Laird, I., Scoltock, J., Forsyth, A., & Yuan, X. (2016). A Unified Framework for Computationally Efficient Power Converter Design Optimisation. In *8th IET International Conference on Power Electronics, Machines and Drives (PEMD 2016)* [cp2016.0133] Institution of Engineering and Technology (IET).
<https://doi.org/10.1049/cp.2016.0133>

Peer reviewed version

Link to published version (if available):
[10.1049/cp.2016.0133](https://doi.org/10.1049/cp.2016.0133)

[Link to publication record in Explore Bristol Research](#)
PDF-document

This is the author accepted manuscript (AAM). The final published version (version of record) is available online via IET at <http://digital-library.theiet.org/content/conferences/10.1049/cp.2016.0133>

University of Bristol - Explore Bristol Research

General rights

This document is made available in accordance with publisher policies. Please cite only the published version using the reference above. Full terms of use are available:
<http://www.bristol.ac.uk/red/research-policy/pure/user-guides/ebr-terms/>

A Unified Framework for Computationally Efficient Power Converter Design Optimisation

I D Laird, J Scoltock †, A J Forsyth †, X Yuan**

**University of Bristol, UK, ian.laird@bristol.ac.uk, †University of Manchester, UK, james.scoltock@manchester.ac.uk*

Keywords: Design optimisation, Unified framework, Computational efficiency, Silicon carbide, Power density

Abstract: Within the field of power converter design, the use of optimisation techniques is a popular method to ensure that the best possible design is achieved for a given objective, e.g. weight or volume reduction. However, there is currently a lack of efficient unified frameworks for optimising the system level design of converters in terms of the device, heatsink and passive component selection. This paper presents an efficient unified design optimisation framework that can be applied to a wide range of power converter design problems. The main advantage of the proposed framework is that it produces optimal solutions with great computational efficiency. The proposed framework was validated by applying to SiC-based converter design problems, achieving power densities of 8.32 kW/kg for a 50-kW DC-DC converter, and 3.50 kW/L for a 5-kW 3-phase DC-AC converter.

1 Introduction

Traditionally the design of a power converter involves an engineer using standard design equations to select each of the system's components on a part-by-part basis. Whilst straightforward, this tends to result in sub-optimal designs, as the interactions between the values of different components are not fully considered. In order to achieve fully optimised designs, and thus meet the performance targets, a holistic design optimisation method is needed that accounts for all the interactions between the various components and/or design variables (e.g. switching frequency) within the system [1]. While topologies and specifications vary greatly from one converter to another, the design variable interactions tend to be very similar. For example, the switching frequency has a direct impact on the cooling requirement of the system, as it is the main factor in determining the switching losses of the semiconductor devices. Similarly, it has a significant effect on the passive components, influencing the current and voltage ripples in a DC-DC converter, as well as the harmonics that an electromagnetic interference (EMI) filter will need to mitigate in a DC-AC converter.

These common interactions found between very different converter topologies means it is possible to create a unified design framework that streamlines the design optimisation process for a range of different topologies, specifications and applications. The benefit of a common design platform, such as this, is that it would unify different design methods and allow component models and databases to be shared. It would

also make comparing various competing topologies much easier, such as determining whether a 2-level, 3-level or multi-level system inverter results in the highest power density, or how many interleaved phases are required to produce the optimal design. In this paper, we present a unifying framework for optimising power converter designs. Although a number of previous works have dealt with the issue of power converter optimisation, e.g. [2–5], there appears to be an absence of a generalised framework which is suitable for the system-level optimisation of a wide range of power converters. The main advantage of the proposed framework is that it can exploit independencies between the design variables which makes it very computationally efficient at finding optimal solutions.

Recently, SiC-based devices have gained significant interest due to the higher efficiency, higher switching frequencies, and more compact designs that they produce relative to established Si-based devices [6,7]. Due to the higher cost of SiC devices, however, there is a greater incentive to ensure that the best possible designs are achieved. In order to evaluate the proposed optimisation framework, it was applied to the design of a 50-kW SiC-based DC-DC converter and a 3-phase 5-kW SiC-based DC-AC converter, where the objectives were to minimise weight and volume, respectively.

Section 2 outlines the optimisation framework and explains how the computational efficiency can be improved. Section 3 shows how the framework was applied to the optimisation of a 50-kW DC-DC converter and a 3-phase 5-kW DC-AC converter, the latter of which is supported with experimental validation. Finally, conclusions are given in section 4.

