALM for (ECP). To apply ALM, we need an auxiliary problem of the form (ECP) instead. Our approach to achieving this works by temporarily using an auxiliary variable $\boldsymbol{\xi} \in \mathbb{R}^m$. This variable will be eliminated in the augmented optimality system.

A. Auxiliary Problem

Consider the following equivalent problem to (QPP):

$$\min_{\hat{\mathbf{x}} := (\mathbf{x}, \boldsymbol{\xi}) \in \mathbb{R}^{(n+m)}} \hat{f}(\hat{\mathbf{x}}) := f(\mathbf{x}) + \frac{\omega}{2} \|\boldsymbol{\xi}\|_2^2 \quad (4a)$$

$$\text{s.t.} \quad \hat{c}(\hat{\mathbf{x}}) := c(\mathbf{x}) + \omega \boldsymbol{\xi} = \mathbf{0}, \quad (4b)$$

$$\hat{\mathbf{A}} \hat{\mathbf{x}} \geq \mathbf{b}, \quad (4c)$$

where $\hat{\mathbf{A}} := [\mathbf{A} \ \mathbf{0}]$. Introducing $\boldsymbol{\lambda} \in \mathbb{R}^m, \boldsymbol{\eta} \in \mathbb{R}_{\geq 0}^p$, the optimality conditions of (4) are (KKT2) and

$$\begin{bmatrix} \nabla f(\mathbf{x}) \\ \omega \boldsymbol{\xi} \end{bmatrix} - \begin{bmatrix} \nabla c(\mathbf{x}) \\ \omega \mathbf{I} \end{bmatrix} \boldsymbol{\lambda} - \hat{\mathbf{A}}^\top \boldsymbol{\eta} = \mathbf{0} \quad (5a)$$

$$c(\mathbf{x}) + \omega \boldsymbol{\xi} = \mathbf{0}. \quad (5b)$$

B. Augmented Optimality System

Since (4) is of form (ECP), we can apply ALM [11, Alg. 17.3]. There are different ways for doing this: Rockafellar uses a method with augmented Lagrangians for inequalities [13], whereas typical codes like LANCELOT [4] use augmented Lagrangians only for the equality constraints. LANCELOT proceeds as such for mainly two reasons: First, Rockafellar's approach results in a non-smooth gradient of the unconstrained problem. Second, convergence theory for augmented Lagrangians to inequality constraints is less understood and requires further assumptions. For these two reasons, we also use augmented Lagrangians only for the equality constraints. To this end, we augment (5) with an auxiliary vector $\mathbf{z} \in \mathbb{R}^m$ and a moderate penalty parameter $\rho > 0$:

$$\begin{bmatrix} \nabla f(\mathbf{x}) \\ \omega \boldsymbol{\xi} \end{bmatrix} - \begin{bmatrix} \nabla c(\mathbf{x}) \\ \omega \mathbf{I} \end{bmatrix} (\boldsymbol{\lambda} + \mathbf{z}) - \hat{\mathbf{A}}^\top \boldsymbol{\eta} = \mathbf{0} \quad (6a)$$

$$c(\mathbf{x}) + \omega \boldsymbol{\xi} + \rho \mathbf{z} = \mathbf{0}. \quad (6b)$$

The intuition for doing so in ALM is similar to what we did in (KKT1) for (QPP), where the Lagrange multipliers were used as a penalty substitute to ensure that the

Algorithm 1 Modified Augmented Lagrangian Method

```
1: procedure MALM( $f, c, \omega, \mathbf{x}_0, \boldsymbol{\lambda}_0, \mathbf{A}, \mathbf{b}, \text{tol}$ )
2:    $\rho \leftarrow \rho_0$ 
3:   for  $k = 1, 2, 3, \dots, k_{\max}$  do
4:     Compute  $\mathbf{x}_k$  and optionally  $\boldsymbol{\eta}_k$  by solving
           
$$\min_{\mathbf{x} \in \mathbb{R}^n} \Psi_k(\mathbf{x}) \quad \text{s.t. } \mathbf{A}\mathbf{x} \geq \mathbf{b}. \quad (9)$$

5:     Update  $\boldsymbol{\lambda}_k \leftarrow \boldsymbol{\lambda}_{k-1} - \frac{1}{\omega + \rho} (c(\mathbf{x}_k) + \omega \boldsymbol{\lambda}_{k-1})$ 
6:     if  $\|c(\mathbf{x}_k) + \omega \boldsymbol{\lambda}_k\|_\infty \leq \text{tol}$  then
7:       return  $\mathbf{x}_k, \boldsymbol{\lambda}_k$  and optionally  $\boldsymbol{\eta}_k$ 
8:     else
9:       Decrease  $\rho \leftarrow c_\rho \rho$  to promote convergence.
10:    end if
11:  end for
12: end procedure
```

gradient of \mathcal{L} matches that of Φ_ω . Likewise, here \mathbf{z} works as a penalty substitute for (5b).

We could use (6) directly in order to form an ALM iteration. That iteration would consist of alternately solving the optimality system (6) together with (KKT2) for $(\mathbf{x}, \boldsymbol{\xi}, \mathbf{z}, \boldsymbol{\eta}, \mathcal{A})$ where $\boldsymbol{\lambda}$ is fixed and updating $\boldsymbol{\lambda} \leftarrow \boldsymbol{\lambda} + \mathbf{z}$, this being equivalent to $\boldsymbol{\lambda} \leftarrow \boldsymbol{\lambda} - \frac{1}{\rho} (c(\mathbf{x}) + \omega \boldsymbol{\xi})$.

