Optimization, sensitivity analysis, and robust design using response surface modeling

Egbert Lenderink*[a] and Peter Stehouwer[b]
[a]Philips Lighting, Advanced Development Lighting, Mathildelaan 1 (EEC-2), 5611 BD Eindhoven, the Netherlands; [b]Center for Quantitative Methods, P.O. Box 414, 5600 AK Eindhoven, the Netherlands

ABSTRACT

Response surface modeling (RSM) is described as a tool to perform optimizations and sensitivity analysis in optical modeling. With this method, the behavior of the system is first mapped out using a limited set of ray-tracing simulations, carefully spanning the full parameter space. This can already be done before the full merit function is known. The relation between design parameters and system performance is then approximated by fitting the simulation results to functional forms. All subsequent optimizations are then performed very time-efficiently on the functional descriptions of the dependencies. This contrasts with direct optimization, where the computationally intensive optical modeling is in the loop of the optimization algorithm, and where changes in merit function, mapping out trade-offs, and determining sensitivities, are very time consuming. The paper will discuss the advantages of RSM with respect to direct optimization and give recommendations for the type of problems that are preferentially addressed by RSM. The method will be illustrated by a case: how optical simulations were used in the design of Lumiramic™ phosphor conversion components for LEDs.

Keywords: response surface modeling, optimization, Lumiramic, ceramic color converter.

1. INTRODUCTION

The use of optimization methods, while already widely used in imaging optics, is now maturing also in the field of illumination optics. Most effort is directed towards search algorithms that find their way through the multidimensional design space as efficiently as possible and thus determine the global minimum of a certain merit function within the design space. Optical simulations are performed at each iteration to determine the value of this merit function and to obtain the set of design parameters for the next iteration.

This works well when the behavior of the system is well known (i.e. when a perfect computer model of the real system exists), when one optical simulation run can be performed in relatively short time, and when the merit function is well known: this means that all contributions to the merit function, as well as their relative weights, are not subject to discussion or major changes. When there is a change to the merit function, for instance by new feedback from customers, the optimization routine has to be run again, demanding a new set of ray trace simulations.

Moreover, the optimization method gives just one answer and does not provide insight in sensitivities and the influence of manufacturing tolerances. Sensitivity analysis is also very important when there is some uncertainty on the computer model. This is the case for instance when processes that are notoriously difficult to model, such as volume scattering, absorption and fluorescence, make up the majority of the optical rays that hit the target. Of course, it is always possible to set up a sensitivity analysis around the found optimum, but this is time consuming as it involves again more ray tracing simulations.

This paper presents an alternative approach called response surface modeling (RSM). This contrasts with the optimization methods described above (which we will call *direct optimization* in the following) in that the computationally intensive optical simulations are kept out of the loop of the search algorithm. Instead, one starts with a well-defined set of optical simulations that carefully spans the full design space range of interest. From these simulations, all relevant response parameters are determined and their response *surfaces* are determined: the relation

* egbert.lenderink@philips.com; phone +31 40 2757732; fax +31 40 2756564; www.lighting.philips.com

Illumination Optics, edited by Tina E. Kidger, Stuart R. David
Proc. of SPIE Vol. 7103, 710302 · © 2008 SPIE · CCC code: 0277-786X/08/$18 · doi: 10.1117/12.797649

between design parameters and response parameters is fit using functional forms. All subsequent analysis is then performed on these functional forms.

This strategy opens the way to types of analyses that are hardly possible or very time-consuming using direct optimization. Although the method starts with one long time-consuming run of optical simulations, all subsequent analyses are performed very time-efficiently. This makes it possible to quickly investigate trade-offs and sensitivities. It becomes possible to turn various knobs in the merit function and to see directly the effect on the found optimum. Not only does this increase the understanding that is gained on system performance, it also helps in setting the specifications in cases where the optimum trade-off is not known a priori. Finally, it brings sensitivities to sources of spread immediately out into the open. RSM fits very well in a Design for Six Sigma (DfSS) approach, where robustness of the design to spread in process and component parameters is already taken into account in a very early stage of product development.

