Automatic Optimization of a Klystron Interaction Structure

Christopher James Lingwood, Graeme Burt, Kester James Gunn, Richard G. Carter, *Senior Member, IEEE*, Rodolphe Marchesin, and Erk Jensen

***Abstract*— The design of klystrons has long been a manual process guided by experience. However, with well-defined specifications and sufficiently rapid simulation methods, it is a good candidate process for automatic optimization techniques. In this paper, such a technique is evaluated and refined using klystron specific techniques, leading to several designs (with different tradeoffs between efficiency and size) each of a structure comparable with the SLAC B-factory klystrons. The most efficient of which, while only 1% more efficient, is 17.1% shorter.**

***Index Terms*— Efficiency, evolutionary algorithm (EA), interaction structure, klystron, klystron design, optimization.**

I. Introduction

With the high demand for higher efficiency klystrons, the optimization of interaction structures is of continued importance. This demand is driven by proposed experiments demanding higher beam powers (such as Compact Linear Collider [1], [2] and European Spallation Source [3]) balanced against budgetary and environmental pressure for lower energy usage [3].

When an interaction structure is designed manually it can take significant additional time to evaluate design changes and ideas. Although the parameter space is simple enough to be evaluated by hand, the process is time consuming and there is a risk of becoming trapped in a local optimum.

The well-defined specifications and input parameters make klystron interaction structures suitable for automatic optimization. This paper will focus on the use of an evolutionary algorithm (EA) to design such a structure from a simple specification. This approach allows a quick evaluation of specification changes.

Some modification to the generic algorithm is required to allow it to converge to a useful set of klystron parameters. This is made necessary by, for instance the ease with which one can design a very short but very inefficient klystron.

Manuscript received February 3, 2013; revised May 14, 2013 and June 13, 2013; accepted June 13, 2013. Date of current version July 19, 2013. This work was supported by the STFC and Thales Group. The review of this paper was arranged by Editor M. Thumm.

C. J. Lingwood, G. Burt, K. J. Gunn, and R. G. Carter are with the Engineering Department, Lancaster University, Lancaster LA1 4YR, U.K. (e-mail: c.lingwood@lancaster.ac.uk; g.burt@lancaster.ac.uk; k.gunn@lancaster.ac.uk; r.carter@lancaster.ac.uk).

R. Marchesin is with Thales Group, Velizy 78140, France (e-mail: rodolphe.marchesin@thalesgroup.com).

E. Jensen is with CERN, Geneva 1217, Switzerland (e-mail: erk.jensen@cern.ch).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TED.2013.2270975

The interaction structure of the B-factory klystron [4] will be used as a test problem for this technique, with all cavity frequencies and drifts as free parameters. The resulting structures are evaluated and compared with the reference design.

II. Generic Optimization Techniques

EAs are an example of a probabilistic, parallel, black-box search technique inspired by natural evolution. Like many optimization methods, they work from a list or population (**P**) of individuals (**X**) which are a related set of decisions (**x**–i.e., simulation inputs) and solutions (**z**–in this case simulation outputs) characterizing the search space. For each iteration the best individuals are selected and used as parents in a reproduction operator to share information and create offspring (ϕ). This allows for complex, epistatic (nonlinear) and multimodal (multiple optima and/or suboptima) optimization problems to be efficiently explored [5], [6].

Basic genetic algorithms (GAs) encode the decision space in binary, and then using the analogy of gene manipulation, cut and splice these often by using the concept of mutation to add an extra random element. Their operation is analogous to evolution; successful given time but unfocused [5].

Real-coded EAs are distinct from GAs in that they manipulate the parameters numerically rather than encoding them as binary information. Using the available (often incomplete) information, various strategies are used to analyze the solution space to identify good areas worthy of further investigation [5], [7]–[9].

Real-coded EAs are used in a wide range of scientific applications and their characteristics are well understood. They are well suited to performing optimizations in problems with a high dimensionality.