2 Overview of optimisation framework

2.1 Basic concept

The proposed framework approaches the task of power converter design as a combinatorial optimisation problem. User inputs to the algorithm include the design's specifications (e.g. input and output voltage, power level or required topology), constraints (e.g. maximum voltage or current ripple) and objective (e.g. minimisation of price, weight or volume). Additionally, the user can also define the design variables, and their potential values, which the algorithm combines to form every possible design within the solution space. Each design variable, examples of which include the semiconductor device(s), inductor core(s), capacitor(s), topology, heatsink(s), modulation strategy and switching frequency, is restricted to a discrete set of values, referred to as a 'value set'.

The most basic approach is to solve the problem via full enumeration, as shown in Figure 1. Each design in the solution space, denoted by the index $s \in S$, where S is an index set and $|S|$ is the number of possible designs, is assessed for feasibility at the specified operating point(s). If it is feasible, then the cost, $J(s)$, based on the design objective, is computed and compared with the cost value of the most optimal design found so far. If it is better, then the current design and its cost are saved for comparison with subsequent designs. The process is repeated until each design within the solution space has been evaluated.

Figure 1: Design optimisation flowchart with full enumeration of possible designs

When full enumeration is employed, the solution space can be visualised as shown in Figure 2(a), where each variable (here labelled a , b , c and d) occupies a level of a nested loop structure. It is obvious that the complexity of the problem grows quickly as the number of design variables, and/or the size of the value sets, increase. For even a modest design problem, with, for instance, 4 design variables, each of which can take one of ten values, there are $10^4 = 10,000$ possible designs.

Although full enumeration guarantees that the best design within the solution space will be found, evaluating every possible design can be time consuming, and will usually involve checking a large number of solutions that are not feasible. For large problems, enumerating and evaluating each of these designs can be very costly from a computational standpoint, and for larger designs can result in the problem becoming intractable. In order to address this issue, one should exploit the independencies between design variables, or groups of variables, wherever possible. By doing so, the size of the solution space, and therefore execution time, can be greatly reduced without changing the number of design variables and the size of their corresponding value sets.

2.2 Strategies for Computational Efficiency

It is necessary to use full enumeration, and hence the nested loop structure shown in Figure 2(a), when the cost function is of the form

$$J = f(a, b, c, d) \quad (1)$$

i.e., when the cost function cannot be separated into multiple sub-functions, or one or more of the constraints are of the form

$$g(a, b, c, d) \leq G \quad (2)$$

i.e., when the constrained quantity (or quantities) depends on all of the design variables. Here, G is a constraint which, although not shown explicitly in Equation (2), can in practise depend on one or more of the design variables. An example of this is a constraint on capacitor current ripple, which will vary in accordance with the type of capacitor that is used.

Figure 2(b) illustrates a basic reduction of the solution space shown in Figure 2(a). For a problem to be formulated with this structure, it requires the cost function to be separable such that all of the sub-functions have at least one common argument, and where none of the sub-functions take all of the design variables as arguments. For the case illustrated in Figure 2(b), the cost function will be of the form

$$J = f_1(a, b) + f_2(a, c, d) \quad (3)$$

Similarly, the constraints, denoted by the indices $1 - n$, need to be expressible in a similar form, i.e.

$$\begin{aligned} g_1(a, b) &\leq G_1, \dots, g_m(a, b) \leq G_m \\ g_{m+1}(a, c, d) &\leq G_{m+1}, \dots, g_n(a, c, d) \leq G_n \end{aligned} \quad (4)$$

where the first m constrained quantities depend only on the variables that are arguments to f_1 , and the next $(n - m)$ constraints depend only on the variables that are arguments to f_2 . Within such a structure, the variable(s) that is common to the sub-functions within the cost function and each of the constraints (being the variable a in this example) are referred to as the ‘parent’ variable(s), while the others are referred to as

Figure 2: Different solution space structures. In (a), there are no independencies between the variables. In (b), there are independencies between groups of ‘child’ variables underneath a group of ‘parent’ variables. In (c), there are independencies between groups of variables, with dependencies within each group.

‘child’ variables. For the example shown in Figure 2(b), the total number of sub-designs that are evaluated is reduced to $10(10 + 10^2) = 1,100$, an 89% reduction compared to the 10,000 designs that are evaluated with full enumeration.

In other instances, it may be possible for the problem to be split into fully decoupled sub-problems. This is illustrated in Figure 2(c), where the sub-problems are shown as having the structure of Figure 2(a). Sub-problems with different internal structures are also conceivable. For a problem to be formulated with this structure, the cost function needs to be separable into sub-functions that do not share any arguments. For the case illustrated in Figure 2(c), the cost function will be of the form

$$J = f_1(a, b) + f_2(c, d) \quad (5)$$

As with the previously described structure, the constraints need to be expressible in a similar form to the cost function, i.e.

$$\begin{aligned} g_1(a, b) &\leq G_1, \dots, g_m(a, b) \leq G_m \\ g_m + 1(c, d) &\leq G_m + 1, \dots, g_n(c, d) \leq G_n \end{aligned} \quad (6)$$

For the example shown in Figure 2(c), the total number of designs is reduced to $10^2 + 10^2 = 200$ designs, a 98% reduction compared to full enumeration.

