

Solution of discrete-time optimal control problems on parallel computers *

Stephen J. WRIGHT

Department of Mathematics, Box 8205, North Carolina State University, Raleigh, NC27695, USA

Received August 1989

Revised January/June 1990

Abstract. We describe locally-convergent algorithms for discrete-time optimal control problems which are amenable to multiprocessor implementation. Parallelism is achieved both through concurrent evaluation of the component functions and their derivatives, and through the use of a parallel band solver which solves a linear system to find the step at each iteration. Results from an implementation on the Alliant FX/8 are described.

Keywords. Discrete-time optimal control problems, Constrained optimization problem, Parallel band solver, Multiprocessor implementation, Shared memory multiprocessor.

1. Introduction

The unconstrained N -stage discrete-time optimal control problem with Bolza objectives has the form

$$\min_u F(u) \stackrel{\text{def}}{=} l_{N+1}(x^{N+1}) + \sum_{i=1}^N l_i(x^i, u^i) \quad (1.1)$$

$$\begin{aligned} x^1 &= a(\text{fixed}) \\ x^{i+1} &= f^i(x^i, u^i), \end{aligned} \quad (1.2)$$

where $x^i \in R^n$, $i = 1, \dots, N + 1$, and $u^i \in R^m$, $i = 1, \dots, N$. The x^i variables are usually referred to as *states* and the u^i as *controls*. This problem may arise, for example as a discretization of the continuous-time problem

$$\begin{aligned} \min_{u(t)} \int_0^1 L(x(t), u(t), t) dt + \phi(x(1)) \\ \text{s.t. } \dot{x} = f(x, u, t), \quad x(0) = x_0, \end{aligned}$$

(Polak [12], Jacobson and Mayne [7], Kelley and Sachs [8]), but it is also of interest in its own right in modelling certain physical processes.

In some situations there may be stagewise constraints on the x^i and u^i ; in their most general form these can be written as

$$g^i(x^i, u^i) \leq 0, \quad i = 1, \dots, N, \quad (1.3)$$

* Research partially performed while the author was visiting Argonne National Laboratory, Argonne, IL, and supported by the National Science Foundation under Grants DMS-8619903, DMS-8900984 and ASC-8714009, and the Air Force Office of Scientific Research under Grant AFOSR-ISSA-87-0092

where $g^i: R^n \times R^m \rightarrow R^{n_i}$. One important special case is the simple control-constrained problem in which we require

$$\|u^i\|_\infty \leq \eta, \quad i = 1, \dots, N, \tag{1.4}$$

where η is some constant. In this paper, we consider the typical case in which N is much larger than m or n , and will describe the implementation of quadratically-convergent algorithms for solving these problems on shared-memory multiprocessors. It will be assumed throughout that

- (i) the functions l_i , f^i and g^i are smooth enough (e.g. twice continuously differentiable) that the theoretical local convergence properties of these methods hold, and
- (ii) the second derivatives of l_i , f^i and g^i with respect to x^i and u^i can be calculated.

A naive application of Newton's method to $F(u)$ would necessitate solving a dense linear system of order mN at each iteration – a task whose complexity is $O((mN)^3)$. However, algorithms have been known for some time for which the calculation of the step in u can be done in $O((m^3 + n^3)N)$ operations, and for which the local convergence is fast. Dunn and Bertsekas [4] have shown recently that the Newton step itself is obtainable in $O((m^3 + n^3)N)$ operations. In section 2.1, we motivate these Newton-like methods somewhat differently, by considering (1.1, 1.2) as a *constrained* optimization problem in which both the u^i and x^i vectors are unknowns. Three variants of the sequential quadratic programming method are presented, the last of which is in fact Newton's method applied to $F(u)$. The basis of all these methods is the solution of a large block-banded symmetric linear system, which can be done in $O((m^3 + n^3)N)$ operations, using the methods mentioned above. These are inherently serial, however, and one of the features of the parallel method which we propose for (1.1, 1.2) and (1.1, 1.2, 1.4) is the solution of this linear system using the parallel band solver from Wright [14]. The band solver can be tailored to take advantage of some of the special features of the linear system in question; this is discussed in Section 2.2. Some test problems from the literature and numerical results obtained on the Alliant FX/8 are given in Section 2.3. In Section 3.1, the solution of the general constrained problem (1.1, 1.2, 1.4) using active-set methods is outlined, while in Section 3.2, application of a gradient-projection-based approach to the simple control-constrained problem is discussed. For this latter application, two algorithms (one from Dunn [3] and one which uses the parallel band solver) are described; some numerical results are given in Section 3.3.

Superscripts are used to distinguish between vector and matrix quantities at different stages $i = 1, 2, \dots, N + 1$. Subscripts on scalars serve this same purpose. Subscripts on vectors and matrices indicate iteration number.

In keeping with the usual conventions, we say that a sequence $\{a_k\}$ is R-quadratically convergent to a_* if there exists a sequence of scalars $\{\eta_k\}$ such that

$$\|a_k - a_*\| \leq \eta_k \tag{1.5}$$

and

$$\limsup \frac{\eta_{k+1}}{\eta_k^2} < \infty.$$

We say that $\{a_k\}$ is Q-quadratically convergent if the inequality in (1.5) can be replaced by an equality.