To compare the individuals the concept of fitness (ƒ a scalar) is used, it being higher for individuals in a population closer to the optimum. For a single-solution maximization problem, this can simply be the scalar solution, where bigger is better. For a multiobjective (vector solution) problem this is more complicated.

A. Selection

Selection is based on the fundamental assumption that good individuals, when combined, will produce good offspring and that there is a correlation between parent and offspring fitness [6].

Here a simple selection method, tournament selection, is used. A predetermined number of individuals are selected at random from the population with the individual with the highest fitness selected. The size of the tournament T strongly affects the selection intensity and the exploitation-exploration balance [6].

B. Reproduction

In an EA, it is mathematical analysis that creates the new decision vectors or offspring. This allows complex analysis and representation of the space [8].

Reproduction of operators act over a population of parents selected using a method described above. In parent centric operators [8] a particular parent is nominated to be the center of reproduction: the mother [7]. Here, a reproduction operator called mother centric recombination (MrCX) [10] is used. It has a control parameter σ_x called the search strength, which defines how far from the parents the offspring will be produced.

For this paper, a search strength of 0.67 is used, which is found to be generally effective for the standard test problems evaluated [7], [10], [11]. These test problems are simply functions with useful properties (such as a high dimensionality and multiple maxima) on which the algorithm is repeatedly applied to find the best settings. While these functions have no specific applicability to a klystron structure beyond the complexity of their solutions, the performance of an algorithm using this value should be sufficient.

C. Pareto Optimization

In any multiobjective simulation there exists a set of solutions which are demonstrably superior to all other solutions. Such a set is known as a nondominated set and is an example of a Pareto front. Once converged, such a set contains individuals where no one characteristic can be further optimized without detrimental effect on the other characteristics. The remaining solution space can be described as dominated. Given the opportunity, a decision maker may be interested in knowing alternative solutions [12]. If this set is removed then a further Pareto front is exposed. If this process is repeated a series of fronts can be exposed, allowing each individual to be assigned a Pareto front index.

The Pareto fronts are found using nondominated sorting that uses the operator

$$\mathbf{z_a} \preceq \mathbf{z_b} \tag{1}$$

where $\mathbf{z_a}$ and $\mathbf{z_b}$ are the solution vectors and $\preceq$ means all components are greater than or equal to.

The effect of (1) on a solution space dominated by a series of points can be seen in Fig. 1.

The fitness operator used in this paper (NSGAII [13]) uses the Pareto front index to assign a scalar fitness, comprised of the front index added to a normalized estimation of the local density of solutions. It is through the latter method that the algorithm attempts to spread the solutions along the Pareto fronts.

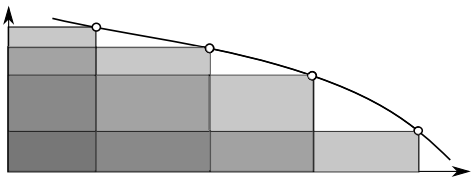

Fig. 1. Effect of (1) on the space behind a series of points. Darker colors: more highly dominated area for a maximization problem. The points shown could represent a Pareto front.

III. Evolutionary Algorithm Setup

When such a generic optimization method is applied to klystron interaction structures a number of issues become apparent both during setup and during optimization. The reasons for, and potential solutions to these problems are detailed here.

Furthermore, not all results from the simulation code will be physical. For instance, for very bad structures it is possible to get a negative efficiency and not all simulations will converge in the number of iterations allowed. These erroneous or unconverged simulations are identified and discarded.

As a generic technique, the EA does not have any expert knowledge of how to improve a klystron interaction structure. The algorithm starts with a population of random guesses in a predefined range. These will have very low efficiency and will not meet the bandwidth specification. Each iteration of the algorithm uses the best structures [as defined by the solutions and NSGAII/ Pareto search pressure focusing (PSPF)] to attempt to predict a number of better structures using MrCX, which are then evaluated. Each cavity and drift are optimized concurrently and as the optimization progresses the range of values for each decision narrows, and the overall efficiency of the population increases. During the optimization a Pareto front emerges, showing the best efficiency for a given length. This front then advances toward shorter, more efficient klystron structures. Almost all structures in the final population will meet the bandwidth specification.

