Taylor & Francis
Taylor & Francis Group

Ship Technology Research
Schiffstechnik

ISSN: (Print) (Online) Journal homepage: https://www.tandfonline.com/loi/ystr20

Practical ship afterbody optimization by multifidelity techniques

Hoyte C. Raven & Joy Klinkenberg

To cite this article: Hoyte C. Raven & Joy Klinkenberg (21 Nov 2023): Practical ship afterbody optimization by multifidelity techniques, Ship Technology Research, DOI: 10.1080/09377255.2023.2275371

To link to this article: https://doi.org/10.1080/09377255.2023.2275371

© 2023 Maritime Research Institute Netherlands. Published by Informa UK Limited, trading as Taylor & Francis Group

Published online: 21 Nov 2023.

Submit your article to this journal

View related articles

View Crossmark data

Full Terms & Conditions of access and use can be found at
https://www.tandfonline.com/action/journalInformation?journalCode=ystr20

SHIP TECHNOLOGY RESEARCH (SCHIFFSTECHNIK)
https://doi.org/10.1080/09377255.2023.2275371

Taylor & Francis
Taylor & Francis Group

OPEN ACCESS Check for updates

Practical ship afterbody optimization by multifidelity techniques

Hoyte C. Raven* and Joy Klinkenberg

Maritime Research Institute Netherlands (MARIN), Wageningen, Netherlands

ABSTRACT
This paper discusses multifidelity methods for CFD-based optimization of ship afterbody designs; aimed at a fast application to a variety of practical cases. Surrogate-based global optimization is used, using a multi-objective genetic algorithm. The surrogates are derived by combining few high-fidelity computations, by free-surface RANS codes, with many low-fidelity computations. For rather slender vessels for which the wave resistance variations over the design space are dominant, a free-surface potential flow code is found very effective as a low-fidelity solver, permitting a large reduction of the number of RANS computations and the associated cost. Examples are shown for model 5415 and a fast displacement vessel. For cases with variation in both viscous and wave resistance, an alternative method is used combining coarse and fine-grid RANS computations; the coarse-grid ones being about 20 times cheaper. Application of this coarse/fine grid multifidelity optimization to a containership and a motor yacht shows its effectivity.

ARTICLE HISTORY
Received 15 February 2023
Accepted 6 August 2023

KEYWORDS
Ship afterbody design; CFD-based optimization; multifidelity; response surface; surrogate-based optimization; coarse/fine grid; potential/RANS; ship resistance

1. Introduction

The detailed design of a ship hull form has a large influence on its resistance and engine power at a given speed. In the past, the hull form design was based on experience and insight; and often supported by model testing some design variations. Mainly since the 1990s, a computational assessment of the hull form has become a standard component of the design process; initially using free-surface potential flow codes, subsequently also Reynolds-Averaged Navier-Stokes (RANS) codes. Both are now being used routinely in ship design. From analysis of the computed flow field and wave pattern, desired modifications of the design are derived, and in a few steps improvements of the design are usually achieved.

Today, also CFD-based hull form optimization methods are being used increasingly in ship design. These can have important benefits for fine-tuning beyond what can be achieved in a few designer-directed steps, and for finding the optimal trade-off between several objectives or design points. There are many recent publications on such optimization methods. However, methods proposed often do not fit in well with the demands of practical projects in our practice; as extensive preceding work is often needed for the parametrization of the hull form or for prescribing constraints, and the computational stage may also be time-consuming.

In our practice we have to deal with a large variety of ship types; for each project there is just little time available, perhaps 2 weeks at most for the computational work; and extensive design knowhow is available which we want not to replace, but to incorporate in and supplement by optimization procedures. Therefore, the methods we have developed differ in some regards from other published approaches. We use completely general parametric deformations of an initial hull form, we involve designer-selected hull form parameters that directly address flow properties of interest, and we pay much attention to the required computational effort and wall-clock time for the entire process.

The present paper considers the optimization of a ship's afterbody design, aiming at minimum resistance. The design of the afterbody is important as it is responsible for the viscous pressure resistance and often a substantial part of the wave resistance. Both resistance components are quite variable and strongly dependent on hull form details. Furthermore, the stern wave making and viscous flow interact, in a way that is roughly, but not precisely, known. Therefore, most likely there is often room for further improvement of ship afterbody design, and a detailed and accurate optimization procedure would be an asset. To reduce the computational effort of such an optimization we have investigated multifidelity optimization methods.

The paper is set up as follows. Section 2 discusses the hull form parametrization. After a brief description of

CONTACT Hoyte C. Raven raven032@planet.nl Maritime Research Institute Netherlands (MARIN), P. O. Box 28, 6700AA Wageningen, Netherlands
*Retired from MARIN.

© 2023 Maritime Research Institute Netherlands. Published by Informa UK Limited, trading as Taylor & Francis Group
This is an Open Access article distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivatives License (http://creativecommons.org/licenses/bync-nd/4.0/), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the original work is properly cited, and is not altered, transformed, or built upon in any way.

the computational tools in Section 3, the optimization method is described in Section 4. Then we discuss multifidelity (MF) methods (Section 5), and the two classes of MF methods we use: one combining a free-surface potential-flow code with a RANS code (Section 6), and one using coarse and fine-grid RANS computations (Section 7). Examples of application to practical cases are shown. After a discussion in Section 8, conclusions are summarized.

2. Hull form parametrization

A principal component of a hull form optimization system is a parametrization of the hull form variations. This plays a crucial role, determining the relevance of the design space, the required number of parameters, the need to specify constraints, the feasibility and efficiency of surrogate-based techniques, and the complexity of the response surfaces. In general terms, we want to have complete flexibility to define specific shape variations by as few parameters as possible; while ideally each member of the family of hull forms is an acceptable, faired shape satisfying some basic constraints.

A variety of hull form parametrization methods has been proposed. In several papers, a full parametric representation of the entire hull form is used, and a selection of the parameters is varied in the optimization. In our practice, usually, an initial hull form is proposed by a yard, and is to be improved by detailed modifications. This would require to first reproduce the initial form by a parametric one, and then choose the parameters to be varied. However, these parametric models often appear to involve a large number of parameters. For example, Han et al. (2012) needed 80 parameters just to describe the forebody. Feng et al. (2021) describe a significant amount of work to represent some containerships by parametric models. Next, sorting out which of these parameters to vary seems a large amount of work; not practical for the variety of cases we have to address.

The alternative is parametric deformations of an initial hull. Various methods have been proposed also here; for example, (modified) Lackenby shifts (Kim and Yang 2010; Han et al. 2012), additive polynomial patches (Chun 2010), Radial Basis Functions (Kim and Yang 2010) or Free-Form Deformations (Brizzolara et al. 2015). In general, these methods have the disadvantage of being unfamiliar to a designer, and they often lack flexibility.

Instead, a 'parametric blending' as proposed in (Hoekstra and Raven 2003) is used at MARIN successfully since 20 years. We start from the NURBS representation of the initial hull form in a Computer-Aided Design (CAD) system. If the designer acknowledges a certain variation mode as promising for the case at hand, a modified hull form is created having the maximum deformation to be considered for that mode. This is simply defined in the CAD system in the familiar way, starting from the original and moving NURBS control points. This defines a 1-parameter family of hull forms, given by the control points moved by a fraction $0 \leq \mathrm{par}_1 \leq 1$ from the original to the modified position. The same is done for other deformation modes of interest, leading to n basis hull forms in addition to the original, and thereby an n-dimensional family of hull forms given by the parameters $\mathrm{par}_1 \ldots \mathrm{par}_n$. Immediate visual inspection of all these hull forms can be done in the CAD system.

This method was originally developed as part of the CAD system GMS, and has later been reimplemented in RHINO (as proprietary plugins). It has been used both for optimization and for systematic hull form variations, in the RapidExplorer and ParnassosExplorer systems (Raven 2022). It has several advantages:

- There is complete flexibility in the deformations selected.
- Specific desired hull form changes can be achieved with just 1 parameter, rather than several.
- All hulls are smooth and faired, and implicitly constraints can already be incorporated in the definition of the deformed basis hull forms; such as maintaining a flat of side, flat of bottom, or hard points. Also, the basis hull forms can be defined so as to have approximately the right displacement and Longitudinal Centre of Buoyancy (LCB), such that at most a small final adjustment is needed after the optimization.
- The method is entirely familiar to CAD engineers, who create a deformed basis hull in less than an hour.

The variations applied to the cases discussed in this paper, while sometimes apparently similar, still depend too much on the particular case to be usefully automated. The limited amount of manual CAD work inherent to our approach is therefore acceptable.

With this ability to define whatever parametric deformations, how to select them? Here we aim at choosing deformation modes that directly affect the physical effects of interest for the case at hand. If in an initial computation for the original design, we observe an undesirable flow feature, we try to address that with physically relevant hull form changes, and choose these as the basic variation modes; perhaps supplemented by a few more general variations. This limits the dimensionality of the design space and leads to clear trends and relatively simple response surfaces, suitable for the optimization methods we will apply. Thus, existing design and hydrodynamic know-how can be fully exploited in the definition of the optimization space; and we base ourselves on that as far as we can. Instead, if fixed, geometrically-

oriented parameters are used, as in many of the mentioned alternative parametrizations, these would often produce the desired physical effect only in combination; asking for variation of many parameters, many computations and yielding unclear trends.

3. Computational tools and objectives

Three computational tools are used in this paper, and they are briefly described here for completeness.

3.1 RAPID

RAPID (Raven 1992, 1996) is a free-surface potential-flow code for computing the wave pattern and wave resistance of a ship. It uses a Rankine source panel distribution on the hull and on a (wavy) plane at a small distance above the wave surface. The full nonlinear inviscid free-surface boundary conditions are imposed. The problem is solved by iteration, in which the wave surface is updated and the dynamic trim and sinkage of the hull are adjusted until convergence. A computation normally takes 1–10 min on a desktop PC, depending on the Froude number. The code is used on a large scale at MARIN and elsewhere, without interruption since 1994.

As a potential flow is supposed, viscous effects are neglected. Therefore, while the computed wave pattern and flow are accurate along the greater part of the hull, near the stern deviations occur, in particular for fuller hull forms and wetted transom stern flows. Consistently with the potential-flow assumption, a dry-transom flow is always supposed.

In some other papers, free-surface potential flow codes were found to perform poorly in optimization (Grigoropoulos et al. 2017; Liu et al. 2022). However, these were methods imposing linearized free-surface boundary conditions. The linearization leads to a substantial loss of accuracy in wave pattern, transom stern flow and resistance, so their judgement does not apply to the present code.

3.2 PARNASSOS

Parnassos (Hoekstra 1999; Van der Ploeg et al. 2000) is a RANS code for computing the viscous flow around ship hulls. The steady RANS equations are discretized on a structured multiblock mesh, and solved in fully coupled form by a multiple-sweep marching iteration. This gives the code some particular properties, e.g. computing a full-scale flow without wall functions is straightforward. Today it is used routinely for calculating the viscous resistance and form factor for all hull forms tested at MARIN.

For free-surface flows, a free-surface fitting approach is used. The steady problem is solved by iteration instead of time stepping. This requires that the free-surface boundary conditions are cast in a particular form (Raven and Van Brummelen 1999; Van Brummelen et al. 2001). The resulting method is accurate and extremely efficient in computation time. However, due to occasional robustness problems of the iterative process in practical applications, for free-surface flows the use of Parnassos is less common today, and Refresco is used instead, even though the required computation time is at least an order of magnitude larger.

