Preprint ANL/MCS-P1015-1202, December, 2002<br>Mathematics and Computer Science Division<br>Argonne National Laboratory

# A GLOBALLY CONVERGENT LCL METHOD FOR NONLINEAR OPTIMIZATION 

MICHAEL P. FRIEDLANDER* AND MICHAEL A. SAUNDERS ${ }^{\dagger}$


#### Abstract

For optimization problems with nonlinear constraints, linearly constrained Lagrangian (LCL) methods sequentially minimize a Lagrangian function subject to linearized constraints. These methods converge rapidly near a solution but may not be reliable from arbitrary starting points. The well known example MINOS has proven effective on many large problems. Its success motivates us to propose a globally convergent variant. Our stabilized LCL method possesses two important properties: the subproblems are always feasible, and they may be solved inexactly. These features are present in MINOS only as heuristics.

The new algorithm has been implemented in MATLAB, with the option to use either the MINOS or SNOPT Fortran codes to solve the linearly constrained subproblems. Only first derivatives are required. We present numerical results on a nonlinear subset of the COPS, CUTE, and HS testproblem sets, which include many large examples. The results demonstrate the robustness and efficiency of the stabilized LCL procedure.


Key words. large-scale optimization, nonlinear programming, nonlinear inequality constraints, augmented Lagrangian

AMS subject classifications. 49M37, 65K05, 90C30

1. Introduction. For optimization problems with nonlinear constraints, linearly constrained Lagrangian (LCL) methods sequentially minimize a Lagrangian function subject to linearized constraints. As currently defined, these methods converge rapidly near a solution but may not be reliable from arbitrary starting points. The well known example MINOS [31] has proven effective on many large and small problems, especially within the GAMS [8] and AMPL [18] environments, and is widely used in industry and academia. Its success motivates us to propose a globally convergent variant of the LCL method.

Our stabilized LCL algorithm solves a sequence of linearly constrained subproblems. Each subproblem minimizes an augmented Lagrangian function within a linear manifold that describes a current approximation to the nonlinear constraints. This manifold is nominally a linearization of the constraint space but may be a relaxed (i.e., larger) space at any stage, particularly during early iterations. Few conditions are imposed on the nature of the subproblem solutions; consequently, the subproblems may be solved with any of a variety of optimization routines for linearly constrained problems, providing much flexibility.

The stabilized LCL method possesses two important properties: the subproblems are always feasible, and they may be solved inexactly. These features are present in MINOS only as heuristics. The method may be regarded as a generalization of sequential augmented Lagrangian methods (see, for example, $[23,1,17]$ ). The theory

[^0]we develop provides a framework that unifies Robinson's LCL method [35] with the bound-constrained Lagrangian (BCL) method used, for example, by LANCELOT [11]. In the context of our theory, the proposed algorithm is actually a continuum of methods, with LCL and BCL methods at opposite ends of a spectrum. The stabilized LCL algorithm exploits this connection between BCL and LCL methods, preserving the fast local convergence properties of LCL methods while inheriting the global convergence properties of BCL methods. This connection is explored in more detail by Friedlander [19].

Our focus is on large-scale problems. We implemented the stabilized LCL method using the reduced-gradient part of MINOS [30] and the sequential quadratic programming code SNOPT [22] to solve the linearly constrained subproblems. These solvers are most efficient on problems with few degrees of freedom. Also, they use only first derivatives, and consequently our implementation requires only first derivatives. We discuss how the stabilized LCL method might be used with first- or second-derivative linearly constrained solvers.
1.1. The optimization problem. The proposed method solves nonlinearly constrained optimization problems of the form

where $f: \mathbb{R}^{n} \mapsto \mathbb{R}$ is a linear or nonlinear objective function, $c: \mathbb{R}^{n} \mapsto \mathbb{R}^{m}$ is a vector of nonlinear constraint functions, $A$ is a matrix, and $l$ and $u$ are vectors of bounds. We assume that $A$ and the derivatives of $c$ are sparse and that the problem (NP) is feasible. We recognize that not all optimization problems are feasible. This possibility is addressed in $\S 3.3$, where we explain how the proposed algorithm reveals an infeasible optimization problem and discuss properties of the points to which it converges.

One of the strengths of our method is that it does not explicitly require secondorder information. However, the fast convergence rate of the algorithm relies on sufficient smoothness of the nonlinear functions, indicated by the existence of second derivatives. We make that assumption:

Assumption 1.1. The functions $f$ and $c$ are twice continuously differentiable on an open neighborhood containing the region

$$
l \leq\binom{ x}{A x} \leq u
$$

Note that second derivatives could be used if they were available, thus accelerating the solutions of the subproblems and changing the properties of the solutions obtained by the algorithm. We discuss this further in $\S 3.4$.
1.2. The LCL approach. The acronym LCL is new. Methods belonging to this class typically have been referred to in the optimization literature as sequential linearized constraint (SLC) methods (cf. [23, 32]). The term SLC was chosen for compatibility with the terms sequential quadratic programming (SQP) and sequential linear programming (SLP). Those methods also sequentially linearize the constraints. The term linearly constrained Lagrangian, however, emphasizes that the Lagrangian
itself, and not an approximation, is used in the subproblems. Moreover, there is a useful relationship (which we exploit) between LCL and BCL methods, and this is hinted at by the nomenclature.

The first LCL methods were proposed independently in 1972. Robinson [35] and Rosen and Kreuser [38] describe similar algorithms based on minimizing a sequence of Lagrangian functions subject to linearized constraints. Robinson is able to prove that, under suitable conditions, the sequence of subproblem solutions converges quadratically to a solution of (NP). A strength of this method is that efficient large-scale methods exist for the solution of the linearly constrained subproblems formed at each iteration. Any suitable example of these subproblem solvers may be called as a black box.
1.3. Other work on stabilizing LCL methods. Other approaches to stabilizing LCL algorithms include two-phase methods proposed by Rosen [37] and Van Der Hoek [40]. In these approaches, a Phase 1 problem is formed by moving the nonlinear constraints into the objective by means of a quadratic penalty function. The solution of the Phase 1 problem is used to initialize Robinson's method (Phase 2). With a sufficiently large penalty parameter, the Phase 1 solution will yield a starting point that allows Robinson's method to converge quickly to a solution. These twophase methods choose the penalty parameter arbitrarily, however, and do not deal methodically with infeasible linearizations.

In 1981, Best et al. [2] describe a variant of the two-phase method whereby the Phase 1 penalty parameter is gradually increased by repeated return to the Phase 1 problem if the Phase 2 iterations are not converging. This two-phase method differs further from Rosen's and Van Der Hoek's methods in that the Phase 2 iterations involve only those equality constraints identified as active by the Phase 1 problem. The authors are able to retain local quadratic convergence of the Phase 2 LCL iterations while proving global convergence to a stationary point. A drawback of their method is that it requires a fourth-order penalty term to ensure continuous second derivatives of the penalty objective. This requirement may introduce significant numerical difficulty for the solution of the Phase 1 problem (though probably a quadratic-penalty term would suffice in practice).

Both two-phase methods share the disadvantage that the Phase 1 penalty problems need to be optimized over a larger subspace than the subsequent LCL phase. We seek a method that retains the linearized constraints as part of the subproblem, in order to keep the number of degrees of freedom small; and, as in Robinson's 1972 method, we allow the subproblem to determine the final set of active constraints.
1.4. The generic problem. For the theoretical development of a stabilized LCL method, we consider a simplified, generic formulation of (NP) and take the optimization problem to be

| (GNP) | $\operatorname{minimize}_{x \in \mathbb{R}^{n}} \quad f(x)$ |
| :--- | :--- |
| subject to $\quad c(x)$ | $=0$ |
|  | $x \geq 0$, |

where $c: \mathbb{R}^{n} \mapsto \mathbb{R}^{m}$. Section 4 returns to the formulation (NP) in its discussion of the implementation of the stabilized LCL method.

We define the augmented Lagrangian function corresponding to (GNP) as

$$
\begin{equation*}
\mathcal{L}(x, y, \rho)=f(x)-y^{T} c(x)+\frac{1}{2} \rho\|c(x)\|_{2}^{2} \tag{1.1}
\end{equation*}
$$

where $x$, the $m$-vector $y$, and the scalar $\rho$ are independent variables. Let $g(x)$ denote the gradient of the objective function $f(x)$, and let $J(x)$ denote the Jacobian matrix of the constraint vector $c(x)$. Denote by $H(x)$ and $H_{i}(x)$ the Hessian matrices of $f(x)$ and $[c(x)]_{i}$, respectively, where $[\cdot]_{i}$ refers to the $i$ th component of a vector. Define

$$
\begin{equation*}
\widehat{y}(x, y, \rho)=y-\rho c(x) \tag{1.2}
\end{equation*}
$$

The derivatives of $\mathcal{L}$ with respect to $x$ may be written as follows:

$$
\begin{align*}
\nabla_{x} \mathcal{L}(x, y, \rho) & =g(x)-J(x)^{T} \widehat{y}(x, y, \rho)  \tag{1.3}\\
\nabla_{x x}^{2} \mathcal{L}(x, y, \rho) & =H(x)-\sum_{i=1}^{m}[\widehat{y}(x, y, \rho)]_{i} H_{i}(x)+\rho J(x)^{T} J(x) \tag{1.4}
\end{align*}
$$

We assume that problem (GNP) is feasible and has at least one point $\left(x_{*}, y_{*}, z_{*}\right)$ that satisfies the first-order Karush-Kuhn-Tucker (KKT) optimality conditions.

Definition 1.2 (First-Order Optimality Conditions). A triple $\left(x_{*}, y_{*}, z_{*}\right)$ is a first-order KKT point for (GNP) if for any $\rho \geq 0$ all of the following hold:

$$
\begin{align*}
c\left(x_{*}\right) & =0  \tag{1.5a}\\
\nabla_{x} \mathcal{L}\left(x_{*}, y_{*}, \rho\right) & =z_{*}  \tag{1.5b}\\
\min \left(x_{*}, z_{*}\right) & =0 . \tag{1.5c}
\end{align*}
$$

Note that (1.5c) implies

$$
\begin{align*}
& x_{*} \geq 0  \tag{1.6a}\\
& z_{*} \geq 0, \tag{1.6b}
\end{align*}
$$

so that $x_{*}$ and $z_{*}$ must be primal and dual feasible, respectively.
Let $\eta_{*}>0$ and $\omega_{*}>0$ be specified as primal and dual convergence tolerances. We regard the point $(x, y, z)$ to be an acceptable solution of (GNP) if it satisfies (1.5) to within these tolerances. Specifically, we identify $(x, y, z)$ as an approximate solution of (GNP) if

$$
\begin{align*}
\|c(x)\| & \leq \eta_{*}  \tag{1.7a}\\
\nabla_{x} \mathcal{L}(x, y, \rho) & =z  \tag{1.7b}\\
\|\min (x, z)\|_{\infty} & \leq \omega_{*} . \tag{1.7c}
\end{align*}
$$

Note that (1.7c) relaxes the nonnegativity conditions (1.6) by the same tolerance $\omega_{*}$. In practice, we might choose to relax (1.6a) to $x \geq-\delta_{*} e$, for some $\delta_{*}>0$. However, we ignore this detail for now.

For theoretical purposes, we assume that strict complementarity and the secondorder sufficiency conditions hold at each $\left(x_{*}, y_{*}, z_{*}\right)$. We define these conditions as follows.

Definition 1.3 (Strict Complementarity). The point $\left(x_{*}, y_{*}, z_{*}\right)$ satisfies strict complementarity if it satisfies (1.5) and $\max \left(x_{*}, z_{*}\right)>0$.

Definition 1.4 (Second-Order Sufficiency). The point $\left(x_{*}, y_{*}, z_{*}\right)$ satisfies the second-order sufficiency conditions for (GNP) if it satisfies (1.5) and strict complementarity and if for any $\rho \geq 0$,

$$
\begin{equation*}
p^{T} \nabla_{x x}^{2} \mathcal{L}\left(x_{*}, y_{*}, \rho\right) p>0 \tag{1.8}
\end{equation*}
$$

for all $p \neq 0$ satisfying

$$
\begin{equation*}
J\left(x_{*}\right) p=0 \text { and }[p]_{j}=0 \text { for all } j \text { such that }\left[x_{*}\right]_{j}=0 \tag{1.9}
\end{equation*}
$$

(and $\left[z_{*}\right]_{j}>0$ ).
AsSumption 1.5. The point $\left(x_{*}, y_{*}, z_{*}\right)$ satisfies the second-order sufficiency conditions for (GNP).
1.5. The canonical LCL method. The software package MINOS solves the nonlinearly constrained problem (GNP) by minimizing a sequence of augmented Lagrangian functions subject to linearized constraints. Define the constraint linearization at the point $x_{k}$ as

$$
\bar{c}_{k}(x)=c\left(x_{k}\right)+J\left(x_{k}\right)\left(x-x_{k}\right)
$$

Algorithm 1 outlines what we regard to be a canonical LCL method. It forms the basis for the MINOS algorithm and is based on solving the linearly constrained subproblems

$$
\begin{array}{|ccc}
\hline\left(\mathrm{LC}_{k}\right) & \underset{x}{\operatorname{minimize}} & \mathcal{L}\left(x, y_{k}, \rho_{k}\right) \\
& \text { subject to } & \bar{c}_{k}(x)=0 \\
& x \geq 0
\end{array}
$$

which are parameterized by the latest estimates $x_{k}$ and $y_{k}$, and a fixed penalty parameter $\rho_{k} \equiv \bar{\rho}$ (which may be set to zero). The linear constraints $\bar{c}_{k}(x)=0$ are the linearization of $c$ at the point $x_{k}$.

Empirically, a positive penalty parameter $\bar{\rho}$ has proven a helpful addition to Robinson's method, but for other problems it has been ineffective. A theoretical understanding of when and how to modify the penalty term has been lacking.
1.6. Notation. The symbol $x_{*}$ is used in two senses: as a limit point of the sequence $\left\{x_{k}\right\}$, and as the primal solution of (GNP). We distinguish between the two cases when the context is not clear. Denote by $\widehat{g}(x)$ the vector of components of $g(x)$ corresponding to inactive bounds at $x_{*}$, so that if $\mathcal{I}=\left\{i \in 1, \ldots, n \mid\left[x_{*}\right]_{i}>0\right\}$, $\widehat{g}(x)=[g(x)]_{\mathcal{I}}$ (where $[\cdot]_{\mathcal{I}}$ is a shorthand notation for a subvector formed from the indices in $\mathcal{I})$. Similarly, let $\widehat{J}(x)$ denote the corresponding columns of the Jacobian matrix.

Unless otherwise specified, the function $\|x\|$ represents the Euclidean norm of the vector $x$. When the arguments are vectors, define the function $\min (\cdot, \cdot)$ componentwise. The following notation is used throughout:
$(x, y, z) \quad$ primal variables, dual variables, and reduced costs for (GNP),
$\left(x_{*}, y_{*}, z_{*}\right) \quad$ optimal variables for (GNP),
$\left(x_{k}, y_{k}, z_{k}\right) \quad$ the $k$ th estimate of $\left(x_{*}, y_{*}, z_{*}\right)$,
$\left(x_{k}^{*}, \Delta y_{k}^{*}, z_{k}^{*}\right) \quad$ solution of the $k$ th subproblem,
$y_{k}^{*} \quad y_{k}+\Delta y_{k}^{*}$; an updated multiplier estimate,
$f_{k}, g_{k}, c_{k}, J_{k} \quad$ functions and gradients evaluated at $x_{k}$,
$f_{*}, g_{*}, c_{*}, J_{*} \quad$ functions and gradients evaluated at $x_{*}$.
The augmented Lagrangian function is particularly important for our analysis. We often use the shorthand notation

$$
\begin{equation*}
\mathcal{L}_{k}(x) \equiv \mathcal{L}\left(x, y_{k}, \rho_{k}\right)=f(x)-y_{k}^{T} c(x)+\frac{1}{2} \rho_{k}\|c(x)\|^{2} \tag{1.10}
\end{equation*}
$$

```
Algorithm 1: Canonical LCL
    Input: \(x_{0}, y_{0}, z_{0}\)
    Output: \(x_{*}, y_{*}, z_{*}\)
    [Initialize parameters]
        Set the penalty parameter \(\bar{\rho} \geq 0\). Set positive convergence tolerances
        \(\omega_{*}, \eta_{*} \ll 1\);
    \(k \leftarrow 0\);
    converged \(\leftarrow\) false;
    repeat
        [Solve the LC subproblem]
            Solve \(\left(\mathrm{LC}_{k}\right)\) to obtain a point \(\left(x_{k}^{*}, \Delta y_{k}^{*}, z_{k}^{*}\right)\). If there is more than one
            such point, choose \(\left(x_{k}^{*}, \Delta y_{k}^{*}, z_{k}^{*}\right)\) closest in norm to ( \(x_{k}, 0, z_{k}\) );
            \(y_{k}^{*} \leftarrow y_{k}+\Delta y_{k}^{*}\);
        [Update solution estimates]
            \(x_{k+1} \leftarrow x_{k}^{*}, y_{k+1} \leftarrow y_{k}^{*}, z_{k+1} \leftarrow z_{k}^{*} ;\)
        [Test convergence]
            if \(\left(x_{k+1}, y_{k+1}, z_{k+1}\right)\) satisfies (1.7) then converged \(\leftarrow\) true;
        \(\rho_{k} \leftarrow \bar{\rho} ; \quad\left[\right.\) keep \(\rho_{k}\) fixed]
        \(k \leftarrow k+1 ;\)
    until converged;
    \(x_{*} \leftarrow x_{k}, y_{*} \leftarrow y_{k}, z_{*} \leftarrow z_{k} ;\)
    return \(x_{*}, y_{*}, z_{*}\);
```

when $y_{k}$ and $\rho_{k}$ are fixed. The algorithms we discuss are structured around major and minor iterations. Each major iteration solves a subproblem and generates an element of the sequence $\left\{\left(x_{k}, y_{k}, z_{k}\right)\right\}$. Under certain (desirable) circumstances, this sequence converges to a solution $\left(x_{*}, y_{*}, z_{*}\right)$. For each major iteration $k$, there is a corresponding set of minor iterations converging to $\left(x_{k}^{*}, \Delta y_{k}^{*}, z_{k}^{*}\right)$, the solution of the current subproblem. In our development and analysis of a stabilized LCL method, we are primarily concerned with the "outer"-level algorithm. Unless stated otherwise, "iterations" refers to major iterations.
2. An Elastic LC Subproblem. The original LCL method introduced by Robinson [35] sets $\bar{\rho}=0$ in Algorithm 1. A positive penalty parameter could be used (as it is in MINOS [31]) and may help convergence from difficult starting points.

