



ELSEVIER

Mathematics and Computers in Simulation 42 (1996) 135–142

MATHEMATICS
AND
COMPUTERS
IN SIMULATION

Parameter estimation in dynamic systems

Walter J.H. Stortelder *

CWI, Department of Numerical Mathematics, PO Box 94079, 1090 GB Amsterdam, The Netherlands

Abstract

In this paper we present an approach to the solution of parameter estimation problems in systems described mathematically by differential algebraic equations. The numerical solution of these equations is compared with measurements from experiments. By adaptation of the parameters in the differential algebraic equations we try to fit the solution to the measurements. For the fit we use a (weighted) least squares criterion. Not only the final estimate of the unknown parameters, but also additional information about the reliability is derived.

A computer program, with an interactive graphical user interface, has been developed to steer through the computation in order to influence the precise formulation of the model and the data and to see the numerical results by direct visualisation.

1. Introduction

In many areas of the experimental sciences, such as chemistry, biology or environmental engineering, where mathematical models are used to describe a process under consideration and experiments are performed to validate these models, an increasing demand for mathematical support has arisen. The models concern time-dependent processes with numerous state variables and many interactions, and feedback loops between the variables. They can be described mathematically by a set of differential algebraic equations (DAEs). It happens frequently that these model equations contain a number of unknown parameters, e.g. reaction constants (growth rates, etc.). The aim is to determine these unknown parameters such that an optimal fit is accomplished between the measurements from the laboratory or the plant and the theoretical results obtained by solving the set of DAEs.

After the mathematical formulation of the problem, we present the numerical techniques which we use to cope with this problem. An additional section about the statistical background is followed by a case study from biochemistry.

2. Mathematical formulation

The model equations are denoted by a system of differential equations (DAEs) as

* E-mail: walterst@cwi.nl. Supported by the Dutch Technology Foundation (STW).

$$E \frac{dy}{dt} = f(t, y, \theta), \quad \text{with } y(t_0, \theta) = y_0(\theta), \quad (1)$$

where t is the time, θ is an m -dimensional vector of unknown parameters, $y(t, \theta)$ is an n -dimensional state vector depending on t and θ , $f(t, y, \theta) : \mathbb{R} \times \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}^n$ and the diagonal, $n \times n$ -matrix E denotes which equations are algebraic: E_{ii} equals 1 if the i th equation is a differential equation and 0 if it is an algebraic equation.

In order to determine the unknown parameters, some measurements, say N , are available on the process under consideration. Each measurement is characterised by a triple

$$(c_i, t_i, \hat{y}_i), \quad 1 \leq i \leq N,$$

where c_i indicates which component of the state vector, y , has been measured, t_i is the time of the measurement and \hat{y}_i is the measured value. It is possible to have more measurements at the same time: $t_i = t_{i+1}$. A necessary condition reads $N \geq m$.

An obvious approach is to estimate the unknown parameters such that the (weighted) sum of squared discrepancies between the calculated and the measured values,

$$S(\theta) = \sum_{i=1}^N w_i^2 (y_{c_i}(t_i, \theta) - \hat{y}_i)^2, \quad (2)$$

is minimal. The calculated values are computed by means of an automatic solver for the DAEs based on the backward differentiation formulas (BDF, see for example [3] or [4]). The weights, w_i , are based on the accuracy of the measurements and have dimension $1/[\hat{y}_i]$. In the case that the errors in the measurements are statistically independent and we assume a Gaussian distribution of the errors with standard deviation, σ_i , we take $w_i = 1/\sigma_i$ and weighted least squares coincides with the maximum likelihood estimate.

3. Variational equations

In order to use a gradient-based minimisation procedure we solve, besides the set of DAEs (1), the corresponding set of variational equations with respect to the unknown parameter vector:

$$E \frac{d}{dt} \frac{\partial y}{\partial \theta} = \frac{\partial f}{\partial y} \frac{\partial y}{\partial \theta} + \frac{\partial f}{\partial \theta}, \quad \text{with } \frac{\partial}{\partial \theta} y(t_0, \theta) = \frac{\partial}{\partial \theta} y_0(\theta). \quad (3)$$

The variational equations can easily be derived from (1) by a computer algebra package (we use Maple V). The solution of (3) yields the gradient $\partial y(t, \theta)/\partial \theta$, which will be used later for the minimisation of the weighted sum of squared discrepancies (cf. (2)).