Although structures which exploit independencies can significantly reduce the number of designs, and therefore improve the execution time, it is difficult to automatically identify these structures. Additionally, it may be necessary to simplify the expressions within the cost function and/or constraints, in order to create independencies between the variables. This will be further discussed in the next section in relation to the DC-DC and DC-AC converter case studies.

3 Case studies: DC-DC and DC-AC

The optimisation strategy was applied to two different case studies – a 50 kW DC-DC converter and a 5 kW three-phase DC-AC converter – in order to verify its performance. The objective function was formulated to minimise the mass/weight of the DC-DC converter and the volume of the

DC-AC converter. Both the DC-DC and DC-AC converter optimisation programs were implemented in MATLAB.

3.1 DC-DC case study

The DC-DC case study specifications are for a 200 – 600 V, 50 kW, SiC-based converter. The topology under consideration is the Interleaved Boost Converter (IBC) with decoupled phases [8]. Table 1 provides the basic information relating to the design variables, which are all restricted to finite sets. [9] provides details of a previously built converter which is constructed from a similar range of components. It is assumed that the semiconductor devices and inductors are conduction-cooled via a liquid-cooled heat-sink. The inductors are foil-wound with amorphous metal cores, and are potted in aluminium cans. Film capacitors are used at the output. The heat sink model is based on a liquid-cooled concept with an inlet temperature of 60°C. Constrained quantities are summarised in Table 2. Note that ‘design dependent’ refers to the fact that the constraints on capacitor current ripple depend on the type and number of parallel-connected capacitors.

For the chosen set of design variables, full enumeration, as illustrated in Figure 2(a), results in 997,920 designs, which when executed on a desktop PC equipped with a 3.4 GHz Intel i7-3770 CPU and 16.0 GB of RAM results in an execution time of 304.0 minutes. A summary of the resulting design is provided in Table 3. With a weight of 6.01 kg, the optimised design has a power density of 8.32 kW/kg. The numerical value of the optimal output capacitance is 14.0 μ F, and the numerical value of the inductance (per phase) is 23.7 μ H. Note that the heat sink is sized on the assumption that the inductors and modules are mounted to it, and its weight is estimated accordingly.

As mentioned in the previous section, it is possible to significantly reduce the execution time by exploiting independencies between design variables. For the DC-DC converter under consideration, this is achieved using a version of the structure that is shown in Figure 2(c). The implementation relies on the fact that, with the exception of the output voltage ripple and output capacitor current ripple, the selection of the output capacitors has a negligible impact on

any of the constrained quantities. As a result, the constraints can be split into two - the first group being those that are negligibly impacted by the capacitor selection, and which involve all variables except the capacitors, and the second being those that relate to the capacitors, and which are set according to the optimal values of the first set of variables. Moreover, since the cost function is simply a measure of the weight, in which the weights of the different components are summed, the cost function is naturally of the form given in Equation (5).

With the efficient implementation strategy, the number of sub-designs that are evaluated is reduced to 18,199, or 1.8% of the number of designs that were evaluated using full enumeration. In terms of execution time, the efficient implementation arrives at an optimal design in 5.8 minutes, which is a 98% reduction relative to full enumeration. It should be noted that the optimal design that is achieved with the efficient implementation is the same as with full enumeration, which validates the accuracy of the approach.

Design variable	Range of values
Number of phases	1, 2,
Modules per phase	1, 2, 3
Type of module	CREE CAS100H12AM1, CAS300M12BM2 (2 total)
Switching frequency	60, 64, ..., 140 kHz
Output capacitor type	EPCOS B3277 (11 in total)
Number of output capacitors	1, 2, ..., 5
Inductor core type	FINEMET F3CC series (12 total)
Turns per inductor	10, 12, ..., 20

Table 1: DC-DC converter case study - value sets for major design variables

Quantity	Constraint
Efficiency	95 %
Output voltage ripple	2 %
Conduction mode	Continuous
Output capacitor current ripple	(Design dependent)
Junction temperature	150 °C
Core temperature	150 °C
Winding temperature	150 °C