It is necessary to select a number of parameters when setting up the EA. While poor selection of these parameters leads to poor convergence, the problems are often easily identified. For instance, a tendency to become stuck in local maxima would indicate too high a search pressure, perhaps solved by relaxing the selection of parents. The parameters, with the exception of search strength, do not need to be tightly optimized to allow the algorithm to work well.

Once set up, it is simple to reconfigure the algorithm to work on a new structure. Only the allowable ranges of the decisions (cavity frequencies and drifts) and the simulation input file they are mapped to are structure specific.

A. Pareto Search Pressure Focusing

Generic multiobjective EAs that use the Pareto front concept are often demonstrated in publications using mathematical test functions. These often have not only a clear transformation between decision and solution space, but also and critically, every point on the first Pareto front is equally optimal.

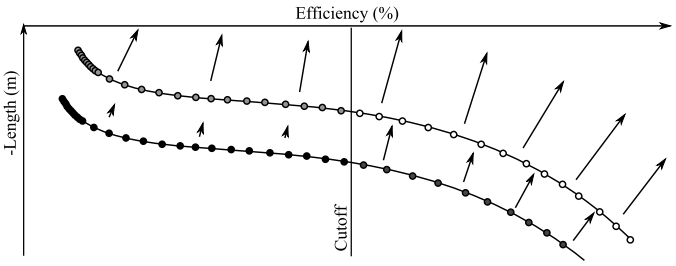

Fig. 2. Demonstration of the effect of PSPF applied to efficiency on the length, efficiency plot (note the sign of the y axis). Shade: fitness from high (white) → low (black). The magnitude of the arrows represents the relative search pressure (derived from the fitness).

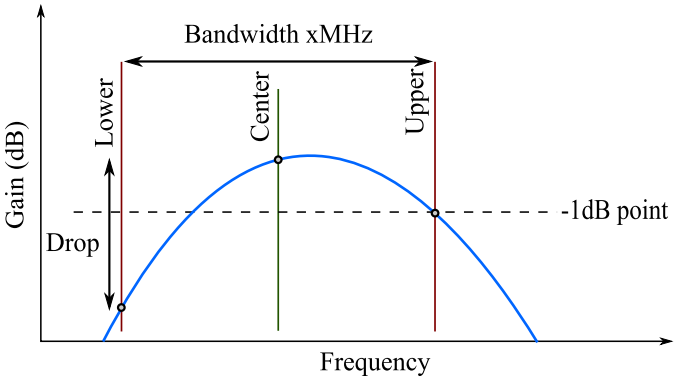

Fig. 3. More suitable definition of bandwidth for use in an optimizer.

This interest in the full extent of the front does not apply to klystron interaction structures.

This algorithm is intended to target short, efficient klystrons. However there exists a significant set of decisions which result in a short, very inefficient klystron. Using a pure Pareto front concept these are equally as optimal as a longer, significantly more efficient klystron. This causes the algorithm to waste time on areas of the solution space that are of no interest. As a result, without modification this concept does not usefully optimize a klystron interaction structure.

As fitness is defined by Pareto front index (II-C) in NSGAII and the Pareto concept still has useful behavior, a refinement, PSPF is proposed (Fig. 2).

An arbitrary cutoff point is defined either as a fraction of the maximum of a single dimension or a fixed point. Individuals above the cutoff in this dimension have their fitness unchanged, but individuals below this point are defined as less fit. This decreases the search pressure below the cutoff point. This single dimension focusing effect is felt through the fronts, so individuals above the cutoff are always preferentially selected.

Application of this technique to efficiency is shown in Fig. 2. Two Pareto fronts are shown (with the upper being the fittest) for which each individual would be equally fit. For each of these fronts individuals with an efficiency below the cutoff are defined as being less fit, reducing the search pressure in this region. This causes the algorithm to ignore the short, inefficient klystrons as desired. The variable cut-off has very useful behavior, as solutions found above the cutoff in efficiency, are at least initially the exception rather than the rule. A variable cutoff allows an increase in search pressure for the most efficient individuals from the start.

