VIII International Conference on Computational Methods in Marine Engineering
MARINE 2019
R. Bensow and J. Ringsberg (Eds)

AN ASSESSMENT OF MULTIFIDELITY PROCEDURES FOR SHIP HULL FORM OPTIMISATION

H.C. RAVEN, T.P. SCHOLCZ

Maritime Research Institute Netherlands (MARIN)
P.O. Box 28, 6700 AA Wageningen, Netherlands
e-mail: h.c.raven@marin.nl

Keywords: Ship design, optimisation, multifidelity

Summary: *Progress in the assessment of multifidelity techniques for ship hull form optimisation at MARIN is described. Surrogate-based optimisation is used, with response surfaces found from a combination of free-surface potential-flow and free-surface RANS computations. A substantial level of correspondence of the trends of the low and high-fidelity methods appears to be required. Applications of multifidelity stern design optimisation for a fast displacement ship and a containership are reviewed. Alternative low-fidelity formulations are considered.*

1 INTRODUCTION

Today's computational tools offer great possibilities for improving a ship's hull form design. At MARIN, RANS methods are frequently used in practical ship design projects since about the year 2000. The design steps were usually based on experience and physical insight, and then checked by computations. This is an effective procedure which is still used, but it is more and more often supplemented by a formal optimisation stage. Optimisation approaches have been introduced that fit in with this procedure, in that the setup of the design space is done based on the same experience and insight, but the optimiser permits a wider exploration and a fine-tuning of the design.

A significant part of the process to minimise the still-water resistance, for forebody wave making in particular, is normally done using a free-surface potential-flow code. The optimisation procedure used [1] is a surrogate-based method, in which the optimiser, typically a genetic algorithm, acts on a response surface found from an initial set of computations. The method is quite effective and fast, and frequently permits significant resistance reductions in just a few days work. However, for ship afterbody design a potential-flow solver is often insufficient, and free-surface RANS codes are used. But because of their computational effort and the limited time available in practical projects, normally the number of parameters varied is small, and the feeling remains that larger resistance gains would be possible.

To alleviate the computational burden of RANS-based ship afterbody optimisation, we consider multifidelity optimisation techniques. By combining a cheap low-fidelity (LF) method and a more expensive high-fidelity (HF) method (in our case, a free-surface RANS code), we hope to achieve efficiency improvements. As a LF method we have mainly used the free-surface potential flow code. Besides its large computational efficiency, this choice has additional benefits. The RANS computations just give frictional and pressure resistance, and

may not clearly indicate the origin of resistance changes. By the parallel work with a potential-flow code, we get useful additional insight in the viscous and wave resistance components and their trends. However, for stern shape optimisation we clearly have to investigate the limits of the use of the potential-flow code as a LF method; and some alternative choices are considered in this paper.

After summarising the tools used, in Section 3 we briefly recall the multifidelity formulation, in particular cokriging, and show that a considerable degree of correspondence of low and high-fidelity results is generally required to benefit from the multifidelity approach. Section 4 reviews some experiences with multifidelity methods for resistance minimisation. We briefly summarise the case from [2], stern design of a fast displacement ship, for which a substantial gain in computational efficiency was obtained. Another case is described here, for a large container ship, which turned out quite differently. Some alternative choices for LF methods are addressed in Section 5.

The research is in progress and additional results will be presented at the symposium.

2 CONTEXT AND TOOLS USED

2.1 Hull form variation

A first essential component is the shape parametrisation. As described in [1], we use a design space defined by parametric deformations of an initial hull form. Specific for the case at hand, we choose design changes considered promising, based on analysis of computations for the initial hull form. We aim at selecting independent deformation modes that are clearly related with the flow features we want to improve. In this way, we normally have a fairly small number of hull form parameters to be varied, most of them hydrodynamically relevant for the design problem. Compared with this, choosing design parameters based on just geometric properties usually leads to a much larger number of parameters.

For each deformation mode, the most deformed hull form is created in the CAD system. The hull form variations are then obtained as a parametric blending of the initial hull form and all of these deformed ones. This blending is generated in the Rhino CAD system, using proprietary plug-ins based on earlier developments [3].

2.2 Surrogate-based optimisation

For the actual optimisation we use primarily the DAKOTA optimisation toolbox [4], which offers a variety of options for response surface generation, optimisation and sensitivity analysis. Once the design space is defined, we generate a set of points spread over that space, using Design of Experiments methods. Usually we choose a Latin Hypercube Sampling, except for small dimensions of the design space. For each point, i.e. each hull form variation defined by the parameters, we run the flow code for all the conditions (e.g. speeds, drafts) to be considered. We feed the results to the Surrogate-Based Global Optimisation algorithm, which for each flow condition generates a response surface based on the computed results; and performs a global optimisation using a genetic algorithm, acting on the response surfaces only. This is a very quick process, normally less than an minute. When conflicting design objectives are present, the multiobjective optimisation problem results in an estimated Pareto front. A selection of points on the front is again evaluated using the flow code to obtain the

true results. If needed, response surfaces can then be updated using the new data points, and the optimisation repeated. The approach, used routinely now in combination with the potential-flow solver, is most effective; in particular as the choice of design parameters that we make often leads to well-defined response surfaces.

