

Mixed Variable Optimization of a Load-Bearing Thermal Insulation System Using a Filter Pattern Search Algorithm

Mark A. Abramson

Air Force Institute of Technology

Department of Mathematics and Statistics

2950 Hobson Way, Building 640

Wright Patterson AFB, Ohio, 45433 USA

`Mark.Abramson@afit.edu`, <http://www.caam.rice.edu/~abramson>

May 5, 2003

Key words: Optimization, nonsmooth optimization, thermal insulation, heat intercepts, categorical variables, mixed variable programming, pattern search algorithm, filter algorithm, nonlinear constraints

Abstract: This paper describes the optimization of a load-bearing thermal insulation system characterized by hot and cold surfaces with a series of heat intercepts and insulators between them. The optimization problem is represented as a mixed variable programming (MVP) problem with nonlinear constraints, in which the objective is to minimize the power required to maintain the heat intercepts at fixed temperatures so that one surface is kept sufficiently cold. MVP problems are more general than mixed integer nonlinear programming (MINLP) problems in that the discrete variables are categorical; *i.e.*, they must always take on values from a predefined enumerable set or list. Thus, traditional approaches that use branch and bound techniques cannot be applied.

In a previous paper, a linearly constrained version of this problem was solved numerically using the Audet-Dennis generalized pattern search (GPS) method for MVP problems. However, this algorithm may not work for problems with general nonlinear constraints. A new algorithm that extends that of Audet and Dennis by incorporating a filter to handle nonlinear constraints makes it possible to solve the more general problem. Additional nonlinear constraints on stress,

Report Documentation Page			Form Approved OMB No. 0704-0188		
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 05 MAY 2003		2. REPORT TYPE		3. DATES COVERED 00-00-2003 to 00-00-2003	
4. TITLE AND SUBTITLE Mixed Variable Optimization of a Load-Bearing Thermal Insulation System Using a Filter Pattern Search Algorithm				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Air Force Institute of Technology, Department of Mathematics and Statistics, 2950 Hobson Way Building 640, Wright Patterson AFB, OH, 45433-7765				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES The original document contains color images.					
14. ABSTRACT see report					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES 24	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

mass, and thermal contraction are added to that of the previous work in an effort to find a more realistic feasible design. Several computational experiments show a substantial improvement in power required to maintain the system, as compared to the previous literature. The addition of the new constraints leads to a very different design without significantly changing the power required. The results demonstrate that the new algorithm can be applied to a very broad class of optimization problems, for which no previous algorithm with provable convergence results could be applied.

1 Introduction

In a thermal insulation system, heat intercepts are often used to minimize the heat flow from a hot to a cold surface. Figure 1 illustrates an example of such a system of fixed length L , in which power is applied to maintain each intercept i at a specified cooling temperature $\bar{T}_i, i = 1, 2, \dots, n$ (we use $\bar{T}$ here to distinguish it from the general temperature variable T used later). An insulator of thickness x_i is placed between each pair of intercepts $i - 1$ and i , with the convention that $i = 0$ and $i = n + 1$ represent the cold and hot surfaces, respectively, so that $\bar{T}_0 = \bar{T}_C$ and $\bar{T}_{n+1} = \bar{T}_H$. Note that each insulator in Figure 1 may have a different cross-sectional area. The design variables for the system include the number and cooling temperatures of the intercepts, and the insulator types and thicknesses. Furthermore, we assume that the system must be load-bearing, meaning that the insulators act as mechanical supports; thus, only solid materials can be used.




Figure 1: Schematic of a Thermal Insulation System

Variations of this type of system used in cryogenic engineering applications, such as superconducting magnetic energy storage systems and space borne magnets, have been studied

by several authors. Hilal and Boom [22] used a gradient-based optimizer to minimize power for $n = 1, 2$, and 3 intercepts with two choices of insulators of constant cross-sectional area, but without mixing insulator types within the system. Hilal and Eyssa [23] studied the same problem, but with variable cross sectional area for the mechanical supports. In considering systems with more general types of insulation, Chato and Khodadidi [14] sought to minimize entropy, similar to the formulation given by Bejan [10]. Other related attempts to optimize the design of these systems are found in [30], [35], and [41]. An actual implementation of these types of systems for the Large Hadron Collider (LHC) project is discussed in three technical reports [18, 24, 31].

While all of these studies vary in geometry and fidelity of the underlying models, none of them optimizes with respect to the number of intercepts and types of insulators. These variables are referred to as *categorical*, meaning that they must take on values only from a predefined list or discrete set of categories, or else the design cannot be evaluated. The term appears more commonly in the field of statistics, in which each observation of a sample is assigned a category based on what interval its value lies in (e.g., see [16]). While categorical variables can take on numerical values, the values may not have any real meaning or inherent ordering (e.g., 1 = steel, 2 = aluminum, etc.). The difficulty these variables present is that traditional methods (most notably, any branch and bound approach) for solving mixed integer problems cannot be applied because this restriction on the categorical variables cannot be relaxed.

Kokkolaras, Audet, and Dennis [26] extended the Hilal and Boom model to include the number of intercepts and type of insulators as actual variables in the model, thus allowing a mixture of different insulator types in the system. This required an extension of the theory [7] to be able to deal with categorical variables. They achieved a 65% reduction in required refrigeration power from that of [22] by using the Audet-Dennis mixed variable generalized pattern search (MGPS) algorithm [7] to solve this bound constrained mixed variable programming (MVP) problem. However, other than choosing the specific list of possible insulators, they did not consider certain load-bearing aspects of the problem, such as thermal expansion, system mass, and stress, because these are modelled as nonlinear constraints.

Since the MGPS algorithm does not handle nonlinear constraints, a new GPS algorithm [2] is applied, in which a filter is added to the MGPS algorithm to handle the nonlinear constraints. This paper describes the formulation of the optimization problem (Section 2), the extensions that have been made to the Audet-Dennis algorithm (Section 3), the implementation of the algorithm for solving the problem (Section 4), and computational results (Section 5).