B. Bandwidth Targeting

Optimization of the bandwidth presents two problems: traditional computation of the bandwidth is computationally expensive if a small signal analytical solution is not used; optimization beyond the specification bandwidth is unnecessary.

The algorithm must also preferentially select individuals closer or equal to the specification bandwidth. To achieve this and solve the above problems the output power is calculated in three positions, the center and the upper and lower bandwidth points as shown in Fig. 3. The objective used for optimization is the sum of the drop between the center ($P_{=}$) and the bandwidth points ($P_{+/-}$) in dB (Δ_{gain}); see (2). This objective is allowed to saturate at the point that the specification is met. For a bandwidth specified at 1 dB all the individuals within specification report a Δ_{gain} of 2 dB and as such are identically fit in this dimension, the intended behavior.

$$\Delta_{\text{gain}} = \begin{cases} 10 \cdot \left[\log_{10}\left(\frac{p_-}{p_=}\right) + \log_{10}\left(\frac{p_+}{p_=}\right)\right] & \text{when } < -2 \\ -2 & \text{when } > -2. \end{cases} \tag{2}$$

Pareto search pressure focusing is used to strongly target the structures within specification; without this a full range of bandwidths would be targeted. A simple rejection method would not be effective, as the initial population would be full of klystrons which do not meet the specification, which would then all be discarded.

The saturation of this objective leads to further valuable behavior. At the start of the optimization it behaves as a true objective in a $|\mathbf{z}|$ objective problem, but when the algorithm has converged and almost all the individuals meet the bandwidth specification it behaves as a $|\mathbf{z}| - 1$ objective problem. This requires fewer individuals to adequately describe the solution space at convergence. It should be noted that this approach is only applicable to narrowband tubes, as broadband tubes may have several peaks and troughs in their frequency response.

C. Reflected Electrons

As the klystron interaction structure is optimized the output cavity will remove increasing amounts of kinetic energy from the beam. If the electrons leave the output gap too slowly the build of space charge can potentially cause back-streaming; shortening the tube life or creating a feedback mechanism and causing oscillation. This effect limits the retarding potential achievable, limiting output power and so efficiency.

Significantly higher efficiencies are available when reflected electrons are tolerated, and the algorithm if not restrained will explore this portion of the decision space.

Rejecting simulations with reflected electrons is the simplest approach, with the advantage that in the initial population rejection is unlikely. However a structure that reflects the electrons could differ only slightly from a structure which

TABLE I
REFERENCE DESIGN PARAMETERS [15]

Parameter		Units
Operating Frequency (f)	476	MHz
Output power saturation (P_{out})	1.200	MW
Efficiency @ saturation (η)	>60	%
Saturated Gain (G)	>43	dB
Beam Voltage (V_b)	85	kV
Beam Current (I_b)	23	A
1 dB Bandwidth	3	MHz

does not. Another approach is to use the velocity of the slowest electron is as an objective. Using PSPF an absolute cutoff point can be defined, below which solutions are deemed undesirable. However, in practice this proved to complicate the results unnecessarily so the rejection method is used here.

D. B-factory klystron

The detailed specification of the klystron designed for the Asymmetric Storage Ring B factory at SLAC is in the public domain and will be referred to in this paper as the reference design in Table I [14]. However, depending on publication the beam parameters differ. For the purposes of this paper V_b and I_b stated in Demmel *et al*. [15] are taken as authoritative due to its inclusion of simulation data. Additionally, AJDisk [16] (a refinement of JapanDisk, which was used in part to design that klystron) predicts reflected electrons for lower beam voltages.

A short group delay is also specified to allow it to rapidly correct both amplitude and phase to damp oscillations in the accelerating cavities [4]. To take into account this part of the specification the short group delay would be added as an objective. As this is mainly determined by the Q of the intermediate cavities, and this will not be varied, this specification will be ignored in this paper.