The paper is built up as follows. In the next section, the method of RSM is described in further detail. Then, in section 3, it is shown how RSM is used in organizing the optical simulations in the earliest stages of the design of Lumiramic™ [1]: a novel phosphor technology, developed by Philips to minimize the spread in color points of white LEDs. The final section summarizes the findings, lists the main benefits of the RSM way of working and the type of problems that are most advantageously tackled by RSM.

2. RESPONSE SURFACE MODELING

The presented approach elaborates and extends on both response surface modeling (RSM) and what is called design and analysis of computer experiments (DACE); see for instance [2] and [3]. The approach consists of three steps being design of computer experiments (DoCE), response surface modeling (RSM), and design optimization, which will be treated in more detail in the sequel.

2.1 Design of computer experiments

First step is to define the design parameter space to be explored and to generate a set of designs to be simulated. This set should be such that accurate response surfaces can be constructed for all relevant simulation outputs. The art of efficiently setting up physical experiments is called *Design of Experiments (DoE);* see [4]. Computer experiments differ from classical DoE in the sense that experiments are not subject to random noise, which makes classical design schemes like fractional factorial schemes less suitable in a computer simulation setting. To have maximal signal to noise ratio, classical DoE focuses on parameter settings near the perimeter of the design region and builds in replicates (more than 1 experiment at the same design parameter settings).

In Design of Computer Experiments (DoCE) it generally gives better results to use space filling schemes in which the design points are distributed equally throughout the design space. For more background and references on this subject we refer to [5,6].

Deliverable of this step is a list of design parameter settings for a batch of simulation runs.

2.2 Response surface modeling

Second step is to obtain accurate response surface models for each of the relevant simulation outputs in terms of the design parameters. The batch of simulation runs, obtained in the first step, is performed using the simulation software. All output values of the simulation runs are collected and fit functions are determined for the relations between design parameters and output (response) parameters. For linear or slightly non-linear relations first or second order polynomial models are usually sufficient. In case of highly non-linear relations Kriging models [3] can be used, that are actually sophisticated interpolation algorithms that can fit any shape of response surface.

Very important part of this step is model validation. Perhaps the best way to assess model prediction capabilities is to run the simulation tool at a number of additional design points, not part of the original set, and to calculate the root mean squared error (RMSE) between RSM results and simulation results. However, taking into account that the simulations are typically very time-consuming, this is a costly option. We have good experiences with the use of (*leave-one-out) cross-validation* to assess the prediction capabilities of a response surface model. Given the simulated data of n design points, this technique loops through all the design points. In every cycle of the loop the current design point is held apart and an RSM model is constructed for the remaining $(n - 1)$ design points. The prediction capabilities of this RSM model are then assessed on the design point held out: the error between RSM and simulation result is determined. This error is

determined for all design points and in this way the *cross-validation RMSE* is calculated without having to perform additional simulations.

Note that these validations do not reflect the absolute accuracy of the calculated results with respect to real-life experimental results. They show how well the RSM models can substitute the more time-consuming computer simulations.

In this step it becomes clear if RSM is indeed the appropriate tool for the problem at hand. For problems with many different design parameters, the number of simulations required for the desired accuracy may run out of bounds. In that case, it would again be computationally more efficient to do well-chosen direct optimizations. This has to be weighed against the increased understanding that an RSM analysis offers. A good compromise for these problems would be to use RSM to get an approximate understanding where the interesting regions in the design parameter space are, and then use direct optimization to find the local optima with the required accuracy.

2.3 Design optimization

Now we have validated RSMs for the relevant simulation outputs, they can be used for prediction, robust design, optimization and sensitivity analysis. By applying Monte Carlo simulation techniques on the RSMs the robustness of a certain design to random variations due to manufacturing variation can be evaluated. Design optimization consists of finding values for some design parameters satisfying a set of constraints and minimizing/maximizing some kind of properly chosen merit function. Many software packages providing RSM functionality also offer some kind of mathematical optimization techniques. It is often very useful to exploit optimization to calculate so-called trade-off curves between two competing design requirements involving re-calculation of the optimal design for a number of successive design parameter values.