2 Basic Model of Thermal Insulation Systems

In this section, we reformulate the Kokkolaras *et al.* model (as extended from that of [22]) to include constraints on thermal expansion, stress, and mass. Much of this discussion comes directly from the presentation in [26], although some of the discussion of stress is similar to that of [23]. The new model can be expressed as

$$\begin{aligned} & \min_{(n, I, x, \bar{T}) \in X} f(n, I, x, \bar{T}) \\ & \text{subject to } g(n, I, x, \bar{T}) \leq 0, \end{aligned} \quad (1)$$

with the following nomenclature:

- n is the number of heat intercepts, with the convention that the cold and hot walls are numbered 0 and $n + 1$, respectively.
- $I \in \mathcal{I}^{n+1}$ is the set of insulators used, where I_i represent the insulator type between intercepts $i - 1$ and i , and $\mathcal{I}$ denotes the finite set of possible insulator types.
- $x \in \mathbb{R}^n$ is the vector whose i -th component is the thickness of the i -th insulator, with the convention that $x_{n+1} = L - \sum_{i=1}^n x_i$.
- $\bar{T} \in \mathbb{R}^n$ is the vector whose i -th component is the temperature of the i -th intercept, with the convention that $\bar{T}_0 = \bar{T}_C$ and $\bar{T}_{n+1} = \bar{T}_H$.
- The feasible region X is defined by the following linear and categorical constraints:

$$n \in \{1, 2, \dots, n_{\max}\}, \quad (2)$$

$$I \in \mathcal{I}^{n+1}, \quad (3)$$

$$\sum_{i=1}^n x_i \leq L, \quad (4)$$

$$x_i \geq 0, \quad i = 1, \dots, n \quad (5)$$

$$\bar{T}_{i-1} \leq \bar{T}_i \leq \bar{T}_{i+1}, \quad i = 1, \dots, n. \quad (6)$$

A difficulty in solving this problem is that the dimension of the vectors I , x , and $\bar{T}$ depend on the variable n . For any value of n , there are $n + 1$ other categorical variables and $2n$ continuous variables, yielding a total of $3n + 2$ variables.

2.1 Objective Function

The objective function represents the total refrigeration power of the system; thus,

$$f = \sum_{i=1}^n P_i, \quad (7)$$

where P_i is the power applied to intercept i , $i = 1, 2, \dots, n$. The required power to keep intercept i at a fixed temperature $\bar{T}_i$ is given by

$$P_i = C_i \left(\frac{\bar{T}_H}{\bar{T}_i} - 1 \right) (q_{i+1} - q_i), \quad (8)$$

where C_i (a function of temperature T) is the thermodynamic cycle efficiency coefficient at intercept i , and q_i represents the heat flow from intercept i to $i - 1$.

In general, heat flow q through a volume is governed by Fourier's law,

$$qdx = Ak(T)dT, \quad (9)$$

where A (a function of the spatial coordinates in the z - y plane perpendicular to the x coordinate) is the cross-sectional area of the volume, and k (a function of temperature) is the effective thermal conductivity of the volume.

For the problem at hand, the heat flow q_i from intercept i to $i - 1$ is given by

$$q_i = \frac{A_i}{x_i} \int_{\bar{T}_{i-1}}^{\bar{T}_i} k(T; I_i) dT, \quad i = 1, 2, \dots, n+1, \quad (10)$$

where A_i denotes the cross-sectional area of insulator i , and the thermal conductivity k is a function of both the temperature T and the type of insulator I_i and is assumed to be isotropic.

By substituting (10) into (8), the objective function can be expressed by (7), with

$$P_i = C_i \left(\frac{\bar{T}_H}{\bar{T}_i} - 1 \right) \left[\frac{A_{i+1}}{x_{i+1}} \int_{\bar{T}_i}^{\bar{T}_{i+1}} k(T; I_{i+1}) dT - \frac{A_i}{x_i} \int_{\bar{T}_{i-1}}^{\bar{T}_i} k(T; I_i) dT \right]. \quad (11)$$

2.2 Nonlinear Constraints

The inequality in (1) represents the new nonlinear constraints added to the MVP problem of [26], which we use to numerically test the new filter MGPS (FMGPS) algorithm. Although we add three types of constraints; namely mass, stress, and thermal expansion, we model the stress constraint as an implicit equality constraint so that we can eliminate insulator cross-sectional areas as design variables.

The first constraint concerns the overall mass of the system. This constraint may actually be budgetary in nature, as larger amounts of insulation material increase the overall cost. Since the weight of each insulator can be expressed as a product of the density of the insulator material and its volume, we have the constraint,

$$\sum_{i=1}^n \rho_i(I_i) A_i x_i \leq m_{\max}, \quad (12)$$

where $\rho_i(I_i)$ represents the density of material used as insulator i , and $m_{\max}$ is the maximum allowable mass of the system. Note that setting each volume to $A_i x_i$ assumes that the area of the insulator is constant throughout the interval; however, this is not an unreasonable assumption.

We also assume that the system must be capable of bearing a specified load F . The stress, defined as load per unit area, must not be allowed to exceed a certain level. In this case, we constrain the stress applied to each insulator to be no greater than the tensile yield strength of that insulator (thus, we assume that the load is suspended from the system, rather than resting on top of it). The tensile yield strength at insulator i , denoted by $\sigma_i(T; I_i)$, is a function of both the insulator type and temperature, and its minimum over the temperature range of the insulator $\bar{\sigma}_i$ must act as the constraint limit for the allowable stress. Therefore, we have constraints of the form

$$\frac{F}{A_i} \leq \bar{\sigma}_i \doteq \min \{ \sigma_i(T; I_i) : \bar{T}_{i-1} \leq T \leq \bar{T}_i \}, \quad i = 1, 2, \dots, n+1. \quad (13)$$

The difficulty with the constraints given in (12) and (13) is that they treat the areas A_i , $i = 1, \dots, n$, as additional design variables. However, there is a convenient and perfectly legitimate way around this problem. First, observe that the direct relationship between A_i and P_i in (11) means that decreasing the cross-sectional area of any insulator reduces the power applied to the corresponding intercept, which is exactly the goal of the optimization problem. Thus, at an optimal point, each A_i should be as small as possible. Furthermore, as each A_i is made smaller, the stress constraint in (13) becomes binding – and must be so at optimality. Therefore, we can assume that (13) holds with equality and make the substitution $A_i = \frac{F}{\bar{\sigma}_i}$ to eliminate A_i as a variable in (12) and in the objective function. This yields a stress-mass constraint of the form

$$\sum_{i=1}^n \rho_i(I_i) \frac{x_i}{\bar{\sigma}_i} \leq \frac{m_{\max}}{F}. \quad (14)$$

The final constraint is one associated with the system's thermal expansion, or as in our cryogenic application, thermal contraction. In addition to moving the heat intercepts out of their optimal position, thermal contraction causes additional stress on the materials and, if excessive, can cause other difficulties, such as deformations in the material. The development of this constraint as presented here is adapted from [9].

