towards global optimisation

edited by

l. c. w. dixon and g. p. szegö

TOWARDS GLOBAL OPTIMISATION

Proceedings of a Workshop at the University of Cagliari, Italy, October 1974

Editors

L.C.W. DIXON

The Numerical Optimisation Centre
The Hatfield Polytechnic
Hatfield, U.K.

and

G.P. SZEGÖ

Professor of Mathematics
University of Venice
Venice, Italy

N·H
P.C.

1975

NORTH-HOLLAND PUBLISHING COMPANY – AMSTERDAM · OXFORD
AMERICAN ELSEVIER PUBLISHING CO., INC. – NEW YORK

North-Holland ISBN: 0 7204 2463 1
American Elsevier ISBN: 0 444 10955 2

PUBLISHERS:
NORTH-HOLLAND PUBLISHING COMPANY – AMSTERDAM
NORTH-HOLLAND PUBLISHING COMPANY, LTD. – LONDON

SOLE DISTRIBUTORS FOR THE U.S.A. AND CANADA:
AMERICAN ELSEVIER PUBLISHING COMPANY, INC.
52 VANDERBILT AVENUE, NEW YORK, N.Y. 10017

Library of Congress Cataloging in Publication Data
Main entry under title:

Towards global optimisation.

(North-Holland mathematical library)
1. Mathematical optimization. I. Dixon, Laurence Charles Ward. II. Szegö, G. P.
QA402.5.T7 519.4 74-28995
ISBN 0-444-10955-2

Printed in The Netherlands

PREFACE

This volume contains, amongst others, most of the papers that were presented at the Workshop on Global Optimisation and Related Topics held at the University of Cagliari, Italy from 4 October to 9 October 1974. This conference was organised by the Cagliari section of the UMI (Unione Matematica Italiana) with the support of the Regione Sardegna.

The speakers at the Workshop were F Archetti, C G Broyden, L C W Dixon, E Fagiuoli, M Gaviano, J Gomulka, J W Hardy, S E Hersom, P Mazzoleni, E Spedicato, C Sutti, G P Szegő and G Treccani.

Many papers presented in this book represent results obtained during the initial stages of a joint research program on Global Minimisation sponsored by the CNR in Italy and by the SRC at the Numerical Optimisation Centre, The Hatfield Polytechnic, Hatfield, United Kingdom.

The papers in this book are divided into five sections. The first section consists of a single paper by S E Hersom. In this he first discusses some important aspects of modelling industrial problems for use with optimisation algorithms. This discussion is based on the experience obtained at the Numerical Optimisation Centre, The Hatfield Polytechnic, in carrying out their consultancy activity with British industry. He continues by describing the NOC Optima package of algorithms that have been developed to solve these problems.

This is followed by the main section of the book which is devoted to the Global Optimisation problem. Whilst no available technique can claim to guarantee to obtain the solution of all such problems, the first paper of this section is a review of the state of the art in which an attempt is made to describe the ideas and assumptions behind the various previously published work in this field and indicate when they might have a chance of success. This review is followed by eight papers giving new results on the Global Optimisation problem.

The third section is devoted to unconstrained local optimisation. Many people approach the global problem by running a technique that locates a local minimum from a number of random starting points. Features that affect the efficiency of a local algorithm are, therefore, of considerable interest to the global optimisers as well as in their own right. Eight papers are included on this problem. These include a review of the recent developments in efficient variable metric methods by E Spedicato and a discussion on when it is advisable to use specialised algorithms for the very common Sums of Squared terms problem is presented by J J McKeown, and extended to the parameter estimation problem in dynamic systems in the next paper by Stott & James. Other new results on the local problem are given in the papers by Dixon, Sutti and Fagiuoli.

The treatment of constraints is discussed in Section 4. Most practical problems contain constraints of one type or another and hence eventually this treatment will have to be included in global optimisation packages. This section is fairly short and contains a review paper followed by three describing new results.

The final section is devoted to the presentation of some recent results on the application of the theory of dynamic systems to the investigation of various problems in minimisation theory. In particular, the first two papers by Castillo and Rodriguez apply the theory of discrete semi-dynamical systems to the investigation of the convergence properties of iterative algorithms. They follow a method proposed by G P Szegö and G Treccani in the volume "Minimisation Algorithms", edited by G P Szegö and published by Academic Press, New York. In these two papers the authors propose two different sets of axioms and describe the system in different spaces. The final paper by G Treccani provides a description of the critical points of a continuously differentiable function and contains results that have been available before in Italian but not translated into English.

