Numerical Optimization 1984

edited by
Paul T. Boggs
National Bureau of Standards
Richard H. Byrd and Robert B. Schnabel
University of Colorado

siam Philadelphia/1985

Proceedings of the SIAM Conference
on Numerical Optimization
Boulder, Colorado
June 12–14, 1984

Supported in part by the
Air Force Office of Scientific Research,
Army Research Office,
National Science Foundation,
and Office of Naval Research.

NONLINEAR PROGRAMMING, EXACT PENALTY FUNCTIONS AND PROJECTION TECHNIQUES FOR NON-SMOOTH FUNCTIONS*

ANDREW R. CONN†

Abstract. We present a personal overview of various approaches to solving nonlinear programs with nonlinear constraints that make use of the l_1 exact penalty function.

The advantages, disadvantages and related remaining difficulties of these approaches will be considered.

Finally some recent research and extensions are given.

1. Introduction. Most of this article will be devoted to considerations of the problem

$$\text{NLP}: \quad \underset{x \in \mathbf{R}^n}{\text{minimize}} \; f(x)$$

$$\text{subject to} \quad \phi_i(x) \geq 0, \quad i \in I$$

$$\psi_i(x) = 0, \quad i \in E$$

where I and E are index sets and the f, ϕ_i's and ψ_i's are assumed to be twice continuously differentiable.

We will predominately be concerned with small and medium problems that are well-scaled.

It is instructive to begin by considering the state-of-the-art for unconstrained optimization.

At least for small and medium problems it is reasonable to claim that we currently have algorithms that are both robust and efficient.

For a useful, up-to-date and lucid review of the status of unconstrained optimization the reader is referred to Schnabel [1982]. The salient points that I would like to make here are:

1) The algorithms and the second-order sufficiency conditions for unconstrained optimization closely match.

2) If the Hessian is available one uses a (modified) Newton method. Otherwise secant approximations are usually preferred with BFGS (Broyden [1970], Fletcher [1970], Goldfarb [1970], Shanno [1970]) being the most

*This research was partially supported by Natural Sciences and Engineering Research Council of Canada grants A-8639. This paper was typeset using software developed at Bell Laboratories and the University of California at Berkeley. Final copy was produced on Autologics APS-μ5 typesetter.

† Dept. of Comp. Sci., University of Waterloo, Waterloo, Ontario, Canada N2L 3G1.

popular. However, finite difference approximations to the Hessian are also often appropriate.

3) Convergence is global (that is, convergence from an arbitrary starting point) with an asymptotic rate of convergence which is normally Q-quadratic, in the case of (modified) Newton, or Q-superlinear, in the case of secant methods.

4) Finally, it is worthwhile mentioning that we are beginning to have useful models for singular problems, eg. Schnabel and Frank [1983].

By the way of contrast let us consider the state-of-the-art for nonlinearly constrained optimization.

We first note that there are no inherent difficulties with the theory. Moreover, it is constructive to take the point of view that the theoretical results concerning optimality conditions are derived via the implicit function theorem. Thus one implicitly determines a reduced space and invokes the theory previously obtained for unconstrained optimization and applies it to the resulting unconstrained problem in the reduced space. Consequently, it is the *projected* Hessian and the *projected* gradient that give the requisite second and first-order conditions.

Thus, in a certain sense the constrained problem may be considered easier than the unconstrained problem, since each linearly independent active constraint corresponds to a loss of one degree of freedom. This is obviously likely to result in simplifications in the extreme cases of n linearly independent active constraints or linear equality constraints. However, more generally, in the case of nonlinear constraints, one does not know the true reduced space until one is at the solution *even if the correct active set is identified.*

The preceding comments explain why *linearly* constrained optimization is considerably simpler than nonlinear constrained problems. Indeed, it is reasonable to say that the state-of-the-art of linearly constrained optimization is closely comparable to that of unconstrained optimization.

By contrast one might enquire as to whether we are asking too much in desiring a comparable status in the case of nonlinear constraints.

We will attempt to consider this question in some detail below. In doing so, we will look at the state-of-the-art, concentrating on the l_1 exact penalty function.

2. Exact Penalty Functions and NLP. We begin this section with the following quote from the 11th International Symposium on Mathematical Programming, Bonn 1982. (Fletcher [1983], p. 89).

"Exact penalty functions can be used in two ways; either to provide a transformation of the nonlinear programming problem to an unconstrained minimization problem or as a criterion function (i.e. merit function) for use with other direct methods for nonlinear programming."

It is my contention that the second usage is a mistake - it is successful only in so far as it comes close to the first usage. It is hoped that the arguments that follow will justify this claim to the reader.

The currently most popular global methods for nonlinearly constrained nonlinear optimization are the so-called sequential quadratic programming (SQP) methods, with the l_1 exact penalty function for a merit function and (approximate) line searches.

In these methods the search direction at iteration k is determined from the quadratic programming problem

$$\text{SQP} : \underset{d}{\text{minimize}}\ f(x^k) + \nabla f(x^k)^T d + 1/2\, d^T W^k d$$

$$\text{subject to} \quad \phi_i(x^k) + \nabla \phi_i (x^k)^T d \geq 0, \quad i \in I,$$

$$\psi_i(x^k) + \nabla \psi_i (x^k)^T d = 0, \quad i \in E,$$

where W^k is the Hessian of the Lagrangian function, or some approximation to it.

The derivation of the quadratic program is based upon the following observations in the case where all the constraints are equality constraints.

Let

$$L(x, \lambda) = f(x) - \Psi^T \lambda,$$

be the Lagrangian function, with

$$\Psi(x) = [\psi_1(x)\ \psi_2(x) \cdots \psi_t(x)]^T,$$

$$\lambda = [\lambda_1, \lambda_2 \cdots \lambda_t]^T,$$

and

$$E = \{1,2, \ldots, t\}.$$

The method of Lagrange implies one must solve the system of $n + t$ equations in $n + t$ unknowns, x and λ, given by

$$\nabla_{[x\ \lambda]} L(x, \lambda) = 0.$$

Newton's method for solving such a system gives

$$\begin{bmatrix} W^k & -A^k \\ -[A^k]^T & 0 \end{bmatrix} \begin{bmatrix} \delta x^k \\ \delta \lambda^k \end{bmatrix} = -\begin{bmatrix} \nabla_x L(x^k, \lambda^k) \\ -\Psi(x^k) \end{bmatrix}, \tag{2.1}$$

where

$$W^k = \nabla^2_{xx} L(x^k, \lambda^k),$$

$$\delta x^k = x^{k+1} - x^k,$$

$$\delta \lambda^k = \lambda^{k+1} - \lambda^k,$$

and

$$A^k = [\nabla \psi_1(x^k)\ \nabla \psi_2(x^k) \cdots \nabla \psi_t(x^k)].$$

Now, let us consider the quadratic programming problem

$$EQP1 : \underset{d}{\text{minimize}} \quad d^T \nabla f(x^k) + d^T W^k d$$

$$\text{subject to} \quad [A^k]^T d + \Psi(x^k) = 0\,.$$

At a Kuhn-Tucker point of EQP1,

$$W^k d + \nabla f(x^k) = A^k \mu.$$

Thus, identifying δx^k with d and $\lambda^k + \delta\lambda^k$ (equals λ^{k+1}), with μ, we are able to determine an iteration of Newton's method for solving the Lagrangian system directly from the EQP1.

However, whereas the system of equations (2.1) does not generalise naturally in the case of inequality constraints, one readily obtains SQP as a generalisation of EQP1.

Any SQP approach must be able to consider the following (see for example Fletcher [1983] or Powell [1983]),

i) infeasible or unbounded quadratic subproblems,
ii) unbounded multiplier approximations,
iii) unbounded Hessian approximations,
iv) the Maratos effect - that is, a stepsize of one produces ascent on the l_1 merit function arbitrarily close to the solution, x^*, of the original constrained problem.

It is my opinion that all these complications are a manifestation of the fact that the model problem for the search direction does not come directly from the merit function, but is more closely related to the Lagrangian system.

Recently, much work on SQP methods has been devoted to overcoming the above mentioned difficulties.

Methods to overcome the Maratos effect include the second-order correction (see, for example Coleman and Conn [1982b]) and the Watchdog Technique of Chamberlain et al (Chamberlain, Lemaréchal, Pedersen and Powell [1982]).

Methods to overcome inconsistent quadratic subproblems include relaxations of the constraints, (see, for example, Biggs [1975] and Murray and Wright [1978] and the approach of Fletcher, (Fletcher [1985]).

Fletcher's method avoids most of the difficulties mentioned above. It is useful to consider his development as follows.