Table 2: DC-DC converter case study - list of constrained quantities

Design variable	Optimal value
Number of phases	2
Modules per phase	1
Type of module	CAS300M12BM2
Switching frequency	108 kHz
Output capacitor type	B32776E8146+000
Number of output capacitors	1
Inductor core type	F3CC0125
Turns per inductor	10

Table 3: DC-DC converter case study - optimised values for major design variables

Figure 3: Weight breakdown of optimised DC-DC converter (Total mass = 6.01 kg)

3.2 DC-AC case study

The specifications for the DC-AC converter are for a 3-phase 330 V (line-to-line), 400 Hz output at 5 kW from a DC link of 600 V where the design is optimised for volume, rather than weight as in the DC-DC case study. The topology chosen for optimisation is the standard 6-switch, 2-level structure controlled by sinusoidal pulse width modulation (SPWM). The line and EMI filters connected to the output of the switching block are arranged as an LCLCL filter network on each phase where L_{AI} functions as the line filter. The following CL component blocks combine with L_{AI} to form the differential mode (DM) and common mode (CM) EMI filters respectively, as shown in Figure 4. The design variables for this case study along with their range of values are shown in Table 4.

Figure 4: Topology used in DC-AC case study

The key constrained properties for the design are shown in Table 5. Whilst the algorithm is able to consider fan-forced cooling, the optimisation process was limited to natural convection heatsink designs. Details of these constraints are shown in Table 5.

Full enumeration of the design variables results in more than 2.2×10^{12} designs. This is an intractable problem, so in order to achieve an optimised design with reasonable computational efficiency, a structure similar to that shown in Figure 2(b) was implemented. With the new structure the switching frequency remains the design variable common to all stages of the design (i.e. the 'parent' variable) while the devices and heatsinks are selected independently from the passives, each of these representing a separate group of 'child' variables. The design of the passives remains a nested structure with the line, DM and CM filters all dependent on each other. By adopting this

structure and eliminating switching frequencies that will not produce valid designs, the number of designs that require evaluation is reduced to 224,822, which is a reduction of over 99.9%.

A summary of the results of the optimisation algorithm can be seen in Table 6. A volume breakdown of the optimised design is shown in Figure 5 where the ‘DM’ and ‘CM’ volume percentages refer to the additional L and C components that are added to, but not including, the line filter inductor in order to form the full DM and CM filters respectively. As can be seen the majority of the volume is taken up by the line filter which has to be sized to meet the ripple current requirement. As a result the following DM and CM filters can be made much smaller and still met the EMI requirements. With a total volume of 1.427 L, this results in a power density of 3.50 kW/L.

Design variable	Range of values
Device type	Cree C2M MOSFETs (5 in total)
Switching frequency	10, 11, ... , 300 kHz
Line and DM inductor core type	Ferroxcube gapped double ETD cores (7 in total)
CM inductor core type	TDK/EPCOS & Ferroxcube toroids (44 in total)
DM capacitor type	Kemet & Vishay X1 class film (21 in total)
CM capacitor type	Kemet, TDK & Vishay Y2 class film (68 in total)
DC link capacitor type	AVX, Kemet & Vishay DC film (36 in total)
Heatsink types	Aavid thermalloy (14 in total)

Table 4: DC-AC converter case study value sets for major design variables

Quantity	Constraint
Efficiency	98 %
Output phase current ripple	10 %
DC link voltage ripple	0.5 %
Device junction temperature	125 °C
EMI limit standard	DO-160E category L, M and H

Table 5: DC-AC converter case study list of constrained quantities

Design variable	Optimal value
Device type	C2M0040120D
Switch frequency	63 kHz
Line inductor	ETD59/31/22, 71 turns, 2.4 mm gap
DM inductor	ETD29/16/10, 18 turns, 0.6 mm gap
CM inductor	TX36/23/15-3E5, 9 turns
DM capacitor	474R3220(1)A1(2), 8 in parallel
CM capacitor	B32024A3224M, 2 in parallel
DC link capacitor	MKP1848 530 094K2
Heatsink	000EK*, 40.15 mm

Table 6: DC-AC converter case study list of optimised values of design variables

Figure 5: Volume breakdown of optimised DC-AC converter (Total volume = 1.427 L)

For this case study, the optimised design given in Table 6 was experimentally validated. The control was provided by a TI TMS320F2812PGFA DSP. Due to controller limitations the converter was operated at 64 kHz instead of the optimal value of 63 kHz. The measured voltage, current and power levels are shown in Table 7. The simulation framework predicted that the efficiency at 64 kHz would be 98.17% which is very close to that of the measured efficiency of 97.37%.

