10/5-15-95 JSC (1)

SANDIA REPORT

SAND95-8224 • UC-406

Unlimited Release

Printed April 1995

Optimal Control of a CVD Reactor for Prescribed Temperature Behavior

J. C. Meza, T. D. Plantenga

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94551
for the United States Department of Energy
under Contract DE-AC04-94AL85400

Approved for public release; distribution is unlimited.

Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

NOTICE: This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, nor any of the contractors, subcontractors, or their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government, any agency thereof or any of their contractors or subcontractors. The views and opinions expressed herein do not necessarily state or reflect those of the United States Government, any agency thereof or any of their contractors or subcontractors.

This report has been reproduced from the best available copy.

Available to DOE and DOE contractors from:

Office of Scientific and Technical Information
P. O. Box 62
Oak Ridge, TN 37831

Prices available from (615) 576-8401, FTS 626-8401

Available to the public from:

National Technical Information Service
U.S. Department of Commerce
5285 Port Royal Rd.
Springfield, VA 22161

DISCLAIMER

Portions of this document may be illegible in electronic image products. Images are produced from the best available original document.

SAND95-8224

Unlimited Release

Printed April 1995

Optimal Control of a CVD Reactor for Prescribed Temperature Behavior

J. C. Meza and T.D. Plantenga
Scientific Computing Department
Sandia National Laboratories
Livermore, CA 94551-0969

ABSTRACT

The optimal control of a CVD reactor is an important issue in the semiconductor industry. In this report we examine the issues of using large-scale optimization methods for the solution of the CVD optimal control problem. We consider using the TWAFER analysis code to compute the temperature distributions inside a CVD reactor and coupling it with a nonlinear optimization code to find the optimal power curves which achieve a specified target temperature in minimum time. A model problem is numerically tested and analyzed to give insight into our proposed method. The test results indicate that this approach can yield excellent predictions for the optimal power curves.

3/4

MASTER

1. Introduction

An important resource in the semiconductor industry is the use of chemical vapor deposition (CVD) furnaces for the fabrication of wafers. These furnaces are used to heat up a silicon wafer to a specified temperature so that certain chemical reactions may occur, thereby creating a circuit pattern on the wafer. It is extremely important that the specified temperature is achieved and that the temperature is as uniform as possible throughout the entire wafer so that the correct reaction rates occur. There are many different designs for CVD furnaces, each with its own set of advantages and disadvantages. In each case, however, optimal control of the power to the furnace heating elements to meet certain design criteria is an important issue.

Ideally, a furnace should ramp up so that the target temperature of the silicon wafers is attained as fast as possible. A typical scenario starts at 300° C. and increases to 1000° C., which is then maintained while the wafers “cook”. During the process of increasing the heater powers to reach the target temperature, it is important that the temperature variation across each individual wafer not exceed a certain tolerance or the wafer may develop stress fractures. The maximum temperature difference that can be tolerated is usually a known function of temperature. Thus, the problem of ramping up the temperature as fast as possible can be viewed as an optimal trajectory problem with known constraints.

The goal of this research is to predict this optimal temperature trajectory through nonlinear optimization methods. There is also great practical value in developing new smaller batch furnaces that ramp up temperatures faster than a large furnace. This is a problem in optimal design, for which computing an optimal heater power trajectory is an underlying subproblem. It is expected that this effort will require manipulation of many basic elements of the furnace design, such as the placement and spacing of the silicon wafers, and the size and geometry of the furnace and its heating elements. In this report we focus on optimizing the transient performance of a given furnace design by determining

heater power control profiles.

The approach we have taken is to apply nonlinear optimization algorithms to solve for the heater element powers used in a simulation of a CVD reactor furnace. The actual simulations are conducted by the TWAFER analysis software of Houf and Grçar [7], which computes the temperature distribution within a furnace by a finite volume technique. One of the criteria of this project is to regard the TWAFER code as a “black-box”, which implies that the code is not available for modification.

