

Sensitivity of night cooling performance to room/system design: surrogate models based on CFD

Kim Goethals¹, Ivo Couckuyt², Arnold Janssens¹ and Tom Dhaene²

¹Research Group of Building Physics, Construction and Services, Ghent University, Ghent, Belgium

²Internet Based Communication Networks and Services Research Group, INTEC, IBBT-Ghent University, Belgium

Keywords: night cooling, convective heat transfer, sensitivity, surrogate modelling

ABSTRACT

Night cooling, especially in offices, attracts growing interest. Unfortunately, building designers face considerable problems with the case-specific convective heat transfer by night. The BES programs they use actually need extra input, from either costly experiments or CFD simulations. Alternatively, up-front research on how to engineer best a generic night cooled office – as in this work – can thrust the application of night cooling. A fully-automated configuration of data sampling, geometry/grid generation, CFD solving and surrogate modelling, generates several surrogate models. These models relate the convective heat flow in a night cooled landscape office to the ventilation concept, mass distribution, geometry and driving force for convective heat transfer. The results indicate that cases with a thermally massive floor have the highest night cooling performance.

1. Introduction

Growing interest lies in passive cooling techniques, especially night cooling. After all, night cooling improves the summer comfort and minimizes the need for mechanical cooling. At night, natural or mechanical ventilation cools down the building fabric. The following day, the thermal mass absorbs the heat – reducing and delaying peak temperatures. For optimal performance three basic elements are necessary: the supply of cold air, the ability to store heat and the related heat transfer. Especially the convective heat transfer during nighttime plays a key role. Unfortunately, today's customary design tools, i.e. Building Energy Simulation (BES) programs, cannot grasp this case-specific convective heat transfer, at least not without extra input. Setting up new costly experiments just because of this is impossible. Meanwhile conflating Computational Fluid Dynamics (CFD) with BES in building design remains unattainable, not least because of the large computation time. Surely, researchers and software developers look for ways out for before-mentioned shortcomings. However, one way which pays off immediately, is to investigate with CFD how to engineer best a generic night cooled office. Yet, apart from the numerical and modelling errors in CFD, the computational expense remains the main issue. Without a doubt, parameterizing the design problem only amplifies this unmanageability. One indeed blunt solution is to just deploy more hardware. Another, more popular way is applying approximation methods to produce a so-called surrogate model, which either comes close to some part of the (unknown) reference model (i.e. local modelling) or is accurate over the complete design space (i.e. global modelling). In general, starting from a set of design parameters, an adaptive sampling algorithm selects iteratively new sample points. A computationally expensive simulator then evaluates these sample points, while the outcomes eventually serve to update the surrogate model. Sure enough, such approach enables to mimic, at a now manageable cost, the (unknown) response as closely as needed over (some part of) the design space. Certainly as an up-front investment such a study is worthwhile. The workflow of surrogate modelling is always pretty much the same, but interpretation of each step challenges even experts

in the field, let alone laymen like building engineers. Guidance on selecting and setting up such techniques or perhaps even a ready-made computer code is no luxury for them. One such convenient computer tool connecting the two worlds is the Matlab SURrogate MOdelling (SUMO) toolbox (Gorissen et al., 2010). Successful applications of this toolbox are plentiful: e.g. optimization of microwave filters and identification of electrical properties of textile antennas (Couckuyt et al., 2010) and blood flow data modelling (Degroote et al., 2011). Underlying study is just another such application, now in building engineering. The aim is to optimize the night cooling design. To this end, a global Surrogate-Based Optimization (SBO) procedure explores the design space. A fully-automated framework of sampling by SUMO, geometry and grid generation by Gambit (Fluent Inc., 2006), CFD solving by Fluent (Fluent Inc., 2006) and surrogate modelling by SUMO, generates several surrogate models. These surrogate models relate the convective heat flow in a night cooled landscape office to several room/system design parameters, which are subdivided into ventilation concept, mass distribution, geometry and driving force for convective heat transfer. Strictly speaking, these surrogate models merely guide the sampling towards the global optimum. Yet, they can also provide additional rough-hewn insight into the global behaviour.

2. Experimental design

2.1 Simulation experiment setup

2.1.1 Annex 20 2-D case as a starting point

Landscape offices usually have a large longitudinal section compared to the crosscut. This often brings along the use of line-shaped diffusers and band windows. All in all, this leads, roughly speaking, to 2-D airflow patterns, indeed influenced by 3-D eddies. Meanwhile, 3-D simulations render underlying study computationally infeasible. So, it is not a bad choice to limit the problem to a 2-D case. This study starts from the well-researched 2-D Annex 20 case (Lemaire, 1993). This simple case dates back to the scale model experiments of (Schwenke, 1975) in the mid-seventies of the previous century, but not until the early nineties it found acceptance as a valuable benchmark.

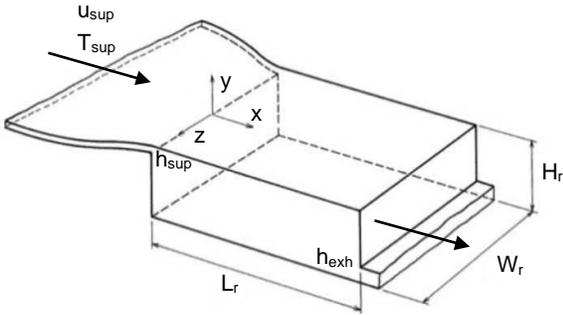


Figure 1 Blueprint of Annex 20 2-D experiment setup

Figure 1 depicts this Annex 20 2-D setup. Basically, it concerns a rather long ($L_r/H_r=3.0$) and wide ($W_r/H_r=1.0$ or 4.7) ventilated room having on either side a wall-to-wall opening. The air supply opening on the left side is rather large compared to practical diffusers ($h_{sup}/H_r=0.056$), but this relaxes the number of grid points near this opening in simulations. The height of the exhaust opening h_{exh} is to the height of the room H_r as 0.16 is to 1 . The air supply temperature T_{sup} equals 20°C , the air change rate n is 10.2h^{-1} . In the Annex 20 2-D2 case, a constant heat flux is added along the floor – which is indeed raised in succeeding experiments (i.e. summer cooling at a range of Richardson numbers). Critical factor is the impact of the Richardson number (Eq. (1)) on the jet penetration. The streamlines at the midplane act as a reference. Also underlying study relies on these data to validate the CFD simulation approach.

$$Ri = \frac{\beta \cdot g \cdot h_{sup} \cdot (T_w - T_{sup})}{u_{sup}^2} \quad (1)$$

2.1.2 Parameterization

This simple reference case enables a straightforward parameterization. For that matter, Table 1 states all considered design parameters with their respective categories/continuous interval. As previously mentioned, this sensitivity study identifies among the design parameters four subsets: ventilation concept, mass distribution, geometry and driving force for convective heat transfer. Mind you, future simulations also include the length of the room and the inclination of the air supply.

