Proceedings of the 2003 IEEE Congress on Evolutionary Computation, presented in the special session on Evolutionary Design Optimization.

Design Optimization for a Novel Class of High Power Microwave Sources

Laurence D. Merkle
Rose-Hulman Institute of Technology
Department of Computer Science and Software Engineering
5500 Wabash Ave, CM 103
Terre Haute, IN 47803
laurence.d.merkle@rose-hulman.edu

John W. Luginsland
SAIC
3465 Box Hill Corporate Center Dr.
Abingdon, MD 21009
john.w.luginsland@saic.com

Abstract- Significant benefits would follow from improving the signal growth rates of certain high-power microwave (HPM) sources, including the relativistic klystron oscillator (RKO). Optimization of the growth rate via analytical and standard numerical techniques is intractable because of the high dimensionality of the design space and the existence of many local optima. Instead, the growth rate is optimized using a real-valued evolutionary algorithm (EA), which performs mutation, selection, and recombination on a population of candidate design parameters.

Practical application of EAs requires the availability of a computationally efficient model of design quality. Two models of the RKO are developed relating the growth rate of the microwave output power to the design parameters. Both models have computationally efficient implementations, and one of them generalizes easily to a novel multi-cavity class of RKO devices, which has significantly better growth rates than standard two-cavity RKOs.

Many design optimization problems of interest involve physical constraints. The GENOCOP evolutionary algorithm includes features which support the incorporation of physical constraints in the problem specification through the maintenance of separate search and reference populations, where the latter consists entirely of feasible individuals. It provides "blind" operators to recombine individuals from the two populations to produce new reference population individuals. However, the use of these blind operators can result in unnecessary modification of the search individual, and domain specific recombination operators can result in improved effectiveness. As with any optimization technique, GENOCOP also allows the use of either the penalty function or repair method for evaluation of infeasible individuals. Computational experiments are performed comparing the effectiveness of each possible combination of these constraint handling techniques.

1 Introduction

The relativistic klystron oscillator (RKO) is a high power microwave (HPM) source under development within the Air Force Research Laboratory [1]. Significant benefits would result from an increase in the device's signal growth rate. However, optimization of the growth rate via analytical and standard numerical techniques is intractable because of the high dimensionality of the design space and the existence of many local optima. Instead, the growth rate is optimized using the real-valued Evolutionary Algorithm (EA) GENOCOP [2]. Design optimization via EA pays off in two ways: better designs and improved understanding of models.

GENOCOP has been used widely and successfully for real-valued (continuous) parameter optimization problems. It offers a wide variety of selection, recombination, and mutation operators that have been shown to be effective in many practical applications. It supports constrained minimization and maximization problems, including those with non-linear equality constraints, as well as both linear and non-linear inequality constraints. It does so by maintaining separate search and reference populations. The former may contain both feasible and infeasible individuals, while the latter contains only feasible individuals. Highly fit search individuals are occasionally recombined with reference individuals to produce new reference individuals.

The next section gives brief descriptions of EAs and the RKO. The following three sections describe a series of computational experiments involving the use of GENOCOP to optimize the signal growth rate of the RKO. The first set of experiments was performed as a feasibility study using a very simple mathematical model of the RKO. The results of those experiments were encouraging, but highlighted the need for a more accurate mathematical model. One difference between the simple model and the more accurate one is that the solution space of the latter model is constrained. The second set of experiments uses GENOCOP's constraint-handling capabilities accordingly. Those constraint-handling capabilities are intentionally general, so that GENOCOP can be applied to any constrained optimization problem. However, this suggests the possibility that greater effectiveness can be achieved using domain-specific constraint-handling techniques. This hypothesis led to the third set of computational experiments. The remaining sections summarize conclusions based on those experiments and discuss ongoing related research.

2 Background

2.1 Evolutionary Algorithms

This section provides a very brief introduction to EAs, which are algorithms inspired by theories of evolution. Data structures called individuals are used to represent possible