The B-factory klystron was specified in Demmel *et al*. [15] as both a six and seven cavity klystron, with the latter using a second harmonic cavity as the fourth cavity. This paper focuses on the seven cavity structure as the more complex and efficient of the two cases.

It should be noted that this paper considers only the comparison of simulation results, not experimental performance. Demmel *et al*. [15] reported that a model in Field Charge Interaction (FCI) [17] (Field Charge Interaction Program, a 2.5-D particle in cell code) predicts that the B factory klystron saturates at 1.33 MW (67%) with an input a power level around 40 W [15]. Fig. 4 shows same parameters in AJDisk predict a higher efficiency of 71% at 45 W input power. As FCI is a more detailed simulation, taking into account for instance radial movement, this difference is not unexpected. Overlaid on Fig. 4 is the transfer curve of the structure Ξ_{eff} (the most efficient structure found in the following optimization). For this optimization an input power of 40 W is used.

The model in AJDisk is found to converge at 70 disks and 25 steps. Multiple rings are not used as convergence can be unpredictable. This leads to a short runtime, there being three evaluations to find Δ_{gain}. To enable the simulations to be performed in a reasonable time a CONDOR [18] pool of ~30

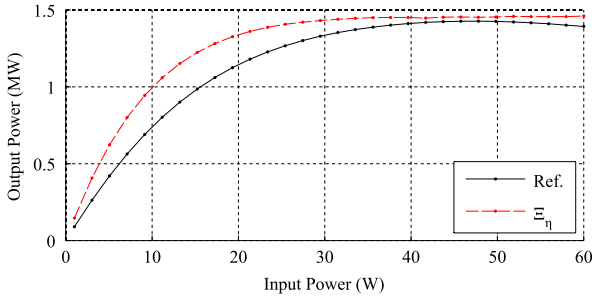

Fig. 4. AJDisk simulated transfer curve for the B-Factory klystron reference design and the most efficient structure found (Ξ_η).

Windows desktops is used. Each of the optimized structures presented are obtained in around 12 hours using this resource, critically no user input is required after initialization.

For every cavity, the drift length from the previous cavity and the frequency are set as free parameters (with the exception of the input cavity which has no associated drift length). An analytical equation exists for determining optimal output coupling (Q_e^7) which yields good results and is used in this paper [14]. No such relation exists for the input coupling (Q_e^1) which is left as a free parameter. The R/Q, Q_0 and gap lengths (and so gap coupling factors) are kept constant throughout the optimization. This helps simplify the decision space to expedite convergence and simplify analysis.

The optimization is therefore performed with 14 free parameters set between broad limits. All other klystron parameters remain constant throughout the optimization process. The decision vector is therefore $\mathbf{x} = [f_1, f_2, \ldots, f_n, d_2, \ldots, d_n, Q_e^1]$ (where f_n is the frequency of the n^{th} cavity and d_n the drift length before it). The objectives used are efficiency (%), interaction length (m) and Δ_{gain} (dB); so $\mathbf{z} = [\eta, \ell, \Delta_{\text{gain}}]$. Slowest electron velocity is used to reject simulations where electrons are reflected. It would be trivial to expand the solution space to target additional tube properties, for instance gain. However, this would complicate the solution space for both optimization and analysis. This paper focuses on the simpler case for clarity.

The optimization is run for 10 000 evaluations with a population of 50. The 14 parents are selected using a tournament of 20 and the mother is selected using a tournament of 40. These settings are used throughout this paper.

IV. ANALYSIS

A. PSPF Tuning

Many of the parameters involved in the definition of the EA are investigated in other work. However, the precise position of the cutoff point for PSPF is not well known as it is most probably problem specific. The cutoff point is defined here as a fraction of the maximum efficiency at the current iteration, and is therefore not constant.

