

Chapter 8

Parameter Estimation from Laser Flash Experiment Data

Louise Wright, Xin-She Yang, Clare Matthews, Lindsay Chapman,
and Simon Roberts

Abstract. Optimisation techniques are commonly used for parameter estimation in a wide variety of applications. The application described here is a laser flash thermal diffusivity experiment on a layered sample where the thermal properties of some of the layers are unknown. The aim is to estimate the unknown properties by minimising, in a least squares sense, the difference between model predictions and measured data. Two optimisation techniques have been applied to the problem. Results suggest that the classical nonlinear least-squares optimiser is more efficient than particle swarm optimisation (PSO) for this type of problem. Results have also highlighted the importance of defining a suitable objective function and choosing appropriate model parameters.

8.1 Introduction

Many components that operate in a high-temperature corrosive environment, such as engine parts and turbine blades, use coatings to increase their operational lifetime. In some cases these coatings are grown on the component by reaction (e.g. oxide layers), and in other cases they are separate substances applied to the surface of the component before it is put into operation. It is often difficult to obtain samples of the coating on its own, since the coating is often too thin and too fragile to be removed from the component in pieces of a usable size.

Accurate prediction of the behaviour and, in particular, the lifetime of such components in operation can avoid unexpected component failure and hence reduce downtime and maintenance costs. Models for prediction of component lifetime

Louise Wright · Xin-She Yang · Clare Matthews · Lindsay Chapman · Simon Roberts
Mathematics and Scientific Computing,
National Physical Laboratory,
Teddington, Middlesex TW11 0LW, UK

generally require a coupled thermal-mechanical analysis to predict stresses caused by differential thermal expansion and oxide layer growth. The thermal part of this analysis requires knowledge of the thermal conductivity of each material present within the component, including coatings.

This chapter describes the application of an optimisation process to a finite volume model of the laser flash experiment using a layered sample. The optimisation minimises the difference between the measured data and model predictions by adjusting model parameters, including the thermal conductivity of one of the layers. The aim is to demonstrate that the thermal conductivity of a layer within a sample can be obtained using optimisation techniques.

The process described shows the steps required for application of optimisation techniques to a real-world problem: data preparation, model development, choice of objective function and parameters, and choice of an appropriate optimisation method. The work reported also illustrates that each of these steps may be revisited repeatedly before a fit-for-purpose model is achieved.

The laser flash experiment will be summarised in section 8.2, and the initial model used to simulate the experiment will be defined in section 8.3. The initial optimisation results obtained will be discussed in section 8.4, and subsequent alterations to the model will be explained in section 8.4.2. The final optimisation results will be discussed in section 8.4.3. Our concluding remarks are given in section 8.5.

8.2 The Laser Flash Experiment

The laser flash experiment measures the thermal diffusivity of materials. Thermal diffusivity is a measure of how quickly heat travels through a material and has units of $\text{m}^2 \text{s}^{-1}$. Thermal diffusivity α is related to density ρ , thermal conductivity λ , and specific heat capacity c_p by the equation

$$\alpha = \frac{\lambda}{\rho c_p}, \quad (8.1)$$

and so if the density and the specific heat capacity of a material can be obtained from other experiments, the thermal conductivity can be calculated from the thermal diffusivity.

The laser flash experiment generates a set of temperature measurements gathered over time. The model used to determine the thermal diffusivity from the measurements is based on a number of assumptions, including the assumption that the material is uniform and isotropic. These assumptions are clearly not true for layered and coated materials such as the components described above. Since it is generally difficult to obtain samples of the coating that are sufficiently large to use in the laser flash experiment, a method of obtaining the thermal properties of each layer within a layered sample would enable the properties of the coating to be determined.

The experiment exposes one circular face of a cylindrical sample of material to a pulse of laser light, and measures the temperature rise of the centre of the opposite circular face. The sample is placed in a furnace so that measurements can be carried

out at well-controlled temperatures. The sample is put in the furnace before the experiment starts and the experiment is not started until the furnace temperature and sample temperature are judged to be equal. The sample is supported by three small pins to minimise conductive losses, and the furnace is held in near-vacuum conditions to minimise convective losses.

The laser flash is a pulse of laser light lasting less than 1 ms. The laser power is adjusted so that the maximum temperature rise in the sample caused by the laser flash is typically between 3 K and 4 K. This temperature rise gives a good signal-to-noise ratio on the detected signal, but is sufficiently small that radiative losses can be approximated well by linearisation and can be taken into account in a straightforward manner when calculating the thermal diffusivity. The laser power used is not known by the user and cannot be obtained from the equipment.

The temperature change of the rear face of the sample is measured throughout the experiment by an infra-red (IR) temperature sensor. The temperature sensor has a finite spot size and so the measurement is an average over an area rather than a value at a single point.

In order to shield the temperature sensor from the laser flash, a guard cap with a window in it is placed over the end of the sample. The guard cap should not be in contact with the material sample since the conductive heat losses from the sample to the cap will affect the measured temperature and hence the calculated thermal diffusivity value.