3.3 REFRESCO

Refresco (Crepier 2017; Pereira et al. 2017) is a more general finite-volume RANS code, computing steady or unsteady viscous flows. The momentum and pressure correction equations are discretized on an unstructured mesh, and solved in segregated form by a SIMPLE-type solution algorithm. Typically, grids are generated using the HEXPRESS tool. Free-surface capturing by a Volume-of-Fluid method is used, asking for a time-stepping solution process. Computations are done on a high-performance computing cluster with large-scale parallellization. The code is used extensively for a variety of practical applications.

4. Optimization method

In the first publications on ship hull form optimization, the optimization algorithm was usually coupled directly with the CFD code: the search algorithm prescribed the next parameter values to be evaluated, and the CFD code computed the corresponding objective. This made the effectivity of the optimizer crucial for the effort. Many different optimization algorithms have been proposed. One distinction is between local and global optimizers. Both have been applied in ship hull form design, sometimes in combination.

Because of the possible existence of local minima and the smaller sensitivity to any numerical noise in the computed objectives, we favour the use of a global optimization method. In addition, such methods are effective in finding Pareto-optimal solutions for multiple objectives. However, global methods are invariably less efficient than local optimizers. For example, Brizzolara et al. (2015) use a genetic algorithm in a 20-parameter optimization and need to evaluate some 7000 hull forms for convergence; while Kim and Yang (2010) compute about 2000 hull forms for a 9-parameter optimization. Chun (2010) compares an SQP method with some different Particle Swarm methods, finding that the best of the latter is acceptably efficient. Still, a direct coupling of a CFD code with a global optimizer often results in needlessly overresolving the variation of the objectives in the design space.

Today, therefore, surrogate-based methods are more popular. In these, the objective function is

interpolated or approximated algebraically by a 'response surface'. In the optimization process, each objective evaluation is then just an algebraic operation rather than running a CFD code, giving a vast reduction of the computational effort. Peri and Campana (2008) partially used surrogate models in their early multifidelity methods. Kim et al. (2011) apply them successfully in hull form design. Scholcz et al. (2015) study the accuracy of some different response surface prescriptions and also conclude that surrogate-based optimization gives a large increase in efficiency.

The surrogate-based global optimization (SBGO) we have set up (Raven and Scholcz 2017), originally for the potential-flow code, is briefly described here. We have based the system on the DAKOTA package, a large set of optimization tools developed by Sandia National Laboratories (Adams et al. 2015); of which we apply several components.

Step 1: a Design of Experiments (DoE) is created; a set of points spread over the design space in a particular fashion. In most cases, we have used a Latin Hypercube Sampling (LHS), often supplemented with some corner points of the design space. Each point represents a hull form given by its parameter values.

Step 2: For all of these hull forms the flow code is run, yielding the value of the objective(s). For multi-objective problems involving multiple speeds or drafts, a computation is needed for each of these.

Step 3: For each objective a response surface is generated, interpolating or approximating the dependence of the objective on the parameter values. In most cases, we use a Kriging approximation, generated using Dakota.

Step 4: Optimization is done by running a Multi-Objective Genetic Algorithm (MOGA), using the response surfaces to provide the value of the objective functions at each point requested by the optimizer. This surrogate-based optimization takes negligible computation time, so the choice of the optimizer is immaterial, as long as it finds the global optimum reliably. The basic MOGA serves this purpose well. The result of this phase is a list of Pareto-optimal hull forms.

Step 5: The Pareto front is entirely based on the response surfaces, the accuracy of which may not be sufficient everywhere. Therefore, next the CFD code is run for some of the supposedly optimal hull forms spread along the front. If the computed objective values differ significantly from those from the response surfaces, these new points are added to the table of DoE results, and the response surfaces are regenerated, thus making them more accurate in the region close to the optimum.

Step 6: with the updated response surfaces we return to step 4 and run the optimization. If a significant change is found, this loop may be repeated until convergence of the front. We note that in earlier papers this iteration on the response surfaces has frequently been omitted (Kim et al 2011; Scholcz et al. 2015) but we have incidentally found significant shifts of the front in the next iteration. On the other hand, continued iteration should be limited to prevent a clustering of points resulting in overfitting of the Kriging approximation.

The required size of the DoE depends on the case. An approach to find it is by 'Incremental LHS'; determining the accuracy of the response surfaces by cross validation, and doubling the DoE size until it is sufficient. Clearly, 'adaptive sampling' techniques (Bonfiglio et al. 2017; Pellegrini et al. 2018) are much more systematic, but strongly consecutive, which is not practical for short-term projects if a large computational facility is available.

The only time-consuming part is the flow computations (steps 2 and 5). In case of wave resistance minimization using the potential-flow code, all computations for the DoE can normally be done overnight on a single desktop PC. All the other steps together are typically done in a few hours; even though we avoid a complete automation, such that we can eliminate any computational outliers and supervise the generation of response surfaces.

The procedure described is frequently applied in practice since 2016 and has been found most effective. Some examples have been published before. In (Raven and Scholcz 2017), it was used for minimizing the wave resistance of a product carrier at two speeds. Five design parameters were introduced, each of them addressing an aspect of the wave pattern computed for the initial hull form. The optimization required to evaluate 127 hull forms. Wave resistance reductions of 14% and 19% were achieved. Bigini et al. (2022) apply the same framework, parametrization and potential flow code successfully for optimization of a motor yacht hull form. By additional RANS computations, the quantitative validity of the potential-flow predictions was confirmed.

5. Multifidelity optimization

The potential flow code provides a good prediction of much of the wave making but usually has deviations for stern waves. Therefore, for optimizing the stern design more precise flow models need to be used; specifically, free-surface RANS codes. However, an optimization based on those asks a computational effort incomparable to that mentioned in the previous section, even with surrogate-based global optimization. For containership design projects, today often a multi-objective optimization for e.g. 4 different draft/speed combinations is requested. With 5 design parameters, at least some 120 hull form variations would need to be computed, thus asking for 480

RANS computations – a time-consuming and expensive matter. Therefore, further computational acceleration is desired.

Such acceleration is sought in multifidelity techniques. When applied in an SBGO procedure, this means that the response surfaces will be constructed based on a limited number of high-fidelity (HF) computations, supplemented by a large number of cheaper low-fidelity (LF) computations. If properly set up, equal accuracy of the response surfaces should be reached with reduced computational cost. For the low-fidelity computations, either a simpler model of the flow can be considered, or the same RANS model but with a coarser grid.

As a first illustration, Figure 1 shows the response surfaces for a 3-parameter variation of the afterbody of a fast displacement ship. The resistance is plotted as a function of par2 and par3, the three planes are for constant values par1 = 0, 0.5, 1. The upper three planes have been obtained from 31 HF computations (using a free-surface RANS code). These data points, indicated by markers, are for variable par1, so they are contained between the surfaces. The three planes below are the corresponding results from an LF solver, 150 potential flow calculations. Clearly, the trends against the design parameters are quite similar, but not identical; and the LF resistance level is a few percent lower. If we subtract the low-fidelity results from the high-fidelity ones, we obtain a difference function shown at the bottom of the figure, which is a very simple function of the design parameters.

Thus, if the low-fidelity approximation has a similar dependence on the hull form parameters as the HF method, we postulate

$$Rt_HF(par) = Rt_LF(par) + \Delta Rt(par) \qquad (1)$$

with Rt the objective, e.g. the resistance. Most of the trends are given by Rt_LF, found from a large number of cheap low-fidelity computations; while the simple function ΔRt can be well deduced from a small number of data points – therefore, needing a small number of HF and LF evaluations.

Instead of the 'additive' formulation given above, we find it often useful to introduce a scale factor ρ multiplying the LF result:

$$Rt_HF(par) = \rho.Rt_LF(par) + \Delta Rt(par) \qquad (2)$$

which we denote as a 'scaled additive' form. Here we derive ρ from the slope of the correlation line of HF and LF objectives.

The same expression applies to the 'CoKriging' method (Forrester et al. 2007), in which both Rt_LF

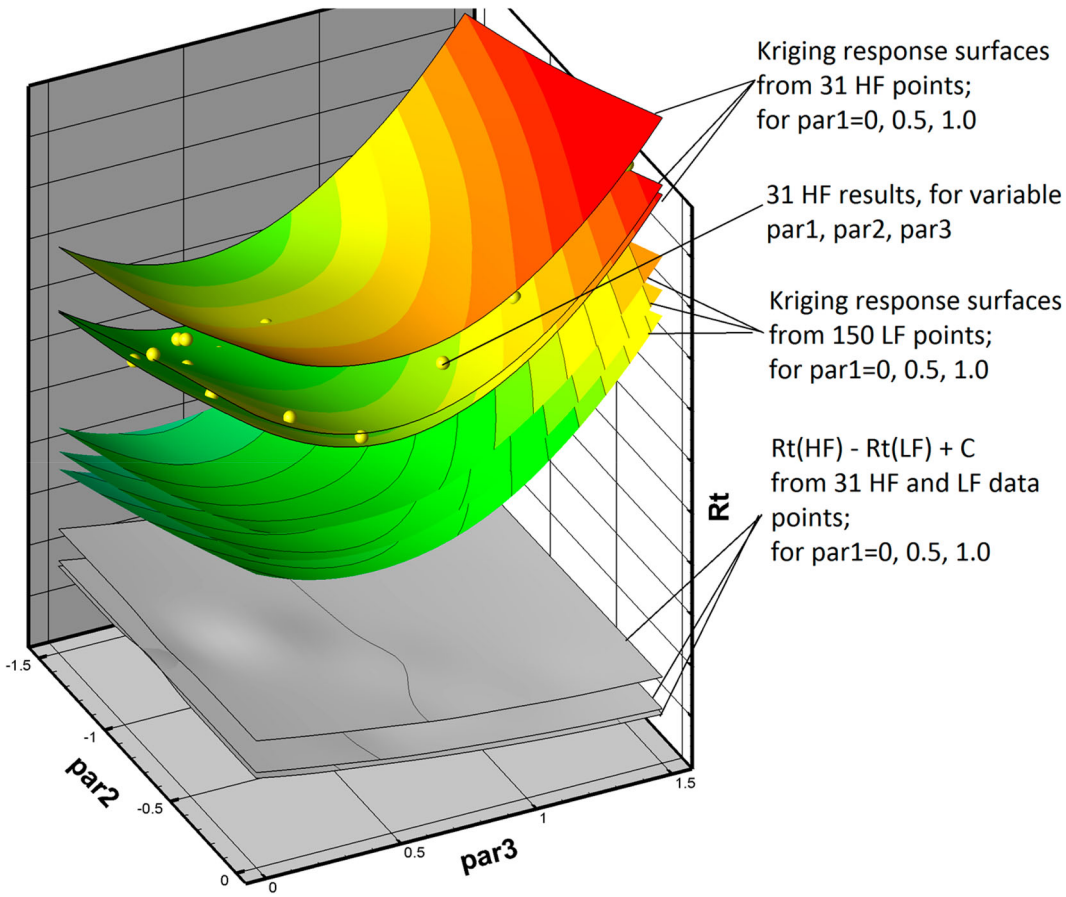

Figure 1. Response surfaces for fast displacement ship. Each plane is for constant par1. Upper 3 surfaces are for HF method, middle 3 surfaces for LF, lower 3 surfaces show the difference ΔRt (increased by a constant to show it in the same graph).

and ΔRt are approximated by Kriging, but the scale factor ρ is determined together with the Kriging approximation of ΔRt by Maximum Likelihood Estimation. This is more advanced than the scaled additive approach and aims at choosing the scale factor ρ such that the best Kriging approximation of ΔRt is obtained. Besides these methods, in principle other mappings from LF to HF can be considered, e.g. a multiplicative formulation, but we have not yet studied the possibilities of those.

