OPEN ACCESS

University of BRISTOL

Simpson, N., & Mellor, P. (2015). Exploiting Cloud Computing in the Multi-Physics Design and Optimisation of Electromagnetic Devices. In IEEE 16th Workshop on Control and Modeling for Power Electronics (COMPEL) 2015. IEEE.

Peer reviewed version

Link to publication record in Explore Bristol Research
PDF-document

University of Bristol - Explore Bristol Research

General rights

Exploiting Cloud Computing in the Multi-Physics Design and Optimisation of Electromagnetic Devices

N. Simpson and P. H. Mellor
Department of Electrical and Electronic Engineering
University of Bristol
Bristol, UK
nick.simpson@bristol.ac.uk

Abstract—**Cloud computing is a type of distributed computing that enables public, on-demand network (internet) access to a large pool of shared computing resources, making it easier than ever to outsource computing tasks on a pay-as-you-go basis. In this paper, a cloud computing service is used to complement an existing high-fidelity multi-physics power inductor design workflow by enabling multiple instances of the design analysis to be executed in parallel on virtual computers in a cloud computing environment. This method accelerates the design development cycle and can allow a more thorough evaluation of the design space within a short time-frame, at a moderate cost and can potentially lead to improved designs. The modified workflow is demonstrated by the design optimisation of a high-energy-density filter inductor for an automotive application.**

I. INTRODUCTION

In order to meet the ever increasing demand for higher power-density, energy-efficiency and lower-cost of electrical machines and devices, multiple coupled physical phenomena must be accounted for at the design stage, [1]. As a minimum the design synthesis should capture the close interrelation between the electromagnetic and thermal behaviour, [2]. This would ideally be achieved using high fidelity electromagnetic and thermal models which are coupled through a loss estimation. The bulk of the losses are generally temperature dependent requiring the electromagnetic and thermal models to be iteratively evaluated, [3]. This significantly increases computation time and tends to prohibit the use of such models, particularly for iterative design optimisation. The problem is often addressed by employing computationally efficient modelling methods, [4], or by compromising the detail and accuracy of the models, [2], [5]–[7], in favour of short solution times. Thereby, enabling a design optimisation to be performed using commonplace desktop computing hardware, [4], [7]. An alternative is to adopt a distributed computing approach, [8], where high-fidelity models are evaluated using many networked computers in parallel to reduce solution times. This approach has been adopted in areas such as Electromagnetic Compatibility (EMC), [9], electronic chip-package design, [10] and photonics, [11]. It has also been adopted in the electromagnetic and later coupled electromagnetic and thermal design optimisation of a Permanent Magnet (PM) machine,

[12], [13]. However, other efforts are not widely reported in the area of electrical machine and wound passive component design, [12], despite the idea being known for a number of years, [14], and tools being available, [15], [16]. This is perhaps due to the perceived complexity of integrating the use of such distributed computing systems in existing workflows.

Cloud computing is a type of distributed computing that enables public, on-demand network (internet) access to a large pool of shared computing resources, [8], making it easier than ever to outsource computing tasks on a pay-as-you-go basis. In this paper, the use of cloud computing to complement an existing multi-physics power inductor design workflow is presented, [17]. The implementation of the computationally efficient power inductor model is described along with its role in the current workflow. Limitations of the present single desktop computer arrangement, such as reaching computational feasibility limits and excessive design times, are alleviated by using a simple algorithm to execute multiple instances of the multi-physics analysis on virtual computers in a cloud computing environment. The modified workflow is demonstrated by the design optimisation of a high-energy-density filter inductor for an automotive application. An indication of the economic costs of using the cloud based service is given along with a discussion of advantages and future improvements to the implementation.