Since different insulators at different temperatures exhibit different contraction behaviors, we must treat thermal contraction of each insulator separately as a change in its thickness. The constraint can then be expressed as a weighted sum, where each insulator's weight is simply its thickness divided by the total length of the system; *i.e.*,

$$\sum_{i=1}^n \left(\frac{\Delta x_i}{x_i} \right) \left(\frac{x_i}{L} \right) \leq \frac{\delta}{100}, \quad (15)$$

where δ is a limit on the percent total contraction of the system.

For insulator i , $e(T; I_i)$ denotes the unit thermal contraction (a function of temperature) from intercept i to any point between intercept i and $i - 1$ and is computed by the formula,

$$e(T; I_i) = \int_T^{\bar{T}_i} \lambda_i dT, \quad (16)$$

where $T \in [\bar{T}_{i-1}, \bar{T}_i]$, and λ_i is the linear coefficient of thermal expansion for insulator type i . Pre-computed values for $e(T; I_i)$ are available in lookup tables for a wide range of temperatures and for several material types [9, 34]. Total thermal contraction for an insulator i is given by

$$\frac{\Delta x_i}{x_i} = \frac{\int_{\bar{T}_{i-1}}^{\bar{T}_i} e(T; I_i) k(T; I_i) dT}{\int_{\bar{T}_{i-1}}^{\bar{T}_i} k(T; I_i) dT}; \quad (17)$$

thus, the nonlinear thermal expansion constraint is given by

$$\sum_{i=1}^n \left[\frac{\int_{\bar{T}_{i-1}}^{\bar{T}_i} e(T; I_i) k(T; I_i) dT}{\int_{\bar{T}_{i-1}}^{\bar{T}_i} k(T; I_i) dT} \right] \left(\frac{x_i}{L} \right) \leq \frac{\delta}{100}. \quad (18)$$

The resulting optimization problem can now be expressed by the objective function in (7) and (11), modified to eliminate the A_i variables, together with the linear constraints defined in (2)–(6) and the nonlinear constraints defined by (14) and (18); namely,

$$\min_{(n, I, x, \bar{T})} P = F \sum_{i=1}^n C_i \left(\frac{\bar{T}_H}{\bar{T}_i} - 1 \right) \left[\frac{1}{\bar{\sigma}_{i+1} x_{i+1}} \int_{\bar{T}_i}^{\bar{T}_{i+1}} k(T; I_{i+1}) dT - \frac{1}{\bar{\sigma}_i x_i} \int_{\bar{T}_{i-1}}^{\bar{T}_i} k(T; I_i) dT \right]$$

subject to

$$\begin{aligned} & \sum_{i=1}^n \left[\frac{\int_{\bar{T}_{i-1}}^{\bar{T}_i} e(T; I_i) k(T; I_i) dT}{\int_{\bar{T}_{i-1}}^{\bar{T}_i} k(T; I_i) dT} \right] \left(\frac{x_i}{L} \right) \leq \frac{\delta}{100}, \\ & \sum_{i=1}^n \rho_i(I_i) \frac{x_i}{\bar{\sigma}_i} \leq \frac{m_{\max}}{F}, \\ & \sum_{i=1}^n x_i \leq L, \\ & x_i \geq 0, \quad i = 1, \dots, n \\ & \bar{T}_{i-1} \leq \bar{T}_i \leq \bar{T}_{i+1}, \quad i = 1, \dots, n, \\ & n \in \{1, 2, \dots, n_{\max}\}, \\ & I \in \mathcal{I}^{n+1}. \end{aligned}$$

3 Pattern Search Algorithms for Mixed Variable Optimization

Generalized Pattern Search (GPS) is a class of derivative-free optimization algorithms that was first defined and analyzed for unconstrained problems by Torczon [40], and extended to problems with bound [27] and linear [28] constraints by Lewis and Torczon. Audet and Dennis [7] extended the work of Lewis and Torczon to handle MVP problems with bound constraints. This algorithm was demonstrated numerically in [26]. Audet and Dennis provide additional theoretical results under a hierarchy of more general smoothness conditions [8]. A straightforward extension of [7] for linearly constrained MVP problems under the relaxed smoothness assumptions is described in [2].

For problems with nonlinear constraints, Lewis and Torczon [29] describe an augmented Lagrangian GPS approach in which GPS is used to generate approximate solutions to a sequence of bound constrained augmented Lagrangian subproblems (see [15]). As an alternative, Audet and Dennis [6] introduce a filter GPS algorithm. The latter cannot quite guarantee convergence to a first-order stationary point as the former can, but it applies directly, rather than sequentially, to more general problems, and it avoids the use of penalty parameters and Lagrange multipliers.

The new FMGPS algorithm used here is an extension of the mixed variable GPS algorithm of Audet and Dennis, in which a filter is added to handle the nonlinear constraints. It represents a generalization of all the work cited here, and is described in detail in [4] and [2].

3.1 Filters

Fletcher and Leyffer [19] developed the filter algorithm as a way to globalize sequential quadratic programming (SQP) and sequential linear programming (SLP) methods without the need for a merit or penalty function, which would require the user to specify the relative weighting of optimality versus feasibility. Proofs of convergence to a first-order stationary point are given in [20] for SQP and [21] for SLP methods.

In filter algorithms, the goal is to minimize two functions, the objective f and a continuous aggregate constraint violation function h that satisfies $h(x) \geq 0$ with $h(x) = 0$ if and only if x is feasible. The function h is often set to $h(x) = \|C(x)_+\|$, where $\|\cdot\|$ is a vector norm and $C(x)_+$ is the vector of constraint violations at x ; *i.e.*, for $i = 1, 2, \dots, m$, $C_i(x)_+ = C_i(x)$ if $C_i(x) > 0$; otherwise, $C_i(x)_+ = 0$. In particular, if the squared 2-norm is used, then h inherits whatever smoothness properties C possesses [6].

Simply defined, a *filter*, denoted $\mathcal{F}$, is a set of trial points such that no point dominates any other in the set with respect to its objective and constraint violation function values. In other words, given any two points x and y in the filter, either $f(x) < f(y)$ or $h(x) < h(y)$,

but not both.