Similarly, the editors took the opportunity of inviting Professor J Mockus of the University of Vilnius, USSR, to contribute his paper to the book, in which he describes the Bayesian approach to Optimisation, which has been developed in a series of papers available in Russian but not readily available in English.

The editors would like to record their thanks to all the authors who have contributed to this volume, and especially to Professor G Aymerich, Rector of the University of Cagliari and Director of the Institute of Mathematics of the Faculty of Science at the University of Cagliari for acting as host to the Workshop, and to Professor O Montaldo, Director of

the Cagliari Section of the Unione Matematica Italiana and Director of the Institute of Mathematics of the Faculty of Engineering at the University of Cagliari, for his active sponsorship and participation in the proceedings.

We would also like to express our sincere thanks to Dr Paulo Dolci of the University of Cagliari for all his work in the organisation of the Workshop and to Mrs Mary Callan for the painstaking typing of the final masters from which this book has been produced.

L C W Dixon
G P Szegö

10 October 1974

C O N T E N T S

PART ONE: APPLICATIONS

PAPER 1

THE PRACTICE OF OPTIMISATION

by S E Hersom, The Numerical Optimisation Centre, The Hatfield Polytechnic

ABSTRACT: In this paper some of the experience gained at the Numerical Optimisation Centre during our consultancy activities is discussed. Certain important features in the building of models for optimisation purposes are outlined and the algorithms developed at the Centre for use on problems of a general nature in this field are described. These algorithms are available in the NOC OPTIMA package.

1 INTRODUCTION

Since 1968 the Numerical Optimisation Centre has been developing optimisation algorithms for the prime purpose of solving practical problems. This has entailed a considerable effort both in the theoretical justification, and in the development, of computer procedures, and has resulted in the acquisition of much hard-earned experience in a variety of situations. As one outcome of their activity, the NOC has armed itself with a set of routines, collectively known as NOC - OPTIMA, which it applies to these practical problems as they arise.

Details of these routines are given later in this paper, but it is necessary in order to understand the reasons for developing this particular set, to describe both the types of problems which have arisen and the mathematical theory on which the algorithms are based. Accordingly a group of somewhat simplified case studies are presented from which many necessary and desirable features of the algorithms can be deduced.

These case studies are drawn from various engineering industries, as, indeed, are most of the other practical problems with which the NOC has had to deal. There has not been any deliberate selection of problems to achieve this bias. No doubt there are good historical and other reasons which have contributed to this state of affairs, but the main reason is thought to be the comparative ease with which people engaged in engineering and the associated sciences can formulate their problems in a manner suitable for the application of optimisation techniques. In fact, one of the chief lessons to be learnt from the NOC's experience is the importance of correctly formulating the problem. Since this invariably depends on the mathematical model of some practical situation, a section of this paper is devoted to the necessary and desirable features of these models, and these features complement those of the algorithms.

One of the case studies points the way in which development should proceed, and this is towards the solution of problems which have several local optima in the region of interest, ie. the necessity to develop Global Minimisation Techniques. It is in this context that the author welcomes the effort being made by other workers as is evidenced by the contributions to this conference.

2 PRACTICAL PROBLEMS

When a problem can be posed as the determination of a set of values for the elements of a vector $\underline{x}$ such that the values of some scalar objective function $f(\underline{x})$ has a maximum or minimum value, then an optimisation algorithm might well be employed to arrive at the solution. Such problems frequently arise in the general field of engineering design and the following is a selection of such applications.

2.1 Machine Tool Design (ref. 1)

Figure 1 gives a diagrammatic sketch of a machine tool removing material from a workpiece. The tool itself (a cutter) is rotated at a speed, v. The workpiece is moved relative to the tool at a speed fv and so if the depth of cut is d, the rate of removal of material is proportional to vfd. The problem posed was: if the rate of wear of the tool could be measured, how could the machine be operated at a minimum cost?

First of all, therefore, a measure of cost had to be established. This was taken to be the cost per unit volume of material removed from the workpiece and it comprised two terms, one being the cost of an operator attending the machine and the other that of repairing or replacing the tool. The cost, C, therefore has the form

$$C = (R + T_c/T_1)/vfd$$

where R is the operator's rate, T_c is the cost of the tool and T_1 is its life. Since the real-time measurement of this tool life was only hypothetical at this stage a plausible mathematical expression had to be given to T_1 in order to simulate an actual situation. After some study of practical results and other people's experience, the client proposed that T_1 should be proportional to $1/d^{0.6}f^{1.7}v^{3}$.