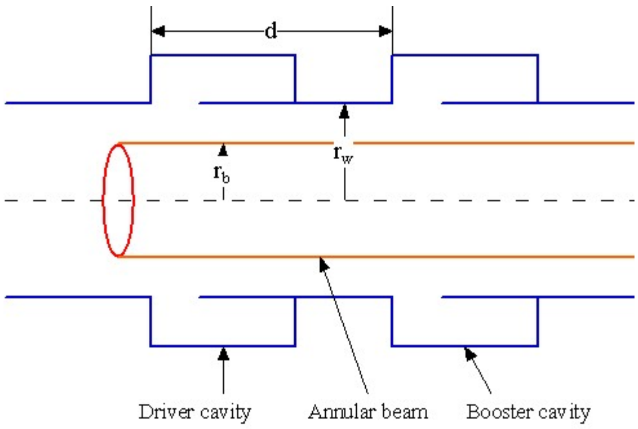

Figure 1: Relativistic Klystron Oscillator

solutions to a problem. The most distinguishing characteristic of EAs is that they manipulate collections of individuals, called populations, and that the operations performed on each individual may depend on the other individuals in the population. This is because the operations are inspired by the concepts of natural evolution, including fitness, selection, mutation, and recombination. Back and Schwefel formalize these ideas to some extent in the context of the three most prominent EA paradigms [3]. Merkle offers a more general and rigorous version that precisely captures the essential nature of EAs [4].

The concepts of fitness and selection are closely related, as are the fitness functions and selection operators that are inspired by those concepts and used in EAs. The fitness function assigns each individual a "fitness" based on some evaluation of the solution that it represents. In the case of optimization problems, which are perhaps the most common application area for EAs, the fitness function is related somehow to the objective function of the problem. The specific relationship is one of the design issues involved in applying EAs. Once the fitness of each individual has been assigned, a selection operator randomly selects individuals from the current population to copy into the next population. More fit individuals have higher probabilities of selection.

Mutation operators randomly alter individuals as a means of exploring the search space in neighborhoods of known good solutions. Recombination operators randomly select features from two or more individuals to create new individuals. In successful applications, the combined effect of the selection, mutation, and recombination operators is to gradually produce populations of individuals that represent very good solutions to the underlying problem, in analogy to the principle of survival of the fittest.

3 Relativistic Klystron Oscillator

Figure 1 shows a cross-section of the RKO, illustrating the features of its operation that are relevant to this research. An annular electron beam is injected into a cylindrical waveguide (the beam enters the figure from the left). A static axial magnetic field is applied, which has the effect of restricting electron motion in the radial and azimuthal directions (i.e. the electrons can be thought of as being restricted to moving approximately along the magnetic field lines, parallel to the centerline of the device). Electromagnetic oscillations are created in an annular cavity called the "driver cavity" by an external RF source. The resulting oscillating voltage across the cavity gap induces corresponding oscillations in the electron beam velocity. A second annular cavity called the "booster cavity" is downstream from the driver cavity. The cavities have weak electromagnetic coupling, meaning that the electromagnetic oscillations in the driver cavity propagate through the waveguide, and although they are attenuated along the way, they induce oscillations in the booster cavity [5]. The location of the booster cavity is chosen carefully so that its oscillations are in phase with the beam's oscillations. As a result, the latter are enhanced by the former. More thorough descriptions of the device and current experimental efforts are available elsewhere (e.g. [6, 7]).

At least two significant benefits would result from an increase in the RKO's RF signal growth rate. First, a sufficient increase would allow elimination of the external RF source, which contributes a substantial fraction of the system's total weight. Second, the electron beam is a limited duration pulse, and a higher signal growth rate would allow extraction of a greater fraction of the energy in the pulse.

4 Feasibility Study

The computational experiments in the feasibility study use a (lumped parameter) circuit model of the RKO in the absence of the electron beam to determine the "cold tube" resonant frequency, and use the relationship of that frequency to ω_0, the "natural frequency" (i.e. the resonant frequency of the driver and booster cavities), to determine the electromagnetic coupling coefficient C between the cavities. Finally, a dispersion relation involving the design parameters and C determines the resonant frequency of the RKO in the presence of the electron beam, along with its growth rate.

4.1 RKO Circuit Model

The circuit model is shown in Figure 2. It can be viewed as a series of three components. The left portion of the figure is an LRC circuit that models the driver cavity. The voltage across the capacitor models the gap voltage of the cavity. Likewise, the right portion of the figure is an LRC circuit that models the booster cavity. The LC series element in the "bridge" of the circuit models the transmission characteristic of the waveguide, while the other capacitor in the bridge models the attenuation.

The parameters L, R, and C are determined by ω_0 and the quality Q of the cavities, while the bridge parameters L_1, C_1, and C_2 are determined by the gap separation based on energy principles. Given these parameters, the admittance Y of the circuit is determined by the standard circuit

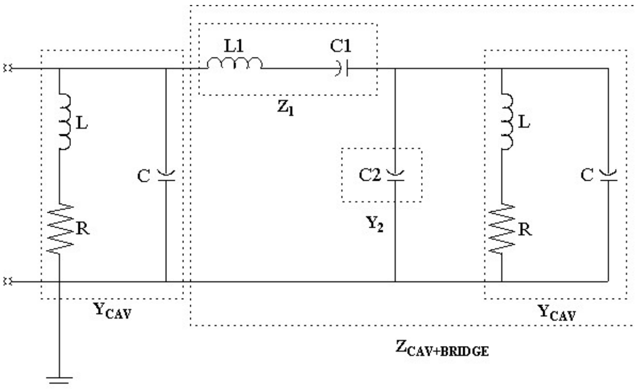

Figure 2: Two-cavity RKO Circuit Model

analysis indicated by Equations 1.