The results obtained in this paper where obtained by using the design optimization tool COMPACT developed by CQM [7].

3. LUMIRAMIC™ CONVERSION COMPONENTS FOR LEDS

3.1 Introduction to Lumiramic

Lumiramic is the Philips trademark for ceramic phosphor components for LEDs. Instead of the conventional phosphor technology, where a slurry of phosphor powder is deposited on a blue LED in a wet-chemical process, the phosphor is now a polycrystalline ceramic plate [1], which is a separate component that can already be characterized separately from the LED. These characterizations allow predicting the performance of the resulting white LED much better than is possible with the wet-chemical process. Optimal combinations of Lumiramic and blue LED can be determined beforehand, to obtain minimal spread in color point of the final product. See the schematic drawing in Figure 1 to show where the Lumiramic is placed in a typical Philips Lumileds LED package. It is used together with thin-film flip chip (TFFC) InGaN dice [8], that do not have any bond wires protruding from the top of the die and that provide a flat mechanical interface for Lumiramic tile attach.

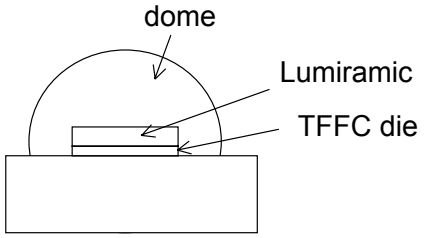

Figure 1. Schematic drawing (not to scale) that shows where the Lumiramic is located in a Luxeon® Rebel package. It is used together with a TFFC (thin-film flip chip) blue InGaN die.

Optical simulations were used to get an understanding of the optical processes inside the Lumiramic, such as absorption, fluorescence, and scattering, and how they determine the performance of the final product in terms of color point and efficiency. The optical simulations were performed using ASAP and LightTools. The blue die was taken as a Lambertian surface source of blue light, which was then traced through the Lumiramic model. The spatial distribution of the

absorption of blue light and the generation of phosphor emission was explicitly simulated. All rays were then traced to the point where they leave the LED package or where they are lost in any parasitic loss mechanism.

At the point in the design process where optical modeling was invoked, there were a number of obstacles for direct optimization methods. First of all, Lumiramic was designed to be a platform for a wide range of LED packages for different customers, end-users, and applications. There was no such thing as a well-defined merit function. Second, the optical models for the materials were not known very well and were difficult to model in the simulation software. This means that even if there had been a well-defined merit function, a single iterative optimization method would have given just one answer, probably wrong, and not contributed anything to the understanding of the system.

For these main reasons, an RSM approach was invoked. We wanted to understand the system, its sensitivities (how accurately do we need to know and control all material parameters), and its trade-offs. We wanted to respond quickly when merit functions became clear for certain LED applications.

3.2 Setting up the problem

After a first set of exploratory simulations, a list was made of the most important design parameters that determine the performance of a Lumiramic LED. They were separated into those design parameters that relate to the Lumiramic itself, and those that are part of the LED or the LED package around it. Then it was determined which of the output parameters must be determined by computationally intensive ray trace simulations, and which by more straightforward calculations. This resulted in the diagram of relations as in Figure 2.