We recognize two particular causes of failure for the LCL method:

- The linearized constraints may be infeasible, so that the LCL iterations are not defined;
- A near-singular Jacobian $J_{k}$ (we only assume nonsingularity of the Jacobian at limit points-cf. Assumption 3.2) might lead to an arbitrarily large value of $\left\|x_{k}^{*}-x_{k}\right\|$ regardless of the values of $y_{k}$ and $\rho_{k}$ in the subproblem objective.
To remedy both deficiencies we modify the linearized constraints used by the LCL method, allowing some degree of flexibility in their satisfaction. We introduce a set of nonnegative elastic variables, $v$ and $w$, into the constraints and introduce a penalty
on these variables into the subproblem objective. Hence, we define the subproblem as

$$
\begin{array}{|cc|}
\hline\left(\mathrm{ELC}_{k}\right) & \operatorname{minimize}_{x, v, w} \\
\mathcal{L}_{k}(x)+\sigma_{k} e^{T}(v+w) \\
\text { subject to } & \bar{c}_{k}(x)+v-w=0 \\
& x, v, w \geq 0
\end{array}
$$

where $e$ is a vector of ones. This elastic subproblem is always feasible. Its solution yields a 5 -tuple $\left(x_{k}^{*}, \Delta y_{k}^{*}, z_{k}^{*}, v_{k}^{*}, w_{k}^{*}\right)$ that satisfies the first-order KKT conditions

$$
\begin{align*}
v, w & \geq 0  \tag{2.1a}\\
\bar{c}_{k}(x)+v-w & =0  \tag{2.1b}\\
\nabla \mathcal{L}_{k}(x)-J_{k}^{T} \Delta y & =z  \tag{2.1c}\\
\min (x, z) & =0  \tag{2.1d}\\
\|\Delta y\|_{\infty} & \leq \sigma_{k} . \tag{2.1e}
\end{align*}
$$

Note that $\nabla \mathcal{L}_{k}(x)$ involves $y_{k}$ and $\rho_{k}$.
The term $\sigma_{k} e^{T}(v+w)$ is the $\ell_{1}$-penalty function, and together with the nonnegativity constraints $v, w \geq 0$ it is equivalent to a penalty on the one-norm of $(v-w)$. We find later that the bound (2.1e) is crucial for the global convergence analysis of the proposed method.

We note that the elastic LC subproblem can be equivalently stated as

$$
\begin{array}{ll}
\left(\mathrm{ELC}_{k}^{\prime}\right) & \underset{x}{\operatorname{minimize}} \quad \mathcal{L}_{k}(x)+\sigma_{k}\left\|\bar{c}_{k}(x)\right\|_{1} \\
& \text { subject to } x \geq 0
\end{array}
$$

with solution $\left(x_{k}^{*}, z_{k}^{*}\right)$. This immediately reveals the stabilized LCL method's intimate connection with both the augmented Lagrangian function and the BCL method. Far from a solution, the $\ell_{1}$-penalty term $\sigma_{k}\left\|\bar{c}_{k}(x)\right\|_{1}$ gives the method an opportunity to deviate from the constraint linearizations. Near a solution, it keeps the iterates close to the linearizations. For values of $\sigma_{k}$ over a threshold value, the linearized constraints are satisfied exactly, as required by the LCL method.
2.1. The $\ell_{1}$-penalty function. For any given subproblem of the stabilized LCL method, the penalty term $\sigma_{k} e^{T}(v+w)$ may or may not equal zero, indicating that the linearized constraints may not always be satisfied. In contrast, the MINOS or the canonical LCL subproblems must always satisfy the linearized constraints. Thus, the set of active linearized constraints of the stabilized LCL subproblem is always a subset (though not necessarily strict) of the the canonical LCL subproblem. Fletcher [16] makes the same observation in connection with his $\mathrm{S} \ell_{1} \mathrm{QP}$ method. The global convergence properties of the stabilized LCL method do not require independent constraint gradients or bounded multipliers for each subproblem (these are required only at limit points of the sequence generated by the algorithm).

Recovering the BCL subproblem. Set $\sigma_{k}=0$. Then $\left(\mathrm{ELC}_{k}\right)$ and $\left(\mathrm{ELC}_{k}^{\prime}\right)$ reduce to the equivalent bound-constrained minimization problem

| $\left(\mathrm{BC}_{k}\right)$ | $\underset{x}{\operatorname{minimize}} \mathcal{L}_{k}(x)$ |
| :--- | :--- |
|  | subject to $x \geq 0$, |

where the bounds on the variables $v$ and $w$ have been eliminated because they no longer appear in the objective. The subproblem $\left(\mathrm{BC}_{k}\right)$ is used by the BCL method (see, for example, Hestenes [24], Powell [34], Bertsekas [1], and Conn et al. [9, 10]).

Recovering the LCL subproblem. The $\ell_{1}$-penalty function is exact. If the linearization is feasible and $\sigma_{k}$ is larger than a certain threshold, $v$ and $w$ will be zero and the minimizers of the elastic problem $\left(\mathrm{ELC}_{k}\right)$ will coincide with the minimizers of the inelastic problem $\left(\mathrm{LC}_{k}\right)$. Exact penalty functions have been studied by $[1,16,26]$ among others. See the book by Conn et al. [12] for a more recent discussion.

We are particularly interested in this feature when the iterates generated by the stabilized LCL algorithm are approaching a solution $\left(x_{*}, y_{*}, z_{*}\right)$. Recovering the canonical LCL subproblem as the iterates approach a solution ensures that the stabilized LCL method inherits LCL's fast local convergence properties.

To prove that the condition

$$
\begin{equation*}
\left\|\Delta y_{k}^{*}\right\|_{\infty}<\sigma_{k} \tag{2.2}
\end{equation*}
$$

is sufficient to force the elastic variables to zero, we require two conditions: (i) the inelastic subproblem $\left(\mathrm{LC}_{k}\right)$ must satisfy the second-order sufficiency conditions at a solution $x_{k}^{*}$; and (ii) $x_{k}^{*}$ must be a regular point. Assumptions 1.5 and 3.2 guarantee that both these conditions are met. For $x_{k}^{*}$ close to $x_{*}$, Assumption 1.5 guarantees that $\left(\mathrm{LC}_{k}\right)$ satisfies the second-order conditions. Assumption 3.2 guarantees the regularity of $x_{k}^{*}$ when it is near $x_{*}$. Lemma 2.1 establishes the threshold value of $\sigma_{k}$.

Lemma 2.1. Suppose that $\left(x_{k}^{*}, \Delta y_{k}^{*}, z_{k}^{*}\right)$ satisfies the second-order sufficiency conditions for $\left(\mathrm{LC}_{k}\right)$. Then if (2.2) holds, $\left(x_{k}^{*}, \Delta y_{k}^{*}, z_{k}^{*}\right)$ also solves $\left(\mathrm{ELC}_{k}\right)$.

Proof. See Luenberger [26, p. 389]. $\square$
2.2. Early termination of the subproblems. Poor values of $x_{k}, y_{k}$, or $\rho_{k}$ may imply subproblems whose accurate solutions are far from a solution of (GNP). We therefore terminate subproblems early by relaxing (2.1d) and (2.1e) by an amount $\omega_{k}$. However, we enforce the nonnegativity condition on $x$ (implied by (2.1d)):

$$
\begin{align*}
x, v, w & \geq 0  \tag{2.3a}\\
\bar{c}_{k}(x)+v-w & =0  \tag{2.3b}\\
\nabla \mathcal{L}_{k}(x)-J_{k}^{T} \Delta y & =z  \tag{2.3c}\\
\|\min (x, z)\|_{\infty} & \leq \omega_{k}  \tag{2.3d}\\
\|\Delta y\|_{\infty} & \leq \sigma_{k}+\omega_{k} . \tag{2.3e}
\end{align*}
$$

Each subproblem is required to return a solution satisfying the linear and nonnegativity constraints and, as discussed in connection with (1.6a), in practice (2.3a) and/or (2.3b) would be relaxed by a fixed tolerance $\delta$.
3. The Stabilized LCL Algorithm. Algorithm 2 outlines the stabilized LCL method. Its structure closely parallels the BCL algorithm described in [10]. Based on the current primal infeasibility, each iteration of the algorithm is regarded as either "successful" or "unsuccessful." In the "successful" case, the solution estimates are updated by using information from the current subproblem solution. If the iteration is "unsuccessful," the subproblem solutions are discarded, the current solution estimates are held fixed, and the penalty parameter $\rho_{k}$ is increased in an effort to reduce the primal infeasibility in the next iteration. In order for the linearized constraints not to continue interfering with the penalty parameter's ability to reduce the primal

```
Algorithm 2: Stabilized LCL.
    Input: \(x_{0}, y_{0}, z_{0}\)
    Output: \(x_{*}, y_{*}, z_{*}\)
    [Initialize parameters]
        Set \(\bar{\sigma}>\underline{\sigma}>0\). Set constants \(\tau_{\rho}, \tau_{\sigma}>1\). Set the initial penalty parameters
        \(\rho_{0}>1\) and \(\sigma_{0} \gg 1\). Set positive convergence tolerances \(\omega_{*}, \eta_{*} \ll 1\) and initial
        tolerances \(\omega_{0}>\omega_{*}\) and \(\eta_{0}>\eta_{*}\). Set constants \(\alpha, \beta>0\) with \(\alpha<1\);
    \(k \leftarrow 0 ;\)
    converged \(\leftarrow\) false;
    repeat
        Choose \(\omega_{k} \geq \omega_{*}\) such that \(\lim _{k \rightarrow \infty} \omega_{k}=\omega_{*} ;\)
        [Solve the LC subproblem]
            Solve \(\left(\mathrm{ELC}_{k}\right)\) to obtain a point \(\left(x_{k}^{*}, \Delta y_{k}^{*}, z_{k}^{*}\right)\) satisfying (2.3). If there is
            more than one such point, compute the one closest in norm to ( \(x_{k}, 0, z_{k}\) );
            \(y_{k}^{*} \leftarrow y_{k}+\Delta y_{k}^{*} ;\)
        if \(\left\|c\left(x_{k}^{*}\right)\right\| \leq \max \left(\eta_{*}, \eta_{k}\right)\) then
            [Update solution estimates]
                \(x_{k+1} \leftarrow x_{k}^{*}\);
                \(y_{k+1} \leftarrow y_{k}^{*}-\rho_{k} c\left(x_{k}^{*}\right)\left(\equiv \widehat{y}\left(x_{k}^{*}, y_{k}^{*}, \rho_{k}\right)\right) ; \quad\left[\right.\) or \(\left.y_{k+1} \leftarrow y_{k}^{*}\right]\)
                \(z_{k+1} \leftarrow z_{k}^{*} ; \quad\left[\right.\) or \(\left.z_{k+1} \leftarrow g_{k+1}-J_{k+1}^{T} y_{k+1}\right]\)
            [Update penalty parameter and elastic weight]
                \(\rho_{k+1} \leftarrow \rho_{k} ; \quad\left[\right.\) keep \(\left.\rho_{k}\right]\)
                \(\sigma_{k+1} \leftarrow \max \left\{\underline{\sigma}, \min \left(\left\|\Delta y_{k}^{*}\right\|_{\infty}, \bar{\sigma}\right)\right\} ; \quad\left[\right.\) reset \(\left.\sigma_{k}\right]\)
            [Test convergence]
                if \(\left(x_{k+1}, y_{k+1}, z_{k+1}\right)\) satisfies (1.7) then converged \(\leftarrow\) true;
            \(\eta_{k+1} \leftarrow \eta_{k} / \rho_{k+1}^{\beta} ; \quad\) [decrease \(\eta_{k}\) ]
        else
            [Keep solution estimates]
                \(\left\llcorner x_{k+1} \leftarrow x_{k} ; y_{k+1} \leftarrow y_{k} ; z_{k+1} \leftarrow z_{k} ;\right.\)
            [Update penalty parameter and elastic weight]
                \(\rho_{k+1} \leftarrow \tau_{\rho} \rho_{k} ; \quad\) [increase \(\rho_{k}\) ]
                \(\sigma_{k+1} \leftarrow \sigma_{k} / \tau_{\sigma} ; \quad\) [decrease \(\sigma_{k}\) ]
            \(\eta_{k+1} \leftarrow \eta_{0} / \rho_{k+1}^{\alpha} ; \quad\) [may increase or decrease \(\eta_{k}\) ]
        \(k \leftarrow k+1 ;\)
    until converged;
    \(x_{*} \leftarrow x_{k} ; y_{*} \leftarrow y_{k} ; z_{*} \leftarrow z_{k} ;\)
    return \(x_{*}, y_{*}, z_{*}\);
```

infeasibility, the algorithm relaxes the linearizations by reducing the elastic penalty parameter $\sigma_{k}$.

The two salient features of this algorithm are that it is globally convergent and that it is asymptotically equivalent to the canonical LCL method. In $\S 3.1$ we demonstrate the global convergence properties of the algorithm by proving results analogous to Lemma 4.3 and Theorem 4.4 in [10]. In $\S 3.2$ we demonstrate that the algorithm eventually reduces to the canonical LCL method and hence inherits that method's asymptotic convergence properties.
3.1. Global convergence properties. We make the following assumptions.

Assumption 3.1. The sequence of iterates $\left\{x_{k}^{*}\right\}$ lies in the closed and bounded set $\mathcal{B} \subset \mathbb{R}^{n}$.

Assumption 3.2. The matrix $\widehat{J}\left(x_{*}\right)$ has full row rank at every limit point $x_{*}$ of the sequence $\left\{x_{k}^{*}\right\}$.

The first assumption guarantees that any sequence of iterates generated by the algorithm always has some convergent subsequence. The second assumption is commonly known as the linear independence constraint qualification (LICQ) (see, for example, Mangasarian [27], or for a more recent reference, Nocedal and Wright [32]).

Let $x_{*}$ be any limit point of the sequence $\left\{x_{k}^{*}\right\}$. At all points $x$ for which $\widehat{J}(x)$ has full row rank we define the least-squares multiplier estimate, $\widetilde{y}(x)$, as the solution of the linear least-squares problem

$$
\begin{equation*}
\underset{y}{\operatorname{minimize}}\left\|\widehat{g}(x)-\widehat{J}(x)^{T} y\right\|^{2} \tag{3.1}
\end{equation*}
$$

Note that the definitions of $\widehat{g}, \widehat{J}$, and hence $\widetilde{y}$ require a priori knowledge of the bounds active at $x_{*}$. We emphasize that $\widetilde{y}$ is used only as an analytical device and its computation is never required. Assumption 3.2 guarantees the uniqueness of $\widetilde{y}$ at every limit point of the sequence $\left\{x_{k}^{*}\right\}$.
3.1.1. Convergence of LC subproblem solutions. In this section we prove that the sequence of LC subproblem solutions generated by Algorithm 2 converges to a KKT point of (GNP).