$$\begin{aligned} Z_1 &= j\omega L_1 + 1/(j\omega C_1) \\ Y_2 &= j\omega C_2 \\ Y_{CAV} &= j\omega C + 1/(R + j\omega L) \\ Z_{CAV+BRIDGE} &= Z_1 + 1/(Y_{CAV} + Y_2) \\ Y &= Y_{CAV} + 1/Z_{CAV+BRIDGE} \end{aligned} \quad (1)$$

The resonant frequencies of the circuit satisfy $Im[Y] = 0$. One of them is ω_0, and the other corresponds to the higher of the RKO's two cold tube resonant frequencies, which is the one of interest in practice.

4.2 Electromagnetic coupling

Given the resonant frequencies of the cold tube, the electromagnetic coupling constant C between the cavities is determined by the relationship

$$\omega = \omega_0 \left[1 + \frac{j}{2Q} \pm \frac{C}{2} \right] \quad (2)$$

where Q is the electromagnetic quality of the cavities.

4.3 Dispersion Relation

A similar relationship holds in the presence of the beam

$$\omega = \omega_0 \left[1 + \frac{j}{2Q} + \frac{C}{2} \sqrt{1 + \frac{Z}{CR} \sin(k_p d) e^{-j\theta}} \right] \quad (3)$$

where $R = V_0/I_0$ is the beam impedance, Z is the cavity impedance (an independent design parameter), θ is the phase velocity associated with the finite velocity of the beam due to the beam voltage V_0, and k_p is the wave number associated with the plasma oscillation and is inversely proportional to I_0. Thus, the signal growth rate $-Im[\omega]$, which is the quantity to be optimized, is defined as a function of V_0, I_0, and d.

Table 1: Comparison of signal growth rates of EA-identified design and previously considered designs.

V_0 (kV)	I_0 (kA)	d (cm)	Growth Rate (nsec^{-1})
Previously considered designs:			
400	12	8.4	-1.87e6
600	24	8.4	-1.45e6
400	12	11.0	-2.76e6
Design identified by the EA:			
400	25	9.4	-9.79e6

4.4 Computational Experiments

The feasibility study is a straightforward application of GENOCOP to the optimization of the signal growth rate using the model just described, with each individual in the EA consisting of the values of the three independent variables V_0, I_0, and d. Linear constraints (domains) are applied to each of the independent variables: $400kV \leq V_0 \leq 650kV$, $5kA \leq I_0 \leq 35kA$, and $2cm \leq d \leq 50cm$. Standard GENOCOP operators are used, including exponential ranking, uniform mutation, boundary mutation, non-uniform mutation, whole non-uniform mutation, whole arithmetic crossover and simple arithmetic crossover. Fifty independent experiments were performed, each with 70 individuals for 500 generations. GENOCOP performs 2 evaluations per individual per generation, so this set of experiments required 3.5 million total evaluations. The wall clock time on a 233 MHz Pentium II operating under MS Windows NT 4.0 was approximately 1 hour.

The signal growth rate is "optimized" easily. Many experiments found high growth-rate designs near V_0 = 400 kV, I_0 = 25 kA, d = 9.4 cm. The model predicts the signal growth rate of this design to be nearly an order of magnitude better than that of previously considered designs (see Table 1). High-fidelity simulation using the MAGIC Particle-in-cell (PIC) software [8] confirms the high signal growth rate of the design identified by the EA. In the simulation of one previously considered design, the signal begins to grow rapidly shortly after 150 ns (see Figure 3), while in the simulation of the design identified by the EA, it begins to grow rapidly about 40 ns earlier (see Figure 4). However, the simulation also predicts the formation of a virtual cathode (a region in which the electron density becomes so high that other electrons are repelled back towards the current source) and corresponding shutoff of the device.

5 Multi-cavity RKO

The success of the feasibility study suggests that EAs may be useful in optimizing other aspects of the RKO design as well. The beam current at which virtual cathode formation occurs depends on a number of other design parameters, but it can be predicted analytically and thus viewed as a non-linear inequality constraint on the design optimization problem. This section describes a set of computational experiments based on a dispersion relation model.