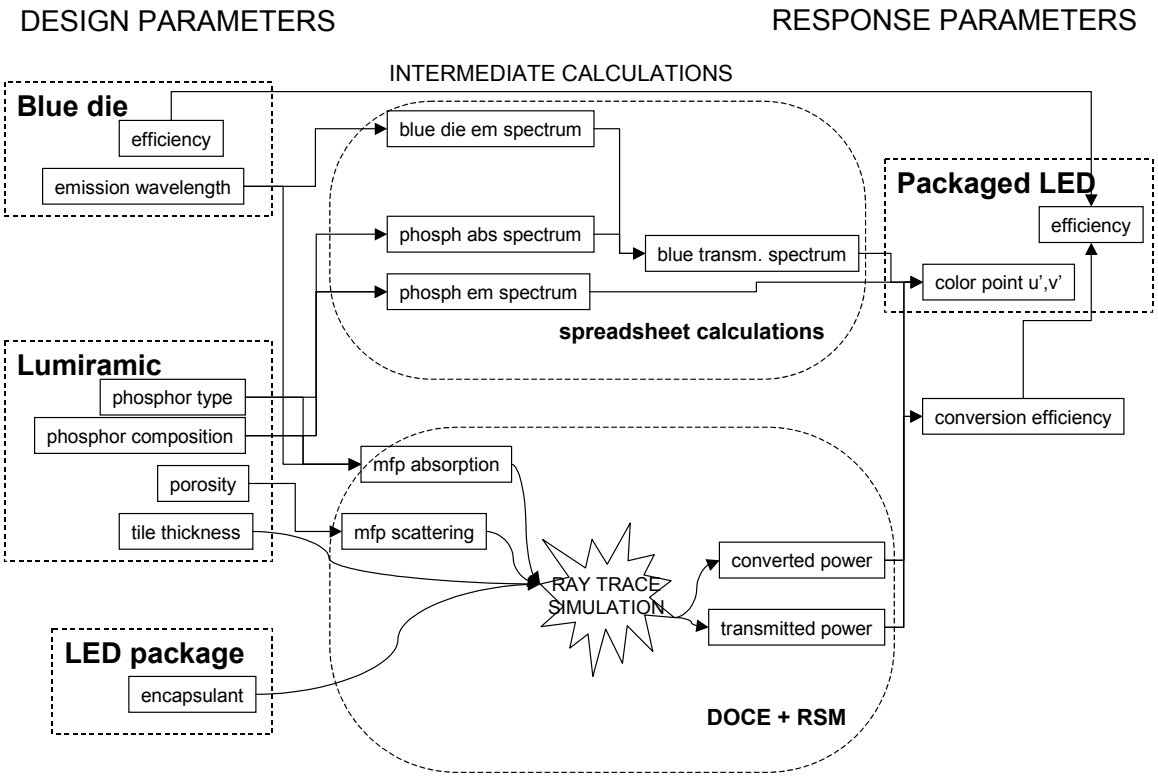

Figure 2. Schematic drawing of the relations between the most important design and response parameters, and how calculations were used to determine the dependencies.

For the more straightforward calculations, it makes sense to leave them out of the DOCE/RSM method. For instance, the effect of the emission wavelength of the blue die has the following two effects on the final result:

- the mean free path (mfp) for absorption depends on the blue die wavelength in a known way, through the excitation spectrum of the phosphor;
- there is a predictable shift in the *u'* color coordinate of the blue emission spectrum.

These two effects were kept outside the RSM models, and incorporated in spreadsheet models.

3.3 Design of computer experiments and building the RSMs

The DOCE was built as closely as possible around the computationally most intensive part of the calculation, which is the ray trace. For the DOCE, the following parameters were used as inputs: mfp for absorption, mfp for scattering, and tile thickness. The outputs were: the transmitted (blue) power and the emitted (phosphor) power, both scaled to 1 W input from the original blue source in the ray trace simulation. These two power values can be translated to a color point and to a conversion efficiency value in a relatively straightforward way.

Although the encapsulant of the LED package does play a role in determining the response parameters, it was kept constant in each of the DOCEs; a number of different DOCEs were performed, one for each LED package of interest, and separate response models were determined per LED package.

For each input parameter, a region was selected that was as wide as possible within all practical limits. Then, the COMPACT software tool [7] was used to generate a design space filling set of simulations. Numbers of simulations were typically between 50 and 120, depending on the desired accuracy. The software tool allows starting with a smaller set of simulations, and then generating additional simulations if the accuracy with the smaller set is not satisfactory. Each individual simulation could last around 30 minutes, leading to a total simulation time of 25 to 60 hours; although this may seem long, the big advantage is that all these calculations can be performed in a batch mode overnight or in the weekend, without the attention of an operator required.