In constructing a filter for GPS, we put two additional restrictions on $\mathcal{F}$. First, we set a bound on aggregate constraint violation, $h_{\max}$, so that each point $y \in \mathcal{F}$ satisfies $h(y) < h_{\max}$. Second, we include only infeasible points in the filter and track feasible points separately. This is done in order to avoid a problem with what Fletcher and Leyffer [19] refer to as “blocking entries”, in which a feasible filter point with lower function value than a nearby local minimum prevents convergence to both that minimum and a global minimum. Tracking feasible points outside of the filter circumvents this uncommon but plausible scenario. With these modifications, we now refer to a point y as *filtered* if it is dominated by any point in the filter, satisfies $h(y) \geq h_{\max}$, or is feasible and has objective function value greater than the incumbent best feasible point. A point that is not filtered is referred to as *unfiltered*.

3.2 The Filter MGPS Algorithm

A basic pattern search algorithm is characterized by two phases per iteration – a global SEARCH and a local POLL, in which trial points lying on a carefully constructed mesh are evaluated. The goal of the SEARCH is to adequately sample the variable space, so as to quickly identify a promising region containing a good local minimizer. In this step, any finite strategy may be employed (including none) to identify trial points, as long as they lie on the mesh. This gives the user great flexibility in choosing how to select points. For example, any of the following strategies may be used in the SEARCH step to identify trial points to evaluate.

- Randomly select mesh points using a Latin hypercube search [33, 38, 39] or orthogonal arrays [36];
- Use a popular heuristic, such as a few iterations (*i.e.*, generations) of a genetic algorithm;
- Optimize a less expensive surrogate function and map the resulting numerical solution to its closest mesh point (see [12, 13] as examples).

The latter option is popular in practice for optimization problems whose objective and constraint functions are expensive to evaluate. An initial investment of function evaluations is needed to form good surrogate objective and constraint functions, but then optimization of the surrogate problem may be quickly achieved with little expense, and the resulting point mapped to the mesh should yield an improved design. If an improved design is not found, the surrogate functions are updated and then re-optimized during the next iteration. Several studies have shown success using this approach [5, 11, 12, 13].

If the SEARCH step does not produce an *improved mesh point* (*i.e.*, a point with a lower objective function value than the current iterate), the POLL step is executed, in which the

mesh neighbors of the current iterate (called the poll set) are evaluated. This step is performed so that convergence of a subsequence of iterates is guaranteed to produce a limit point satisfying certain necessary conditions for optimality. The careful construction of the mesh enables us to converge to such a point. The POLL step can be terminated early if an improved mesh point is found.

For filter GPS algorithms, the mesh and poll set are defined in terms of a poll center, as opposed to the current iterate. That is, at each iteration k , the poll center $p_k \in \{p_k^F, p_k^I\}$ is chosen as either the incumbent best feasible point p_k^F or the incumbent least infeasible point p_k^I (I and F denote infeasible and feasible, respectively). Instead of seeking a trial point with a better objective function value than the incumbent, the filter GPS algorithm searches for an unfiltered point that can be added to the filter – ideally, an improved incumbent best feasible or least infeasible point.

For mixed variable problems, each iterate $x_k = (x_k^c, x_k^d)$, where x_k^c is the vector of continuous variables and x_k^d is the vector of discrete or categorical variables. The mesh is constructed as the direct product of the discrete variable space X^d with the union of a finite number of lattices in the continuous variable space X^c , but translated from the poll center; *i.e.*,

$$M_k = X^d \times \bigcup_{i=1}^{i_{\max}} \{p_k^c + \Delta_k D^i z \in X^c : z \in \mathcal{Z}_+^{|D^i|}\}, \quad (19)$$

where the mesh size parameter Δ_k is a positive real number that controls the fineness of the mesh, and D^i denotes an $n \times n_{D^i}$ (where $n_{D^i} > n$) matrix whose columns positively span the continuous variable space. Typically, these are chosen as $[I, -I]$ or $[I, -e]$, where I is the identity matrix and e is the vector of ones.

The poll set is the union of the continuous mesh neighbors $P(p_k)$ of p_k with the user-defined set of discrete neighbor points $\mathcal{N}(p_k)$ of p_k . The continuous mesh neighbors can be expressed as

$$P_k(p_k) = \{p_k^c + \Delta_k d \in X^c : d \in D_k^i\}, \quad (20)$$

where $D_k^i \subseteq D^i$ is the set of poll directions at iteration k . For convenience, we use the notation, $h_k^I = h(p_k^I) > 0$, $f_k^I = f(p_k^I)$, and $f_k^F = f(p_k^F)$. The flexibility in selecting either of the two poll centers does not affect the convergence theory [3, 6], but it can cause convergence to a different limit point. Audet and Dennis [6] (see Example 7.1 there) show that choosing one or the other exclusively in can result in unfavorable results in practice. They suggest that one good strategy may be to alternate between the two poll centers each time a mesh isolated filter point is found.

We should note that polling around other points in the filter is certainly allowed; however, other points do not possess the same convergence properties as p_k^F and p_k^I in the limit. Therefore, any such polling should be regarded as part of the SEARCH step, so that the convergence theory is preserved.

In the MGPS method of Audet and Dennis [7], if the POLL step fails to yield an improved mesh point, an EXTENDED POLL step is invoked around any discrete neighbor whose objective function value is sufficiently close to that of the current iterate (*i.e.*, “almost” an improved mesh point). This is done with the goal of generating an improved mesh point.

With the addition of nonlinear constraints to the problem, the filter approach seeks a point that improves either the objective function f or the constraint violation function h . Given the current poll center p_k and user-specified extended poll triggers $\xi_k^f \geq \xi > 0$ and $\xi_k^h \geq \xi > 0$ for f and h , respectively (for some positive constant ξ), we perform an EXTENDED POLL step around any feasible discrete neighbor whose objective function value is within ξ_k^f of that of the best feasible point, or around any infeasible discrete neighbor whose constraint violation function value is within ξ_k^h of that of the least infeasible point (without exceeding $h_{\max}$).

The bi-loss graph in Figure 2 depicts a filter with its possible poll centers for the next iteration. The best feasible and least infeasible points are indicated, and the feasible solutions lie on the vertical axis (labelled f). The dashed lines have been added to indicate the areas for which an EXTENDED POLL step is triggered. If a feasible discrete neighbor has an objective function value higher on the axis than the current incumbent, but lower than the horizontal dashed line, an EXTENDED POLL step is performed around this discrete neighbor. Similarly, an EXTENDED POLL step is performed if a filtered discrete neighbor lies to the right of the current least infeasible solution, but left of the vertical dashed line.