Further, inspection of (3) shows that its Jacobian corresponds with m copies of the Jacobian of the original model equations (1). This means that the variational equations inherit the stiffness character of the original equations. In case of a BDF method, when a certain order and step strategy has been provided for numerical solution of (1), the same strategy can be used to solve (3) numerically.

The use of a computer algebra package to derive the variational equations in combination with the same order and step strategy, leads faster to a more accurate gradient than is possible by finite differences.

4. Minimisation

Denoting the vector of weighted discrepancies by the column vector

$$Y(\theta) = w_i(y_{c_i}(t_i, \theta) - \widehat{y}_i)_{i=1, \dots, N}, \tag{4}$$

we can write the sum of squares (2) as

$$S(\theta) = \|Y(\theta)\|^2 = Y^T(\theta)Y(\theta). \tag{5}$$

For a given value of θ , the vector $Y(\theta)$ can be computed by numerical integration of the model equations (1). The variational equations (3) lead to the calculation of the $N \times m$ matrix

$$A(\theta) = \frac{\partial Y(\theta)}{\partial \theta} = \left(w_i \frac{\partial}{\partial \theta} y_{c_i}(t_i, \theta) \right)_{i=1, \dots, N}. \tag{6}$$

The minimisation of (5) is done by an iterative procedure. Suppose θ is a trial vector and its correction is given by $\delta\theta$. The squared sum of the improved parameter vector can be approximated by a quadratic function of $\delta\theta$:

$$\begin{aligned} S(\theta + \delta\theta) &= Y^T(\theta + \delta\theta)Y(\theta + \delta\theta) \\ &\approx (Y(\theta) + A(\theta)\delta\theta)^T(Y(\theta) + A(\theta)\delta\theta) \\ &= Y^T(\theta)Y(\theta) + 2\delta\theta^T A^T(\theta)Y(\theta) + \delta\theta^T A^T(\theta)A(\theta)\delta\theta. \end{aligned} \tag{7}$$

Its minimum is given by the normal equations

$$A^T(\theta)A(\theta)\delta\theta = -A^T(\theta)Y(\theta). \tag{8}$$

The last formula is the starting point of a Gauss–Newton method. It is obvious that the Gauss–Newton procedure will fail if the matrix $A(\theta)$ is (almost) singular. A well-known remedy is the use of the Levenberg–Marquardt method, which changes (8) to a limited extent into

$$(A^T(\theta)A(\theta) + \lambda I_m)\delta\theta = -A^T(\theta)Y(\theta), \tag{9}$$

where λ is adjusted to the condition of the matrix $A(\theta)$. The Levenberg–Marquardt method can be seen as a combination of the Gauss–Newton and steepest descent.

To solve $\delta\theta$ from (9), we use the singular value decomposition (SVD) of the matrix $A(\theta)$ as

$$A(\theta) = U(\theta)\Sigma(\theta)V^T(\theta), \tag{10}$$

with $U(\theta)$ and $V(\theta)$ unitary $N \times m$ and $m \times m$ matrices, respectively, and $\Sigma(\theta)$ is a diagonal matrix, containing the singular values in a non-decreasing order. Substitution of (10) in (9) leads to the following expression for the correction of the parameter vector:

$$\delta\theta = -V(\theta) \left(\Sigma^2(\theta) + \lambda I_m \right)^{-1} \Sigma(\theta)U^T(\theta)Y(\theta). \tag{11}$$