We need the following lemma to bound the errors in the least-squares multiplier estimates relative to the error in $x_{k}$. The lemma simply demonstrates that $\widetilde{y}(x)$ is Lipschitz continuous in a neighborhood of $x_{*}$.

Lemma 3.3. Let $\left\{x_{k}\right\}, k \in \mathcal{K}$ be a subsequence converging to $x_{*}$ and suppose that Assumptions 1.1 and 3.2 hold. Then there exists a positive constant $\alpha$ such that $\left\|\widetilde{y}\left(x_{k}\right)-\widetilde{y}\left(x_{*}\right)\right\| \leq \alpha\left\|x_{k}-x_{*}\right\|$ for all $k \in \mathcal{K}$ sufficiently large.

Proof. See Lemmas 2.1 and 4.4 of [9].
To prove the global convergence properties of Algorithm 2, we first describe the properties of any limit point that the algorithm generates. We are not claiming (yet!) that the algorithm is globally convergent, only that if it does converge, then the set of limit points generated must satisfy some desirable properties. The following lemma is adapted from Lemma 4.4 of [9].

Lemma 3.4. Let $\left\{\omega_{k}\right\}$ and $\left\{\rho_{k}\right\}$ be sequences of positive scalars, where $\omega_{k} \rightarrow 0$. Let $\left\{x_{k}\right\}$ be any sequence of $n$-vectors and $\left\{y_{k}\right\}$ be any sequence of m-vectors. Let $\left\{\left(x_{k}^{*}, \Delta y_{k}^{*}, z_{k}^{*}\right)\right\}$ be a sequence of vectors satisfying (2.3a), (2.3c), and (2.3d). Let $x_{*}$ be any limit point of the sequence $\left\{x_{k}^{*}\right\}$, and let $\mathcal{K}$ be the infinite set of indices associated with that convergent subsequence. Suppose that Assumptions 1.1, 3.1, and 3.2 hold. Set $y_{k}^{*}=y_{k}+\Delta y_{k}^{*}, \widehat{y}_{k}=\widehat{y}\left(x_{k}^{*}, y_{k}^{*}, \rho_{k}\right)$, and $y_{*}=\widetilde{y}\left(x_{*}\right)$. The following properties then hold:

1. There are positive constants $\alpha_{1}, \alpha_{2}$, and $M$ such that

$$
\begin{align*}
\left\|\widehat{y}_{k}-y_{*}\right\| & \leq \beta_{1} \equiv \alpha_{1} \omega_{k}+M\left\|x_{k}^{*}-x_{k}\right\|\left\|y_{k}^{*}-y_{k}\right\|+\alpha_{2}\left\|x_{k}^{*}-x_{*}\right\|,  \tag{3.2}\\
\rho_{k}\left\|c\left(x_{k}^{*}\right)\right\| & \leq \beta_{2} \equiv \beta_{1}+\left\|y_{k}^{*}-y_{k}\right\|+\left\|y_{k}-y_{*}\right\| \tag{3.3}
\end{align*}
$$

for all $k \in \mathcal{K}$ sufficiently large.
2. As $k \in \mathcal{K}$ gets large, if $\left\|y_{k}^{*}-y_{k}\right\| \rightarrow 0$, or if $\left\|y_{k}^{*}-y_{k}\right\|$ is bounded and $\left\|x_{k}^{*}-x_{k}\right\| \rightarrow 0$, then

$$
\widehat{y}_{k} \rightarrow y_{*} \quad \text { and } \quad z_{k}^{*} \rightarrow z_{*} \stackrel{\text { def }}{=} \nabla_{x} \mathcal{L}\left(x_{*}, y_{*}, 0\right) .
$$

3. If, in addition, $c_{*}=0$, then $\left(x_{*}, y_{*}, z_{*}\right)$ is a first-order $K K$ point for (GNP). Proof. From the definition of $\widetilde{y}\left(x_{k}^{*}\right)$, the least-squares multiplier estimates,

$$
\begin{align*}
\left\|\widetilde{y}\left(x_{k}^{*}\right)-\widehat{y}_{k}\right\| & =\left\|\left(\widehat{J}\left(x_{k}^{*}\right) \widehat{J}\left(x_{k}^{*}\right)^{T}\right)^{-1} \widehat{J}\left(x_{k}^{*}\right) \widehat{g}\left(x_{k}^{*}\right)-\widehat{y}_{k}\right\| \\
& =\left\|\left(\widehat{J}\left(x_{k}^{*}\right) \widehat{J}\left(x_{k}^{*}\right)^{T}\right)^{-1} \widehat{J}\left(x_{k}^{*}\right)\left(\widehat{g}\left(x_{k}^{*}\right)-\widehat{J}\left(x_{k}^{*}\right)^{T} \widehat{y}_{k}\right)\right\|  \tag{3.4}\\
& \leq\left\|\left(\widehat{J}\left(x_{k}^{*}\right) \widehat{J}\left(x_{k}^{*}\right)^{T}\right)^{-1} \widehat{J}\left(x_{k}^{*}\right)\right\| \cdot\left\|\widehat{g}\left(x_{k}^{*}\right)-\widehat{J}\left(x_{k}^{*}\right)^{T} \widehat{y}_{k}\right\| .
\end{align*}
$$

By assumption, $\widehat{J}\left(x_{*}\right)$ has full row rank. Continuity of $J$ then implies that

$$
\left(\widehat{J}\left(x_{k}^{*}\right) \widehat{J}\left(x_{k}^{*}\right)^{T}\right)^{-1} \widehat{J}\left(x_{k}^{*}\right)
$$

exists for all $k \in \mathcal{K}$ large enough. Then there exists a positive scalar $\alpha_{1}$ such that

$$
\begin{equation*}
\left\|\left(\widehat{J}\left(x_{k}^{*}\right) \widehat{J}\left(x_{k}^{*}\right)^{T}\right)^{-1} \widehat{J}\left(x_{k}^{*}\right)\right\| \leq \frac{\alpha_{1}}{\sqrt{n}} \tag{3.5}
\end{equation*}
$$

where $n$ is the dimension of the vector $x$. Substituting (3.5) into (3.4),

$$
\begin{equation*}
\left\|\widetilde{y}\left(x_{k}^{*}\right)-\widehat{y}_{k}\right\| \leq \frac{\alpha_{1}}{\sqrt{n}}\left\|\widehat{g}\left(x_{k}^{*}\right)-\widehat{J}\left(x_{k}^{*}\right)^{T} \widehat{y}_{k}\right\| . \tag{3.6}
\end{equation*}
$$

We now show that $\left\|\widehat{g}\left(x_{k}^{*}\right)-\widehat{J}\left(x_{k}^{*}\right)^{T} \widehat{y}_{k}\right\|$ is bounded. By hypothesis, $\left(x_{k}^{*}, y_{k}^{*}, z_{k}^{*}\right)$ satisfies (2.3c). Together with (1.3),

$$
\begin{align*}
z_{k}^{*} & =\nabla \mathcal{L}_{k}\left(x_{k}^{*}\right)-J_{k}^{T} \Delta y_{k}^{*} \\
& =g\left(x_{k}^{*}\right)-J\left(x_{k}^{*}\right)^{T}\left(y_{k}-\rho_{k} c\left(x_{k}^{*}\right)\right)-J_{k}^{T} \Delta y_{k}^{*} \\
& =g\left(x_{k}^{*}\right)-J\left(x_{k}^{*}\right)^{T}\left(y_{k}+\Delta y_{k}^{*}-\rho_{k} c\left(x_{k}^{*}\right)\right)+\left(J\left(x_{k}^{*}\right)-J_{k}\right)^{T} \Delta y_{k}^{*}  \tag{3.7}\\
& =g\left(x_{k}^{*}\right)-J\left(x_{k}^{*}\right)^{T} \widehat{y}_{k}+\left(J\left(x_{k}^{*}\right)-J_{k}\right)^{T}\left(y_{k}^{*}-y_{k}\right),
\end{align*}
$$

where $y_{k}^{*} \stackrel{\text { def }}{=} y_{k}+\Delta y_{k}^{*}$ and $\widehat{y}_{k} \stackrel{\text { def }}{=} \widehat{y}\left(x_{k}^{*}, y_{k}^{*}, \rho_{k}\right)=y_{k}^{*}-\rho_{k} c\left(x_{k}^{*}\right)$. For $k \in \mathcal{K}$ large enough, $x_{k}^{*}$ is sufficiently close to $x_{*}$ so that

$$
\begin{equation*}
\left\|\left[z_{k}^{*}\right]_{\mathcal{I}}\right\| \leq\left\|\min \left(x_{k}^{*}, z_{k}^{*}\right)\right\| \tag{3.8}
\end{equation*}
$$

where $\mathcal{I}$ is the index set of inactive bounds at $x_{k}^{*}$, as defined in §1.6. Because $x_{k}^{*}$ and $z_{k}^{*}$ both satisfy (2.3d), (3.8) implies that

$$
\begin{equation*}
\left\|\left[z_{k}^{*}\right]_{\mathcal{I}}\right\| \leq \sqrt{n} \omega_{k} \tag{3.9}
\end{equation*}
$$

Combining (3.7) and (3.9),

$$
\begin{equation*}
\left\|\widehat{g}\left(x_{k}^{*}\right)-\widehat{J}\left(x_{k}^{*}\right)^{T} \widehat{y}_{k}+\left(\widehat{J}\left(x_{k}^{*}\right)-\widehat{J}_{k}\right)^{T}\left(y_{k}^{*}-y_{k}\right)\right\| \leq \sqrt{n} \omega_{k} \tag{3.10}
\end{equation*}
$$

But, from the triangle and Cauchy-Schwartz inequalities, we have

$$
\begin{align*}
\left\|\widehat{g}\left(x_{k}^{*}\right)-\widehat{J}\left(x_{k}^{*}\right)^{T} \widehat{y}_{k}\right\| \leq\left\|\widehat{g}\left(x_{k}^{*}\right)-\widehat{J}\left(x_{k}^{*}\right)^{T} \widehat{y}_{k}+\left(\widehat{J}\left(x_{k}^{*}\right)-\widehat{J}_{k}\right)^{T}\left(y_{k}^{*}-y_{k}\right)\right\| \\
+\left\|\widehat{J}\left(x_{k}^{*}\right)-\widehat{J}_{k}\right\|\left\|y_{k}^{*}-y_{k}\right\| . \tag{3.11}
\end{align*}
$$

Also, the continuity of $J$ implies that there exists a positive constant $M$ such that $\left\|\widehat{J}\left(x_{k}^{*}\right)-\widehat{J}_{k}\right\| \leq M \frac{\sqrt{n}}{\alpha_{1}}\left\|x_{k}^{*}-x_{k}\right\|$. Together, (3.11) and (3.10) imply that

$$
\begin{equation*}
\left\|\widehat{g}\left(x_{k}^{*}\right)-\widehat{J}\left(x_{k}^{*}\right)^{T} \widehat{y}_{k}\right\| \leq \sqrt{n} \omega_{k}+M \frac{\sqrt{n}}{\alpha_{1}}\left\|x_{k}^{*}-x_{k}\right\|\left\|y_{k}^{*}-y_{k}\right\| \tag{3.12}
\end{equation*}
$$

and so we have derived a bound on $\left\|\widehat{g}\left(x_{k}^{*}\right)-\widehat{J}\left(x_{k}^{*}\right)^{T} \widehat{y}_{k}\right\|$, as required.
We now derive (3.2). From the triangle inequality,

$$
\begin{equation*}
\left\|\widehat{y}_{k}-y_{*}\right\| \leq\left\|\widetilde{y}\left(x_{k}^{*}\right)-\widehat{y}_{k}\right\|+\left\|\widetilde{y}\left(x_{k}^{*}\right)-y_{*}\right\| . \tag{3.13}
\end{equation*}
$$

Using inequality (3.12) in (3.6), we deduce that

$$
\begin{equation*}
\left\|\widetilde{y}\left(x_{k}^{*}\right)-\widehat{y}_{k}\right\| \leq \alpha_{1} \omega_{k}+M\left\|x_{k}^{*}-x_{k}\right\|\left\|y_{k}^{*}-y_{k}\right\| \tag{3.14}
\end{equation*}
$$

and Lemma 3.3 implies that there exists a constant $\alpha_{2}$ such that

$$
\begin{equation*}
\left\|\widetilde{y}\left(x_{k}^{*}\right)-y_{*}\right\| \leq \alpha_{2}\left\|x_{k}^{*}-x_{*}\right\| \tag{3.15}
\end{equation*}
$$

for all $k \in \mathcal{K}$ large enough (recall that $y_{*} \equiv \widetilde{y}\left(x_{*}\right)$ ). Substituting (3.14) and (3.15) into (3.13), we obtain $\left\|\widehat{y}_{k}-y_{*}\right\| \leq \beta_{1}$ as stated in (3.2).

We now prove (3.3). From the definition of $\widehat{y}_{k}$, rearranging terms yields

$$
\begin{equation*}
\rho_{k} c\left(x_{k}^{*}\right)=y_{k}^{*}-\widehat{y}_{k} . \tag{3.16}
\end{equation*}
$$

Taking norms of both sides of (3.16) and using (3.2) yields

$$
\begin{aligned}
\rho_{k}\left\|c\left(x_{k}^{*}\right)\right\| & =\left\|y_{k}^{*}-\widehat{y}_{k}\right\| \\
& =\left\|y_{k}-y_{*}+y_{*}-\widehat{y}_{k}+y_{k}^{*}-y_{k}\right\| \\
& \leq\left\|\widehat{y}_{k}-y_{*}\right\|+\left\|y_{k}-y_{*}\right\|+\left\|y_{k}^{*}-y_{k}\right\| \\
& \leq \beta_{1}+\left\|y_{k}-y_{*}\right\|+\left\|y_{k}^{*}-y_{k}\right\| \\
& \equiv \beta_{2}
\end{aligned}
$$

and so Part 1 of Lemma 3.4 is proved.
Now suppose that $\left\|y_{k}^{*}-y_{k}\right\| \rightarrow 0$ as $k \in \mathcal{K}$ goes to infinity. Because $\left\{x_{k}^{*}\right\}$ and $\left\{x_{k}\right\}$ are in the compact set $\mathcal{B},\left\|x_{k}^{*}-x_{k}\right\|$ is bounded. We conclude from (3.2) that $\widehat{y}_{k} \rightarrow y_{*}$ as $k \in \mathcal{K}$ goes to infinity. We also conclude from the continuity of $J$ that $\left\|J\left(x_{k}^{*}\right)-J_{k}\right\|$ is bounded, so that

$$
\begin{equation*}
\lim _{k \in \mathcal{K}}\left\|\left(J\left(x_{k}^{*}\right)-J_{k}\right)^{T}\left(y_{k}^{*}-y_{k}\right)\right\|=0 \tag{3.17}
\end{equation*}
$$

On the other hand, suppose that $\left\|y_{k}^{*}-y_{k}\right\|$ is uniformly bounded and that $\lim _{k \in \mathcal{K}} \| x_{k}^{*}-$ $x_{k} \|=0$. We then conclude from (3.2) that $\widehat{y}_{k} \rightarrow y_{*}$ as $k \in \mathcal{K}$ goes to infinity and (3.17) holds. Because $\lim _{k \in \mathcal{K}}\left(x_{k}^{*}, \widehat{y}_{k}\right)=\left(x_{*}, y_{*}\right)$,

$$
g\left(x_{k}^{*}\right)-J\left(x_{k}^{*}\right)^{T} \widehat{y}_{k} \rightarrow g_{*}-J_{*}^{T} y_{*}
$$

and so (3.7) and (3.17) together imply that

$$
\begin{equation*}
z_{k}^{*} \rightarrow z_{*} \equiv \nabla_{x} \mathcal{L}\left(x_{*}, y_{*}, 0\right) \tag{3.18}
\end{equation*}
$$

as $k \in \mathcal{K}$ goes to infinity. Thus we have proved Part 2 of Lemma 3.4.
Now suppose that

$$
\begin{equation*}
c_{*}=0 \tag{3.19}
\end{equation*}
$$

Each $x_{k}^{*}$ and $z_{k}^{*}$ satisfies (2.3d). Then $\lim _{k \in \mathcal{K}}\left(x_{k}^{*}, z_{k}^{*}\right)=\left(x_{*}, z_{*}\right)$, and $\omega_{k} \rightarrow 0$ implies

$$
\begin{equation*}
\min \left(x_{*}, z_{*}\right)=0 \tag{3.20}
\end{equation*}
$$

Therefore, (3.18)-(3.20) imply that $\left(x_{*}, y_{*}, z_{*}\right)$ satisfies (1.5) and so it is a first-order KKT point for (GNP). Part 3 is thus proved, and the proof is complete.