The measurement data set used in this work is shown in Fig. 8.1. This data set is used as target data in the work reported here, meaning that the aim of the optimisation work was to generate model results that fit these data well. The measurement was carried out at a furnace temperature of 947.15 K, and Fig. 8.1 shows the change in temperature relative to the furnace temperature. This data set was chosen because i) the ambient temperature was sufficiently high that radiative losses would be significant, and so determination of emissivity would be a possibility, and ii) the same data set had been studied previously [1], giving values to which the calculated results could be compared.

The measurement data set consists of temperature change measurements every 1.568 ms. The measurements are continued for 2.373952 s after the laser flash, giving a total of 1515 measurements for time $t \geq 0$ (the first measurement being at $t = 0$). It is assumed that the sample temperature has fully stabilised by the time that the laser is fired, and so the temperature measured at $t = 0$ is taken to be the ambient (furnace) temperature. The time axis is scaled such that the laser was fired at $t = 0$. The small peak shortly after $t = 0$ is caused by energy from the laser flash that has not been absorbed by the sample being measured by the temperature sensor.

The simplest form of data analysis of these data [9] is based on an analytical solution to the transient heat flow equation that assumes a uniform sample, an instantaneous uniform laser flash, and no heat losses from the sample. This approach leads to a 1-D model for the heat flow, and solution of this model gives an expression of the form

$$\Delta T = \Delta T_m \left(1 + 2 \sum_{n=1}^{\infty} (-1)^n \exp(-n^2 \pi^2 \alpha t / L^2) \right) \quad (8.2)$$

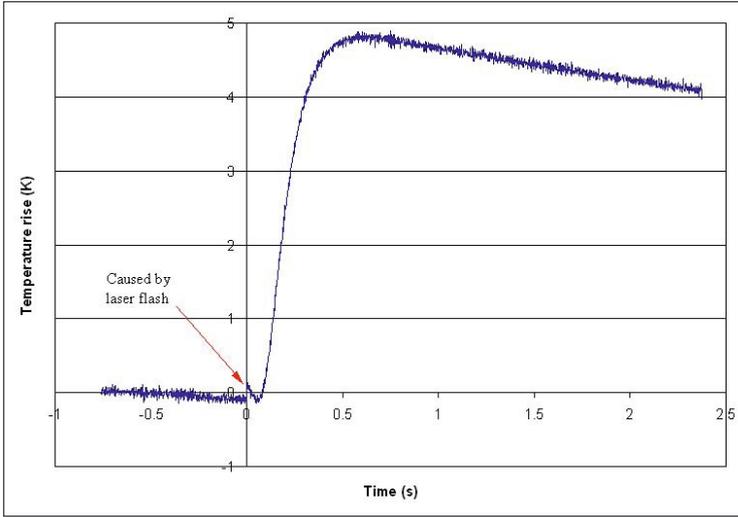


Fig. 8.1 Measured data set used for the work reported here

where $\Delta T(t)$ is the temperature rise of the rear face at time t , ΔT_m is the maximum temperature rise, and L is the thickness of the sample. Defining $t_{1/2}$ as the time taken for the temperature rise to reach half of its maximum value, which can be determined from the measured temperature values, gives

$$\sum_{n=1}^{\infty} (-1)^n \exp(-n^2 \pi^2 \alpha t_{1/2} / L^2) + \frac{1}{4} = 0. \quad (8.3)$$

This is a nonlinear equation that gives α in terms of known values. Solving the equation gives

$$\alpha = 0.138785 \frac{L^2}{t_{1/2}} \quad (8.4)$$

Subsequent work [2, 5, 7, 11] has developed corrections to the simple one-dimensional model to allow for the finite duration of the laser pulse and for radiant heat losses (including those from the curved faces). The methods of data processing that include corrections still make a number of assumptions, including spatial uniformity and isotropy of sample properties, insignificant temperature-dependency of material properties during the experiment, spatial uniformity of the laser flash, and absence of conductive and convective heat losses. The first of these assumptions is clearly not the case for the layered samples of interest in this work. It will be shown in section 8.3 that the final assumption is not valid either.

For the purposes of the modelling work, the sample is assumed to be perfectly cylindrical with plane parallel circular faces of radius 6 mm. The data shown in Fig. 8.1 were gathered during the measurement of a layered sample. It is assumed

that the sample consists of two distinct layers. Each layer is assumed to be uniform and isotropic. The known sizes and properties of the layers in the sample are listed in table 8.1. The emissivity, ϵ , is required for implementation of radiative boundary conditions. A dash in a cell indicates that the property is unknown and is to be determined using optimisation methods to minimise the difference between the measured data and the model predictions. As has been mentioned above, the power of the laser that generates the flash is unknown and also must be determined using optimisation.