If this parameter is set too low there will be too great a number of klystrons with very low efficiencies. If it is set too high then the algorithm is at risk of becoming stuck in local optimums. A single run is used for each PSPF fraction in Fig. 5, to investigate this using a statistically significant

TABLE II

COMPARISON OF THE REFERENCE DESIGN [14] AND THE MOST EFFICIENT Ξ_η, THE SHORTEST Ξ_ℓ (WITHIN SPECIFICATION), AND A STRUCTURE WITH THE REFERENCE EFFICIENCY $\Xi_=$

	f_1	f_2	f_3	f_4	f_5	f_6	f_7
Ref. (MHz)	**478.00**	**473.20**	**482.30**	**938.80**	**499.00**	**491.50**	**475.86**
Ξ_η (MHz)	473.23	487.19	478.70	934.32	491.08	488.58	475.00
Δ (MHz)	-4.77	13.99	-3.60	-4.48	-7.92	-2.92	-0.85
$\Xi_=$ (MHz)	473.19	487.02	478.50	934.38	491.10	489.52	474.95
Δ (MHz)	-4.81	13.82	-3.80	-4.42	-7.90	-1.98	-0.90
Ξ_ℓ (MHz)	472.60	487.67	479.55	934.54	491.53	486.79	474.86
Δ (MHz)	-5.40	14.47	-2.75	-4.26	-7.47	-4.71	-0.99

(a) Cavity frequencies

	Q_e^1
Ref.	**176**
Ξ_η	215
Δ (%)	22
$\Xi_=$	215
Δ (%)	22
Ξ_ℓ	216
Δ (%)	23

(b) Input Q_e

	d_2	d_3	d_4	d_5	d_6	d_7
Ref. (mm)	**445**	**495**	**333**	**546**	**610**	**254**
Ξ_η (mm)	309	279	446	259	679	252
Δ (%)	-30.6	-43.7	34.0	-52.5	11.3	-0.6
$\Xi_=$ (mm)	298	262	428	277	630	256
Δ (%)	-32.9	-47.2	28.6	-49.4	3.4	0.7
Ξ_ℓ (mm)	318	276	252	276	477	254
Δ (%)	-28.4	-44.2	-24.4	-49.5	-21.8	0.1

(c) Drift lengths

	η	ℓ	Δ_{gain}
Ref.	**70.0 %**	**2.682 m**	**-5.01 dB**
Ξ_η	71.9 %	2.224 m	-2.00 dB
Δ	2.7 %	-17.1 %	-
$\Xi_=$	70.9 %	2.150 m	-2.00 dB
Δ	1.3 %	-19.8 %	-
Ξ_ℓ	60.1 %	1.853 m	-2.00 dB
Δ	-14.2 %	-30.9 %	-

(d) Output parameters

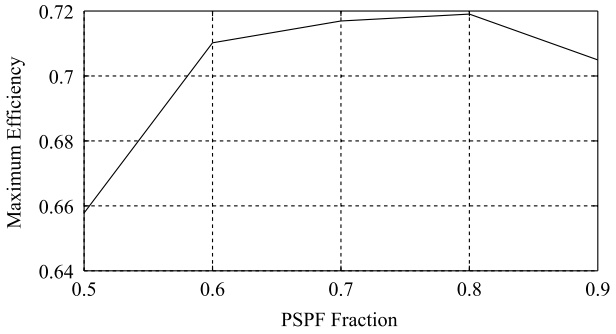

Fig. 5. Effect of changing the PSPF fraction on the maximum efficiency after 10 000 iterations.

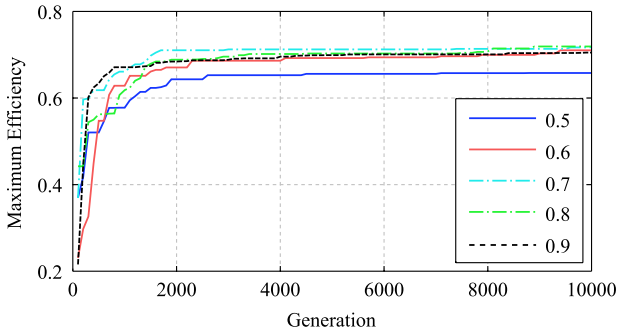

Fig. 6. Effect of changing the PSPF factor on the evolution of the maximum efficiency.

sample of klystron optimizations would be prohibitively time consuming.