Essentially, instead of using a (possibly unobtainable) solution to the quadratic subproblem, SQP, to determine a search direction, followed by a line search on the l_1 penalty function, Fletcher uses an l_1 penalty formulation for the quadratic programming problem with a trust region. Thus in a certain sense the merit function and the model problem are more compatible and, for example, since one is using a penalty function approach to solve the quadratic programming problem, it is not essential that the constraints be consistent.

However, one should note that in the sense I will describe below, the quadratic programming problem is still the "wrong" subproblem. Furthermore, the trust region is updated on the basis of the overall l_1 merit function, not the l_1 function for the subproblem. The reader is referred to Fletcher [1985] for more details.

Presumably, any method that emulates the theoretical optimality conditions for nonlinearly constrained optimization and is comparable with the techniques available for unconstrained optimization will, asymptotically at least, make use of (approximations to) the projected Hessian of the Lagrangian along the null space associated with the active constraint gradients. Furthermore, at least for non-singular problems, the algorithm must be globally convergent with an asymptotic rate that is better than linear.

As a consequence, it is clear that *asymptotically* the curvature of the active constraints is crucial. From the sequential quadratic programming point of view an equivalent statement is that the Hessian term of the quadratic programming subproblem is crucial. SQP methods use $W^k = \nabla^2_{xx} L(x^k, \lambda^k)$. In the light of the comments just made, this is reasonable in a neighbourhood of x^* provided good multiplier estimates, λ^k, are available. However, it seems sensible to ask how reasonable it is to use the same approximation globally?

We will now give details of a method that derives its search direction directly from the l_1 merit function and, although *asymptotically* it approximates the projected Hessian of the Lagrangian in the neighbourhood of a solution, globally it uses quite different approximations, in general.

3. The Method of Coleman and Conn. The basis of the method described in Coleman and Conn [1982a] and [1982b] is the following.

Consider the problem NLP with corresponding l_1 exact penalty function $p(x, \mu)$, defined by $p(x, \mu) = \mu f(x) - \sum_{i \in I} \min \{0, \phi_i(x)\} + \sum_{i \in E} | \psi_i(x) |$.

Suppose, without loss of generality that $I = \{1, \ldots, t_i\}$ and that $E = \{1, \ldots, t_e\}$. Furthermore, assume that at the point y, for given, small positive ε,

$$| \phi_i(y) | < \varepsilon \quad i \in A_i \subseteq I$$

$$| \psi_i(y) | < \varepsilon \quad i \in A_e \subseteq E$$

with A_i and A_e maximal, in the sense that $| \phi_i(y) | > \varepsilon$ whenever $i \notin A_i$, and $| \psi_i(y) | > \varepsilon$ whenever $i \notin A_e$. Such constraints are termed near-active, and we note that they are ε-feasible.
Clearly

$$p(x, \mu) = r(x) + \sum_{i \in A_e} | \psi_i(x) | - \sum_{i \in A_i} \min (0, \phi_i(x)),$$

where

$$r(x) = \mu f(x) + \sum_{i \notin A_e} | \psi_i(x) | - \sum_{i \notin A_i} \min (0, \phi_i(x)).$$

The first observation we would like to make is that $r(x)$ is *differentiable* in a neighbourhood of y. In fact, under the assumption that f, ψ_i's and ϕ_i's are twice continuously differentiable, $r(x)$ is twice continuously differentiable in a neighbourhood of y. Thus a second-order Taylor's expansion of $r(x)$ exists about y and is a valid model. Consequently, we consider the following nonlinear subproblem.

$$\text{QPQC:} \quad \underset{d}{\text{minimize}}\ r(y) + d^T \nabla_x r(y) + 1/2\, d^T \nabla^2_{xx} r(y)\, d$$

$$\text{subject to } \psi_i(y) + d^T \nabla_x \psi_i(y) + 1/2\, d^T \nabla^2_{xx} \psi_i(y)\, d = \psi_i(y), \quad i \in A_e,$$

$$\phi_i(y) + d^T \nabla_x \phi_i(y) + 1/2\, d^T \nabla^2_{xx} \phi_i(y)\, d = \phi_i(y), \quad i \in A_i.$$

We remark that the constraints signify that the near-active constraints do not change *up to second-order* on moving from y to $y + d$.

Thus solving QPQC is indeed a reasonable extension of a Newton method, applied to the unconstrained minimization of a twice continuously differentiable

function, to the piecewise differentiable function $p(x, \mu)$. Moreover, as we will see shortly, the necessary asymptotic conditions for a second-order nonlinearly constrained nonlinear programming algorithm are maintained. Furthermore it is not difficult to anticipate the extension to secant type methods.

We remark that the subproblem QPQC enables us to model *directly* the l_1 penalty function, locally, up to second-order, in the sense that in a direction d that satisfies the given (quadratic) constraints, the second-order change in p is indeed given by $d^T \nabla r + 1/2\, d^T \nabla^2 r\, d$.

Unfortunately, it would appear that the problem of minimizing a quadratic function subject to quadratic constraints is almost as difficult as solving the general NLP problem. Consequently, it is necessary to make further observations.

In a standard way (see, for example, Gill, Murray and Wright [1981]), assuming full-rank, let us define a matrix of "near-active" constraint gradients, A by

$$A = [\nabla \gamma_1(y), \ldots, \nabla \gamma_t(y)]$$

where

$$\gamma_j = \phi_{k_j}, \quad k_j \in A_i$$
$$= \psi_{k_j}, \quad k_j \in A_e$$

and

$$|A_i| + |A_e| = t.$$

We also require an associated $n \times (n - t)$ matrix Z satisfying

$$A^T Z = 0,$$
$$Z^T Z = I_{n-t}.$$

Thus the columns of Z form an orthonormal basis for the space orthogonal to the space spanned by the gradients of the near-active constraints.

We will now solve alternative problems to QPQC depending upon whether the projected "gradient" of $p(x, \mu)$, $Z^T \nabla r$ is "small" or not. Note that $Z^T \nabla r$ is an *exact* representation of the first-order change of p along the space spanned by the columns of Z, even though p need not be differentiable over the entire space.

In the case where $Z^T \nabla r$ is not considered small QPQC is replaced by the straightforward quadratic programming problem that is obtained by ignoring the curvature terms in the constraints, namely

$$\text{QP1:} \quad \underset{d}{\text{minimize}}\ r(y) + d^T \nabla r(y) + 1/2\, d^T \nabla^2 r(y)\, d$$

$$\text{subject to } d^T \nabla \gamma_i(y) = 0, \quad i = 1, \ldots, t.$$

The replacement is justified with the following remarks.

$Z^T \nabla r$ "large" implies that the first-order change of the penalty function alc
the space orthogonal to the gradients of the active constraints is large in
neighbourhood of the point y. Consequently, the curvatures of the active constr
are not significant. That is, they are "swamped" by the changes in the reduced
- it is easy to obtain sufficient decrease of p. Sufficient decrease is all t

function, to the piecewise differentiable function $p(x, \mu)$. Moreover, as we will see shortly, the necessary asymptotic conditions for a second-order nonlinearly constrained nonlinear programming algorithm are maintained. Furthermore it is not difficult to anticipate the extension to secant type methods.

We remark that the subproblem QPQC enables us to model *directly* the l_1 penalty function, locally, up to second-order, in the sense that in a direction d that satisfies the given (quadratic) constraints, the second-order change in p is indeed given by $d^T \nabla r + 1/2\, d^T \nabla^2 r\, d$.

Unfortunately, it would appear that the problem of minimizing a quadratic function subject to quadratic constraints is almost as difficult as solving the general NLP problem. Consequently, it is necessary to make further observations.

In a standard way (see, for example, Gill, Murray and Wright [1981]), assuming full-rank, let us define a matrix of "near-active" constraint gradients, A by

$$A = [\nabla \gamma_1(y), \ldots, \nabla \gamma_t(y)]$$

where

$$\gamma_j = \phi_{k_j}, \quad k_j \in A_i$$
$$= \psi_{k_j}, \quad k_j \in A_e$$

and

$$|A_i| + |A_e| = t.$$

We also require an associated $n \times (n - t)$ matrix Z satisfying

$$A^T Z = 0,$$
$$Z^T Z = I_{n-t}.$$

Thus the columns of Z form an orthonormal basis for the space orthogonal to the space spanned by the gradients of the near-active constraints.

We will now solve alternative problems to QPQC depending upon whether the projected "gradient" of $p(x, \mu)$, $Z^T \nabla r$ is "small" or not. Note that $Z^T \nabla r$ is an *exact* representation of the first-order change of p along the space spanned by the columns of Z, even though p need not be differentiable over the entire space.