Table 1 Overview of parameters for sensitivity study

Parameter	Type	Min	Max
-	Ventilation concept	Single side/cross/under floor	
-	Location of isothermal plane	Floor/ceiling	
H_{sup} (m)	Geometry	0+BL	2.5+BL (4.0)
H_{exh} (m)		0+BL	2.5-BL
$h_{sup/exh}$ (m)		1.5	10
n (h^{-1})	Driving force	1.5	10
T_w ($^\circ\text{C}$)		16	25

The first subset, i.e. ventilation concept, which matches its single input parameter, makes a clear-cut distinction: cross, single sided or under floor ventilation. Another such subset concurring with a single variable, is the distribution of mass. Here, the choice between floor and ceiling means that the declared surface is at a higher temperature than the supply air while the remaining surfaces behave adiabatically (Figure 2). As a matter of fact, starting from steady boundary conditions, the one isothermal plane represents a thermally heavy finishing while the adiabatic surfaces correspond to light structures. The subset geometry comprises the distance of the

air supply to the zero point $(0,0)$ H_{sup} , the distance of the exhaust opening to the zero point $(0,0)$ H_{exh} and the height of the air supply/exhaust $h_{sup/exh}$ (Figure 2). Now, continuous numerical intervals apply, in contrast with the previous two categorical variables (i.e. without numerical meaning) (Table 1). Note that, for programming simplicity, the bounds of H_{sup} and H_{exh} are corrected for the boundary layer thickness BL. For that same reason, H_{sup} is limited to $L_r-0.5\text{m}$ in case of under floor ventilation. Also mark in Figure 2 that underlying study constantly makes use of a quite long exhaust channel, simply to incorporate the possible recirculation flow. Otherwise, the CFD solver often stalls. The last subset, i.e. driving force for convective heat transfer, takes into account the relative magnitudes of the forced convection component, defined by the air change rate n , and of its natural convection equivalent, represented by the temperature of the isothermal plane T_w . Note that the supply air temperature T_{sup} equals 15°C in every simulation. Last-mentioned input variables are, just like the ones part of the subset geometry, quantitative variables, each defined by a continuous interval. For that matter, stating the type of these variables is not trivial. After all, it determines whether or not the variable can be part of the surrogate modelling process. In this study, the ventilation concept and the mass distribution, i.e. two categorical subsets, are outsiders. Therefore, this study defines 3 (cf ventilation concept) times 2 (cf mass distribution) base cases and lets SUMO built for each one of them a surrogate model which incorporates the subsets geometry and driving force for convective heat transfer.

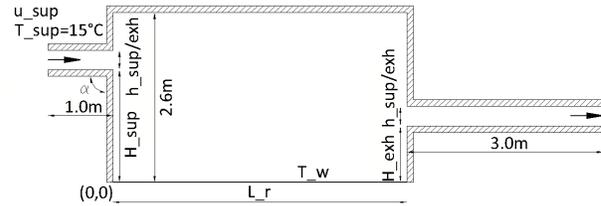


Figure 2 Schematic simulation experiment setup

2.2 GAMBIT + FLUENT + SUMO = surrogate model

Applying the SUMO toolbox implies that Matlab (MathWorks, 2011) controls the actions. The SUMO toolbox itself is a set of Matlab scripts which controls the whole process and takes care of data sampling and surrogate modelling. Logically, a custom-made Matlab function then interacts with Gambit and Fluent by reading, manipulating and eventually evaluating journal files (i.e. text files containing program-specific commands). All in all, this enables a fully-automated repeating workflow of sampling by SUMO, geometry and grid generation by Gambit, CFD solving by Fluent and surrogate modelling by SUMO. Table 2 presents this sequence in more detail. In general, the workflow starts with SUMO sampling input parameter values. Subsequently, the Matlab function reads in a generic Gambit journal file, fills in the geometrical parameters and evaluates the script. Gambit is launched, the grid built and, before closing, exported to a Fluent compatible mesh file. Once the Matlab function has registered this mesh file, it reads in the generic Fluent journal file 'steady/isothermal', fills in the parameters involving the subset driving force for convective heat transfer and evaluates the script: a first Fluent simulation takes off. The moment that the Matlab function registers the corresponding saved Fluent file, it reads in a second generic Fluent journal file, i.e. 'steady/non-isothermal', fills in the necessary quantities and evaluates it.