II. DESCRIPTION OF THE EXISTING WORKFLOW

In the existing workflow, a parametrically defined, computationally efficient multi-physics model of an E-core power inductor is used to manually assess the design parameter space for a given specification. This allows infeasible solutions to be ruled out based on previous experience and a wider understanding of the application. Parametric studies are then performed using a simple algorithm to systematically vary a parameter of interest and assessing the impact on performance. For example, the effect of rectangular strip or edge wound conductors on ac losses can be studied. At this point further infeasible solutions are identified and discounted. A design optimisation is then performed using the resulting reduced set of design parameters and an appropriate algorithm, Fig. 1.
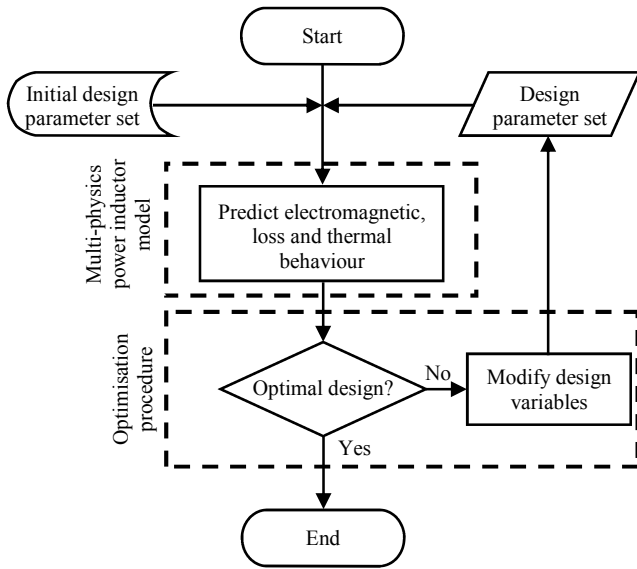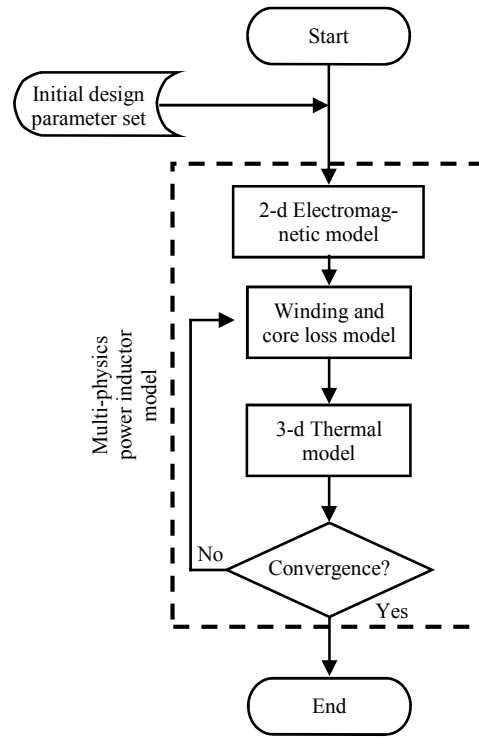
Fig. 1. Optimisation workflow.

Fig. 2. Structure of the multi-physics power inductor model.

– – Line of symmetry

Fig. 3. Example of the 2-D electromagnetic FE model with 9 turns in 3 layers.

A. Multi-Physics Power Inductor Model

The multi-physics inductor model is composed of individual electromagnetic and thermal Finite Element (FE) models coupled by winding and core loss estimations, as illustrated in Fig. 2, [17]. The models are fully parametrically defined and appear as a "black-box" problem which takes a set of design parameters, D_p, containing geometric, material, boundary condition and specification data as input and returns a high-fidelity approximation of the electromagnetic, loss and thermal behaviour of the inductor as output.

1) Electromagnetic Model: A parametric time-harmonic non-linear 2-D FE model is used to predict the electromagnetic behaviour and losses, Fig. 3, [18]. Due to symmetry, only one half of the inductor cross-section is modelled to reduce computation time. The 2-D nature of the model prohibits transposed conductors from being directly catered for. The contribution of the end-winding to inductance and ac loss is assumed negligible and is confirmed by 3-D FE analysis in a post-design validation step. The inductance is estimated as the ratio of change in flux linkage, Ψ, to change in winding current, I, (1) which requires two FE analyses to be performed at a winding current of $I \pm \delta I$.

$$L = \frac{d\Psi}{dI} \tag{1}$$

2) Winding Loss: The temperature and frequency dependent ac winding loss is determined from the ratio of effective ac to dc resistance, R_{ac}/R_{dc}, [19], [20]. The operating temperature of the inductor is emulated by scaling the resistivity, ρ, of the winding material with the well established linearised equation, (2), where T_{op}, T_0 and α are the operating temperature, reference temperature and temperature coefficient of resistance respectively.