Figure 2: A Filter for GPS with Extended Poll Triggers.

Similar to the MGPS algorithm of [7], the EXTENDED POLL STEP generates a finite sequence of EXTENDED POLL centers during a single iteration, beginning with the discrete neighbor y_k . However, the progress of this step is also different from that of [7], since we seek an unfiltered point, as opposed to a simple decrease in the objective function.

At each step in the EXTENDED POLL, there are three possible outcomes. First, if the EXTENDED POLL yields an unfiltered point, then the point is treated exactly the same as any other unfiltered point; namely, the point is added to the filter, the filter is updated, and

the mesh is coarsened or left unchanged. Second, if the EXTENDED POLL fails to find a point with a better objective or constraint violation function value than the current extended poll center, then the EXTENDED POLL step around the current discrete neighbor is terminated, and extended polling begins around the next qualified discrete neighbor. Once all such points have been tested with the same result, the EXTENDED POLL step is terminated, and the mesh is refined.

However, the third possible outcome requires a new construction; that is, if extended polling around a discrete neighbor fails to yield an unfiltered point, but improves the objective or constraint violation function value, *relative to the discrete neighbor point*, then the EXTENDED POLL step should be allowed to continue, since finding an unfiltered point is still promising. This is done by establishing a temporary *local filter*. At iteration k , for each discrete neighbor y_k , a local filter $\mathcal{F}_k^L(y_k)$ is simply a filter relative to the current EXTENDED POLL step. It is populated initially with only the point y_k and with $h_{\max}^L = \min(h_k^I + \xi_k^h, h_{\max})$. The EXTENDED POLL step then generates a finite sequence of extended poll centers, where each is chosen either as the best feasible or least infeasible point, relative to the local filter. Extended polling with respect to y_k proceeds, with points being added to the local filter as appropriate, until no more unfiltered mesh points can be found with respect to the new local filter, or until an unfiltered point is found with respect to the main filter. Once either of these conditions is satisfied (which is guaranteed to occur for a fixed mesh size), the EXTENDED POLL step ends, and the main filter is appropriately updated with the points of the local filter, which is then discarded. The mesh size parameter Δ_k , which has been kept constant throughout the step, is then updated, depending on the success of the SEARCH, POLL, and EXTENDED POLL steps in finding an unfiltered point with respect to the main filter.

If a trial point (*i.e.*, any point evaluated in the three phases) is dominated by any point in the filter, it is said to be “filtered”. If all the trial points are filtered, then the current poll centers are retained and the mesh is refined by setting $\Delta_{k+1} < \Delta_k$, typically $\Delta_{k+1} = \Delta_k/2$. Otherwise, the trial point is added to the filter, the filter is updated to remove any points that are now dominated by the new point, and the mesh is coarsened by setting $\Delta_{k+1} \geq \Delta_k$. Note that this includes the choice to leave the mesh unchanged. More precise rules that govern how Δ_k must be refined or coarsened to maintain convergence properties of the algorithm are given in [40], [8], and [7].

The EXTENDED POLL step is summarized by the algorithm in Figure 3, and the FMGPS Algorithm is summarized in Figure 4

4 Computational Model

This section further describes the optimization problem in terms of modelling decisions, material data sources, and other problem setup issues.




Figure 3: Extended Poll Step for the FMGPS Algorithm

4.1 Material Data

The types of insulators were chosen as the same as in [26]; namely, nylon, teflon, fiberglass epoxy (both normal and plane), 6063-T5 aluminum, 1020 low-carbon steel, and 304 stainless steel. For each of these materials, a substantial amount of engineering data was required. Thermal conductivity and contraction data were obtained from lookup tables in [9] and [34], while material densities were found in [37] and [34], and tensile yield strength data were obtained from [17] and [32].

Thermodynamic cycle efficiency coefficients C_i , $i = 1, 2, \dots, n$ (see (8)) are dependent on temperature as follows:

$$C_i = \begin{cases} 5, & \text{if } T < 4.2K \\ 4, & \text{if } 4.2K \leq T < 71K \\ 2.5, & \text{if } T \geq 71K. \end{cases} \quad (21)$$

In order to make the model as accurate and efficient as possible, cubic splines were used to fit all of the data found in lookup tables, including thermal conductivity, thermal contraction, and tensile strength data. Numerical integrations were performed by applying a composite Simpson's Rule, with nodes matching those of the cubic spline. This eliminates truncation error, since Simpson's Rule is exact for cubic polynomials [25].




Figure 4: FMGPS Algorithm

4.2 Choosing Discrete Neighbors

The neighborhood structure that the user chooses to incorporate essentially determines the definition of a minimizer. That is, when the solution of an MVP problem is found, it is with respect to the user-specified discrete set of neighbors. If a neighborhood structure is chosen so that all other sets of discrete variable values are neighbors, a more global solution can be obtained, but often at extraordinary computational cost. On the other hand, severely restricting the size of the set of neighbors will save significant computational cost, but a local optimizer with a higher objective function value is likely to be obtained.

In order to make proper comparisons, the set of neighbors we use for this problem is exactly the same as was used by Kokkolaras *et al.* [26]. It includes designs in which the following occur:

- The type of insulator between any two heat intercepts is changed to any other type, while insulator thicknesses and intercept temperatures remain constant.

- An intercept and the insulator above it are removed, while thicknesses of the remaining insulators are increased proportionally (rounded to the nearest integer multiple of the current mesh size) to fill the remaining space.
- A new intercept and an insulator underneath it are added with the following properties:
 - The type of insulator is the same as the one below it,
 - The cooling temperature is set to the average of the two intercepts adjacent to it, rounded to the nearest integer multiple of the current mesh size,
 - The thickness of both the new insulator and the insulator below it are both set to half of that of latter, rounded to the nearest integer multiple of the current mesh size.

Note that rounding to the nearest integer multiple of the current mesh size is necessary to ensure that the trial point lies on the mesh.

5 Computational Results

The FMGPS algorithm described in Section 3.2 has been implemented in a Matlab[®] code, called NOMADm [1], and applied to this problem. We now present results of several NOMADm runs and show that the approach significantly improves the design of Hilal and Eyssa [23], and is comparable to the results of Kokkolaras *et al.* [26], even though the problem takes on the additional nonlinear load-bearing constraints.

Table 1 shows the data parameters chosen for the results that follow. The first four have values identical to the choices of Kokkolaras *et al.* [26], while the remaining parameters are unique to this problem.