The TWAFER code is a heat transfer analysis code that can compute the temperature distribution at all of the elements of the furnace. TWAFER has the capability to model both conduction and radiation. It does this by satisfying energy conservation laws on a finite volume model of the furnace components. Mathematically, it solves a discretized form of the nonlinear partial differential heat equation for the temperatures at each node. Some of the nodes correspond to heater elements, which introduce power into the system by electrical resistance heating.

In the following section, we first discuss the various optimization formulations that can be taken. In Section 3 we discuss one particular approach we have used and give an example solution from a one dimensional model problem.

2. Optimization Formulation

The straightforward approach to formulating the optimal trajectory problem as an optimization problem is to minimize the time to reach a target temperature. This problem can be phrased as:

Minimize the time to reach the target temperature,

Subject to temperature variations across each wafer being less than the tolerance.

Figure 1 shows the kind of temperature profiles across a wafer that an optimal solution

should generate (the graph was produced from TWAFER output for an extrapolated guess of the optimal heater power curves). The solid lower line is the temperature at the center of one wafer and the dashed upper line is the maximum deviation from this temperature that is allowed. The dotted lines in between are temperature profiles at various points on the wafer, all lying within the tolerance. The figure does not show the heater power curves.

Figure 1: Example of optimal wafer temperature profiles

An algorithm to compute the optimal power curves is conceptually easy to implement. We choose an initial guess for the electrical currents into each heater, let TWAFER solve for temperatures throughout the furnace, and then check the temperature variation across

the wafers. If the temperature variations are within the prescribed tolerance and the time is minimum, then stop; otherwise, choose new heater powers and continue iterating. The optimization variables are the heater powers, which the TWAFER analysis code accesses by reading an input file. This scheme allows TWAFER to operate as a black-box, generating output temperature profiles for a given input choice of heater powers.

The brief description above ignores a number of important issues that may affect algorithm performance and robustness. For instance, an efficient optimization method needs derivatives of the problem, which are not provided by the TWAFER code. Can we compute derivative estimates by finite difference approximations in an accurate and efficient manner? Alternatively, can we use automatic differentiation tools [6] to look inside the TWAFER source code and generate analytic derivatives? We must also anticipate the problem being ill-conditioned with respect to some unknowns because of the time stepping solution process used by TWAFER. As one example, the final temperature and time to reach it is probably relatively insensitive to the heater power values at time zero. These early variables are connected to the final temperatures by a long chain of intermediate values internal to the TWAFER code. Will the optimization algorithm be robust enough to handle the ill-conditioning, or will we have to adjust the derivatives or impose additional constraints to obtain physically meaningful solutions?

Some further insight into the optimization process is afforded by considering the “all-at-once” approach to multi-disciplinary optimization described in [4] (although here we have just one discipline). The idea is to pull the internal TWAFER variables out into full view and look at the total mathematical problem. To cast this in standard optimal control notation, let us designate the “control” or “decision” variables by the vector u , and the “state” variables by the vector x . In our case, the components of u are the heater power profiles, each a function of time t . The state variables are temperature profiles at various points in the furnace (node points determined by the TWAFER finite volume algorithm)

and again are functions of time. Schematically, our optimal control problem looks like

$$\begin{aligned} \min_{u(t)} \quad & t_f \\ \text{subject to} \quad & c(x(t), u(t)) = 0, \\ & h(x(t)) \geq 0. \end{aligned}$$

Here, t_f is the unknown time to reach the target temperature, $h(x(t))$ is a set of inequalities representing the restrictions on temperature variation across each wafer, and $c(x(t), u(t))$ is a set of “state” equations inside the TWAFER code that model the heat flow in the CVD furnace.

Now let’s think about what our optimization procedure does in terms of this problem description. We choose some values for $u(t)$. This is input to TWAFER, which computes $x(t)$ such that $c(x(t), u(t)) = 0$ is satisfied (in practice, this is only happening at certain discretized points in time, but we don’t need to worry about that yet). Then we check the inequalities $h(x(t))$ and test if t_f is minimized. If necessary, we choose a new $u(t)$ and iterate again.