C. Elimination of the Auxiliary Vector

Instead, we propose to eliminate $\boldsymbol{\xi} = \boldsymbol{\lambda} + \mathbf{z}$ to obtain

$$\nabla f(\mathbf{x}) - \nabla c(\mathbf{x})(\boldsymbol{\lambda} + \mathbf{z}) - \mathbf{A}^\top \boldsymbol{\eta} = \mathbf{0} \quad (7a)$$

$$c(\mathbf{x}) + \omega \boldsymbol{\lambda} + (\omega + \rho) \mathbf{z} = \mathbf{0}. \quad (7b)$$

As in ALM, we solve (7) and (KKT2) with an iteration of two alternating steps:

- 1) Keep the value of $\boldsymbol{\lambda}$ fixed, and solve (7) and (KKT2) for $(\mathbf{x}, \mathbf{z}, \boldsymbol{\eta}, \mathcal{A})$.
- 2) Update $\boldsymbol{\lambda}$ as $\boldsymbol{\lambda} \leftarrow \boldsymbol{\lambda} + \mathbf{z}$.

Analogous to ALM, the first step can be realized by minimizing an augmented Lagrangian function for \mathbf{x} at fixed $\boldsymbol{\lambda}$ subject to $\mathbf{x} \in \mathcal{B}$, whereas in the second step \mathbf{z} can be expressed in terms of \mathbf{x} from (7b). Using this, the method can be expressed in Algorithm 1, where

$$\Psi_{k+1}(\mathbf{x}) := \mathcal{L}(\mathbf{x}, \boldsymbol{\lambda}_k, \mathbf{0}) + \frac{0.5}{\omega + \rho} \|c(\mathbf{x}) + \omega \boldsymbol{\lambda}_k\|_2^2 \quad (8)$$

is the augmented Lagrangian function, with $\mathcal{L}(\mathbf{x}, \boldsymbol{\lambda}, \mathbf{0}) \equiv f(\mathbf{x}) - \boldsymbol{\lambda}^\top \cdot c(\mathbf{x})$ from (KKT1).

Practical values are $\text{tol} = 10^{-8}$, $c_\rho = 0.1$, $\rho_0 = 0.1$. Care must be taken that Ψ_k in (9) attains a lower bound on the feasible set. To this end, practical methods use box constraints [4, eq. 3.2.2] or a trust-region [4, eq. 3.2.4].

In order to minimize (9), one can use any numerical method for linear inequality constrained nonlinear minimization; e.g. an interior-point method like IPOPT [15] or an active set method like SNOPT [5].

We refer to [11, eq. 17.21] for details on how the quasi-Newton direction for the quadratic penalty function can

be computed in a more numerically stable fashion from a saddle-point linear equation system.

D. Discussion

1) *True Generalization of ALM:* MALM is a true generalization of ALM, because they differ only by the parameter ω . If $\omega = 0$ then MALM is identical to ALM. Both methods enjoy the same convergence properties and approach the same limit point. In contrast to ALM, by selecting $\omega > 0$, we show below that MALM converges to critical points of (QPP) with the given ω .

2) *Benefit:* MALM solves the penalty function Φ_ω in (QPP) by minimizing a sequence of penalty functions Ψ_k . When does this make sense? If we select $\rho \gg \omega$. Thereby, the penalty functions Ψ_k have less steep valleys and hence can often be minimized more efficiently in comparison to a single minimization of Φ_ω . The numerical experiments below verify this claim.

E. Global and Local Convergence

We prove convergence by showing that the iterates of MALM are the same as those of ALM applied to (4). The analysis is unfortunately awkward due to notation for f, c, \mathbf{x} and $\hat{f}, \hat{c}, \hat{\mathbf{x}}$.

The below theorem assumes that all sub-problems (9) are solved exactly, and that computations are performed in exact arithmetic. Throughout this subsection, MALM means the callback-function in Algorithm 1, wherein any black-box method can be used to solve (9).

Theorem 1 (Global Convergence): Let $\omega > 0$, choose $\Omega \subset \mathbb{R}^n$ bounded, c bounded in Ω , and let f, c be twice continuously differentiable in Ω . Consider the call $\mathcal{C} := \text{MALM}(f, c, \omega, \mathbf{x}_0, \boldsymbol{\lambda}_0, \mathbf{A}, \mathbf{b})$. Suppose all iterates $\{\mathbf{x}_k\}_{k \in \mathbb{N}_0}$ of \mathcal{C} live in Ω . If ρ_0 is sufficiently small then $\{\mathbf{x}_k\}_{k \in \mathbb{N}_0}$ converges to a critical point of (QPP).

See Appendix I for the proof. The theorem text defines the instance $\mathcal{I} := (f, c, \omega, \mathbf{x}_0, \boldsymbol{\lambda}_0, \mathbf{A}, \mathbf{b})$. The proof works by identifying the iterates of \mathcal{C} with iterates of a different call \mathcal{D} . Before that, convergence of \mathcal{D} is asserted by a result from the literature. Section 4 in [3] discusses how the assumption $\{\mathbf{x}_k\}_{k \in \mathbb{N}_0} \subset \Omega$ is not strong but can be ensured; e.g. by construction such that $\mathcal{B} \subseteq \Omega$.

The global convergence analysis of MALM shows that MALM converges like ALM on a different problem. However, when applying both methods onto the *same* instance, then MALM converges locally faster than ALM, assuming the primal solution is locally unique, first and second derivatives of f, c are Lipschitz continuous, and strict complementarity. This is shown in Appendix II.