In the case where $Z^T \nabla r$ is not considered small QPQC is replaced by the straightforward quadratic programming problem that is obtained by ignoring the curvature terms in the constraints, namely

$$\text{QP1:} \quad \underset{d}{\text{minimize}}\ r(y) + d^T \nabla r(y) + 1/2\, d^T \nabla^2 r(y)\, d$$
$$\text{subject to } d^T \nabla \gamma_i(y) = 0, \quad i = 1, \ldots, t.$$

The replacement is justified with the following remarks.

$Z^T \nabla r$ "large" implies that the first-order change of the penalty function along the space orthogonal to the gradients of the active constraints is large in the neighbourhood of the point y. Consequently, the curvatures of the active constraints are not significant. That is, they are "swamped" by the changes in the reduced space - it is easy to obtain sufficient decrease of p. Sufficient decrease is all that is

required, since, clearly, one is not in any asymptotic region by virtue of the fact that the projected gradient is not small.

One could equally justify dropping the $\nabla^2 r$ term (i.e. replacing QP1 by an LP) and indeed in certain contexts (e.g. large sparse problems) this might be reasonable. However, our intent is to produce a quadratic model function in a certain reduced space. What is being said here is that the ideal reduction is on a nonlinear manifold but, owing to the nature of the projected gradient, a linear approximation to the manifold will suffice.

We note that at iteration k of our intended algorithm (which corresponds to $y = x^k$*), the second-order terms of SQP,* W^k *are not equivalent to* $\nabla^2 r(x^k)$*, the second-order terms in QP1, unless i)* $\lambda^k = 0$ *and ii)* x^k *is* ε*-feasible for NLP.*

If $Z^T \nabla r$ is small a different approach is taken. We first consider the fact that at a Kuhn-Tucker point for QPQC, x^k, say, there exist multipliers λ_i such that

$$\nabla r(x^k) + \nabla^2 r(x^k)\, d = \sum_{i \in A_u} \lambda_i \, [\nabla \gamma_i(x^k) + \nabla^2 \gamma_i(x^k)\, d] \tag{3.1}$$

where $A_u = A_i \cup A_e$.

[Note that equation (3.1) could be considered to be an approximation to

$$\nabla r(x^k + d) = \sum_{i \in A_u} \lambda_i \nabla \gamma_i(x^k + d) \qquad].$$

We next observe that $Z^T \nabla r$ being small is equivalent to being in the neighbourhood of a stationary point in the restricted space determined by the near-active constraints. Thus one may obtain a "good" approximation for λ_i by solving $\nabla r(x^k) = \sum_{i \in A_u} \lambda_i \nabla \gamma_i(x^k)$ in the least squares sense. Using the usual active set strategy (see, for example Gill, Murray and Wright [1981]), one may conclude that a λ_i "out-of-kilter" - i.e. $\lambda_i < 0 \;\; i \in A_i$, implies that one has not correctly identified the active set [or at least, one is far from the stationary point, since indeed, for a nonlinear problem, one may have the right active constraints but the current gradient values may be far from the values at a stationary point]. In this case, there is essentially one significant degree of freedom obtained by "dropping" a single active constraint corresponding to an out-of-kilter multiplier λ_j, say; we take the resulting projected gradient direction. [This is equivalent to replacing QPQC with the linear programming problem

$$\begin{aligned} &\underset{\|d\|=1}{\text{minimize}} && r(x^k) + d^T \nabla r(x^k) \\ &\text{subject to} && d^T \nabla \gamma_i(x^k) = 0, \quad i \in A_{u/\{j\}}, \; d^T \nabla \phi_j(x^k) > 0 \;]. \end{aligned}$$

Otherwise, we have a good approximation $\hat{\lambda}_i$ for the λ_i and all the λ_i's ($\hat{\lambda}_i$'s) are in-kilter. Thus (3.1) may be replaced, approximately, by

$$\nabla r(x^k) + [\nabla^2 r(x^k) - \sum_{i \in A_u} \hat{\lambda}_i \nabla^2 \gamma_i(x^k)]\, d = \sum_{i \in A_u} \lambda_i \nabla \gamma_i(x^k). \tag{3.2}$$

But, (3.2) is satisfied by a Kuhn-Tucker point of the quadratic program

$$\begin{aligned} \text{QP2:} \quad &\underset{d}{\text{minimize}} && r(x^k) + d^T \nabla r(x^k) + 1/2\, d^T [\nabla^2 r(x^k) - \sum_{i \in A_u} \hat{\lambda}_i \nabla^2 \gamma_i(x^k)]\, d \\ &\text{subject to} && d^T \nabla \gamma_i(x^k) = 0, \quad i \in A_u. \end{aligned}$$

We note that

a) $W^k \neq \nabla^2 r(x^k) - \sum_{i \in A_u} \hat{\lambda}_i \nabla^2 \gamma_i(x^k)$, in general.

b) We use multiplier estimates $\hat{\lambda}_i$ only when good estimates are available [- otherwise they are not necessary].

c) In the asymptotic region, by definition, there are no clearly violated constraints (i.e. $|\psi_i(x^k)| > \varepsilon$ or $\phi_i(x^k) < -\varepsilon$), $|\,|Z_k^T \nabla r(x^k)|\,|$ is small (where the subscript k on Z indicates that Z is determined at x^k), and W^k is equivalent to $\nabla^2 r(x^k) - \sum_{i \in A_u} \hat{\lambda}_i \nabla^2 \gamma_i(x^k)$ with *the additional proviso that the $\hat{\lambda}_i$'s are likely to be good approximations to the true multipliers.* [We use the term 'equivalent to' to denote the fact that in general the quadratic programming multipliers are not identical with the multipliers obtained from the least squares estimates, except in the limit].

In the neighbourhood of a stationary point [- projected gradient small, multiplier estimates in-kilter] we would like the near-active constraints to be active.

In other words, we require that

$$\gamma_i(x^k + d^k + v^k) = 0, \quad i \in A_u. \tag{3.3}$$

Approximating (3.3) by

$$\gamma_i(x^k + d^k) + \nabla \gamma_i(x^k)^T v^k = 0, \quad i \in A_u$$

we choose

$$v^k = -A_k^{\dagger} \Gamma(x^k + d^k), \tag{3.4}$$

where $A_k^{\dagger}$ denotes the pseudo-inverse of the matrix A evaluated at x^k and

$$\Gamma(y) = [\gamma_1(y), \gamma_2(y), \ldots, \gamma_t(y)]^T, \; A_u = [1,2,\ldots,t].$$

[Remark: the *gradients* of the constraints are not re-evaluated at $x^k + d^k$.]

We note that ultimately Newton steps will be taken, in the sense that $d^k + v^k$ gives descent on $p(x,\mu)$. That is, the Maratos effect does not take place.

For more details, the interested reader is urged to consult the references Coleman and Conn [1982a] and [1982b] already cited above.

Finally, for those deeply attached to quadratic programs we note that the asymptotic (Newton) step given by $x^{k+1} = x^k + \delta^k$, where $\delta^k = d^k + v^k$, with d^k defined by QP2 and v^k defined by (3.4), is the unique solution to the quadratic programming problem

$$\text{QP3:} \quad \underset{\delta}{\text{minimize}} \; r(x^k) + \delta^T \nabla r(x^k) + 1/2\,(\delta^T Z_k)(Z_k^T B_k Z_k)(Z_k^T \delta)$$

$$\text{subject to } A^T \delta + \Gamma(x^k - Z_k(Z_k^T B_k Z_k)^{-1} Z_k^T \nabla r(x^k)) = 0,$$

where

$$B_k = \nabla^2 r(x^k) - \sum_{i \in A_u} \hat{\lambda}_i \nabla^2 \gamma_i(x^k), \quad \text{and } Z_k^T B_k Z_k$$

is positive definite. Moreover, we note that under the second-order sufficiency

conditions in the neighbourhood of a local constrained optimum, our *projected* Hessian $Z_k^T B_k Z_k$ is guaranteed to be positive definite.

4. Discussion. In attempting to construct a method for nonlinearly constrained optimization that is ideal, in the sense that one is able to do as well as unconstrained techniques, we require, conceptually at least, direct reduction.

The method of Lagrange can be derived from (an implicit) reduction. However, commonly used approximations based upon the method of Lagrange are quite distinct from approximations one might anticipate from direct reduction.

This is most easily appreciated by considering a particular example.