To avoid cluttering of the expressions in Section 2, we omit the transpose superscript 'T' on vectors where this does not cause confusion. For example the expression $y^T Q y$ will be written as $y Q y$, and $y^T x$ will be yx . (The superscript is retained for matrix transposes.) Another notational convenience in the remainder of the paper is to use u , x , and p defined by

$$\begin{aligned} u^T &= (u^{1T}, u^{2T}, \dots, u^{NT}) \\ x^T &= (x^{2T}, x^{3T}, \dots, x^{N+1T}) \\ p^T &= (p^{2T}, p^{3T}, \dots, p^{N+1T}). \end{aligned}$$

(p^{i+1} is the Lagrange multiplier vector for the ‘constraint’ (1.2).) In keeping with omission of the transpose on vectors, (u, x, p) will be used for $\begin{pmatrix} u \\ x \\ p \end{pmatrix}$.

2. The unconstrained problem

2.1 Algorithms with fast local convergence

Specialized algorithms for solving (1.1) which are both quadratically convergent and for which the cost of calculating successive iterates increases only linearly with N have been known for some time. The best-known examples are the differential dynamic programming (DDP) algorithm of Jacobson and Mayne [7] and the method described by Polak [12, pp. 92–102]. The DDP algorithm is quadratically convergent in u , since it produces steps which are asymptotically similar to Newton steps. It was used in preference to Newton’s method because it was apparently believed for some time that the Newton step is much more expensive to calculate. This is not true; Dunn and Bertsekas [4] show that the Newton step can be obtained by using dynamic programming to solve a linear-quadratic optimal control problem. Polak’s approach also solves a linear-quadratic subproblem.

In this section, we consider the treatment of (1.1, 1.2) as a problem in constrained optimization, where x and u are both variable and the state equations (1.2) are constraints. Application of the sequential quadratic programming (SQP) method yields, under the usual second-order sufficiency assumptions, Q-quadratic convergence in the triple (u, x, p) , where the vector p is made up of Lagrange multipliers p^{i+1} which correspond to the ‘constraint’ $x^{i+1} = f^i(x^i, u^i)$. Modifications to the SQP algorithm produce variants which are Q-quadratically convergent in (u, x) and in u alone; the latter is in fact Newton’s method. These variants are successively less parallelizable, and since there is little practical difference in the convergence rates (all are R-quadratically convergent in u), SQP should certainly be considered as an alternative to Newton’s method, at least in the neighbourhood of a solution. We take up this point later.

We assume throughout that there are vectors u_* and x_* which satisfy second-order sufficient conditions to be a solution of (1.1). These are defined after we have introduced some notation. Further, we assume that all iterates are sufficiently close to u_* and x_* that the usual local convergence properties hold. Convergence to a solution from an arbitrary starting point is of course desirable from a practical point of view – a Levenberg-style modification is discussed in Dunn and Bertsekas [4] which would help in this aim. We put this issue aside here, since we are mainly concerned with the basic algorithm and its implementation in a parallel setting.

The Lagrangian for (1.1) is, using the notation already introduced for Lagrange multipliers,

$$\mathcal{L}(u, x, p) = l_{N+1}(x^{N+1}) + \sum_{i=1}^N l_i(x^i, u^i) - \sum_{i=1}^N (x^{i+1} - f^i(x^i, u^i))p^{i+1}. \quad (2.1)$$

First-order conditions for (u, x, p) to be a solution of (1.1) are that

$$\frac{\partial \mathcal{L}}{\partial u^i} = \frac{\partial l_i}{\partial u^i} + \frac{\partial f^i}{\partial u^i} p^{i+1} = 0, \quad i = 1, \dots, N \quad (2.2)$$

$$\frac{\partial \mathcal{L}}{\partial x^i} = \frac{\partial l_i}{\partial x^i} - p^i + \frac{\partial f^i}{\partial x^i} p^{i+1} = 0, \quad i = 2, \dots, N \quad (2.3)$$

$$\frac{\partial \mathcal{L}}{\partial x^{N+1}} = \frac{\partial l_{N+1}}{\partial x^{N+1}} - p^{N+1} = 0 \quad (2.4)$$

$$\frac{\partial \mathcal{L}}{\partial p^{i+1}} = -x^{i+1} + f^i(x^i, u^i) = 0, \quad i = 1, \dots, N. \quad (2.5)$$

(Here we have fixed $x^1 = a$ and eliminated it from the problem.) We introduce the following notation, some of which has been used elsewhere (Dunn [3], Dunn and Bertsekas [4]), although defined there in terms of Hamiltonians rather than the Lagrangian.

$$\begin{aligned} Q^i &= \frac{\partial^2 \mathcal{L}}{\partial (x^i)^2}; & R^i &= \frac{\partial^2 \mathcal{L}}{\partial x^i \partial u^i}; & S^i &= \frac{\partial^2 \mathcal{L}}{\partial (u^i)^2}; \\ A^i &= \frac{\partial f^i}{\partial x^i}; & B^i &= \frac{\partial f^i}{\partial u^i}; & t^i &= \frac{\partial \mathcal{L}}{\partial x^i}; & r^i &= \frac{\partial \mathcal{L}}{\partial u^i}. \end{aligned} \quad (2.6)$$