$$\rho|_{T_{op}} = \rho|_{T_0} \left(1 + \alpha \left(T_{op} - T_0 \right) \right) \tag{2}$$

$$\frac{R_{ac}}{R_{dc}}\bigg|_{T_{op}} = \frac{\left(\frac{R_{ac}}{R_{dc}}\bigg|_{T_0} - 1 \right)}{\left(1 + \alpha \left(T_{op} - T_0 \right) \right)^{\beta+1}} + 1 \tag{3}$$

$$\beta = \frac{\log \left(\frac{R_{ac}}{R_{dc}}\bigg|_{T_0} - 1 \right) - \log \left(\frac{R_{ac}}{R_{dc}}\bigg|_{T_{max}} - 1 \right)}{\log \left(1 + \alpha \left(T_{max} - T_0 \right) \right)} - 1 \tag{4}$$

The R_{ac}/R_{dc} ratio is estimated at T_{op} for a fixed frequency using an analytical loss function, (3) and (4), informed by two R_{ac}/R_{dc} values obtained from the electromagnetic FE analysis at the reference temperature and the maximum ex-

pected operating temperature, T_{max}, respectively, where β is a curve fitting term, [3]. Hence, two FE analyses are required to determine the ac winding loss over the expected operating temperature range at each frequency of interest.

3) Core Loss: The core loss arising from hysteresis, eddy current and excess loss effects is estimated using the Bertotti iron loss model, (5), [21], where f and B are the frequency and peak flux density respectively. The coefficients k_h, k_e and k_a are derived from measured loss data provided by the electrical steel manufacturer. Only one FE analysis is required at each frequency of interest to estimate the iron loss.

$$W_{Fe} = k_h f B^{\alpha_{fe}} + k_e f^2 B^2 + k_a f^{1.5} B^{1.5} \tag{5}$$

B. Thermal Model

Fig. 4. Example of the 3-D thermal FE model.

A parametric steady-state 3-D FE model is used to predict the thermal behaviour, Fig. 4. The inductor is assumed to be potted and housed within a case mounted to a cold-plate at a fixed temperature. Hence, the model is reduced to a pure conduction problem with boundary conditions. The heat extraction from the cold plate is assumed to dominate over convection and radiation effects. An adiabatic boundary condition is applied to each surface with the exception of the base where a fixed temperature boundary condition is applied to emulate the cold-plate. The model incorporates slot liner and calibrated interface gaps in order to better represent a real world device and improve temperature predictions, [1], [17], [22]. Due to symmetry only one quarter of the inductor is modelled. The winding region represents a composite material composed of conductor, electrical insulation and encapsulant insulation materials. In order to simplify the thermal model definition and reduce computation time, the winding region is represented as a lumped homogeneous material with equivalent anisotropic thermal properties, [22], [23]. This lumped region is split into subregions representing each winding conductor, Figs. 3 and 4, and a unique loss expression informed by the electromagnetic FE analyses, Section II-A2, is applied as a local loss source. This enables the temperature variation resulting from non-uniform loss distribution to be accounted for resulting in an improved prediction of the hot-spot temperature and location.

The core loss, Section II-A3, is applied as a global loss source within the core region. Only one execution of the thermal FE model is required, the iterative update of the loss with temperature is performed internally by the modelling software, [24].

C. Parallel Model Execution

Table I details the set of electromagnetic FE analyses, along with the appropriate winding current, frequency and temperate cases, required to characterise the electromagnetic and loss behaviour over the temperature range of interest at a fixed frequency, f_0, Sections II-A1 to II-A3. The electromagnetic model assumes pure sinusoidal excitation current, however, non-sinusoidal current can be approximated by decomposing the waveform into its Fourier components and evaluating the FE model at each significant harmonic along with the dc average and superimposing the result, [20], [25]. Hence, the number of FE analyses required, N_{EM}, to characterise the electromagnetic and loss behaviour is given by, (6), where N_h is the number of significant harmonics.

$$N_{EM} = 2 + 2N_h \tag{6}$$

TABLE I
ELECTROMAGNETIC FE ANALYSES

Characteristic	Winding Current	Frequency	Temperature
Inductance and winding loss (dc)	$I_{pk} + \delta I$	0	T_0
Inductance and winding loss (dc)	$I_{pk} - \delta I$	0	T_0
Winding and Core loss	I_{rms}	f_0	T_0
Winding loss	I_{rms}	f_0	T_{max}

The electromagnetic FE analyses are independent and as such can be performed in parallel. A simple queueing algorithm is used to control the timing of the model executions as illustrated in Fig. 5. Initially all of the required FE models are generated and one model is executed for each available processor core, N_{core}, typically in the range 1–8 for a desktop computer. When a model is complete, the results are written to a file and the next model is executed. The process repeats until all N_{EM} FE analyses are complete. This ensures a high utilisation of modern multi-core processors and minimises the computation time of the electromagnetic and loss modelling stage. The time required to predict the electromagnetic, loss and temperature behaviour of an inductor with sinusoidal current excitation is in the range 60 - 90 seconds using $N_{core} = 4$ processor cores. The computation time is difficult to approximate as it is dependent on many factors including the size and mesh density of the models as well as their non-linear behaviour.