Mathematically, therefore, the problem resolves into the determination of the minimum of C by the appropriate selection of f and v (d was assumed constant since, once the machine had been set to a given depth of cut, it would not be practical to adjust it during any one particular operation). Simple mathematics showed that f and v should be made very large, and so it was at this point other considerations had to be introduced.

Firstly f, which is in practice a setting in a gear-box, must lie within the upper and lower limits as provided in the machine. Further, the power obtainable from the motor had an upper limit P_{MAX} and if the stress in the tool was too great, the tool would shear. Finally, the direction of

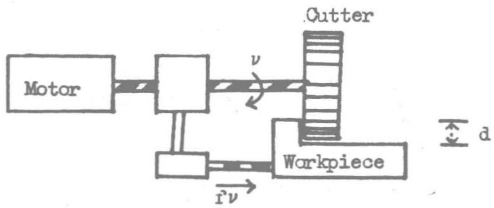

Fig. 1 Machine Tool

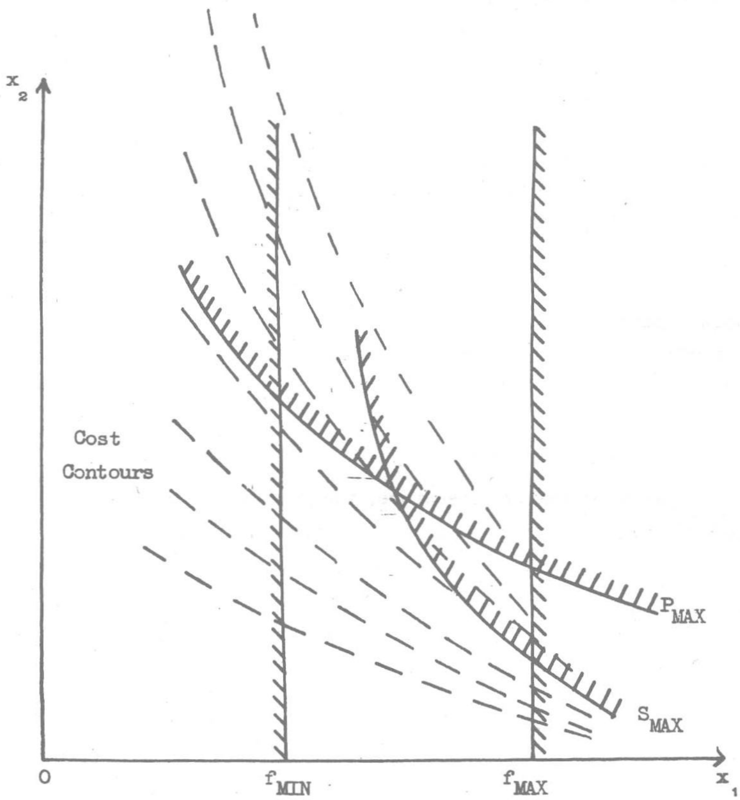

Fig. 2 Machine Tool Cost Contours and Constraints

cutting must always be positive.

Mathematically this problem can be expressed as

maximise $f(\underline{x}) = (R + T_C/T_1)/(x_1 x_2 d)$

where $\underline{x}^T = (f, v)$

$T_1 = k/(d^{0.6} x_1^{1.7} x_2^3)$

subject to $x_1 \leqslant f_{MAX}$

$x_1 \geqslant f_{MIN}$

$x_1^{0.8} x_2 d \leqslant P_{MAX}$

$x_1 x_2^2 \leqslant S_{MAX}$

$x_2 \geqslant 0$

Fig. 2 shows the situation in the $x_1 x_2$ plane. The feasible region for the point (x_1, x_2) is that between the upper and lower limits for x_1, and below one or other of the constraint curves $x_1^{0.8} x_2$ = constant and $x_1 x_2^2$ = constant. The lines of equal costs are also indicated. The optimisation problem is therefore to find the cost contour which has least cost and which just enters the feasible region.

2.1.1 Comments

Once the decision on the form of the cost function and those of supporting functions (the tool life, stress, etc.) had been taken, the resulting expression was comparatively simple mathematically. The complications arose on account of the constraints. Note also that the "obvious" constraint, $x_2 \geqslant 0$ must also be included since otherwise an automatic optimisation routine might find a condition, with x_2 negative, producing a negative life for the tool and other nonsensical conditions.