There are two points to make about this method. First, the quantity t_f is not something we get directly for each choice of heater powers. We must look at the temperature profiles $x(t)$ generated by TWAFER and decide when they reach the target temperature, and we might have to do something to make sure the profiles are not drastically overshooting the target (remember, the furnace is supposed to reach the target temperature and stay there while the wafers “cook”).

Second, our optimization method is not solving the problem in an “all-at-once” sense; that is, it does not let $u(t)$ and $x(t)$ vary independently in reaching a total solution. We choose $u(t)$ freely, but $x(t)$ is determined completely by $u(t)$ so that $c(x(t), u(t)) = 0$ is always satisfied. This is a consequence of the black-box nature of TWAFER, which limits us to using the values of $x(t)$ that it computes. There are good reasons to believe this is an inefficient way to solve nonlinear constrained optimization problems (e.g., [5, p. 317]),

but for the moment there is no alternative.

Before concluding this section, let us mention a different optimization formulation for the CVD furnace problem proposed by Norman [9] [10]. He did not try to solve the full optimal control problem as we are. Norman assumed the user provides a near-optimal guess of the temperature profile at one point in the furnace, then he formulates an optimization problem that calculates the heater control curves necessary to produce this given temperature profile. From his solution the temperatures at all other points in the furnace are known, so the wafer temperature variations can be checked. If the variations are too large, then a new guess of the one fixed temperature profile is made and the algorithm repeats. Norman's method never tries to minimize t_f , although an engineer could get lucky and guess a temperature profile that is optimal. The advantage of this simplification is that each optimization subproblem is linear and therefore much easier to solve. We have decided to try and solve the full optimal control problem and determine power curves that truly minimize the furnace ramp time.

This section has outlined a candidate method for finding the optimal control trajectory, but our examination has raised many questions about the method. The full CVD furnace problem is highly nonlinear and will have thousands of variables. As mentioned, it will probably be ill-conditioned, and the optimization objective t_f must be computed indirectly. Before tackling these difficulties, we decided to first investigate a simplified heat transfer problem that does not require the use of TWAFER but still captures the essence of the design optimization problem.

3. Optimizing a Model Problem

Let us consider a homogeneous circular rod of unit length and constant cross-sectional area. One end is perfectly insulated (no heat loss), while the temperature at the other end is the control variable we want to optimize. Heat is transferred through the rod by conduction,

and lost to the surrounding environment by radiation (a nonlinear phenomenon). The object is to choose a controlling temperature that raises the temperature at the insulated end of the rod to a target value T_{target} as fast as possible. To simulate the variation tolerance of the CVD problem, we insist that the temperature difference between the midpoint of the rod and the insulated end never exceed a fixed tolerance ΔT_{max} .

Figure 2: Radiating rod problem geometry

To write down a mathematical description, let the rod lie along the x -axis between zero and one, with the insulated end at $x = 1$ (see Figure 2). Let $T(t, x)$ be the absolute temperature (in degrees Kelvin) at time t and point x . This is the state variable of the system. The state equation is an energy conservation law of the form

$$\frac{\rho c_P}{k} \frac{\partial T}{\partial t} = \frac{\partial^2 T}{\partial x^2} - \frac{2\sigma\epsilon}{kR} (T^4 - T_{amb}^4),$$

where ρ , c_P , k , and ϵ are material constants, R is the radius of the rod, σ is the Stefan-Boltzmann constant, and T_{amb} is the fixed ambient temperature surrounding the rod. For the rest of the paper we will write this equation more compactly as

$$a_1 T_t = T_{xx} - a_2 (T^4 - T_{amb}^4),$$

where the constants have been collapsed into a_1 and a_2 .

