

# Design Optimisation of Coilguns

G.Hainsworth and D.Rodger

University of Bath, BATH BA2 7AY, UK.

**Abstract**— In recent years, many current filament coilgun models have been interfaced to optimisation routines. This has provided a valuable insight into the operation of coilguns, and has generally yielded kinetic energies superior to those obtained purely by analytical methods. Due to concern over the presence of local optima, the use of a global optimisation scheme, Simulated Annealing, has been investigated. In its basic form, Simulated Annealing requires a vast number of simulation runs, and is therefore impractical for coilgun optimisation. However, a variation on the algorithm, Very Fast Simulated Reannealing (VFSR) converges much faster and has produced good results.

## I. INTRODUCTION

Computer models are very useful for designing coilguns. Usually these models have consisted of inductance matrix methods based on current filaments, although recently Finite Element modelling has been used [1]. Without optimisation, design would usually proceed as follows. By making certain assumptions, the required time constants for each gun stage can be estimated. This will give an indication of the size and shape of the coils. Calculations of rise times will give some approximation to the point at which to trigger each stage. Certain combinations can then be tested with the simulation model, but this soon becomes tedious. Low efficiencies tend to result, mainly because it is difficult to find good trigger positions due to the interaction between gun stages. The situation becomes harder for guns with just a few stages, since everything is very transient, and all of the stages behave differently.

The design process is made considerably easier and yields much better results when an optimiser is used. This iteratively calls the simulation model to find the best design [2]. The optimisation technique normally chosen is some form of steepest descent algorithm. This requires an initial design, and will find the nearest (local) optimum, where the derivatives with respect to each parameter are zero. Since it does not guarantee finding the global optimum, it is normal to begin the search from the best

design found analytically. Further runs may be necessary from different start points, to confirm that the result is reasonable.

There are however, global optimisation schemes, such as Simulated Annealing, which if used correctly, do guarantee finding the global optimum. This has the advantage that only one run is required, and no starting design is necessary. Unfortunately, most forms of global optimisation require many times the number of simulations needed by a local optimiser, making the run time excessive. A variation on Simulated Annealing is discussed which reduces this time, and with other factors taken into account (such as the need for multiple runs from different start points), can produce results in a comparable time.

The computer code developed consists of three main blocks:

1. An initialisation routine, which instructs the optimiser over which parameters it has control, and the ranges in which they must lie. It also sets the value of those parameters that are not to be included in the optimisation.
2. The cost function, which uses a current filament model to simulate the gun performance for a given set of parameters. The results of the simulation are used to generate the cost.
3. The optimiser, which iteratively calls the cost function with sets of parameters, attempting to minimise the cost.

These are discussed in the following sections.

## II. COILGUN PARAMETERS

The parameters fall into two categories, those that are fixed, and those that are available for optimisation. The fixed parameters are as follows:

- Number of stages (drive coils)
- Projectile and drive coil material data
- Payload mass
- Projectile and drive coil initial temperatures
- Projectile and drive coil packing factors

Manuscript received 13 April 1994.

This work is funded by DRA under agreement no. CB/FRN/9/4/2112/095/MA

In some applications, it may be desirable to optimise some of the above parameters, but under normal circumstances they would be considered fixed. The parameters available for optimisation are:

- Drive coil inner radius (bore)
- Projectile coil inner and outer radius, and length

and for each stage:

- Drive coil outer radius and length
- Drive coil wire cross sectional area
- Coil trigger positions
- Distance between coils
- Power supply polarity, initial voltage and capacitance.

These parameters have been chosen to simplify the optimisation process as much as possible. For example, instead of specifying the coil length and gap between coils, it would have been possible to use absolute start and end positions for each coil. However, this would make optimisation harder, since the resulting parameters are less independent and changes in one cause more complex effects.

The trigger positions are also made relative, by specifying the distance from the back of the projectile to the centre of the coil, as a proportion of the length of that coil. This helps since it maintains a sensible trigger position as the optimiser changes the length of the coils.

The optimiser does not have direct control of the number of turns on the coils, instead this is calculated using the coil size and packing factor. This helps produce coils that are practical to build.

The software requires an initial value, and an allowable range for each parameter. The initial value is of no real significance, it just acts as a start point for the random generator. A flag is used to indicate whether to include each parameter in the optimisation. It is also possible to fix parameters to be equal – for example all of the drive coils could have the same length.

### III. COST FUNCTION

Each set of parameters produced by the optimiser is simulated using a current filament coilgun model. A flag is passed back to the optimiser to indicate whether any complex constraint has been violated, for example, acceleration or stresses exceeding their permissible limit. The results of the simulation run are used to generate the cost. This may be the negative of the kinetic energy (since the optimiser minimises), or may be a more complex function

involving stage efficiencies. Optimising for velocity tends to result in trivial solutions (projectile mass being reduced to zero), unless the payload mass is specified.