III. NUMERICAL EXPERIMENTS

We present two numerical test problems with nonlinear equality constraints and linear inequalities. The first is instructional, the second is(2). The sub-problems in (9) are solved with IPOPT. This is a line-search primal-dual interior-point method. For cost comparison, we compare the number of IPOPT iterations. Since MALM

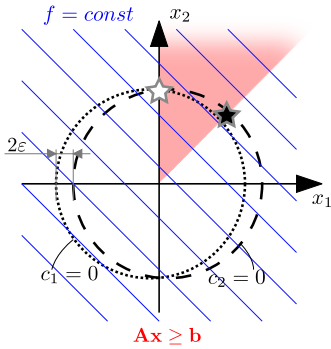


Fig. 2. Geometry of the Circle Problem, with level sets of f, c_1, c_2 .

solves a sequence of problems (9), sub-solvers with better warm-starting capabilities would likely converge in fewer iterations. For examples with equality constraints only, we refer to our experiments in [9].

A. Circle Problem

1) *Setting*: This problem considers the instance (3) for various values of ε , with initial guess $\mathbf{x}_0 = [2 \ 1]^T$, $\boldsymbol{\lambda}_0 = \mathbf{0}$. Fig. 2 shows the geometry of the instance: Level sets of f, c_1, c_2 are blue, dotted, and dashed, respectively. The feasible region to the inequality constraints is indicated in red. The figure shows two points $\mathbf{x}_A := [0 \ \sqrt{2}]^T$, $\mathbf{x}_B := [1 \ 1]^T$ as white and black star, respectively.

The instance can be interpreted in either of two ways:

- (A) Either we meant c in a precise sense, meaning we wish to find a solution to $c(\mathbf{x}) = \mathbf{0}$ and, if non-unique, select the point that minimizes $f(\mathbf{x})$.
- (B) Or we actually meant c in a rough sense, meaning we wish to minimize f subject to $\|\mathbf{x}\|_2^2 = 2 + \mathcal{O}(\varepsilon)$.

Both problems are reasonable in their own right: For example, (A) makes sense when we have to solve a complicated equation system and want to find a desirable solution. On the other hand, (B) makes sense when our constraints suffer from errors, e.g. measurement errors or consistency errors, such as by discretization. For example, imagine a discretized optimal control problem, where c inherits consistency errors that have the size of ε .

Crucially, both solutions $\mathbf{x}_A, \mathbf{x}_B$ can be characterized sharply with a suitable problem statement: \mathbf{x}_A is the solution of (ECP). Less obvious, \mathbf{x}_B can be computed as the solution of (QPP) when choosing ω suitable with respect to ε . Here, a suitable choice is $\omega = \mathcal{O}(\varepsilon)$. To see this, notice that (KKT1) admits a well-scaled solution $\|\mathbf{x}\|_2, \|\boldsymbol{\lambda}\|_2 = \mathcal{O}(1)$ and $\|c(\mathbf{x})\| = \mathcal{O}(\varepsilon)$ when this selection for ω is made.

Lastly, we stress that for this instance the solution \mathbf{x}_A has an ill-conditioned KKT system with a dual solution $\|\boldsymbol{\lambda}\| = \mathcal{O}(1/\varepsilon^2)$, whereas \mathbf{x}_B is well-behaved, i.e. its KKT equations are well-conditioned.

2) *Computational Results*: We solve the instance with MALM and QPM, for various values of ε, ω , including 0. We implement QPM by solving (QPP) directly in IPOPT with objective Φ_ω . Recall that MALM=ALM for $\omega = 0$

TABLE I

SOLUTION OF THE CIRCLE PROBLEM WITH RESPECT TO ε, ω .
SMALLER VALUES MEAN CLOSER CONVERGENCE TO EITHER POINT.
CELLS IN THE LOWER LEFT CONVERGE TO \mathbf{x}_A , CELLS IN THE UPPER RIGHT TO \mathbf{x}_B .

		ε				
		$1.0e-1$	$1.0e-2$	$1.0e-4$	$1.0e-6$	0.0
ω	$1.0e-1$	1.1e-0 2.6e-3	1.1e+0 4.3e-3	1.1e+0 4.4e-3	1.1e+0 4.4e-3	1.1e+0 4.4e-3
	$1.0e-2$	1.2e-1 9.6e-1	1.1e+0 3.7e-4	1.1e+0 4.4e-4	1.1e+0 4.4e-4	1.1e+0 4.4e-4
	$1.0e-4$	3.8e-3 1.1e+0	1.2e-1 9.6e-1	1.1e+0 4.4e-6	1.1e+0 4.4e-6	1.1e+0 4.4e-6
	$1.0e-6$	3.5e-3 1.1e+0	1.3e-3 1.1e+0	1.1e+0 3.7e-8	1.1e+0 4.4e-8	1.1e+0 4.4e-8
	$1.0e-8$	3.5e-3 1.1e+0	3.7e-5 1.1e+0	1.2e-1 9.6e-1	1.3e+0 7.1e-9	1.3e+0 7.1e-9
	0.0	3.5e-3 1.1e+0	3.7e-5 1.1e+0	1.2e-1 9.6e-1	1.3e+0 7.1e-9	1.3e+0 0.0

and that QPM is not applicable (n.a.) when $\omega = 0$, since Φ_ω is undefined.

We also analyze the limit points \mathbf{x}_∞ (which are identical for both tested methods throughout all tests) for each ε, ω , by measuring the quantities

$$e_A := \|\mathbf{x}_\infty - \mathbf{x}_A\|_2, \quad e_B := \|\mathbf{x}_\infty - \mathbf{x}_B\|_2.$$

Table I shows the quantities e_A, e_B for respective ε, ω . Dividing the table into a lower left and an upper right triangle, we see that solutions in the lower triangle rather converge to \mathbf{x}_A while those on the diagonal and in the upper right converge to \mathbf{x}_B .