Table 8.1 Thicknesses and thermal properties of the layers. A dash in a cell indicates that the property is unknown and is to be determined using optimisation

Material	P92 steel	Oxide
Thickness (mm)	2.0942	0.2265
Density (kg m^{-3})	7871	5015
Specific heat capacity ($\text{J kg}^{-1} \text{K}^{-1}$)	1473.2	934.8
Thermal conductivity ($\text{W m}^{-1} \text{K}^{-1}$)	45.181	-
Emissivity	0.8	-

Previous analysis and simulation relating to this sample have been described in an NPL report [1]. All properties used in the work reported here have been taken from that report or from the references therein. The full chemical composition of the steel is given in the earlier report. The oxide layer consists of two components, magnetite and iron/chromium spinel, but they have been treated as a single uniform substance in order to provide a simpler model for initial investigations. The model could easily be extended to account for more complex layered structures.

8.3 Mathematical Model

8.3.1 Governing Equations

The model considers heat flow within the sample and assumes that the heat flow within the rest of the equipment is either irrelevant or can be taken into account via an appropriate choice of boundary conditions. Cylindrical polar coordinates $\mathbf{r} = \{r, \theta, z\}$ and total temperature (rather than temperature change relative to furnace temperature) will be used throughout.

The temperature distribution within the sample obeys the transient heat flow equation

$$\rho(\mathbf{r})c_p(\mathbf{r})\frac{\partial T(\mathbf{r},t)}{\partial t} = \nabla \cdot (\lambda(\mathbf{r})\nabla T(\mathbf{r},t)) + Q(\mathbf{r},t), \quad (8.5)$$

where \mathbf{r} denotes a position within the sample, $T(\mathbf{r},t)$ is the temperature at a point \mathbf{r} and time t , and $Q(\mathbf{r},t)$ is a heat source term that is used to account for the laser flash.

The domain is defined as $0 \leq r \leq R$, $0 \leq \theta \leq 2\pi$, $0 \leq z \leq L$ where R is the radius of the sample and L is its thickness. It is assumed that the problem is axisymmetric so that variation with θ can be neglected. This reduces the model to two dimensions, making it simpler and quicker to solve.

The domain is split into two layers in the z direction. The properties of each layer are isotropic and uniform and the two layers are assumed to have a perfect thermal bond. The layers are of thickness z_1 (P92) and z_2 (oxide), with $z_1 + z_2 = L$. Then the material properties of the two layers are dependent only on z and are given by

$$\lambda(z) = \begin{cases} \lambda_1 & 0 \leq z \leq z_1 \\ \lambda_2 & z_1 < z \leq L \end{cases} \quad (8.6)$$

$$\rho(z) = \begin{cases} \rho_1 & 0 \leq z \leq z_1 \\ \rho_2 & z_1 < z \leq L \end{cases} \quad (8.7)$$

$$\varepsilon(z) = \begin{cases} \varepsilon_1 & 0 \leq z \leq z_1 \\ \varepsilon_2 & z_1 < z \leq L \end{cases} \quad (8.8)$$

$$c_p(z) = \begin{cases} c_{p1} & 0 \leq z \leq z_1 \\ c_{p2} & z_1 < z \leq L \end{cases} \quad (8.9)$$

The source term Q is assumed to affect a uniform layer (thickness Δz) at the front of the sample directly. It is assumed that the flash is of equal intensity over its finite duration. If the duration is t_0 and the intensity is I then

$$Q(z, t) = \begin{cases} I & 0 \leq z \leq \Delta z, 0 \leq t \leq t_0 \\ 0 & \text{otherwise} \end{cases} \quad (8.10)$$

A value for t_0 is known from the experiment and a value for Δz will be assigned when the numerical solution method is described. The value of t_0 used in this work was 0.8 ms.

The initial conditions assume that the sample is uniformly at the ambient temperature T_0 at time 0, so that

$$T(r, z, 0) = T_0, 0 \leq r \leq R, 0 \leq z \leq L. \quad (8.11)$$

8.3.2 Boundary Conditions

The simplest modelling assumption for boundary conditions is that all cooling is due to radiation only. A straight-line fit to the cooling section of the curve shown in Fig. 8.1 shows a cooling rate of approximately 0.42 K s^{-1} . A simplified model assuming instantaneous uniform temperature change throughout the sample during cooling shows that the maximum radiative heat loss for a sample of this type at an ambient temperature of 947.15 K and a sample temperature of 952.15 K is approximately 0.02 K s^{-1} (this value is likely to be an overestimate due to the model assumptions). The difference of an order of magnitude in cooling rates suggests that the sample must be losing heat via some mechanism in addition to radiation.

Considering the experimental set-up, the most likely source of extra heat loss from the sample is contact between the guard cap and the sample. This contact could be caused by thermal expansion of the sample, the sample being too large for the holder, or poor positioning of the sample within the holder.

The cross section of the guard cap is shown in Fig. 8.2, including dimensions. The cap has a window at its centre through which the sample temperature is measured and which is transparent to infrared radiation. The radiative losses of the sample pass through this window. It is assumed that there is a perfect thermal bond between the surface of the guard cap marked with a heavy line and the equivalent portion of the sample, and that the guard cap is uniformly at the ambient temperature. It is assumed that the window is not in contact with the sample since in reality it is slightly offset from the main part of the guard cap. These assumptions avoid the need to include the heat flow within the guard cap in the model and enable the conductive heat losses to be modelled as a boundary condition.