Multifidelity (MF) methods for ship hull form optimization have been studied in several recent papers. Peri and Campana (2008) consider MF methods, either combining potential and RANS methods or using different panel densities. They use a Particle Swarm method, driven by the LF solver plus a correction to HF. This correction is modelled by Kriging or other surrogate models, and repeatedly updated when new HF data are added during the optimization. Therefore, no surrogate is used for the LF results, and many computations are needed. In an example combining potential flow and RANS, the trends of both solvers are opposite, and single-fidelity might have been more efficient. In Peri et al. (2010) however, the same MF combination is used successfully to optimize the forebody of the 5415 model, showing the potential to reduce the number of HF evaluations strongly. Peri et al. (2012) apply this MF method to optimize a catamaran hull form. The MF formulation reduces the required number of HF computations by 40%.

Initially, no surrogate-based MF methods were used, but most later developments do. Pellegrini et al. (2017) apply an additive MF formulation, using response surfaces formulated in terms of Radial Basis Functions. LF and HF data are from a linearized free-surface potential flow code with different panel densities. Using adaptive sampling they solve a 4-parameter SWATH optimization problem in 117 LF and 27 HF calculations, which seems really efficient. Bonfiglio et al. (2016, 2017) use CoKriging. In Bonfiglio et al. (2016) a linearized free-surface potential flow code and a RANS solver are combined. For optimization of a SWATH design with 8 parameters, using adaptive sampling in the final stage, they use 1100 LF and 107 HF computations, again most efficient.

Wackers et al. (2020) have studied an additive multifidelity formulation with LF and HF data obtained from the same RANS solver but with different grid density. For one of the solvers considered, these were obtained by automatic grid adaptation. The response surfaces were defined by Radial Basis Functions, with fewer basis points than data points, in a way to filter out the noise, found mostly in the lower-fidelity data. They find that their procedure works better with 3 fidelity levels than with two, a conclusion that likely depends on details of the noise filtering used. In a 2-parameter optimization of model 5415, they use 8 HF, 7 intermediate-fidelity and 30 LF computations.

Liu et al. (2022) study hull form optimization using a CoKriging method. Various analytical test problems are studied to determine its properties and dependence on the number of HF evaluations. The method is then applied to optimization of the 5415 model at Fr = 0.28, combining the Neumann–Michell method as LF, and a RANS code as HF solver. The LF code is more than 200 times faster than the RANS code but yields a rather poor prediction of the wave pattern for the basis hull form, which makes it somewhat uncertain for predicting design trends. The final optimization is done with only 50 LF and 30 HF computations, for a 5-parameter design problem.

In our own developments (Raven 2018; Raven and Scholcz 2019), initially we have also used CoKriging to generate the multifidelity response surfaces. An implementation in the routine MultiFiCoKriging, part of the OpenMDAO package, was used. An improvement compared with an additive MF formulation was not observed, and the routine showed an unpredictable dependence on the input. Later it was found that the Maximum Likelihood Estimation in the Kriging process was done using a local optimizer, leading to inferior results in some cases. This has probably affected our judgement of the method.

CoKriging is often promoted by quoting an example from Forrester et al. (2007), which shows a magic improvement of the approximation of an algebraic function by the use of a somewhat similar LF function. However, this is a constructed case with a very particular choice of the LF function; for many other LF functions, even more similar to the HF function, the advantage of CoKriging compared with single-fidelity Kriging is soon lost (Raven and Scholcz 2019). Therefore, while for certain less well-correlated LF and HF functions CoKriging could lead to a more accurate approximation of the difference function than the scaled additive method we use now, the advantage will likely not be that substantial. In general, a good correlation of LF and HF results remains essential for the success of a multifidelity method, as demonstrated by Toal (2015).

We shall now discuss the two classes of multifidelity optimization for ship hull forms that we use in practical ship design projects: the combination of potential flow and RANS codes, and of coarse and fine grid RANS computations.

6. Multifidelity optimization using potential flow and RANS codes

The first multifidelity optimization method we have developed is a rational extension of the method of

Section 4. The free-surface potential flow solver in general cannot be relied on for optimizing the stern design, but for rather slender ships it often predicts the stern wave pattern fairly accurately. Therefore, it is of interest to study whether it still can serve as a low-fidelity solver in a multifidelity optimization.

Thus, the HF result is the total resistance from a free-surface RANS code. The LF approximation is the wave resistance predicted by the potential flow code, plus a viscous resistance estimate found from the actual wetted surface area and a fixed form factor 1 + k, which is estimated or computed for the initial design. Therefore, this multifidelity procedure can be expected to be limited to ships and hull form variations for which wave resistance variations are dominant and the viscous resistance is largely constant.

6.1 Model 5415

The first example is model 5415, a naval vessel design frequently used as a test case; for Fr = 0.28 at full scale. In previous optimization studies for this model, often many hull form parameters were varied (Grigoropoulos et al. 2017), and sometimes the most advanced techniques have been applied to select the hull form variation modes (Serani et al. 2018). Instead, we apply the approach outlined before.

Figure 2 compares the wave patterns of the original design, predicted by the nonlinear potential flow code RAPID and by the RANS solver, REFRESCO in this case. For the forebody waves, there is very little difference, except for the quicker decay of the waves in the RANS code, caused by the larger numerical damping and coarser grid away from the hull. The stern wave system is somewhat weaker in the RANS results due to viscous effects, with some forward shift caused by the wave propagation over the viscous wake. Further aft, most of the amplitude difference is again due to the numerical damping of the RANS code. The agreement promises a useful application of the potential flow code.

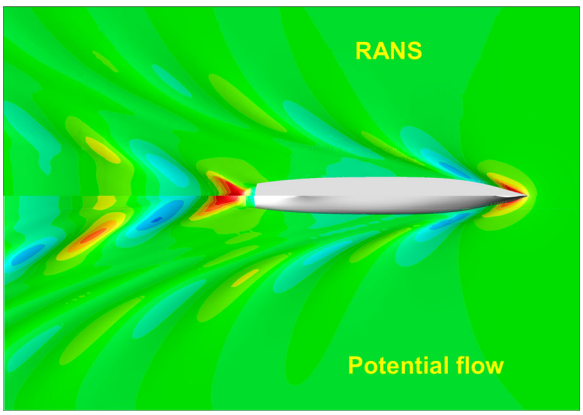

Figure 2. Wave pattern of model 5415, Fr = 0.28, full scale. Top half: RANS (REFRESCO); bottom half: potential flow (RAPID).

We recall here that a *linearized* potential-flow code cannot properly deal with the immersed transom and may not be useful here. A similar wave pattern comparison in Serani et al. (2022) clearly shows the typical deviations caused by the linearization; which in their case also leads to a poor LF-HF correlation.

To select the hull form variations to be studied, we notice from Figure 2 that the largest contribution to the wave resistance comes from the stern wave system. These waves have at least the same amplitude as the bow wave system but are less divergent, therefore their resistance contribution is larger. So we address the stern wave making in the first place. Based on this wave pattern and experience with similar cases, we introduce only 3 parameters:

- one lifting the transom edge, to reduce the rooster tail and stern wave amplitude;
- one increasing the S-shape of the stern buttocks, with a concave part just ahead of the transom. If properly shaped this often reduces the transverse stern wave components;
- one increasing the stern deadrise, which may make the stern waves more divergent but might counteract the effect of the S-shape.

In this design space, we perform RANS computations for 27 hull forms (3 * 3 * 3). Each takes 4.7 h on 240 cores, for a grid of 5.2 M cells. For the potential-flow code, we create a DoE consisting of the same 27, plus a Latin Hypercube Sampling of 100 hull forms. Each calculation takes 2.5 min on a single desktop PC for the converged nonlinear solution, using 12,500 panels; so all are ready in 5.5 h. Therefore, an HF evaluation takes 25,000 times more calculation time than an LF one.

Figure 3 shows the correlation of LF and HF resistance values, which is quite satisfactory; taking into

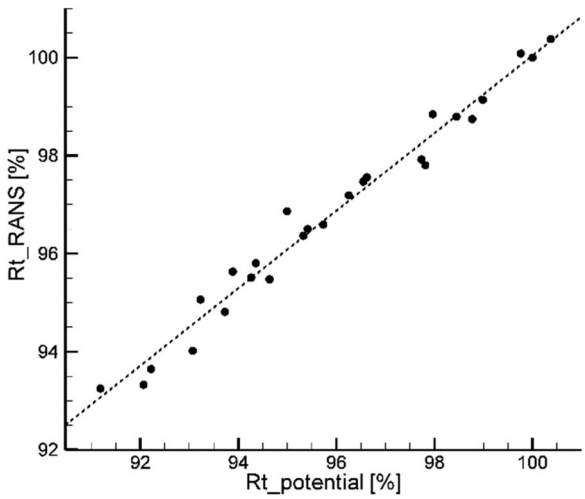

Figure 3. Correlation of resistance values from potential and RANS code, for 5415 case.

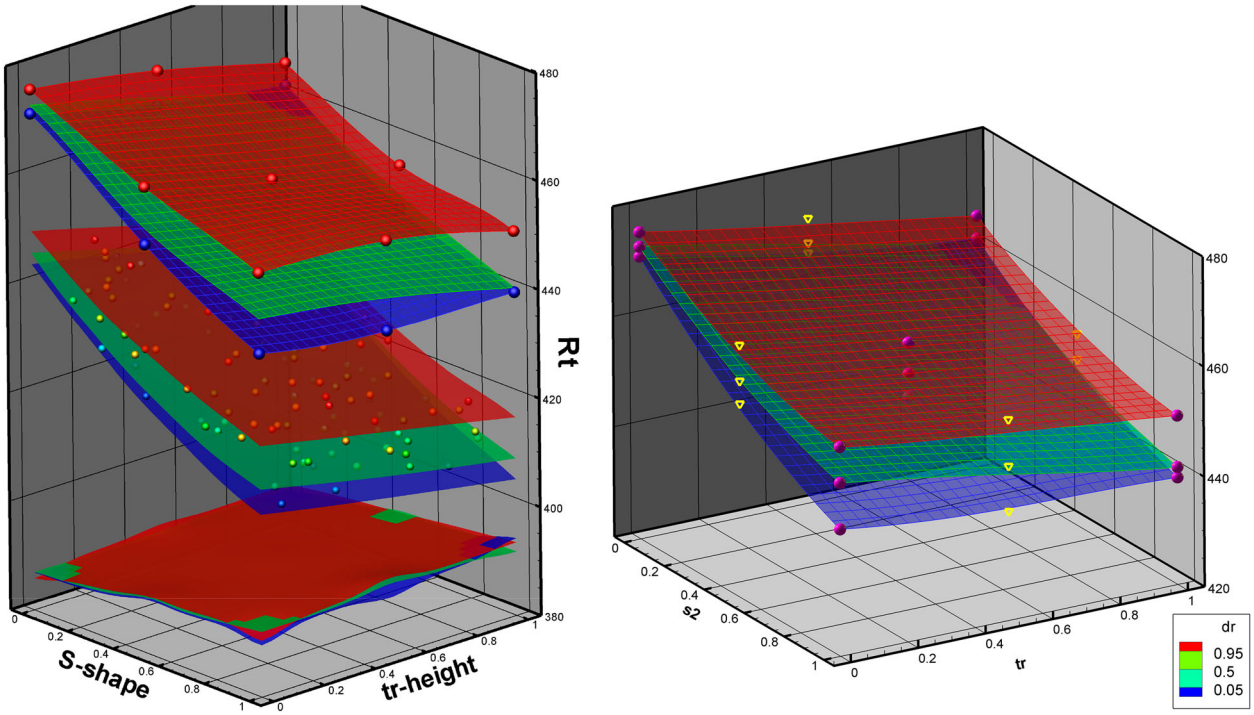

Figure 4. Model 5415, full scale, Fr=0.28. (a) Response surfaces for HF (upper 3 surfaces) and LF (middle 3), against S-shape and transom height parameters, for 3 values of the deadrise parameter. Bottom 3 planes are the same for ΔRt (plus a constant). (b) Multifidelity response surface deduced from just 15 HF points (spherical markers). Triangles are the twelve unused HF points, for validation.