Obviously if the cost function is  $-E_k$  there is no point in optimising the power supply values, since these would go to their upper limit.

The current filament software used has been verified against the Finite Element code, Mega [3].

### IV. SIMULATED ANNEALING

Simulated Annealing (SA) operates by generating a random set of parameters (a state) and evaluating the cost function at that point. The resulting cost is used to decide whether to accept or reject that state. If it is accepted, it becomes the 'saved' state. A new state is then generated and the process repeated. The significance of SA is in the methods used to generate and govern acceptance of these states.

#### A. State generation

Standard SA uses a Gaussian distribution, centred on the saved state to generate new states. Initially, the saved state is chosen randomly.

#### B. State acceptance

If the new state is better than the saved one (ie. lower cost), it is accepted. Obviously, if this were the only criteria, the system would quickly find a local minimum and stay there. There must be some mechanism for 'hill climbing' to enable escape from local minima so that the global one can be found. Therefore, if the new state has a higher cost than the saved one, it may still be accepted. This is governed by the difference in the two costs,  $\Delta E$ .

$$\text{Acceptance Probability } h(\Delta E) = \exp\left(\frac{-\Delta E}{T}\right) \quad (1)$$

Here,  $T$  is referred to as the system 'Temperature' and is used to control the annealing process. When the temperature is very high,  $h(\Delta E) \approx 1$ , resulting in all states being accepted, causing the system to jump around randomly. When the temperature is very low,  $h(\Delta E) \approx 0$  causing all bad states to be rejected. The annealing process requires the system to start at a sufficiently high temperature, and then cooled in a controlled fashion. The cooling schedule used is:

$$T(t) = \frac{T_0}{\ln t} \quad (2)$$

where  $T$  is the current temperature,  $T_0$  is the initial temperature, and  $t$  is a measure of time based on the number of states accepted. It can be shown [4] that cooling at this rate will statistically guarantee finding the global minimum. Cooling faster than (2) results in a form of simulated quenching, which does not guarantee finding the global minimum.

## V. IMPROVEMENTS ON SIMULATED ANNEALING

Standard SA would not be suitable for optimisation of coilguns, since it requires a vast number of cost function evaluations. This can be seen by examining the cooling schedule (2). One improvement is to use a wider distribution for generating the states. This makes it easier for the system to jump out of local minima, thus permitting faster cooling. Fast annealing (FA) [5] uses a Cauchy distribution and requires a cooling schedule of:

$$T(t) = \frac{T_0}{t} \quad (3)$$

a significant improvement over (2).

However, the optimisation scheme chosen is a recent development, Very Fast Simulated Reannealing (VFSR) [6]. This has many advantages over SA and FA. VFSR takes account of the differing sensitivities of the parameters in a multi-dimensional system, by allocating a temperature to each. This temperature controls the width of the generating probability distribution. It is cooled in a similar way to the cost function acceptance temperature, except that the measure of time used is based on the number of states generated, rather than the number accepted. A special probability distribution is used for generation, which, with the other changes enables an exponential cooling schedule to be used:

$$T(t) = T_0 \exp(-t) \quad (4)$$

Periodically, reannealing takes place, which involves adjustment of all of the temperatures, based on the derivatives of the cost function with respect to each of the parameters. A more detailed explanation of VFSR is given in [4].

## VI. RESULTS

To evaluate the suitability of VFSR to coilgun optimisation, its performance has been compared with a linear optimiser. Both are given the task of optimising a four stage gun for maximum kinetic energy, with control over the following eight parameters:

DOR	Drive coil outer radius	(same for all 4 coils)
DLN	Drive coil length	(same for all 4 coils)
TR1-4	Trigger position for each stage	
WCS	Wire cross sectional area	(same for all 4 coils)
PLN	Projectile coil length	

The remaining parameters are fixed, the important ones being:

Drive coil inner radius:	27mm
Projectile inner radius:	19mm
Projectile outer radius:	25mm
Gap between coils:	20mm
Power supply initial energy:	100kJ
Projectile type:	Copper tube

Three runs were performed with each optimiser. VFSR was run with differing start temperatures: 0, 2.5 and 5kJ (temperature has the same units as the cost function). The linear optimiser was run from different start points, shown in figure 1. There is no need to do this for VFSR, since the start point is chosen randomly.