Example 4.1

$$\underset{x=[u,v,w]}{\text{minimize}}\ u + v^2 + w^2$$

$$\text{subject to}\ -u + 3v - vw = 1,$$

This example is chosen because direct substitution gives rise to an unconstrained quadratic problem, although the constrained problem is a QPQC.

Thus using the explicit reduction

$$u = 3v - vw - 1,$$

one obtains the equivalent *unconstrained* problem

$$\underset{(v,w)}{\text{minimize}}\ 3v + v^2 - vw + w^2 - 1,$$

with corresponding Hessian matrix,

$$H = \begin{pmatrix} 2 & -1 \\ -1 & 2 \end{pmatrix}$$

with eigenvalues 1 and 3, and gradient $(g_v, g_w) = (3 + 2v - w, -v + 2w)$.

Thus, one step of Newton's method, from any starting point produces the unconstrained minimum, $(v, w) = (-2, -1)$, from which it can be deduced that the optimum value of u and λ is, $u = -9, \lambda = -1$, with corresponding objective function value -4. For, example, starting at $(v^0, w^0) = (0, 0)$

$$\begin{pmatrix} v^1 \\ w^1 \end{pmatrix} = -H^{-1} \begin{pmatrix} g_v \\ g_w \end{pmatrix} = - \begin{pmatrix} 2/3 & 1/3 \\ 1/3 & 2/3 \end{pmatrix} \begin{pmatrix} 3 \\ 0 \end{pmatrix} = \begin{pmatrix} -2 \\ -1 \end{pmatrix}.$$

We note that, referring to the original constrained problem, although the Hessian of the Lagrangian is not positive definite, the (orthogonally) projected Hessian is positive definite at (0,0) [with eigenvalues 0.113 and 3.696].

Let us now consider Newton's method on the Lagrangian equations

$$L(u, v, w, \lambda) = u + v^2 + w^2 - \lambda(-u + 3v - vw - 1).$$

Newton's method for the system of equations $\nabla_{x,\lambda} L(u, v, w, \lambda) = 0$ gives

conditions in the neighbourhood of a local constrained optimum, our *projected* Hessian $Z_k^T B_k Z_k$ is guaranteed to be positive definite.

4. Discussion. In attempting to construct a method for nonlinearly constrained optimization that is ideal, in the sense that one is able to do as well as unconstrained techniques, we require, conceptually at least, direct reduction.

The method of Lagrange can be derived from (an implicit) reduction. However, commonly used approximations based upon the method of Lagrange are quite distinct from approximations one might anticipate from direct reduction.

This is most easily appreciated by considering a particular example.

Example 4.1

$$\underset{x=[u,v,w]}{\text{minimize}}\ u + v^2 + w^2$$

$$\text{subject to } -u + 3v - vw = 1,$$

This example is chosen because direct substitution gives rise to an unconstrained quadratic problem, although the constrained problem is a QPQC.

Thus using the explicit reduction

$$u = 3v - vw - 1,$$

one obtains the equivalent *unconstrained* problem

$$\underset{(v,w)}{\text{minimize}}\ 3v + v^2 - vw + w^2 - 1,$$

with corresponding Hessian matrix,

$$H = \begin{pmatrix} 2 & -1 \\ -1 & 2 \end{pmatrix}$$

with eigenvalues 1 and 3, and gradient $(g_v, g_w) = (3 + 2v - w, -v + 2w)$.

Thus, one step of Newton's method, from any starting point produces the unconstrained minimum, $(v, w) = (-2, -1)$, from which it can be deduced that the optimum value of u and λ is, $u = -9, \lambda = -1$, with corresponding objective function value -4. For, example, starting at $(v^0, w^0) = (0, 0)$

$$\begin{pmatrix} v^1 \\ w^1 \end{pmatrix} = -H^{-1} \begin{pmatrix} g_v \\ g_w \end{pmatrix} = - \begin{pmatrix} 2/3 & 1/3 \\ 1/3 & 2/3 \end{pmatrix} \begin{pmatrix} 3 \\ 0 \end{pmatrix} = \begin{pmatrix} -2 \\ -1 \end{pmatrix}.$$

We note that, referring to the original constrained problem, although the Hessian of the Lagrangian is not positive definite, the (orthogonally) projected Hessian is positive definite at (0,0) [with eigenvalues 0.113 and 3.696].

Let us now consider Newton's method on the Lagrangian equations

$$L(u, v, w, \lambda) = u + v^2 + w^2 - \lambda(-u + 3v - vw - 1).$$

Newton's method for the system of equations $\nabla_{x,\lambda} L(u, v, w, \lambda) = 0$ gives

$$\begin{bmatrix} 0 & 0 & 0 & 1 \\ 0 & 2 & \lambda^k & -3+w^k \\ 0 & \lambda^k & 2 & v^k \\ 1 & -3+w^k & v^k & 0 \end{bmatrix} \begin{bmatrix} \delta u^k \\ \delta v^k \\ \delta w^k \\ \delta \lambda^k \end{bmatrix} = - \begin{bmatrix} 1 + & \lambda^k & \\ 2v^k - & 3\lambda^k & + \lambda^k w^k \\ 2w^k + & \lambda^k v^k & \\ u^k - & 3v^k & + v^k w^k + 1 \end{bmatrix}$$

Suppose

$$[u^0, v^0, w^0, \lambda^0] = [0, 0, 0, 0].$$

Then one obtains

$$\delta\lambda_0 = -1, \quad \delta w_0 = 0, \quad \delta v_0 = -3/2 \text{ and } \delta u_0 = -11/2$$

as the unique solution to the above system.

We also note that $\psi_1(u^1, v^1, w^1) = 0$, but although $\lambda^1 = \lambda^*$,

$$\nabla f(u^1, v^1, w^1), \neq \lambda^1 \nabla \psi_1(u^1, v^1, w^1),$$

and $f(u^1, v^1, w^1) = -13/4$.

The problem is that

$$\nabla_{x,\lambda} L(u, v, w, \lambda) = 0$$

is not linear in u, v, w, λ, *even if we have* $\lambda = \lambda^*$, whereas, we already observed, the corresponding reduced system is linear.

For example, starting at $[-1, 0, 0, -1]$ the Newton iteration on the Lagrangian system gives $\delta\lambda_0 = 0$, $\delta w_0 = -1$, $\delta v_0 = -2$, $\delta u_0 = -6$, and thus $[u^1, v^1, w^1, \lambda^1] = [-7, -2, -1, -1]$, an infeasible point with corresponding objective value -2.

We now consider the method of Coleman and Conn starting at the same point, $[u^0, v^0, w^0, \lambda^0] = [-1, 0, 0, -1]$. We have the following computations:

$$r(x) = f(x),$$

$$A = \begin{bmatrix} -1 \\ 3 \\ 0 \end{bmatrix}, \quad Z = \begin{bmatrix} 3/\sqrt{10} & 0 \\ 1/\sqrt{10} & 0 \\ 0 & 1 \end{bmatrix}, \quad \nabla r = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix},$$

and $Z^T \nabla r = \begin{pmatrix} 3/\sqrt{10} \\ 0 \end{pmatrix}$, whose norm is not considered small. Hence

$$Z^T \nabla^2 r \, Z = \begin{pmatrix} 2/10 & 0 \\ 0 & 2 \end{pmatrix}$$

and thus we obtain $d^1 = (-9/2, -3/2, 0)^T$. Minimizing along d^1 $p(x, \mu) = p(u, v, w, \mu)$, with say $\mu = 1$ [actually, the solution in this case is independent of μ since the nonlinear contribution to the constraint is zero], we obtain $(u^1, v^1, w^1) = (-11/2, -3/2, 0)$.

In order to understand more clearly the difference between direct reduction and the method of Lagrange, we need to examine more closely the derivation of the Lagrangian technique via the implicit function theorem.

The basic idea is as follows (for more details the reader is referred to Avriel [1976], for example).