Table 1: Thermal Insulation Problem: Model Parameters

Parameter	Symbol	Value
Hot surface temperature	$\bar{T}_H$	300 K
Cold surface temperature	$\bar{T}_C$	4.2 K
System total length	L	100 cm
Maximum number of intercepts	$n_{\max}$	10
Load placed on the system	F	250 kN
Maximum total system mass	$m_{\max}$	10 kg
Maximum system thermal contraction	δ	5%

To match the setup of [26] as much as possible, runs were performed with an initial mesh size of $\Delta_0 = 10$ and terminated when the condition $\Delta_k \leq .15625$ was achieved. The mesh refinement strategy used by [26] could not be duplicated by NOMADm; thus, in our runs, we refine the mesh by simply dividing the mesh size parameter Δ_k in half. Coarsening of the mesh was not performed.

Extended poll triggers for the objective and constraint violation function were set at one and five percent, respectively, the former being consistent with [26]. When the filter logic of the FMGPS is applied for the nonlinear constraints, polling is performed around the best feasible point.

Also consistent with [26], no SEARCH is used, and the initial design consisted of one intercept placed exactly in the middle of the system and set at 150 K, with a nylon insulator on the cold side and a teflon insulator on the hot side.

5.1 Validation

The software and function files were validated by mimicking the designs of [22] and [23], running these problems, and comparing the designs. In both of these previous papers, the authors applied their optimizer to cases of 1-3 heat intercepts with insulators made of either 304 stainless steel or plane-cloth fiberglass epoxy. For stainless steel, our results matched theirs almost exactly. For fiberglass epoxy, there were some slight differences when more than one intercept is used, but these were noted in [26] as well. These differences were most likely caused by different methods in computing the objective function integrals. As in [26], we used cubic splines to fit thermal conductivity data that was available only in tabular form for specific temperatures. However, rather than apply a Matlab[®] integration routine, we applied our own implementation of Simpson's rule, which is exact for cubic polynomials. The inaccuracy is more visible in the epoxy results because thermal conductivity data was only available at four temperatures, as opposed to the 18 different temperatures available for stainless steel.

When we recomputed the results of [26] to validate our mixed variable logic, we converged to a different design with a similar low objective function value. Table 2 shows the differences between the two runs, where the materials cited there are abbreviated by the following: N = nylon, E = epoxy (normal), E_p = epoxy (plane), and T = teflon. Note that, although power is optimized in our software, we report normalized power at termination, in which power is multiplied by the system length L and divided by the smallest cross-sectional area of any insulator. Previous authors have expressed results this way so that designs can be compared, independent of these two parameters. We keep this convention for the same purpose.

While the numerical integration issue just described can lead to small deviations, we suspect that the difference in mesh refinement strategies led to a different local optimizer along a different path. In [26], with a starting mesh size of $\Delta_0 = 10$, the mesh refinement strategy was to divide the current mesh size by 2^ℓ , where $\ell = 1, 2, \dots$ is incremented each

time the mesh is refined. Since NOMADm is currently incapable of incrementing ℓ , our mesh refinement consists of simply dividing the current mesh size by two (*i.e.*, $\ell = 1$).

Table 2: Thermal Insulation Problem: MVP Validation

Problem:	Original		FMGPS ReRun	
Power ($\frac{PL}{A}$):	25.294 W/cm		25.589 W/cm	
Insulators Used:	NNNNNNNEEET		NNNNNNNTEETT	
i	x_i (cm)	$\bar{T}_i$ (K)	x_i (cm)	$\bar{T}_i$ (K)
1	0.3125	4.2188	4.5313	6.125
2	5.4688	7.3438	6.7188	10.55
3	3.9062	10	4.8437	14.35
4	6.5625	15	4.2188	17.994
5	5.7812	20	7.3438	24.969
6	5.1562	25	9.8438	36.006
7	13.2812	40	24.948	71.094
8	21.4062	71.0938	12.135	116.88
9	8.5938	101.25	7.5	156.88
10	9.2188	146.25	6.4063	198.44
11	20.3125		11.5105	

5.2 Adding the Nonlinear Constraints

The nonlinear constraints were added to the runs in three steps. First, we simply tested the two designs from Table 2 versus the new nonlinear constraints and found that the thermal contraction constraint for both designs was violated by approximately 8% (*i.e.*, the observed thermal contraction of 5.4% exceeded the 5% threshold by 8%). This suggests that a new design having a different material configuration should be expected as the new constraints are incorporated. Second, the implicit constraint on stress (given by (13) with equality), was added to allow variable cross-sectional areas, and thus match the formulation of Hilal and Eyssa [23]. We refer to this as the *partial model*. Finally, the mass and thermal contraction constraints were added to complete the full model. By doing so, the resulting change in required power represents the cost of satisfying the additional load-bearing constraints.

Table 3 shows the results for the partial and full models (columns 3 and 4), along with the design found by Hilal and Eyssa [23] (column 2). For each run, the minimal required

power is normalized by multiplying by the total system length and dividing by the smallest cross-sectional area of an insulator. Following the power and the material configuration for each design, insulator thicknesses and heat intercept temperature settings are listed.

Table 3: Thermal Insulation Problem: Results

Problem:	Hilal & Eyssa		Partial Model		Full Model	
Power ($\frac{PL}{A}$):	53.2 W/cm		24.551 W/cm		23.768 W/cm	
Insulators Used:	$E_p E_p E_p$		EEEEEEEEEEEEEE		EEEEEEEEEEEEEE	
i	x_i (cm)	$\bar{T}_i$ (K)	x_i (cm)	$\bar{T}_i$ (K)	x_i (cm)	$\bar{T}_i$ (K)
1	22.0	8.38	7.1875	6.5875	0.625	4.25
2	23.8	36.3	11.406	12.938	8.125	7.7375
3	24.8	116.6	15.625	25.85	7.9688	12.369
4	29.4		29.531	71.094	7.8125	18.094
5			6.875	100.31	12.344	29.912
6			5	127.66	26.094	71.094
7			2.5	143.13	8.125	105.94
8			2.5	159.06	5.3125	135.47
9			4.6875	188.59	5	165.94
10			5	222.5	5.625	202.03
11			9.688		12.9682	

We can see immediately that the addition of the implicit stress constraint results in a variable cross-sectional area design that requires over 50% less (normalized) power than the design found by Hilal and Eyssa [23]. This savings was expected because the newer formulation allows for varying the number of heat intercepts and the mixing of insulator types. A similar (65%) savings is achieved by Kokkolaras *et al.* [26] in optimizing the constant cross-sectional area formulation of Hilal and Boom [22], but with the additional categorical design variables.