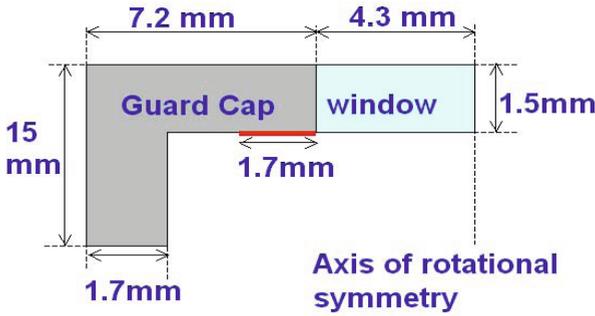


Fig. 8.2 Sketch of the guard cap geometry (not to scale). The dark area marks the region of contact between guard cap and sample

The curved surfaces of the sample are assumed to be perfectly insulated. The temperature gradient along the axis of symmetry must be zero for axisymmetry to be valid. The flat face exposed to the laser is assumed to lose heat radiatively. These boundary conditions can be expressed as

$$\left. \frac{\partial T}{\partial r} \right|_{r=0} = 0, \quad 0 \leq z \leq L, \quad (8.12)$$

$$\left. \frac{\partial T}{\partial r} \right|_{r=R} = 0, \quad 0 \leq z \leq L, \quad (8.13)$$

$$\left. \frac{\partial T}{\partial z} \right|_{z=0} = \varepsilon \sigma (T^4(r, 0, t) - T_0^4), \quad 0 \leq r \leq R \quad (8.14)$$

$$\left. \frac{\partial T}{\partial z} \right|_{z=L} = -\varepsilon \sigma (T^4(r, L, t) - T_0^4), \quad 0 \leq r \leq r_w \quad (8.15)$$

$$T(r, L, t) = T_0 \quad r_w < r \leq R \quad (8.16)$$

where r_w is the radius of the window in the guard cap and σ is the Stefan-Boltzmann constant.

8.3.3 Numerical Methods

These equations, boundary conditions, initial conditions, and material properties fully define the two-dimensional model. The model cannot be solved analytically. A numerical approximation technique must be used instead.

The technique used to solve the model numerically is based on the finite volume method. The structure and approach are described in detail elsewhere [1, 6]. The work reported here has used a version of TherMol [1, 6], an NPL software package for multiphysics applications focussing on the diffusion equation, as the basis for the model. The software used has been adapted from a three-dimensional implementation of TherMol.

The finite volume mesh uses two different volume sizes Δz in the z direction, one for each material. The oxide layer had $\Delta z = 0.0453$ mm and the P92 steel had $\Delta z = 0.0419$ mm. The latter value was used as the value of Δz in the definition of $Q(\mathbf{r}, t)$ in (8.10). A uniform volume size of $\Delta r = 0.1$ mm was used in the r direction.

The finite volume model calculates the change in temperature relative to the initial temperature. This approach means that the rounding errors caused by the use of finite-precision arithmetic have little effect. The only feature of the model that requires the use of the true temperature is the calculation of radiative losses, and the appropriate formulation is used in that section of the software.

An explicit time integration method has been used for the transient calculations for simplicity and ease of implementation. The time step was chosen by trial and error for a typical set of parameter values, and was then divided by 10 to ensure that the model would run for more extreme parameter value choices. No numerical stability problems have been encountered during the work.

The results of interest from the calculations were the temperature changes of the rear face averaged over the spot size of the temperature sensor. It was assumed that the temperature sensor spot size was the same size as the window of the guard cap. The results were output at the same time intervals as the measurements to enable direct comparison.

The software TherMOL has been used before with an optimisation routine to determine unknown properties of the laser flash experiment [1]. The work reported used a one-dimensional model with radiative cooling only and an extension of the Nelder-Mead algorithm, COBYLA [10], able to handle constraints on the parameter values. When applied to the data shown in Fig. 8.1, the optimisation process gave a thermal conductivity of $2.1 \text{ W m}^{-1} \text{ K}^{-1}$ for the oxide layer, and a laser power intensity of $1.6 \times 10^8 \text{ W m}^{-2}$, but the model results did not fit the cooling part of the temperature curve at all well.

8.4 Optimisation Results

The initial optimisation aimed to minimise the relative difference between the measured temperatures and the model predictions by varying the thermal conductivity of the oxide layer, λ_2 , the laser power I , and the emissivity of the oxide layer ε_2 . For any optimisation problem, the formulation of the right objective function is important. Here we intend to minimise errors, but errors can be defined as relative errors and absolute errors, which implies two different ways of defining the objective function.

8.4.1 Initial Optimisation Results

The objective function was initially defined as the root mean square average of the relative differences between the measured data and the calculated values:

$$\sqrt{\frac{1}{N} \sum_{n=1}^N \left(1 - \frac{\bar{T}(t_n; \lambda_2, I, \varepsilon_2)}{T_n} \right)^2}, \quad (8.17)$$

where $N = 1515$ is the number of data points. T_n is the measured temperature rise (i.e. $T - T_0$) at time t_n , and $\bar{T}(t_n; \lambda_2, I, \varepsilon_2)$ is the calculated averaged surface temperature rise over the spot size of the temperature sensor for a given set of values of the model parameters $\lambda_2, I, \varepsilon_2$.