A first so-called steady/non-isothermal Fluent case is launched. At the end, the Matlab function performs a double check. First, it verifies whether the number of iterations performed during loop A exceeds 6000. In the second place, the Matlab function checks, again by means of a generic journal file, whether or not the heat flux imbalance is smaller than the heat flux at the isothermal plane. If not, the Matlab function extracts, with another journal file, the simulated temperature at the exhaust and feeds it to a new ‘steady/non-isothermal’ Fluent simulation. On the other hand, once one of before-mentioned conditions is fulfilled, the Matlab function moves to loop B. Now, the workflow continues with unsteady RANS simulations (controlled in a similar fashion). To advance to loop C, again one of two conditions needs to be fulfilled. Either the number of time steps in loop B exceeds 6000 or the heat flux imbalance is at the most 10% of the convective heat flux at the isothermal plane. Loop C introduces solution adaptive grid refinement, which reduces the numerical error in the solution, with minimal numerical cost. At this stage, severe convergence criteria hold. When the convective heat flux at the warm surface is higher than $0.1W.m^{-2}$, the heat flux imbalance needs to be smaller than $0.05W.m^{-2}$ and 1% of the surface heat flux. Otherwise, the heat flux imbalance can go up to 10% of the surface heat flux. Only then, the final Fluent simulation, now without grid adaptation, starts. The time-averaged convective heat flux at the isothermal plane obtained from this last simulation is then used by SUMO to update the surrogate model. When some simulation time is left and the sample budget is not used up, the sequence starts over again. At first sight, the sequence of different Fluent simulations looks laborious. However, experience indicates that this procedure reaches much faster a converged solution. The preliminary steady simulations provide a reasonable estimate of the airflow. Meanwhile, loop B corrects for unsteadiness before the grid adaptive simulations come in. Moreover, adjusting the backflow temperature on the fly is no luxury. After all, a badly chosen backflow temperature often slows down the convergence rate or leads even to a wrong solution. Furthermore, in this study, only the heat imbalance acts as a convergence criterion as monitoring residuals and target quantities leaves considerable room for interpretation while getting to a small heat imbalance proves to be a more severe target.

Table 2 Flow of Gambit, Fluent and SUMO

Action	Software	Progress condition
data sampling/ surrogate modelling	SUMO	
generate geometry/grid	Gambit	
perform CFD calculations: steady/isothermal (1x), steady/non-isothermal	Fluent (loop A)	#iterations(loopA)>6000 or heat imbalance<100%
perform CFD calculations: unsteady/non-isothermal	Fluent (loop B)	#iterations(loopB)>6000 or heat imbalance<10%
perform CFD calculations: unsteady/non-isothermal/ adaptive grid	Fluent (loop C)	If heat flux>0.1W.m ⁻² : heat imbalance<1% and imbalance flux<0.05W.m ⁻² Else: heat imbalance<10%
perform CFD calculation: unsteady/non-isothermal	Fluent	

2.3.1 Geometry/grid generation

Feeding the geometry into Gambit is a straightforward procedure. Drawing up a proper grid, however, usually involves a great deal of trial and error. A way to sidestep this handicap in a fully-automated process is solution adaptive grid refinement. Table 2 already sketched the principle. First, a sequence of CFD simulations on a generic grid provides an estimate of the airflow. Then, the CFD software refines the grid on the fly, which, in the end, leads to an optimal grid distribution. Incidentally, this grid adaptation procedure also saves on computational resources. This study puts three distinct parameterized Gambit journal files forward. These actually tie in with the investigated ventilation concepts (single sided, cross, under floor). Once read into Gambit, each one of them brings forth the case-specific geometry and the corresponding generic grid. To homogenize the resulting grids, this study opts for a clear zoning (Figure 3). Otherwise, the few size functions can mess up the grid. First of all, size functions start from the points connecting the air supply/exhaust with the room. In this way, the grid resolution luckily increases near these openings. Obviously, this implies the use of an unstructured or pave mesh; which is preferably limited in space. For that reason, this study restricts the use of this mesh type to the immediate proximity of the air supply/exhaust. However, these zones are sufficiently wide so that the enclosed pave mesh fits the adjoining structured or map meshes. As matter of fact, the innermost zone comprises a regular structured grid, i.e. made out of square elements. The remaining zones include rectangular elements. On the one hand, this last-mentioned mesh type allows size functions in the air supply/exhaust to continue the x-wise expansion of the cells (y-wise in the supply in case of under floor ventilation). On the other, the height of the cells in the middle zones can gradually diminish towards the surface boundaries. Here, a dense boundary layer mesh is present, simply to fulfil the near-wall modelling conditions ($y^+ < 1$ and $Re_y < 200$). To refine (or coarsen) the grid, Fluent makes use of a gradient adaptation function. This approach assumes that maximum discretisation errors occur in high-gradient regions, which, as a consequence, need refinement. What’s more, it relies on the rate at which the gradient of the solution changes (Laplacian), which proves especially helpful for smooth solutions. For that matter, this Laplacian needs to be normalized. After all, re-adjusting the coarsen/refine thresholds during an automatic dynamic adaptation process would be a hopeless task. Further, neighbouring cells are not allowed to differ by more than one level of refinement, simply to avoid excessive cell volume variations, while the minimum cell volume is bounded. The solution variable used for all this, is the turbulent kinetic energy, which refines in particular in the jet and near-wall regions.

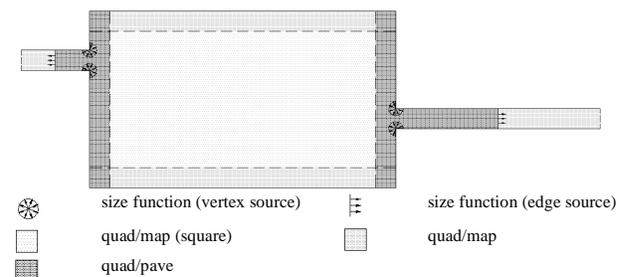


Figure 3 Scheme of the grid

2.3.2 CFD solver settings

Firstly, the pressure-based solver takes care of the linearization and the solution of the discretised equations. The PISO algorithm is responsible for the pressure-velocity coupling. The second-order upwind scheme guarantees the interpolation of the convection terms; the PRESTO! scheme the interpolation of pressure. Similarly, a time discretisation scheme is necessary. After all, this study goes on time-dependent simulations. The steady boundary conditions are assumed to lead to practically time-independent solutions; which allows the use of a first-order (implicit) time discretisation scheme. Further, to model turbulence, this study uses the RNG k- ϵ model. After all, this model performs best, in terms of accuracy, computing efficiency and robustness, over a wide range of flow regimes. This versatility is indeed no luxury for underlying study: the driving forces for convective heat transfer – and, thus, the flow regimes – come in a wide variety. Meanwhile, the simulations obviously need to resolve the near-wall region. Otherwise, the prediction of the convective heat transfer goes wrong. In case of the used High Reynolds Number model, the two-layer approach urges itself. In this case, the RNG k- ϵ model takes care of the fully-turbulent region while the one-equation model of Wolfstein resolves the viscosity-affected near-wall region. The boundary conditions depend for the most part somehow on the parameter values set by the SUMO toolbox (Table 1). Like this, the air supply boundary conditions, i.e. velocity and turbulent intensity, go largely on the height of the air supply $h_{\text{sup/exh}}$ and the air change rate n . Meanwhile, the supply air temperature T_{sup} equals 15°C in all simulations. For that matter, uniform values of these quantities apply to the ‘velocity inlet boundary condition’, located at one end of the air supply duct. A ‘pressure outlet boundary condition’ represents the exhaust. The gauge pressure at the opening is set as zero using the Dirichlet condition. In case of backflow, the air supply temperature of the exhaust equals the updated mass-averaged temperature near this opening while the turbulence intensity assumedly equals the one of the air supply. Furthermore, all surfaces but one (floor or ceiling) behave adiabatically. This one surface gets its temperature directly from the SUMO toolbox. As the boundary conditions remain constant during each simulation, the walls obviously have no mass. Further, the physical properties of the air need to be defined. The viscosity, the heat capacity and the thermal conductivity are piece-wise functions of temperature while the variable air density is calculated with the incompressible ideal gas model. Finally, radiation is not part of the simulations.