Like the real furnace problems we want to solve, this system model is a second-order nonlinear parabolic PDE. It requires an initial condition and two boundary conditions. In addition, the problem statement requires that the system be approximately steady state at the target temperature. We enforce this condition by saying that when the target temperature is reached, the partial derivative of T with respect to time is zero at one

particular point. Let t_f be the final time, and denote the controlling temperature function at the left end by $u(t)$. Then we have

$$\begin{aligned}
 (3.1) \quad & \min_{u(t)} \quad t_f \\
 & \text{subject to } a_1 T_t = T_{xx} - a_2(T^4 - T_{amb}^4) \quad (\text{nonlinear state equation}), \\
 & \quad T(0, x) = T_0 \quad (\text{initial condition}), \\
 & \quad T(t, 0) = u(t) \quad (\text{boundary condition at } x = 0), \\
 & \quad \frac{\partial T(t, 1)}{\partial x} = 0 \quad (\text{boundary condition at } x = 1), \\
 & \quad T(t_f, 1) = T_{target} \quad (\text{stop at the target temperature}), \\
 & \quad \frac{\partial T(t_f, 1)}{\partial t} = 0 \quad (\text{reach steady state at } T_{target}), \\
 & \quad |T(t, 0.5) - T(t, 1)| \leq \Delta T_{max} \quad (\text{limit on temperature variation}).
 \end{aligned}$$

Now we may discretize the space and time variables, apply a collocation principle, and convert the partial derivatives into algebraic quantities. We use an implicit time stepping scheme so that the discretized PDE is numerically stable. This converts the PDE state equation and its initial and boundary conditions to a single large system of nonlinear equations. The discretized optimization problem is therefore an algebraic system with nonlinear equality constraints and linear inequalities.

The numerical solution of the discretized version of (3.1) is particularly difficult because we are trying to minimize the final time t_f . The value of t_f influences our discretization scheme, determining either the spacing between time steps or the total number of steps. A standard method for handling this situation [2, pp. 71–75, 225–233] is to treat t_f as an independent variable and determine its optimal value by an outer loop. We make a guess of t_f and fix it at that value, then try to find an optimal control $u(t)$ that satisfies all the constraints. Information from this subproblem is used to refine the guess of t_f , and the process is repeated.

Let us think about a subproblem with t_f fixed. As the control temperature $u(t)$ is increased at $x = 0$, it creates a temperature distribution across the rod. Although this

distribution cannot be solved for analytically, it seems reasonable to assume the temperature will decrease monotonically from the heated end at $x = 0$ to the insulated end at $x = 1$. There will be a temperature drop between the points $x = 0.5$ and $x = 1$, whose size varies with $u(t)$. Now if t_f is fixed at too small a value, it will be impossible to reach T_{target} without violating the inequality constraints; i.e., the temperature drop will exceed ΔT_{max} . We don't know in advance how large to make t_f ; therefore, we need to relax some constraint to make sure we have a sensible subproblem for every possible value of t_f . One simple idea might be to dispense with the requirement to reach the target T_{target} . Instead, we could find a control trajectory that gets as close to T_{target} as possible in the time available, without violating the inequality constraints. Our optimization subproblem for fixed t_f would then be

$$(3.2) \quad \left. \begin{array}{ll} \min_{u(t)} & (T(t_f, 1) - T_{target})^2 \\ \text{subject to} & a_1 T_t = T_{xx} - a_2(T^4 - T_{amb}^4) \quad (\text{nonlinear state equation}), \\ & T(0, x) = T_0 \quad (\text{initial condition}), \\ & T(t, 0) = u(t) \quad (\text{boundary condition at } x = 0), \\ & \frac{\partial T(t, 1)}{\partial x} = 0 \quad (\text{boundary condition at } x = 1), \\ & \frac{\partial T(t_f, 1)}{\partial t} = 0 \quad (\text{reach steady state at } T_{target}), \\ & |T(t, 0.5) - T(t, 1)| \leq \Delta T_{max} \quad (\text{limit on temperature variation}). \end{array} \right\}$$

The system (3.2) can be discretized and solved, but it becomes underdetermined if t_f is chosen too large. In this case T_{target} is reached easily without violating any of the inequality constraints. The constraints are inactive – effectively not part of the problem. The underdetermined system has infinitely many solutions, some of which don't make any physical sense (for instance, $u(t)$ becomes negative at certain times, or fluctuates wildly). We could try to eliminate the nonphysical solutions by imposing other constraints, but there is a simpler approach.