D. Limitations of the Existing Workflow

Despite the computationally efficient implementation of the inductor model, Section II-A, the number of model analyses
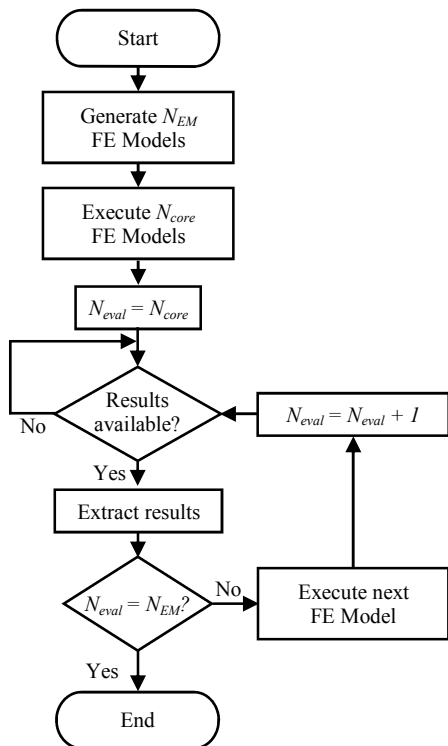
Fig. 5. Queueing algorithm used to control parallel execution of FE models.

Configuration Number	Number of CPUs	Memory (RAM)	Memory (Storage)	Price per Hour [USD]
1	4	7.5	User specified	0.43
2	8	15.0	User specified	0.86
3	36	60.0	User specified	3.44

computing services are available, however the Amazon Elastic Compute Cloud (EC2) is used in this work, [26]. The algorithm begins by requesting and starting N_{comp} identical preconfigured virtual computers referred to as compute instances, Fig. 6. At this point the compute time is charged on a per-hour basis for each instance until they are terminated, Table II, [26]. The inductor model, associated parametric study or optimisation algorithm and data, such as parameter ranges, material properties and specification details are uploaded to each compute instance. The design tool is executed on the N_{comp} compute instances simultaneously. The algorithm polls for results, if found, they are downloaded and the next study is executed on the now idle compute instance. This process repeats until all of the studies or optimisations, N_{EXEC}, are complete. When a compute instance becomes idle and there are no more executions required it is terminated.

IV. FILTER INDUCTOR DESIGN CASE STUDY

The benefits of incorporating additional computing resources into the workflow are illustrated by the deign optimisation of a high-energy-density filter inductor intended to be integrated into the power electronic stage and water cooling circuit of an automotive power converter. The full specification is given in Table III. In order to reduce material costs and simplify manufacture, aluminium conductors arranged in a single layer strip winding are assumed. The six design parameters, D_p, to be optimised are given in Table IV.

Parameter	Value	Unit
Current rating, I_{rms}	200	A
Operating frequency, f_{op}	400	Hz
Inductance, L_t	80	μH
Maximum operating temperature, T_t	180	^{o}C
Target energy density (complete assembly), E_{dt}	1.2	J/kg

required by the parametric studies and optimisation lead to a typical computation time of between 5 and 12 hours. These studies are performed sequentially by necessity and can therefore lead to excessive design times. However, if multiple parametric studies and optimisations could be performed simultaneously, using a simple to implement approach, then the design time can be significantly reduced. This could enable a more thorough evaluation of the design space to be performed and lead to further design improvements.

III. MODIFIED WORKFLOW - EXPLOITING CLOUD COMPUTING

Cloud computing, in the context of this work, provides on-demand, pay-as-you-go, network (internet) access to a large pool of shared computing resources, [8]. Virtualisation software allows users to request the creation of virtual computers with specific software and hardware configurations which run on the underlying physical hardware, Table II, [26]. In essence, the existing workflow can be complemented by enabling parametric studies and optimisations to be executed on many virtual computers in parallel on a pay-per-hour basis.