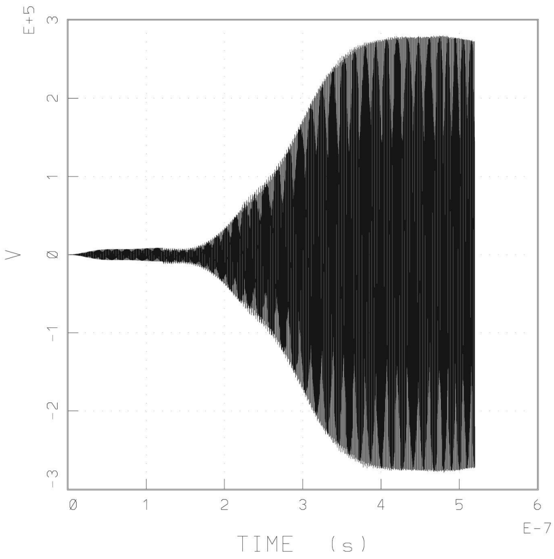

Figure 3: Booster cavity gap voltage for a previously considered design, as determined by MAGIC simulation

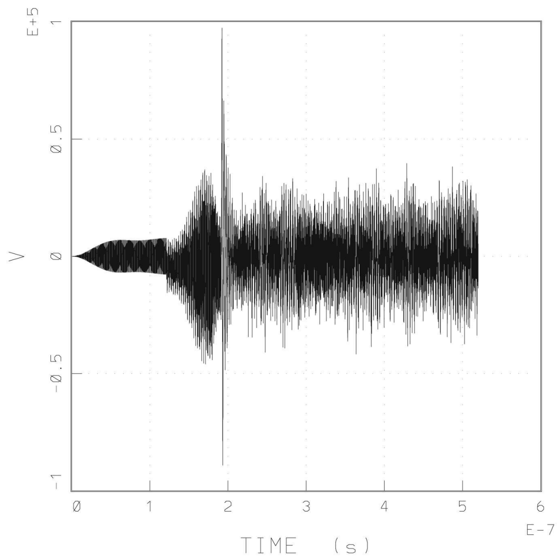

Figure 4: Booster cavity gap voltage for the design identified by the EA, as determined by MAGIC simulation

5.1 Dispersion Relation

The model used in these experiments is similar to the two-cavity model developed by Luginsland [5] in several respects:

- it describes the evolution of the gap voltages V_m in the frequency domain;
- it considers cavity resonances (oscillations within individual cavities);
- it considers electromagnetic coupling (oscillations induced in one cavity by oscillations from another cavity propagating through the waveguide in either direction);
- it considers beam coupling (oscillations induced indirectly in a downstream cavity by oscillations from an upstream cavity modulating the beam); and
- it assumes small signal, modal, steady-state solutions (so that the superposition principle applies to beam modulation, and thus the beam coupling between various pairs of cavities is independent).

On the other hand, the model used in these experiments is more general than Luginsland's original model in two respects. First, it allows cavities to have distinct natural frequencies, qualities, and impedances. Similarly, it allows the drift regions to have distinct radii, lengths, and loss coefficients. Finally, it allows $n \geq 1$ cavities. Thus, each candidate design consists of a vector of the following independent variables, for which the domains are given in Table 2:

- Beam voltage V_0,

Table 2: Independent variable domains.

Quantity	Lower Bound	Upper Bound
V_0	300 kV	650 kV
I_0	5 kA	35 kA
r_i	0.1 cm	12 cm
$r_o - r_i$	0.1 cm	1 cm
$f_{0,i}$	1 GHz	2 GHz
Q_i	50	500
$Q_i Z_i$	50 Ohms	377 Ohms
d_i	2 cm	50 cm
$\chi_{r,i}$	0	1
$\chi_{c,i}$	0	1

- Beam current I_0,
- Beam inner radius r_i,
- Beam thickness $r_o - r_i$,
- Cavity natural frequencies $f_{0,1}, \ldots, f_{0,N}$,
- Cavity qualities $Q_1, \ldots, Q_N$,
- Cavity impedances $Q_1 Z_1, \ldots, Q_N Z_N$,
- Drift space lengths $d_1, \ldots, d_{N-1}$,
- Drift space radius multipliers $\chi_{r,1}, \ldots, \chi_{r,N-1}$, and
- Drift space electromagnetic coupling multipliers $(\chi_{c,1}, \ldots, \chi_{c,N-1})^T$