We note that our objective is just to raise the temperature of the rod as fast as possible

before leveling out to a steady state value at t_f . For a given t_f each subproblem should tell us how high we can raise the temperature without violating the inequality constraints. This information can be used by an outer loop to find the particular t_f which reaches exactly T_{target} . Thus, a better defined optimization subproblem for fixed t_f is

$$(3.3) \quad \left. \begin{array}{ll} \max_{u(t)} & T(t_f, 1) \\ \text{subject to} & a_1 T_t = T_{xx} - a_2(T^4 - T_{amb}^4) \quad (\text{nonlinear state equation}), \\ & T(0, x) = T_0 \quad (\text{initial condition}), \\ & T(t, 0) = u(t) \quad (\text{boundary condition at } x = 0), \\ & \frac{\partial T(t, 1)}{\partial x} = 0 \quad (\text{boundary condition at } x = 1), \\ & \frac{\partial T(t_f, 1)}{\partial t} = 0 \quad (\text{reach steady state at } T_{target}), \\ & |T(t, 0.5) - T(t, 1)| \leq \Delta T_{max} \quad (\text{limit on temperature variation}). \end{array} \right\}$$

The system (3.3) is easy to discretize and has a unique solution. It turns out that $T(t_f, 1)$ varies almost linearly as a function of t_f for reasonable radiation loss values, so the outer loop does not need much information to figure out the optimal value of t_f .

4. Results for Optimal Heating of the Model Problem

We implemented a discretized version of (3.3) and solved it numerically using constrained optimization techniques. The problem was coded in accordance with the CUTE [1] optimization interface standard. Although this interface is somewhat cumbersome to use, it automatically computes analytic derivatives and allows the use of various sophisticated optimization codes.

The state variable $T(t, x)$ was discretized uniformly in space and time to generate the set of collocation points

$$T\left(\frac{i}{N}t_f, \frac{j}{M}\right), \quad \text{for } i = 1, \dots, N \quad \text{and} \quad j = 0, \dots, M.$$

This gave a total of $N(M + 1)$ discrete variables, with N of them acting as control temperatures.

The radiation term in the nonlinear state equation requires that $T(t, x)$ be measured in an absolute temperature scale such as degrees Kelvin. We found that a scaling transformation which keeps variables roughly between zero and one significantly improved the numerical accuracy of the algorithm. Therefore, our implementation was in terms of the scaled quantities

$$y_{i,j} \cong \frac{T(\frac{i}{N}t_f, \frac{j}{M}) - T_0}{T_{target} - T_0}.$$

This equality is approximate in the sense that the variables $y_{i,j}$ were computed from finite difference approximations of the partial differential equations.

We employed a second-order accurate Crank-Nicholson finite differencing scheme to approximate the nonlinear state equation $a_1 T_t = T_{xx} - a_2(T^4 - T_{amb}^4)$. This scheme is not the most accurate and it does not properly match the dissipative nature of heat flow, but it is an implicit scheme that is numerically stable for all discretizations. The primary purpose of this investigation was to demonstrate the feasibility of finding t_f using (3.3), and an implicit scheme makes this possible with reasonably small values of N and M . For points in the interior of the rod (i.e., $i \geq 1$ and $j = 1, \dots, M - 1$), the approximate state equation is

$$(4.1) \quad \begin{aligned} (\mu + a_1)y_{i+1,j} + (\mu - a_1)y_{i,j} = & \frac{\mu}{2}(y_{i+1,j+1} + y_{i+1,j-1} + y_{i,j+1} + y_{i,j-1}) \\ & - \frac{a_2[(T_{target} - T_0)y_{i+1,j} + T_0]^4}{2(T_{target} - T_0)} \\ & - \frac{a_2[(T_{target} - T_0)y_{i,j} + T_0]^4}{2(T_{target} - T_0)} - \frac{a_2 T_{amb}^4}{T_{target} - T_0}, \end{aligned}$$