2.3 Free-surface potential-flow code

The potential-flow code used is RAPID [5], a panel code that iteratively solves the steady nonlinear free-surface flow problem. The complete inviscid free-surface conditions are imposed, and the dynamic trim and sinkage are incorporated. The computation time is just a few minutes on a desktop PC. Validations have shown that the wave pattern calculated is quite accurate, except for stern waves: those are usually overestimated due to the neglect of viscous effects, the more so for fuller hull forms and lower Froude numbers. Also, for transom sterns only a dry-transom flow can be well modelled.

2.4 Free-surface RANS codes

Two different free-surface RANS codes have been used in the work so far. PARNASSOS [6] is a free-surface fitting code, working on a multiblock structured mesh that deforms to fit the free surface. The steady RANS equations are solved by a multiple-sweep marching iteration, solving the fully coupled momentum and continuity equations. Also the free surface is solved by iteration, using a particular formulation of the steady free-surface boundary conditions [7]. The code is most efficient, but may have occasional robustness issues. Also it is not suitable for large-scale parallellisation. In optimisation this is no large disadvantage though, as on a cluster many computations (e.g. the Design of Experiments) are done simultaneously.

REFRESCO [8,9] is an unstructured-mesh finite-volume code, using free-surface capturing by a volume-of-fluid formulation. The unsteady momentum and pressure equations are solved by a SIMPLE-type algorithm. It is far more general than PARNASSOS, and more robust, but also more time-consuming. In addition, the lower inherent accuracy of the unstructured-mesh discretisation asks for much denser grids. Therefore, if hundreds of computations would need to be done in a practical optimisation, the computational effort would be a disadvantage.

3 A TEST OF MULTIFIDELITY SURROGATES

Suppose we have a set of N_{LF} data points in the design space with results from a low-fidelity method; and for a part of these points, a smaller number N_{HF} , we also have results from a high-fidelity method. We want to exploit both sets to make the best approximation of the high-fidelity function F_{HF} over the design space.

We use a Cokriging formulation, see e.g. [10], defining

$$F_{HF}(\boldsymbol{par}) = \rho . F_{LF}(\boldsymbol{par}) + \Delta F (\boldsymbol{par})$$

in which ***par*** is the vector of parameter values. In the implementation used, first a Kriging approximation of the low-fidelity data set is made. Next, another Kriging approximation is made of the difference ΔF ***(par),*** along with determination of the scale factor ρ.The idea is that ΔF is a smaller quantity with a simpler distribution, therefore it should be better approximated by the limited data set of N_{HF} values ΔF.

Forrester et al [10] show some examples, among which a 1-parameter case:

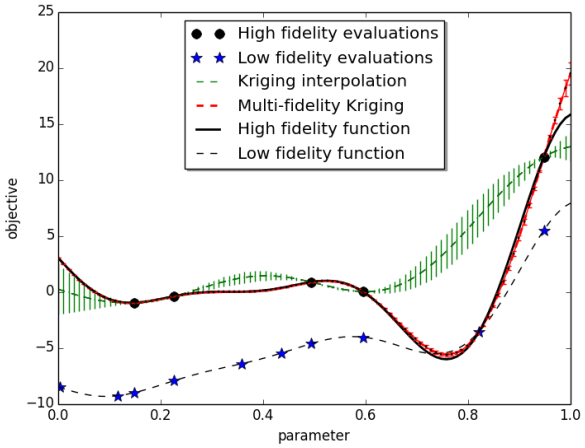

Figure 1 Forrester example, with LF and HF functions, data points, and approximations by Kriging and Cokriging.

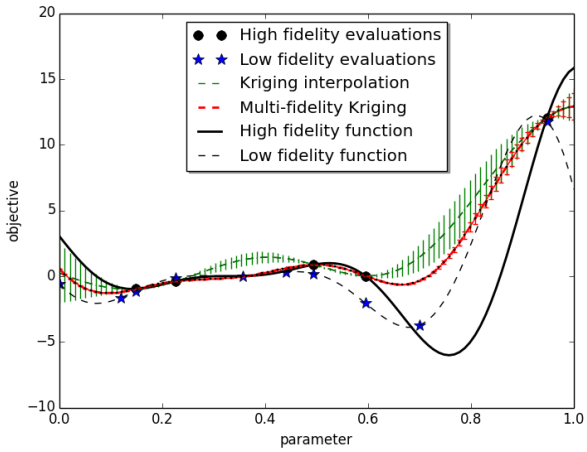

Figure 2 Same with modified LF function.