The model has the form

$$L_m(\omega)V_m + \sum_{n<m} \Gamma_{m,n} V_n + \sum_{n \neq m} C_{m,n} V_m = 0 \qquad (4)$$

where the remainder of this section defines and discusses L_m, $\Gamma_{m,n}$, and $C_{m,n}$. First,

$$L_m(\omega) = \frac{\omega^2}{\omega_{0,m}^2} - \frac{j\omega}{\omega_{0,m}Q_m} - 1 \tag{5}$$

is the damped harmonic oscillator operator for cavity m. If cavity m were isolated from the rest of the system, the equation $L_m V_m = 0$ would hold. Next,

$$\begin{aligned}\Gamma_{m,n} &= \frac{Z_m}{R}\sin\left\{\sum_{r=n}^{m-1} k_{p,r}(x_{r+1}-x_r)\right\} \\ &\quad \exp\left[-\frac{j\omega_{0,n}}{\beta c}(x_m - x_n)\right]\end{aligned} \tag{6}$$

is the beam coupling coefficient between cavities the upstream cavity n and the downstream cavity m. Finally, $C_{m,n} = C_{n,m}$ is the electromagnetic coupling coefficient between cavities m and n. However, as stated above, the model assumes that cavity coupling is weak and occurs through a cutoff waveguide. This implies that only nearest neighbor electromagnetic coupling is significant, i.e. unless $|n-m| = 1$, $C_{m,n} = 0$. Thus, Equation 4 can be reduced to

$$L_m(\omega)V_m + \sum_{n<m} \Gamma_{m,n}V_n + C_{m-1}V_m + C_{m+1}V_m = 0 \tag{7}$$

where

$$C_0 = 0 \text{ , and otherwise} \tag{8}$$

$$\begin{aligned} C_m &= C_{m,m-1} \\ &= \chi_{c,m}\exp\left[\frac{2.405}{r_{w,m}}(x_{m+1}-x_m)\right. \\ &\quad \left.\times\sqrt{1-\left(\frac{2\pi\omega_{0,m}r_{w,m}}{0.383c}\right)^2}\right]\end{aligned} \tag{9}$$

5.2 Physical Constraints

The model includes three physical constraints. First, each drift space must be cutoff (to satisfy the assumption of the model):

$$r_{w,m} < \frac{0.383c}{f_{0,m}} \tag{10}$$

Second, the beam must fit within each drift region:

$$r_0 < r_{w,m} \tag{11}$$

Finally, the beam current must not exceed the limiting current of any drift space [9]:

$$\begin{aligned} I_0 &\leq 17000\left[\left(1+\frac{V_0}{mc^2}\right)^{\frac{2}{3}} - 1\right]^{\frac{3}{2}} \\ &\quad \times\left[1 - 2\left(\frac{r_i^2}{r_0^2 - r_i^2}\log\frac{r_0}{r_i} - \log\frac{r_{w,m}}{r_0}\right)\right]^{-1}\end{aligned} \tag{12}$$

5.3 Computational Method

The evolution of the cavity voltages $V = (V_1, V_2, , V_N)^T$ is thus described by $[A(\omega)]V = 0$, where

$$A(\omega)_{ij} = \begin{cases} \Gamma_{i,j} & \text{if } i > j+1 \\ \Gamma_{i,j} - C_j & \text{if } i = j+1 \\ L_i(\omega) & \text{if } i = j \\ -C_i & \text{if } i = j-1 \\ 0 & \text{if } i < j-1 \end{cases} \tag{13}$$

The resonant frequencies ω of the device satisfy $\det[A(\omega)] = 0$. The determinant is a polynomial of degree $2N$ in ω. The mode's growth rate is $-Im[\omega]$, which is to be maximized.

The computational determination of the feasibility of a candidate design is implemented as non-linear inequality constraints in GENOCOP as follows. The first set of constraints combines the physical constraints on the drift space and the beam radii:

$$\left(0.95\frac{0.383c}{f_{0,m}}\right) - (r_0 + 0.2cm) \geq 0 \tag{14}$$

If the first set of constraints is satisfied, then the drift space radii are computed:

$$r_{w,m} = \chi_{r,m}\left(0.95\frac{0.383c}{f_0, m}\right) + (1-\chi_{r,m})(r_0 + 0.2cm) \tag{15}$$

The second set of constraints ensures that the limiting currents are not exceeded:

$$\begin{aligned} &17000\left[\left(1+\frac{V_0}{mc^2}\right)^{\frac{2}{3}} - 1\right]^{\frac{3}{2}} \times \\ &\left[1 - 2\left(\frac{r_i^2}{r_0^2 - r_i^2}\log\frac{r_0}{r_i} - \log\frac{r_{w,m}}{r_0}\right)\right]^{-1} - I_0 \geq 0\end{aligned} \tag{16}$$

Thus, the evaluation of candidate designs includes the following steps:

1. Repair the design if necessary (other possibilities are addressed later).
   (a) Reduce the beam radius if it (almost) exceeds a cutoff radius.
   (b) Reduce the beam current if it exceeds a limiting current.
2. Compute the electromagnetic coupling coefficients (C's).
3. Compute the beam coupling coefficients (Γ's).
4. Compute the harmonic operator coefficients (L's).
5. Construct the $N \times N$ matrix $A(\omega)$. The elements are polynomials in ω, and are represented by their coefficients.