where $\mu = \Delta t / (\Delta x)^2 = t_f M^2 / N$ is the mesh size parameter. The adiabatic boundary condition can be incorporated into the state equation for points at the right end of the rod

by using

$$(4.2) \quad (\mu + a_1)y_{i+1,M} + (\mu - a_1)y_{i,M} = \mu(y_{i+1,M-1} + y_{i,M-1}) - \frac{a_2[(T_{target} - T_0)y_{i+1,M} + T_0]^4}{T_{target} - T_0} - \frac{a_2[(T_{target} - T_0)y_{i,M} + T_0]^4}{T_{target} - T_0} - \frac{2a_2T_{amb}^4}{T_{target} - T_0}.$$

The initial condition $T(0, x) = T_0$ simply means that

$$(4.3) \quad y_{0,j} = 0.$$

The equations (4.1)-(4.3) account for the nonlinear PDE and its initial and boundary conditions. They constitute a set of NM nonlinear algebraic equations in $N(M+1)$ unknowns. The steady state condition adds another equation

$$(4.4) \quad y_{N,M} = y_{N-1,M},$$

and the scaled temperature variation inequalities take the form

$$(4.5) \quad y_{i,M/2} - y_{i,M} \leq \frac{\Delta T_{max}}{T_{target} - T_0}, \quad \text{for } i = 1, \dots, N,$$

assuming that M is an even number. Thus, up to $N-1$ of the inequalities can be active at a solution. If all N inequalities are active, then the problem is overdetermined and potentially inconsistent, but in practice this situation did not occur.

This completes the conversion of (3.3) into a discretized algebraic optimization problem. For clarity, the full optimization procedure is outlined below in Algorithm 4.1.

Algorithm 4.1 selects trial values of t_f according to a quadratic model based on three optimization subproblem solutions. Note that the quadratic model is not updated with derivative information, just simple iterative refinement. Derivatives could be estimated by finite differences, but the relationship between t_f and $T(t_f, 1)$ is smoothly varying enough that this is not necessary.

Algorithm 4.1. *Numerical procedure for optimizing the heated rod problem*

1. Make three initial guesses of t_f , called t_f^1 , t_f^2 and t_f^3
2. For each guess t_f^k solve the discretized form of (3.3); that is,

$$\max T_0 + y_{N,M}(T_{target} - T_0)$$
 subject to equations (4.1)-(4.5) and with t_f fixed at t_f^k
 Let the maximum temperature found be called T_{max}^k
3. Fit a quadratic curve through the three points (t_f^1, T_{max}^1) , (t_f^2, T_{max}^2) , and (t_f^3, T_{max}^3)
4. Set $k \leftarrow 4$
5. From the quadratic model calculate a value t_f^k that maps to the point (t_f^k, T_{target})
6. Solve the discretized form of (3.3) using t_f^k , and call the solution T_{max}^k
7. If the solution $T_{max}^k = T_{target}$, then stop
8. Make a new quadratic model using (t_f^k, T_{max}^k) and two previously calculated points
9. Set $k \leftarrow k + 1$ and go to step 5.

As a typical example for a conducting metal rod, we set $a_1 = 1$, $a_2 = 10^{-11}$, $T_0 = T_{amb} = 300^\circ \text{ C.}$, and sought to raise the temperature to $T_{target} = 600^\circ \text{ C.}$ with no temperature variations greater than $\Delta T_{max} = 1^\circ \text{ C.}$ We found that discretizing the system (3.3) using constants of $N = 24$ and $M = 12$ was sufficient to compute $T_{max}^k = T(t_f^k, 1)$ to an accuracy better than one degree. Starting from the three guesses $t_f^1 = 25$, $t_f^2 = 50$ and $t_f^3 = 75$ seconds, Algorithm 4.1 gave the results shown in Table 1.