Supposing one has the underlying problem

$$\underset{x \in R^n}{\text{minimize}} \quad f(x)$$

$$\text{subject to } \psi_i(x) = 0, \quad i = 1, \ldots, m, \quad m \leq n.$$

By suitable rearrangement one may define λ_i^*, $i = 1, \ldots, m$ uniquely as the solution to the system of linear equations

$$\sum_{i=1}^{m} \frac{\partial \psi_i(x^*)}{\partial x_j} \lambda_i = \frac{\partial f(x^*)}{\partial x_j}, \quad j = 1, \ldots, m. \tag{4.1}$$

Now, by applying the implicit function theorem to the constraint equations, there exist real functions $h_j(x_{m+1} \cdots x_n)$ such that

$$x_j = h_j(x_{m+1}, \ldots, x_n), \quad j = 1, \ldots, m.$$

Some straightforward algebra then gives

$$\sum_{k=1}^{m} \left[\frac{\partial f(x^*)}{\partial x_k} - \sum_{i=1}^{m} \lambda_i^* \frac{\partial \psi_i(x^*)}{\partial x_k} \right] \frac{\partial h_k(x_{m+1}^* \cdots x_n^*)}{\partial x_j}$$

$$+ \frac{\partial f(x^*)}{\partial x_j} - \sum_{i=1}^{m} \lambda_i^* \frac{\partial \psi_i(x^*)}{\partial x_j} = 0, \quad j = m + 1, \ldots, n. \tag{4.2}$$

We note that (4.1) are the first m equations of the standard Lagrangian system. (4.1) along with (4.2) determine the next $n - m$ equations, namely

$$\frac{\partial f(x^*)}{\partial x_j} - \sum_{i=1}^{m} \lambda_i^* \frac{\partial \psi_i(x^*)}{\partial x_j} = 0, \quad j = m + 1, \ldots, n, \tag{4.3}$$

but, the derivation depends upon (4.1) being satisfied.

Asymptotically, (4.1) is approximately satisfied. It is approximately satisfied by Coleman and Conn's algorithm *whenever multiplier estimates are used in the neighbourhood of a stationary point.* It is in this sense that I mean that the system of equations corresponding to the Lagrangian has no global interpretation consistent with direct reduction. Moreover, I consider the (implicit function theorem) derivation of the Lagrangian technique to be the derivation that is consistent with the approach I wish to take. That is, the relationship between constrained theory and constrained algorithms should enhance the possibility of producing methods as satisfactory as those currently available for unconstrained problems.

Arguably, the system of equations one should be solving via Newton's method ought to involve iterations of the type

$$B^k \begin{bmatrix} \delta x \\ \delta \lambda \end{bmatrix} = -g^k, \tag{4.4}$$

where

$$g_j^k = \sum_{i=1}^{m} \frac{\partial \psi_i(x^k)}{\partial x_j} \lambda_i^k - \frac{\partial f(x^k)}{\partial x_j}, \quad j = 1, \ldots, m$$

$$g_j^k = -\left[\sum_{s=1}^{m} g_s^k \frac{\partial h_s(x_{m+1}^k \cdots x_n^k)}{\partial x_j}\right] + \frac{\partial f(x^k)}{\partial x_j} - \sum_{i=1}^{m} \lambda_i^k \frac{\partial \psi_i(x^k)}{\partial x_j}, \quad j = m+1, \ldots, n$$

$$g_j^k = -\psi_{j-n}(x^k), \quad j = n+1, \ldots, n+m$$

and

$$B_{ij} = \frac{\partial(g_i^k)}{\partial \alpha_j}, \quad i = 1, \ldots, n+m, \quad j = 1, \ldots, n+m$$

where h_j, is defined by

$$x_j = h_j(x_{m+1}, \ldots, x_n), \quad j = 1, \ldots, m,$$

and

$$\alpha_j = x_j, \quad j = 1, \ldots, n$$
$$= \lambda_{j-n}, \quad j = n+1, \ldots, n+m.$$

The difficulty, of course, is that the h_s, $s = 1, \ldots, m$, are *implicit* functions. Furthermore, the argument is, that if we wish to emulate the reduction technique globally, the local behaviour distant from the asymptotic neighbourhood of x^*, is described by the model (4.4) and not the model (4.1), (4.3) and feasibility.

An alternative approach is to reject the reduction point of view but just state that ultimately

$$\nabla_{[x\lambda]} L(x^*, \lambda^*) = 0.$$

The difficulty with this is that one has abandoned the optimization model and replaced it by a nonlinear system of equations model for which there is no natural merit function. In addition, a solution to the Lagrangian system need not solve the given problem. It is generally recognized that the system of equations problem is more difficult than the unconstrained optimization problem for exactly the reason that in order to globalise the method it is necessary to have a merit function. Usually the merit function chosen is the least squares function (see, for example Dennis and Schnabel [1983], chapter 6).

Thus, from one point of view one might say that Lagrangian/SQP techniques are a global approach based upon an asymptotic result - one is using an equations model that can be rather bad globally. The standard approach to globalising the nonlinear equation methods is to use a merit function, usually the l_2 function. Obviously, SQP methods use a merit function closer to the l_1 function, namely the l_1 penalty function $p(x, \mu)$.

In contrast, the claim is that Coleman and Conn always use a model that is not locally invalid, irrespective of whether one is in the asymptotic region or not.

However, one interpretation of the problem that *globally* involves the least squares estimates of the Lagrangian multipliers and the corresponding Lagrangian Hessian, at least in the case of equality constraints, is given below.

Suppose one considers the problem

$$\text{EQP} \quad : \quad \underset{x \in R^n}{\text{minimize}} \, f(x)$$

$$\text{subject to } \gamma_i(x) = 0, \quad i \in E.$$

A reasonable objective is to determine a new point x^+ via a quadratic model of some sort, such that $f(x^+) < f(x)$ and $\Gamma(x^+) = \Gamma(x)$. [Γ denotes the vector whose components are the constraint functions γ_i].

It is not obvious as to how this can be achieved because of the nonlinearity of Γ. However, suppose one tries to construct a suitable quadratic model on the linear manifold $M = x + \langle Z \rangle$. The question is, what function should be modelled in M? One possible answer follows.

Let $x + Zw$ denote any point in M and let $g(w)$ represent the function we wish to model. If $\bar{x}(w)$ denotes the "nearest" point to $x + Zw$ such that $\Gamma(\bar{x}(w)) = \Gamma(x)$ then it is reasonable to define $g(w) = f(\bar{x}(w))$. Unfortunately $g(w)$ is not explicitly available. A compromise is to define $g(w) = f(u(w))$, where $u(w) = x + Zw + A^{\dagger}\Delta\,\Gamma(w)$ and $\Delta\,\Gamma(w)$ denotes $\Gamma(x) - \Gamma(x + Zw)$

If one then uses the model function

$$g(w) = f(u(w)) = f(x + Zw - A^{\dagger}\Delta\Gamma(w)),$$

then $g(w)$ represents the change in f evaluated on the linear manifold with a linear least squares correction to reflect the nonlinearity of the constraints.

We note that the gradient and the Hessian of $g(w)$ evaluated at $w = 0$ are given by

$$\nabla_w g(0) = Z^T \nabla_x f(x)$$

and

$$\nabla^2_{ww} g(0) = Z^T [\, \nabla^2_{xx} f(x) - \sum_{i \in E} \lambda_i \, \nabla^2_{xx} \gamma_i(x)] Z,$$

where $\lambda = A^{\dagger^T} \nabla_x f(x)$.

These observations were brought to my attention by T. F. Coleman and are essentially contained in Fletcher [1981], chapter 12.

Although the above is a reasonable global interpretation of the Lagrangian multiplier and Hessian it is not, in my opinion, one which recommends itself algorithmically. In particular, the range space step amounts to an undamped Newton step on the constraints. Such a step seems particularly undesirable when the projected gradient onto the linear approximation to the manifold is large. On the other hand, it does not appear computationally desirable to perform a line search on the range space step.

5. Secant Methods. Having motivated and discussed the merits and difficulties

of various Newton-like approaches to the nonlinearly constrained problem we would now like to consider the generalisation of the Coleman/Conn approach.

Our aim is to develop a method that

i) uses secant approximations throughout,

ii) approximates only projected Hessians

and

iii) maintains a comparable (i.e. better than linear) rate of convergence.

The objectives seem reasonable since

a) the second-order sufficiency conditions suggest it is possible,

b) if the number of active constraints is large, considerable savings may be realised,

and c) in the neighbourhood of a local minimum the *projected* Lagrangian Hessian is positive definite.

In other words, we are in a situation that parallels the unconstrained case. However, there are some differences. These include the fact that inherited positive definiteness is not so easy to maintain as in the unconstrained case [see below] and there are some advantages to recurring more than the projected Hessian. This latter observation is a consequence of our earlier remark that if the problem has non-linear constraints one doesn't know the true projected space even if the correct active set is identified.

Details of the asymptotics of such a secant approach are given in Coleman and Conn [1984].

The update suggested is natural in the sense that it is a straightforward generalisation derived from projecting the unconstrained secant method. For consistency with the details given in Coleman and Conn [1984] we shall describe a method based upon Davidon-Fletcher-Powell. Corresponding results are readily obtained for a projected Broyden-Fletcher-Goldfarb-Shanno update.

We now define some essential notation.