A bit surprising is the slight decrease in power when the final two constraints are added (see columns 3 and 4), particularly because this is the point at which a linearly constrained problem becomes a nonlinearly constrained one, and the new algorithm's filter logic is applied. Recall that the theory ensures convergence to a first-order stationary point for the partial model (since all its constraints are linear), but does not do so for the full model. In spite of this, the new algorithm still finds a better feasible design. It is indeed possible that the FMGPS algorithm generates a different sequence and simply terminates near a better local minimizer.

Furthermore, when we restart the process with this solution as the initial point, no movement away from the solution occurs. We also tried starting at a few other “common sense” initial designs, including both (now infeasible) solutions from Table 2, and from a design with 10 evenly spaced intercepts with all normal epoxy insulators; however, all of these produced designs that required more power than the full model design given in Table 3. Thus we believe our solution to be fairly robust.

Figure 5 illustrates the performance of the FMGPS algorithm on the full model, where the power required for the incumbent best design is plotted versus the number of function evaluations. The lower plot is a magnification of the upper one. The “L”-shaped plot is very typical behavior of derivative-free methods, since good stopping rules for these methods are difficult. The “stair steps” seen in the right-hand plot indicate varying length polling sequences.

We should note that the power values shown on the vertical axes of these plots do not match the data in the Table 3 because they represent two different things. The objective function is to minimize power, as measured in Figure 5, but the required power shown in Table 3 is normalized (hence the $(\frac{PL}{A})$ notation), so as to allow comparisons with the results of Hilal and Eyssa [23].

Figure 6 depicts the progression of the filter during the run of the full model, where the plots in the right column are magnifications of those on the left. Each of the three rows represents a “snapshot” taken after 150, 200, and 500 respective function evaluations were performed. Although the algorithm terminated after more than 9000 function evaluations, changes in the filter after 500 function evaluations could not be detected within the resolution of the plot. This is consistent with the long and shallow progression of the best objective function value seen in Figure 5. Clearly, better stopping rules would be useful.

In the filter plots, the asterisks represent a subset of the best feasible points found up to that point, while the “stair step” lines represent the boundary between the filtered and unfiltered points. In this run, the nonlinear constraints were scaled by dividing each by its right-hand side and then subtracting one from both sides. Thus in the left column plots, the choice of $h_{\max} = 1$ represents a 100% constraint violation.

Table 3 also shows that the new constraints yield significantly different insulator configurations than that of Kokkolaras *et al.*[26] – that of all normal cloth fiberglass epoxy. This is consistent with the raw data we used, which shows epoxy to have low thermal conductivity and higher resistance to stress than nylon or teflon.

However, some of the other materials (including nylon and teflon) have better thermal contraction properties than epoxy, which also has a low thermal stress threshold and which would be tested by any thermal contraction. Modelling thermal stress as an additional constraint that depends on thermal contraction would be an interesting extension to this problem and might result in a completely different material configuration.

The results shown here demonstrate that, although the FMGPS algorithm can be expensive when applied to mixed variable problems, it successfully generated much-improved




Figure 5: Iteration History for the Thermal Insulation System Design Problem

designs for this problem. The design of [23] has been significantly improved, and the addition of constraints on stress, mass, and thermal contraction yields a more realistic feasible design with essentially no additional power required over that of [26].

Acknowledgments: This work was done while the author was a doctoral student at Rice University, under the direction of John Dennis and Charles Audet. The author wishes to thank them for their guidance, support, and excellent suggestions that have improved the presentation, as well as Andrew Meade and Michael Kokkolaras for their numerous suggestions in improving the insulation system model. The author also wishes to thank Dr. Robert Mohling of Technology Applications, Inc. for his assistance in obtaining some hard-to-find materials data used in the numerical computations.

The views expressed in this document are those of the author and do not reflect the official policy or position of the United States Air Force, Department of Defense, or United States Government.




Figure 6: Filter Progression for the Full Model

References

- [1] Mark A. Abramson. Nonlinear optimization with mixed variables and derivatives—Matlab® (NOMADm). Software. Available from Mark Abramson’s NOMADm Page, <http://www.caam.rice.edu/~abramson/NOMADm.html>, 2002.
- [2] Mark A. Abramson. *Pattern Search Algorithms for Mixed Variable General Constrained Optimization Problems*. PhD thesis, Rice University, Department of Computational and Applied Mathematics (CAAM), Houston, TX, August 2002. Also available as CAAM Technical Report TR02-11.
- [3] Mark A. Abramson, Charles Audet, and J. E. Dennis, Jr. Generalized pattern searches with derivative information. Technical Report TR02-10, Department of Computational and Applied Mathematics, Rice University, Houston Texas, 2002.
- [4] Mark A. Abramson, Charles Audet, and J. E. Dennis, Jr. Pattern search filter algorithms for mixed variable programming. Technical report, Department of Computational and Applied Mathematics, Rice University, Houston Texas, 2002. In preparation.