Notice from Table 1 that $T(t_f, 1)$ was very nearly a linear function of t_f . We could have extrapolated the solution $t_f = 47.2$ seconds using just the first two guesses. However, a more realistic CVD furnace problem may not have a linear relation, so the quadratic curve fitting in Algorithm 4.1 is still important for general problems.

The optimal control solution of each discretized subproblem had exactly $N - 1$ of the

k	t_f^k	maximum $T(t_f^k, 1)$
1	25.0 seconds	464.6° C.
2	50.0 seconds	616.1° C.
3	75.0 seconds	752.6° C.
4	47.2 seconds	600.0° C.

Table 1: Solution for radiating rod problem

inequality constraints (4.5) active. Thus, each solution was characterized by having $N(M + 1) = 312$ equalities for the 312 unknowns. In every subproblem it was the final inequality $T(t_f, 0.5) - T(t_f, 1) \leq \Delta T_{max}$ that was inactive. This is reasonable, since the quasi-steady state condition forced the temperature at $T(t_f, 1)$ to stop increasing and level off.

We used two different large-scale constrained optimization algorithms to solve the discretized subproblems. LANCELOT [3] is a commercial software package that uses an augmented Lagrangian penalty function to handle nonlinear constraints. ETR [8] [11] is our own code that uses an SQP method to solve problems with equality constraints. We were able to use ETR because we know which inequality constraints are active and which are not. Both codes use analytic first and second derivatives of the problem to compute a solution.

The subproblems turned out to be quite difficult for the LANCELOT algorithm to solve. We deliberately initialized all the scaled state and control variables to zero to make the optimization subproblems find a solution from scratch. The ETR code took less than a second to solve each subproblem, while LANCELOT took over 200 seconds (this was the best performance of the eleven preconditioning options provided in LANCELOT). All results were obtained with a Silicon Graphics workstation under the IRIX 5.2 operating system using a 150 MHz MIPS R4400 processor. We suspect the poor performance of LANCELOT

stems from the use of the augmented Lagrangian function on an ill-conditioned problem.

The optimization subproblems were harder to solve with a finer mesh or bigger values of the radiation loss constant a_2 . In both cases this was due to ill-conditioning. If we think of the discrete state variables arranged in a rectangular grid of N by M points, then each state equation (4.1) couples the temperature at a point to the temperature at five of its neighbors. Our goal is to choose values of $y_{i,0}$ that produce a maximum value of $y_{N,M}$, but the influence of the control variables is felt at $y_{N,M}$ only after propagating through all intermediate points, one state equation at a time. Thus, $y_{N,M}$ is not very sensitive to changes in variables near $y_{1,0}$. Increasing N or M makes this worse. Increasing a_2 causes more heat to be lost as radiation before it can conduct to the right end, which again decreases the influence that power applied at $y_{1,0}$ has on $y_{N,M}$. The ill-conditioning is an inherent feature of this problem, although performance was improved when variables were started close to a suspected solution instead of at zero.

5. Conclusions

Our investigation of the model problem in §3 and §4 gives us several clues for how to proceed with solving real furnace problems using TWAFER. Instead of directly minimizing the final time t_f , a more tractable problem is to maximize the final temperature for fixed t_f and modify t_f in an outer loop. It is important to impose some condition that makes the temperature profile level off at t_f , especially if each subproblem seeks to maximize the final temperature. In the model problem the optimum solution has most of the temperature variation inequality constraints active, particularly those at early times. This kind of information is very helpful in solving large optimization problems. Finally, we observed that ill-conditioning is a significant problem, indicating the need for a sophisticated optimization algorithm and careful interpretation of the results.

In conclusion, we believe the problem of finding an optimum heater power control

trajectory for a CVD furnace is challenging, but attainable. Our next step is to find a way to get derivatives from the TWAFER code, then couple it with a suitable optimization algorithm.