The response surface models were generated for the direct outputs of the ray trace simulations (the converted and the transmitted powers), but it proved to increase the speed of optimizations to generate response surface models also for color point and conversion efficiency. The response parameters are then determined not directly from the DOCE, but after combining the direct output values of the DOCE with the spectral spreadsheet calculations.

Typical results are shown in Figure 3. Shown is a typical relation between thickness and *v'* color coordinate for three different phosphor compositions, differing only in the resulting mfp for absorption, where all other parameters are kept constant. The RSM for *v'* color coordinate was too complicated to be described by a simple polynomial, therefore a Kriging interpolation model was used. The software uses slide bars to vary the other parameters and get the new relation between thickness and *v'* instantly on screen while varying the parameters not explicitly shown in the graph. The right panel shows the result of the cross validation of the RSM for the *v'* coordinate for all simulations in the DOCE.

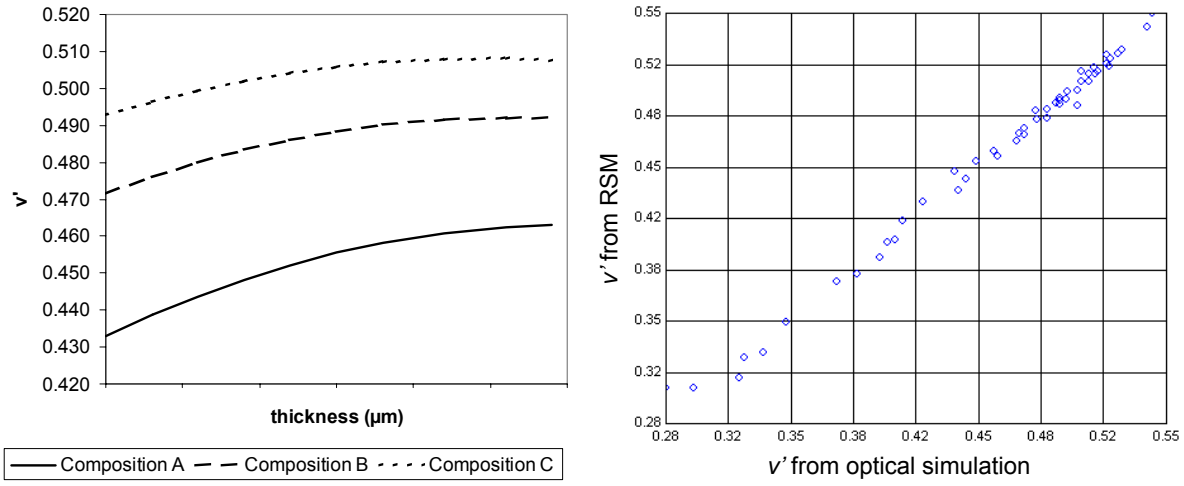

Figure 3. Example results from the analysis of the DOCE. *Left:* typical relations between tile thickness and *v'* coordinate for three different phosphor compositions, based on an RSM model. *Right*: the cross validation result used to validate the RSM model with respect to the detailed optical simulations.

Curves as in the left panel of Figure 3 allow for extremely quick optimizations. The designer does not have to know the exact color target at the time the simulations are run. When the designer is presented with a color target, he can immediately read from these curves what the initial guess for thickness and phosphor composition is.

It should be noted that there is a remaining uncertainty when comparing the simulation results to real life, due to the fact that the relation between the mfp for scattering, which is a microscopic quantity that cannot be measured directly, and the output of experiments designed to measure the scattering strength of the material (such as porosity determinations and angle-resolved laser scattering), is known only up to a certain level of accuracy. But this uncertainty does not affect the quality of the simulations. Once a more accurate relation between mfp for scattering and a measured material property is established, one can put more accurate values for the mfp in the model calculations, but the models themselves remain the same; it is certainly not necessary to re-do the computationally intensive run of ray trace simulations.

3.4 Analysis and application of the RSM models in the design process

The DOCE and RSM method was applied a number of times during the initial design stages of Lumiramic product development. It has contributed greatly to understanding the system, much more than direct optimizations would. Here are a few examples where the RSM methodology showed its value.