Assumption 1. There exist vectors u_* and x_* which satisfy second-order sufficient conditions for (1.1), that is,

(i) there exists at least one vector p_* such that the equations (2.2–2.5) hold (with all quantities evaluated at (u_*, x_*)),

(ii) for some p_* satisfying (i), and all nonzero directions (v, y) such that

$$\begin{aligned} y^1 &= 0 \\ y^{i+1} &= A^i_* y^i + B^i_* v^i \end{aligned}$$

we have that

$$(v, y) \nabla^2 \mathcal{L}(u_*, x_*, p_*)(v, y) > 0,$$

that is,

$$\frac{1}{2} y^{N+1} Q_*^{N+1} y^{N+1} + \frac{1}{2} \sum_{i=1}^N (y^i Q_*^i y^i + 2 y^i R_*^i v^i + v S_*^i v^i) > 0$$

(where the subscript ‘*’ denotes evaluation at (u_*, x_*, p_*)).

It is well known (see Fletcher [5]) that in the neighbourhood of a point (u_*, x_*) satisfying Assumption 1, SQP is equivalent to applying Newton’s method to (2.2–2.5). In the following, we use v^i , y^i and q^i to denote components of the Newton steps in the control, state and multiplier vectors (δu , δx and δp) respectively. Linearization of (2.2–2.5) gives the following system to be solved to find these quantities:

$$r^i + R^{iT} y^i + S^i v^i + B^{iT} q^{i+1} = 0, \quad i = 1, \dots, N, \quad (2.7)$$

$$t^i + R^i v^i + Q^i y^i - q^i + A^{iT} q^{i+1} = 0, \quad i = 1, \dots, N, \quad (2.8)$$

$$t^{N+1} + Q^{N+1} y^{N+1} - q^{N+1} = 0, \quad (2.9)$$

$$[-x^{i+1} + f^i(x^i, u^i)] - y^{i+1} + A^i y^i + B^i v^i = 0, \quad i = 1, \dots, N. \quad (2.10)$$

Using the notation

$$\hat{r}^i = \frac{\partial l_i}{\partial u^i}, \quad \hat{t}^i = \frac{\partial l_i}{\partial x^i}, \quad \hat{s}^i = -x^{i+1} + f^i(x^i, u^i), \quad p_+ = p + q,$$

(2.7–2.10) can be written alternatively as

$$\hat{r}^i + R^{iT} y^i + S^i v^i + B^{iT} p_+^{i+1} = 0, \quad i = 1, \dots, N, \quad (2.11)$$

$$\hat{t}^i + R^i v^i + Q^i y^i - p_+^i + A^{iT} p_+^{i+1} = 0, \quad i = 2, \dots, N, \quad (2.12)$$

$$\hat{t}^{N+1} + Q^{N+1} y^{N+1} - p_+^{N+1} = 0, \quad (2.13)$$

$$\hat{s}^i - y^{i+1} + A^i y^i + B^i v^i = 0, \quad i = 1, \dots, N. \quad (2.14)$$

The (locally equivalent) SQP form of (2.7–2.10) is of course

$$\min_{\substack{v, y \\ y^1=0}} \sum_{i=1}^N r^i v^i + t^i y^i + \frac{1}{2} (y^i Q^i y^i + 2 y^i R^i v^i + v^i S^i v^i) + t^{N+1} y^{N+1} + \frac{1}{2} y^{N+1} Q^{N+1} y^{N+1} \quad (2.15)$$

$$\text{s.t. } y^{i+1} = A^i y^i + B^i v^i + \hat{s}^i, \quad i = 1, \dots, N.$$

where the q^i are the Lagrange multipliers for the (linear) constraints at the solution of (2.15). The system (2.11–2.14) has a similar quadratic-programming representation, with the optimal Lagrange multipliers being p_+^{i+1} instead of q^{i+1} .

We can summarize the resulting algorithm as follows:

Algorithm I (SQP/Lagrange-Newton).

$$k \leftarrow 0$$

repeat

1. Solve (2.7–2.10) (or the alternatives) for v^i , y^i and q^{i+1} (or p_+^{i+1})
2. Set

$$\begin{aligned} u^i &\leftarrow u^i + v^i, & i = 1, \dots, N \\ x^i &\leftarrow x^i + y^i, & i = 2, \dots, N+1 \\ p^{i+1} &\leftarrow p^{i+1} + q^{i+1} = p_+^{i+1}, & i = 1, \dots, N \end{aligned}$$

3. $k \leftarrow k + 1$

until convergence.

The following convergence result is a standard one for SQP methods; its proof is omitted.

Theorem 2.1. *Let u_* and x_* satisfy the second-order sufficient conditions for (1.1) with optimal Lagrange multipliers p_* . Denote the sequence of iterates produced by Algorithm I as (u_k, x_k, p_k) . Then if the initial estimates (u_0, x_0, p_0) are sufficiently close to (u_*, x_*, p_*) , the estimate*

$$\|(u_{k+1} - u_*, x_{k+1} - x_*, p_{k+1} - p_*)\| = O(\|(u_k - u_*, x_k - x_*, p_k - p_*)\|^2)$$

holds, that is, the sequence $\{(u_k, x_k, p_k)\}$ is Q -quadratically convergent.