The conclusions of Lemma 3.4 pertain to any sequence $\left\{\left(x_{k}^{*}, \Delta y_{k}^{*}, z_{k}^{*}\right)\right\}$ satisfying the approximate first-order conditions (2.3). Algorithm 2 generates such a sequence and also generates auxiliary sequences of scalars $\left\{\omega_{k}\right\},\left\{\rho_{k}\right\}$, and $\left\{\sigma_{k}\right\}$ in such a way as to guarantee that the hypotheses of Lemma 3.4 hold. We demonstrate in Theorem 3.6 that the condition of Part 3 of Lemma 3.4 holds. Therefore, every limit point of the sequence $\left\{\left(x_{k}^{*}, \widehat{y}_{k}, z_{k}^{*}\right)\right\}$ is a first-order KKT point for (GNP).
3.1.2. Convergence of $\left\|y_{k}\right\| / \rho_{k}$. Before laying out the global convergence properties of the stabilized LCL method, we need to show that if $\rho_{k} \rightarrow \infty$ then the quotient $\left\|y_{k}\right\| / \rho_{k}$ converges to 0 . This property is required (and used by Conn et al. [9, 10]) in lieu of assuming that $\left\|y_{k}\right\|$ remains bounded.

Lemma 3.5. Suppose that $\rho_{k} \rightarrow \infty$ as $k$ increases when Algorithm 2 is executed. Then $\left\|y_{k}\right\| / \rho_{k} \rightarrow 0$.

Proof. The multiplier update of Algorithm 2 (Step 5) is $y_{k+1} \leftarrow y_{k}^{*}-\rho_{k} c\left(x_{k}^{*}\right)$. If this is replaced by $y_{k+1} \leftarrow y_{k}-\rho_{k} c\left(x_{k}^{*}\right)$, Lemma 4.1 of Conn et al. [10] applies. The construction of the forcing sequence $\eta_{k}$ (Steps 8 and 11) is therefore sufficient to guarantee that $\left\|y_{k}\right\| / \rho_{k} \rightarrow 0$. Note that the norm of the difference between the two updates is given by $\left\|y_{k}^{*}-y_{k}\right\| \equiv\left\|\Delta y_{k}^{*}\right\|$. This difference is bounded because $\Delta y_{k}^{*}$ satisfies (2.3e). Therefore, $\left\|y_{k}\right\| / \rho_{k} \rightarrow 0$ in Algorithm 2 as $\rho_{k} \rightarrow \infty$.
3.1.3. Main convergence result. With Lemmas 3.4-3.5 in hand, we are now able to prove global convergence of the stabilized LCL method.

ThEOREM 3.6. Let $\left\{\left(x_{k}^{*}, y_{k}^{*}, z_{k}^{*}\right)\right\}$ be the sequence of vectors generated by Algorithm 2 with tolerances $\omega_{*}=0$ and $\eta_{*}=0$. Let $x_{*}$ be any limit point of the sequence $\left\{x_{k}^{*}\right\}$ and let $\mathcal{K}$ be the infinite set of indices associated with that convergent subsequence. Then, under the assumptions of Lemma 3.4, Parts 1, 2, and 3 of that lemma hold.

Proof. Algorithm 2 generates positive scalars $\rho_{k}$ and, by Steps 1, 8 , and 11, generates positive scalars $\omega_{k} \rightarrow 0$ and $\eta_{k} \rightarrow 0$. Step 2 of the algorithm generates a sequence $\left\{\left(x_{k}^{*}, y_{k}^{*}, z_{k}^{*}\right)\right\}$, where $y_{k}^{*} \equiv y_{k}+\Delta y_{k}^{*}$. Each $\left(x_{k}^{*}, \Delta y_{k}^{*}, z_{k}^{*}\right)$ satisfies (2.3). Therefore, the hypotheses of Lemma 3.4 hold, and Part 1 of the lemma follows immediately.

Note that each $x_{k}^{*}$ satisfies (2.3a), and so $x_{k}^{*} \geq 0$ for all $k$. Thus, $x_{*} \geq 0$. Moreover, because $\tau_{\sigma}>1$ and $\bar{\sigma}$ is finite, Steps 6 and 10 of Algorithm 2 ensure that $\sigma_{k}$ is uniformly bounded. We then need to consider the four possible cases. For all $k \in \mathcal{K}$,

1. $\rho_{k}$ is uniformly bounded, and $\sigma_{k} \rightarrow 0$ as $k$ gets large;
2. $\rho_{k}$ is uniformly bounded, and $\sigma_{k}$ is uniformly bounded away from zero;
3. $\rho_{k} \rightarrow \infty$ and $\sigma_{k} \rightarrow 0$ as $k$ gets large;
4. $\rho_{k} \rightarrow \infty$ and $\sigma_{k}$ is uniformly bounded away from zero.

For the remainder of this proof, we consider only $k \in \mathcal{K}$.
We dismiss Case 1 because it cannot be generated by the algorithm. (As $k$ gets large, $\sigma_{k} \rightarrow 0$ only if Step 10 is executed infinitely many times, contradicting the finiteness of $\rho_{k}$.)

Case 2 implies that Step 6 of Algorithm 2 is executed for all $k$ large enough. Thus, $x_{k+1}=x_{k}^{*}$ for all large $k$, and hence $x_{k}^{*} \rightarrow x_{*}$ implies $x_{k} \rightarrow x_{*}$. Therefore, $\left\|x_{k}^{*}-x_{k}\right\| \rightarrow 0$. Because each $\Delta y_{k}^{*}$ satisfies (2.3e), $y_{k}^{*}$ satisfies

$$
\begin{equation*}
\left\|y_{k}^{*}-y_{k}\right\|_{\infty} \leq \omega_{k}+\sigma_{k} \tag{3.21}
\end{equation*}
$$

Because $\sigma_{k}$ and $\omega_{k}$ are uniformly bounded, Part 2 of Lemma 3.4 holds. In addition, $\left\|c\left(x_{k}^{*}\right)\right\| \leq \eta_{k}$ for all $k$ large enough, and so $\eta_{k} \rightarrow 0$ implies that $c\left(x_{k}^{*}\right) \rightarrow 0$. By continuity of $c, c_{*}=0$. Thus, Part 3 of Lemma 3.4 holds.

Now consider Case 3. Because $\sigma_{k} \rightarrow 0$ and $\omega_{k} \rightarrow 0$, (3.21) implies that $\left\|y_{k}^{*}-y_{k}\right\| \rightarrow$ 0 as $k$ increases. Then Part 2 of the lemma holds. To show that $c\left(x_{k}^{*}\right) \rightarrow 0$, divide both sides of (3.3) by $\rho_{k}$ to obtain

$$
\left\|c\left(x_{k}^{*}\right)\right\| \leq \underbrace{\frac{\alpha_{1} \omega_{k}}{\rho_{k}}}_{(a)}+\underbrace{\frac{1}{\rho_{k}}\left\|y_{k}^{*}-y_{k}\right\|\left(M\left\|x_{k}^{*}-x_{k}\right\|+1\right)}_{(b)}+\underbrace{\frac{\alpha_{2}}{\rho_{k}}\left\|x_{k}^{*}-x_{*}\right\|}_{(c)}+\underbrace{\frac{1}{\rho_{k}}\left\|y_{k}-y_{*}\right\|}_{(d)} .
$$

Term (a) clearly goes to zero as $\rho_{k}$ increases. Because $y_{k}^{*}$ and $y_{k}$ satisfy (3.21), and because $x_{k}^{*}$ and $x_{k}$ belong to the compact set $\mathcal{B},(b)$ and $(c)$ go to zero as $\rho_{k}$ increases. By Lemma 3.5, $\left\|y_{k}\right\| / \rho_{k} \rightarrow 0$, and so ( $d$ ) goes to 0 . We conclude that $\left\|c\left(x_{k}^{*}\right)\right\| \rightarrow 0$ as $k$ increases, as required.

In Case 4, both Steps 6 and 10 are executed infinitely often. But because $\sigma_{k}$ is uniformly bounded, so is $\left\|y_{k}^{*}-y_{k}\right\|$. As in Case $2,\left\|x_{k}^{*}-x_{k}\right\| \rightarrow 0$ as $k$ get large, and so Part 2 of Lemma 3.4 holds. The rest of the analysis for this case is the same as for Case 3.
3.1.4. Finite termination. Note that the convergence test takes place only if Step 3 of Algorithm 2 tests true; i.e., if $\left\|c\left(x_{k}^{*}\right)\right\| \leq \eta_{k}$ (because $\eta_{*}=0$ ). To guarantee that the algorithm will eventually terminate as the iterates $x_{k}, y_{k}$, and $z_{k}$ converge, we need to guarantee that Steps 4 and 7 execute infinitely often. The forcing sequence $\eta_{k}$ is intimately tied to this occurrence. For example, if $\eta_{k} \equiv 0$, then we would not normally expect Step 3 to evaluate true (except in rare occasions when $c\left(x_{k}^{*}\right)=0$ ). The forcing sequence defined by Steps 8 and 11 of Algorithm 2 is suggested by Conn et al. $[9,10]$. The following corollaries show that this forcing sequence has the desired property and summarize the global convergence properties of Algorithm 2. Unlike for the previous results in this section, we now need to strengthen our assumptions and require that only a single limit point exist.

Corollary 3.7 (Global convergence). Let $\left\{\left(x_{k}, y_{k}, z_{k}\right)\right\}$ be the sequence of vectors generated by Algorithm 2. Let $x_{*}$ be the single limit point of the sequence $\left\{x_{k}^{*}\right\}$. Suppose that Assumptions 1.1, 3.1, and 3.2 hold. Then

$$
\lim _{k \rightarrow \infty}\left(x_{k}, y_{k}, z_{k}\right)=\left(x_{*}, y_{*}, z_{*}\right)
$$

and $\left(x_{*}, y_{*}, z_{*}\right)$ is a first-order KKT point for (GNP).
Proof. Let $\left\{\left(x_{k}^{*}, y_{k}^{*}, z_{k}^{*}\right)\right\}$ be the sequence of vectors generated by Step 2 of Algorithm 2 and set $\widehat{y}_{k}=\widehat{y}\left(x_{k}^{*}, y_{k}^{*}, \rho_{k}\right)$. By Lemma 3.4 and Theorem 3.6,

$$
\lim _{k \rightarrow \infty} \widehat{y}_{k}=y_{*} \quad \text { and } \quad \lim _{k \rightarrow \infty} z_{k}^{*}=z_{*} .
$$

Moreover, $\left(x_{*}, y_{*}, z_{*}\right)$ is a first-order KKT point for (GNP). Suppose Step 4 is executed infinitely often. The result then follows immediately because $x_{k}, y_{k}$, and $z_{k}$ are updated infinitely often and form a convergent sequence from $x_{k}^{*}, \widehat{y}_{k}$, and $z_{k}^{*}$.

We now show by contradiction that Step 4 does occur infinitely often. Suppose instead that it does not. Then there exists a $k_{1}$ large enough so that Steps 9 and 10 are executed for all $k>k_{1}$. Consider only iterations $k>k_{1}$. Then $y_{k} \equiv \bar{y}$ and
$\rho_{k} \rightarrow \infty$. From (3.16),

$$
\begin{align*}
\rho_{k}\left\|c\left(x_{k}^{*}\right)\right\| & =\left\|y_{k}^{*}-\widehat{y}_{k}\right\| \\
& =\left\|\left(y_{k}^{*}-\bar{y}\right)+\bar{y}-\widehat{y}_{k}\right\|  \tag{3.22}\\
& \leq\left\|y_{k}^{*}-\bar{y}\right\|+\|\bar{y}\|+\left\|\widehat{y}_{k}\right\| .
\end{align*}
$$

The vector $y_{k}^{*} \equiv y_{k}+\Delta y_{k}^{*}$ and each $\Delta y_{k}^{*}$ satisfies (2.3e). Thus, $\left\|y_{k}^{*}-\bar{y}\right\| \leq \sigma_{k}+$ $\omega_{k}$. Moreover, $\lim _{k \rightarrow \infty} \widehat{y}_{k}=y_{*}$ and $y_{*}$ is bounded (see Assumption 3.2). Then, from (3.22), there exists some constant $L>0$, independent of $k$, such that

$$
\begin{equation*}
\rho_{k}\left\|c\left(x_{k}^{*}\right)\right\| \leq L \tag{3.23}
\end{equation*}
$$

for all $k$. But the test at Step 3 fails at every iteration, so that

$$
\begin{equation*}
\eta_{k}<\left\|c\left(x_{k}^{*}\right)\right\| . \tag{3.24}
\end{equation*}
$$

Combining (3.23) and (3.24), we find that

$$
\begin{equation*}
\rho_{k} \eta_{k}<\rho_{k}\left\|c\left(x_{k}^{*}\right)\right\| \leq L \tag{3.25}
\end{equation*}
$$

From Step 11, $\eta_{k+1}=\eta_{0} / \rho_{k+1}^{\alpha}$, so

$$
\begin{equation*}
\rho_{k} \eta_{k}=\rho_{k} \frac{\eta_{0}}{\rho_{k}^{\alpha}}=\eta_{0} \rho_{k}^{1-\alpha} \tag{3.26}
\end{equation*}
$$

Substituting (3.26) into (3.25), we find that $\eta_{0} \rho_{k}^{1-\alpha}<L$ for all $k$. This is a contradiction under the hypothesis that $\alpha<1$ and $\rho_{k} \rightarrow \infty$. Therefore, Step 4 must occur infinitely often.

The following result simply asserts that Algorithm 2 will eventually exit when, as in practice, $\omega_{*}$ and $\eta_{*}$ are positive.

Corollary 3.8 (Finite Termination). Suppose that the convergence tolerances $\omega_{*}$ and $\eta_{*}$ are strictly positive. Then, under the assumptions of Corollary 3.7, Algorithm 2 terminates after a finite number of iterations.

Proof. Let $\left\{\left(x_{k}^{*}, y_{k}^{*}, z_{k}^{*}\right)\right\}$ and $x_{*}$ be as defined in Theorem 3.6. Set $\widehat{y}_{k}=\widehat{y}_{k}\left(x_{k}^{*}, y_{k}^{*}, \rho_{k}\right)$. By that theorem,

$$
\begin{aligned}
& \lim _{k \rightarrow \infty} \widehat{y}_{k}=y_{*} \\
& \lim _{k \rightarrow \infty} z_{k}^{*}=z_{*} \stackrel{\text { def }}{=} \nabla_{x} \mathcal{L}\left(x_{*}, y_{*}, 0\right)
\end{aligned}
$$

and $\left(x_{*}, y_{*}, z_{*}\right)$ is a first-order KKT point for (GNP). Then, $\left(x_{*}, y_{*}, z_{*}\right)$ must satisfy (1.7). By the continuity of $c, \lim _{k \rightarrow \infty}\left\|c\left(x_{k}^{*}\right)\right\| \rightarrow c_{*}=0$, and because $\eta_{*}>0$,

$$
\left\|c\left(x_{k}^{*}\right)\right\|<\eta_{*} \leq \max \left(\eta_{k}, \eta_{*}\right)
$$

for all $k \in \mathcal{K}$ large enough. Consequently, Step 7 is executed infinitely often and

$$
\lim _{k \rightarrow \infty}\left(x_{k}, y_{k}, z_{k}\right) \rightarrow\left(x_{*}, y_{*}, z_{*}\right)
$$

Because $\omega_{*}>0$ and $\eta_{*}>0,\left(x_{k}, y_{k}, z_{k}\right)$ satisfies conditions (1.7) for some $k$ large enough.
3.2. Local convergence properties. In this section we show that the stabilized LCL algorithm preserves the local convergence characteristics of Robinson's original LCL algorithm. Moreover, it can retain fast local convergence under inexact solutions to the subproblems.

Bertsekas [1] and Conn et al. [9, 10] show how to construct a forcing sequence $\left\{\eta_{k}\right\}$ to guarantee that $\left\|c\left(x_{k}^{*}\right)\right\| \leq \eta_{k}$ will eventually always be true so that the iterates $x_{k}$, $y_{k}$, and $z_{k}$ are updated (see Step 4 of Algorithm 2) for all iterations after some $k$ large enough. The penalty parameter $\rho_{k}$ then remains uniformly bounded-an important property. These results rely on a relationship between $\left\|c\left(x_{k}^{*}\right)\right\|$ and $\rho_{k}$, namely (3.3). We know from the BCL convergence theory that the convergence rate approaches superlinear as $\rho_{k}$ grows large (cf. [1] and [9, 10]). Because $\eta_{k}$ is reduced at a sublinear rate, $\left\|c\left(x_{k}^{*}\right)\right\|$ will eventually go to zero faster than $\eta_{k}$, at which point it is no longer necessary to increase $\rho_{k}$. Thus, we can be assured that Algorithm 2 does not increase $\rho_{k}$ without bound.