account we are considering stern variations. Next, we generate response surfaces for HF and LF separately. Figure 4(a) shows these, plotted against the S-shape and transom height parameters, for three values of the deadrise parameter. The trends are remarkably similar. In an additive MF formulation, the difference function ΔRt = Rt_HF - Rt_LF appears to be a simple and fairly smooth function of the parameters, as shown in the same figure. Thereby, if the number of HF points is reduced to 15, the difference function found is nearly identical. Figure 4(b) shows the multifidelity response surfaces obtained from just 15 HF points plus 127 LF points. The pink markers are the HF data used, the yellow ones are the 12 unused points. For these, the MF surface has an RMS error of 0.33% of Rt, or 4% of the range of Rt. Thus, with just 15 HF computations we have obtained a fairly accurate response surface; most efficient for a 3D design space.

The optimum clearly is outside the permitted parameter ranges, and larger gains could be made. For the best form inside this design space, the resistance reduction is 6.8%, for this full-scale case.

To compare this with model-scale optimizations done in other papers, we have to correct for the larger viscous resistance contribution in that case. For an estimated form factor 1 + k and assuming all reduction achieved is in wave resistance, this would mean some 4.4% reduction of the total resistance for model scale. We note that Liu et al. (2022), using a CoKriging technique with 30 HF points and optimizing in a 5-parameter space, achieve a resistance reduction of 5.1% at the same Froude number; while Wackers et al. (2020), using 2 parameters found by extensive preceding work, achieve 4.5% reduction at Fr = 0.30, using 7 HF and 8 intermediate-grid results. So our straightforward approach, combining design considerations with multifidelity optimization in the way shown, has produced a comparable resistance reduction, achieved in a limited amount of time.

The resulting hull form is compared with the original in Figure 5. The S-shape parameter is found quite effective and is at maximum value. Lifting the transom helps somewhat while increasing deadrise increases resistance. Overall the new hull form is quite acceptable and fair. Figure 6 shows that a substantial reduction of the stern wave system has been achieved.

6.2 Fast displacement ship

This test case, already described in (Raven 2018), concerns a slender fast displacement hull with a bulbous bow and a flat, wide, immersed transom stern. Block coefficient is 0.57, L/B = 5. The stern is to be optimized for Fr = 0.27 and 0.37. A comparison of the wave patterns of the potential flow and RANS code again showed a large similarity. Three parameters were defined: one increasing the deadrise at the stern, the other two giving a moderate S-shape to the stern buttocks. For potential flow, a dense DoE of 150 hull forms was calculated. 31 free-surface RANS computations were done using PARNASSOS, on a grid of 4.9 M cells per symmetric half. Each took about 24 h on

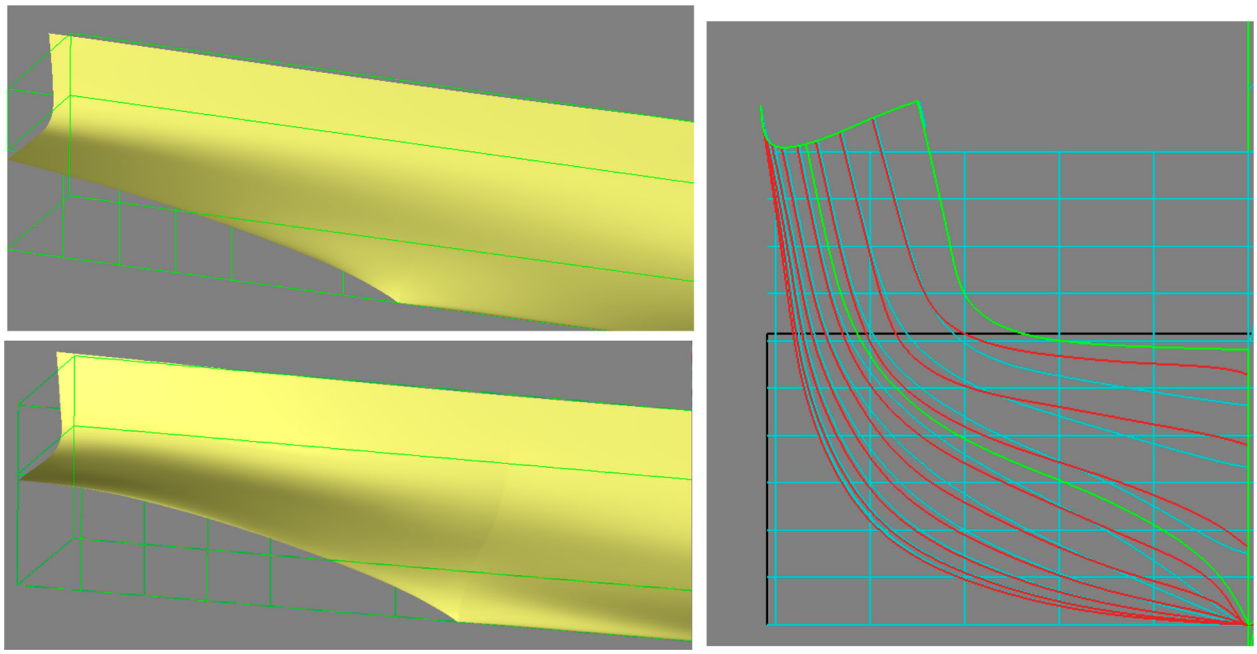

Figure 5. Original (top) and improved afterbody of model 5415. In bodyplan, light lines (blue) are original, dark lines (red) improved design.

3 cores of an HPC cluster – 2000 times more than the LF code. Response surfaces, generated separately for LF and HF results, for Fr = 0.27 have been shown in Figure 1 and show near-identical trends, while for Fr = 0.37 the LF results had somewhat weaker trends than HF.

The difference function ΔRt is again a simple function of the parameters for both speeds, and could be derived nearly as well from just 7 points in this 3D space; therefore just requiring 7 HF computations for each speed. A multi-objective optimization was done next, using these MF response surfaces derived from 7 HF and 150 LF results; and the resulting Pareto front checked by additional HF computations. As shown in Figure 7 this confirmed the accuracy of the front and the MF response surfaces. The LF optimization alone produced a similar front but with smaller gains; also it indicated similar parameter values and a very similar reduction of the stern wave system here.

For a selected point on the front the total resistance is reduced by 3.6% for the lower speed, 2.6% for the higher; corresponding with 18% and 6% of the wave resistance, respectively. As the original hull was the result of quite some design work already, this further refinement is significant.

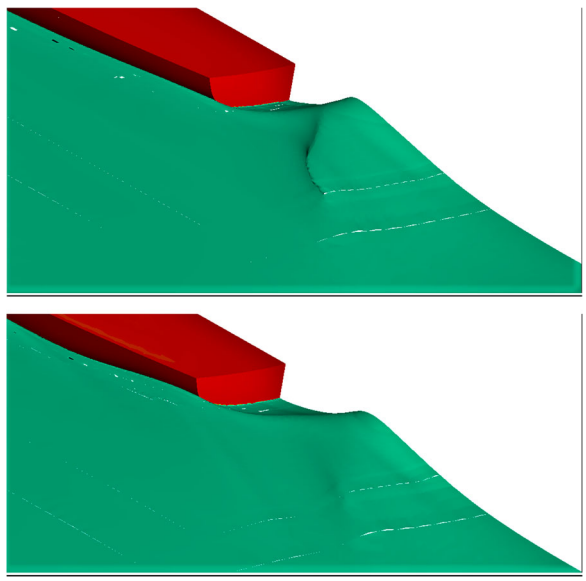

Figure 6. Stern wave of original (top) and improved design (bottom), as computed by RANS code.

6.3 Containership

For the third example, a containership stern, there is less chance of success. For the lower Froude number

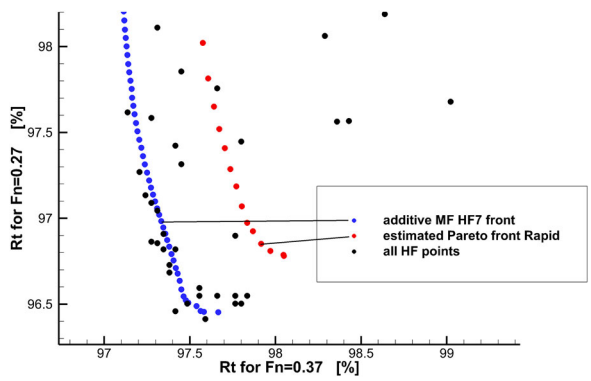

Figure 7. Pareto plot for fast displacement ship case. Resistance as percentage of that of original design. Right front (red markers) is Pareto front from potential-flow code, left front (blue markers) that from MF optimization using 7 HF points. Black markers are all RANS computations done.

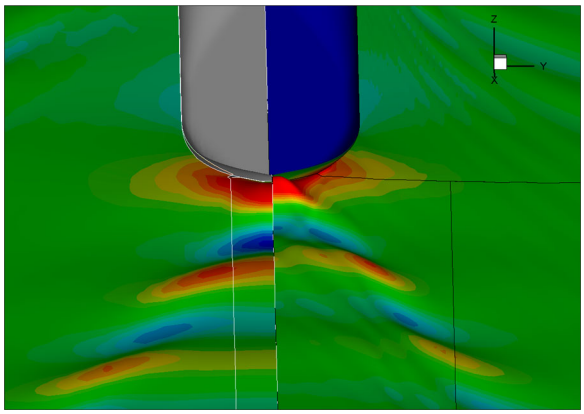

Figure 8. Perspective view of computed stern wave patterns for containership, original design. Left side: as computed by potential-flow code; right side: computed by RANS code.

of 0.20, the stern waves are shorter compared with the thick boundary layer and wake at the stern, so more viscous effects on the wave making can be expected. In the comparison of wave patterns in Figure 8, large differences in predicted stern wave are observed. The pronounced transverse waves of the potential-flow solution are largely absent in the RANS solution, which instead has more diverging components. Whether this will result in a sufficient correspondence of design trends to make multifidelity optimization effective, is to be seen.