Two optimisation algorithms were used: the Levenburg-Marquardt algorithm within a trust-region [3, 4] and a particle swarm optimisation algorithm (PSO) [8]. The Levenburg-Marquardt algorithm is an efficient local optimiser and will find global minima for smooth unimodal surfaces. The work reported here uses the Matlab Optimisation Toolbox function `lsqnonlin`, an implementation of the Levenburg-Marquardt algorithm designed for minimisation of least-squares functions such as (8.17). The PSO implementation used was developed at NPL. The technique is a global optimiser and can be used to check that the local optimiser's results are globally optimal. The PSO is expected to take more time to converge to the optimal solution than the local optimiser.

The initial optimisation results gave a poor fit to the measurement data, particularly at later times. A typical set of results, generated by parameter values identified as optimal by both algorithms, is shown in Fig. 8.3. The lower of the plots shows a close-up of the first 0.15 seconds. The measured temperature changes here are very close to zero, and the model is a very good fit to these values. The small values of T_i for these initial times mean that the relative errors are very large, and so the fit of the model to the data at small times dominates the overall fit.

This dominance of small measured temperature change values suggests that the root mean square of differences,

$$\sqrt{\frac{1}{N} \sum_{n=1}^N (T_n - \bar{T}(t_n; \lambda_2, I, \varepsilon_2))^2}, \quad (8.18)$$

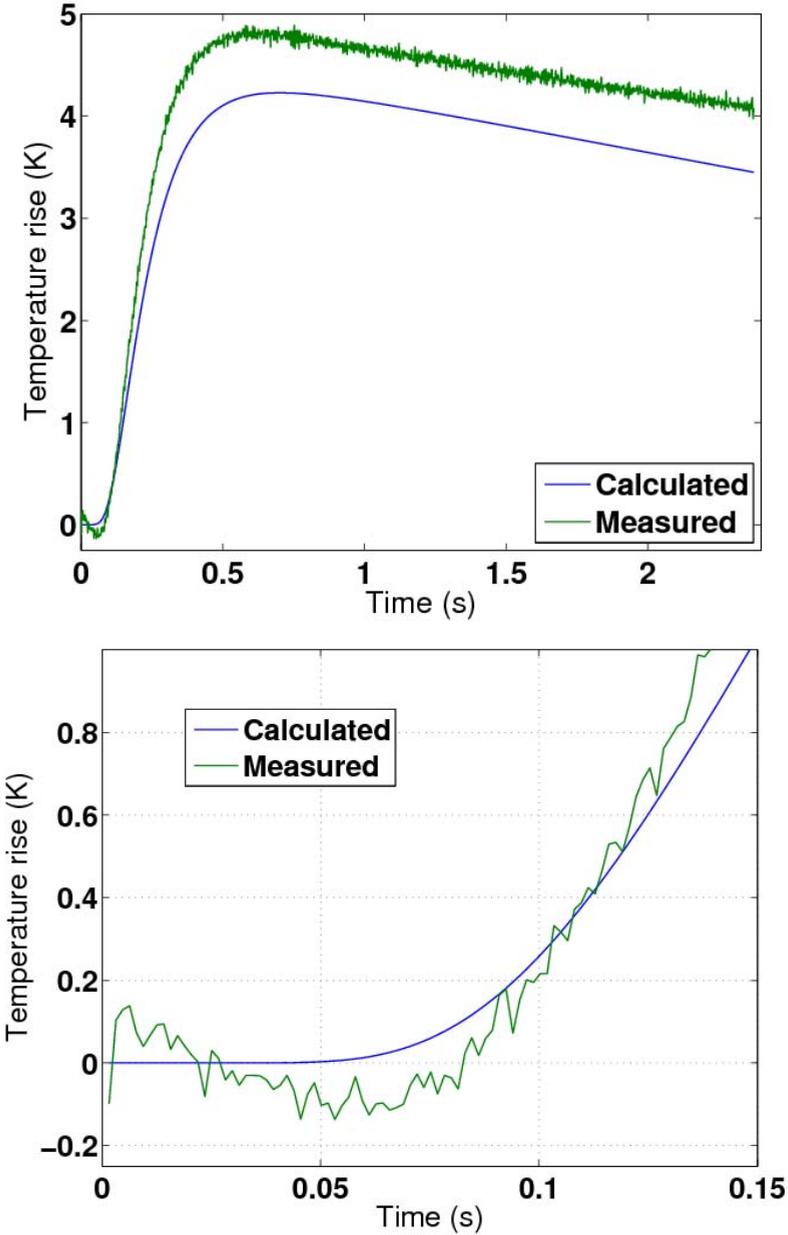


Fig. 8.3 An inappropriate objective function leads to a poor fit of model results to measured data (upper figure), but an unnecessarily good fit to values around 0 (lower figure)

would be a better choice of objective function. This function has been used to generate all results in the rest of this chapter. Normally this sum would be weighted according to the uncertainties associated with the measurements, but here it is assumed that the uncertainties were the same for all measurements.