When optimzing for efficiency, we found that the routine always went for the lower limit of the thickness that we had set. However, it was found that in production, much robustness would be gained if the thickness was allowed to increase. The question was how much efficiency would be sacrificed for that. A simple calculation of efficiency vs. thickness for all other parameters constant indicated a quite strong effect. However, when the system was optimized toward the same target color point for each thickness value, the thickness dependence of efficiency reduced by a factor of more than 2. The latter analysis was indispensable for getting to the right conclusions. It involved five optimization runs, but because they were performed on the RSM models and the ray trace was not in the loop, these optimizations together took only approx. 1 minute of computation time, and this analysis could therefore be performed for several sets of boundary conditions.

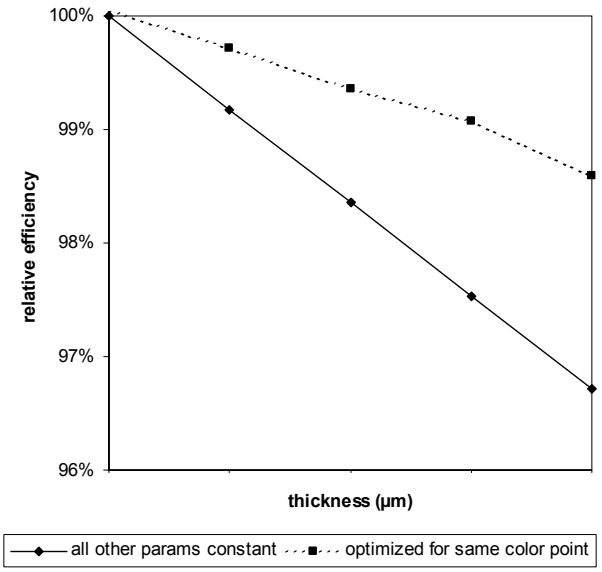

Figure 4. Efficiency effect of thickness. Shown as the effect of thickness when all other parameters are kept constant (solid line) and as the effect when an optimization toward the same color point is performed for each thickness value (dashed line). The efficiency of the global optimum was set as 100%.

The RSM models could also be embedded in spreadsheet programs where they are combined with other calculations. As shown in Figure 2, the effect of blue die wavelength was kept outside the RSMs and incorporated in afterward calculations, by substituting the corresponding mfp for absorption for the combination of blue die wavelength and phosphor composition, and by inserting the blue die emission spectrum in the spectrum calculator. In this way we built a

color point prediction tool that could be used "downstream" in the project where people had no expertise with the ray tracing software. Example output of this tool is shown in Figure 5. It allows finding the correct combinations of Lumiramics and blue dice that gives maximum yield on the color bins centered on the black body line, and determining the sensitivity of given Lumiramic designs for spread in blue die emission wavelength.

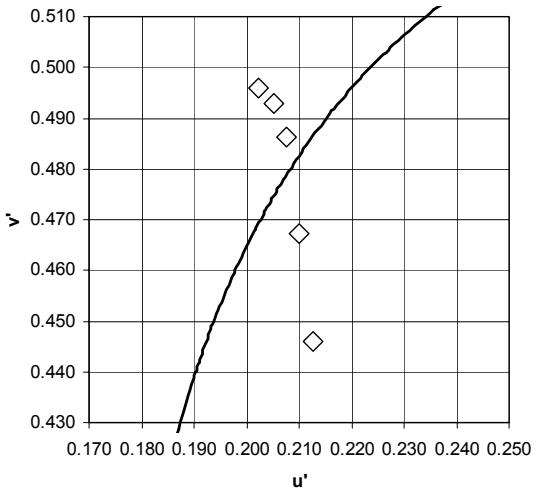

Figure 5. Example output of a spreadsheet tool that computes the color points (diamond symbols) of the same Lumiramic when combined with blue dice of five different emission wavelengths. Solid line is the black body locus.

Because the tool is already built up around RSMs, that use the same conceptual language as models that are derived from DOE experimental methods, the tool could easily be updated once results from real-life experiments became available.