2.4 SUMO settings

2.4.1 Data collection strategy

Building a surrogate model implies populating a dataset. Classical Design Of Experiments (DOE) techniques for computer-based experimentation, such as orthogonal array sampling, Latin hypercube sampling, try to cover the design space as uniformly as possible. Assumedly, this maximizes the accuracy of the resulting one-shot surrogate model. However, now there is a growing body of opinion that sequential design, also known as adaptive sampling, is a better way. Adaptive sampling addresses the main shortcoming of DOE: the up-front number of sample points and their distribution over the parameter space are not necessarily appropriate. In adaptive sampling, the (repeating)

sequence of (1) generating a (multiple) candidate sample point(s), indeed based on an intermediate surrogate model, (2) performing a simulation(s), (3) generating/updating the surrogate model as best one can, until some target is reached, prevents oversampling. However, not all adaptive sampling strategies are equally suited. For example, for global SBO (used in this study), it is crucial that the adaptive sampling strategy strikes the correct balance between exploration (i.e. enhancing the general accuracy of the surrogate model) and exploitation (i.e. enhancing the accuracy of the surrogate model in the region of the (intermediate) optimum). One such popular strategy is the expected improvement algorithm (Jones et al., 1998), on which this study builds. To introduce expected improvement, this study reverts to a graphical illustration. Figure 4 shows how an intermediate surrogate model, based on ten sample points, approximates an unknown one-variable reference model. The function value at any point x is treated as the realization of a normally distributed variable $Y(x)$, with mean $\mu(x)$ and variance $\sigma^2(x)$ (i.e. a Gaussian process). Close to the sample points, the prediction is accurate. However, the surrogate model completely misses the ball in the data-sparse region on the right-hand side of the plot (i.e. where coincidentally the global optimum lies). The probability density function (PDF) at e.g. $x=2.4$ indicates a large uncertainty in the function value. What’s more, the tail of the PDF extends below the line $y=f_{\text{min}}$, which means that the function value can improve on the current best (simulated) function value f_{min} . The corresponding shaded area under the PDF matches the cumulative distribution $P(Y(x) \leq f_{\text{min}})$, or better, the probability of improvement $\text{PoI}(x)$. Probability of improvement may then already indicate the possibility of a better minimum, it does not quantify how large the improvement will be. By contrast, the first moment of the shaded area, better known as expected improvement, does! Expected improvement $E[I(x)]$ equals every possible improvement at x , i.e. $I(x)$, multiplied by the associated likelihood (Eq. (2)). In short, the expected improvement is large where $Y(x)$ is likely smaller than f_{min} and/or where there is a high uncertainty in the prediction value itself. On the other hand, at sample points, the prediction variance $\sigma^2(x)$ equals zero and, thus, the expected improvement is nil.

$$E[I(x)] = \int_{-\infty}^{f_{\text{min}}} I \cdot Y(x) dY \quad \text{where } I(x) = \max[f_{\text{min}} - Y(x), 0] \quad (2)$$

The optimization process normally includes all parameters from the start. However, the CFD simulations in this study are so expensive that it is better to increase sequentially the number of parameters. This leaves the opportunity to check (and validate) the intermediate optimization results and associated surrogate models. For that matter, this explains the presence of collinear points in the plots to come. In fact, the sampling starts from a Latin hypercube design which includes only two parameters, i.e. H_{sup} and H_{exh} . From there one, adaptive sampling comes in, initially superadding two more parameters: n and T_w . Later on, $h_{\text{sup/exh}}$ applies also to the parameter space. To determine the next sample point in this iterative process, the competitive DIviding RECTangles (DIRECT) algorithm optimizes the expected improvement criterion. When the DIRECT algorithm fails to find a unique sample, the optimization of the fall back criterion starts. The model then looks for the parameter combination that has the largest prediction variance.