$F_{HF} = (6x-2)^2 \sin(12x-4)$ $\quad$ $F_{LF} = 0.5\, F_{HF} + 10(x-0.5)$, for $0<x<1$.

Fig. 1 shows there is a significant difference between these functions. Kriging on just the 5 HF points (green line with error bars) misses the minimum altogether. Still, with just 10 LF points added, cokriging approximates the high-fidelity function extremely well. The uncertainty implied by the Kriging approximations, shown as 'error bars', is also much reduced for cokriging. However, the actual error is 4-5 times larger than this uncertainty, both for Kriging and for Cokriging. We remark that this uncertainty is used in some 'adaptive sampling' procedures, in which points are added where maximum improvement is expected; and a better indication of the actual error would be desirable.

However, in this example, for $\rho=2$ it follows $\Delta F = -20x+10$, just a linear function of x without any Kriging approximation involved. We can even choose $-F_{LF}$ as the LF function, with trend completely opposite to F_{HF}, and the Cokriging approximation is still perfect. The example fits the Cokriging formulation nicely but is not general at all. As Forrester et al state, this example is 'somewhat contrived', nevertheless it is copied in several other papers, not always indicating its limited validity.

So for the same $F_{HF} = (6x-2)^2 \sin(12x-4)$ let us choose another LF function,
$F_{LF} = (6x-2)^2 \sin(12x-3)$, which has a small shift of the minimum ---
a most probable situation in practice. Fig. 2 shows that, notwithstanding its similar shape and values, the LF function hardly helps to construct an approximation of the HF function: Cokriging is just marginally better than simple Kriging on the 5 HF points. While the error is much larger than in Fig. 1, the uncertainty is not; therefore it is hardly indicative of the error, and of the location where the largest errors occur. Other small changes to the LF function can even make single-fidelity Kriging more accurate than Cokriging. The 1D example in [10] thus is not representative of what can be expected in case the LF function has its minimum slightly shifted or has other general deviations from the HF function.

This simple example illustrates that in general a significant degree of correspondence between HF and LF functions is needed, such that the difference function ΔF is simpler to model than the HF function itself [10]. Moreover, the relation of the indicated uncertainty and the actual

error is weak and variable; the error here being much larger.

4 MULTIFIDELITY OPTIMISATION COMBINING POTENTIAL FLOW AND RANS CODES

The multifidelity formulations that we have applied mostly used the free-surface potential flow code as the LF method. This method tends to overestimate the stern wave system, which in reality is reduced by viscous effects in a not easily predicted way, and to a different degree for different wave components. The resistance trends for stern variations as predicted by RAPID might thus be different from those from a RANS code, making this a critical test for a multifidelity formulation.

4.1 Stern optimisation for a fast displacement ship

We briefly summarise here the results for the test case from [2], optimisation of the stern of a fast displacement vessel, for Fr = 0.27 and 0.37. We introduced 3 parameters: two stern buttock shape modifications, one parameter changing the deadrise at the stern.

For the LF method, a Design of Experiments of 150 hull forms was generated in this 3D parameter space. For each hull form, computations have been made for both speeds using RAPID. This DoE size surely is an overkill, but could be run overnight on a single PC nonetheless. Total resistance was estimated by adding (1+k) times flat plate friction to the computed wave resistance. Next, an overlapping smaller DoE of 32 hull forms was evaluated by PARNASSOS, for both speeds.

The LF and HF resistance values appeared well correlated, the correlation lines having a

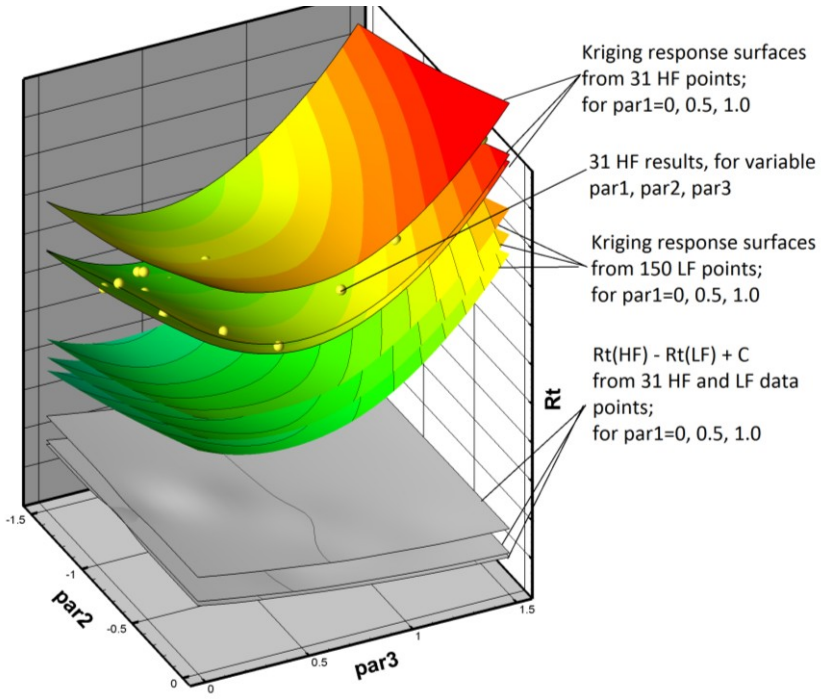

Figure 3 Response surfaces for fast displacement ship. Upper 3 surfaces are for HF method, middle 3 surfaces for LF, lower 3 surfaces show their difference (but increased by a constant to show in the same graph). Each plane is for a constant *par1*.