5. Statistical background

We assume that the errors in the measurements are statistically independent, scaled via the weights in such a way that they have equal variance (σ^2) and come from a Gaussian distribution. In most practical

cases this is quite a reasonable assumption. The (weighted) experimental errors in the measurements are then given by $Y(\theta)$, as in (4). This means that the covariance matrix of the experimental errors is given by

$$\mathbf{E}\left(Y(\theta)Y^T(\theta),\right) = \sigma^2 I_N. \quad (12)$$

In combination with (11) this gives the covariance matrix of $\delta\theta$:

$$\mathbf{E}(\delta\theta\delta\theta^T) = \sigma^2 \left(A^T A\right)^{-1} = \sigma^2 V \Sigma^{-2} V^T. \quad (13)$$

Upon convergence of the Levenberg–Marquardt algorithm we obtain a final estimate of θ , denoted by $\bar{\theta} = \theta + \delta\theta$. We can write the corresponding final sum of squares as

$$S(\bar{\theta}) = Y^T(\theta + \delta\theta)Y(\theta + \delta\theta) \approx Y^T(\theta)Y(\theta) - \delta\theta^T \left(V \Sigma^2 V^T\right) \delta\theta, \quad (14)$$

where the matrices V and Σ are from the SVD (see (10)).

From standard statistics (see [2]) we know that $S(\bar{\theta})/\sigma^2$ and $\delta\theta^T(V \Sigma^2 V^T)\delta\theta/\sigma^2$ have a χ^2 -distribution with $N - m$ and m degrees of freedom, respectively. An unbiased estimate of σ^2 is given by

$$s^2 = S(\bar{\theta})/(N - m). \quad (15)$$

The confidence region at level α is the ellipsoidal region

$$\delta\theta^T \left(V \Sigma^2 V^T\right) \delta\theta \leq \frac{m}{N - m} S(\bar{\theta}) \mathcal{F}_\alpha(m, N - m), \quad (16)$$

where $\mathcal{F}_\alpha(m, N - m)$ is the upper α quantile for Fisher's \mathcal{F} -distribution with m and $N - m$ degrees of freedom.

The independent confidence interval for each estimate is given by

$$((\bar{\theta}_i) - \delta\theta_i^*, (\bar{\theta}_i) + \delta\theta_i^*), \quad (17)$$

with

$$\delta\theta_i^* = \sqrt{\frac{m}{N - m} S(\bar{\theta}) \mathcal{F}_\alpha(m, N - m) (V \Sigma^{-2} V^T)_{ii}}.$$

In case the weights are not known a priori, these quantities can be estimated together with the parameters if some assumptions are made. This case of unknown weights in combination with dependent measurement errors is treated in Appendix A.

It should be emphasised that all the results from this section are obtained by linearisation. For most non-linear problems this can give some insight in the neighbourhood of the minimum, but if we restrict ourselves to this information it can at the same time be very misleading. More information can be obtained from sampling many points in the ellipsoidal region, this may show that the confidence region is not an ellipse but a non-convex and irregular region.

6. Restrictions with respect to the parameters

For many practical reasons it may occur that there are restrictions with respect to the parameters to be estimated (e.g. reaction constants are always non-negative). Most of the simpler linear restrictions can be taken into account by a re-parametrisation, but that is not sufficient for the general case.

Suppose we have K restrictions for the m unknown parameters. The restrictions are, in general, non-linear so that we write

$$R(\theta) \leq 0, \tag{18}$$

where $R(\theta)$ is a K -dimensional vector function. The restrictions imply that a subset of the m -dimensional parameter space is excluded. In case of a constraint optimisation problem, we start the numerical procedure as if we were dealing with the un-constraint case (starting with an initial θ s.t. $R(\theta) \leq 0$) which results in a $\delta\theta$ (see (11)). Then we check whether after the correction the constraints are still fulfilled: $R(\theta + \delta\theta) \leq 0$. When some of the constraints are violated, there will be a non-empty subset $\{i_1, \dots, i_k\} \subset \{1, \dots, K\}$; k is the number of violated or active constraints.