One variant of this approach is to update the multipliers p^{i+1} using the formulae (2.3–2.4), rather than setting $p^{i+1} \leftarrow p_+^{i+1}$. The resulting algorithm is

Algorithm II.

$$k \leftarrow 0$$

repeat

1. Solve (2.7–2.10) (or the alternatives) for v^i , y^i and q^{i+1} (or p_+^{i+1})
2. Set

$$\begin{aligned} u^i &\leftarrow u^i + v^i, & i = 1, \dots, N \\ x^i &\leftarrow x^i + y^i, & i = 2, \dots, N+1 \end{aligned}$$

3. Evaluate A^i , t^i , $i = 2, \dots, N$ (at the new (u, x) iterate)
4. $p^{N+1} \leftarrow t^{N+1}$; **for** $i = N, \dots, 2$ $p^i \leftarrow t^i + A^{iT} p^{i+1}$;
5. $k \leftarrow k + 1$

until convergence.

Step 4 is the solution of the adjoint equation, an essential component of the methods in [4,12]. No extra function evaluations are needed here, since A^i and t^i can of course be re-used in Step 1 of the following iteration. However, Step 4 does seem to introduce an inherently sequential element into the algorithm (although some speedup of this step is possible in a parallel environment, using techniques discussed in Section 2.2). The convergence of Algorithm II is slightly different from that of Algorithm I. Denoting the updated value of (u, x) obtained in Step 3 of iteration k as (u_k, x_k) , and the corresponding p obtained in Step 4 by p_k , it is easy to show using an inductive argument based on (2.3–2.4) that

$$\|p_k - p_\star\| = O(\|(u_k - u_\star, x_k - x_\star)\|).$$

This observation leads directly to the following result:

Theorem 2.2. *Let u_\star, x_\star and p_\star be as in Theorem 2.1, and suppose that the starting point (u_0, x_0) is sufficiently close to (u_\star, x_\star) and that the initial p_0 is obtained by performing Steps 3 and 4 of Algorithm II at (u_0, x_0) . Then the sequence of iterates generated by Algorithm II satisfies*

$$\|(u_{k+1} - u_\star, x_{k+1} - x_\star)\| = O(\|(u_k - u_\star, x_k - x_\star)\|^2).$$

Proof.

$$\begin{aligned} & \|(u_{k+1} - u_\star, x_{k+1} - x_\star)\| \\ & \leq \|(u_{k+1} - u_\star, x_{k+1} - x_\star, p_k + q_k - p_\star)\| \\ & = O(\|(u_k - u_\star, x_k - x_\star, p_k - p_\star)\|^2) \text{ by Theorem 1} \\ & = O(\|(u_k - u_\star, x_k - x_\star)\|^2) \text{ by the observation above. } \square \end{aligned}$$

A second variant is obtained by evaluating $x^i, i = 2, \dots, N + 1$ from the original state equations (2.5), rather than by setting $x^i \leftarrow x^i + y^i$:

Algorithm III.

$k \leftarrow 0$

repeat

1. Solve (2.7–2.10) (or the alternatives) for v^i, y^i and q^{i+1} (or p_+^{i+1})
2. Set

$$u^i \leftarrow u^i + v^i, \quad i = 1, \dots, N$$

3. $x^1 \leftarrow a$; **for** $i = 1, \dots, N$ $x^{i+1} \leftarrow f^i(x^i, u^i)$;
4. Evaluate $A^i, t^i, i = 2, \dots, N$ (at the new (u, x) iterate)
5. $p^{N+1} \leftarrow t^{N+1}$; **for** $i = N, \dots, 2$ $p^i \leftarrow t^i + A^{iT} p^{i+1}$;
6. $k \leftarrow k + 1$

until convergence.

Clearly, Step 3 is inherently serial, since each x^{i+1} may depend nonlinearly on x^i . The linear-quadratic subproblem to be solved in Step 1 simplifies a little because of the choice of x and p ; now $\hat{s}^i = 0, i = 1, \dots, N$, and $t^i = 0, i = 1, \dots, N + 1$. Making these substitutions in (2.15) and comparing with Dunn and Bertsekas [4, equation (19)], we see that Algorithm III is simply Newton's method applied to $F(u)$! This is of course known to be quadratically convergent in u . We can also see this directly by using an inductive argument to show that x_k obtained in Step 3 has

$$\|x_k - x_\star\| = O(\|u_k - u_\star\|),$$

and also that p_k obtained in Steps 4 and 5 has

$$\| p_k - p_* \| = O(\| (u_k - u_*, x_k - x_*) \|) = O(\| u_k - u_* \|).$$

Application of Theorem 2.2 then yields the following result, which we state for completeness:

Theorem 2.3. *Let u_* , x_* , p_* be as in Theorem 2.1, and suppose that the starting vector u_0 is sufficiently close to u_* and that the initial x_0 and p_0 are obtained by executing Steps 3, 4 and 5, with $u = u_0$. Then the sequence of iterates generated by Algorithm III satisfies*

$$\| u_{k+1} - u_* \| = O(\| u_k - u_* \|^2).$$

In passing, we note from [4, equations (10b) and (18a)] that the condition (ii) in Assumption 1 is equivalent to positive definiteness of the Hessian $\nabla^2 F(u^*)$, and hence second-order sufficient conditions for (1.1) (regarded as a constrained problem) are equivalent to sufficient conditions for u_* to be a strict local minimizer of $F(u)$.