slope near 1.0. Fig. 3 illustrates that response surfaces, derived separately from the LF and HF computations, have essentially the same shape. Their difference ΔR_T ***(par)***, here assuming that $\rho = 1$, thus is a very simple function of the parameters, as shown in Fig. 3. Instead of using all 31 HF data, a much lower number would sufficiently define it. In [2], it has been checked that the response surfaces for the multifidelity formulations actually did not change much if just 7 HF data points were used in this 3D space; both from visual inspection and from the errors in a set of test points. Instead, a simple Kriging of only the HF data lost more of its accuracy at this reduction of the number of HF data. Subsequently, surrogate-based optimisation was done using these response surfaces; and here again, the multifidelity methods well indicated the optimal hull form even with just 7 HF data points.

In this way, the optimisation of an already refined design indicated some subtle stern shape modifications that reduced the total resistance by 2.6% and 3.6% at the two speeds. However, already the LF method indicated the optimum design and the achievable stern wave reduction very well. To really test the capabilities of MF optimisation, a more challenging case was taken up next.

4.2 Optimisation of a containership stern

The next test case in our search for the limit of multifidelity optimisation is a large containership, at $Fr = 0.20$. The somewhat larger block coefficient and the lower Froude number cause larger viscous effects on the wave making here. The potential flow code predicts a substantially stronger transverse stern wave component than the RANS code, while the diverging components are comparable.

Four parameters are defined for the stern shape modifications. Two of them introduce a slight S-shape to the stern buttock lines, long and short respectively. The 3rd parameter lifts the transom, while the 4th modifies the transom V-shape, decreasing and increasing stern deadrise. A Latin Hypercube Sampling of 240 hull forms is evaluated using RAPID as LF code, followed by surrogate-based global optimisation. The optimum hull form from just the LF method has maximum S-shape of the buttocks, minimum V-shape of the sections, promising 6.2% total resistance reduction owing to a significant reduction of the transverse stern wave (Fig. 8).

Next, a set of 40 hull forms is evaluated using the HF-code (PARNASSOS). Response surfaces are generated for both HF and LF data, and one view is shown in Fig.4. Clearly, they are very different, those from LF having a clear and strong trend and simple shape, those from HF substantially more convoluted. Fig.5 shows the degree of correlation of LF and HF resistance values; a cloud of points without much direction, but showing 8% of variation from LF, 2 % from HF. It seems doubtful whether multifidelity will work in this case.

Upon a detailed inspection of the computed flow fields, the transom flow regime, wetted or dry, is found to play an important role here. The stern variations are such that it is not clear a priori what type of transom flow will be obtained, until it is found from the RANS solutions. In the LF method however, a dry-transom flow was enforced in all cases, since a wetted transom flow cannot be properly modelled by a potential flow. If now we distinguish both transom flow regimes in Fig. 5, we note that the dry-transom results do display some correlation with the LF results: low resistance in RAPID usually means low resistance in PARNASSOS, although the variations are smaller: for these cases we are actually reducing

transverse stern waves, but as these are smaller due to viscous effects, the trend is weakened. For the wetted transom flow cases though, RAPID assumes a dry transom and predicts a high resistance for most, but some have actually a low resistance.

This change in the transom flow regime is not easily anticipated from just hull shape parameters. While in most views of the design space the occurrence of wetted or dry transoms seems arbitrarily spread, the exception is Fig.6, which shows that lower values of the S-shape parameter and higher values of the V-shape tend to promote a wetted transom.

Response surfaces from the HF data are then derived separately for both types of transom flow, and are shown in Fig.7, now based on an increased set of 77 HF data points. We observe that design trends more or less reverse in the middle of the design space at the change of the transom flow regime, for these 2 parameters. Also we observe that the wetted-transom cases actually have a slightly lower resistance. The initial design has a wetted transom and thereby is rather favourable. Because the design space has been set up with trends for dry-transom flow in mind, it is hard to make any gains without redefining the design space. This is why most points in Fig.5 are worse than the initial design. The lowest resistance found is 0.7% lower than that of the original. Fig.8 shows the substantial reduction of transverse stern waves promised by the potential-flow code, and the small reduction actually achieved: we have gone from a wetted transom to a dry transom, which in itself would increase the stern wave amplitude, and which here cancels most of the promised reduction. While the viscous losses in the recirculating flow aft of the transom have been eliminated, the net gain is small.