All in all, the models are used on a regular basis, sometimes even "in-line" in discussions with our internal customers around the table, allowing sliders to be moved and parameters to be set and showing the results with no noticeable time lag on the projection screen.

4. SUMMARY AND CONCLUSIONS

This paper describes the method of response surface modeling (RSM) as an alternative to direct optimization algorithms. Its value is illustrated by its use in the optical simulations used in the early phases of design of Lumiramic, the ceramic phosphor components that are built into Philips Lumileds LEDs to create white LEDs with unprecedented low spread in color point.

It is shown that the RSM method is most valuable in cases where:

1. optical simulations are very time-consuming, even a single run, and where there are inherent uncertainties in the optical model, for instance when difficult-to-model processes as scattering and fluorescence are not present in the systems as perturbations, but as main effects that strongly determine system performance
2. the merit function is initially not well defined, especially concerning the relative weights and trade-offs between different response parameters
3. designing for robust performance is required and where the effect of spread in process and component parameters must be taken into account explicitly in the design

It gives the following advantages, compared to direct optimization methods:

1. more efficient use of people's and computer's time; the simulations are time-consuming, but they can be performed in a single long batch without attention from an operator; the subsequent analyses that do require the attention of a development engineer do not involve any slow computations

2. multidisciplinary and integral optimization, as the RSM framework and language is not restricted to the optical modeling but fits in a full-system Design for Six Sigma approach
3. the resulting RSM models can also be incorporated in spreadsheets that can be used by engineers that are not trained to use the simulation software
4. the RSM models provide more insights in parameter behavior, trade-offs and sensitivities, and thus provide the development team with more confidence that indeed the best solution is found, within all system constraints and robustness requirements.

RSM modeling is less beneficial when:

1. the system cannot be described in a limited set of parameters; this is the case e.g. when designing free-form optics; the number of DOCE simulations necessary to get sufficient accuracy goes up when the number of design parameters increases
2. the required accuracy of the answer is higher than the RSM models can get in a reasonable number of DOCE simulations; although in this case one can use RSM to find the approximate optimum and use this as a starting point in a direct optimization algorithm.

ACKNOWLEDGMENTS

We gratefully acknowledge the following colleagues who contributed to this study: Johan Ansems, Lars Casper, and Teus Tukker (Philips Lighting), who built the optical models and performed the ray tracing simulations; Erwin Stinstra (CQM), for his assistance in applying the RSM methodology to the Lumiramic project; Mart Peeters, Niels van der Veen and Jack Boerekamp (Philips Lighting), Helmut Bechtel and Peter Schmidt (Philips Research) for their input on Lumiramic technology.

REFERENCES

[1] Bechtel, H., Schmidt, P., Busselt, W., and Schreinemacher, B., "Lumiramic^{TM}: a new phosphor technology for high performance solid state light sources", Proceedings of the SPIE 7058, accepted for publication (2008).
[2] Myers, R.H., "Response surface methodology – current status and future directions", Journal of Quality Technology 31, 30-74 (1999).
[3] Sachs, J., Welch, W.J., Mitchel, T.J., and Wynn, H.P., "Design and analysis of computer experiments", Statistical Science 4, 409-435 (1989).
[4] Montgomery, D.C. 1984), *Design and Analysis of Experiments (Second Edition)*, John Wiley & Sons, New York, 1984.
[5] Stehouwer, H.P. and den Hertog, D., "Simulation-based design optimisation: methodology and applications", Proceedings of the First ASMO UK / ISSMO Conference on Engineering Design Optimization, Ilkley, UK (1999).
[6] den Hertog, D. and Stehouwer, H.P., "Optimizing color picture tubes by high-cost nonlinear programming", European Journal on Operational Research 140, 197-211 (2002).
[7] COMPACT, www.cqm.nl
[8] Shchekin, O.B., Epler, J.E., Trottier, T.A., Margalith, T., Steigerwald, D.A., Holcomb, M.O., Martin, P.S. and Krames, M.R., "High performance thin-film flip-chip InGaN-GaN light-emitting diodes", Applied Physics Letters 89, 071109 (2006).