In general, if the target data contains values close to zero, absolute differences may be a better choice of objective function than relative differences. Whilst minimisation of relative errors can be a good way to combine results of different types, the measured data may not be the best choice of scaling factor when measured values are close to zero. The problems experienced illustrate the importance of considering the construction of the objective function carefully.

In order to check the sensitivity of the converged solution to the initial parameter estimates, five sets of optimisation runs were carried out using each optimisation algorithm. The runs started from randomly-generated points within the parameter search space. The optimal solutions found and the number of function evaluations required to find them are summarized in Table 8.2.

The results in Table 8.2 show that both algorithms converge to the same optimal solution repeatedly. The PSO showed more variation within the five runs than `lsqnonlin`, leading to a higher standard deviation for each of the parameters. The size of the standard deviation is strongly linked to the stopping criteria of the optimisation algorithms. The deviations listed below are consistent with the algorithms effectively arriving at the same solution. As expected, PSO required more function evaluations (about 327) to converge than the efficient local optimiser (about 32). As stated in section 8.3, earlier work had found values of $\lambda_2 = 2.1 \text{ W m}^{-1} \text{ K}^{-1}$ and $I = 1.6 \times 10^8 \text{ W m}^{-2}$, which is a change of about 25% in the value of λ_2 .

Table 8.2 Summary of optimisation results. Means and standard deviations of parameters calculated from 5 runs. Note that intervals specified here are \pm one standard deviation

	PSO	LSQnonlin
Evaluations	327 ± 56	32 ± 3
Mean/Std	$\lambda_2 = 2.83 \pm 0.23 \text{ W m}^{-1} \text{ K}^{-1}$ $I = [1.69 \pm 0.12] \times 10^8 \text{ W m}^{-2}$ $\varepsilon_2 = 0.092 \pm 0.073$	$2.82 \pm 0.005 \text{ W m}^{-1} \text{ K}^{-1}$ $[1.69 \pm 0.04] \times 10^8 \text{ W m}^{-2}$ 0.0000 ± 0.002

The covariance matrix, V_a , associated with these parameter estimates has been calculated from the goodness of fit and Jacobian matrix, using the equation

$$V_a = (J^T J)^{-1} \frac{1}{N - m} \sum_{n=1}^N (T_n - \bar{T}(t_n; \lambda_2, I, \varepsilon_2))^2 \quad (8.19)$$

where J is the Jacobian matrix and $m = 3$ since there are three parameters. The Jacobian matrix has been estimated using finite difference approximations since the objective function of this model is a black box. The standard uncertainties associated with the parameter estimates are given by the square root of the diagonal

entries of the resulting matrix. The standard uncertainties were, in the order λ_2 , I , ε_2 , $0.047 \text{ W m}^{-1} \text{ K}^{-1}$, $7.2 \times 10^5 \text{ W m}^{-2}$, and 0.17. This reflects the high degree of uncertainty about the emissivity. The associated estimate of the goodness of fit was $8.3 \times 10^{-3} \text{ K}$.

The results of the model obtained by using the optimal parameter values are shown in Fig. 8.4. These model results are clearly a better fit to the measured data than those shown in Fig. 8.3, illustrating the benefit of changing the objective function. The results are also a better fit to the measured values than the results obtained during the work described in [1], illustrating that the new model simulates the experiment better than the original version.

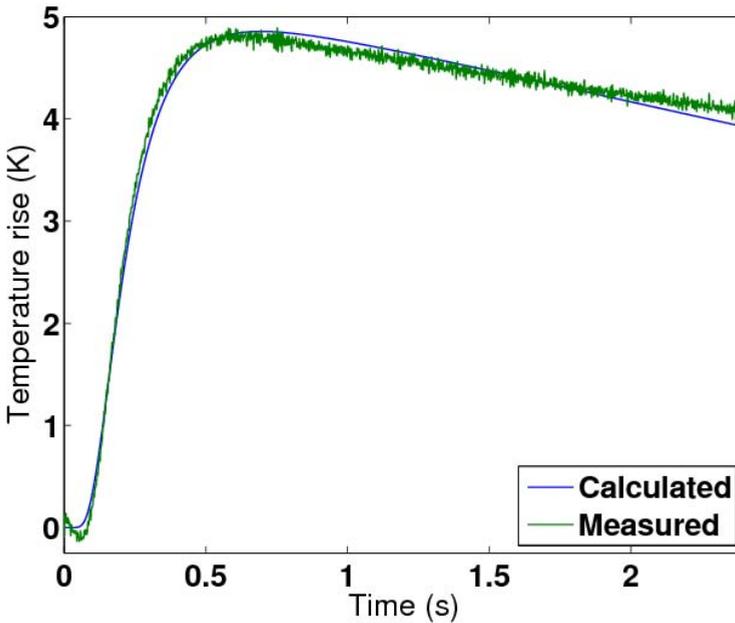


Fig. 8.4 Model results calculated using the best solutions found ($\lambda_2 = 2.8229$, $I = 1.6903 \times 10^8$ and $\varepsilon = 0.0$)

There is still a discrepancy in the cooling part of the curve: the model results appear to cool too fast. In addition, the value of ε that has been found is unexpectedly low (physically the value must be between 0 and 1, and was expected to be close to 0.8). These observations suggest that the conductive losses through the guard cap over-estimate the true cooling, and that the conductive losses dominate the radiative losses to the point where emissivity cannot be determined.