REFERENCES

- [1] Bongartz, I., A.R. Conn, N.I.M. Gould, and Ph.L. Toint. "CUTE: constrained and unconstrained testing environment." Research Report. IBM T.J. Watson Research Center, Yorktown Heights, NY, 1993.
- [2] Bryson, A.E., Jr., and Y.-C. Ho. *Applied Optimal Control*. Revised Printing. United States: Hemisphere Pub., 1975.
- [3] Conn, A.R., N.I.M. Gould, and Ph.L. Toint. *LANCELOT : A Fortran Package for Large-Scale Nonlinear Optimization (Release A)*. Springer Series in Computational Mathematics 17. Berlin: Springer-Verlag, 1992.
- [4] Cramer, D.J., J.E. Dennis, P.D. Frank, R.M. Lewis, and G.R. Shubin. "Problem formulation for multidisciplinary optimization." *SIAM J. Opt.* 4 (1994): 754-776.
- [5] Fletcher, R. *Practical Methods of Optimization*. 2nd ed. Chichester, UK: Wiley & Sons, 1987.
- [6] Griewank, A., and G.F. Corliss. *Automatic Differentiation of Algorithms: Theory, Implementation, and Applications*. Philadelphia: SIAM, 1991.
- [7] Houf, W.G., and J.F. Grcar. "TWAFER: A model for determining temperatures in batch LPCVD reactors." Technical Report SAND93-8223. Sandia National Laboratories, Livermore, CA, 1993.
- [8] Lalee, M., J. Nocedal, and T. Plantenga. "On the implementation of an algorithm for large-scale equality constrained optimization." *submitted to SIAM J. Optimization*, 1993.
- [9] Norman, S.A. "Optimization of transient temperature uniformity in RTP systems." *IEEE Trans. Electron Dev.* 39 (1992): 205-207.
- [10] Norman, S.A. "Wafer temperature control in rapid thermal processing." Diss. Dept. of Electrical Engineering, Stanford University, 1992.
- [11] Plantenga, T.D. "Large-scale nonlinear constrained optimization using trust regions." Diss. Dept. of Electrical Engineering and Computer Science, Northwestern University, 1994.

UNLIMITED RELEASE

INITIAL DISTRIBUTION:

Dr. John E. Dennis, Jr.
Department of Mathematical Sciences
Rice University
P.O. Box 1892
Houston, TX 77251-1892

Dr. Matthias Heinkenschloss
Interdisciplinary Center for Applied Math
Virginia Polytechnic Institute
Blacksburg, VA 24061-0531

Mary Beth Hribar
Department of Electrical Engineering and
Computer Science
Northwestern University
Evanston, IL 60208-3118

Dr. Jorge More
Argonne National Laboratory
Mathematics and Computer
Science Division
Argonne, IL 60439-4803

Luis Vicente
Department of Mathematical Sciences
Rice University
P.O. Box 1892
Houston, TX 77251-1892

MS1110	R. A. Allen, 1422
MS1110	Robert Leland, 1422
MS1111	John N. Shadid, 1421
MS9001	J. C. Crawford, 8000;
	Attn: E. E. Ives, 5200
	J. B. Wright, 5300
	R. J. Detry, 8200
	W. J. McLean, 8300
	L. A. Hiles, 8400
	P. N. Smith, 8500
	L. A. West, 8600
	R. C. Wayne, 8700
	T. M. Dyer, 8800
	D. L. Crawford, 8900

MS9004	M. E. John, 8100
MS9042	W. G. Houf, 8745
MS9042	R. J. Kee, 8745
MS9042	C. D. Moen, 8745
MS9042	P. Spence, 8745
MS9201	L. D. Brandt, 8112
MS9201	R. J. Gallagher, 8114
MS9214	W. E. Mason, 8117
MS9214	J. C. Meza, 8117 (20)
MS9214	T. D. Plantenga, 8117 (10)
MS9022	Mail Distribution Team, 8533-1 for OSTI (10)
MS9022	Mail Distribution Team, 8533-1/Technical Library, MS0899, 13414
MS0899	Technical Library, 13414 (4)
MS9018	Central Technical Files, 8523-2 (3)