The four hull shape variation modes selected are apparently comparable to the previous cases, but look quite different for the present hull form: we choose a long and short S-shape of the stern buttocks, lifting of the transom edge, and variation of the V-shape of the stern. A DoE of 240 hull forms was computed using the potential-flow code. Optimization just based on those promises 6.2% resistance reduction by going for a flat stern (minimum V-shape) and maximum S-shape.

Next, a DoE of 40 hull forms has been computed using the RANS code (PARNASSOS). An initial comparison of the LF and HF response surfaces indicated a complete lack of agreement, and a correlation of LF and HF resistance values in Figure 9 seems absent. No

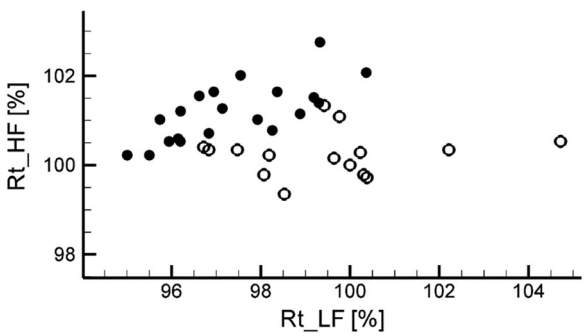

Figure 9. Containership case, correlation of HF and LF total resistance. Black dots: dry-transom flow in RANS solution; circles: wetted-transom flow.

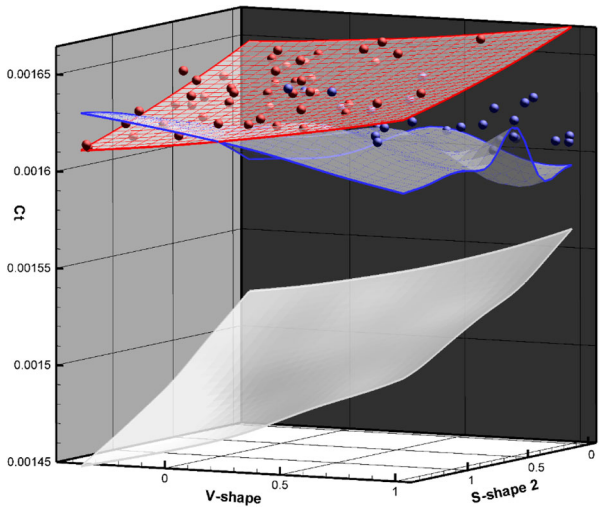

Figure 10. Response surfaces against V-shape and short S-shape parameters, for fixed other parameters. Upper plane (red mesh and markers): dry transom. Middle plane (blue mesh and markers): wetted transom. Lower plane is from LF method. Markers are for variable par2 and par3.

significant improvements are found by the HF code, even for cases that are much better in the LF results. Therefore, this case turns out completely different than the previous ones.

An explanation was found in the transom flow regime. For these stern variations, the type of transom flow is not evident a priori, and the RANS code predicts a dry-transom flow in some of the cases, wetted transom in others. In the potential-flow solver, however, a dry transom is modelled in all cases, since without viscous effects there is no valid model of a wetted-transom flow. Thus, for cases having a dry-transom flow in the RANS solution, the transom flow regimes in HF and LF solution correspond and we do observe some correlation (Figure 9); but the trend is about half as strong in the HF solution, likely connected with the weaker stern wave system. For wetted-transom flows the correlation is absent and the LF solution does not indicate the hull shapes that result in some resistance reduction.

Subsequently, additional HF computations have been made, to a total of 77, and HF response surfaces have been generated separately for the cases with wet and dry transoms. From Figure 10 we observe that these have completely different trends, indicating that different design guidelines need to be applied for the two transom flow regimes; and asking for a redefinition of the design space here. The lower plane in the figure is the response surface from the LF method, which has a slope corresponding with, but larger than, that of the HF dry-transom surface.

Dry transoms primarily occur here for more strongly S-shaped buttocks and smaller V-shape parameters, i.e. flat sterns. Instead, more V-shape and less S-shape tends to promote a wetted transom, as the more contracted shape leads to a higher pressure

level aft of the stern. In the present case, the lowest resistance levels are actually found for wetted transoms. Also, the original hull form had a wetted transom flow. This explains the smaller gains made in the HF results, and also the much weaker transverse wave system in Figure 8.

Therefore, we have clearly met a limit of the applicability of potential flow as an LF method for multifidelity optimization of a ship's afterbody design. As anticipated, this method can serve in case the viscous resistance is essentially constant in the design space. But the exchange of wave resistance and viscous losses that occurs in the wet-dry transom transition cannot be captured.

In summary, for optimizing the afterbody of rather slender vessels at a not too low Froude number and with dry-transom flow, multifidelity optimization combining the free-surface potential flow solver and a RANS code has been found to be effective and most efficient. For other cases, we will need another LF method, as described in the next section.

7. Multifidelity optimization using coarse and fine-grid RANS computations

If significant viscous-resistance variations occur in the design space, these need to be included also in the LF model. The second multifidelity method we have set up, therefore, only uses a RANS solver but with different grid densities. A large number of LF computations on coarse grids, and a limited number of HF computations on fine grids, are combined to derive multifidelity response surfaces, to be used in surrogate-based global optimization. Again this will require that the trends with design variations are sufficiently corresponding. In this case, this correspondence is not limited by differences in the flow model, but by numerical errors and their variation over the design space.

7.1 Grid generation

We intend to use the REFRESCO code here. For the generation of the unstructured grids, NUMECA Hexpress is used. In this tool, first a background grid with large cell sizes is set up, which is then locally refined with some refinement levels: the number of refinements (halving the cell size) compared to this base background grid. The refinement is done on the surfaces of the hull, including additional refinement at the transom edge. Furthermore, refinement boxes are defined, where the grid is also refined. For example, some of these boxes are there to capture the Kelvin wedge, to correctly represent the wave pattern. Thereafter, the grid is snapped to the surfaces and a buffer layer is constructed, before the viscous layer is defined. This typically results in an unstructured grid of around 7 million cells for a bare hull, in a computational domain of length 6 Lpp, half width 2 Lpp and height 4 Lpp. For full scale, wall functions are used, with a y^+ of around 200.

Coarsening of a grid can be done in several ways of which two have been explored, denoted Coarse I and Coarse II. In Coarse I, the initial grid cell size was doubled in each direction. Almost all other refinements were set one level less. This means that in general 4x less cells in each direction were specified. In Coarse II, the initial cell size was also doubled in each direction, but the refinement levels were kept the same. This means only 2x less cells in each direction. To still obtain a sufficiently small grid size, only refinement boxes close to the ship are kept in place. The Kelvin wedge has been completely excluded from grid refinements. In this way, the pressure and shear stresses on the hull should be better captured compared to Coarse I.

The first test case was a set of three hull forms, consisting of the original, a variant and an intermediate shape. Figure 11 shows the resulting normal and coarse grids, including the number of cells. The variation of the calculated resistance is shown in Table 1. Both the time step and the number of outer loops were kept equal in these computations. The wall clock time is given in Table 2.

Clearly, the low fidelity results, calculated on a grid with about 1/10 of the number of cells, have a larger resistance. Coarse I shows an increase of about 15% and Coarse II of about 7.5%. This increase makes it quantitatively unacceptable for most purposes. However, the relative variation of resistance between the design variants is almost equal at least for Coarse II. As in general its results were much better, only Coarse II is used in the remainder of this paper.

We thus notice that the coarse-grid generation is rather critical. Moreover, if other objectives need to be deduced from the same LF computations, e.g. the wake field, additional requirements for the coarse-grid generation hold.

7.2 Acceleration of LF computations

To reduce the overall effort and justify the added complexity, we have studied several other possibilities to further reduce the computation time of the LF computations. For the coarse grids, initially, the reduction of the calculation time was roughly proportional to the reduction of the number of cells, see Table 2. Normally in CFD the calculation time scales with the number of cells to a higher order, so the coarse-grid computations seemed unduly expensive. Acceleration was obtained by increasing the time step size proportionally with the mean cell size, and by reducing the number of outer loops per time step as enabled by the faster convergence. With the settings found, the coarse grid computations have about 10 times fewer cells and take about 20–25

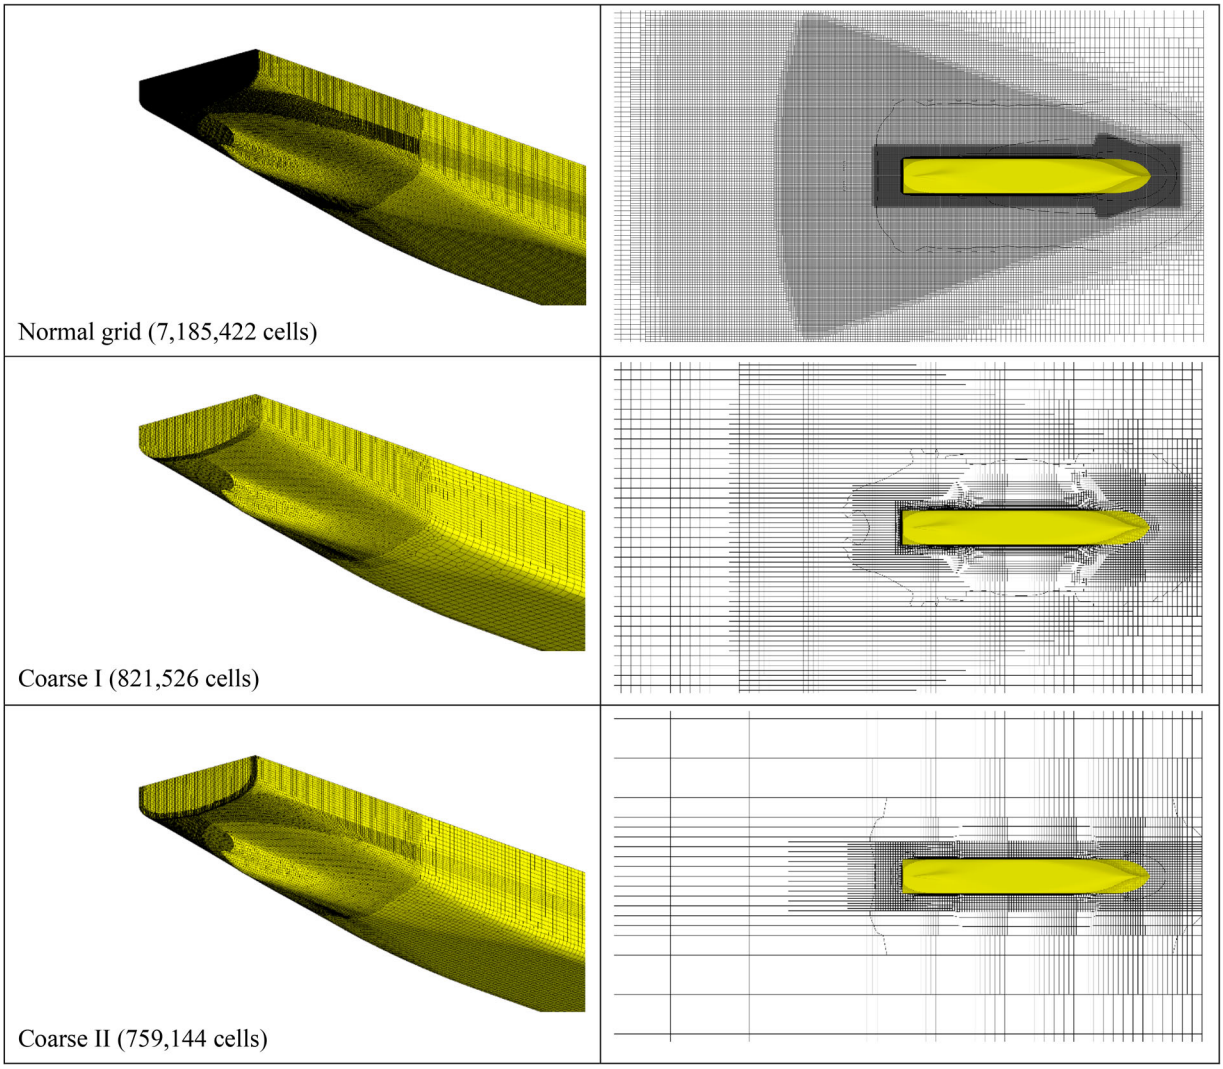

Figure 11. Surface grids (left) and cross cut at waterline (right) for the three methods of meshing.

times less CPU. Additional acceleration options may exist in the trim and sinkage adjustment; for example, by fixing sinkage at the value of the initial hull and just solving for the (more variable) trim angle.