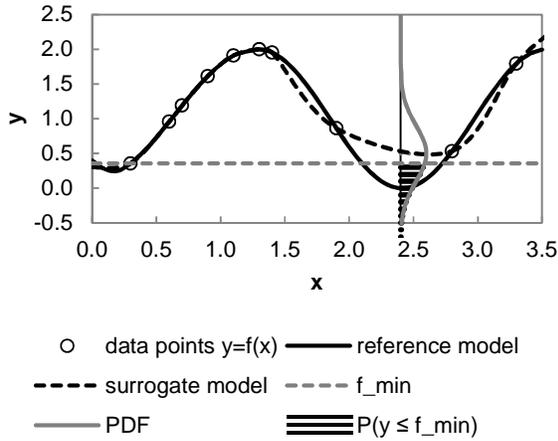


Figure 4 Graphical illustration of expected improvement

2.4.2 Modelling strategy

Once the data collection strategy is fixed, it is necessary to choose how to build a surrogate model from the scattered data. Generally, interpolation surrogate models approximate best deterministic computer experiments. Such a model is usually a linear combination of polynomial terms, which model the trend over the parameter space, and special ‘basis function’ terms, which ‘pull’ the surrogate model through the observed data (Jones et al., 1998). In particular the form of the basis function, which quantifies the correlation of (nearby) points, determines the usability of the corresponding interpolation technique. In particular kriging stands out, because its basis function includes tunable parameters. This in fact statistical feature of the basis function not only allows to compute an interpolator (or ‘predictor’), but also enables to estimate the potential error in the predictor. Precisely this bonus is necessary to apply the expected improvement algorithm. Logically, this study deploys (ordinary) kriging – simply called kriging in the remainder. Basically, kriging tries to construct, from n sample points \mathbf{x}_i (where \mathbf{x}_i is a d -dimensional vector $\mathbf{x}_i=(x_{i1} \ x_{i2} \ \dots \ x_{id})'$) and the n corresponding function values $y_i=y(\mathbf{x}_i)$, a surrogate model with minimal prediction variance. Yet, there is more to kriging than meets the eye. Actually, kriging assumes that the function being studied is a realization of a Gaussian stochastic process $Y(\mathbf{x}_i)=\mu+Z(\mathbf{x}_i)$. At this, the constant mean μ (i.e. a constant polynomial) replaces the regression terms while the zero-mean, stationary, Gaussian stochastic process $Z(\mathbf{x}_i)$, with variance σ^2 and parametric correlation function Eq. (3), corresponds to the basis function terms.

$$Corr[Y(\mathbf{x}_i),Y(\mathbf{x}_j)]=\exp\left(-\sum_{\ell=1}^d \theta_{\ell} |x_{i\ell} - x_{j\ell}|^{p_{\ell}}\right) \quad (3)$$

This correlation function implies that the random variables $Y(\mathbf{x}_i)$ and $Y(\mathbf{x}_j)$ – and, thus, the function values $y(\mathbf{x}_i)$ and $y(\mathbf{x}_j)$ – correlate better when their sample points \mathbf{x}_i and \mathbf{x}_j are closer, i.e. their Euclidean distance in the ℓ^{th} input dimension of the sample points \mathbf{x}_i and \mathbf{x}_j is smaller. The first (tunable) correlation parameter in Eq. (3), i.e. $\theta_{\ell} (\geq 0)$, indicates the importance of input dimension ℓ : the higher θ_{ℓ} is, the faster the correlation decreases with distance. The other one, i.e. p_{ℓ} , determines the smoothness of the function in the ℓ^{th} input dimension. For example, $p_{\ell}=1$ yields the exponential correlation function; $p_{\ell}=2$ gives the so-called Gaussian

correlation function. Yet, this study sets p_{ℓ} equal to two. As a consequence, the behaviour of the function only depends on the parameters μ , σ^2 and θ_{ℓ} , of which the (estimated) values should maximize the probability (i.e. the likelihood) of the sampled data. Whilst deriving the predictor value for an additional point \mathbf{x}^* – which is either an old point or a new one –, the aim is obviously once again maximizing the – now ‘augmented’ – likelihood. The derivation of this kriging predictor is out of scope. Therefore, this work restricts itself to stating the standard formula, i.e. the Best Linear Unbiased Predictor (BLUP) (Eq. (4)):

$$y(\mathbf{x}^*) = \mu + \mathbf{r}^T \mathbf{R}^{-1} (\mathbf{y} - \mathbf{1}\mu) \quad (4)$$

where $\mathbf{r} = \{\text{Corr}[Y(\mathbf{x}^*), Y(\mathbf{x}_1)] \dots \text{Corr}[Y(\mathbf{x}^*), Y(\mathbf{x}_n)]\}^T$ is the vector of correlations between \mathbf{x}^* and the n sample points, \mathbf{R} is the $n \times n$ correlation matrix whose $(i,j)^{\text{th}}$ entry is given by Eq. (3) and $\mathbf{1}$ denotes the n -dimensional vector with ones. Now, the Generalized Least Squares (GLS) determines the constant mean μ and the process variance σ^2 while the Sequential Programming (SQP) algorithm determines θ_{ℓ} so that the probability of the sampled data is at his maximum. For that matter, the predictor is more reliable when the ‘augmented’ likelihood drops off dramatically as one moves away from the optimal value of $y(\mathbf{x}^*)$. Actually, this closely relates to the mean squared error of the predictor, i.e. $\sigma^2(\mathbf{x}^*)$ (Eq. (5)). This last-mentioned formula has the intuitive property that at any sampled point it equals zero.

$$\sigma^2(\mathbf{x}^*) = \sigma^2 \left(1 - \mathbf{r}^T \mathbf{R}^{-1} \mathbf{r} + \frac{(\mathbf{1} - \mathbf{r}^T \mathbf{R}^{-1} \mathbf{r})^2}{\mathbf{1}^T \mathbf{R}^{-1} \mathbf{1}} \right) \quad (5)$$

3. Results

3.1.1 Validation by means of the Annex 20 2-D2 case

As validation data for the case(s) at hand is unavailable, this study retakes the similar Annex 20 2-D2 case, adopts the simulation approach described in 2.3 and looks if, as in the measurements, the jet deflects at a Richardson number of 0.02. To determine this deflection point, the distance from the zero point in the x -direction at which the x -wall shear stress at the ceiling becomes negative, i.e. x_{re} , is plotted as a function of the Richardson number. Figure 5 indicates that the CFD approach and the adopted grid density of 2730cells.m⁻² approximate the deflection point better than the simulation studies of Annex 20. So, the CFD approach and grid density are considered to be apt.