Table II shows the sum of the number of all inner iterations of QPM and MALM for respective ε, ω . We see a trend for each of the two methods: QPM converges in a few iterations when ω is moderate. However, when ε, ω both decrease, the iteration count blows up. The trend for MALM is different. MALM converges reliably for all ε, ω in the upper right triangle, including those where ε, ω are very small. The last row shows ALM. ALM converges quickly when $\varepsilon = 0$, but its iteration count blows up for positive decreasing values of ε . In two instances ALM did not converge (n.c.) within 1000 iterations.

3) *Interpretation of Results*: Table I confirms that, depending on parameters ε, ω , we either solve for \mathbf{x}_A or \mathbf{x}_B . Table II indicates that \mathbf{x}_A cannot be computed numerically efficiently, as expected due to the almost linearly dependent constraints. In this regard we find that ALM does not converge in the last row when $\varepsilon = 10^{-4}$.

In the optimal control context we are interested in solving problems in the sense of (B), hence we now investigate the upper right triangle in Table II. We see that both methods converge for all of these instances. Yet, when ε, ω are small, but strictly positive, MALM outperforms QPM. This is relevant because there are problems from optimal control discretizations with inconsistencies from discretization errors. The present test

TABLE II

TOTAL NUMBER OF IPOPT ITERATIONS FOR MALM AND QPM FOR THE CIRCLE PROBLEM WITH RESPECT TO ε, ω . FEWER ITERATIONS MEAN BETTER COMPUTATIONAL EFFICIENCY.

		ε				
		1.0e-1	1.0e-2	1.0e-4	1.0e-6	0.0
ω	1.0e-1	28	22	22	19	19
	1.0e-2	14	13	13	13	13
1.0e-4	1.0e-2	36	28	16	23	20
	1.0e-4	12	16	16	16	16
1.0e-6	1.0e-4	21	56	32	29	23
	1.0e-6	16	36	43	43	43
1.0e-8	1.0e-6	29	68	45	39	31
	1.0e-8	16	35	138	134	134
0.0	1.0e-8	34	60	n. c.	52	40
	0.0	n. c.	n. c.	n. c.	429	374
		34	60	n. c.	52	40
		n. a.	n. a.	n. a.	n. a.	n. a.

problem models these inconsistencies quantitatively with the parameter ε . In the case of discretizations, we wish to drive $\omega, h \searrow 0$ to decrease the discretization error, which brings us into the lower right region of the table.

B. Optimal Control Problem

1) *Setting*: We solve the instance (2) for various values of N, ω for $\mathbf{x}_0 = \mathbf{0}, \boldsymbol{\lambda}_0 = \mathbf{0}$. Recall that an instance represents the discretization (1) with mesh size $h = \frac{5}{N}$, which only converges to the analytical solution when $h, \omega \searrow 0$ together, as was discussed along Fig. 1.

As for the former experiment, we solve the instance with MALM and QPM, for various values of N, ω . We also analyze the limit points \mathbf{x}_∞ (which we confirmed to be identical for both tested methods throughout all tests) for each N, ω , by computing the quantities

$$\delta J := f(\mathbf{x}) - J^* \equiv J(y_h, u_h) - J^*,$$

$$r := \|c(\mathbf{x})\|_2 \equiv \left\| -\dot{y}_h + \frac{1}{2}y_h^2 + u_h \right\|_{L^2(0, \pi/2)}.$$

Recall from the circle test problem that there are two interpretations (A) and (B) for the instance. As is clear from the context, we wish to compute a solution of kind (B). However, if we choose ω very small on a coarse mesh then the iteration will converge to a solution of kind (A).

2) *Computational Results*: Table III shows the quantities $\delta J, r$ for respective N, ω . Dividing the table into a lower left and an upper right triangle, we see that solutions in the lower left achieve good feasibility but at the sacrifice of optimality, whereas solutions in the upper right are not sufficiently feasible. When being limited by computation time to solve on a moderately sized mesh, then accordingly ω should not be chosen too small. In contrast to ALM, which always sets $\omega = 0$, then MALM allows for choosing ω of moderate size.

Table IV shows the sum of the number of all inner iterations of MALM and QPM for respective N, ω . We see the same trend as for the circle problem: QPM converges in a few iterations when ω is moderate. In

TABLE III

SOLUTION OF THE OPTIMAL CONTROL PROBLEM WITH RESPECT TO N, ω . FOR A GIVEN MESH SIZE N , THE VALUE FOR ω IS SUITABLE WHEN δJ (OPTIMALITY GAP) AND r (FEASIBILITY RESIDUAL) HAVE SIMILAR MAGNITUDE.