From Fig. 5 given the small size of the sample, it is difficult to distinguish any advantages between a PSPF cutoff fraction of 0.6–0.9. There is an indication that a value of 0.7–0.8 will be most successful, although given the probabilistic nature of the algorithm this could be misleading. It should be noted that

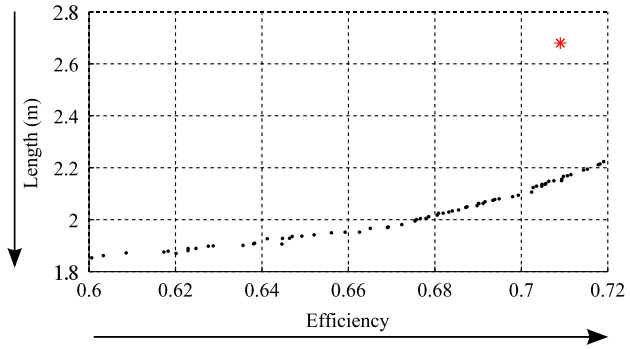

Fig. 7. Tradeoff between klystron length and achievable efficiency. Arrows: optimization direction. Star: reference design [14].

0.7 is more successful until around 7000 evaluations (Fig 6). A factor of 0.8 is chosen for the remaining work.

B. Optimization

After 10 000 generations all of the structures in the final Pareto front have the target bandwidth as PSPF and the saturation of Δ_{gain} ensure this is the case. Although Table II(d) shows that the reference design does not meet the bandwidth specification.

A curve showing the tradeoff between length and efficiency (Fig. 7), is as expected found to show that if a longer klystron can be tolerated then a more efficient structure can be made. The B-factory Klystron is plotted as a red star. This structure is clearly dominated by the solutions found automatically, although notably a significantly more efficient structure is not found.

As the optimum structure depends on the tradeoff between efficiency and length, three structures are selected for comparison. The most efficient Ξ_η, the shortest Ξ_ℓ (within specification) and a structure closest to the efficiency of the reference structure $\Xi_=$.

Table II(a) and II(c) show that the final cavity d_7 and f_7 differ by less than 1% and 1 MHz, respectively, from the reference design in Ξ_η, Ξ_ℓ and $\Xi_=$. This is not unexpected as Q_e^7 is fixed and the cavity is tuned to take energy out of a well-formed bunch. It would have been unlikely that the reference design was not well optimized in this respect, hence dramatic improvement is not expected.

d_2, d_3 and d_5 of the optimum structures are substantially shorter than the reference design leading to a significantly shortened overall structure for each. Only d_4 and d_6 are longer and only d_4 (the drift before the 2nd harmonic cavity) significantly so in Ξ_η and $\Xi_=$. d_4 and d_6, which range between 44.2% to 34.0% and 21.8% to 11.3%, respectively, for Ξ_ℓ and Ξ_η, have the least impact on bunching as it is these drifts which have the smallest impact on efficiency.

The slightly lower input power and slightly increased efficiency lead to an improved gain of 45.4 dB compared with 44.8 dB in the reference design. A higher Q_e^1 of 214 compared with 176 appears to be preferable for a shorter tube.

V. Conclusion

The design of a klystron interaction structure can be automated with the use of an EA. The problems specific to klystron optimization, such as the conflict between very short inefficient klystrons and longer efficient klystrons, were effectively solved by using the additional techniques suggested, such as PSPF. For narrow band tubes the approximation of bandwidth to Δ_{gain} was successful.

If a reduction in length was desired, the reduction of drift four and six was optimal (cavity four being the 2nd harmonic cavity). In the context of a comparison of 1-D simulations, an alternative design was found (although not experimentally verified) for the B-factory klystron, which met the specification and was shorter than the published length by 19.8%.