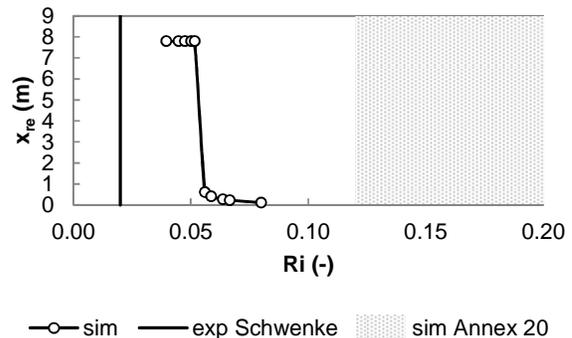


Figure 5 Reduced penetration depth x_{re} as a function of the Richardson number Ri

3.1.2 Pareto optimality

The actual analysis begins by looking at the Pareto optimality of the respective cases. After all, this study looks first and foremost for optimal solutions. The surrogate models are merely a bonus. In particular the design of fan-assisted night cooling is faced with the trade-off between the convective heat flow by night and the energy use of the fans. So, it is interesting to assess the accompanying Pareto optimality. Figure 6 depicts for the six combinations of ventilation concept and thermal mass distribution the convective heat flux-air change rate Pareto fronts. Each Pareto front is based on one of the six surrogate models and indicates the maximum attainable convective heat fluxes for the full range of n and $(T_w - T_{sup}) = 10^\circ\text{C}$ for each base case. As expected, the higher the air change rate, the larger the convective heat flow. What's more, the cases with a so-called massive floor clearly outperform the ones with the ceiling at a higher temperature. They produce convective heat flows which are four to five times as large. In effect, the cold jet truly adheres to the ceiling only in case of predominantly forced convection (i.e. for low T_w and high n). Otherwise, gravity pulls the jet down. That same reasoning also explains why optima of the cases with a warm ceiling rise, relatively speaking, more with increasing air change rates. Now, looking merely at the three best performing cases reveals that the cross ventilation concept results in the overall largest convective heat flow. As a matter of fact, in case of cross ventilation, the cold air coming out of the supply, generally strikes the warm floor with a reasonable speed and then continues to flow along that same floor. On the other hand, single sided ventilation, in particular with the air supply on top of the exhaust, may short-circuit the airflow, leaving most of the room air untouched. Even primarily forced convection regimes, of which Figure 7 depicts an example, suffer from this deficit. With under floor ventilation, conversely, the jet is directed away from the warm floor: here, most of the jet momentum is lost by the time the cold air hits the floor. What's more, there is again a considerable risk that the jet never strikes the warm floor. For example, the forced convection case in Figure 8 exhibits a clear short-circuit between the air supply and the exhaust.

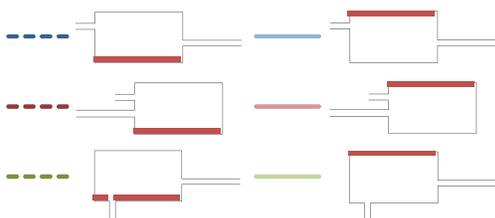
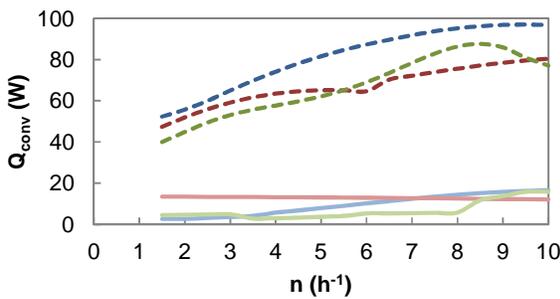


Figure 6 Convective heat flow-air change rate Pareto front (for references to colour, the reader is referred to the digital version)

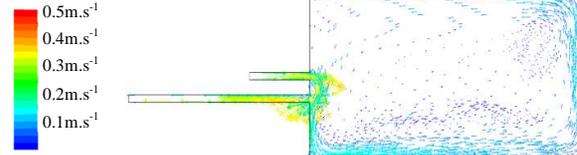


Figure 7 Velocity vectors of case 'single sided, warm floor, $H_{sup}=1.24\text{m}$, $H_{exh}=0.87\text{m}$, $h_{sup/exh}=0.12\text{m}$, $n=9.75\text{h}^{-1}$, $T_w=18.51^\circ\text{C}$ ' (for references to colour, the reader is referred to the digital version)

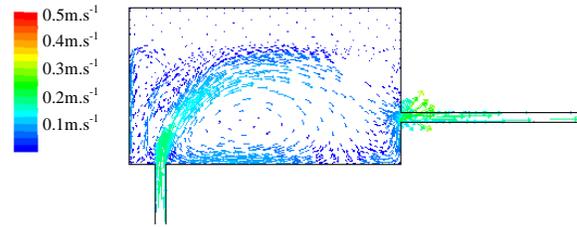


Figure 8 Velocity vectors of case 'under floor, warm floor, $H_{sup}=0.43\text{m}$, $H_{exh}=0.70\text{m}$, $h_{sup/exh}=0.17\text{m}$, $n=9.53\text{h}^{-1}$, $T_w=17.50^\circ\text{C}$ ' (for references to colour, the reader is referred to the digital version)

3.1.3 Massive floor cases: points of interest

As night cooling benefits in particular from adding a massive floor, the following section merely zooms in on these special cases. Figure 9, Figure 10 and Figure 11 display several contour plots of the surrogate models of the three studied ventilation concepts with a massive floor. The contour plots show how the convective heat flow relates to the distance of the air supply to the zero point H_{sup} (x-axis of each contour plot) and the distance of the exhaust to the zero point H_{exh} (y-axis of each contour plot). Each individual contour plot corresponds to one combination of fixed values of the remaining three parameters, i.e. the height of the air supply/exhaust $h_{sup/exh}$, the air change rate n and the temperature of the isothermal plane T_w (i.e. a slice). The x-axis of the graph itself includes nine fixed discrete values of $h_{sup/exh}$ while in the y-direction two axes prevail: one for T_w and an underlying one for n . Meanwhile, the dots are projections of the sample points. Their dimension indicates the distance of these sample points to the slice. By the way, mind that the surrogate models rely on the negative of the convective heat flow. After all, the expected improvement algorithm directs the solution procedure towards the global minimum. Further note that the scales in the respective graphs differ slightly. Finally, each graph of a ventilation concept clearly exhibits very specific features. Therefore, the following discusses the figures one by one. Figure 9, which depicts the results of the cross ventilation concept, states the obvious: to get a high convective heat flow, a high air change rate and a large temperature difference between the supply air and the isothermal plane is primordial. In particular such mixed convection flow regimes are sensitive to the position of the air supply and the one of the exhaust. Yet, locating the air supply at the top and the exhaust at the bottom clearly leads to the best night cooling. As previously mentioned, this is, on the one hand, thanks to the concerted action of the jet x-momentum and the gravitational acceleration and, on the other, due to the absence of a recirculation flow in the lower right corner of the room. Conversely, the position of in particular the air supply is of no account in primarily natural/forced convection regimes ($T_w=298.15\text{K}$, $n=1.5\text{h}^{-1}$ / $T_w=289.15\text{K}$, $n=10\text{h}^{-1}$).