In the modified Newton method of Coleman and Conn described in section 3 the direction d^k can be considered to be defined by

$$d^k = -Z_k(Z_k^T G_k Z_k)^{-1} Z_k^T \nabla r^k$$

where

$$\begin{aligned} G_k &= I && : \textit{Case}\,1 \\ &= \nabla^2 r^k && : \textit{Case}\,2 \\ &= \nabla^2 r^k - \sum_{i \in A_u} \hat{\lambda}_i \nabla^2 \gamma_i^k && : \textit{Case}\,3 \end{aligned}$$

Case 1: corresponds to taking the projected negative gradient direction (the projected gradient is small and at least one multiplier approximation is out-of-kilter).

Note: in this case Z_k corresponds to a basis for the null space of the gradients of the t active constraints with the out-of-kilter constraint removed i.e. Z_k has $n - t + 1$ columns.

Case 2: corresponds to the projected gradient not being considered small.

Case 3: corresponds to a small projected gradient with all multiplier estimates in-kilter.

In the secant context we wish to redefine d^k by $d^k = -Z_k B_k^{-1} Z_k^T \nabla r^k$ where B_k denotes the current approximation to the *projected* Hessian $Z_k^T G_k Z_k$ with G_k defined as above. We emphasize that if we have t active constraints at iteration k, B_k is an $(n-t) \times (n-t)$ [or, in case 1, an $(n-t+1) \times (n-t+1)$] matrix.
Let

$$s^k = Z_k^T [x_k^+ - x_k] \text{ and } y^k = Z_k^T[\nabla \tilde{p}(x_k^+) - \nabla \tilde{p}(x_k)] \tag{5.1}$$

where

$$\nabla \tilde{p}(x) = \begin{cases} \nabla r(x) - \sum_{i \in A_u} \hat{\lambda}_i^k \nabla \gamma_i(x) & \text{, if } \hat{\lambda}^k \text{ available} \\ & \text{[i.e. } Z_k^T \nabla r(x) \text{ small] and in-kilter,} \\ \nabla r(x) & \text{, otherwise} \end{cases}$$

and

$$x_k^+ = x^k + d^k.$$

Clearly s^k is the *projected* null space step and y^k is the *projected* difference in the "gradient" of p. [The superscript k on x has sometimes become a subscript in an attempt to avoid cluttering the superscripts.]

The rank two update that corresponds to Davidon-Fletcher-Powell is given by

$$B_{k+1} = B_k + \frac{[y^k - B_k s^k][y^k]^T + [y^k][y^k - B_k s^k]^T}{[s^k]^T y^k} - \frac{[s^k]^T[y^k - B_k s^k] y^k [y^k]^T}{([s^k]^T y^k)^2}$$

What follows is then an adaption of Broyden, Dennis and Moré [1973], the asymptotic analysis in Coleman and Conn [1982a] and the characterization of Dembo, Eisenstat and Steihaug [1982].

Under the basic assumptions (see Coleman and Conn [1984] for more details)

a) $f, \gamma_i \in C^2$,
b) second-order sufficiency conditions hold at x^*,
c) $\{x^k\} \in W$, a compact set,
d) linear independence of the active constraints at x^*, and
e) the correct active set is identified,

we prove the following theorem Coleman and Conn [1984].

Theorem 5.1
If

$$x^k \to x^*, \ ||x^{k+1} - x^*|| = O(||x^k - x^*||),$$

and

$$||r_p^k|| + ||r_\Gamma^k|| = o(||Z_k^T \nabla r^k|| + ||\Gamma^k||),$$

then

$$||x^{k+1} - x^*|| = o(||x^{k-1} - x^*||),$$

where

$$r_p^k \triangleq Z_k^T \nabla r^k + H_k Z_k^T \delta^k, \quad r_\Gamma^k \triangleq \Gamma^k + A_k^T \delta^k,$$

$$\delta^k = x^{k+1} - x^k, \quad H^k = Z_k^T \nabla^2 \tilde{p} Z_k,$$

and x^* is a stationary point of $p(x, \mu)$.

Remark: The statement of Theorem 5.1 is compatible with the approach taken by Dembo, Eisenstat and Steihaug, adapted to the constrained problem. It takes some straightforward algebra to show that

$$||r_p^k|| = O(||[B_k - Z_*^T \nabla^2 \tilde{p}(x^*) Z_*] Z_k^T [x^{k+1} - x^k]||)$$

Thus,

$$||r_p^k|| = o(||Z_k^T \nabla r^k||)$$

implies that,

$$||[B_k - Z_*^T \nabla^2 \tilde{p}(x^*) Z_*] Z_k^T [x^{k+1} - x^k]|| = o[||x^{k+1} - x^k||],$$

using $Z_k^T \nabla r_p^k = -B_k Z_k^T [x^{k+1} - x^k]$ and $||B_k||$ is uniformly bounded below. This is more like the Broyden, Dennis and Moré statement. However, note that there is an omission in this latter interpretation. The above only considers the null space component of the step. When the range space step is considered a two step superlinear convergence rate results. We remark that r_p^k reflects the accuracy to which the system $\nabla^2 \tilde{p}_k Z_k^T \delta^k = -Z_k^T \nabla r^k$ is satisfied, whereas r_Γ^k reflects the accuracy to which $A_k^T \delta^k = -\Gamma^k$ is solved. Not surpisingly (since there are no constraints) this component is absent from the Broyden, Dennis and Moré statements.

The result of Theorem 5.1 is closely related to that given by Powell [1978]: however, the conditions given do not presuppose a particular algorithm class.

It is then shown (Coleman and Conn [1984], Theorem 3.6) that under appropriate conditions, the secant method above, necessarily satisfies $||r_p^k|| + ||r_\Gamma^k|| = o(||Z_k^T \nabla r_k|| + ||\Gamma^k||)$, and consequently an asymptotic two-step superlinear convergence rate is attained.

More recently, other researchers have considered asymptotic results pertaining to secant methods for nonlinear constrained optimization problems that update projected Hessians (see e.g. Nocedal and Overton [1984], Fontecilla [1983] and Byrd [1984]).

There remains one difficulty with globalising the above secant results - the question of inherited positive definiteness. If one wished to extended directly the unconstrained results one requires that $s^T y$ is strictly positive where, as above [equation (5.1)] , s and y are the *projected* differences in null space step and gradient. In Coleman and Conn [1984] this is proven in a neighbourhood of

$(x^*, Z_*^T \nabla^2 \tilde{p}(x^*) Z_*)$ - in other words, asymptotically. Furthermore, it is easy to prove globally for convex problems.

The usual "trick" in unconstrained optimization is to use the line search, since at the minimum along the line, $\delta^T g^{k+1} = 0$, and $\delta^T g^k < 0$ implies that $\delta^T y > 0$, where δ is the difference in x (i.e. the search direction), g is the objective function gradient and y is the gradient difference. One undesirable feature of this approach is that although our (quadratic) model is a local one our condition depends upon a non-local result, namely that the *average* curvature is positive. Thus, for example, if one is in the neighbourhood of a local maximum and one begins descending, this may necessitate a very long step indeed - in fact to a region where the previous Hessian approximation is totally invalid. It is this non-local characterization of the line-search resolution of the problem that causes difficulty in the nonlinearly constrained context. Our projective affine space determined by the gradients of the active constraints at the end of the step may bear no relation to the affine space at the start of the step.

Thus the approach one might take is to use the line search to ensure that $[d^k]^T Z_k Z_k^T \nabla \tilde{p}$ is small enough, since $[d^k]^T Z_k Z_k^T \nabla \tilde{p} < 0$. The difficulty is then to ensure that $p(x^k + \tau d^k + \alpha v^k) < p(x^k)$ where the scalar τ is determined by the line search and the scalar α is determined by the second-order correction ($\alpha = 0$ or 1 in our present implementation). This approach is currently under investigation. However, I feel that it is likely to be inferior to an approach that can intelligently abandon positive definiteness and make use of directions of negative curvature since this latter approach is more consistent with the idea of using a local quadratic model (c.f. trust region methods).

6. Miscellanea.

Degeneracy: Degeneracy is a difficulty whose importance is gradually being recognised. In particular there are classes of problems for which degeneracy (often, as near-degeneracy) occurs frequently. There are some similarities between degeneracy and singular optimization/nonlinear equation problems. Two approaches that are currently being considered are special perturbations and less degenerate subproblems.