Bertsekas suggests constructing the sequence $\eta_{k}$ as

$$
\begin{equation*}
\eta_{k+1}=\gamma\left\|c\left(x_{k}^{*}\right)\right\| \tag{3.27}
\end{equation*}
$$

for some $\gamma<1$. Within Algorithm 2, this would lead to the following update rule:

$$
\rho_{k+1}= \begin{cases}\rho_{k} & \text { if }\left\|c\left(x_{k}^{*}\right)\right\| \leq \gamma\left\|c\left(x_{k}\right)\right\|  \tag{3.28}\\ \tau_{\rho} \rho_{k} & \text { if }\left\|c\left(x_{k}^{*}\right)\right\|>\gamma\left\|c\left(x_{k}\right)\right\| .\end{cases}
$$

As $\rho_{k}$ gets larger, the convergence rate gets arbitrarily close to superlinear, so that the first case of (3.28) is always satisfied, and $\rho_{k}$ becomes constant for all $k$ large enough. We prefer not to use rule (3.27) because it may be too strict. Any intermediate (and nonoptimal) iterate $x_{k}^{*}$ could be feasible or nearly feasible for (GNP), so that $\left\|c\left(x_{k}^{*}\right)\right\|$ could be very small. Then $\eta_{k+1}$ would be smaller than warranted on the following iteration. The forcing sequence suggested by Conn et al. [9, 10] does not suffer from this defect and has been proven by them to keep $\rho_{k}$ bounded. We have used this update in Algorithm 2 (see Steps 8 and 11).

For this analysis and the remainder of this section, we assume that $\rho_{k}$ is uniformly bounded, so that $\rho_{k}=\bar{\rho}$ for all $k$ greater than some $\bar{k}$. Hence, we drop the subscript on $\rho_{k}$ and simply write $\bar{\rho}$. We consider only iterations $k>\bar{k}$.

We begin by discussing the local convergence rates of the Algorithm 2 under the assumption that the elastic variables are always zero-that is, the linearized constraints are always satisfied. Next, we show that after finitely many iterations the elastic penalty parameter $\sigma_{k}$ will always be large enough to guarantee that this assumption holds. In this way, we demonstrate that stabilized LCL becomes equivalent to MINOS (and to canonical LCL) as it approaches the solution.
3.2.1. Convergence rates. Robinson's [35] local convergence analysis applies to the canonical LCL algorithm under the special case in which $\rho_{k} \equiv 0$ (cf. (2.2)) and each subproblem is solved to full accuracy (i.e., $\omega_{k} \equiv 0$ ). He proved that one can expect fast convergence from a good enough starting point. In particular, under Assumptions 1.1, 1.5, and 3.2, we can expect an R-quadratic rate of convergence (see Ortega and Rheinboldt [33] for an in-depth discussion of root-convergence rates). For a sufficiently good starting point, Robinson [36] proves that the subproblems ( $\mathrm{LC}_{k}$ ) are always feasible. He also shows that near a solution, the solutions to the LC subproblems, if parameterized appropriately, form a continuous path converging to $\left(x_{*}, y_{*}, z_{*}\right)$.

In a later paper, Bräuninger [6] shows how the fast local convergence rate can be preserved with only approximate solutions of the subproblems (again, with $\rho_{k} \equiv 0$ ). The subproblems are solved to a tolerance that is tightened at a rate that matches the decrease in the square of the primal and dual infeasibilities. Our proposed LCL algorithm uses a similar strategy.

Robinson's local convergence analysis also applies to the canonical LCL algorithm when $\rho_{k} \equiv \bar{\rho}>0$. One can see this by considering the following optimization problem:

$$
\begin{array}{ll}
\underset{x}{\operatorname{minimize}} & f(x)+\frac{1}{2} \bar{\rho}\|c(x)\|^{2}  \tag{3.29}\\
\text { subject to } & c(x)=0, \quad x \geq 0
\end{array}
$$

The solutions of (3.29) are identical to the solutions of (GNP). The Robinson LCL subproblem objective corresponding to (3.29) is given by

$$
R_{k}(x)=f(x)+\frac{1}{2} \bar{\rho}\|c(x)\|^{2}-y_{k}^{T} c(x)
$$

The canonical LCL subproblem objective is $\mathcal{L}_{k}(x) \equiv \mathcal{L}\left(x, y_{k}, \rho_{k}\right)$, and so $\mathcal{L}_{k}(x) \equiv$ $R_{k}(x)$ for all $k$ because $\rho_{k} \equiv \bar{\rho}$. We then observe that the canonical LCL subproblem corresponding to (GNP), with a penalty parameter $\rho_{k} \equiv \bar{\rho}$, is equivalent to the Robinson LCL subproblem corresponding to problem (3.29), with $\rho_{k} \equiv 0$. The convergence characteristics of the canonical LCL algorithm are therefore the same as those demonstrated by Robinson [35]. (However, while the asymptotic convergence rate remains R-quadratic, we expect a different asymptotic error constant.)

Under the assumption that the elastic variables are always equal to 0 and that $\bar{\rho}$ is finite, the steps executed by Algorithms 1 and 2 are identical, and the subproblems $\left(\mathrm{ELC}_{k}\right)$ and $\left(\mathrm{LC}_{k}\right)$ are also identical. The only difference is the multiplier update formulas:

$$
\begin{array}{ll}
\text { Canonical LCL update } & y_{k+1}=y_{k}^{*} \\
\text { Stabilized LCL update } & y_{k+1}=y_{k}^{*}-\bar{\rho} c\left(x_{k}^{*}\right) \tag{3.30b}
\end{array}
$$

which differ only by the vector $\bar{\rho} c\left(x_{k}^{*}\right)$. We may think of this vector as a perturbation of the LCL multiplier update (3.30a). Moreover, Robinson [35] shows that this perturbation converges to 0 at the same rate as $\left\{x_{k}^{*}\right\}$ converges to $x_{*}$. Therefore, it does not interfere with the convergence rate of the stabilized LCL iterates. Robinson's local convergence analysis then applies to the stabilized LCL method.

We summarize the convergence results in Theorem 3.9. Note that the function

$$
F(x, y, z)=\left[\begin{array}{c}
c(x) \\
\nabla_{x} \mathcal{L}(x, y, \rho)-z \\
\min (x, z)
\end{array}\right]
$$

captures the first-order optimality conditions of (GNP), in the sense that

$$
F\left(x_{*}, y_{*}, z_{*}\right)=0
$$

if and only if $\left(x_{*}, y_{*}, z_{*}\right)$ is a first-order KKT point for (GNP). Thus, $\|F(x, y, z)\|$ is a measure of the deviation from optimality. For the next theorem only, define

$$
r=\left(\begin{array}{l}
x \\
y \\
z
\end{array}\right) \quad \text { and } \quad F(r)=F(x, y, z)
$$

Theorem 3.9 (Robinson [35]; Bräuninger [6]). Suppose Assumptions 1.1-1.5 and 3.2 hold at $r_{*}$. Moreover, suppose $\omega_{k}=O\left(\left\|F\left(r_{k}\right)\right\|^{2}\right)$ for all $k \geq 0$. Then there is a positive constant $\delta$ such that if

$$
\left\|r_{0}-r_{*}\right\|<\delta
$$

the sequence $\left\{r_{k}\right\}$ generated by Algorithm 2 converges to $r_{*}$. Moreover, the sequence converges $R$-quadratically, so that for all $k \geq 0$,

$$
\begin{equation*}
\left\|r_{k}-r_{*}\right\| \leq Q\left(\frac{1}{2}\right)^{2^{k}} \tag{3.31}
\end{equation*}
$$

for some positive constant $Q$. Also,

$$
\begin{equation*}
\left\|r_{k+1}-r_{k}\right\| \leq M\left\|F\left(r_{k}\right)\right\| \tag{3.32}
\end{equation*}
$$

for some positive constant $M$.
Robinson does not state (3.32) as part of a theorem, but it is found in the proof of (3.31).
3.2.2. Asymptotic equivalence to MINOS. Much of the efficiency of LCL methods, including MINOS, derives from the fact that they eventually identify the correct active set, and each subproblem restricts its search to the subspace defined by a linear approximation of the constraints. This approximation can be very accurate near the solution. The stabilized LCL subproblems do not restrict themselves to this subspace. In early iterations we do not expect, nor do we wish, the method to honor these linearizations. The elastic variables give the subproblems an opportunity to deviate from this subspace. In order to recover LCL's fast convergence rate, however, it is not desirable to allow deviation near the solution.

We show below that as the stabilized LCL iterations approach a solution of (GNP), the solutions of the stabilized LCL subproblems eventually always have the elastic variables equal to zero. Hence, $\bar{c}_{k}\left(x_{k}^{*}\right)=0$ and $v_{k}^{*}, w_{k}^{*}=0$, so that each $x_{k}^{*}$ satisfies the constraints of the canonical LCL (and MINOS) subproblem, and the objective of $\left(\mathrm{ELC}_{k}\right)$ at $\left(x_{k}^{*}, v_{k}^{*}, w_{k}^{*}\right)$ is equivalent to $\left(\mathrm{LC}_{k}\right)$.

As discussed in $\S 3.2 .1$, Theorem 3.9 applies to Algorithm 2, and so for all $k$ large enough, (3.32) yields

$$
\begin{equation*}
\left\|y_{k}^{*}-y_{k}\right\| \leq M\left\|F\left(x_{k}, y_{k}, z_{k}\right)\right\| \tag{3.33}
\end{equation*}
$$

where $y_{k}^{*} \equiv y_{k}+\Delta y_{k}^{*}$ and $M$ is some positive constant. By Corollary 3.7, $\left(x_{k}, y_{k}, z_{k}\right) \rightarrow$ $\left(x_{*}, y_{*}, z_{*}\right)$, and because $\left\|F\left(x_{*}, y_{*}, z_{*}\right)\right\|=0$ and $F$ is continuous,

$$
\begin{equation*}
\left\|F\left(x_{k}, y_{k}, z_{k}\right)\right\|<\frac{\sigma}{M} \tag{3.34}
\end{equation*}
$$

for all $k$ large enough ( $\underline{\sigma}$ is defined in Algorithm 2). Combining (3.33) and (3.34), we conclude that $\left\|y_{k}^{*}-y_{k}\right\|<\underline{\sigma}$, or equivalently,

$$
\begin{equation*}
\left\|\Delta y_{k}^{*}\right\|<\underline{\sigma} \tag{3.35}
\end{equation*}
$$

for all $k$ large enough. However, Step 6 of Algorithm 2 guarantees that $\underline{\sigma} \leq \sigma_{k}$ for all $k$, and so from (3.35), $\left\|\Delta y_{k}^{*}\right\|<\sigma_{k}$ for all $k$ large enough. Lemma 2.1 then implies that $\sigma_{k}$ will be sufficiently large that the optimal elastic variables will be equal to 0 .
3.3. Infeasible problems. Not all optimization problems are feasible. The user of an optimization algorithm may formulate a set of nonlinear constraints $c(x)=0$ for which no nonnegative solution exists. Detecting infeasibility of the system $c(x)=0$, $x \geq 0$, is equivalent to verifying that the global minimizer of

$$
\begin{array}{ll}
\underset{x}{\operatorname{minimize}} & \frac{1}{2}\|c(x)\|^{2} \\
\text { subject to } & x \geq 0 \tag{3.36}
\end{array}
$$

yields a positive objective value. Detecting such infeasibility is a useful feature, but it is a very difficult problem and is beyond the purview of this paper.

We analyze the properties of the stabilized LCL algorithm when it is applied to an infeasible problem with convergence tolerances $\omega_{*}=\eta_{*}=0$. We show that Algorithm 2 converges to a point that satisfies the first-order optimality conditions of the minimum-norm problem (3.36).

Theorem 3.10. Let $x_{*}$ be any limit point of the sequence of vectors $\left\{x_{k}^{*}\right\}$ generated by Algorithm 2, and let $\mathcal{K}$ be the infinite set of indices associated with that subsequence. Suppose that (GNP) is infeasible. Then, under the assumptions of Lemma 3.4,

$$
\lim _{k \in \mathcal{K}} J\left(x_{k}^{*}\right)^{T} c\left(x_{k}^{*}\right)=z_{*} \stackrel{\text { def }}{=} J_{*}^{T} c_{*},
$$

and $\left(x_{*}, z_{*}\right)$ is a first-order KKT point for (3.36).
Proof. The pair $\left(x_{*}, z_{*}\right)$ satisfies the first-order KKT conditions of (3.36) if

$$
\begin{align*}
J_{*}^{T} c_{*} & =z_{*} \\
\min \left(x_{*}, z_{*}\right) & =0 \tag{3.37}
\end{align*}
$$

Because (GNP) is infeasible, there exists a constant $\delta>0$ such that $\delta<\|c(x)\|$ for all $x \geq 0$. Moreover, Steps 8 and 11 of Algorithm 2 generate a sequence $\left\{\eta_{k}\right\}$ converging to 0 , and so $\eta_{k}<\delta$ for all $k$ large enough. Consider only such $k$. Then, $\eta_{k}<\delta<\left\|c\left(x_{k}^{*}\right)\right\|$, and Step 10 is executed at every $k$, so that $\rho_{k} \rightarrow \infty$ and $\sigma_{k} \rightarrow 0$. Moreover, $x_{k}$ and $y_{k}$ are not updated, so that for some $n$-vector $\bar{x}$ and $m$-vector $\bar{y}$,

$$
\begin{equation*}
x_{k} \equiv \bar{x} \quad \text { and } \quad y_{k} \equiv \bar{y} \tag{3.38}
\end{equation*}
$$

Note that Algorithm 2 generates $x_{k}^{*}$ satisfying (2.3). Therefore, $x_{k}^{*} \geq 0$ for all $k$, and so $\lim _{k \in \mathcal{K}} x_{k}^{*}=x_{*}$ implies

$$
\begin{equation*}
x_{*} \geq 0 \tag{3.39}
\end{equation*}
$$

From (2.3d), (3.7), and (3.38),

$$
\begin{equation*}
g\left(x_{k}^{*}\right)-J\left(x_{k}^{*}\right)^{T}\left(\bar{y}-\rho_{k} c\left(x_{k}^{*}\right)\right)-J(\bar{x})^{T} \Delta y_{k}^{*} \geq-\omega_{k} e \tag{3.40}
\end{equation*}
$$

or, after rearranging terms,

$$
\begin{equation*}
\underbrace{g\left(x_{k}^{*}\right)-J\left(x_{k}^{*}\right)^{T} \bar{y}}_{(a)}-\underbrace{J(\bar{x})^{T} \Delta y_{k}^{*}}_{(b)}+\rho_{k} J\left(x_{k}^{*}\right)^{T} c\left(x_{k}^{*}\right) \geq-\omega_{k} e \tag{3.41}
\end{equation*}
$$

By hypothesis, all iterates $x_{k}^{*}$ lie in a compact set, and so $(a)$ is bounded because $g$ and $J$ are continuous and $\bar{y}$ is constant. Also, (b) is bounded because $\bar{x}$ is constant,
and from (2.3e) we have $\left\|\Delta y_{k}^{*}\right\|_{\infty} \leq \sigma_{k}+\omega_{k}$. Then, because $\omega_{k} \rightarrow 0$ and $\rho_{k} \rightarrow \infty$, (3.41) implies that $J\left(x_{k}^{*}\right)^{T} c\left(x_{k}^{*}\right) \geq 0$ for all $k$ large enough. Otherwise, (3.41) would eventually be violated as $\rho_{k}$ grew large. Then,

$$
\begin{equation*}
z_{*} \stackrel{\text { def }}{=} \lim _{k \in \mathcal{K}} J\left(x_{k}^{*}\right)^{T} c\left(x_{k}^{*}\right)=J_{*}^{T} c_{*} \geq 0 . \tag{3.42}
\end{equation*}
$$

All $x_{k}^{*}$ lie in a compact set, there exists some constant $L>0$ such that

$$
\begin{equation*}
\left\|x_{k}^{*}-\bar{x}\right\| \leq \frac{L \alpha_{1}}{\sqrt{n} M} \tag{3.43}
\end{equation*}
$$

where $M$ and $\alpha_{1}$ are as defined in Lemma 3.4 and $n$ is the number of elements in the vector $x_{k}^{*}$. Substituting (3.43) into (3.12) and using (3.21), we have

$$
\begin{equation*}
\left\|\widehat{g}\left(x_{k}^{*}\right)-\widehat{J}\left(x_{k}^{*}\right)^{T} y_{k}^{*}+\rho_{k} \widehat{J}\left(x_{k}^{*}\right)^{T} c\left(x_{k}^{*}\right)\right\| \leq \sqrt{n}\left\{\omega_{k}+L\left(\sigma_{k}+\omega_{k}\right)\right\} \tag{3.44}
\end{equation*}
$$

Dividing (3.44) through by $\rho_{k}$, we obtain

$$
\begin{equation*}
\left\|\frac{1}{\rho_{k}}\left(\widehat{g}\left(x_{k}^{*}\right)-\widehat{J}\left(x_{k}^{*}\right)^{T} y_{k}^{*}\right)+\widehat{J}\left(x_{k}^{*}\right)^{T} c\left(x_{k}^{*}\right)\right\| \leq \frac{\sqrt{n}\left\{\omega_{k}+L\left(\sigma_{k}+\omega_{k}\right)\right\}}{\rho_{k}} . \tag{3.45}
\end{equation*}
$$

The quantity $\widehat{g}\left(x_{k}^{*}\right)-\widehat{J}\left(x_{k}^{*}\right)^{T} y_{k}^{*}$ is bounded for the same reasons that $(a)$ and $(b)$ above are bounded. Taking limits of both sides of (3.45), $\rho_{k} \rightarrow \infty$ and $\omega_{k}, \sigma_{k} \rightarrow 0$ imply that $\widehat{J}\left(x_{k}^{*}\right)^{T} c\left(x_{k}^{*}\right) \rightarrow 0$. By continuity of $J$ and $c, \widehat{J}_{*}^{T} c_{*}=0$. Equivalently, we may write

$$
\begin{equation*}
\left[J_{*}^{T} c_{*}\right]_{j}=0 \quad \text { if } \quad\left[x_{*}\right]_{j}>0 \tag{3.46}
\end{equation*}
$$

for $j=1, \ldots, n$. Therefore (3.39), (3.42) and (3.46) together imply that $\left(x_{*}, z_{*}\right)$ satisfies conditions (3.37), as required.