		N				
		50	250	500	1000	2000
ω	1.0e-2	-1.7e-1	-1.7e-1	-1.7e-1	-1.7e-1	-1.7e-1
	1.0e-3	9.2e-2	4.3e-2	4.3e-2	4.2e-2	4.3e-2
1.0e-4	1.0e-3	4.3e-2	-8.6e-3	-9.6e-3	-9.8e-3	-9.9e-3
	1.0e-4	2.7e-2	5.3e-3	4.6e-3	4.4e-3	4.4e-3
1.0e-5	1.0e-4	7.6e-2	1.9e-2	1.2e-2	-8.7e-3	-7.2e-3
	1.0e-5	8.8e-3	6.1e-4	6.1e-4	6.0e-4	5.7e-4
1.0e-6	1.0e-5	7.9e-2	2.3e-2	1.6e-2	1.2e-2	1.0e-2
	1.0e-6	2.5e-3	6.4e-5	6.2e-5	6.3e-5	6.6e-5
0.0	1.0e-6	8.0e-2	2.3e-2	1.6e-2	1.2e-2	1.0e-2
	0.0	6.5e-4	1.8e-5	7.5e-6	1.1e-5	1.5e-5
		7.2e+0	7.2e+0	7.2e+0	7.2e+0	7.2e+0
		0.0	0.0	0.0	0.0	0.0

TABLE IV

TOTAL NUMBER OF IPOPT ITERATIONS FOR MALM AND QPM FOR THE OPTIMAL CONTROL PROBLEM WITH RESPECT TO N, ω . FEWER ITERATIONS MEAN BETTER COMPUTATIONAL EFFICIENCY.

		N				
		50	250	500	1000	2000
ω	1.0e-2	39	53	64	69	69
	1.0e-3	17	24	23	26	44
1.0e-4	1.0e-3	47	63	74	90	89
	1.0e-4	38	44	43	44	68
1.0e-5	1.0e-4	61	66	80	93	121
	1.0e-5	57	92	85	133	124
1.0e-6	1.0e-5	64	78	80	93	113
	1.0e-6	102	161	266	242	252
0.0	1.0e-6	81	78	110	97	126
	0.0	163	222	328	278	224
		n. c.	n. c.	n. c.	n. c.	n. c.
		n. a.	n. a.	n. a.	n. a.	n. a.

contrast, when h, ω both decrease then the iteration count of QPM increases. In contrast, MALM converges reliably for all N, ω in the upper right triangle, including those where N is very large and ω very small. The last row shows that ALM does not convergence (n.c.) within 500 iterations for any mesh size.

C. Interpretation of the Results

For this test problem, Table III demonstrates that the numerical solution converges to the optimal control solution when $h, \omega \searrow 0$ together.

Table IV shows that QPM converges fast when N is small and ω is moderate. In contrast to this, for large N and small ω , MALM is clearly more efficient in terms of total IPOPT iteration count. However, Table III reveals that large N and small ω are a necessity for the numerical computation of accurate optimal control solutions.

Importantly, both experiments make clear that ALM is unsuitable for solving applications with inconsistent constraints: For the experiment depicted in Fig. 2, ALM converges to \mathbf{x}_A (white star) whenever $\varepsilon > 0$. For optimal

control problems, the magnitude of ε models consistency errors of discretizations, in which case \mathbf{x}_B is the sought solution of a well-conditioned KKT equation system, whereas \mathbf{x}_A has ill-conditioned KKT equations (due to unboundedness of $\boldsymbol{\lambda}$ as $\varepsilon \searrow 0$). Table IV shows that ALM fails to converge for the discretized control problem because it attempts to solve an ill-conditioned problem. Table III shows further that the exact minimizer is undesired here anyways because $\omega = 0$ results in a bad balance between the goals of minimizing both r and δJ , illustrated in Fig. 1.

IV. CONCLUSIONS

We presented a modified augmented Lagrangian method (MALM), generalized to nonlinear non-convex optimization problems with additional linear inequality constraints. We proved global convergence for our generalized method for non-convex nonlinear programs. A local rate-of-convergence result was given in the appendix. The result shows that MALM inherits all the local convergence results of ALM while the regularization in $\omega > 0$ also yields a slight local convergence improvement.

Our numerical experiments demonstrate that MALM outperforms QPM when minimizing quadratic penalty programs (QPP) in those situations where ω is very small, in a similar manner as when ALM outperforms QPM when solving equality-constrained programs (ECP). The experiments further show that there are problem instances where it is beneficial to solve a problem of type (QPP) rather than (ECP), one important class arising from integral penalty discretizations of optimal control problems.

In this paper we have assumed that the sub-problems (9) are solved to high accuracy. Future work could extend the approach to inexact iterations and sub-iterations to mild tolerances. This could reduce computations at sub-iterations where the dual is far from converged.

APPENDIX I

PROOF OF THEOREM 1

From \mathcal{I} , construct the instance $\mathcal{J} := (\hat{f}, \hat{c}, 0, \hat{\mathbf{x}}_0, \hat{\boldsymbol{\lambda}}_0, \hat{\mathbf{A}}, \mathbf{b})$ from (4) with $\boldsymbol{\xi}_0$ arbitrary.

Define the call $\mathcal{D} := \text{MALM}(\mathcal{J})$, which means solving instance \mathcal{J} with ALM; cf. Section II-D.1. To understand the proof, notice that due to differing call arguments of MALM the notation differs with respect to line 4 in Algorithm 1: \mathcal{C} computes \mathbf{x}_k , whereas \mathcal{D} computes $\hat{\mathbf{x}}_k$.