To optimise the hull form in this situation, a multifidelity formulation with a potential-flow code as the LF method seems inadequate. Either another LF formulation needs to be used, or we simply optimise using the HF method. A much larger number of HF computations would then be needed; not only because of the more complicated response in the present case, but also because the different design trends for wet and dry transoms would necessitate introducing a larger number of design parameters to find the optimum design.

4.3 Practical ship design project

We mention the approach we have used in a practical project to illustrate how multifidelity methods can also be applied somewhat informally. It concerned a slender vessel at a rather high speed, with significant stern wave making. A dry-transom flow was found from an initial RANS computation using REFRESCO. In a first stage, RAPID has been used to estimate design trends and useful modification directions. This helped to preselect relevant parameters to optimise. For two stern design parameters for which the validity of the potential-flow code was less certain, a set of 13 RANS computations was carried out. The response surface derived from it was found to be somewhat complex in shape, and a set of 75 potential-flow computations was made in addition to clarify the picture. A fair agreement of the trends from the LF and HF code could then be observed, and a multifidelity response surface was derived and used to locate the optimum. While the multifidelity approach was no breakthrough, the complementary use of both tools, and exploiting the correspondence of trends, led to a more reliable answer than if just these HF data would have been available.

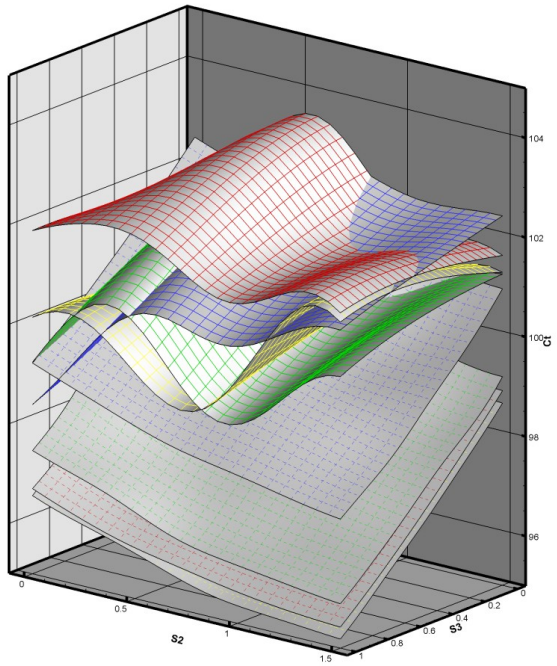

Figure 4 Response surfaces for containership case. Upper surfaces are HF results, lower are LF. Surfaces have constant V-shape parameter.

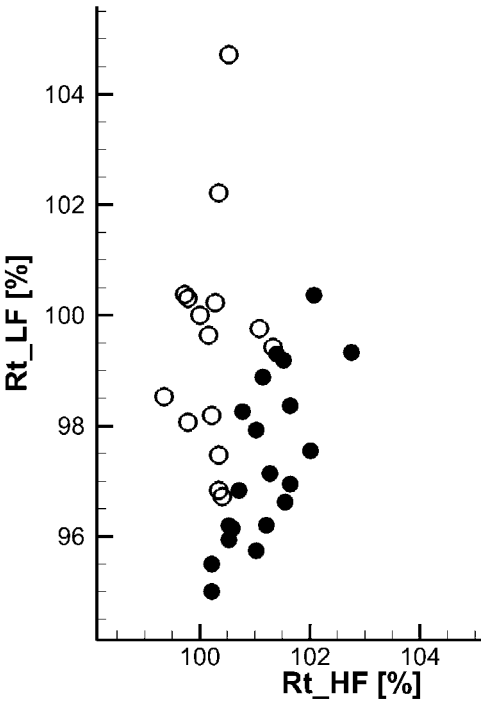

Figure 5 Correlation of HF and LF resistance values. Solid markers: dry transom. Open markers: wetted transom.

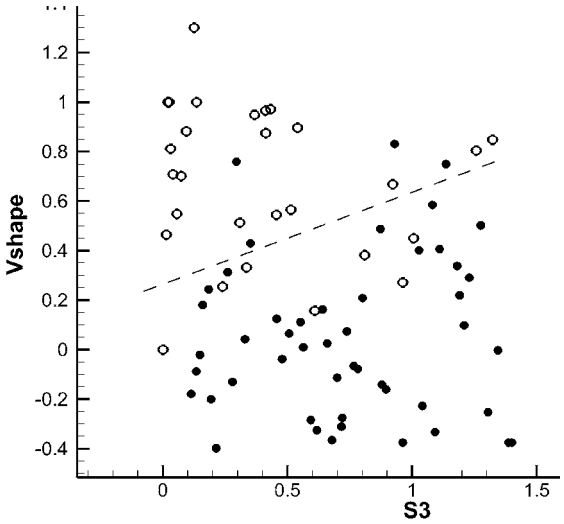

Figure 6 Dry (solid markers) and wetted (open markers) transom flows against parameters S3 and V-shape.