In all these algorithms, Step 1 typically uses most of the computational effort, since it does the evaluation of most of the functions and derivatives, and also solves a large linear system to find the various components of the step. The evaluations are independent of each other and can be carried out concurrently. Solution of the linear system is more complicated. We describe the options available for this part of the problem with reference to the formulae (2.7–2.10), although, as has been noted, a few terms disappear in Algorithms II and III.

The simplest option is to assemble the equations into a large, symmetric (but indefinite), block-banded linear system, and solve it directly. On shared-memory machines, the parallel band solver discussed in [14] can be used. We explore this possibility in Section 2.2, and show in Section 2.3 that this option is the fastest in many situations.

Polak’s recurrence [12] proceeds by assuming nonsingularity of each S^i , $i = 1, \dots, N$. Equation (2.7) is used to eliminate v^i from the problem. The following algorithm is then used to generate matrices K^i and vectors b^i ($i = 2, \dots, N + 1$) such that

$$q^i = K^i y^i + b^i. \tag{2.16}$$

(We omit details of the derivation.)

Algorithm P.

$$K^{N+1} \leftarrow Q^{N+1}, b^{N+1} \leftarrow t^{N+1};$$

for $i = N, \dots, 2$

$$K^i \leftarrow \tilde{A}^{iT} [I - K^{i+1} \tilde{J}^i]^{-1} K^{i+1} \tilde{A}^i + \tilde{Q}^i$$

$$b^i \leftarrow \tilde{A}^{iT} [I - K^{i+1} \tilde{J}^i]^{-1} (K^{i+1} \tilde{f}^i + b^{i+1}) + \tilde{t}^i$$

where

$$\tilde{A}^i = A^i - B^i T^i R^{iT}$$

$$\tilde{J}^i = -B^i T^i B^{iT}$$

$$\tilde{Q}^i = Q^i - R^i T^i R^{iT}$$

$$\tilde{t}^i = t^i - R^i T^i r^i$$

$$\tilde{f}^i = \hat{s}^i - B^i T^i r^i$$

and

$$T^i = (S^i)^{-1}.$$

Problem 2 (Kelley and Sachs [8]): ($m = 1, n = 1$) This problem is a discretization of the continuous problem

$$\min_{u(t)} \int_0^{0.3} 0.5e^{-t} ((x(t) - 1.5)^2 + (u(t) - 3)^2) + \frac{u^4(t)}{40} dt$$

$$\text{s.t. } \dot{x}(t) = u(t) - t^2, \quad x(0) = 1.$$

Given a value of N , set $h = 0.3/N$ and look for u^i to approximate $u((i - 0.5)h)$ and x^i to approximate $x((i - 1)h)$. The integral can be replaced by a trapezoidal rule, and the state equation by the midpoint rule approximation

$$x^{i+1} = x^i + hu^i - h^3(i - 0.5)^2.$$

Problem 3 (Di Pillo, Grippo and Lampariello [2]): ($m = 1, n = 2$) In the terminology of (1.1, 1.2), and choosing $h = 1/N$,

$$l_1(x^1, u^1) = \frac{5h}{2} [(x_1^1)^2 + (x_2^1)^2] + 5h(u^1)^2$$

$$l_{N+1}(x^N) = \frac{5h}{2} [(x_1^{N+1})^2 + (x_2^{N+1})^2]$$

$$l_i(x^i, u^i) = 5h [(x_1^i)^2 + (x_2^i)^2 + (u^i)^2]$$

$$f^i(x^i, u^i) = \begin{bmatrix} x_1^i + 5h \left[(1 - (x_2^i)^2)x_1^i - x_2^i + u^i \right] \\ x_2^i + 5hx_1^i \end{bmatrix}.$$

Problem 2A: ($m = 1, n = 1$) This is the same as Problem 2, except that each function and derivative evaluation is done 100 times. Hence the computation time is (artificially) dominated by evaluations rather than by the linear algebra.

The timings for a single iteration with for N set to 1000, are shown in *Table 1*. For the first three problems, the cost of the linear algebra tends to dominate. This can be deduced by comparing the times for the Dunn–Bertsekas-based algorithms II (fully concurrent evaluations) and III(S) (fully serial evaluations). In going from Algorithm II to Algorithm I, the speedup of about 2 is entirely due to the use of the parallel bandsolver. This speedup is similar to the relative times for an 8-processor implementation of the parallel bandsolver and one-processor Gaussian elimination, as obtained in [14]. (The computational overhead in operating on the partitioned system is quite high, so that although the parallel algorithm makes efficient use of the multiple processors, it does not attain a linear speedup relative to the best serial algorithm. This is discussed in detail in [14].) Slightly better speedup is obtained for Problem 2, in which the system (2.17–2.18) has smaller bandwidth than in Problems 1 and 3, and again this is consistent with the results in [14].

Table 1
Times for one iteration of each algorithm on the Alliant FX/8 in seconds, $N = 1000$.