Heatsink and device temperatures were measured using a Flir E30 thermal camera, the results of which are shown in Table 7. The simulation framework predicted a heatsink temperature of 99.63 °C and a device junction temperature of 104.8 °C. One reason that partially explains the disparity between the results is the fact that the thermal contact resistance between the devices and the heatsink was not accounted for in the model. Nevertheless, the temperature results are within the safe operating limits of the devices and thus the error is acceptable.

Parameter	Value	Parameter	Value
DC-link voltage	599.76 V	Output phase voltage	197.65 V
Input current	7.494 A	Output phase current	7.499 A
Input power	4.490 kW	Output power	4.3766 kW
Input power factor	0.999	Output power factor	0.943
Heatsink temperature	113 °C	Ambient temperature	28 °C
Device case temperature	137 °C	Efficiency	97.37 %

Table 7: Measured experimental results from optimised DC-AC converter

The measured output waveforms of a single phase are shown in Figure 6. A small degree of distortion can be observed which can be analysed by looking at the DM and CM EMI frequency spectrums as shown in Figure 7 and Figure 8 respectively. In both frequency spectrums, harmonics are observed at multiples of the fundamental, switching and resonant frequencies of the filters. The figures show that the EMI filters meet the DO-160E standard requirements which begins at 150 kHz. The distortion in Figure 6 is due to the lower order harmonics that fall outside the DO-160E standard and hence weren't optimised for.

Figure 6: Experimental results for phase A output voltage and current waveforms

Figure 7: Experimental results for the DM noise frequency spectrum with DO-160E EMI standard

Figure 8: Experimental results for the CM noise frequency spectrum with DO-160E EMI standard

4. Conclusion

This paper has presented an efficient unified design optimisation framework that can be applied to a range of converter topologies. It has discussed the framework's operation and how it can be structured to achieve better

computational efficiency. Two case studies were presented, a 50 kW DC-DC converter and a 5 kW DC-AC converter, that showed the implementation and performance of the design optimisation framework. The framework achieved power densities of 8.32 kW/kg for the 50-kW SiC-based DC-DC converter, and 3.50 kW/L for the 5-kW SiC-based 3-phase DC-AC converter. Experimental results were also given for the DC-AC converter case study, verifying that the design proposed by the framework was able to meet all the required specifications.

Acknowledgements

The authors would like to thank the UK EPSRC National Centre for Power Electronics under Grant EP/K035096/1 and EP/K035304/1 for supporting the related research.

References

- [1] J.W. Kolar, et al., "PWM converter power density barriers", *Power Conversion Conference, PCC'07*, Nagoya, Japan, 2 – 5 April, 2007, pp. 9-29.
- [2] S. Balachandran and F.C.Y. Lee, "Algorithms for power converter design optimisation", *IEEE Transactions on Aerospace and Electronic Systems*, vol. AES-15, no. 3, pp. 422 - 432, May 1981.
- [3] R. Lai, et al., "A systematic topology evaluation methodology for high-density three-phase PWM AC-AC converters", *IEEE Transactions on Power Electronics*, vol. 23, no. 6, pp. 2665-2680, Nov. 2008.
- [4] J. Biela, U. Badstuebner, and J.W. Kolar, "Design of a 5-kW, 1-U, 10-kW/dm² resonant DC-DC converter for telecom applications", *IEEE Transactions on Power Electronics*, vol. 24, no. 7, pp. 1701-1710, Jul. 2009.
- [5] I. Laird, X. Yuan and N. McNeill, "A holistic approach to optimise the power density of a silicon carbide (SiC) MOSFET based three-phase inverter", *Power Electronics and Drive Systems, PEDS'15*, Sydney, Australia, 9 – 12 June, 2015, pp. 473-478.
- [6] J. Millan, et al., "A survey of wide bandgap power semiconductor devices", *IEEE Transactions on Power Electronics*, vol. 29, no. 5, pp. 2155-2163, May 2014.
- [7] J. Biela et al., "SiC versus Si - evaluation of potentials for performance improvement of inverter and DC-DC converter systems by SiC power semiconductors", *IEEE Transactions on Industrial Electronics*, vol. 28, no. 7, pp. 2872-2882, Jul. 2011.
- [8] H.-B. Shin et al., "Generalised steady-state analysis of multiphase interleaved boost converter with coupled inductors", *IEE Proceedings on Electric Power Applications*, vol. 152, no. 3, pp 584 – 594, May 2005.
- [9] G. Calderon-Lopez and A.J. Forsyth, "High power density DC-DC converter with SiC MOSFETs for electric vehicles", *IET International Conference on Power Electronics, Machines and Drives, PEMD 2014*, Manchester, U.K., 8 – 10 April, 2014, pp. 1 – 6.