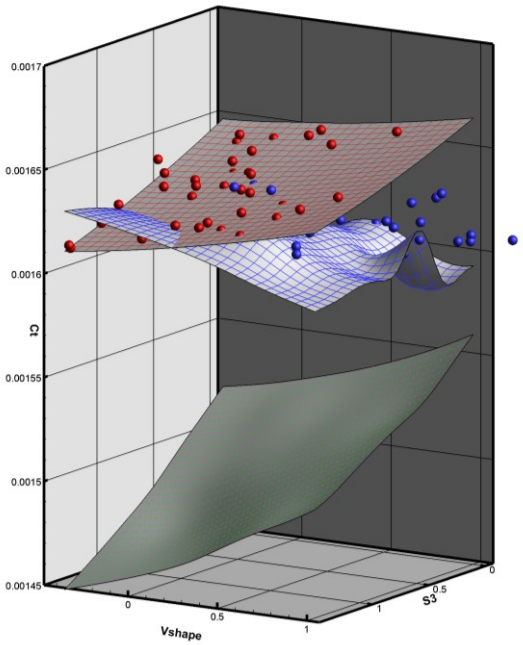

Figure 7 Response surfaces against parameters 1 and 4, for fixed par2 and par3. Red mesh and markers: dry transom. Blue: wetted transom. Lower plane is from LF method.

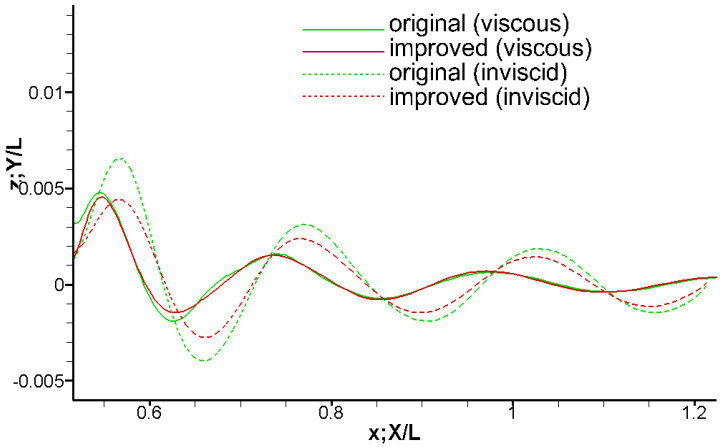

Figure 8 Wave cut along centreline aft of transom. Dashed lines from LF, solid lines from HF method.

5. ALTERNATIVE LOW-FIDELITY METHODS

The free-surface potential flow method has the clear advantage of a very low computation time, little setup time and a clear physical meaning of the results. Therefore, we will surely use it for a class of ships for which correspondence of trends may be expected. However, as expected, it has limitations. Alternative low-fidelity methods can be envisaged.

5.1 Linearised free surface solutions

A possible alternative LF method has been mentioned in [2]. In free-surface computations using PARNASSOS, the wave surface is iteratively updated, and in each iteration, steady free-surface boundary conditions linearised in the wave elevation updates are imposed. If computations for all variations are started from the final result of the original design, the linearised solution (1st iteration) may already be accurate enough, at least to indicate the design trends. The advantage is a significant reduction of the computation time, and an improved robustness as no grid deformations are required.

For the case of Section 4.1, the resistance values so obtained were well correlated with the HF data; and this linearised solution could well serve as a LF method. For the containership of Section 4.2, for the majority of the hull forms the linearised solution already indicated whether the transom would be dry or wetted; but nonetheless, the correlation of resistance values was comparable to that in Fig. 5. This does not offer a suitable alternative for that case.

5.2 Coarse-grid RANS computations using PARNASSOS

While a sufficient grid density in the RANS computations is needed for an accurate prediction, coarser-grid solutions may still predict the trends of resistance correctly; as correlated bias errors in different solutions drop out. In that case, a set of dense-grid solutions may be supplemented by a larger number of coarse-grid computations in a multi-fidelity formulation. At equal total computational cost this might yield better accuracy than just a larger set of dense-grid computations.

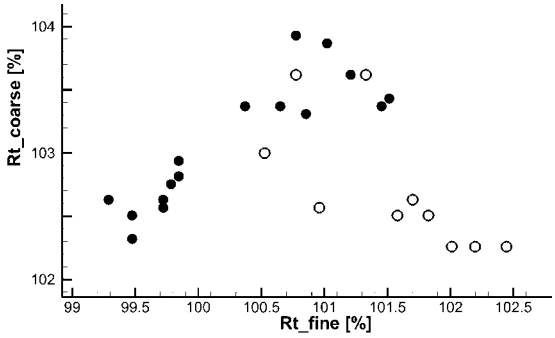

Figure 9 Correlation of LF and HF resistance values. Solid markers: equal transom flow type. Open markers: different transom flow type.