- [5] Charles Audet, Andrew J. Booker, J. E. Dennis, Jr., Paul D. Frank, and Douglas W. Moore. A surrogate-model-based method for constrained optimization. In *AIAA/USAF/NASA/ISSMO Symposium on Multidisciplinary Analysis and Optimization*, volume 2000-4891. AIAA, 2000.
- [6] Charles Audet and J. E. Dennis, Jr. A pattern search filter method for nonlinear programming without derivatives. Technical Report TR00-09, Department of Computational and Applied Mathematics, Rice University, Houston Texas, 2000.
- [7] Charles Audet and J. E. Dennis, Jr. Pattern search algorithms for mixed variable programming. *SIAM Journal on Optimization*, 11(3):573–594, 2001.
- [8] Charles Audet and J. E. Dennis, Jr. Analysis of generalized pattern searches. *SIAM Journal on Optimization*, 13(3):889–903, 2003.
- [9] Randall F. Barron. *Cryogenic Systems*. McGraw-Hill, New York, 1966.
- [10] Adrian Bejan. A general variational principle for thermal insulation system design. *International Journal of Heat and Mass Transfer*, 22(2):219–228, February 1979.
- [11] A. Booker, J. E. Dennis, Jr., P. Frank, D. Serafini, and V. Torczon. Optimization using surrogate objectives on a helicopter test example. In John Burns and Eugene Cliff, editors, *Optimal Design*, Philadelphia, 1998. SIAM.
- [12] A. J. Booker, J. E. Dennis, Jr., P. D. Frank, D. W. Moore, and D. B. Serafini. Managing surrogate objectives to optimize a helicopter rotor design — further experiments. Technical Report 98-4717, AIAA, St. Louis, September 1999.
- [13] A. J. Booker, J. E. Dennis, Jr., P. D. Frank, D. B. Serafini, V. Torczon, and M. W. Trosset. A rigorous framework for optimization of expensive function by surrogates. *Structural Optimization*, 17(1):1–13, 1999.
- [14] J. C. Chato and J. M. Khodadidi. Optimization of cooled shields in insulations. *ASME Transactions, Journal of Heat Transfer*, 106(4):871–875, November 1984.
- [15] Andrew R. Conn, Nicholas I. M. Gould, and Phillipe L. Toint. A globally convergent augmented Lagrangian algorithm for optimization with general constraints and simple bounds. *SIAM Journal on Numerical Analysis*, 28(2):545–572, April 1991.
- [16] Morris H. DeGroot. *Probability and Statistics*. Addison-Wesley Publishing Company, Reading, Massachusetts, 2nd edition, 1986.
- [17] T. F. Durham, R. Michael McClintock, and R. P. Reed. *Cryogenic Materials Data Handbook*. Office of Technical Services, Washington, DC, 1962.
- [18] Gerard Ferlin, Berthold Jenninger, Philippe Lebrun, Guillermo Peon, Germana Rid-done, and Balazs Szeless. Comparison of floating and thermalized multilayer insulation systems at low boundary temperature. Technical Report 21, Large Hadron Collider Project, 1996.
- [19] Roger Fletcher and Sven Leyffer. Nonlinear programming without a penalty function. *Mathematical Programming*, 91(2):239–269, January 2002.

- [20] Roger Fletcher, Sven Leyffer, and Philippe L. Toint. On the global convergence of an SLP-filter algorithm. Technical Report NA/183, Dundee University, Department of Mathematics, 1998.
- [21] Roger Fletcher, Sven Leyffer, and Philippe L. Toint. On the global convergence of an SQP-filter algorithm. *SIAM Journal on Optimization*, 13(1):44–59, 2002.
- [22] M. A. Hilal and R. W. Boom. Optimization of mechanical supports for large superconductive magnets. In K. D. Timmerhaus, R. P. Reed, and A. F. Clark, editors, *Advances in Cryogenic Engineering*, volume 22, pages 224–232. Plenum Press, New York, 1977.
- [23] M. A. Hilal and Y. M. Eyssa. Minimization of refrigeration power for large cryogenic systems. In K. D. Timmerhaus and H. A. Snyder, editors, *Advances in Cryogenic Engineering*, volume 25, pages 350–357. Plenum Press, New York, 1980.
- [24] Berthold Jenninger, Guillermo Peon, and Balazs Szeless. Characterization of net type thermal insulators at 1.8 K low boundary temperature. Technical report, Large Hadron Collider Project, 1996.
- [25] David R. Kincaid and E. Ward Cheney. *Numerical Analysis*. Brooks/Cole Publishing Company, Pacific Grove, California, 3rd edition, 2002.
- [26] Michael Kokkolaras, Charles Audet, and J. E. Dennis, Jr. Mixed variable optimization of the number and composition of heat intercepts in a thermal insulation system. *Engineering and Optimization*, 2(1):5–29, 2001.
- [27] Robert M. Lewis and Virginia Torczon. Pattern search algorithms for bound constrained minimization. *SIAM Journal on Optimization*, 9(4):1082–1099, 1999.
- [28] Robert M. Lewis and Virginia Torczon. Pattern search algorithms for linearly constrained minimization. *SIAM Journal on Optimization*, 10(3):917–941, 2000.
- [29] Robert M. Lewis and Virginia Torczon. A globally convergent augmented Lagrangian pattern search algorithm for optimization with general constraints and simple bounds. *SIAM Journal on Optimization*, 12(4):1075–1089, 2002.
- [30] Q. Li, X. Li, G. E. McIntosh, and R. W. Boom. Minimization of total refrigeration power of liquid neon and nitrogen cooled intercepts for SMES magnets. In R. W. Fast, editor, *Advances in Cryogenic Engineering*, volume 35, Part A, pages 833–840. Plenum Press, New York, 1989.
- [31] M. Mathieu, V. Parma, T. Renaglia, P. Rohmig, and L. R. Williams. 293 K – 1.9 K supporting systems for the large hadron collider (LHC) cryo-magnets. Technical Report 163, Large Hadron Collider Project, 1998.
- [32] R. Michael McClintock and Hugh P. Gibbons. *Mechanical Properties of Structural Materials at Low Temperatures*. Washington, DC, June 1960.
- [33] M. D. McKay, W. J. Conover, and R. J. Beckman. A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics*, 21(2):239–245, 1979.

- [34] Metals and Ceramics Information Center, Batelle, Columbus Laboratories. *Handbook on Materials for Superconducting Machinery: Mechanical, Thermal, Electrical, and Magnetic Properties of Structural Materials*, 1974.
- [35] Z. Musicki, M. A. Hilal, and G. E. McIntosh. Optimization of cryogenic and heat removal system of space borne magnets. In R. W. Fast, editor, *Advances in Cryogenic Engineering*, volume 35, Part B, pages 975–982. Plenum Press, New York, 1989.
- [36] A. B. Owen. Orthogonal arrays for computer experiments, integration, and visualization. *Statistica Sinica*, 2:439–452, 1992.
- [37] Joseph Edward Shigley and Larry D. Mitchell. *Mechanical Engineering Design*. McGraw-Hill, New York, 4th edition, 1983.
- [38] M. Stein. Large sample properties of simulations using latin hypercube sampling. *Technometrics*, 29(2):143–151, 1987.
- [39] Boxin Tang. Orthogonal array-based latin hypercubes. *Journal of the American Statistical Association*, 88(424):1392–1397, 1993.
- [40] Virginia Torczon. On the convergence of pattern search algorithms. *SIAM Journal on Optimization*, 7(1):1–25, 1997.
- [41] Mashito Yamaguchi, Takao Ohmori, and Akira Yamamoto. Design optimization of a vapor-cooled radiation shield for LHe cryostat in space use. In R. W. Fast, editor, *Advances in Cryogenic Engineering*, volume 37, Part B, pages 1367–1375. Plenum Press, New York, 1991.