A. Executing Parametric Studies and Optimisations in a Cloud Computing Environment

An adaptation of the queueing algorithm, Section II-C and Fig. 5, and an Application Programming Interface (API) are used to interact with the cloud computing service in a simple manner and execute the desired parametric studies or optimisations, as illustrated in Fig. 6. Numerous cloud

A. Optimisation Routine

The inductor specification is described in terms of a target inductance, L_t, peak operating temperature, T_t, and energy density, E_{dt}. The energy density is defined as the energy stored in the magnetic field per unit mass of the complete inductor assembly, including the case, where L, M, I_{pk} and I_{rms} are the inductance, mass and peak and rms winding current respectively, (7).
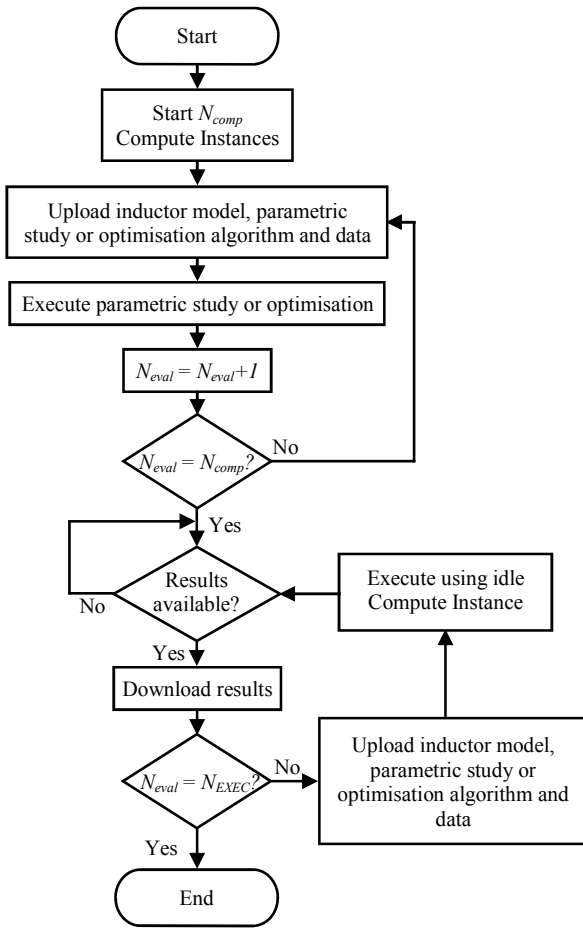
Fig. 6. Simplified queueing algorithm used to execute the design tool in a cloud computing environment.

TABLE IV
INDUCTOR DESIGN PARAMETERS

Parameter	Value Range	Unit
Core thickness[*], T_{core}	$1 \leq T_{core} \leq 50$	mm
Number of turns, N_t	$1 \leq N_t \leq 75$	n/a
Air-gap length (normalised), L_{ag}	$0.01 \leq L_{ag} \leq 0.99$	n/a
Active length[*], L_{act}	$1 \leq L_{act} \leq 300$	mm
Conductor current density, J_c	$1 \leq J_c \leq 40$	A/mm^2
Aspect ratio[*], AR	$0.1 \leq AR \leq 10$	n/a

[*] refer to Figs. 3 and 4

$$E_d = \frac{LI_{pk}^2}{4M} = \frac{LI_{rms}^2}{M} \tag{7}$$

$$E(D_p) = \left| \frac{L_t - L_c(D_p)}{L_t} \right| \\ + \left| \frac{E_{dt} - E_{dc}(D_p)}{E_{dt}} \right| \\ + \left| \frac{T_t - T_c(D_p)}{T_t} \right| \tag{8}$$

The multiple design objectives are treated as a single global objective by formulating the objective function, (8), as a sum of three normalised error terms representing the deviation of a candidate design from the inductance, L_t, energy density, E_{dt}, and peak operating temperature, T_t, targets, respectively Table III. A Particle Swarm Optimisation (PSO) routine is employed to systematically vary the design parameters, D_p, within predetermined bounds and minimise the objective function, $E(D_p)$, (8). The objective function is multi-valued, meaning that a number of design parameter sets, D_p, can equally satisfy the design specification. Therefore, the design tool is used to generate a set of 150 design candidates and the results studied.