An example of the former is given, in the context of multifacility location problems, by Calamai and Conn [1985]. The basic idea is the following: in linear programming, perturbations have the attribute that an optimum for the perturbed problem lies at a vertex. Unfortunately, this is not true, in general, for nonlinear problems. However, for specially structured problems the perturbation can be carefully chosen so that the only points of interest for the perturbed problem are easily obtainable vertices. This is important if, for example, one considers that for a general nonlinear problem with $n + m$ active constraints at the solution (that is, at least m redundancies), a perturbed problem (with arbitrarily small perturbations) may have *no* active constraints at the (perturbed) solution.

An example of the latter approach in the context of nonlinear l_1 problems (but which has broad application) is given by Bartels and Busovača [1984]. The basic idea is as follows: suppose one is optimal for the l_1 problem:

$$\underset{x}{\text{minimize}} \sum_{i=1}^{m} |f_i(x)|$$

that is presumed non-degenerate.
At optimality

$$\sum_{i \notin A} \sigma_i \nabla f_i(x) + \sum_{i \in A} \lambda_i \nabla f_i(x) = 0,$$

with $|\lambda_i| \leq 1$, $i \in A$, where A is the index of active constraints and $\sigma_i = sgn f_i(x)$ (see, for example Bartels, Conn and Sinclair [1978]).

The difficulty in the degenerate case is that, because of the linear dependence of the active constraints, the λ_i are no longer uniquely defined. However, even in this instance, multipliers, λ_i, at an optimal point do exist that satisfy $|\lambda_i| \leq 1$. Consequently, these conditions may be added *explicitly* as constraints. In other words one recognises optimality by solving

$$\sum_{i \notin A} \sigma_i \nabla f_i(x) + \sum_{i \in A} \lambda_i \nabla f_i(x) = 0$$

$$\text{subject to } -1 \leq \lambda_i \leq 1.$$

Moreover, if no solution exists, an optimal point has not been found, but a descent direction can be readily constructed.

Soft Constraints: In engineering practice it is not uncommon to have soft constraints - that is, constraints that are only approximate in the sense that one is willing in certain circumstances to allow some of them to be (moderately) violated. One may consider penalty functions to be appropriate for such environments since the problem for which some of the constraints are soft may in fact be infeasible as stated, although meaningful solutions may well be attainable.

Choice of the Penalty Parameter: I feel there is still no significant progress in determining an ideal value for exact l_1 penalty function(s) penalty parameter(s). The difficulty is that, ignoring the computational cost, one has no real idea as to what an ideal penalty parameter choice is *globally*. One is not even sure as to whether it should be constant or changing. Although changing parameters destroy most conventional global convergence proofs, the asymptotic bounds may be an inappropriate choice far from the neighbourhood of the solution.

Large Sparse Problems: This is an active research area and is still relatively open. However, there may be some particular advantages in using the l_1 exact penalty function when there are just a few constraints that destroy an otherwise highly structured problem. For example, suppose one has a problem for which all but one of the constraints involves only a few of the variables. If this single constraint is explicitly included in the penalty function, with a low priority for activation, then one is able to exploit the separable structure of the remaining constraints. For similar reasons, the l_∞ exact penalty function is also attractive.

7. Applications - new and old. As was mentioned in more detail in [Conn 1982], I would like to emphasize that there are problems for which one might not normally consider exact penalty function techniques but for which they are very useful indeed. In particular, linear programming problems, quadratic programming problems, problems with special structure (for example, continuous multifacility location problems) and problems with linear constraints.

I would like to end this article by describing briefly three examples of current

research that incorporates extensions of the ideas described above.

The first topic relates to the comments made earlier in §4 concerning the merit function for SQP methods viewed as a method for solving the Lagrangian system of equations.

i) System of equations.

Find, if possible, a solution to the system of equations

$$\text{SE:} \qquad \begin{aligned} f_1(x) &= 0, \\ f_2(x) &= 0, \\ &\vdots \\ f_m(x) &= 0. \end{aligned} \qquad ,x \in \mathbb{R}^n \text{ , } m \text{ a general positive integer}$$

One approach is to solve SE via the nonlinear l_1 technique of Bartels and Conn [1981] modified to handle degeneracy as in Bartels and Busovača [1984]. One merely notes here that in the case where some of the functions are (almost) linear the proposed nonlinear l_1 technique seems particularly appropriate and furthermore the direction model subproblem and the merit function $\sum_{i=1}^{m} |f_i(x)|$ are consistent. The above is ongoing research with my colleagues R.H. Bartels and S. Busovača.

ii) Non-smooth problems.

This is joint work with P.F. O'Neill and details are given in Conn and O'Neill [1983].

The problem is

$$\text{NSP:} \qquad \begin{aligned} &\underset{x}{\text{minimize}}\ f(x) \\ &\text{subject to}\ \phi_i(x) \geq 0, \quad i \in I \\ &\qquad\qquad \psi_i(x) = 0, \quad i \in E \end{aligned}$$

where the f, ϕ_i's and ψ_i's may be **non-differentiable** or even **discontinuous.**

However, the f, ϕ_i's and ψ_i's cannot be too pathological since the basic approach is to separate the underlying problem into a sequence of related smooth problems.

Essentially, we "partition" the problem into "cells" such that inside each cell we have continuous (differentiable) problems .

A penalty function approach is then especially desirable since almost all of the cell subproblems will only be "loosely" solved.

One then addresses the following questions

a) Which problems are so partitionable?
b) How does one use "solutions" to the continuous subproblems to solve the original problem?
c) What does one mean by a solution to the original problem?

We have designed and implemented algorithms based on these ideas.

An example of one of several applications is the compound alternative feasible set problem viz.

$$\text{minimize } f^0(x)$$

$$\text{subject to}$$

$$f^i(x) \geq 0, \quad i \in I^1$$

$$\text{or } f^i(x) \geq 0, \quad i \in I^2$$

$$\cdot$$

$$\cdot$$

$$\text{or } f^i(x) \geq 0, \quad i \in I^m$$

i.e. minimize f^0 subject to satisfying *at least* one of the systems I^j, $1 \leq j < m$, being satisfied.

iii) Semi-infinite Programming

This is joint work with N.I.M. Gould and details are given in Conn and Gould [1984].

The problem is

$$\text{SIP:} \qquad \underset{x \in \mathbb{R}^n}{\text{minimize}} f(x)$$

$$\text{subject to } \phi_i(x, t) \geq 0, \quad \text{for all } t \in T_i \subseteq \mathbb{R}^{p_i} \quad i = 1, \ldots, m,$$

where p_i is some integer.

Thus, one has a finite number of variables but an *infinite number of constraints.*

Our aim was to produce a *global* method. The basis was a generalisation of the l_1 exact penalty function approach to nonlinear programming. For simplicity of exposition, I will consider the convex case, i.e. f convex, $\phi_i(x, t)$ concave in x, T_i convex and $\phi_i(x, t)$ convex in t over T_i.

Let

$$\Omega_i(x) \subset T_i,$$

be defined by

$$\Omega_i(x) = \{t \mid \phi_i(x, t) \leq 0\}.$$

We note that, if the T_i's were discrete sets SIP becomes simply a nonlinear programming problem with corresponding l_1 exact penalty function, written somewhat eccentrically as,

$$p(x, \mu) = \mu f(x) - \sum_{i=1}^{m} \sum_{t \in \Omega_i(x)} \phi_i(x, t),$$

$$\left[\text{since } \min [0, \phi_i(x, t)] = \begin{cases} \phi_i(x, t) & t \in \Omega_i(x) \\ 0 & t \notin \Omega_i(x) \end{cases}\right]$$

This suggests a generalisation obtained by considering the limit as the discrete T_i tends to continuous T_i. Thus we solve SIP by using *the exact penalty function*

$$\rho(x, \mu) = \mu f(x) - \sum_{i=1}^{m} \frac{\int_{\Omega_i(x)} \phi_i(x, t)\, dt}{\int_{\Omega_i(x)} dt}.$$

[One can motivate the necessity of the support function $\int_{\Omega_i(x)} dt$ by considering the case when $\phi(x, t)$ is constant and negative on a region Ω_i.]

More generally, in the non-convex case $\rho(x, \mu)$ is discontinuous. However, the problem only occurs because the region of integration splits and is readily overcome by employing a simple trick. The interested reader is referred to Conn and Gould [1984] for details.

The algorithm is then very much like that described for the nonlinear programming problem. ρ is partitioned as $\rho = \rho_1 + \rho_2$ where ρ_1 is that part of ρ that is differentiable in a prescribed neighbourhood of the current point and ρ_2, in contrast, contains all discontinuities in the first derivative. One then attempts to minimize ρ by minimizing ρ_1 subject to ρ_2 not changing, up to the chosen order, *if possible.* Otherwise the derivative of ρ_1 can be written as a linear combination of the gradients of the $\nabla_x \phi_i(x, t_i)$'s that are "active" - i.e. those that determine ρ_2.