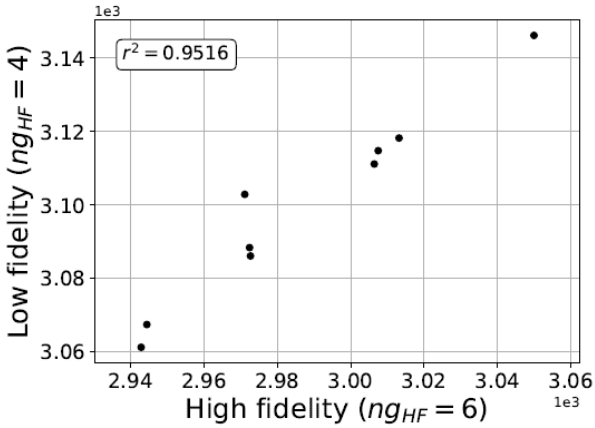

Figure 10 Correlation of coarse and fine grid resistance values. Refresco study.

For the same container ship afterbody considered in Section 4.2, we try to use a coarse-grid RANS solution as a LF method. The grid spacing in longitudinal and transverse directions is doubled, yielding a coarse grid with just 1.7 M cells, but still having full resolution of the boundary layer.

Fig.9 shows the correlation of the coarse and fine grid resistance coefficients. Taking all points together the correlation is poor. However, if we just consider the cases that have the same transom flow regime on coarse and fine grids, the correlation is reasonable; both for dry and wetted transoms. But there is a number of cases that have a dry transom on the fine grid, but wetted on the coarse grid; and again, the change of flow regime destroys the correlation, and thereby, the practicality of the coarse-grid solution. Apparently, the coarse grid here is at least locally too coarse for this purpose.

5.3 Coarse-grid RANS computations using ReFRESCO

A similar test has been done for the case of another large containership, for Froude number 0.185. Two parameters are used to modify the stern, one changing the level of the transom edge, the other the V-shape at the stern. Here all stern flows were of dry-transom type.

The RANS code used here is REFRESCO, and unstructured grids are generated using HEXPRESS. While coarse grids can be generated in a variety of ways, it is considered best to aim for some overall geometrical similarity, by coarsening the initial grid in each direction while maintaining the surface refinement levels [9]. As high-fidelity data, results on a grid of 10.8 million cells are used. For low-fidelity, grids with 4.6 million cells are adopted. Mean computation time on the coarse grids is found to be 32% of that on the fine grids.

Fig.10 illustrates the correlation of resistance values found on both grids, for a 9-point full factorial DoE. There is a reasonable correlation ($R^2 = 0.95$), the variation on the finer grids is about 1.5 times as large as on the coarse grids, and the resistance level is some 4% lower. The coarse-grid results themselves are clearly less accurate, but they may well be helpful to determine the trends.

Toal [11] has given guidelines for when a MF Kriging method is preferred above single-fidelity Kriging:

- The correlation between the low and high-fidelity solver is high: $R^2 > 0.9$;
- Between 10% and 80% of the total computing time is spent to LF computations;
- The fraction of HF data replaced by LF data should be $f_r > 1.75/(1+1/C_r)$; in which C_r is the ratio of computation times (here 0.32).

Following this we replace 5 of the 9 HF evaluations by LF evaluations, keeping the total computational cost roughly equal; i.e. we do 17 LF and 4 HF computations. An accuracy test using a set of 12 additional HF test points was then done (Fig.11). Disappointingly, there is no clear improvement from the use of the MF method. The RMS error in the test points is slightly larger than for SF Kriging. As the overall computation time is roughly equal, this is not helpful. Possibly, for larger grid density ratio's better results could be obtained.

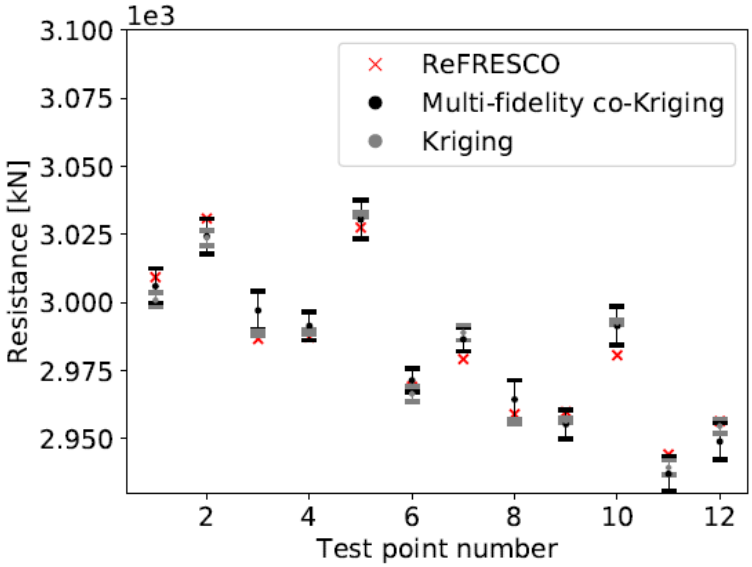

Figure 11 Validation of response surfaces from single and multi-fidelity Kriging. Error bars indicate 2σ confidence interval.