Theorem 3.10 describes a useful feature of Algorithm 2. When applied to an infeasible problem, the algorithm converges to a solution of (3.36) - or at least to a first-order point. One important caveat deserves mention: if the convergence tolerance $\eta_{*}$ is small (it usually will be), Algorithm 2 may never terminate. We need to insert an additional test to provide for the possibility that (GNP) is infeasible. For example, the test could force the algorithm to exit if $\rho_{k}$ is above a certain threshold value and $\left\|c\left(x_{k}^{*}\right)\right\|$ is no longer decreasing. Any test we devise is necessarily heuristic, however; it is impossible to know for certain whether a larger value of $\rho_{k}$ would force $\left\|c\left(x_{k}^{*}\right)\right\|$ to be less than $\eta_{*}$. We discuss this point further in $\S 4.6$.
3.4. Second-order optimality. The stabilized LCL method imposes few requirements on the manner in which the LC subproblems are solved. Our implementation (see Section 4) uses MINOS or SNOPT to solve the LC subproblems. These are active-set solvers suitable for optimization problems with few expected degrees of freedom at the solution and in which only first derivatives are available. However, second derivatives might be readily available for some problems. Also, some problems are expected to have many degrees of freedom at the solution. In either case, an interior-point solver (requiring second derivatives) may be more appropriate for the solution of the subproblems.

Lemma 3.4 and Theorem 3.6 assert that iterates generated by the stabilized LCL algorithm converge to first-order KKT points. A subproblem solver that uses
second-derivatives may be able to guarantee convergence to second-order points. If we augment the convergence criteria for the solution of each subproblem to include second-order conditions, we can show that Algorithm 2 generates iterates converging to points satisfying the second-order sufficiency conditions for (GNP). The following assumption strengthens the first-order conditions (2.1).

Assumption 3.11. Let $x_{*}$ be any limit point of the sequence $\left\{x_{k}^{*}\right\}$, and let $\mathcal{K}$ be the infinite set of indices associated with that convergent subsequence. For all $k \in \mathcal{K}$ large enough, the following conditions hold at each $\left(x_{k}^{*}, y_{k}^{*}, z_{k}^{*}\right)$ : For some $\delta>0$, independent of $k$,

1. (Strict Complementarity)

$$
\begin{equation*}
\max \left(x_{k}^{*}, z_{k}^{*}\right)>\delta e ; \tag{3.47}
\end{equation*}
$$

2. (Second-Order Condition) For any $\rho \geq 0$,

$$
\begin{equation*}
p^{T} \nabla_{x x}^{2} \mathcal{L}\left(x_{k}^{*}, y_{k}^{*}, \rho\right) p>\delta \tag{3.48}
\end{equation*}
$$

for all $p \neq 0$ satisfying

$$
\begin{equation*}
J\left(x_{k}^{*}\right) p=0 \text { and }[p]_{j}=0 \text { for all } j \text { such that }\left[x_{k}^{*}\right]_{j}=0 . \tag{3.49}
\end{equation*}
$$

Condition (3.48) implies that the reduced Hessian of $\mathcal{L}$ is uniformly positive definite at all $x_{k}^{*}$.

The following result extends Theorem 3.6 to consider the case in which iterates generated by Algorithm 2 satisfy Assumption 3.11. Conn et al. [10] show a similar result for their BCL method.

Theorem 3.12. Suppose that Assumptions 1.1, 3.1, 3.2, and 3.11 hold. Let $\left\{\left(x_{k}^{*}, y_{k}^{*}, z_{k}^{*}\right)\right\}$ be the sequences of vectors generated by Algorithm 2. Let $x_{*}$ be any limit point of the sequence $\left\{x_{k}^{*}\right\}$, and let $\mathcal{K}$ be the infinite set of indices associated with that convergent subsequence. Set $\widehat{y}_{k}=\widehat{y}\left(x_{k}^{*}, y_{k}^{*}, \rho_{k}\right)$. Then

$$
\begin{equation*}
\lim _{k \in \mathcal{K}}\left(x_{k}^{*}, \widehat{y}_{k}, z_{k}^{*}\right)=\left(x_{*}, y_{*}, z_{*}\right) \tag{3.50}
\end{equation*}
$$

and $\left(x_{*}, y_{*}, z_{*}\right)$ is an isolated local minimizer of (GNP).
Proof. It follows immediately from Theorem 3.6 that (3.50) holds and that $\left(x_{*}, y_{*}, z_{*}\right)$ is a first-order KKT point for (GNP). It only remains to show that $\left(x_{*}, y_{*}, z_{*}\right)$ satisfies the second-order sufficiency conditions (see Definition 1.4).

By hypothesis, $x_{k}^{*}$ and $z_{k}^{*}$ satisfy Part 1 of Assumption 3.11 for all $k \in \mathcal{K}$. Therefore, their limit points satisfy

$$
\max \left(x_{*}, z_{*}\right) \geq \delta e>0
$$

and so $x_{*}$ and $z_{*}$ satisfy strict complementarity (Definition 1.3 ). We now show that $x_{*}$ and $y_{*}$ satisfy the second-order sufficiency conditions for (GNP).

Let $p$ be any nonzero vector satisfying (3.49) for all $k \in \mathcal{K}$ large enough. Then

$$
\begin{equation*}
p^{T} \nabla_{x x}^{2} \mathcal{L}\left(x_{k}^{*}, y_{k}^{*}, \rho_{k}\right) p=p^{T}\left(H\left(x_{k}^{*}\right)-\sum_{i=1}^{m}\left[\widehat{y}_{k}\right]_{i} H_{i}\left(x_{k}^{*}\right)\right) p \tag{3.51}
\end{equation*}
$$

for all $k \in \mathcal{K}$ large enough. Part 2 of Assumption 3.11 and (3.51) imply that

$$
\begin{equation*}
p^{T}\left(H\left(x_{k}^{*}\right)-\sum_{i=1}^{m}\left[\widehat{y}_{k}\right]_{i} H_{i}\left(x_{k}^{*}\right)\right) p>\delta, \tag{3.52}
\end{equation*}
$$

where $\delta$ is some positive constant. If we take the limit of (3.52), the continuity of $H$ and $H_{i}$ (see Assumption 1.1) and (3.50) imply that

$$
\begin{equation*}
p^{T} \nabla_{x x}^{2} \mathcal{L}\left(x_{*}, y_{*}, \rho\right) p=p^{T}\left(H\left(x_{*}\right)-\sum_{i=1}^{m}\left[y_{*}\right]_{i} H_{i}\left(x_{*}\right)\right) p \geq \delta>0 \tag{3.53}
\end{equation*}
$$

for all $\rho \geq 0$ and for all $p \neq 0$ satisfying (1.9). Therefore, $\left(x_{*}, y_{*}, z_{*}\right)$ satisfies the second-order sufficiency conditions for (GNP), as required. $\square$
4. Implementation. The practical implementation of an algorithm invariably requires many features that are not made explicit by its theory. In this section we discuss some important details of our implementation of the stabilized LCL method. The algorithm has been implemented in Matlab, version 6 [28] and is called LCLOPT. It uses the Fortran codes MINOS [30,31] and SNOPT [21] to solve the linearly constrained subproblems. We now turn our attention back to the more general problem (NP), first presented in §1.1, and leave (GNP) behind.
4.1. Problem formulation. LCLOPT does not solve (NP) directly, but rather solves the equivalent problem

$$
\begin{array}{|ll}
\hline(\mathrm{NPi}) \quad & \begin{array}{l}
\operatorname{minimize} \\
x, s \\
\\
\\
\\
\\
\\
\text { subject to }
\end{array}\binom{c(x)}{A x}-s=0, \quad l \leq\binom{ x}{s} \leq u
\end{array}
$$

The formulation of problem ( NPi ) is chosen to match the problem formulation used by SNOPT. It is also closely related to that used by MINOS. As in those methods, our implementation distinguishes between variables in the vector $x$ that appear and do not appear nonlinearly in the objective or the constraints; variables that appear only linearly are treated specially. The following discussion ignores this detail in order to keep the notation concise.

The linearly constrained subproblems corresponding to ( NPi ) take the form

$$
\begin{array}{|ccl}
\hline\left(\mathrm{ELCi}_{k}\right) & \operatorname{minimize}_{x, s, v, w} & \mathcal{L}_{k}(x)+\sigma_{k} e^{T}(v+w) \\
& \text { subject to }\binom{c_{k}+J_{k}\left(x-x_{k}\right)+v-w}{A x}-s=0, & l \leq\binom{ x}{s} \leq u, \\
& 0 \leq v, w .
\end{array}
$$

4.2. The main algorithm. The computational kernel of LCLOPT resides in the solution of each LC subproblem, and the efficiency of the implementation ultimately relies on the efficiency of the subproblem solver. The main tasks of the outer level are to form the subproblems, update solution estimates, update parameters, and test for convergence or errors.
4.3. Solving the LC subproblems. LCLOPT can use either MINOS or SNOPT to solve $\left(\mathrm{ELCi}_{k}\right)$. For linearly constrained problems, MINOS uses a reduced-gradient method, coupled with a quasi-Newton approximation of the reduced Hessian of the the problem objective. SNOPT implements a sparse SQP method and maintains a limited-memory, quasi-Newton approximation of the Hessian of the problem objective. (In both cases, the problem objective will be the objective of $\left(\mathrm{ELCi}_{k}\right)$.) For linearly constrained problems, SNOPT avoids performing an expensive Cholesky factorization
of the reduced Hessian for the quadratic programming subproblem in each of its own major iterations, and thus realizes considerable computational savings over problems with nonlinear constraints [22].

Both MINOS and SNOPT are available as libraries of Fortran 77 routines. We implemented MEX interfaces [29] written in C to make each of the routines from the MINOS and SNOPT libraries accessible from within Matlab. The subproblem solvers evaluate the nonlinear objective function (there are no nonlinear constraints in $\left.\left(\mathrm{ELCi}_{k}\right)\right)$ through a generic MEX interface, funObj.c. This routine makes calls to a Matlab routine to evaluate the nonlinear objective $\mathcal{L}_{k}$. In turn, the routine for $\mathcal{L}_{k}$ makes calls to routines (available as Matlab or MEX routines) to evaluate the original nonlinear functions $f$ and $c$.
4.4. Computing an initial point. MINOS and SNOPT both ensure that all iterates remain feasible (to within a small tolerance) with respect to the bounds and linear constraints in $\left(\mathrm{ELCi}_{k}\right)$, which includes the bounds and linear constraints in $(\mathrm{NPi})$. LCLOPT is therefore able to restrict the evaluation of the nonlinear functions $f$ and $c$ to points in the latter region. A user of LCLOPT may thus introduce bounds and linear constraints into ( NPi ) to help guard against evaluation of the nonlinear functions at points where they are not defined.

Before entering the first iteration of the stabilized LCL method, LCLOPT solves the following quadratic proximal-point (PP) problem:

$$
\begin{array}{|ll}
\hline(\mathrm{PP} 2) & \underset{x}{\operatorname{minimize}} \\
\frac{1}{2}\|x-\widetilde{x}\|_{2}^{2} \\
& \text { subject to } l \leq\binom{ x}{A x} \leq u
\end{array}
$$

where $\widetilde{x}$ is a vector provided by the LCLOPT user. The solution of (PP2) is used as the initial point $x_{0}$ for the algorithm. The objective function of the PP problem helps find an $x_{0}$ reasonably close to $\widetilde{x}$, while the constraints ensure that $x_{0}$ is feasible with respect to the bounds and linear constraints of (NPi). If (PP2) proves infeasible, $(\mathrm{NPi})$ is declared infeasible and LCLOPT exits immediately with an error message.

An alternative PP problem is based on the one-norm deviation from $\widetilde{x}$ :

$$
\begin{array}{|cl}
\hline(\mathrm{PP} 1) \quad \underset{x}{\operatorname{minimize}} & \|x-\widetilde{x}\|_{1} \\
& \text { subject to } l \leq\binom{ x}{A x} \leq u
\end{array}
$$

An advantage of (PP1) is that it can be reformulated and solved as a linear program, and its solution is therefore expected to lie on more constraint vertices. It is a correspondingly easier problem to solve for reduced-space solvers. SNOPT provides the option of solving either (PP1) or (PP2). LCLOPT can take advantage of this by initializing $x=\widetilde{x}$ and passing SNOPT the optimization problem

$$
\begin{array}{ll}
\underset{x}{\operatorname{minimize}} & 0 \\
\text { subject to } & l \leq\binom{ x}{A x} \leq u
\end{array}
$$

with the parameter Proximal Point set to either 1 or 2. (For MINOS, the constraints would have to be reformulated and a set of elastic variables introduced.)

The computational results presented in $\S 5$ were derived by using (PP2) to compute $x_{0}$. As suggested by Gill et al. [22], a loose optimality tolerance on (PP2) is used to limit the computational expense of its solution: reducing the number of iterations and (typically) the number of superbasic variables.
4.5. Early termination of the LC subproblems. The global convergence results for the stabilized LCL algorithm (cf. Lemma 3.4 and Theorem 3.6) assume that the optimality tolerances $\omega_{k}$ for the subproblems converge to 0 . This loose requirement allows much flexibility in constructing the sequence $\left\{\omega_{k}\right\}$.

The solution estimates may be quite poor during early iterations. We expect slow progress during those iterations, even if they are solved to tight optimality tolerances. A loose tolerance may help limit the computational work performed by the subproblem solver during these early iterations. Near a solution, however, we wish to reduce the optimality tolerance quickly in order to take advantage of the fast local convergence rate predicted by Theorem 3.9.

To construct the sequence $\left\{\omega_{k}\right\}$, we replace Step 1 of Algorithm 2 by

$$
\begin{align*}
\omega & \leftarrow \min \left(\omega_{k},\left\|F\left(x_{k}, y_{k}, z_{k}\right)\right\|_{\infty}^{2}\right) \\
\omega_{k+1} & \leftarrow \max \left(0.5 \omega, \omega_{*}\right), \tag{4.1}
\end{align*}
$$

where $\omega_{0}$ can be set by a user to any value between 0.5 and $\omega_{*}$. The update (4.1) guarantees that $\omega_{k} \rightarrow \omega_{*}$, as required.