After determining the subset $\{i_j\}$, we compute the $k \times m$ matrix B , again by using computer algebra:

$$(B)_{jl} = \partial R_{i_j} / \partial \theta_l. \tag{19}$$

For notational convenience we introduce a k -dimensional vector $r(\theta)$ which contains all the vector elements $R_{i_j}(\theta)$ for $j \in \{1, \dots, k\}$. If we write down the normal equations with linearised constraints and denote the Lagrange multipliers by q , we get

$$\begin{pmatrix} A^T A & B^T \\ B & \emptyset \end{pmatrix} \begin{pmatrix} \delta\theta \\ q \end{pmatrix} = - \begin{pmatrix} A^T Y(\theta) \\ r(\theta) \end{pmatrix}. \tag{20}$$

Again we would not use the standard Gauss–Newton, but choose the Marquardt–Levenberg method. This leads to the correction

$$\delta\theta = -V \left(\Sigma^2 + \lambda I_m \right)^{-1} \left[\Sigma U^T Y(\theta) + (BV)^T q \right], \tag{21}$$

where the Lagrange multipliers, q , are given by

$$q = \left(BV \left(\Sigma^2 + \lambda I_m \right)^{-1} (BV)^T \right)^{-1} \left(BV \left(\Sigma^2 + \lambda I_m \right)^{-1} \Sigma U^T Y(\theta) - r(\theta) \right).$$

It may take some iterations to fulfil all the restrictions. Numerical experiments showed that for a feasible set of restrictions two or three iterations are sufficient.

7. A case study from biochemistry

As an example we treat the following set of chemical reactions originating from biochemistry:



It is a simple but illustrative example. The treatment of real-life problems with many more state variables and parameters is outside the scope of this paper and will be reported elsewhere.

The state variables in the reaction scheme are the concentrations of the enzyme, $[E]$, substrate, $[S]$, and complex, $[C]$. The mathematical description of the problem is given by

$$\frac{d[S]}{dt} = -k_1[E][S] + k_2[C], \quad \frac{d[C]}{dt} = k_1[E][S] - k_2[C] - k_3[C], \quad [E] + [C] = [E_0] + [C_0]. \tag{22}$$

Table 1
Initial and final parameter values for the case study plus additional statistics

	θ_{ini}	θ_{end}	$\delta\theta^*$
k_1	6.0	0.683	0.0762
k_2	0.8	0.312	0.0677
k_3	1.2	0.212	0.00544
$S(\theta)$	0.848	0.00051	




Fig. 1. The calculated concentrations of the complex (C) for the initial (left graph) and final parameter vector (right graph) and the measurements (X).

The initial values are $[S_0] = 1.0$, $[C_0] = 0.0$ and $[E_0] = 1.0$, the vector of unknown parameters yields $\theta^T = (k_1, k_2, k_3)$. During the experiment the complex concentrations have been measured. These can be found in Appendix B.

For this numerical test we take all the weights equal to 1 and assume the errors in the measurements to be independent. The initial parameter vector, θ_{ini} , the final estimate, θ_{end} , the corresponding sum of squared discrepancies (see Eq. (2)) and the confidence limits ($\delta\theta^*$ from (17)) are listed in Table 1.

The solution of the DAEs describing the reaction mechanism (Eq. (22)) with the initial and the final estimates of the parameters is shown in Fig. 1 for the complex concentrations.

8. Conclusions

In this paper we presented an approach to parameter estimation problems in systems of differential algebraic equations. The described solution method has been implemented and can be applied in many sciences where mathematical modelling of time-dependent processes is involved. The treated case study shows the usefulness of this implementation for experimental sciences.

Appendix A

Here we describe a more general situation, which is indispensable if we do not have an estimate for the measurement errors and expect them to be dependent. We consider dependence between errors of different measurements taken at the same time.

In most practical cases the weights, with $w_i^2 = 1/\sigma_i^2$, are unknown and it is impossible to estimate all these weights in addition to the unknown parameters, θ . In case of dependent errors we also have unknown covariances. The best that can be done is to assume that the variances and covariances of the discrepancies only vary with their corresponding components.