6. Reduce $A(\omega)$ to lower triangular form. Specifically, for rows $i = N - 1$ down to 1, and each element $[a(\omega)]_{i,j}$ in row i.
   (a) Multiply by $[a(\omega)]_{i+1,i+1}$.
   (b) Subtract $[a(\omega)]_{i,i+1}[a(\omega)]_{i+1,j}$.

   $\det([A(\omega)])$ is now stored in $[A(\omega)]_{1,1}$ as a polynomial in ω of degree $2N$.
7. Use Laguerre's method to find the roots of $\det(a[\omega])$.
8. Choose the root s.t. $Re[\omega] > 0$ and $Im[\omega]$ is minimized.
9. Assign $Im[\omega]$ as the fitness of the candidate design.

5.4 Computational Experiments

Fifty independent experiments of 100,000 generations each are performed, with 5 (necessarily feasible) individuals in the reference population and 20 (possibly feasible) individuals in the search population. Because GENOCOP performs 2 evaluations per individual per generation, this series of experiments require a total of 250 million evaluations. The wall clock time using a 750 MHz P-III running Red Hat Linux 7 is approximately 14 hours.

The experiments identified high growth-rate, non-intuitive, and dissimilar designs. In comparison to a 10-cavity version of one good two-cavity design, for which the growth rate is 1.30 nsec^{-1}, the best growth rate in these experiments is 2.10 nsec^{-1}. The enhanced growth rates of the 10-cavity design allow pure oscillator operation (the two-cavity design requires injection-locked operation, i.e. an external RF source for the driver cavity). The parameters in most of the designs differ significantly between cavities, and between drift spaces. The best designs from various experiments are dissimilar except for the beam voltage (511 $\pm$ 38.5 kV) and the cavity frequencies (1.54 $\pm$ 0.09 GHz). The dissimilar results suggest the EA designs may be far from the global optimum.

6 Constraint Handling

In the context of the GENOCOP algorithm, constraints are relevant in at least three respects. First, constraints necessarily affect either the evaluation of individuals, the search space, or both (the same is true in the context of any optimization algorithm). The second respect in which constraints are relevant to the GENOCOP algorithm is one of the primary characteristics that distinguishes GENOCOP from most other EAs: it is possible to include the physical constraints in the GENOCOP problem specification. Doing so guarantees that all individuals in the reference population satisfy the physical constraints, along with all of the individuals in the initial search population. Finally, and typically overlooked, the constraints in the GENOCOP problem specification affect the recombination of search individuals with reference individuals.

The remainder of this section briefly identifies the standard EA techniques for handling constrained optimization problems and how they apply in the context of this research. It then discusses GENOCOP's standard technique for recombining search individuals with reference individuals, a feature of that technique that may lead to reduced effectiveness in many applications, and a novel technique to address that effect. The section concludes with a discussion of computational experiments related to constraint handling in this application.

6.1 Standard EA Constraint Handling Techniques

Standard EA techniques for handling constrained optimization problems include penalty functions, repair operators, and specialized operators that preserve feasibility. There are efficiency and effectiveness tradeoffs among these approaches that should be taken into consideration in any application of an EA to a constrained optimization problem.

The penalty function approach essentially treats the original constrained optimization problem as an unconstrained optimization problem in which individuals that do not satisfy the constraints of the original problem are "penalized" through the fitness function. For the problem addressed in this research, the vast majority of candidate designs have fitnesses better than zero. Thus, the obvious way to apply the penalty function approach to this problem is to assign zero fitness to any candidate designs for which the limiting current is exceeded, as well as any candidate designs for which the beam diameter is greater than (the arbitrarily specified fraction of) the minimum waveguide diameter. Note that a penalty function does not affect individuals that satisfy physical constraints. In particular, if the physical constraints are included in the GENOCOP problem description, then the penalty function will not affect the evaluation of any individuals in the reference population.

The repair operator approach treats the problem as an unconstrained optimization problem, but individuals that do not satisfy the constraints of the original problem are mapped to individuals that do (and that are in some sense "close" to the infeasible individuals) and then evaluated. The fitness of the repaired design is then assigned to the candidate design. The repair operation can be viewed as a form of efficient local optimization, so when the repaired individual is used only for the purpose of fitness evaluation, this approach is a Baldwinian memetic algorithm. When the repaired individual replaces the infeasible individual in the population, the approach is a Lamarkian memetic algorithm. The obvious way to apply the repair operator approach to fitness evaluation in this problem is to reduce the current to the limiting current if necessary, and to reduce the beam diameter to fit in waveguide if necessary.