B. Optimised Filter Inductor Design

Figs. 7 to 9 show the resulting inductance, energy density and maximum operating temperature of the 150 design candidates identified by the design tool. Each design candidate is within 5 %, 4.6 % and 2.7 % of the inductance, energy density and temperature targets and could be considered equally optimal from this limited data. However, when the predicted total loss is plotted against the predicted mass for each design, Fig. 10, an approximation to the Pareto front is identified. Designs which lie on this line are non-dominated meaning that an improvement in one characteristic, mass or loss, cannot be realised without a corresponding compromise in the other, [27]. Six design candidates were identified using the existing single desktop workflow, Section II and Fig. 10, in approximately the same time it took to identify the 150 design candidates using the cloud computing service. The additional information afforded by the high number of optimisation executions allows a design to be selected, as shown in Figs. 7 to 10, which exhibits the best compromise between mass and loss and satisfies secondary considerations which are beneficial to the inductor design but are not expressly specified. The resulting inductor design is within 5% ($76 \mu H$), 1.7% ($177 °C$) and 1.7% (1.18) of the inductance, temperature and energy density targets respectively, the optimised design parameters are given in Table V, [17].

TABLE V
OPTIMISED INDUCTOR DESIGN PARAMETERS

Parameter	Value	Unit
Core thickness, T_{core}	10.7	mm
Number of turns, N_t	5	n/a
Air-gap length, L_{ag}	1.26	mm
Active length, L_{act}	133.6	mm
Conductor current density, J_c	17.6	A/mm^2
Aspect ratio, AR	2.24	n/a

C. Computation Time and Cost

The computation time required to perform the multi-physics analysis for one PSO particle, one PSO generation and one optimised design are given in Table VI along with the approximate cost in USD. Hardware configuration 1 featuring 4 processor cores and 7.5 GB of RAM was used for each compute
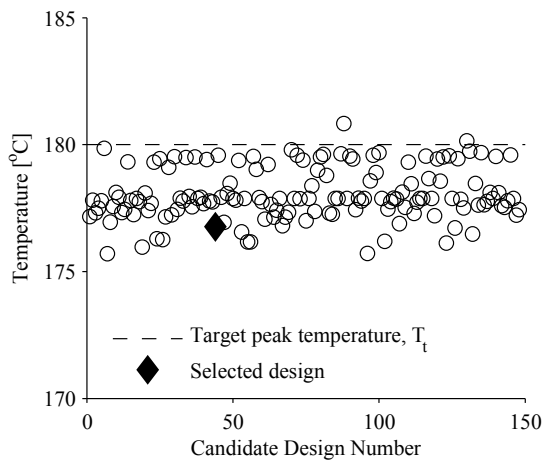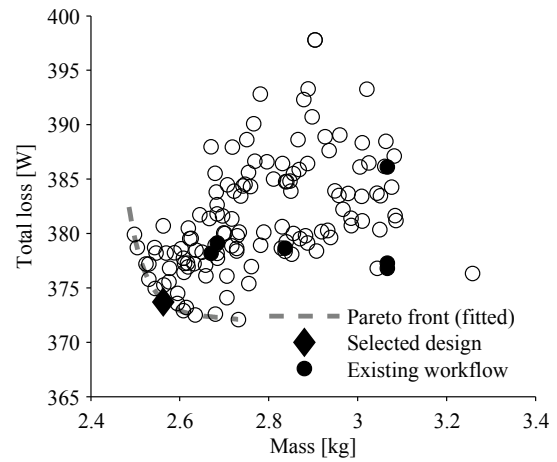
Fig. 7. Inductance

Fig. 10. Loss as a function of mass

during each design tool evaluation, N_{DTE}, and averaging over the number of particles and generations. The timing data does not include the overheads incurred during the initialisation of the compute instances which are dependent on a number of factors including the desired hardware configuration, Table II, and resource availability. However, in this case they were relatively small compared to the overall computation time. The 150 design tool evaluations were completed in approximately 68 real-time hours translating to approximately 1688 compute hours at a cost of approximately 725 USD.

Fig. 8. Energy density

TABLE VI
AVERAGE COMPUTATION TIME AND COST

Evaluation Type	Computation* Time	Computation* Cost [USD]
One particle	90 seconds	-
One generation (36 particles)	13.50 minutes	0.10
One optimised design (50 generations)	11.25 hours	4.84

*assuming hardware configuration 1, Table II

D. Computation Time and Cost Reduction

The generation of a Pareto optimal set using the PSO algorithm is contingent on the stochastic nature of the PSO and performing a sufficiently large set of optimisations, [27], [28]. Hence this method is computationally inefficient and tends to result in a low quality approximation to the Pareto front, where the Pareto points are bunched close together, [27]. A Pareto set can be identified more efficiently using alternative optimisation techniques such as Differential Evolution, [4], or Non-dominated Sorting Genetic Algorithm II (NSGA-II), [29] and can reduce both time and economic cost. Implementing such a change in optimisation approach would require the inductor model, Fig. 2, to be executed in the cloud, controlled by a local optimisation routine which would add significant complexity to the implementation. This

Fig. 9. Temperature

instance, Table II. The timing information was collected by recording the start and end time of each multi-physics analysis

arrangement is adopted in [12], however, the numerical models are executed on a local distributed computing environment using a specialised compute job and resource management system, [15], [30].