7.3 Comparison of response surfaces

To assess the viability of coarse-fine multifidelity optimization, we have first compared LF and HF response surfaces for some cases. Figure 12 shows these for a containership stern, with variation of the vertical position of the transom edge and the V-shape of the aft sections. The response surfaces are defined by Kriging. The upper response surface in the figure is for coarse grids of 700,000 cells, defined as in Section 7.1; the middle surface is for grids of 4.6 million cells, the lower for 10.8 million cells. The coarser the grid, the higher the resistance again, even by 15%. But the trend against the design variables is quite comparable. The intermediate surface, which has a larger number of points, shows more variability. Whether these are genuine trends or just due to scatter in the data is unclear. How we can systematically distinguish these in general is an important question we have not addressed yet. In any case, the figure promises a useful application of the coarse-grid data. Several other cases have been examined, with similar conclusions.

Table 1. Bare hull resistance of half the ship in kN. In brackets, the relative difference compared to the normal mesh is given.

[kN]	Normal	Coarse I	Coarse II
Original	1337	1554 (+16.2%)	1436 (+7.40%)
Intermediate	1284		1377 (+7.24%)
Variant	1255	1436 (+14.42%)	1351 (+7.65%)

7.4 Multi-objective multifidelity optimization of a containership

The first actual test case is a large containership (not the same as that from Section 6.3). The initial hull

Table 2. Wall clock time of the simulations of the original geometry in seconds.

Normal	Coarse I	Coarse II
22,851 (6u20m51s)	3120 (0u52m00s)	2479 (0u41m19s)

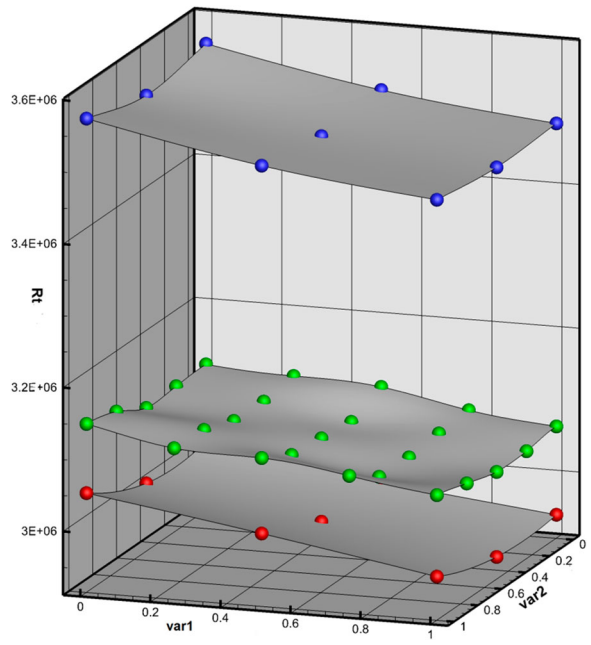

Figure 12. Response surfaces for resistance of a containership as a function of transom height (var1) and V-shape of stern (var2), as computed on 3 grid densities: 0.7 M (upper surface), 4.6M (middle surface) and 10.8 M cells (lower surface).

form is a design from a late stage in a practical project, already refined through CFD work. We want to optimize the afterbody design for two conditions:

A. Draught 16 m, speed 19 kn,
B. Draught 14.5 m, speed 22 kn.

As hull form parameters we choose:

1. the height of the lowest point of the transom edge; varied from 13.8 to 16.9 m above the baseline;
2. The shape of the stern buttocks, varied from somewhat concave to somewhat convex aft, without changing the height of the transom edge;
3. the V-shape of the stern sections;
4. the roundness of the transom edge and the last stern sections.

In the 4D design space we generate a DoE of 12 points for HF, defined by Latin Hypercube Sampling; and 96 for LF, generated by 'Incremental LHS', such that the 12 HF cases are also part of the LF DoE. We use the RANS code Refresco, for the ship at full scale. For the LF computations, a coarse grid of 800,000 cells was used, and trim and sinkage were fixed at the values computed for the original design. The required CPU time per case was 40 min on 120 processors. The HF computations used grids of 7.8 million cells, and trim and sinkage were left free. Each took 7 h on 240 processors, about 20 times more costly than LF.

Figure 13 shows the correlation of the computed LF and HF resistance values. For Condition A this is really good. The correlation line shown has slope of 0.87, so the trends are stronger in the LF results, and the resistance level again is a lot higher. This slope is used as the factor ρ in Equation (2), so ΔRt is just the vertical deviation of the points from the lines in Figure 13, plus a constant. For condition B the correlation is much less convincing and the slope of the line is 1.20.

Response surfaces, derived by Kriging, are shown in Figure 14; where resistances have been normalized with that of the initial design on the coarse grid. Those for LF are well-defined by the 96 LF points. For condition A, the ΔRt surfaces are almost a constant, and are also well defined by the 12 data points. Response surfaces derived from just the 12 HF points in this 4D design space, without the support from the LF data, would be much less reliable. The same holds for condition B, but the poorer correlation of LF and HF data is reflected in ΔRt response surfaces with some more variation and shape. Using just 12 HF data points may lead to some inaccuracy here. Still, the multifidelity response surfaces should be more accurate than the HF surfaces.

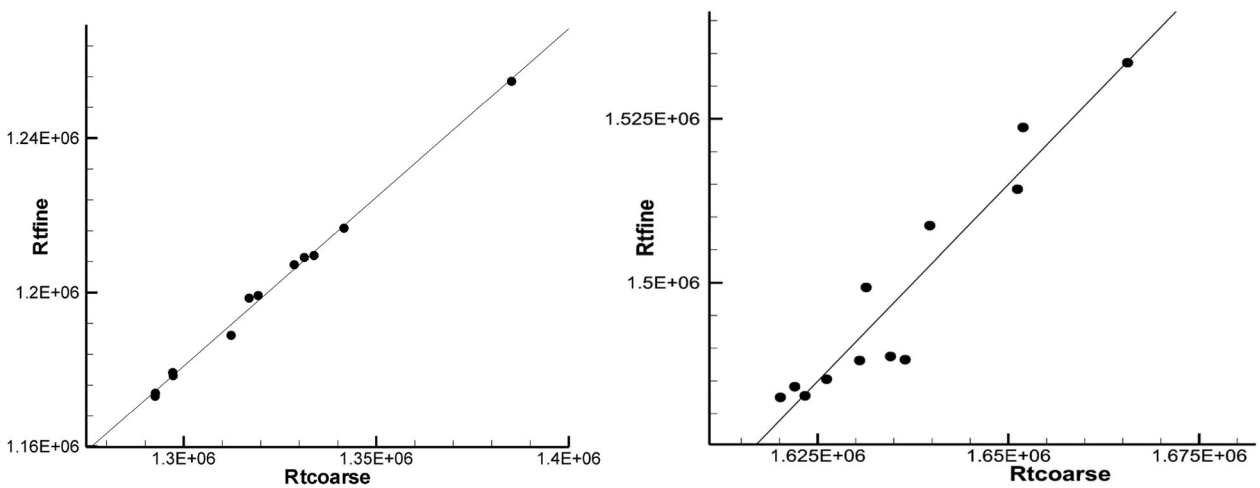

Figure 13. Correlation of resistance computed on coarse and fine grid, for condition A (left) and B (right).

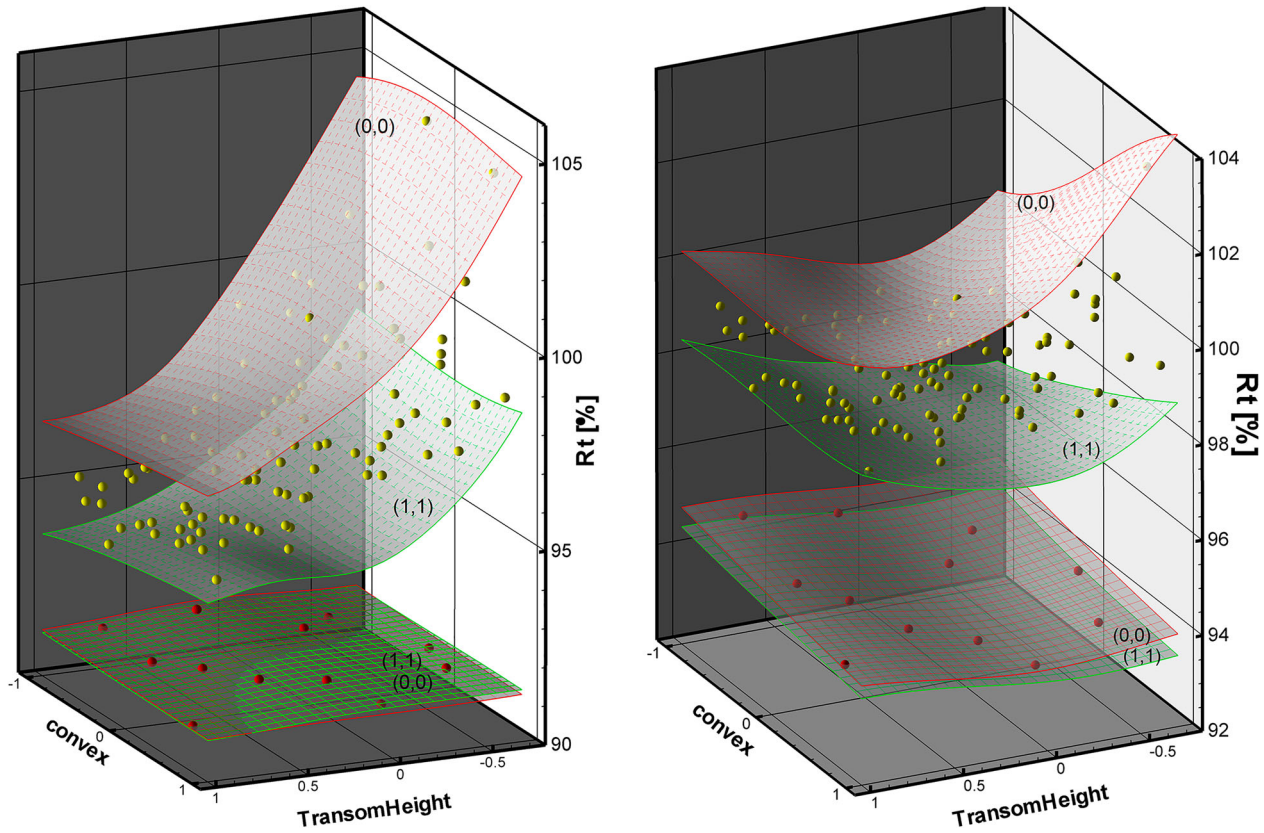

Figure 14. LF response surfaces (upper two surfaces) and ΔRt surfaces (lower surfaces) (increased by a constant). Labels indicate (Vshape, TransomRound) values. Markers are data points, for variable Vshape and TransomRound. Left: condition A; right: condition B.