6 CONCLUSIONS

This paper has outlined our progress in the study of possibilities to accelerate and enhance ship afterbody design optimisation using multifidelity formulations. Main points are summarised as follows.

- Surrogate-based global optimisation is being used in ship hull form design at MARIN. For forebody wave making, this is based on just the free-surface potential-flow code RAPID. For afterbody design, evaluation by free-surface RANS codes is required.
- To reduce the computational effort of such a RANS-based optimisation, multifidelity methods have been tested, combining potential flow and RANS codes to determine the response surfaces. For a fast displacement vessel this was found successful and resulted in a significant reduction of the number of high-fidelity evaluations. For a containership however, the correlation between low and high-fidelity resistance values was found to be poor, as a result of a change of the transom flow regime in the middle of the design space. Wetted transoms appeared to have significantly different trends with the design parameters than dry transoms.
- Using coarse-grid RANS solutions as low-fidelity simulations is a possibility, but finding

the right balance between a good correlation with the HF resistance trends, and a large reduction of the computation time, proved not easy.

- Both an additive multifidelity formulation and cokriging require a large degree of correspondence of the HF and LF trends to really contribute to the accuracy of the response surface. In some 1D examples the uncertainty connected with the Kriging approximation appeared not well correlated with the actual error, and hardly suitable to be used for adaptive sampling.
- While for a class of cases we already use multifidelity optimisation incidentally in practical projects, further developments are required to apply it in daily practice.

ACKNOWLEDGEMENT

This research has been partly funded from the TKI-allowance of the Dutch Ministry of Economic Affairs.

REFERENCES

[1] Raven, H.C., and Scholcz, T.P., Wave resistance minimisation in practical ship design. *Proc. MARINE 2017*, Nantes, France.

[2] Raven, H.C., Minimising ship afterbody wave making using multifidelity techniques. *32nd Symp. Naval Hydrodynamics*, Hamburg, Germany, Aug. 2018.

[3] Hoekstra, M., and Raven, H.C., A practical approach to constrained hydrodynamic optimization of ships. *NAV 2003 Symp.*, Palermo, Italy, June 2003.

[4] Adams, B.M., Bauman, L.E., Bohnhoff, W.J., Dalbey, K.R., Ebeida, M.S., Eddy, J.P., Eldred, M.S., Hough, P.D., Hu, K.T., Jakeman, J.D., Stephens, J.A., Swiler, L.P., Vigil, D.M., and Wildey, T.M., "Dakota, A Multilevel Parallel Object-Oriented Framework for Design Optimization, Parameter Estimation, Uncertainty Quantification, and Sensitivity Analysis: Version 6.0 User's Manual," *Sandia Technical Report SAND2014-4633*, July 2014. Updated November 2015.

[5] Raven, H.C., A solution method for the nonlinear ship wave resistance problem, *PhD Thesis, Delft Univ. Techn.*, 1996.

[6] Hoekstra, M., Numerical simulation of ship stern flows with a space-marching Navier-Stokes method, *PhD Thesis, Delft Univ. Techn*., Netherlands, 1999

[7] Raven, H.C., van der Ploeg, A. and Starke, A.R., "Computation of free-surface viscous flows at model and full scale by a steady iterative approach", *Proceedings 25th Symposium on Naval Hydrodynamics,* St. John's, Newfoundland, Canada, August 2004.

[8] Klaij C. M., Hoekstra, M, and Vaz, G., Design, analysis and verification of a volume-of-fluid model with interface-capturing scheme. *Computers and Fluids*, 2018

[9] P. Crepier, "Ship resistance prediction: Verification and validation exercise on unstructured grids", *MARINE 2017 symposium*, Nantes, 2017.

[10] Forrester, A.I.J., Sóbester, A. and Keane, A.J., "Multi-fidelity optimization via surrogate modelling", *Proceedings Royal Society A*, Dec. 2007.

[11] D. J. Toal, Some considerations regarding the use of multi-fidelity kriging in the construction of surrogate models, *Struct. Multidiscip. Optim.* 51 (6) (2015) 1223{1245. doi:10.1007/s00158-014-1209-5. URL http://dx.doi.org/10.1007/s00158-014-1209-5