When the technique is applicable, the use of specialized operators that preserve feasibility is typically very effective. However, in most interesting constrained optimization problems such operators are either very difficult to construct or very inefficient. The use of specialized mutation and recombination operators that preserve feasibility would be very computationally expensive in this application.

6.2 Recombination of Search and Reference Individuals

One of the distinguishing features of the GENOCOP algorithm is the maintenance of a separate reference population consisting entirely of feasible individuals. A related feature is that highly fit search population individuals are periodically recombined with reference individuals to produce new feasible individuals for inclusion in the reference population. The operator must ensure that specified constraints are satisfied.

Because GENOCOP is designed to be broadly applicable, the operators that it provides are "blind" operators based on convex combinations. In many applications of interest, the constraints may involve a strict subset of the parameters, but GENOCOP's convex combination operators are applied to all of the parameters. Thus, the individuals produced by these operators may differ from the search individuals in ways that do not affect the feasibility of the individual. In such applications, a domain specific operator that modifies only the parameters of the search individual that are involved in the constraints might tend to result in the production of more highly fit reference individuals.

For example, in the problem under investigation in this research, the constraints involve only the natural frequencies of the cavities, the beam's inner and outer radii, the drift space radius multipliers, and the beam current. They do not involve the beam voltage, the qualities and impedances of the cavities, the gap separations, or the drift space electromagnetic coupling multipliers. Thus, given any reference individual and any search individual, a new feasible individual can be constructed by performing a convex combination with respect to the former set of parameters while retaining the latter set of parameters from the search individual.

6.3 Computational Experiments

The computational experiments described in this section investigate the effect on effectiveness of three decisions relating to the handling of constraints for the optimization of the multi-cavity RKO design using the GENOCOP algorithm. The first choice is whether or not to include the physical constraints in the GENOCOP problem description. The second choice is whether to use the penalty function or the repair method approach in the evaluation of individuals. The third choice is whether to use one of the blind operators provided by GENOCOP or a domain specific operator for recombination of search and reference individuals. The domain specific operator performs a convex combination with respect to beam current and beam radius and retains all other parameters of the search individual.

As with the computational experiments described in the previous section, the algorithms used in these experiments are Baldwinian memetic algorithms, i.e. infeasible individuals are repaired for evaluation purposes, but are not replaced by the resulting individuals.

6.4 Computational Experiments

Eight sets of experiments are performed covering all possible combinations of the three choices identified in the previous subsection. For each combination of constraint handling techniques, 50 independent runs of 100,000 generations are performed, each with 5 individuals in the reference population and 20 individuals in the search population. GENOCOP performs 2 evaluations per individual per generation, resulting in a total of 250 million total evaluations. The experiments completed in approximately 14 hours on a 750 MHz P-III running Red Hat Linux 7.

Each experiment found high growth-rate designs. For comparison, a 10-cavity version of one good two-cavity design has a growth rate of 1.30 nsec^{-1}, while the best growth rate identified by the EA has a growth rate of 2.40 nsec^{-1}. The enhanced growth rates of the 10-cavity design are sufficient to allow pure oscillator operation (two-cavity design requires injection-locked operation). The designs identified by the EA are non-intuitive (typical of EA-based design). The best designs from various experiments are dissimilar, suggesting that the global optimum has not been located. The parameters differ significantly between cavities, and between drift spaces, suggesting that it is advantageous to allow these parameters to vary independently.

The growth rates of best designs resulting from each of the three choices are compared using the Kruskal-Wallis test. The domain specific operator for recombination of search and reference individuals is better than blind recombination at the 0.05 level of significance. Repairing infeasible individuals for evaluation is better than the penalty function method at the 0.10 level of significance. The growth rates of the best designs resulting from each of the combinations of choices are also compared using the Kruskal-Wallis test, with the following results:

- Effective combinations:
  - Heuristic recombination and repair to evaluate
  - Specified constraints and heuristic recombination
  - Specified constraints, heuristic recombination, and repair to evaluate
  - Unspecified constraints, blind recombination, and repair to evaluate
- Ineffective combinations
  - Blind recombination and penalty function
  - Specified constraints and blind recombination
  - Specified constraints, blind recombination, and repair to evaluate
  - Specified constraints, blind recombination, and penalty

7 Conclusions

7.1 Feasibility Study

High-fidelity simulation of the 2-cavity RKO design identified by the EA varying only beam voltage, beam current and gap separation confirms the high signal growth rate.