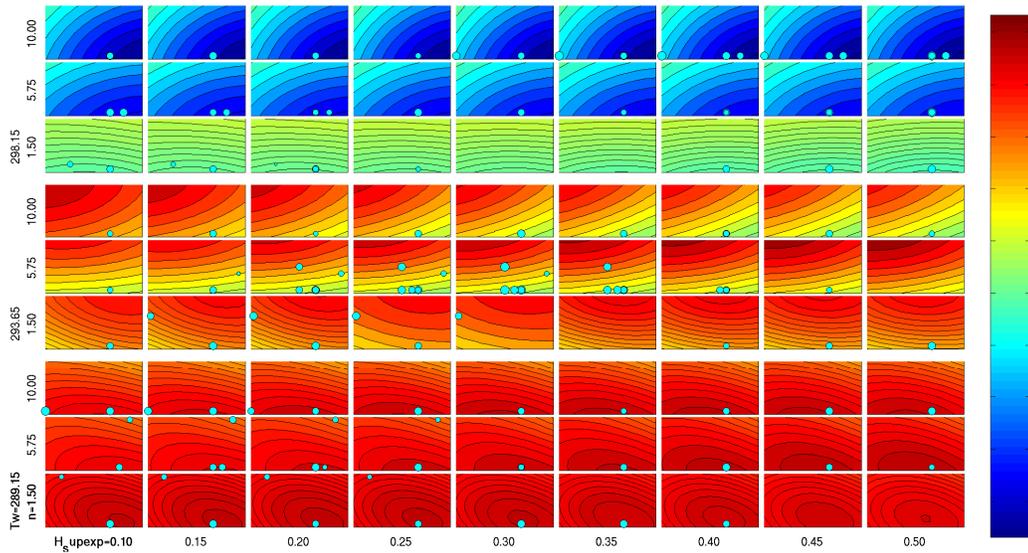


Figure 9 Contour plots of the convective heat flow for case 'cross ventilation concept with the floor at a higher temperature' (for references to colour, the reader is referred to the digital version)

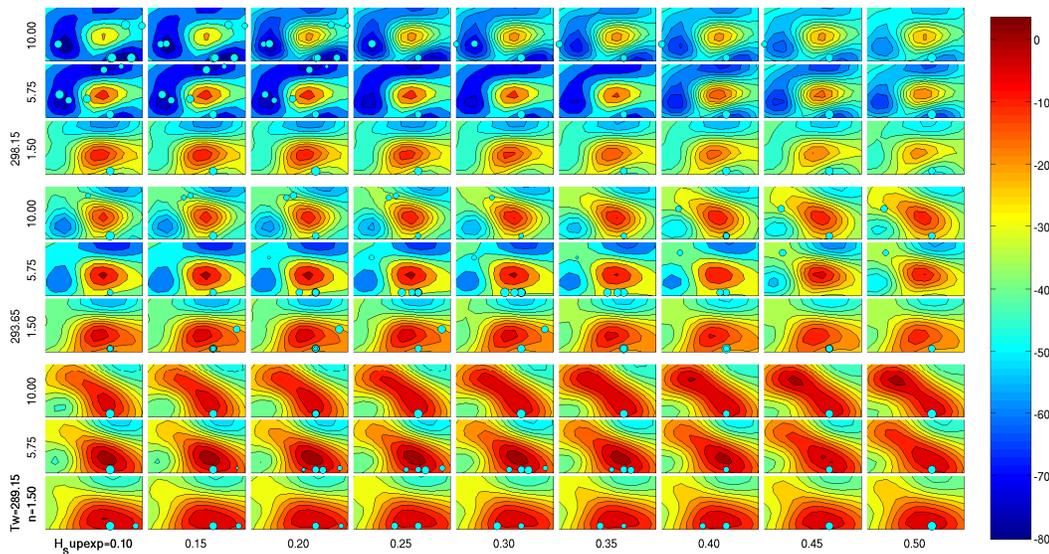


Figure 10 Contour plots of the convective heat flow for case 'single sided ventilation concept with the floor at a higher temperature' (for references to colour, the reader is referred to the digital version)

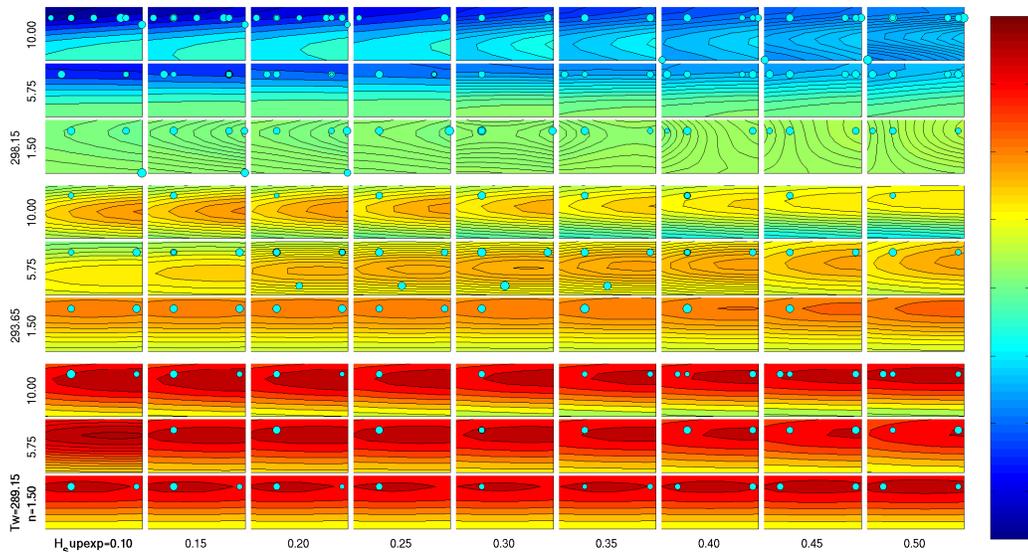


Figure 11 Contour plots of the convective heat flow for case 'under floor ventilation concept with the floor at a higher temperature' (for references to colour, the reader is referred to the digital version)