We first show that *any* limit point $\hat{\mathbf{x}}_\infty$ of \mathcal{D} converges to a critical point of (4). We use [3, Thm 4.6], which requires [3, AS1] feasibility of (4), [3, AS2] twice continuous differentiability of \hat{f}, \hat{c} , [3, AS3] $\hat{\mathbf{x}}_k \in \Omega \times c(\Omega)$ bounded, and [3, AS4] a technical condition discussed below. All requirements are satisfied: AS1 follows from the requirement that \mathcal{B} be non-empty. AS2 is stated in the theorem text. AS3 follows from boundedness of Ω and from boundedness of c over Ω . AS4 requires $\nabla \hat{c}(\hat{\mathbf{x}}) \cdot \mathbf{Z}$ to be of column rank $\geq m$, where \mathbf{Z} is a matrix of

orthonormal columns that span the null-space of $\hat{\mathbf{A}}_{\mathcal{A}}$, i.e. the matrix of sub-rows of $\hat{\mathbf{A}}$ of the active constraints at $\hat{\mathbf{x}}$. Due to the special structure of $\hat{\mathbf{A}}$, we see that \mathbf{Z} has a structure like

$$\mathbf{Z} = \begin{bmatrix} \mathbf{0} & \cdots \\ \mathbf{I} & \cdots \end{bmatrix}.$$

Since $\nabla \hat{\mathbf{x}}(\hat{\mathbf{x}})^\top = [\nabla c(\hat{\mathbf{x}})^\top \mathbf{I}]$ has full row rank, the rank of $\nabla \hat{c}(\hat{\mathbf{x}})^\top \cdot \mathbf{Z}$ is bounded below by the number of columns of \mathbf{Z} , i.e. bounded below by m .

Since any inner method for minimization of (9) can be used, the sequence of iterates $\mathbf{x}_k, \boldsymbol{\lambda}_k$ of calls \mathcal{C}, \mathcal{D} are actually non-unique, depending on whether Ψ_k has multiple local minima.

We show by induction that for each iterate $\mathbf{x}_k, \boldsymbol{\lambda}_k$ of \mathcal{C} there exists an identical sequence of iterates for call \mathcal{D} . Base: For $k = 0$ the proposition holds by construction of the initial guesses. Step: Let the proposition hold for $k - 1$. We now show that the proposition holds for k . The iterate $\hat{\mathbf{x}}_k$ from \mathcal{D} in line 4 necessarily satisfies $\nabla_{\hat{\mathbf{x}}} \Psi_k(\hat{\mathbf{x}}_k) - \hat{\mathbf{A}}^\top \cdot \boldsymbol{\eta}_k = \mathbf{0}$, which is equivalent to (6) after elimination of \mathbf{z} by means of (6b). From the second component of (6a) it follows that

$$\boldsymbol{\xi}_k = \boldsymbol{\xi}(\mathbf{x}_k, \boldsymbol{\lambda}_{k-1}) := \frac{1}{\omega + \rho} (\rho \boldsymbol{\lambda}_{k-1} - c(\mathbf{x}_k)). \quad (10)$$

I.e. $\boldsymbol{\xi}_k$ is uniquely determined in \mathcal{D} , hence $\hat{\mathbf{x}}_k$ takes on the form $\hat{\mathbf{x}}_k = (\mathbf{x}_k, \boldsymbol{\xi}(\mathbf{x}_k, \boldsymbol{\lambda}_{k-1}))$ for some \mathbf{x}_k . Substituting (10) into the first component of (6a) yields $\nabla_{\mathbf{x}} \Psi_k(\mathbf{x}_k) = \mathbf{0}$, which is indeed identical to what \mathbf{x}_k in line 4 of \mathcal{C} satisfies. Thus, $\hat{\mathbf{x}}_k = (\mathbf{x}_k, \boldsymbol{\xi}(\mathbf{x}_k))$ with \mathbf{x}_k from \mathcal{C} is a valid k th iterate of \mathcal{D} . Finally, notice that $\boldsymbol{\lambda}_k$ in \mathcal{C}, \mathcal{D} are identical because

$$-\frac{1}{\rho} (c(\mathbf{x}_k) + \omega \boldsymbol{\xi}(\mathbf{x}_k, \boldsymbol{\lambda}_{k-1})) = -\frac{1}{\rho + \omega} (c(\mathbf{x}_k) + \omega \boldsymbol{\lambda}_{k-1}).$$

APPENDIX II

FASTER LOCAL CONVERGENCE THAN ALM

Theorem 2 (Local Convergence): Consider (ECP) and (QPP), the latter for a given $\omega > 0$, locally around the same initial iterate \mathbf{x}_k . Let $\nabla_{\mathbf{x}} f, \nabla_{\mathbf{x}} c, \nabla_{\mathbf{x}\mathbf{x}}^2 \mathcal{L}$ be locally Lipschitz-continuous and let both problems have a locally unique primal-dual solution $\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\eta}$ in that region, which satisfies strict complementarity. Apply ALM and MALM with fixed penalty parameter ρ to solve either problem, each starting from \mathbf{x}_k . If both methods converge and if \mathbf{x}_k is sufficiently close to the local minimizer of (QPP), then the linear rates of convergence of MALM and ALM satisfy the relation $C_{\text{MALM}} = \frac{\rho}{\rho + \omega} \cdot C_{\text{ALM}} < C_{\text{ALM}}$.

Proof: We use the Taylor series

$$\begin{aligned} & \nabla \Psi_k(\mathbf{x}_k, \boldsymbol{\lambda}_k) \\ &= \mathbf{H} \mathbf{x}_k + \mathbf{g} - \frac{1}{\omega + \rho} \mathbf{J}^\top (\rho \boldsymbol{\lambda}_{k-1} + \mathbf{c}) + R_L(\mathbf{x}_k, \boldsymbol{\lambda}_{k-1}) \end{aligned}$$

with $\mathbf{J}^\top := \nabla c(\mathbf{x}_\infty)$, $\mathbf{H} := \nabla_{\mathbf{x}\mathbf{x}}^2 \mathcal{L}(\mathbf{x}_\infty, \boldsymbol{\lambda}_\infty) + \frac{1}{\omega + \rho} \mathbf{J}^\top \mathbf{J}$, $\mathbf{c} := \mathbf{J} \mathbf{x}_\infty - c(\mathbf{x}_\infty)$ and $\mathbf{g} := \nabla f(\mathbf{x}_\infty)$ has the Lagrange

remainder $\|R_L(\mathbf{x}_k, \boldsymbol{\lambda}_k)\|_2 \leq \frac{L}{\rho+\omega}(\|\mathbf{x}_k - \mathbf{x}_\infty\|_2 + \|\boldsymbol{\lambda}_{k-1} - \boldsymbol{\lambda}_\infty\|_2)^2$, where L is the Lipschitz constant.