References

[1] H. Braun, R. Corsini, J. P. Delahaye, A. de Roeck, S. Döbert, A. Ferrari, G. Geschonke, A. Grudiev, C. Hauviller, B. Jeanneret, E. Jensen, T. Lefèvre, Y. Papaphilippou, G. Riddone, L. Rinolfi, W. D. Schlatter, H. Schmickler, D. Schulte, I. Syratchev, M. Taborelli, F. Tecker, R. Tomás, S. Weisz, and W. Wuensch, "CLIC 2008 parameters," CERN, Meyrin, Switzerland, Tech. Rep. CLIC-Note-764, Oct. 2008.

[2] E. Jensen and I. Syratchev, "CLIC 50 MW L-B and multi-beam klystron," in *Proc. 7th Int. High Energy Density High Power RF Workshop*, Jun. 2005, p. 11.

[3] S. Peggs, "Conceptual design report," ESS, Tunavägen, Sweden, Tech. Rep. ESS-2012-001, Feb 2012.

[4] W. R. Fowkes, G. Caryotakis, E. Doyle, E. Jongewaard, C. Pearson, R. Phillips, J. Sackett, E. Wright, H. Bohlen, G. Huffman, S. Lenci, E. Lien, E. McCune, and G. Miram, "1.2 MW klystron for asymmetric storage ring B factory," in *Proc. PAC*, May 1995, pp. 1–3.

[5] D. E. Goldberg, *Genetic Algorithms in Search, Optimization and Machine Learning*. Reading, MA, USA: Addison-Wesley, 1989.

[6] T. Blickle and L. Thiele, "A comparison of selection schemes used in evolutionary algorithms," *Evol. Comput.*, vol. 4, no. 4, pp. 361–394, 1996.

[7] K. Deb, *Multi-Objective Optimization Using Evolutionary Algorithms*. New York, NY, USA: Wiley, 2001.

[8] K. Deb, "A computationally efficient evolutionary algorithm for real-parameter optimization," *Evol. Comput.*, vol. 10, no. 4, pp. 371–395, Nov. 2002.

[9] E. Zitzler and L. Thiele, "Multiobjective evolutionary algorithms: A comparative case study and the strength Pareto approach," *IEEE J. Evol. Comput.*, vol. 3, no. 4, pp. 257–271, Nov. 1999.

[10] K. J. Gunn, "Generic optimisation of wave energy converter performance," Ph.D. dissertation, Dept. Eng., Lancaster Univ., Lancaster, U.K., 2012.

[11] C. J. Lingwood, "High power high efficiency klystron design," Ph.D. dissertation, Dept. Eng., Lancaster University, Lancaster, U.K., 2010.

[12] N. Srinivas and K. Deb, "Multiobjective optimization using nondominated sorting in genetic algorithms," *J. Evol. Comput.*, vol. 2, no. 3, pp. 221–248, 1994.

[13] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan, "A fast and elitist multiobjective genetic algorithm: NSGA-II," *IEEE Trans. Evol. Comput.*, vol. 6, no. 2, pp. 182–197, Apr. 2002.

[14] G. Caryotakis, R. Phillips, G. Scheitrum, E. Jongewaard, S. Gold, and D. Smithe. (2004). *Klystron Department Lecture Series* [Online]. Available: http://www-group.slac.stanford.edu/kly/Lecture_Series/slac_klystron_lecture_series.htm

[15] E. Demmel, S. Lutgert, B. Offermann, and H. Schwarz, "A new high power klystron for the SLAC B factory," in *Proc. EPAC 1996*, pp. 2152–2154.

[16] A. Jensen, *AJDisk* [Online]. Available: http://www-group.slac.stanford.edu/kly/Lecture_Series/slac_klystron_lecture_series.htm

[17] T. Shintake, "FCI-Field charge interaction program for high power klystron simulations," in *Proc. PAC 1989*, vol. 1, pp. 94–96.

[18] M. Litzkow, M. Livny, and M. Mutka, "Condor—A hunter of idle workstations," in *Proc. 8th Int. Conf. Distrib. Comput. Syst.*, Jun. 1988, pp. 104–111.

Authors' photographs and biographies not available at the time of publication.