Fig. 2 is deceptively simple. It might be imagined that the least cost point might be seen "by eye". It must be pointed out, however, that both the cost contours and the constraint curves have only been sketched in and have not been accurately calculated. In fact, a major reason for adopting optimisation procedures is to avoid the necessity of calculating such curves.

2.2 Electrical Filter (ref. 2)

The design specification of an electrical filter often incorporates upper and lower limits to its response, either in time or across a band of frequencies. Fig. 3 depicts a typical response requirement for the attenuation between two frequencies. By selecting a set of 'design points' in this range of frequencies and by deciding on a plausible circuit configuration then, denoting the values of the circuit elements which have to be determined to meet the specification by the elements of $\underline{x}$ and the difference in the value of the filter characteristic from the its design point by $s_i(\underline{x})$, an 'ideal' design can be defined as one in which all the s's have been reduced simultaneously to zero by some choice of $\underline{x}$. However, the number of design points almost always exceeds the number of circuit elements involved and so the 'ideal' filter is not in general obtainable. A compromise has to be found, and one way is to choose $\underline{x}$ to minimise some function of the s's such as

$$f(\underline{x}) = \sum_i s_i^2 (\underline{x}).$$

In this way it is hoped that the response will be close to the 'ideal'.

This approach has been used in the past and, in fact, it is still being used today, but it has often been found to result in designs which require close tolerance components if the performance of production models is to remain within specification. This means that, as shown in Fig. 3, a small change in one or more components of $\underline{x}$ results in the response curve departing rapidly from the optimum and violating a specification limit. There is, indeed, nothing in the formulation of the problem which ensures that the resulting design will ever be within the specification.

In these circumstances it is necessary to reformulate the problem, and since component tolerances caused the downfall of the first approach, an alternative suggestion is to determine those values of the components which give these components their maximum tolerance whilst ensuring that the response of the filter is always within specification. Posed therefore as an optimisation problem, the objective function is a measure of the tolerances of all the components and, since it is possible in general to maximise only one quantity at a time, this must be just one single measure of these tolerances. The precise measure can depend on the individual circumstances, but a simple method is to give each component a tolerance proportional to one number, K, and then make the objective