We now first consider the case where $p = 0$, i.e. where there are no inequality constraints. Since \mathbf{x}_k is convergent by requirement, \mathbf{H} must be positive semi-definite and, if \mathbf{x}_∞ is locally unique, \mathbf{H} must be positive definite. Clearly, local convergence to a unique point depends quantitatively on uniqueness, hence we imply $\lambda_{\min}(\mathbf{H}) \geq \mu > 0$. For the induced 2-norm it follows that $\|\mathbf{H}^{-1}\|_2 \leq \mu$, hence

$$\begin{aligned} & \left\| \mathbf{x}_k - \mathbf{H}^{-1} \left(\mathbf{g} - \frac{1}{\omega + \rho} \mathbf{J}^\top (\rho \boldsymbol{\lambda}_{k-1} + \mathbf{c}) \right) \right\|_2 \\ & \leq \frac{L}{\mu(\rho + \omega)} \|\boldsymbol{\lambda}_{k-1} - \boldsymbol{\lambda}_\infty\|_2^2. \end{aligned}$$

Inserting the estimate for \mathbf{x}_k into line 5 in Algorithm 1 gives a formula for $\boldsymbol{\lambda}_k$ that only depends on $\boldsymbol{\lambda}_{k-1}$:

$$\boldsymbol{\lambda}_k = \mathbf{M} \cdot \boldsymbol{\lambda}_{k-1} + \mathbf{f} + R_\lambda(\boldsymbol{\lambda}_k) \quad (11)$$

with $\mathbf{M} \in \mathbb{R}^{m \times m}$ below, some $\mathbf{f} \in \mathbb{R}^m$, and $\|R_\lambda(\boldsymbol{\lambda}_k)\|_2 \leq \frac{1}{\mu} \left(\frac{L}{\rho+\omega} \right)^2 \|\boldsymbol{\lambda}_k - \boldsymbol{\lambda}_\infty\|_2^2$. Rearranging reveals

$$\mathbf{M} = \frac{\rho}{\omega + \rho} \left(\mathbf{I} - \frac{1}{\omega + \rho} \mathbf{J} \mathbf{H}^{-1} \mathbf{J}^\top \right).$$

Since Theorem 1 asserts convergence of $\boldsymbol{\lambda}_k$, the second order terms become negligible compared to the first-order terms and can hence be ignored in the limit. Then, (11) is a Banach iteration. Thus, in the limit, the rate of convergence for $\boldsymbol{\lambda}_k$ is linear with contraction $\|\mathbf{M}\|_2 < 1$. The analysis holds regardless of whether $\omega = 0$ or > 0 .

We see that in the limit MALM converges faster than ALM because $\frac{\rho}{\omega+\rho} < 1$ when $\omega > 0$, whereas $\frac{\rho}{\omega+\rho} = 1$ when $\omega = 0$. Hence, MALM yields a stronger contraction for the errors per iteration than ALM. This is in particular an advantage in cases where ALM would converge slowly. For instance, choosing $\rho = 10\omega$ guarantees convergence in the limit with at least a rate of contraction of $\frac{\rho}{\omega+\rho} < 0.91$.

From the above, when dropping the Lagrange remainder terms, we can identify the local rate of convergence by that of the following quadratic model iteration: 1) Solve

$$\min_x \frac{1}{2} \mathbf{x}^\top \mathbf{H} \mathbf{x} + \left(\mathbf{g} + \frac{1}{\omega + \rho} \mathbf{J}^\top (\mathbf{J} \mathbf{x}_{k-1} - \mathbf{c} - \omega \boldsymbol{\lambda}_{k-1}) \right)^\top \mathbf{x}.$$

2) Update $\boldsymbol{\lambda}_k := \boldsymbol{\lambda}_{k-1} - \frac{1}{\omega+\rho} (\mathbf{J} \mathbf{x}_k - \mathbf{c} - \omega \boldsymbol{\lambda}_{k-1})$.

We discuss the case when $p > 0$, i.e. when inequality constraints are present. We use our assumption on strict complementarity, i.e. $i \in \mathcal{A} \Leftrightarrow \boldsymbol{\eta}_i > \beta$ for some real $\beta > 0$. Since \mathbf{x}_k converges by requirement, $\boldsymbol{\lambda}_{k-1}$ converges and thus also $\nabla \Psi_k(\mathbf{x}_k)$ converges. Hence, $\boldsymbol{\eta}_k$ must converge in order to yield $\nabla \Psi_k(\mathbf{x}_k) - \mathbf{A}^\top \boldsymbol{\eta}_k = \mathbf{0}$. Once $\boldsymbol{\eta}_k$ changes less than β at some finite $k_0 \in \mathbb{N}$, the active set \mathcal{A}_k will remain unchanged \mathcal{A}_∞ for all subsequent iterations $k \geq k_0$. We use $\mathbf{A}_\infty := \mathbf{A}_{\mathcal{A}_\infty}$, $\mathbf{b}_\infty := \mathbf{b}_{\mathcal{A}_\infty}$.