8.4.2 Revision of the Model

The results of the initial model suggest that whilst the inclusion of the effects of the guard cap improves the fit to the measured data, the model is still not predicting the cooling part of the measured data correctly. In order to improve the fitting of the cooling part of the data, the perfect bond between the guard cap and the sample is changed to an imperfect thermal bond parameterised by the unknown thermal bond quality parameter β with units $\text{W m}^{-2} \text{K}^{-1}$.

The imperfect bond is defined as an extra layer between the guard cap and the surface of the sample. The bond quality parameter is effectively the thermal conductivity of the extra layer divided by its thickness. The boundary conditions describing the imperfect bond are generated by solving a one-dimensional steady state heat flux equation analytically at each point on the sample surface that is in contact with the guard cap, and imposing continuity of temperature and flux at the boundaries of the extra layer. This approach assumes that the only heat flow within the extra layer only occurs in the z direction, which is valid because the extra layer is not real and is only a simulation of a poor thermal bond.

The imperfect thermal bond boundary condition is implemented as

$$\lambda \left. \frac{\partial T(r, z, t)}{\partial z} \right|_{z=L} = \frac{\beta \lambda / (\Delta z / 2) (T_0 - T(r, L - \Delta z / 2, t))}{\beta + \lambda / (\Delta z / 2)}, \quad (8.20)$$

where $r_w \leq r \leq R$. From this implementation it is clear that $\beta = 0$ is a perfectly insulating boundary, and that as $\beta \rightarrow \infty$ the condition tends towards a perfect thermal bond with $T(r, L, t) = T_0$ as in the initial model.

8.4.3 Optimisation Results with the New Model

Following the successful application of the Levenbug-Marquardt algorithm to the initial version of the model, the optimisations using the new model with the imperfect bond are carried out using `lsqnonlin` only. The new model uses four parameters and so it is expected that more function evaluations will be required to find a converged solution. The optimisation identified the best parameter values as

- $\lambda_2 = 3.55 \text{ W m}^{-1} \text{K}^{-1}$,
- $I = 1.67 \times 10^8 \text{ W m}^{-2}$,
- $\varepsilon_2 = 1.0$,
- $\beta = 1.92 \times 10^4 \text{ W m}^{-2} \text{K}^{-1}$.

These values were obtained from five runs started from randomly-generated initial parameter estimates. The average number of function evaluations required for convergence was 126. The standard deviations of each of the parameter values across the five runs were less than 10^{-3} , suggesting good repeatability. The covariance matrix was calculated using equation (8.19). This calculation was complicated by the Jacobian estimates suggesting that the derivative of the objective function with respect to ε_2 at $\varepsilon_2 = 1$ was zero. This zero sensitivity meant that the emissivity could

not be included in the covariance calculation, so the covariance matrix only considered λ_2 , I , and β . The standard uncertainties found from the matrix were, in the order λ_2 , I , and β , $0.045 \text{ W m}^{-1} \text{ K}^{-1}$, $1.9 \times 10^5 \text{ W m}^{-2}$, and $620 \text{ W m}^{-2} \text{ K}^{-1}$. The associated estimate of the goodness of fit was $5.2 \times 10^{-3} \text{ K}$, an improvement on the previous model.

The model results obtained by using these parameter values are plotted in Fig. 8.5, along with the measured data and the model results shown in Fig. 8.4 (dashed line). These plots show that the use of an imperfect thermal bond improves the fit of the model results to the measured data, particularly for $t > 0.5$. The old and new models are in close agreement for $t < 0.3$, which is the time period where the energy absorbed by the sample from the laser flash is likely to dominate the heat flow and differences in the value of λ_2 are less likely to have an effect as the oxide layer is comparatively thin.