The uncertainty of the response surfaces can be estimated by cross-validation. Assuming independence of the errors in LF and ΔRt response surfaces, we estimate the uncertainty of the MF surface as $\sigma_{MF} = \sqrt{[\rho^2 \ \ \sigma^2_{LF} + \sigma^2_{dif}]}$. Then for condition A the uncertainty of the MF surface is 29% of that of the HF-only surface, about 2.1% of the range of the MF response. For condition B the uncertainty of MF is 58% of that of HF, and 3.6% of the range.

The multifidelity response surfaces are then defined by combining the LF and ΔRt surfaces, and the surrogate-based optimization is done. Figure 15 shows the resulting Pareto front. As this is rather far from the actual HF evaluations done, an iteration is necessary. Additional HF computations have been made for 5 points on the front, plus 9 other points selected for identifying design trends. These new data permitted to determine the RMS errors in the initial HF and MF response surfaces, indicating again that the MF approximation is much better for condition A, and somewhat better for condition B. Still significant errors were found for condition B for the front points. In hindsight, a somewhat larger number of HF data points at least for condition B would have worked better, as was already indicated by the cross-validation results.

All new HF points are then added to the data set, and updated response surfaces are derived. Figure 15 also shows the second and final estimate of the Pareto front.

For comparison, Figure 15 also shows the initial Pareto front estimates derived from an HF-only and LF-only optimization. The HF front, based just on the 12 HF data points, is far off and has significantly different parameter values, so it does not indicate the right

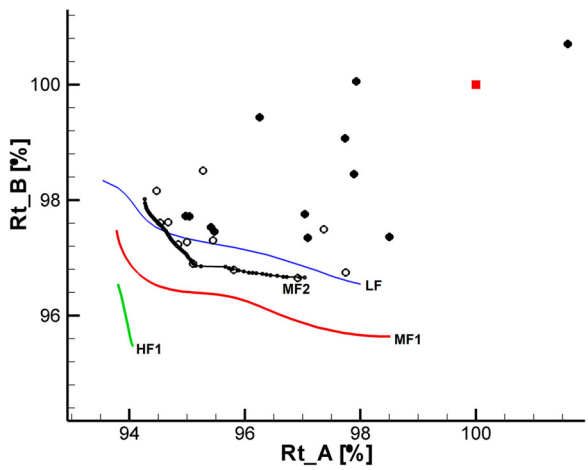

Figure 15. First (MF1) and second (MF2) estimated Pareto fronts from multifidelity optimization. Black dots: HF data of 1st round; black circles: added HF points in 2nd round. HF1: estimated 1st front from HF-only optimization; LF: estimated 1st front from LF-only optimization. Large square is original hull form.

optimal design. Only upon iteration the HF optimization converged rather slowly to the same location as the second front from the MF optimization. Therefore, clearly an HF-only optimization requires a larger number of data points. The LF-only optimization, based on the 96 coarse-grid data points, is far off in absolute resistance level but not bad when expressed as a percentage of the coarse-grid resistance of the original design; and does indicate the right parameter values.

For a selected point on the front, the resistance reduction is 4.9% for condition A, 3.1% for B. As the initial hull form was already a realistic practical design, and in view of the small wave resistance and form factor of modern containerships, this is a satisfying improvement. No constraint on initial stability has been imposed, so this hull form is normally not practically acceptable yet but still provides useful design information. It differs significantly from the initial hull (Figure 16), having maximum V-shape of the stern sections, and the transom edge at an intermediate height, as a compromise between both draft conditions. The convexity of the stern buttocks had little effect but permits to increase the displacement.

For condition A the transom is slightly wetted, as for the initial hull. A substantial reduction of the stern wave making is observed (Figure 17), in particular for the transverse wave components. For condition B, the transom is dry and the transom edge is at a higher position than desired, but due to the stronger V-shape this is not too harmful. Again some reduction of the stern wave making has been obtained. In view of the limited share of wave resistance for this class of ships, the differences in stern wave making do not explain the entire resistance differences, and actually the total-head loss in a transverse plane aft of the transom also was found to be reduced. Therefore, the optimized hull has both a smaller wave resistance and smaller viscous losses.

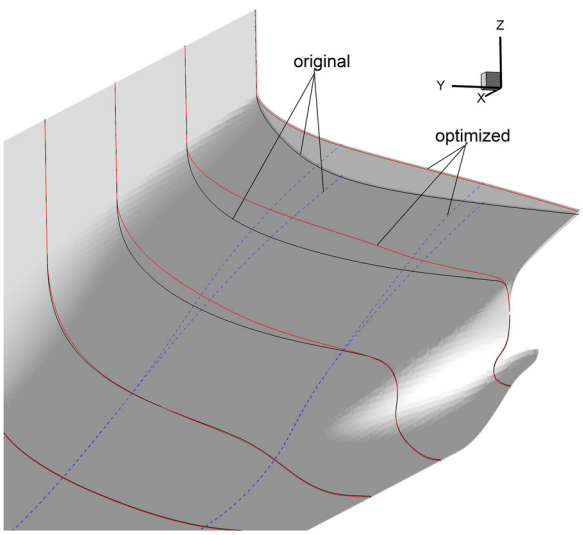

Figure 16. Initial and optimized hull form.

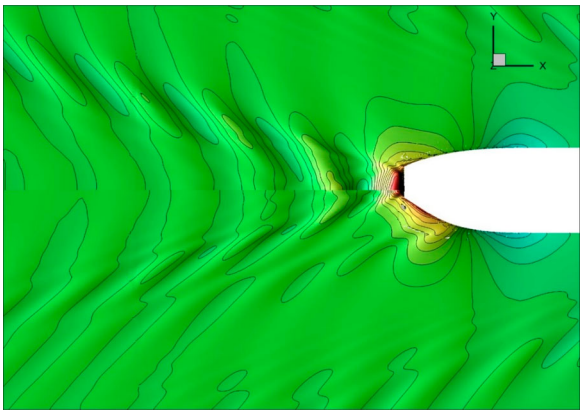

Figure 17. Stern wave pattern for original (top half) and optimized design (bottom half), for condition A.

The complete 2-objective optimization has thus been done using, for each condition: 12 initial HF computations, plus 5 computations for points on the first estimated front; along with 96 LF computations. This adds up to a total computation time corresponding with 22 HF computations per condition. For a single-fidelity optimization based on only HF computations, for this 4D design space, one would probably need some 45 points (a DoE of 40 plus 5 Pareto front points), therefore about twice the computational cost of the multifidelity optimization.

7.5 Multi-objective multifidelity optimization of a motor yacht

The next example, again from a practical project, is a motor yacht with a bulbous bow and a deep and wide transom, at Fr = 0.35 and 0.70. At both speeds the transom was dry. The case could have been dealt with by the potential/RANS multifidelity method of Section 6, but we have used the coarse/fine RANS approach here.

A 3-parameter hull form variation was defined, with variations of the entrance angle, the aft buttock shape and the transom immersion. A DoE of 48 coarse-grid and 12 fine-grid points was generated. Coarse grids had 700,000 and 2.2M cells for the two speeds respectively; fine grids 4.7 and 11.2 M cells. Trim and sinkage were left free for both LF and HF. For the higher speed, some computations did not converge and just 40 coarse-grid and 8 corresponding fine-grid points were available.

The correlation of the resistance values was very good for the lower speed, and the ΔRt function was little more than a constant. For the high speed, the correlation was not particularly good, but still ΔRt was a simple function, and deriving it from just eight points seemed sufficient. The MF optimization yielded the Pareto plot of Figure 18. Along the short front, the transom immersion varies, with a minimum value being best for the lower speed, a somewhat larger

immersion for the high speed. Optimization using only the LF data did not lead to the same front, but to minimum transom immersion for both speeds.

For a selected point on the front additional HF computations have been done, yielding resistances differing by just 0.17% and 0.06% from the estimates based on the MF response surfaces. Large reductions of the total resistance are achieved, of 14% and 4.5% for the two speeds.

Therefore, the MF response optimization worked very well in this case. With just 13 fine-grid calculations per speed for the 3-parameter 2-objective optimization, this was the most efficient.

8. Discussion

In this paper we have proposed and demonstrated multifidelity formulations for use in ship afterbody optimization for minimum resistance. With the cases considered, there are a few aspects to be discussed.

Two different LF methods have been used, suitable for different applications. The potential-flow method is limited to a certain class of cases, the coarse-mesh RANS method seems more generally applicable but is computationally more demanding. These methods have been selected for their ability to well represent the design trends of the HF method. As a matter of fact, we observe a roughly linear correlation $Rt_HF(par) = \rho.\ Rt_LF(par) + \Delta Rt$ in most cases, with ΔRt having little dependence on the parameters. The 'scaled additive' MF formulation we have used fits this behaviour well and is straightforward.

As Toal (2015) demonstrates, the success of an MF method depends quite strongly on the LF-HF correlation. The same is observed in the cases considered, with the example of Section 6.3 requiring a different LF method and that of Section 7.4 suggesting to use more HF points in the DoE for condition B. To obtain sufficient correlation, it is important to make the LF method accurate enough. For instance, the generation of the coarse mesh RANS is not executed by simply coarsening by a factor of 2 in all directions. Some local hull features need more cells compared to others in a coarse mesh, to capture the trends well enough.

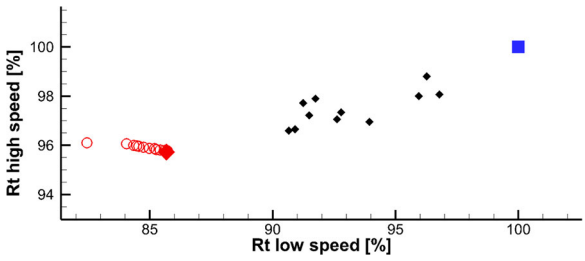

Figure 18. Motor yacht case. Pareto plot, with initial design (large square), fine-grid Design of Experiment (black diamond markers) and estimated Pareto front (circles). Large diamond marker is recomputed Pareto point.

The number of points in the initial DoE is of importance for the overall efficiency. For well-correlated LF and HF functions, Toal recommends to spend a fraction f_r of the total computation time to the LF computations, with $f_r > 1.75/(1 + 1/Cr)$, Cr being the relative cost of an LF computation compared with HF. We moderately exceed this lower bound in all cases.

The required number of HF computations, however, strongly depends on the case and the level of LF-HF correlation. For the cases considered we have indicated our assessment, but more experience is to be collected. Furthermore, it is important to always check the correlation and add more points if required. Use of an extendable DoE along with cross validation can be a good solution.