Following the prescription outlined in $\S 2.2$, we fix at a small value the feasibility tolerance for satisfying the linearized constraints. The feasibility and optimality tolerances for each major iteration are passed to the subproblem solver as run-time parameters.
4.6. Detecting infeasibility and unboundedness. As discussed in §3.3, Algorithm 2 will not exit if the optimization problem is infeasible and the infeasibility tolerance $\eta_{*}$ is small. We declare ( NPi ) infeasible if at any given iteration $k, x_{k}$ is infeasible with respect to the nonlinear constraints and the penalty parameter is greater than some threshold value. In particular, at Step 10, Algorithm 2 exits and (NPi) is declared infeasible if

$$
\begin{aligned}
\max \left(\left\|\left[l_{c}-c_{k}\right]^{+}\right\|_{\infty},\left\|\left[u_{c}-c_{k}\right]^{+}\right\|_{\infty}\right) & >\eta_{*} \\
\rho_{k} & >\bar{\rho},
\end{aligned}
$$

where $l_{c}$ and $u_{c}$ are the lower and upper bounds for the nonlinear constraints and $[\cdot]^{+}$is the positive part of a vector. For the computational results in $\S 5$ the threshold value was set at $\bar{\rho}=10^{8}$.

We also need to consider the possibility that (NPi) is unbounded-i.e., that the objective function $f$ is unbounded below in the feasible region, or that $\|x\| \rightarrow \infty$. As with tests for infeasibility, any test for unboundedness must be ad hoc. We rely on the LC solver to help detect infeasibility. Problem (NPi) is declared unbounded and LCLOPT exits if the point $x_{k}$ is feasible and the LC solver reports $\left(\mathrm{ELCi}_{k}\right)$ as unbounded.
4.7. Summary of the stabilized LCL method. Following is a summary of the stabilized LCL method as implemented in LCLOPT. We assume that $\widetilde{x}$ is given and that the starting tolerances, $\omega_{0}$ and $\eta_{0}$, and parameters, $\rho_{0}$ and $\sigma_{0}$, are set.

1. Apply the LC solver to (PP1) or (PP2) to obtain a starting point $x_{0}$ that is feasible with respect to the bounds and linear constraints and reasonably
close to $\widetilde{x}$. If the PP problem is infeasible, declare (NPi) infeasible and exit. Otherwise, set $k=0$.
2. Evaluate the functions and gradients at $x_{k}$. Linearize the constraints and form $\left(\mathrm{ELCi}_{k}\right)$.
3. Apply the LC solver to $\left(\mathrm{ELCi}_{k}\right)$ with optimality tolerance $\omega_{k}$ to obtain $\left(x_{k}^{*}, \Delta y_{k}^{*}, z_{k}^{*}\right)$. Set $y_{k}^{*}=y_{k}+\Delta y_{k}^{*}$.
4. If $\left(\mathrm{ELCi}_{k}\right)$ is unbounded and $x_{k}^{*}$ is feasible, declare ( NPi ) unbounded and exit. If $\left(\mathrm{ELCi}_{k}\right)$ is unbounded and $x_{k}^{*}$ is infeasible, go to Step 8. Otherwise, continue.
5. If $x_{k}^{*}$ meets the current nonlinear feasibility threshold $\eta_{k}$, continue. Otherwise, go to Step 8.
6. Update the solution estimate: $\left(x_{k+1}, y_{k+1}, z_{k+1}\right) \leftarrow\left(x_{k}^{*}, y_{k}^{*}-\rho_{k} c\left(x_{k}^{*}\right), z_{k}^{*}\right)$. Keep the penalty parameter $\rho_{k}$ fixed and reset the elastic weight $\sigma_{k}$.
7. Test convergence: If $\left(x_{k+1}, y_{k+1}, z_{k+1}\right)$ satisfies the optimality conditions for (NPi), declare the current solution estimate optimal, return $\left(x_{k+1}, y_{k+1}, z_{k+1}\right)$, and exit. Otherwise, go to Step 9.
8. If $\rho_{k}>\bar{\rho}$, declare ( NPi ) infeasible, return $\left(x_{k}^{*}, y_{k}^{*}, z_{k}^{*}\right)$, and exit. Otherwise, discard the subproblem solution (i.e., $\left(x_{k+1}, y_{k+1}, z_{k+1}\right) \leftarrow\left(x_{k}, y_{k}, z_{k}\right)$ ), increase the penalty parameter $\rho_{k}$, and reduce the elastic weight $\sigma_{k}$.
9. Set the next nonlinear feasibility threshold $\eta_{k+1}$ and LC subproblem optimality tolerance $\omega_{k+1}$, so that $\left\{\left(\omega_{k}, \eta_{k}\right)\right\} \rightarrow\left(\omega_{*}, \eta_{*}\right)$.
10. Set $k \leftarrow k+1$. Return to Step 2 .
11. Numerical Results. This section summarizes the results of applying our implementation of the stabilized LCL method, LCLOPT, to a subset of nonlinearly constrained test problems from the COPS 2.0 [14], Hock-Schittkowski [25], and CUTE [5] test suites. Two versions of LCLOPT are applied to each test problem: The first version uses AMPL/MINOS 5.5 [18], version 19981015, to solve the sequence of linearly constrained subproblems; the second version uses SNOPT version 6.1-1(5).

We used the AMPL versions of all problems, as formulated by Vanderbei [41]. A MEX interface to the AMPL libraries makes functions and gradients available in Matlab (see Gay [20] for details on interfacing external routines to AMPL). All runs were conducted on an AMD Athlon 1700XP using 384 MB of RAM, running Linux 2.4.18.

Figure 5 shows the performance profiles, as described by Dolan and Moré [15], of the two versions of LCLOPT (the dotted and dashed lines) and MINOS (the solid line). The three charts of that figure show performance profiles for the total number of nonlinear function evaluations, minor iterations, and major iterations. All the problems selected from the COPS, Hock-Schittkowski, and CUTE test suites are included in each profile. The performance profiles describe the percentage of problems successfully solved (the vertical axes) within a factor $\tau$ of the best-performing solver (the horizonal axes).

By all measures, LCLOPT, using MINOS to solve the subproblems, successfully solves the largest proportion of problems and proves to be the most reliable method. Compared with MINOS, LCLOPT tends to require more minor iterations (a measure of total computational work) but fewer major iterations to reach a solution. We comment further on this fact in $\S 6.1$.

For all problems that can vary in the number of constraints and variables, we describe their dimensions. The following heads are used in Tables 5.1 and 5.4.


Fig. 5.1. Performance profiles. The vertical axes represent the percentage of problems successfully solved within a factor $\tau$ of the best solver. The horizontal axes are based on a log scale. Performance profiles are shown for the number of nonlinear function evaluations, minor iterations, and major iterations. The profiles include the results of the 135 selected test problems.

| Head | Dimension |
| :--- | :--- |
| $m$ | Constraints (linear and nonlinear) |
| $m_{c}$ | Nonlinear constraints |
| $n$ | Variables |
| $n_{c}$ | Variables appearing nonlinearly in $c$ |
| $n_{f}$ | Variables appearing nonlinearly in $f$ |

5.1. Default parameters. Figure 5.2 shows the options files that LCLOPT uses for the LC solvers. These are fixed for all subproblems. Separately, at each major iteration, LCLOPT sets the parameter Optimality Tolerance in MINOS and the parameter Major Optimality Tolerance in SNOPT. These are equivalent to the subproblem optimality tolerance $\omega_{k}$ ( $\S 4.5$ outlines the method for choosing this parameter).

Each test problem supplies a default starting point. This point is used as $\widetilde{x}$ in the proximal-point problem (see $\S 4.4$ ). The initial vector of multiplier estimates $y_{0}$ is set to zero.

Both MINOS and SNOPT provide the option to reuse a quasi-Newton approximation of a Hessian from a previous solve: MINOS approximates the reduced Hessian; SNOPT approximates the full Hessian. We take advantage of this feature for all iterations $k=2,3,4, \ldots$ by setting the MINOS and SNOPT options Start $=$ 'Hot'.

The parameters used by Algorithm 2 are set as follows. The upper and lower bounds of the elastic penalty parameters are $\underline{\sigma}=1$ and $\bar{\sigma}=10^{4}$. The initial elastic weight is $\sigma_{0}=10^{2}$. (Normally, LCLOPT scales this quantity by $1+\left\|y_{0}\right\|_{\infty}$, but the scaling has no effect for these test runs because $y_{0} \equiv 0$.) The penalty scaling factors are $\tau_{\rho}=100^{0.5}$ and $\tau_{\sigma}=10$. As suggested in [10], we set $\alpha=0.1$ and $\beta=0.9$. The initial penalty parameter is $\rho_{0}=10^{5 / 2} / m_{c}$, where $m_{c}$ is the number of nonlinear constraints. The final optimality and feasibility tolerances are $\omega_{*}=\eta_{*}=10^{-6}$. The initial optimality and feasibility tolerances are $\omega_{0}=10^{-3}\left(=\sqrt{\omega_{*}}\right)$ and $\eta_{0}=1$.

In all cases, default options, with the exception of Major Iterations 500 and Superbasics Limit 2000, are used for the MINOS benchmarks.

| BEGIN LCL SUBPROBLEM |  |
| :---: | ---: |
| Scale option | 0 |
| Superbasics limit | 2000 |
| Iterations | 5000 |
| Feasibility tol | $1.0 \mathrm{e}-6$ |
| END LCL SUBPROBLEM |  |
| (a) The MINOS specs file |  |


| BEGIN LCL SUBPROBLEM |  |
| :---: | ---: |
| Scale option | 0 |
| Superbasics limit | 2000 |
| Iterations | 5000 |
| Major iterations | 1000 |
| Minor iterations | 500 |
| Minor feasibility tol | $1.0 \mathrm{e}-6$ |
| Minor optimality tol | $2.5 \mathrm{e}-7$ |
| END LCL SUBPROBLEM |  |
| (b) The SNOPT specs file |  |

FIG. 5.2. The fixed optional parameters for every subproblem solve. The optimality tolerance $\omega_{k}$ is specified by LCLOPT for each $k$.
5.2. The COPS test problems. The COPS 2.0 collection [14] comprises 17 problems. Five problems are excluded for the following reasons:

- 3 problems are unconstrained: bearing, minsurf, and torsion;
- 2 problems cause system errors when called using the AMPL MEX interface: glider and marine.

The dimensions of the COPS test problems can be adjusted. In all cases, the solvers were applied to the largest version of the problem (as specified by the AMPL model) that would not cause the system to age memory to disk. Table 5.1 summarizes the dimensions of the selected problems.

Table 5.1
Dimensions: The 12 selected COPS test problems

| Problem | $m$ | $m_{c}$ | $n$ | $n_{c}$ | $n_{f}$ |
| :--- | ---: | ---: | ---: | ---: | ---: |
| camshape | 1604 | 801 | 800 | 800 | 0 |
| catmix | 1603 | 1600 | 2403 | 2403 | 0 |
| chain | 204 | 1 | 402 | 201 | 402 |
| channel | 800 | 400 | 800 | 800 | 0 |
| elec | 201 | 200 | 600 | 600 | 600 |
| gasoil400 | 4004 | 3200 | 4003 | 4003 | 202 |
| marine | 1208 | 800 | 1215 | 1215 | 344 |
| methanol | 2406 | 1800 | 2405 | 1605 | 1670 |
| pinene | 4006 | 3000 | 4005 | 2405 | 2469 |
| polygon | 1377 | 1225 | 100 | 100 | 100 |
| robot | 2414 | 2400 | 3611 | 3209 | 0 |
| rocket | 2409 | 1200 | 1605 | 1605 | 0 |
| steering | 2011 | 1600 | 2007 | 1204 | 0 |

As shown in Table 5.2, the version of LCLOPT using MINOS for the subproblems solved all 12 problems to first-order optimality. The version using SNOPT solved 11 problems to first-order optimality; the exception was robot, which it declared as having infeasible nonlinear constraints. MINOS solved 10 of the 12 problems to optimality; it declared steering an infeasible problem, and it terminated the solution of elec because of excessive iterations. Feasible points exist for all of the test problems chosen, so we consider all declarations of infeasibility to be errors.

TABLE 5.2
Summary: The 12 selected COPS test problems

|  | LCLOPT |  |  |
| :--- | ---: | ---: | ---: |
|  | (MINOS) | (SNOPT) | MINOS |
| Optimal | 12 | 11 | 10 |
| False Infeasibility |  | 1 | 1 |
| Terminated by iteration limit |  |  | 1 |
| Major iterations | 118 | 179 | 380 |
| Minor iterations | 53950 | 147518 | 61388 |
| Function evaluations | 53081 | 11014 | 63701 |

We note that different local optima appear to have been found for problems camshape, methanol, polygon, and rocket. An excessive number of minor iterations were required by LCLOPT on catmix, elec, and robot with SNOPT as its subproblem solver. Especially during early major iterations, SNOPT was unable to solve the LC subproblems to the required optimality tolerance within the 5000 iteration limit. Rather than terminate with an error message, LCLOPT forces SNOPT to keep working on the same subproblem until it returns a solution within the required optimality tolerance. In practice, a different strategy would be adopted, but our goal here is
to test the robustness of the outer iterations (the stabilized LCL method), not the robustness of the subproblem solvers.
5.3. The Hock-Schittkowski test problems. The HS test suite contains 86 nonlinearly constrained problems [25]. These are generally small and dense problems. We exclude 5 problems from this set for the following reasons:

- 3 problems are not smooth: $h s 67$, $h s 85$, and $h s 87$;
- 2 problems require external functions: $h s 68$ and $h s 69$.

Both versions of LCLOPT solved the same 80 problems to first-order optimality, but both declared hs109 infeasible. MINOS solved 80 problems to first-order optimality but declared hs93 infeasible.

Table 5.3
Summary: The 81 selected Hock-Schittkowski test problems

|  | LCLOPT |  |  |
| :--- | ---: | :---: | :---: |
|  | (MINOS) | (SNOPT) | MINOS |
| Optimal | 80 | 80 | 80 |
| False infeasibility | 1 | 1 | 1 |
| Major iterations | 654 | 648 | 1160 |
| Minor iterations | 7415 | 25290 | 10111 |
| Function evaluations | 12269 | 14712 | 27127 |

On $h s 13$, all the solvers reached different solutions. However, the linear independence constraint qualification does not hold at the solution of this problem-this violates the required assumptions for both LCLOPT and MINOS.

Recall that LCLOPT and MINOS use only first derivatives and hence may not necessarily converge to local solutions of a problem. For example, LCLOPT (in both versions) converged to a known local solution of $h s 16$, but MINOS converged to some other first-order point. In contrast, MINOS converged to the known local solutions of $h s 97$ and $h s 98$, while LCLOPT (in both versions) converged to other first-order points. Similar differences exist for problems $h s 47$ and $h s 77$.
5.4. A selection of CUTE test problems. With the select utility [5], we extracted from the CUTE test suite dated September 7, 2000, problems with the following characteristics ( $*$ is a wild-card character):

| Objective function type | : * |
| :---: | :---: |
| Constraint type | : Q O (quadratic, general nonlinear) |
| Regularity | : R (smooth) |
| Degree of available derivatives | : 1 (first derivatives, at least) |
| Problem interest | : M R (modeling, real applications) |
| Explicit internal variables | : * |
| Number of variables | : |
| Number of constraints | : |

These criteria yield 108 problems. We exclude 66 problems from this set for the following reasons:

- 33 problems do not have AMPL versions: car2, c-reload, dembo7, drugdis, durgdise, errinbar, junkturn, leaknet, lubrif, mribasis, nystrom5, orbit2, reading4, reading5, reading6, reading7, reading8, reading9, rotodisc, saromm, saro, tenbars1, tenbars2,
tenbars3, tenbars4, trigger, truspyr1, truspyr2, zamb2, zamb2-8, zamb2-9, zamb210, and zamb2-11;
- 21 problems cause system errors when evaluated either by the AMPL MEX interface or by MINOS (when invoked from AMPL): brainpc2, brainpc3, brainpc4, brainpc5, brainpc6, brainpc7, brainpc8, brainpc9, bratu2dt, cresc132, csfi1, csfi2, drcav1lq, drcav2lq, drcav3lq, kissing, lakes, porous1, porous2, trainf, and trainh;
- The AMPL versions of 12 problems are formulated with no nonlinear constraints: drcavty1, drcavty2, drcavty3, flosp2hh, flosp2hl, flosp2hm, flosp2th, flosp2tl, flosp2tm, methanb8, methanl8, and res.
The dimensions of 17 of the remaining 42 problems can be adjusted. In all cases, the solvers were applied to the largest problem versions that would not cause the system to page memory to disk. Table 5.4 summarizes the dimensions of the selected problems that can vary in size.