	Prob. 1	Prob. 2	Prob. 3	Prob. 2A
Alg. I	0.382	0.155	0.387	1.24
Alg. II	0.730	0.352	0.731	1.44
Alg. III	0.731	0.407	0.733 *	5.54
Alg. III(S)	0.765	0.462	0.776 *	9.02
Alg. III-P(S)	1.46	0.655	1.47 *	9.21

* = failed to converge.

In Problem 2A, the function evaluations dominate the computation, and we see a progressive improvement in going from the serial algorithm III(S), to the ‘partially parallel’ III, to the ‘fully parallel’ II. The advantage gained by using the parallel band solver (Algorithm I) is smaller in a relative sense than in the other problems. However, it is clear by comparing Algorithm I with Algorithm III(S) that a highly efficient 8-processor implementation has been achieved.

In summary, on an 8-processor machine, we can obtain an algorithm which is two to three times faster than the best serial implementation of Newton’s method, for problems in which the linear algebra dominates, and almost eight times faster than Newton’s method, when the function evaluations dominate.

We stress again that only the ‘no-frills’ locally convergent algorithms have been implemented above. For practical codes, it will be necessary to include features which ensure global convergence. Also, it may be desirable to make use of the chord method, in which the second-derivative matrices are not calculated (and factorized) on every iteration.

3. The constrained problem

3.1 Active set methods

Inclusion of the general constraints (1.3) gives a more difficult problem, which can again be solved in a sequential quadratic programming framework. We sketch the approach in this section. There are two variants of SQP for handling inequality constraints, known in the literature as ‘IQP’ and ‘EQP’. In IQP, the constraints (1.3) are linearized at each step, and the quadratic subproblem (2.15) is solved subject to the additional constraints

$$g^i(x^i, u^i) + y^i \frac{\partial}{\partial x} g^i(x^i, u^i) + v^i \frac{\partial}{\partial u} g^i(x^i, u^i) \leq 0, \quad i = 1, \dots, N. \quad (3.1)$$

The other terms r^i , t^i , Q^i , R^i , and S^i in (2.15) need to be redefined in terms of a Lagrangian with an additional term:

$$\begin{aligned} \mathcal{L}(u, x, p, \mu) = & l_{N+1}(x^{N+1}) + \sum_{i=1}^N l_i(x^i, u^i) - \sum_{i=1}^N (x^{i+1} - f^i(x^i, u^i)) p^{i+1} \\ & + \sum_{i=1}^N g^i(x^i, u^i) \mu^i. \end{aligned}$$

The quadratic programming problem (2.15, 3.1) can be solved using an active set approach, where at each minor iteration a subset of the components of (3.1) are chosen to hold at equality, and the remainder are ignored, at least until a line search is done. Hence (2.15) may be repeatedly solved subject to constraints of the form

$$\begin{aligned} g_j^i(x^i, u^i) + y^i \frac{\partial}{\partial x} g_j^i(x^i, u^i) + v^i \frac{\partial}{\partial u} g_j^i(x^i, u^i) = 0, \\ j \in \Lambda_i \subset \{1, \dots, \nu_i\}, \quad i = 1, \dots, N. \end{aligned} \quad (3.2)$$

Writing down the first order conditions for this subproblem, we obtain a set of linear equations in v^i , y^i , q^i , and μ_j^i ($j \in \Lambda_i$) which is obtained by augmenting (2.7–2.10) with appropriate terms in μ^i , and including the equations (3.2). Again these equations can be solved via recurrence relations (see Ohno [9], or Dunn [3] and Pantoja and Mayne [10] for the case in which the g^i are functions of u^i only). Alternatively, after an appropriate ordering of equations and variables, a block-banded system is again obtained, which can again be solved with the parallel band solver.

In the EQP approach, the active constraint sets Λ_i are fixed during each outer iteration, rather than at each minor iteration as in IQP. The innermost iteration is identical in form to that of IQP, namely (2.15, 3.2), and hence can be solved as before. The relative merits of IQP and EQP have been extensively discussed elsewhere (see for example Powell [13]).

We do not pursue these approaches further here, since they are often unsuitable for large-scale problems. The reason for this is that the active set $\Lambda = \{\Lambda_1, \dots, \Lambda_N\}$ is generally only allowed to change by a small number of components at each (major or minor) iteration. Hence if the initial estimate of Λ is far from the optimal active set, many iterations will be required.

Once the correct active constraint manifold has been identified, these algorithms are R-quadratically convergent in u .

3.2 The simple control-constrained case – a gradient projection method

An alternative to active set methods, which is attractive when the feasible region defined by the constraints (1.3) is convex and has simple geometry, is the two-metric gradient projection algorithm (see Bertsekas [1], Gafni and Bertsekas [6] and Dunn [3]). This method has the advantage that large changes to the active set are allowed at each iteration, while it retains the good local convergence properties of the active set algorithms.