It is useful to compare the proposed multifidelity method with single-fidelity optimization using just HF or just LF data. Compared with HF, the advantage is of course just in computation time, since the optimum found should be the same. Compared with LF only, however, the optimum can differ. For the roughly linear LF-HF correlation that we found for most cases, the location of the optimum in the design space is frequently already predicted by the LF method. This is the case for the examples of Sections 6.2 and 7.4 in this paper. One might then conclude that a LF optimization would suffice for optimization of the afterbody design. However, in any case, the resistance level and the amount of resistance reduction obtained from LF would then not be correct (e.g. Figure 7), and in a multi-objective optimization, the resulting distortion of the Pareto front could also lead to a different Pareto optimum chosen.

Moreover, a ship afterbody optimization using only potential flow, or RANS on very coarse grids, should generally not be trusted. For the other examples in this paper, the LF method does not correctly indicate the optimum, and that of Section 6.3 indicates the misleading result one might get.

Therefore, the addition of a limited set of HF points to the larger set of LF points is needed to verify the optimum found, the correlation between LF and HF and the correspondence of design trends. Having those HF points available, these also serve to derive MF response surfaces, with their increased accuracy and reliability.

9. Conclusions

This paper has discussed practically oriented ship hull form optimization methods, which are being used in practical ship design. Starting from a surrogate-based global optimization using free-surface potential flow calculations, similar methods based on RANS computations have been set up.

In all methods, a parametrization of deformations of a given initial design is used. This is a parametric blending of the initial hull form and several modified hull forms. Where possible, the latter are selected based on physical insight and design knowhow, from analysis of computed results for the initial design; and are simply created in the CAD system. This method is most flexible, and generally leads to effective hull form changes with a limited number of parameters; and often, to simple trends against the parameters, easily identified with few calculations. Thus, the parametrization also contributes to the efficiency of the optimization procedure.

For multi-objective optimization of ship afterbodies, the computational effort for hundreds of RANS computations could be prohibitive. Therefore, multifidelity (MF) techniques have been studied. The combination of a free-surface potential-flow code as a low-fidelity (LF) method, and a RANS code as a high-fidelity (HF) method, has been found to be very efficient for a class of cases. The potential-flow results are easy and cheap to obtain and permit a large reduction in the number of HF computations. This MF method is suitable for fairly slender ships at Fr > 0.25, with dry transom flow, for which the viscous resistance does not change appreciably by the variations applied and wave resistance changes are dominant. The viability of this method can be assessed from an initial comparison of the inviscid and viscous wave pattern and flow. Successful applications to the 5415 case and a fast displacement vessel have been discussed.

It has been shown how for a specific containership this MF method failed due to the occurrence of wetted transom flows in a part of the design space, causing a change of wave resistance and viscous resistance not reflected by the LF results. For ships and stern variations that possibly cause significant changes in the viscous resistance, an alternative multifidelity formulation is needed, and has been found in the use of coarse-grid RANS computations as a LF method. The coarse-grid generation is critical as it is not evident that design trends are well represented. A particular prescription has been found, leading to grids with about 10 times fewer cells and computations asking 20–25 times less CPU time; but with the design trends still corresponding well with those on fine grids, with gradients often differing by 10–20%. In the scaled additive MF formulation, the HF-LF difference often is a near constant function, efficiently determined by few HF computations. In some cases and conditions, we found a poorer correlation of LF and HF results, leading to a poorer performance of the MF method.

The coarse/fine grid multifidelity optimization was applied successfully to a containership and a motor yacht, leading to significant resistance reductions. For both cases, few HF computations sufficed to find the optimum. For the container ship, the MF optimization required about one half the computation time of an optimization based on HF computations alone. This method therefore is efficient, although not as efficient as the potential/RANS MF method; but it is more generally applicable.

In further practical applications, more experience is to be collected. This should provide better insight in the required number of points in the initial DoE's, and of the LF-HF correlation to be expected for various cases. Some attention is also desired for a proper treatment of scatter in the data. With these and other refinements we think that multifidelity techniques for ship hull form optimization may contribute to a further advance of practical ship afterbody design.

Disclosure statement

No potential conflict of interest was reported by the authors.

Funding

This work was supported by the Ministry of Economic Affairs, Netherlands.

References

Adams BM, Bauman LE, Bohnhoff WJ, Dalbey KR, Ebeida MS, Eddy JP, Eldred MS, Hough PD, Hu KT, Jakeman JD, et al. 2015. 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.

Bigini G, Jacoby M, Steenbergen A, Minerva LF, Villa D, Vernengo G. 2022. Large yacht resistance reduction by hydrodynamic multi-objective shape optimization. Ship Tech Res. doi:10.1080/09377255.2022.2054759.

Bonfiglio L, Perdikaris P, Brizzolara S. 2016. Multi-fidelity optimization of high-speed SWATHs. SNAME Maritime Convention, Bellevue, USA.

Bonfiglio L, Perdikaris P, Vernengo G, Karniadakis G. 2017. Multi-fidelity Bayesian optimization of SWATH vessels for improving seakeeping performance, SNAME Maritime Convention, Houston, USA.

Brizzolara S, Vernengo G, Pasquinucci CA, Harries S. 2015. Significance of parametric hull form definition on hydrodynamic performance optimization. 6th Int. Conf. Comp. methods Marine Engineering, MARINE 2015.

Chun HH. 2010. Hull form parametrization technique with local and global optimization algorithms. Int. Conf. Marine Technology, Dhaka, Bangladesh.

Crepier P. 2017. Ship resistance prediction – verification and validation exercise on unstructured grids. 7th Int. Conf. Comp. Methods Marine Engineering (MARINE2017), Nantes, France.

Feng Y, El Moctar O, Schellin TE. 2021. Parametric hull form optimization of containerships for minimum resistance in calm water and in waves. J Mar Sci Appl. 20:670–693. doi:10.1007/s11804-021-00243-w.

Forrester AIJ, Sóbester A, Keane AJ. 2007. Multi-fidelity optimization via surrogate modelling. Proc Royal Soc A. 463(2088):3251–3269.

Grigoropoulos GJ, Campana EF, Diez M, Serani A, Gören O, Sariöz K, Dani'sman DB, Visonneau M, Queutey P, Abdel-Maksoud M, Stern F. 2017. Mission-based hull form and propeller optimization of a transom-stern destroyer for best performance in the sea environment. 7th Int. Conf. Comp. Meth. Marine Eng., MARINE 2017.

Han S, Lee Y-S, Choi YB. 2012. Hydrodynamic hull form optimization using parametric models. Jnl Mar Sci Technol. 17:1. doi:10.1007/s00773-011-0148-8.

Hoekstra M. 1999. Numerical simulation of ship stern flows with a space-marching Navier Stokes method, Dissertation, Delft Univ. Techn.

Hoekstra M, Raven HC. 2003. A practical approach to constrained hydrodynamic optimization of ships. NAV 2003 Symp., Palermo, Italy.

Kim H, Jeong S, Yang C, Noblesse F. 2011. Hull form design exploration based on the response surface method. 21st ISOPE Conf. Hawaii.

Kim H, Yang C. 2010. A new surface modification approach for CFD-based hull form optimization. 9th Int. Conf. Hydrodynamics, Shanghai, China.

Liu X, Zhao W, Wan D. 2022. Multi-fidelity Co-Kriging surrogate model for ship hull optimization. Ocean Eng. 243:110239. doi:10.1016/j.oceaneng.2021.110239

Pellegrini R, Serani A, Broglia R, Diez M, Harries S. 2018. Resistance and payload optimization of a sea vehicle by adaptive multi-fidelity metamodelling. Paper 2018-1904, 2018 AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference. doi:10.2514/6.2018-1904

Pellegrini R, Serani A, Harries S, Diez M. 2017. Multi-objective hull form optimization of a SWATH configuration via design-space dimensionality reduction, multi-fidelity metamodels and swarm intelligence. MARINE 2017 Conf., Nantes, France.

Pereira FS, Eça L, Vaz G. 2017. Verification and validation exercises for the flow around the KVLCC2 tanker at model and full-scale Reynolds numbers. Ocean Eng 129:133–148. doi:10.1016/j.oceaneng.2016.11.005

Peri D, Campana EF. 2008. Variable-fidelity and surrogate modelling in simulation-based design. Proc. 27th Symp. Naval Hydrodynamics, Seoul, South-Korea.

Peri D, Campana EF, Tahara Y, Takai T, Kandasamy M, Stern F. 2010. New developments in simulation-based design with application to high-speed waterjet ship design. Proc. 28th Symp. Naval Hydrodynamics, Pasadena, USA.

Peri D, Kandasamy M, Tahara Y, Wilson W, Miozzi M, Campana EF, Stern F. 2012. Simulation-based design with variable-physics modeling and experimental verification of a waterjet propelled catamaran. Proc. 29th Symp. Naval Hydrodynamics. Gothenburg, Sweden.

Raven HC. 1992. A practical nonlinear method for calculating ship wavemaking and wave resistance. Proc. 19th Symp. Naval Hydrodynamics, Seoul, South-Korea.

Raven HC. 1996. A solution method for the nonlinear ship wave resistance problem. Dissertation, Delft Univ. Technology, Delft, Netherlands.

Raven HC. 2018. Minimising ship afterbody wave making using multifidelity techniques. 32nd Symp. Naval Hydrodynamics, Hamburg, Germany.

Raven HC. 2022. Ship hydrodynamics knowhow derived from computational tools: some examples. J Ocean Eng Mar Energy. doi:10.1007/s40722-022-00256-9

Raven HC, Scholcz TP. 2017. Wave resistance minimisation in practical ship design. Proc. MARINE 2017, Nantes, France.

Raven HC, Scholcz TP. 2019. An assessment of multifidelity procedures for ship hull form optimisation. 8th Int. Conf. Comp. Meth. Marine Eng., Gothenburg,Sweden.

Raven HC, Van Brummelen EH. 1999. A new approach to computing steady free-surface viscous flow problems. 1st MARNET-CFD workshop, Barcelona, Spain.

Scholcz TP, Gornicz T, Veldhuis C. 2015. Multi-objective hull form optimization using Kriging on noisy computer experiments, 6th Int Conf. Comp. Meth. Marine Eng.

Serani A, d'Agostino D, Campana EF, Diez M. 2018. Assessing the interplay of shape and physical parameters by nonlinear dimensionality reduction methods, 32nd Symp. Nav. Hydrodyn., Hamburg, Germany.

Serani A, Ficini S, Broglia R, Diez M, Grigoropoulos G, Bakirtzogou C, Papadakis G, Goren O, Danisman DB, Solak HP, et al. 2022. Resistance and seakeeping optimization of a naval destroyer by multi-fidelity methods, NATO STO-TR-AVT-354.

Toal DJJ. 2015. Some considerations regarding the use of multi-fidelity Kriging in the construction of surrogate models. Struct Multidiscipl Optim. 51(6):1223–1245. doi:10.1007/s00158-014-1209-5

Van Brummelen EH, Raven HC, Koren B. 2001. Efficient numerical solution of steady free-surface Navier-Stokes flow. J Comp Phys. 174:120–137. doi:10.1006/jcph.2001.6880

Van der Ploeg A, Eça L, Hoekstra M. 2000. Combining accuracy and efficiency with robustness in ship stern flow calculation. Proc. 23rd Symp. Naval Hydrodynamics, Val de Rueil, France.

Wackers J, Visonneau M, Serani A, Pellegrini R, Broglia R, Diez M. 2020. Multi-fidelity machine learning from adaptive and multigrid RANS simulations. 33rd Symp. Nav. Hydrodyn., Osaka, Japan.