In this latter situation one can investigate relaxing the "zero rate of change" of *one* of the $\phi_i(x, t_i)$'s to obtain descent.

If one is not able to determine descent in this manner it is possible to conclude that one is at a stationary point of ρ.

Clearly, such an algorithm would require, at *least*, first derivatives of $\rho_1(x, \mu)$, which necessitates the evaluation of at most *mn* integrals.

These derivatives are somewhat cumbersome to calculate. As an illustration consider the simplest case, where $T_i \subset \mathbb{R}$ and $\Omega_i(x) = [a_i(x), b_i(x)]$.

Then one may write

$$\rho(x, \mu) = \mu f(x) - \sum_{i=1}^{m} \Phi_i(x),$$

where

$$\Phi_i(x) = p_i(x) \,/\, q_i(x) \text{ with } p_i(x) = \int_{\Omega_i(x)} \phi_i(x, t)\, dt \text{ and } q_i(x) = \int_{\Omega_i(x)} dt.$$

In this case

$$\nabla_x \Phi_i(x) = \frac{\nabla_x p_i(x)}{q_i(x)} - \frac{p_i(x)\, \nabla_x q_i(x)}{[q_i(x)]^2},$$

where

$$\nabla_x p_i(x) = \int_{a_i(x)}^{b_i(x)} \nabla_x \phi_i(x, t)\, dt$$

and

$$\nabla_x q_i(x) = \nabla_x b_i(x) - \nabla_x a_i(x) = \left\{ \frac{\nabla_x \phi_i(x, a_i(x))}{\nabla_t \phi_i(x, a_i(x))} - \frac{\nabla_x \phi_i(x, b_i(x))}{\nabla_t \phi_i(x, b_i(x))} \right\}.$$

For some problems that arise frequently the structure enables these integrals to be evaluated simply.

Our test results to date correspond to a first-order algorithm tested on small problems to verify the viability of the method. The integrals were evaluated using adaptive quadrature.

Our ongoing research on semi-infinite problems is in investigating second-order methods. We also wish to investigate an approach based upon

$$\rho_\infty(x, \mu) = \mu f(x) - \min_{1 \le i \le m} \min_{t \in \Omega_i(x)} \phi_i(x, t),$$

as well as considering rates of convergence.

8. Conclusion. This article owes much to my splendid colleagues, both at Waterloo and in the mathematical community at large. Much of my interest in exact penalty functions has been motivated by the work of all of us to develop good algorithms for problems with nonlinear constraints. We have made a beginning, but there is much that remains to be done.

My opinions, set out above, are based on the current state-of-the-art as I perceive it. However, they are merely my opinions, and they should be viewed accordingly. If at times they are enlightening, then I am lucky - if they are wrong or misleading, then the blame is entirely mine.

REFERENCES

M. Avriel, "Nonlinear Programming: Analysis and Methods", Prentice-Hall, (1976).

R.H. Bartels and S. Busovača, "Handling degeneracy in a nonlinear l_1 algorithm", presented at the SIAM conference on numerical optimization, Boulder, Colorado, (1984).

R.H. Bartels and A.R. Conn, "An approach to nonlinear l_1 data fitting", in 'Proceedings of the Third Mexican Workshop on Numerical Analysis', J.P. Hennart (ed.), pp. 48-58, Springer-Verlag, (1981).

R.H. Bartels, A.R. Conn and J.W. Sinclair, "Minimization techniques for piecewise differentiable functions: the l_1 solution to an overdetermined linear system", SIAM J. Num. Anal., Vol. 15, pp. 224-241, (1978).

M.C. Biggs, "Constrained minimization using recursive quadratic programming: some alternative subproblem formulations", in 'Towards Global Optimization' L.C.W. Dixon and G.P. Szegö (eds.), North Holland, pp. 341-349, (1975).

C.G. Broyden, "The convergence of a class of double rank minimization algorithms, 2. The new algorithm", J. Inst. Math. Appl., Vol. 6, pp. 222-231, (1970).

C.G. Broyden, J.E. Dennis and J.J. Moré, "On the local and superlinear convergence of quasi-Newton methods", J. Inst. Math. Appl., Vol. 12, pp. 223-245, (1973).

R.H. Byrd, " On the convergence of constrained optimization methods with accurate Hessian information on a subspace", Tech. Rept. CU-CS-270-84, Dept. of Computer Science, University of Colarodo at Boulder, Colorado, (1985).

P.H. Calamai and A.R. Conn, "A projected Newton method for l_p location problems", Tech. Rept. CS-85-07, University of Waterloo, (1985).

R.M. Chamberlain, C. Lemaréchal, H.C. Pedersen and M.J.D. Powell, "The watchdog technique for forcing convergence in algorithms for constrained optimization", in Mathematical Programming Study, Vol. 16, pp. 1-17, (1982).

T.F. Coleman and A.R. Conn, "Nonlinear programming via an exact penalty function: asymptotic analysis", Math. Programming, Vol. 24, pp. 123-136, (1982a).

T.F. Coleman and A.R. Conn, "Nonlinear programming via an exact penalty function: global analysis", Math. Programming, Vol. 24, pp. 137-161, (1982b).

T.F. Coleman and A.R. Conn, "On the local convergence of a quasi-Newton method for the nonlinear programming problem", SIAM J. Num. Anal., Vol. 20 , pp. 755-769, (1984).

A.R. Conn, "Penalty function methods", in 'Nonlinear Optimization 1981', M.J.D. Powell (ed.), pp235-242, Academic Press, (1982).

A.R. Conn and N.I.M. Gould, "An exact penalty function for semi-infinite programming", Tech. Rept. CS-84-17, University of Waterloo, (1984).

A.R. Conn and P.F. O'Neill, "Nondifferentiable optimization and worse", Tech. Rept. CS-83-05, University of Waterloo, (1983).

R.S. Dembo, S.C. Eisenstat and T. Steihaug, "Inexact Newton Methods", SIAM J. Num. Anal., Vol. 19, pp. 400-408, (1982).

J.E. Dennis and R.B. Schnabel, "Numerical Methods for Unconstrained Optimization and Nonlinear Equations", Prentice-Hall, (1983).

R. Fletcher, "A new approach to variable metric algorithms", Comp. J., Vol. 13, pp. 317-322, (1970).

R. Fletcher, "Practical Methods of Optimization, Vol. 2, Constrained Optimization", John Wiley, (1981).

R. Fletcher, "Penalty Functions", in 'Mathematical Programming: The State of the Art, Bonn, 1982', A. Bachem, M. Grötschel and B. Korte (eds.), pp. 87-114, Springer-Verlag, (1983).

R. Fletcher, "l_1 penalty methods for nonlinear constraints", this volume pp. 26-40, (1985).

R. Fontecilla, "The lack of positive definiteness in the Hessian in constrained optimization", Tech. Rept. 1334, Dept. of Computer Science, University of Maryland, (1983).

P.E. Gill, W. Murray and M.H. Wright, "Practical Optimization", Academic Press, (1981).

D. Goldfarb, "A family of variable metric methods derived by variational means", Math. Comp., Vol. 24, pp. 23-26, (1970).

W. Murray and M.H. Wright, "Projected Lagrangian methods based on the trajectories of penalty and barrier functions", Tech. Rept. SOC 78-23, Dept. of O.R., Stanford University, (1978).

J. Nocedal and M. Overton, "Projected Hessian updating algorithms for nonlinearly constrained optimization", SIAM J. Num. Anal. (to appear).

M.J.D. Powell, "The convergence of variable metric methods for nonlinearly constrained optimization calculations", in 'Nonlinear Programming 3', O.L. Mangasarian, R.R. Meyer and S.M. Robinson (eds.), Academic Press, pp. 27-63, (1978).

M.J.D. Powell, "Variable Metric Methods for Constrained Optimization", in 'Mathematical Programming: The State of the Art, Bonn, 1982', A. Bachem, M. Grötschel and B. Korte (eds.), pp. 288-311, Springer-Verlag, (1983).

R.B. Schnabel, "Unconstrained Optimization in 1981" in 'Nonlinear Optimization 1981', M.J.D. Powell (ed.), pp. 3-16, Academic Press, (1982).

R.B. Schnabel and P. Frank, "Tensor methods for solving systems of nonlinear equations", Tech. Rept. CU-CS-243-83, Dept. of Computer Science, University of Colorado at Boulder, (1983).

D.F. Shanno, "Conditioning of quasi-Newton methods for function minimization", Math. Comp., Vol. 24, pp. 647-656, (1970).