Given the above intermezzo, the appropriate model iteration in the limit becomes obvious: 1) Solve

$$\min_x \frac{1}{2} \mathbf{x}^\top \mathbf{H} \mathbf{x} + \left(\mathbf{g} + \frac{1}{\omega + \rho} \mathbf{J}^\top (\mathbf{J} \mathbf{x}_{k-1} - \mathbf{c} - \omega \boldsymbol{\lambda}_{k-1}) \right)^\top \mathbf{x}$$

s.t. $\mathbf{A}_\infty \mathbf{x} = \mathbf{b}_\infty$.

2) Update $\boldsymbol{\lambda}_k := \boldsymbol{\lambda}_{k-1} - \frac{1}{\omega+\rho} (\mathbf{J} \mathbf{x}_k - \mathbf{c} - \omega \boldsymbol{\lambda}_{k-1})$.

This is just a projection of the iteration above. Thus, we can project the iteration for \mathbf{x}_k onto the nullspace of \mathbf{A}_∞ , identifying $\mathbf{x}_k = \mathbf{x}_r + \mathbf{N} \tilde{\mathbf{x}}_k \forall k \geq k_0$, where $\mathbf{x}_r \in \mathbb{R}^n$ has active set \mathcal{A}_∞ , $\tilde{\mathbf{x}}_k \in \mathbb{R}^{n-\dim(\mathcal{A}_\infty)}$ and \mathbf{N} is a matrix of orthogonal columns that span the nullspace of \mathbf{A}_∞ . Defining $\tilde{\mathbf{H}} := \mathbf{N}^\top \mathbf{H} \mathbf{N}$, $\tilde{\mathbf{J}} := \mathbf{J} \mathbf{N}$, and $\tilde{\mathbf{g}}, \tilde{\mathbf{c}}$ appropriately, we arrive at the former unconstrained quadratic model iteration form, but with $\mathbf{H}, \mathbf{g}, \mathbf{J}, \mathbf{c}, \mathbf{x}_k$ replaced by the tilded quantities. Accordingly, the Banach iteration matrix \mathbf{M} is replaced with the matrix

$$\tilde{\mathbf{M}} = \frac{\rho}{\omega + \rho} \left(\mathbf{I} - \frac{1}{\omega + \rho} \tilde{\mathbf{J}} \tilde{\mathbf{H}}^{-1} \tilde{\mathbf{J}}^\top \right).$$

The resulting contraction matrix $\tilde{\mathbf{M}}$ for the Banach iteration of the inequality constrained case has a factor $\frac{\rho}{\omega+\rho}$ in front, just like for the case when $p = 0$. Thus, for $\rho > 0$ the method converges locally faster. ■

REFERENCES

- [1] A. V. Balakrishnan. On a new computing technique in optimal control. *SIAM J. Control*, 6:149–173, 1968.
- [2] J. T. Betts. *Practical Methods for Optimal Control and Estimation Using Nonlinear Programming*. Cambridge University Press, New York, NY, USA, 2nd edition, 2010.
- [3] A. R. Conn, N. Gould, A. Sartenaer, and Ph. L. Toint. Convergence properties of an augmented Lagrangian algorithm for optimization with a combination of general equality and linear constraints. *SIAM J. Optim.*, 6(3):674–703, 1996.
- [4] A. R. Conn, N. I. M. Gould, and Ph. L. Toint. *LANCELOT*, volume 17 of *Springer Series in Computational Mathematics*. Springer-Verlag, Berlin, 1992. A Fortran package for large-scale nonlinear optimization (release A).
- [5] Philip E. Gill, Walter Murray, and Michael A. Saunders. SNOPT: An SQP algorithm for large-scale constrained optimization. *SIAM Rev.*, 47:99–131, 2005.
- [6] W. W. Hager. Multiplier methods for nonlinear optimal control. *SIAM J. Numer. Anal.*, 27(4):1061–1080, 1990.
- [7] M. R. Hestenes. Multiplier and gradient methods. *J. Optim. Theory Appl.*, 4:303–320, 1969.
- [8] M. P. Neuenhofen and E. C. Kerrigan. Dynamic optimization with convergence guarantees. *arXiv:1810.04059*, 2018.
- [9] M. P. Neuenhofen and E. C. Kerrigan. A direct method for solving integral penalty transcriptions of optimal control problems. *Proceedings of the IEEE Conference on Decision and Control 2020 (forthcoming)*. *arXiv:2009.06222*, 2020.
- [10] M. P. Neuenhofen and E. C. Kerrigan. An integral penalty-barrier direct transcription method for optimal control. *Proceedings of the IEEE Conference on Decision and Control 2020 (forthcoming)*. *arXiv:2009.06217*, 2020.
- [11] J. Nocedal and S. J. Wright. *Numerical optimization*. Springer Series in Operations Research and Financial Engineering. Springer, New York, second edition, 2006.
- [12] M. J. D. Powell. A method for nonlinear constraints in minimization problems. In *Optimization (Sympos., Univ. Keele, Keele, 1968)*, pages 283–298. Academic Press, London, 1969.
- [13] R. T. Rockafellar. The multiplier method of Hestenes and Powell applied to convex programming. *J. Optim. Theory Appl.*, 12:555–562, 1973.

- [14] M.H.B.M. Shariff and J.R. Dormand. A modified augmented Lagrangian method for a class of constrained problems. *Journal of Computational and Applied Mathematics*, 151(2):257 – 270, 2003.
- [15] Andreas Wächter and Lorenz T. Biegler. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Math. Program.*, 106(1, Ser. A):25–57, 2006.