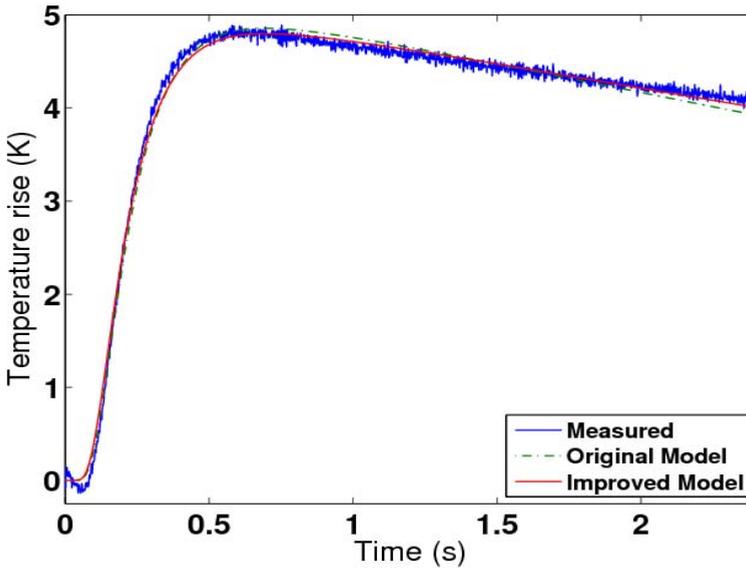


Fig. 8.5 Comparison of new model and original model and their best-fit curves

8.5 Discussion

Using the revised model, the value of λ_2 has increased by about 20 %, and that of ε_2 has gone from 0 to 1, whereas the value of I has not changed significantly. The parameter I defines how much energy goes into the sample during the laser flash, and the good fit of the model predictions to the peak temperature rise suggests that this value has been determined accurately. The value of λ_2 affects how the heat flows within the sample, so it is expected that a change in the boundary conditions would affect the optimal value of λ_2 . Whilst the new value of ε_2 is closer to the

expected value of 0.8, it is by no means certain that this value is a good estimate of the true value. It is likely that the cooling due to the contact with the guard cap still dominates the heat loss, making determination of ε_2 difficult to determine.

It is worth pointing out that the fit, though improved, is still not perfect. Ideally the differences between the model predictions and measured values would lie below the level of the measurement noise, but this level of agreement has not been achieved. There are differences between the model results and the measured values for $0.3 \leq t \leq 0.5$ which suggest that further improvements could be made to the model. Possibilities for improvements include adding circumferential heat losses, considering an imperfect thermal bond between the sample and the oxide, adding extra layers to allow for the multi-phase nature of the oxide, and including a full model of the guard cap so that the conductive losses are modelled more accurately.

It is clear from the results shown here that choosing the best optimisation algorithm can significantly reduce the number of function evaluations, leading to a reduction in computational time. There was no difference between the algorithms in terms of the accuracy of the estimates. The results suggest that the accuracy of the parameter estimates is constrained by the quality of the underpinning model and (were the model to be improved) the uncertainties associated with the measurement data.

Recent trends suggest that metaheuristic algorithms such as PSO are increasingly widely used [8, 12], but popularity does not mean the algorithm is the best choice. In this case study, both PSO and nonlinear least squares provided very good results, but the classical, well-tested nonlinear squares required significantly fewer function evaluations to reach a converged solution. The objective function in this case study was formulated in the least-squares sense, and the results suggest that a unique global optimum exists, which makes the least-squares optimiser more suitable than a metaheuristic algorithm.

Experience gained in this case study and in other applications suggests that it is a good idea to use well-established algorithms when first solving a new optimisation problem. If the well-established algorithms fail, it is worth trying metaheuristic algorithms [12]. This approach can avoid unnecessary and time-consuming trial and error.

References

1. Chapman, L.A., Fry, A.T., Roberts, S.J.: Thermal diffusivity and thermal conductivity measurements on oxide scales. Technical Report MATC(D) 188, National Physical Laboratory (2004)
2. Clark, L.M., Taylor, R.E.: Radiation loss in the flash method for thermal diffusivity. *J. Appl. Phys.* 46, 714–719 (1975)
3. Coleman, T.F., Li, Y.: On the convergence of reflective Newton methods for large-scale nonlinear minimization subject to bounds. *Mathematical Programming* 67, 189–224 (1994)
4. Coleman, T.F., Li, Y.: An interior trust-region approach for nonlinear minimization subject to bounds. *SIAM J. Optim.* 6, 418–485 (1996)

5. Cowan, R.D.: Pulse method for measuring thermal diffusivity at high temperatures. *J. Appl. Phys.* 34, 926–927 (1963)
6. Duncan, B.C., Pilkington, G., Nottay, J.S., Allen, C.R.G., Lawrence, K., Urquhart, J., Roberts, S.J.: Diffusion of moisture in adhesive bonds. Technical Report DEPC-MPR 062, National Physical Laboratory (2007)
7. Dusza, L.: Combined solution of the simultaneous heat loss and finite pulse corrections with the laser flash method. *High Temp. High Press.* 27/28, 467–473 (1995/1996)
8. Kennedy, J., Eberhardt, R.C.: Particle swarm optimization. In: *Proc. IEEE Int. Conf. Neural Networks*, pp. 1942–1948. IEEE Press, Los Alamitos (1995)
9. Parker, W.J., Jenkins, R.J., Butle, C.P., Abbott, G.L.: Flash method for determining thermal diffusivity, heat capacity and thermal conductivity. *J. Appl. Phys.* 32, 1679–1684 (1961)
10. Powell, M.J.D.: A direct search optimization method that models the objective and constraint functions by linear interpolation. In: Gomez, S., Hennart, J.P. (eds.) *Advances in Optimization and Numerical Analysis*. Kluwer Academic, Dordrecht (1994)
11. Taylor, R.E., Clark, L.M.: Finite pulse time effects in flash diffusivity method. *High Temp. High Press.* 6, 65–72 (1974)
12. Talbi, E.G.: *Metaheuristics: From Design to Implementation*. Wiley, Chichester (2009)