We sketch this approach for the important special case in which the constraints have the form (1.4) (see Bertsekas [1] for the details). The problem (1.1, 1.2) is regarded as a function of u alone (i.e. $F(u)$). As in Algorithm III of Section 2, x is obtained from the state equations (1.2) and p is obtained from the adjoint equation (2.3), so that in (2.15), $t^i = 0$ and $\hat{s}^i = 0$. At each iteration, an ‘almost-active set’ $\bar{\Lambda} = \{\bar{\Lambda}_1, \dots, \bar{\Lambda}_N\}$ is constructed, where $\bar{\Lambda}_i$ contains the indices of the components of u^i which at or near their bound $\pm\eta$, and for which the corresponding component of $\nabla F(u)$ has the ‘correct’ sign (i.e. $\partial F/\partial u_j^i > 0$ if $u_j^i \approx -\eta$ and $\partial F/\partial u_j^i < 0$ if $u_j^i \approx \eta$). (2.15) is then solved, using the definitions of Q^i , R^i , S^i and r^i from (2.1) and (2.6), subject to the additional constraints

$$v_j^i = 0, \quad j \in \bar{\Lambda}_i, \quad i = 1, \dots, N. \tag{3.3}$$

This process yields the ‘Newton’ part of the step. The ‘first-order’ part can be obtained by setting

$$v_j^i = -\sigma_j^i \frac{\partial F}{\partial u_j^i}, \quad j \in \bar{\Lambda}_i, \quad i = 1, \dots, N$$

(where $\underline{\sigma} \leq \sigma_j^i \leq \bar{\sigma}$, $\underline{\sigma}$ and $\bar{\sigma}$ being positive constants; we choose $\sigma_j^i \equiv 1$). A search is then carried out along the piecewise-linear path $P(u + \alpha v)$, $\alpha \in [0, 1]$, where P denotes projection of a vector in R^{km} onto the box $\{w \in R^{km} \mid w_i \in [-\eta, \eta]\}$. An Armijo linesearch strategy is used, that is, the values $\alpha = 1, \beta, \beta^2, \dots$ ($\beta < 1$) are tried, until a sufficient decrease criterion is met. For this value of α , u is reset to $P(u + \alpha v)$ for the next iteration. The values of x and p are recalculated accordingly.

Solution of the constrained problem (2.15, 3.3) can be done using a modified version of the recurrences in Algorithm DB (see Dunn [3]), which we state below. Let P^i be the matrix

$$P^i = [e_j^T]_{j \in \bar{\Lambda}_i},$$

where e_j is the j th column of the $m \times m$ identity matrix. Also for any $m \times m$ matrix S , denote $P^i S P^{iT}$ by $[S]_{ii}$.

Algorithm DB(C)

$$\Theta^{N+1} = Q^{N+1}; \quad \beta^{N+1} = t^{N+1};$$

for $i = N, \dots, 1$

$$\begin{aligned} \Gamma^i &\leftarrow -[S^i + B^{iT}\Theta^{i+1}B^i]_{ii}^{-1}P^i(R^{iT} + B^{iT}\Theta^{i+1}A^i) \\ \gamma^i &\leftarrow -[S^i + B^{iT}\Theta^{i+1}B^i]_{ii}^{-1}P^i(r^i + B^{iT}(\Theta^{i+1}\hat{s}^i + \beta^{i+1})) \\ \Theta^i &\leftarrow (Q^i + A^{iT}\Theta^{i+1}A^i) + (R^i + A^{iT}\Theta^{i+1}B^i)P^{iT}\Gamma^i \quad (i \neq 1) \\ \beta^i &\leftarrow A^{iT}\Theta^{i+1}(B^iP^{iT}\gamma^i + \hat{s}^i) + A^{iT}\beta^{i+1} + t^i + R^iP^{iT}\gamma^i \quad (i \neq 1) \end{aligned}$$

The steps P^iv^i and y^i can be recovered as in the loop after Algorithm DB.

An alternative method for solving (2.15, 3.3) is to use a reduced version of the linear system (2.17, 2.18). We can simply delete rows and columns of this matrix which correspond to the indices $j \in \bar{A}_i, i = 1, \dots, N$. The resulting coefficient matrix will have dimension between $2n(N + 1)$ and $2n(N + 1) + mN$. In our implementation we opt for the simpler approach of ‘eliminating’ the variables by replacing the rows and columns in question by zeros (except for a one on the diagonal) and by inserting a zero in the appropriate right-hand side position. This is done because (a) we wish to avoid disturbing the special structure of (2.17) and its associated computational savings, (b) the overhead involved in compressing the elements of (2.17) to give a reduced system would be substantial, and (c) when m and n are similar, the potential reduction in the size of the matrix is not very great.

3.3 Numerically results

We report on results for three implementations of the algorithm described above. In all cases, since gradient projection is being applied to F as a function of u alone, the states x^i must be eliminated using the state equations (1.2), that is, the functions f^i must be evaluated in sequence.

Algorithm GP(B): Step P^iv^i obtained by solving (2.17–2.18) (modified by (3.3)) using the parallel bandsolver. All functions except f^i and l_i evaluated concurrently on 8 processors.

Algorithm GP: Step P^iv^i obtained using Algorithm DB(C). All functions except f^i and l_i evaluated concurrently on 8 processors.

Algorithm GP(S): One-processor implementation of Algorithm GP.

Results are obtained for four problems. All algorithms gave an identical sequence of iterates, and were terminated when the projected gradient norm fell below 10^{-10} . In all cases $N = 1000$, and a starting point of $u = 0$ was used.