In the meantime, the height of the air supply/exhaust is not that important. With single sided ventilation (Figure 10), things are turned upside-down. Now, only for a high surface temperature of the floor, the cases with the air supply at the top and the exhaust at the bottom perform reasonably well. For the rest, the thing is to put the exhaust at a position higher than the one of the air supply. After all, such configuration unlikely leads to a short-circuit of the airflow (remember Figure 7). Hereby, limiting the distance between the two openings and putting them either at the top or the at the bottom of the wall, leads in most cases to the highest convective heat flow. What's more, limiting the height of the air supply and, thus, increasing the jet momentum, pays off. By this, the jet penetrates deeper into the room. At last, under floor ventilation (Figure 11) reveals a confined impact of H_{sup} . That's only logical: H_{sup} represents the position of the air supply in the x-direction of the room. Yet, moving it away from the exhaust as far as possible can do no wrong. So, building designers should primarily focus on the position of the exhaust. As a matter of fact, when the floor is at a high temperature, an exhaust at the top removes best the warm air. On the other hand, especially the regimes ' $T_w=293.65K, n=5.75h^{-1}$ ' and ' $T_w=293.65K, n=10h^{-1}$ ' indicate that also an exhaust at the bottom works. In this case, an exhaust at mid height removes the cold air jet before the air touches the floor. Further, a small temperature difference between the supply air and the warm floor surprisingly asks for an exhaust at the bottom. However, mind that the before-mentioned surrogate models are not accurate over the whole parameter space. One way to show this is by determining the cross validation score. At this, the predictor fits a function using a limited number of sample points, then predicts the function values for the remaining sampled inputs and, finally, determines the associated error. By way of example, Figure 12 shows the cross validation score of the surrogate model of the case 'cross ventilation concept with the floor at a higher temperature': the score of 1.25 based on 54 samples is still quite large. On the other hand, the predicted optima should be roughly correct, which was the aim put first and foremost. Creating a global accurate surrogate model actually necessitates exponentially more simulations (better known as the 'curse of dimensionality'); which is strictly prohibitive for each of the six costly five-dimensional problems presented in this paper. Yet, more simulations will corroborate these findings.

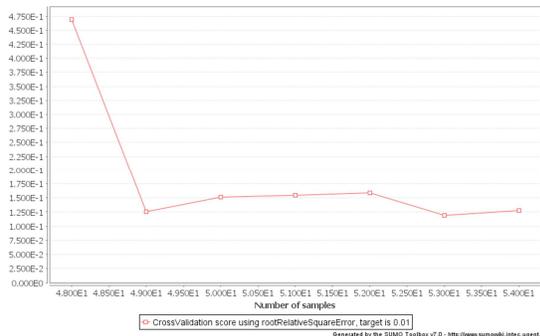


Figure 12 Cross validation score on Q_{conv} for case 'cross ventilation concept with the floor at a higher temperature'

4. Conclusion: optimizing the room/system design makes a difference

This work partly answers how to engineer best a night cooled office. To this end, it did not walk the beaten tracks. As a matter of fact, this study deployed a fully-automated

configuration of data sampling, geometry/grid generation, CFD solving and surrogate modelling. The – indeed tentative – resulting surrogate models mimicked how the convective heat flow related to several room/system design parameters, which were subdivided into ventilation concept, mass distribution, geometry and driving force for convective heat transfer. One particularly striking outcome was that night cooled cases with a massive floor clearly outperformed the ones with a massive ceiling. What's more, in that case, a well-designed cross ventilation concept surpassed single sided and under floor ventilation. The design then preferably included an air supply at the top and an exhaust at the bottom. The air supply/exhaust height was not that important. Such a design could indeed lower the nocturnal energy use of the fans and/or remove more heat. Yet, mind you, these findings do not necessarily apply to the regime by day. For example, it is perfectly possible that in general a case with thermal mass at the ceiling performs better. After all, during daytime, the predominantly natural convection heat transfer at the ceiling is usually larger than the one at the floor, certainly in case of high heat loads. Next to this, these surrogate models can advance BES modelling in two ways. The indicate profitable design solutions for which new convection correlations can be derived (possibly with the aid of CFD-based surrogate modelling). Or, derived more globally accurate surrogate models can be coupled with BES.

Acknowledgements

This work was funded by the Institute for the Promotion of Innovations through Science and Technology in Flanders (project IWT-050154).

References

COUCKUYT, I., DECLERCQ, F., DHAENE, T. & ROGIER, H. 2010. Surrogate-based infill optimization applied to electromagnetic problems. *Advances in design optimization of microwave/rf circuits and systems (special issue)*, 20, 492.

DEGROOTE, J., COUCKUYT, I., VIERENDEELS, J., SEGERS, P. & DHAENE, T. Inverse modelling of an aneurysm's stiffness using surrogate-based optimization of a three-dimensional fluid-structure interaction simulation. ECCOMAS thematic conference: CFD and optimization, 2011 Antalya. ECCOMAS.

FLUENT INC. 2006. Fluent 6.3.26. Lebanon.

GORISSEN, D., COUCKUYT, I., DEMEESTER, P., DHAENE, T. & CROMBECQ, K. 2010. A Surrogate Modeling and Adaptive Sampling Toolbox for Computer Based Design. *Journal of Machine Learning Research*, 11, 2051-2055.

JONES, D., SCHOLAU, M. & WELCH, W. J. 1998. Efficient global optimization of expensive black-box functions. *Journal of Global Optimization*, 13, 455-492.

LEMAIRE, A. D. 1993. Annex 20 - Air flow patterns within buildings. Room air and contaminant flow - evaluation of computational methods. Delft: TNO Building and Construction.

MATHWORKS 2011. Matlab. Natick.

SCHWENKE, H. 1975. Über das Verhalten ebener horizontaler Zuluftstrahlen im begrenzten Raum. *Luft- und Kältetechnik*.