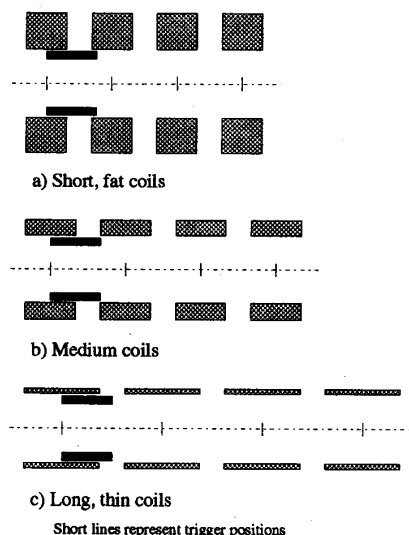


Fig. 1. Start points for linear optimisations

The results of all six runs are given in tables 1 and 2. From this it can be noted that the three linear optimisations have found different local minima, each with kinetic energies of around 25.4kJ. The zero temperature VFSR run has also found a local optimum, as expected. However, the other two VFSR runs have found the same optimum, 27kJ, indicating a reasonable possibility that this may be global, and that the start temperature of 2.5kJ is sufficient

	VFSR 0J	VFSR 2.5kJ	VFSR 5kJ	UNITS
DOR	36.5	35.6	35.7	mm
DLN	30.6	30.7	30.8	mm
TR1	-0.01	0.01	0.00	--
TR2	-0.86	-0.83	-0.83	--
TR3	-1.36	-1.28	-1.31	--
TR4	-1.69	-1.58	-1.61	--
WCS	8.53	8.15	8.12	mm <sup>2</sup>
PLN	54.6	51.4	51.9	mm
KE	26.9	27.0	27.0	kJ

Table 1: Results for three VFSR optimisations

	LIN a)	LIN b)	LIN c)	UNITS
DOR	39.0	31.0	31.6	mm
DLN	31.0	32.8	38.7	mm
TR1	0.03	0.06	0.05	--
TR2	-0.71	-0.59	-0.58	--
TR3	-1.09	-0.96	-0.89	--
TR4	-1.37	-1.15	-1.11	--
WCS	15.31	4.83	5.66	mm <sup>2</sup>
PLN	45.2	43.1	48.7	mm
KE	25.4	25.3	25.4	kJ

Table 2: Results for three linear optimisations

in this case. The zero temperature run has actually found a very good local optimum at 26.9kJ. The cooling effect and different initial temperatures can be clearly seen in figure 2.




Fig. 2. Kinetic energy vs. number accepted

It is unfortunately not possible to confirm that an optimum is global, it is only possible to prove that it is not global, by finding a better one. The results can be checked to ensure that they are local optima, by sweeping each pa-

rameter in turn over its permissible range. Figure 3 shows this for the high temperature VFSR results, whereas figure 4 is for the linear result (starting from short fat coils). The  $x$ -axis shows the deviation from optimum, the units being cm for DOR, DLN and PLN, and  $\times 10 \text{ mm}^2$  for WCS. The reason for the jagged effect on kinetic energy as WCS changes is due to the turns calculation. As the wire size increases, the number of turns on the coils decreases in integer steps, causing corresponding steps in kinetic energy.




Fig. 3. Effect on KE with variation of each parameter - VFSR




Fig. 4. Effect on KE with variation of each parameter - LINEAR

## VII. CONCLUSIONS

Very Fast Simulated Reannealing has successfully been used to optimise coilguns, attaining efficiencies in the 15–50% range. Run times are comparable with those of linear optimisations, and good results are consistently obtained. Further speed increases are possible by increasing the filament size, since VFSR is less sensitive to ‘noise’ on the simulation results.

The need for multiple runs from differing start points is eliminated, since the global optimum is found. A reasonable starting design is not required, making it very simple to run optimisations for different systems. Complex constraints are easily handled and heating or mechanical stresses can be included in the cost function to produce more practical designs.

In the simple example illustrated here, the linear method has performed fairly well, in that the local optima are quite close to the global optimum. This is by no means always true, experience has shown that the linear method can often stop at poor local optima. This is a serious disadvantage in design work.

## REFERENCES

- [1] P.J.Leonard, H.C.Lai, G.Hainsworth, D.Rodger and J.F.Eastham, ‘Analysis of the performance of Tubular Pulsed Coil Induction Launchers’, IEEE Transactions on Magnetics, vol. 29, no. 1, Jan 1993, p.686–690.
- [2] D.A.Bresie, J.A.Andrews and S.W.Ingram, ‘Parametric approach to Linear Induction Accelerator design’, IEEE Transactions on Magnetics, vol. 27, no. 1, Jan 1991, p.390–393.
- [3] G.Hainsworth, P.J.Leonard, D.Rodger and H.C.Lai, ‘A comparison between Finite Element and Mutual Inductance models of Coilguns’, Proceedings of the 4th European symposium on EML technology, 1993.
- [4] L.Ingber and B.Rosen, ‘Genetic algorithms and Very Fast Simulated Reannealing: A comparison’, Mathematical and Computer Modelling, 16, 1992, p.87–100.
- [5] H.Szu and R.Hartley, ‘Fast Simulated Annealing’, Physics Letters A, vol. 122, no. 4, June 1987, p.157–162.
- [6] L.Ingber and B.Rosen, ‘Very Fast Simulated Reannealing (VFSR)’, [ftp research.att.com: /netlib/opt/vfsr.Z], AT&T Bell Labs, Murray Hill, NJ (1992).