V. Conclusion

In this paper, an existing power inductor design workflow, [17], is extended to enable parametric studies and design optimisation to be executed in a publicly accessible, pay-as-you-go cloud computing environment. The extension is demonstrated by the design optimisation of a filter inductor for an automotive application, Section IV. The optimisation is performed 150 times using 25 compute instances in parallel in order to generate a set of design candidates which meet the primary design drivers of inductance, energy density and maximum operating temperature, Table III. The mass and loss of each candidate design are compared, enabling a Pareto front to be identified, Fig. 10. This information is used to select a design that not only meets the inductance, energy density and temperature specification within 5 %, 1.7 % and 1.7 % respectively but enables the design with the best compromise between mass and loss to be selected.

The 150 design candidates were identified in approximately 68 hours, approximately 25 times faster than the existing workflow, at a cost of approximately 725 USD. Alternative optimisation methods could identify a Pareto front in a more robust and efficient manner, [4], [27], [29]. Adding compute resources using the simple algorithm presented here, Section III, accelerates the design development cycle and can allow a more thorough evaluation of the design space within a short time-frame, at a moderate cost and can potentially lead to improved designs. Multiple designs for different applications can be considered simultaneously and the optimisation procedures and modelling methods can be experimented with more quickly, enabling the design tools to be further developed and refined.

Acknowledgement

References

[1] D. Staton, A. Boglietti, and A. Cavagnino, "Solving the more difficult aspects of electric motor thermal analysis in small and medium size industrial induction motors," *Energy Conversion, IEEE Transactions on*, vol. 20, no. 3, pp. 620–628, Sept 2005.

[2] N. Bracikowski, M. Hecquet, P. Brochet, and S. Shirinskii, "Multiphysics modeling of a permanent magnet synchronous machine by using lumped models," *Industrial Electronics, IEEE Transactions on*, vol. 59, no. 6, pp. 2426–2437, 2012.

[3] R. Wrobel, D. Salt, A. Griffo, N. Simpson, and P. Mellor, "Derivation and scaling of ac copper loss in thermal modeling of electrical machines," *Industrial Electronics, IEEE Transactions on*, vol. 61, no. 8, pp. 4412–4420, Aug 2014.

[4] G. Sizov, P. Zhang, D. Ionel, N. Demerdash, and M. Rosu, "Automated multi-objective design optimization of pm ac machines using computationally efficient fea and differential evolution," *Industry Applications, IEEE Transactions on*, vol. 49, no. 5, pp. 2086–2096, Sept 2013.

[5] A. Sarikhani and O. Mohammed, "A multi-physics multi-objective optimal design approach of pm synchronous machines," in *Electrical Machines (ICEM), 2014 International Conference on*, Sept 2014, pp. 968–974.

[6] I. Vese, F. Marignetti, and M. Radulescu, "Multiphysics approach to numerical modeling of a permanent-magnet tubular linear motor," *Industrial Electronics, IEEE Transactions on*, vol. 57, no. 1, pp. 320–326, 2010.

[7] N. Simpson, R. Wrobel, and P. Mellor, "A multi-physics design methodology applied to a high-force-density short-duty linear actuator," in *Energy Conversion Congress and Exposition (ECCE), 2014 IEEE*, Sept 2014, pp. 5168–5175.

[8] N. Sadashiv and S. Kumar, "Cluster, grid and cloud computing: A detailed comparison," in *Computer Science Education (ICCSE), 2011 6th International Conference on*, Aug 2011, pp. 477–482.

[9] H. Li and J. Lu, "Simulation and optimization of electromagnetic compatibility problems using computational electromagnetics and high performance computing: A cloud computing based system," in *Electromagnetic Field Problems and Applications (ICEF), 2012 Sixth International Conference on*, June 2012, pp. 1–4.

[10] D. Gope, V. Jandhyala, X. Wang, D. Macmillen, R. Camposano, S. Chakraborty, J. Pingenot, and D. Williams, "Towards system-level electromagnetic field simulation on computing clouds," in *Electrical Performance of Electronic Packaging and Systems (EPEPS), 2011 IEEE 20th Conference on*, Oct 2011, pp. 167–170.