However, it also predicts the formation of a virtual cathode and corresponding shutoff of the device. This illustrates the ability of design optimization via EA to lead to improved understanding of mathematical models by exploring non-intuitive regions of the search space.

7.2 Multi-cavity RKO

Computational experiments identified high growth-rate, non-intuitive, and dissimilar designs. The enhanced growth rates of the 10-cavity design allow pure oscillator operation and thus a substantial reduction in the total weight of the device. The parameters in most of the designs differ significantly between cavities, and between drift spaces, suggesting that there is value in allowing these parameters to vary independently. The best designs from various experiments are dissimilar suggesting that the EA designs may be far from the global optimum, and that additional design improvements may be possible by improving the effectiveness of the algorithm.

7.3 Constraint Handling

Computational experiments involving eight combinations of constraint handling techniques produce significant differences in the effectiveness of the algorithm. Most significantly, the domain specific operator for recombination of search and reference individuals is more effective than blind recombination. Also, repairing infeasible individuals for evaluation is more effective than penalizing them.

Interestingly, no statistically significant difference is observed as a result of including the physical constraints in the GENOCOP problem specification. This is somewhat surprising given that if the constraints are not included in the problem specification, there is no real distinction between the search population and the reference population. Furthermore, individuals in the initial reference population are not constrained, so the population is unlikely to contain any individuals that satisfy the physical constraints. If the penalty function method is used, then the initial population is likely to have all zero fitnesses.

8 Future Directions

The results of the computational experiments involving constraint handling techniques are difficult to interpret, because GENOCOP provides little diagnostic output. Future research will involve the modification of GENOCOP to provide fitness statistics and diversity measures for both the search and reference populations

The results of the computational experiments result in dissimilar designs, suggesting improving the effectiveness and efficiency of the algorithm might lead to the identification of higher growth rate designs. Future research will involve the use of other memetic techniques (Lamarkian, etc.) and hybridization with local search (e.g. conjugate gradient). A farming model parallel version of GENOCOP is under development. Also, the efficiency of the fitness evaluation would be improved if a technique could be developed to reduce the number of roots found. The Lehmer-Schur algorithm is a promising candidate [**?**].

Several avenues for improvement of the theoretical and computational models are available. One is to consider the limiting currents at cavity gaps. Another is to consider mode competition and sensitivity to design parameters.

Once the effectiveness of the algorithm is improved sufficiently to suggest that a global optimum has been located, the design will be evaluated using a high fidelity PIC simulation.

Acknowledgments

This research has been supported by the Air Force Research Laboratory High Power Microwave Division, as well as the United States Air Force Academy Department of Computer Science.

Bibliography

[1] K. J. Hendricks, P. D. Coleman, R. W. Lemke, M. J. Arman, and L. Bowers, “Extraction of 1 GW of rf power from an injection locked relativistic klystron oscillator,” *Physical Review Letters*, vol. 76, no. 1, pp. 154–157, January 1996.

[2] Z. Michalewicz, *Genetic algorithms + data structures = evolutionary programs*. New York: Springer-Verlag, 1992.

[3] T. Back and H.-P. Schwefel, “An overview of evolutionary algorithms for parameter optimization,” *Evolutionary Computation*, vol. 1, no. 1, pp. 1–23, Spring 1993.

[4] L. D. Merkle and G. B. Lamont, “A random function based framework for evolutionary algorithms,” in *Proceedings fo the Seventh International Conference on Genetic Algorithms*, T. Back, Ed., 1997, pp. 105–112.

[5] J. W. Luginsland, Y. Y. Lau, K. J. Hendricks, and P. D. Coleman, “A model of injection-locked relativistic klystron oscillator,” *IEEE Transactions on Plasma Science*, vol. 24, no. 3, pp. 935–937, June 1996.

[6] J. D. Blahovec, Jr., L. A. Bowers, J. W. Luginsland, G. E. Sasser, and J. J. Watrous, “3-D ICEPIC simulations of the relativistic klystron oscillator,” *IEEE Transactions on Plasma Science*, vol. 28, no. 3, pp. 821–9, June 2000.

[7] K. J. Hendricks and M. D. Haworth, “Experiments on high-power microwave transmission through a belljar,” *IEEE Transactions on Plasma Science*, vol. 30, no. 3, pp. 1215–19, June 2002.

[8] B. Goplen, L. Ludeking, D. Smithe, and G. Warren, “User-configurable MAGIC for elecromagnetic PIC calculations,” *Computer Physics Communications*, vol. 87, no. 1-2, pp. 54–86, May 1995.

[9] R. B. Miller, *An introduction to the physics of intense charged particle beams*. New York: Plenum, 1982.