TABLE 5.4
Dimensions of the variable-size CUTE test problems

| Problem | $m$ | $m_{c}$ | $n$ | $n_{c}$ | $n_{f}$ |
| :--- | ---: | ---: | ---: | ---: | ---: |
| bdvalue | 1000 | 1000 | 1000 | 1000 | 0 |
| bratu2d | 4900 | 4900 | 4900 | 4900 | 0 |
| bratu3d | 512 | 512 | 512 | 512 | 0 |
| cbratu2d | 882 | 882 | 882 | 882 | 0 |
| cbratu3d | 1024 | 1024 | 1024 | 1024 | 0 |
| chandheq | 100 | 100 | 100 | 100 | 0 |
| chemrcta | 2000 | 1996 | 2000 | 1996 | 0 |
| chemrctb | 1000 | 998 | 1000 | 998 | 0 |
| clnlbeam | 1001 | 500 | 1499 | 499 | 1000 |
| hadamard | 257 | 128 | 65 | 64 | 65 |
| manne | 731 | 364 | 1094 | 364 | 729 |
| reading1 | 5001 | 5000 | 10001 | 10000 | 10000 |
| reading3 | 103 | 101 | 202 | 202 | 202 |
| sreadin3 | 5001 | 5000 | 10000 | 9998 | 9998 |
| ssnlbeam | 21 | 10 | 31 | 11 | 22 |
| svanberg | 1001 | 1000 | 1000 | 1000 | 1000 |
| ubh5 | 14001 | 2000 | 19997 | 6003 | 0 |

The version of LCLOPT using MINOS solved 36 of 42 problems to first-order optimality, while the version using SNOPT solved 34 problems to first-order optimality. MINOS solved 34 problems to first-order optimality. Table 5.5 summarizes these results. We note that LCLOPT, in one of its two versions, solves every problem except heart6, which it declares infeasible. With the exception of cresc50, LCLOPT with SNOPT does not seem to suffer (on successful solves) from excessive minor iterations resulting from subproblem restarts, as it does on the COPS problems.
6. Conclusions. The stabilized LCL method developed in this paper is a generalization of the augmented Lagrangian methods discussed in $\S 3$ and it shares the strengths of its predecessors: it is globally convergent (the BCL advantage) and it has fast local convergence (the LCL advantage). The $\ell_{1}$-penalty function brings the two together. Because the stabilized LCL method operates in a reduced space given by the linearized constraints (like the LCL method), it does not suffer from the illconditioning effects that can plague BCL methods.

Table 5.5
Summary: The 42 selected CUTE test problems

|  | LCLOPT |  |  |
| :--- | :---: | :---: | ---: |
|  | (MINOS) | (SNOPT) | MINOS |
| Optimal | 36 | 34 | 34 |
| False infeasibility | 4 | 3 | 3 |
| Terminated by iteration limit | 1 | 3 | 1 |
| Terminated by superbasics limit |  | 1 |  |
| Unbounded/badly scaled |  |  | 3 |
| Final point cannot be improved | 1 | 1 | 1 |
| Major iterations | 400 | 368 | 1149 |
| Minor iterations | 70476 | 278162 | 29021 |
| Function evaluations | 59216 | 57732 | 53069 |

6.1. Importance of early termination. The numerical results presented in $\S 5$ demonstrate that MINOS successfully solved many of the test problems using relatively few minor iterations. MINOS terminates its progress on each of its subproblems after 40 iterations (to avoid a refactorization of the current basis, which by default occurs every 50 iterations). In contrast, LCLOPT attempts to constrain the subproblem iterations by means of an initially loose optimality tolerance (we set $\omega_{0}=\sqrt{\omega_{*}}$ for the runs shown in $\S 5$ ). A potential weakness of this approach vis à vis MINOS is that there is no a priori bound on the number of subproblem iterations. MINOS's aggressive (and heuristic) strategy seems effective in keeping the total minor iteration counts low. This property is particularly important during the early major iterations, when the current solution estimates are poor.

It may be possible to emulate the MINOS strategy and still satisfy the requirement that the subproblem optimality tolerances $\omega_{k}$ converge to zero (cf. Lemma 3.4). For example, LCLOPT might truncate the subproblem solutions after a fixed number of iterations, and only gradually increase the iteration limit on successive major iterations. Especially during early major iterations, such a strategy may keep the accumulated number of subproblem iterations small. During later major iterations, the strategy would still ensure that the subproblem solver returns solutions within the prescribed tolerance $\omega_{k}$.

On the other end of the performance spectrum lies the issue of recovering LCL's fast local convergence rate under inexact solves (cf. §3.2.1). Bräuninger [7] proves that the quadratic convergence rate of Robinson's method is retained when $\omega_{k}$ is reduced at a rate $O\left(\left\|F\left(x_{k}, y_{k}, z_{k}\right)\right\|^{2}\right)$ (cf. Theorem 3.9). The first-order KKT conditions (2.1) for the LCL subproblem can be expressed as

$$
\left(\begin{array}{cc}
\nabla_{x x}^{2} \mathcal{L}_{k}\left(x_{k}\right) & J_{k}^{T}  \tag{6.1}\\
J_{k} &
\end{array}\right)\binom{p}{-y}+O\left(\|p\|^{2}\right)=\binom{-g_{k}+J_{k}^{T} y_{k}}{-c_{k}}
$$

where $p \stackrel{\text { def }}{=} x-x_{k}$, and a first-order Taylor expansion was used to derive the residual term $O\left(\|p\|^{2}\right)$. (We have ignored bound constraints for the moment. Robinson $[35,36]$ shows that the correct active set is identified by the subproblems near a solution.) The nonlinear equations (6.1) are closely related to the linear equations that would be derived from applying Newton's method to (2.1) (again, ignoring bound constraints). In that case, the theory from inexact Newton methods (Dembo et al. [13]) predicts
that the quadratic convergence rate is recovered when the residual error is reduced at the rate $O\left(\left\|F\left(x_{k}, y_{k}, z_{k}\right)\right\|\right)$. The similarity between (6.1) and the Newton equations hints at the possibility of recovering the quadratic convergence rate of the LCL and stabilized LCL methods by reducing $\omega_{k}$ at the rate $O\left(\left\|F\left(x_{k}, y_{k}, z_{k}\right)\right\|\right)$. We note, however, that stronger assumptions may be needed on the smoothness of the nonlinear functions. This issue deserves more study.
6.2. Keeping the penalty parameter small. Preliminary experimentation reveals that a small penalty parameter $\rho_{k}$ can significantly reduce the difficulty of each subproblem solve. BCL methods require that $\rho_{k}$ be larger than some threshold value $\bar{\rho}$. In contrast, LCL methods can converge when $\rho_{k} \equiv 0$ if they are started near a solution (see §3.9).

The challenge here is to find a strategy that can keep $\rho_{k}$ small or reduce it without destabilizing the method. A tentative strategy might be to reduce $\rho_{k}$ only finitely many times. This approach does not violate the hypotheses of Lemma 3.4, and may be effective in practice. A form of this strategy was used for the runs shown in $\S 5$.
6.3. A second-derivative LC solver. We prove in $\S 3.4$ that the stabilized LCL method will converge to second-order stationary points if the LC subproblems are solved to second-order points (for example, by using a second-derivative LC solver). In practice, however, a second-derivative LC solver may be most useful as a means of reducing the overall computational work required by the stabilized LCL method.

The stabilized LCL method is largely independent of the method in which its subproblems are solved. An LC solver using second derivatives is likely to require fewer iterations (and hence less computational work) for the solution of each of the subproblem. We would expect the number of required major iterations to remain constant if each subproblem solution is computed to within the prescribed tolerance $\omega_{k}$. However, we would expect to reduce the number of required major iterations if a MINOS-like strategy is used to terminate the subproblems (see $\S 6.1$ ). Over the same number of iterations, a subproblem solver using second derivatives may make more progress toward a solution than a first-derivative solver.

Any future implementation of the stabilized LCL method would ideally be flexible enough to allow for a variety of solvers to be used for the LC subproblems. The choice of the subproblem solver could then be guided by the characteristics of the optimization problem at hand. In particular, the advent of automatic differentiation makes second derivatives increasingly available for certain problem classes, e.g., within recent versions of GAMS and AMPL, and for more general functions defined by Fortran or C code, notably ADIFOR and ADIC (Bischof et al. [4, 3]). These may be used by SQP and interior methods for nonlinearly constrained (NC) problems (e.g., LOQO Vanderbei [39]). Certain theoretical challenges might be avoided, however, by developing specialized second-derivative LC solvers. Such LC solvers could be extended readily to general NC problems by incorporating them into the stabilized LCL algorithm.

## REFERENCES

[1] D. P. Bertsekas, Constrained Optimization and Lagrange Multiplier Methods, Academic Press, New York, 1982.
[2] M. J. Best, J. Bräuninger, K. Ritter, and S. M. Robinson, A globally and quadratically convergent algorithm for general nonlinear programming problems, Computing, 26 (1981), pp. 141-155.
[3] C. Bischof, A. Carle, P. Hovland, P. Khademi, and A. Mauer, ADIFOR 2.0 Users' Guide, Tech. Report 192, Mathematics and Computer Science Division, Argonne National Laboratory, 9700 S. Cass Avenue, Argonne, Illinois 60439, June 1998.
[4] C. Bischof and L. Roh, ADIC: An extensible automatic differentiation tool for ANSI-C, Software: Practice and Experience, 27 (1997), pp. 1427-1456.
[5] I. Bongartz, A. R. Conn, N. Gould, and P. Toint, CUTE: Constrained and unconstrained testing environment, ACM Trans. Math. Software, 21 (1995), pp. 123-160.
[6] J. Bräuninger, A modification of Robinson's algorithm for general nonlinear programming problems requiring only approximate solutions of subproblems with linear equality constraints, in Optimization Techniques. Proceedings of the 8th IFIP Conference, Part 2, Würzburg, 1977, pp. 33-41.
[7] _ , A globally convergent version of Robinson's algorithm for general nonlinear programming problems without using derivatives, J. Opt. Theory Applic., 35 (1981), pp. 195-216.
[8] A. Brooke, D. Kendrick, and A. Merraus, GAMS: A User's Guide, Scientific Press, South San Francisco, 1988.
[9] A. R. Conn, N. I. M. Gould, A. Sartenaer, and P. L. Toint, Convergence properties of an augmented Lagrangian algorithm for optimization with a combination of general equality and linear constraints, SIAM J. Optim., 6 (1996), pp. 674-703.
[10] A. R. Conn, N. I. M. Gould, and P. L. Toint, A globally convergent augmented Lagrangian algorithm for optimization with general constraints and simple bounds, SIAM J. Numer. Anal., 28 (1991), pp. 545-572.
[11] ——, LANCELOT: A Fortran Package for Large-Scale Nonlinear Optimization (Release A), Springer-Verlag, Berlin, 1991.
[12] —— Trust-Region Methods, MPS-SIAM Series on Optimization, SIAM Publications, Philadelphia, 2000.
[13] R. S. Dembo, S. C. Eisenstat, and T. Steihaug, Inexact Newton methods, SIAM J. Numer. Anal., 19 (1982), pp. 400-408.
[14] E. D. Dolan and J. J. Moré, Benchmarking optimization software with COPS, Tech. Report ANL/MCS-246, Mathematics and Computer Science Division, Argonne National Laboratory, 9700 South Cass Avenue, Argonne, Illinois 60439, 2000. Revised January 2, 2001.
[15] -, Benchmarking optimization software with performance profiles, Tech. Report ANL/MCS-P861-1200, Mathematics and Computer Science Division, Argonne National Laboratory, 9700 South Cass Avenue, Argonne, Illinois 60439, 2001.
[16] R. Fletcher, An $\ell_{1}$ penalty method for nonlinear constraints, in Numerical Optimization, P. T. Boggs, R. H. Byrd, and R. B. Schnabel, eds., Boulder, Colorado, June 1984, SIAM, pp. 26-40.
[17] - Practical Methods of Optimization, John Wiley and Sons, New York, second ed., 1987.
[18] R. Fourer, D. M. Gay, and B. W. Kernighan, AMPL: A Modeling Language for Mathematical Programming, Scientific Press, San Francisco, 1993.
[19] M. P. Friedlander, A globally convergent linearly constrained Lagrangian method for nonlineary constrained optimization, PhD thesis, Stanford University, Stanford, California, August 2002.
[20] D. M. GAY, Hooking your solver to AMPL, Tech. Report 97-4-06, Computing Sciences Research Center, Bell Laboratories, Murray Hill, New Jersey, 07974, April 1997.
[21] P. E. Gill, W. Murray, and M. A. Saunders, User's guide for SNOPT 5.3: A fortran package for large-scale nonlinear programming, Tech. Report 98-1, Systems Optimization Laboratory, Department of Operations Research, Stanford University, Stanford, California 94305, May 1997.
[22] - SNOPT: An SQP algorithm for large-scale constrained optimization, SIAM J. Optim., 12 (2002), pp. 979-1006.
[23] P. E. Gill, W. Murray, and M. H. Wright, Practical Optimization, Academic Press, San Diego, California, 1981.
[24] M. R. Hestenes, Multiplier and gradient methods, J. Opt. Theory Applic., 4 (1969), pp. 303320.
[25] W. Hock and K. Schittkowski, Test Examples for Nonlinear Programming Codes, Lecture Notes in Economics and Mathematical Systems 187, Springer-Verlag, Berlin, Heidelberg, and New York, 1981.
[26] D. G. Luenberger, Linear and Nonlinear Programming, Addison-Wesley, second ed., 1984.
[27] O. L. Mangasarian, Nonlinear Programming, McGraw-Hill, New York, 1969.
[28] MathWorks, MATLAB User's Guide, The MathWorks, Inc., Natick, Massachussetts, 1992.
[29] - MATLAB: External Interfaces, Natick, Massachussetts, 1995.
[30] B. A. Murtagh and M. A. Saunders, Large-scale linearly constrained optimization, Math.

Prog., 14 (1978), pp. 41-72.
[31] B. A. Murtagh and M. A. Saunders, A projected Langrangian algorithm and its implementation for sparse nonlinear constraints, Mathematical Programming Study, 16 (1982), pp. 84-117.
[32] J. Nocedal and S. J. Wright, Numerical Optimization, Springer-Verlag, New York, 1999.
[33] J. M. Ortega and W. C. Rheinboldt, Iterative Solutions of Nonlinear Equations in Several Variables, Academic Press, London, 1970.
[34] M. J. D. Powell, A method for nonlinear constraints in minimization problems, in Optimization, R. Fletcher, ed., Academic Press, London and New York, 1969, ch. 19.
[35] S. M. Robinson, A quadratically-convergent algorithm for general nonlinear programming problems, Math. Prog., 3 (1972), pp. 145-156.
[36] - Perturbed Kuhn-Tucker points and rates of convergence for a class of nonlinearprogramming algorithms, Math. Prog., 7 (1974), pp. 1-16.
[37] J. B. Rosen, Two-phase algorithm for nonlinear constraint problems, in Nonlinear Programming 3, O. Mangasarian, R. Meyer, and S. Robinson, eds., New York, 1978, Academic Press, pp. 97-124.
[38] J. B. Rosen and J. Kreuser, A gradient projection algorithm for non-linear constraints, in Numerical Methods for Nonlinear Optimization, F. A. Lootsma, ed., Academic Press, London, 1972, pp. 297-300.
[39] D. Shanno and R. Vanderbei, An interior-point algorithm for nonconvex nonlinear programming, Computational Optimization and Applications, 13 (1999), pp. 231-252.
[40] G. Van Der Hoek, Asymptotic properties of reduction methods applying linearly equality constrained reduced problems, Mathematical Programming Study, 16 (1982), pp. 162-189.
[41] R. Vanderbei, Benchmarks for nonlinear optimization. http://www.princeton.edu/~rvdb/bench.html, December 2002.

The submitted manuscript has been created by the University of Chicago as Operator of Argonne National Laboratory ("Argonne") under Contract No. W-31-109-ENG-38 with the U.S. Department of Energy. The U.S. Government retains for itself, and others acting on its behalf, a paid-up, nonexclusive, irrevocable worldwide license in said article to reproduce, prepare derivative works, distribute copies to the public, and perform publicly and display publicly, by or on behalf of the Government.


[^0]:    *Mathematics and Computer Science Division, Argonne National Laboratory, 9700 S. Cass Avenue, Argonne, IL 60439-4844 (michael@mcs.anl.gov). This work was supported by the U.S. National Science Foundation grant CCR-9988205 and by the Mathematical, Information, and Computational Sciences Division subprogram of the Office of Advanced Scientific Computing Research, U.S. Department of Energy contract W-31-109-Eng-38.
    ${ }^{\dagger}$ Department of Management Science and Engineering, Stanford University, Stanford, CA 943054026 (saunders@stanford.edu). This work was supported by the U.S. National Science Foundation grant CCR-9988205, and the U.S. Office of Naval Research grants N00014-96-1-0274 and N00014-02-1-0076.