[11] M. Molinari, K. Thomas, and S. Cox, "Electromagnetic design search and optimisation of photonic bandgap devices on distributed computational resources," in *Proceedings of the Fifth IEE International Conference on Computation in Electromagnetics*. The Institution of Engineering and Technology, April 2004, pp. 103–104. [Online]. Available: http://eprints.soton.ac.uk/45810/

[12] W. Jiang, T. Jahns, T. Lipo, W. Taylor, and Y. Suzuki, "Machine design optimization based on finite element analysis in a high-throughput computing environment," in *Energy Conversion Congress and Exposition (ECCE), 2012 IEEE*, Sept 2012, pp. 869–876.

[13] W. Jiang and T. Jahns, "Development of efficient electromagnetic-thermal coupled model of electric machines based on finite element analysis," in *Electric Machines Drives Conference (IEMDC), 2013 IEEE International*, 2013, pp. 816–823.

[14] C. Chan and K. Chau, "Design of electrical machines by the finite element method using distributed computing," *Computers in industry*, vol. 17, no. 4, pp. 367–374, 1991.

[15] D. Thain, T. Tannenbaum, and M. Livny, "Distributed computing in practice: The condor experience," *Concurrency-Practice and Experience*, vol. 17, no. 2-4, pp. 323–356, 2005.

[16] K. Lysiak and J. Polendo, "A generalized matlab-based distributed-computing optimization tool," in *Wireless Communications and Applied Computational Electromagnetics, 2005. IEEE/ACES International Conference on*, April 2005, pp. 170–173.

[17] N. Simpson, R. Wrobel, and P. Mellor, "Multi-physics design of high-energy-density wound components," in *Energy Conversion Congress and Exposition (ECCE), 2015 IEEE*, [In Press].

[18] (2015, July) Finite element method magnetics (femm). Internet. [Online]. Available: http://www.femm.info/wiki/HomePage

[19] C. Sullivan, "Computationally efficient winding loss calculation with multiple windings, arbitrary waveforms, and two-dimensional or three-dimensional field geometry," *Power Electronics, IEEE Transactions on*, vol. 16, no. 1, pp. 142–150, Jan 2001.

[20] M. Mirjafari and R. Balog, "Survey of modelling techniques used in optimisation of power electronic components," *Power Electronics, IET*, vol. 7, no. 5, pp. 1192–1203, May 2014.

[21] M. Popescu, D. Ionel, A. Boglietti, A. Cavagnino, C. Cossar, and M. McGilp, "A general model for estimating the laminated steel losses under pwm voltage supply," *Industry Applications, IEEE Transactions on*, vol. 46, no. 4, pp. 1389–1396, July 2010.

[22] A. Boglietti, A. Cavagnino, and D. Staton, "Determination of critical parameters in electrical machine thermal models," *Industry Applications, IEEE Transactions on*, vol. 44, no. 4, pp. 1150–1159, 2008.

[23] N. Simpson, R. Wrobel, and P. Mellor, "Estimation of equivalent thermal parameters of impregnated electrical windings," *Industry Applications, IEEE Transactions on*, vol. 49, no. 6, pp. 2505–2515, Nov 2013.

[24] (2015, July) Getdp: a general environment for the treatment of discrete problems. Internet. [Online]. Available: http://www.geuz.org/getdp/

[25] M. Gerber, J. Ferreira, I. Hofsajer, and N. Seliger, "A high-density heat-sink-mounted inductor for automotive applications," *Industry Applications, IEEE Transactions on*, vol. 40, no. 4, pp. 1031–1038, July 2004.

[26] (2015, March) Amazon ec2 pricing. Internet. [Online]. Available: http://aws.amazon.com/ec2/pricing/

[27] K. Deb, *Multi-objective optimization using evolutionary algorithms*. John Wiley & Sons, 2001, vol. 16.

[28] R. Poli, J. Kennedy, and T. Blackwell, "Particle swarm optimization," *Swarm Intelligence*, vol. 1, pp. 33–57, 2007.

[29] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan, "A fast and elitist multiobjective genetic algorithm: Nsga-ii," *Evolutionary Computation, IEEE Transactions on*, vol. 6, no. 2, pp. 182–197, Apr 2002.

[30] (2015, July) Htcondor. Internet. [Online]. Available: http://research.cs.wisc.edu/htcondor/