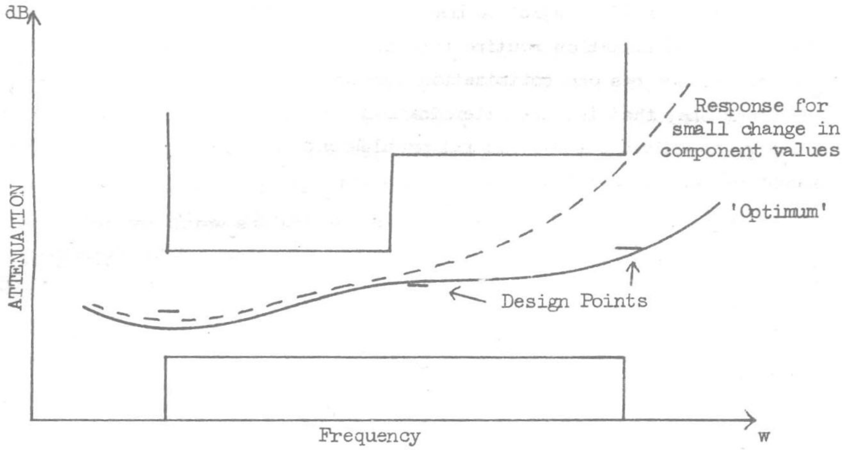

Fig. 3 Electrical Filter

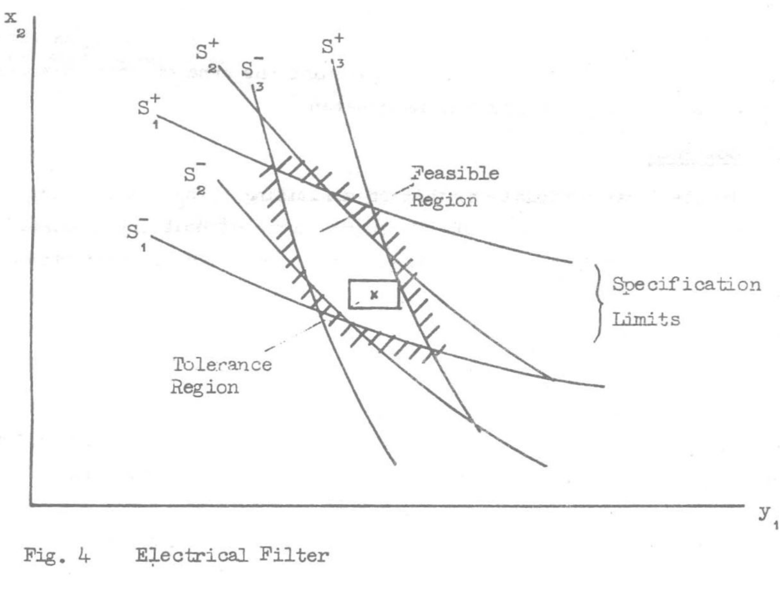

Fig. 4 Electrical Filter

function

$f(\underline{x}) = \mathrm{Max}(K)$ subject to meeting the specification for a fixed $\underline{x}$. The overall optimisation routine then determines $\underline{x}$ which maximises $f(\underline{x})$. The problem now has one optimisation problem nested inside another, but the inner one, that is, the determination of the maximum value of K for a fixed $\underline{x}$, is only a one-dimensional problem and so requires only a modest amount of computing resources.

The specification is now a set of constraints which can be formulated as follows:

We still consider our design points which represent the 'ideal' filter and calculate the $s_i(\underline{x})$ as before. At this point, at the ith frequency, we must have

$$s_i^- \leqslant s_i(\underline{x}) \leqslant s_i^+$$

where s_i^+ and s_i^- are the differences in response from this 'ideal' point and the upper and lower limits respectively. Similar limits apply at all the other frequencies, and so we have the situation, as shown in Fig. 4 for the 2-component case, where the curves labelled s_1^-, s_1^+, s_2^-, s_2^+.... are those for which $s_1(\underline{x}) = s_1^-$, etc. The 'feasible' region, ie. the region in x_1, x_2 for which the filter meets the specification, is shown as that lying within these curves. The maximum tolerance region around $\underline{x}$ is shown, and the position of $\underline{x}$ must be found to maximise the area of this tolerance region whilst keeping its proportions constant.

2.2.1 Comments

In its first formulation, that of minimising $\sum s_i^2$, the problem took the form of minimising a sum-of-squared terms without any constraints. This is a very common form for optimisation problems and one that often gives rise to 'tolerance' considerations, but further discussion on this point is deferred until later.

Perhaps a more important aspect is that after experimenting with the initial formulation a rather more realistic formulation was then derived. Minimising the sums-of-squares of the differences from the 'ideal' filter characteristic is no more than a mathematical technique to produce a 'good' design by numerical methods. When, however, the tolerances are maximised, the problem is much closer to the practical situation and a tangible benefit can be derived. Since tolerances have a direct link with cost, the difference between an optimal design and any other can be readily expressed in monetary terms.

2.3 Transistor Life-Test Model (ref. 3)

The characteristics of a semi-conductor are functions both of the life of the device and of the temperatures at which it has been operating. In a particular application it was required to predict the threshold voltage shift of a transistor, when running at normal temperature, from data acquired from a batch of similar transistors when run at elevated temperatures. This data was obtained by making an initial room-temperature measurement of the average threshold voltage followed by a 24-hour operating period at 100°C. After a return to room temperature for further measurement, the temperature was raised to 120°C for the same dwell time and the cycle of events repeated with 20°C increments up to 300°C. The whole process was repeated with a second batch of transistors but now employing a dwell time of 168 hours.

The model assumed the voltage shift, ΔV, at constant temperature, T, to be given by

$$\Delta V = \Delta V_{SAT} \left[1-\exp \left\{ - (t/\tau)^{\frac{1}{3}} \right\} \right]$$

where

$$\tau = \tau_o \exp \left(e\phi/kT\right)$$

$$\Delta V_{SAT} = \Delta V_o \exp \left(-e\theta/kT\right)$$

and where e and k are the electron charge and Boltzmann's constant respectively. ϕ and θ (activation energies), the time constant τ_o and the voltage ΔV_o are the parameters of the model which are to be determined so that it fits the measurement data, ie. these are the elements of the vector $\underline{x}$. This model was used to simulate the actual experiments and to obtain a series of calculated voltage shifts, $\Delta V(i)$, which could be compared with the actual measured shifts, $\Delta V_d(i)$. The optimisation routine was used to determine the values of the parameters which minimised

$$f = \sum_i \left[\left\{ \Delta V(i) - \Delta V_d(i) \right\} / \Delta V_d(i) \right]^2$$

Since there were 11 measuring points in each test, there were in all 22 terms in the summation.

The model was devised by the client before approaching the NOC and represented the outcome of detailed considerations of the physical