Problem 1 from Section 2.3, with $\eta = 1$. Converges in 3 iterations, with 183 components of u active.

Problem 2 from Section 2.3, with $\eta = 2$. Converges in 3 iterations, with 740 components of u active.

Problem 2 from Section 2.3, with $\eta = 3$. Converges in 5 iterations, with no components of u active.

Problem 2B with $\eta = 2$. Same as problem 2 from Section 2.3, except that function evaluation costs are artificially magnified by doing each evaluation (except for f^i and l_i) 100 times.

Total times to solution on the Alliant FX/8 are given in Table 2. Again, the cost of the linear algebra dominates in the first three columns. The use of the parallel band solver gives a

Table 2
Times to solution on the Alliant FX/8 in seconds, $N = 1000$.

	Prob. 1 ($\eta = 1$)	Prob. 2 ($\eta = 2$)	Prob. 3 ($\eta = 3$)	Prob. 2B ($\eta = 2$)
Alg. GP(B)	1.69	1.01	1.55	2.68
Alg. GP	2.42	1.29	2.67	2.91
Alg. GP(S)	2.59	1.52	3.01	14.4

speedup of somewhat less than that obtained for the corresponding unconstrained problems. This can be attributed to the fact (mentioned above) that the band solver in Algorithm GP(P) does not at present take advantage of the reduced dimension of the Newton step. (Hence the fewer constraints active, the better the speedup – compare the second and third columns.) Best results are obtained for Problem 2B, which is meant to typify applications in which the state equations f^i are easily evaluated (e.g. linear), but evaluation of the objective functions and their derivatives is expensive. A speedup of 5.4 is obtained by the best parallel algorithm GP(B) over the best sequential algorithm GP(S).

3.4 A hybrid approach

As noted above, the need for evaluation of the nonlinear state equations (1.2) between iterations of the gradient projection method presents a possible serial bottleneck in the simple control-constrained case, since this entails evaluating the functions f^i in sequence. A potentially more parallelizable approach would be to use a two-level method. Sequential quadratic programming (IQP variant) could be applied to the original nonlinear problem (1.1, 1.2, 1.4), yielding a subproblem in which the state equations are linear, and in which the bound constraints (1.4) still appear. This subproblem could then be solved using a gradient projection method. The state equations still need to be evaluated between each gradient projection iteration, but since they are now linear, this step can be implemented in parallel using much the same techniques which were used to parallelize the banded system solver.

4. Conclusions and further work

We have discussed multiprocessor implementations of some basic rapidly-convergent algorithms for discrete-time optimal control problems, which attempt to exploit parallelism both in the linear algebra and in evaluation of the functions. Further work remains to be done to implement these methods on message-passing machines (e.g. hypercubes and massively parallel architectures). Additionally, algorithms with better global convergence properties are needed. Simple techniques for ensuring global convergence of Newton's method and gradient projection methods are quite well understood; for SQP more complicated devices are needed. The challenge is to implement the globalization strategies without interfering too much with the parallelism which has been obtained in the basic algorithms above.

References

- [1] D.P. Bertsekas, Projected Newton methods for optimization problems with simple constraints, *SIAM J. Control Optim.* 20 (1982) 221–246.
- [2] G. Di Pillo, L. Grippo and F. Lampariello, A class of structured quasi-Newton algorithms for optimal control problems, in: *Proc. IFAC Applications of Nonlinear Programming to Optimization and Control* (1983).
- [3] J. Dunn, A projected Newton method for minimization problems with nonlinear inequality constraints, *Numer. Math.* 53 (1988) 377–409.
- [4] J. Dunn and D.P. Bertsekas, Efficient dynamic programming implementations of Newton's method for unconstrained optimal control problems, *J. Optim. Theory Appl.*, 63 (1989) 23–38.
- [5] R. Fletcher, *Practical Methods of Optimization (2nd edition)* (Wiley, New York, 1987).
- [6] E.M. Gafni and D.P. Bertsekas, Two-metric projection methods for constrained optimization, *SIAM J. Control Optim.* 22 (1984) 936–964.
- [7] D.H. Jacobson and D.Q. Mayne, *Differential Dynamic Programming* (American Elsevier, New York, 1970).
- [8] C.T. Kelley and E.W. Sachs, A pointwise quasi-Newton method for unconstrained optimal control problems, *Numer. Math.* 55 (1989) 159–176.

- [9] K. Ohno, A new approach to differential dynamic programming for discrete time systems, *IEEE Trans. Automat. Control* AC-23 (1978) 37–47.
- [10] J.F.A. Pantoja and D.Q. Mayne, A sequential quadratic programming algorithm for discrete optimal control problems with control inequality constraints, submitted, 1988.
- [11] L.J. Podrazik and G.G.L. Meyer, A parallel first-order linear recurrence solver, *J. Parallel Dist. Comp.* 4 (1987) 117–132.
- [12] E. Polak, *Computational Methods in Optimization* (Academic Press, New York, 1970).
- [13] M.J.D. Powell, ed., *Nonlinear Optimization 1981* (Academic Press, New York, 1982).
- [14] S.J. Wright, Parallel algorithms for banded linear systems, Argonne National Laboratory Preprint MCS-P64-0289, 1989, *SIAM J. Sci. Statist. Comput.*, to appear.