For convenience of notation we introduce the integers q as the number of measured components, $q \leq n$, and r as the number of measurements taken at the same time, otherwise known as sample size.¹ Now we may form the $r \times q$ matrix D containing the (unweighted) discrepancies, $d_i = y_{ci}(t_i, \theta) - \hat{y}_i$, in such a way that each column is associated with one measured component and each row corresponds to one sample. The matrix entry D_{ij} corresponds to the j th measurement of the i th sample. With the matrix D we introduce the $q \times q$ matrix M , given by

$$M = D^T D. \quad (\text{A.1})$$

In case of independent measurement errors the covariance matrix, V , is diagonal and its non-zero entries equal σ_i^2 ($i \in \{1, \dots, q\}$). A short hand notation of the weighted sum of squares (Eq. (2)) is

$$\text{Tr}(V^{-1}M). \quad (\text{A.2})$$

For the general case (a full matrix V), the probability density (see [1]) is given by

$$p(d_1, \dots, d_N, \theta) = (2\pi)^{-N/2} (\det(V^{-1}))^{r/2} \exp(-\frac{1}{2} \text{Tr}(V^{-1}M)). \quad (\text{A.3})$$

The logarithm of the corresponding likelihood function (LLF) yields

$$\ln \mathcal{L}(\theta) = -N \ln(\sqrt{2\pi}) - \frac{1}{2} r \ln(\det(V)) - \frac{1}{2} \text{Tr}(V^{-1}M). \quad (\text{A.4})$$

Differentiation with respect to the covariance matrix leads to

$$\partial(\ln \mathcal{L}(\theta))/\partial V = -\frac{1}{2} r V^{-1} + \frac{1}{2} V^{-1} M V^{-1}, \quad (\text{A.5})$$

and this expression vanishes if $V = M/r$. The last expression gives a biased and consistent estimate of the covariance matrix. An unbiased estimate is given by

$$\hat{V} = \frac{1}{r - m/q} M. \quad (\text{A.6})$$

If we finally substitute this estimate in the LLF of Eq. (A.4) we obtain

$$\begin{aligned} \ln \mathcal{L}(\theta) &= -\frac{1}{2} N \ln(2\pi) - \frac{1}{2} r \ln \left(\left(\frac{1}{r - m/q} \right)^q \det(M) \right) - \frac{1}{2} \text{Tr}((r - m/q)I_q) \\ &= \frac{1}{2} N \ln((N - m)/2q\pi) - \frac{1}{2} r \ln(\det(M)) + \frac{1}{2} (m - N). \end{aligned}$$

¹ Here we assume that $qr = N$, which is in general not true. The adaptations for non-constant sample sizes are straightforward and therefore omitted.

Table 2
Measurements of the complex concentration ($[C]$)

time (t_i)	$[C]_i$	time (t_i)	$[C]_i$
1.0	0.32	11.0	0.18
2.0	0.38	12.0	0.16
3.0	0.38	13.0	0.15
4.0	0.36	14.0	0.13
5.0	0.33	15.0	0.13
6.0	0.30	16.0	0.11
7.0	0.28	17.0	0.10
8.0	0.25	18.0	0.09
9.0	0.23	19.0	0.07
10.0	0.20	20.0	0.06

Maximising this expression is equal to minimising

$$\ln \tilde{\mathcal{L}}(\theta) = \ln(\det(M)). \quad (\text{A.7})$$

Appendix B

The measurements ($N = 20$) corresponding to the case study of Section 7 give values of the complex concentration, $[C]$, at a series of times as given in Table 2.

References

- [1] Y. Bard, *Nonlinear Parameter Estimation* (Academic Press, New York, 1974).
- [2] D.M. Bates and D.G. Watts, *Nonlinear Regression Analyses and Its Applications* (Wiley, New York, 1988).
- [3] C.W. Gear, *Numerical Initial Value Problems in Ordinary Differential Equations* (Prentice-Hall, Englewood Cliffs, NJ, 1971).
- [4] E. Hairer, S.P. Nørsett and G. Wanner, *Solving Ordinary Differential Equations I*, 2nd Edition (Springer, Berlin, 1992).