Decoupled simulations of offshore wind turbines with reduced rotor loads and aerodynamic damping by Sebastian Schafhirt and Michael Muskulus.

Comments from Tor Anders Nygaard
Institute for Energy Technology, Norway
tor.anders.nygaard@ife.no

General comments

This paper is relevant because it addresses common, simplified approaches of taking the rotor into account for offshore wind turbines. Using existing validated models for the substructure, that also may include code checks, makes a lot of sense, if the rotor can be represented in a way taking the most important interactions with the substructure into account.

The impact of the paper is increased opportunity for engineers working with simplified rotor models to take more details of the rotor into account, without increasing the computational cost. Also, the article gives guidance on the level of uncertainty associated with simplified rotor models relative to a detailed rotor model.

Even though the article does not propose any new models, it provides new insights by building a systematic framework for optimization of linear dashpot dampers used in different flavors of decoupled simulations of offshore wind turbines. I think the article, with some extra effort, has the potential to be even more systematic, and even improve on the main conclusion.

The quality of the article is very good; it is easy to read, the results look plausible, and the discussions explain the results well. Regarding the logical structure of the article, I found it necessary to start with the detailed model, and then work “backwards” towards the simplified models.

I have suggested a “new” model 3b that seems to be a logical consequence of the findings in the paper. Examination of this model would take very little extra effort, if it has not been tested and discarded already.

Specific comments

During the review of the article, please keep in mind that the article and cited articles should contain all the information necessary for others to replicate the computations.

The article compares computational effort for several models, but I suspect there are different levels of parallelization at play. If this is the case, please mention that more specifically. I would like to see apples-to apples comparison of computational speed for the FEDEM model with the decoupled rotor model and pre-computed time-series of rotor forces, relative to the complete integrated FEDEM model. The effort for pre-computation of forces and optimization of the damping coefficients do not need to be included here. Does the relative difference change if this is run on a single processor?

For an offshore wind turbine, the computation of fatigue life at the most critical points in the structure usually means keeping track of ‘effective’ stress (such as von Mises stress) not only at a number of tower and foundation sections, but also at the different azimuth positions of those sections, through a large load case matrix under combined wind and wave loading. This poses a challenge already, as in the UpWind design basis (Fischer et al., 2010), when the load case matrix is lumped into an equivalent series of load cases with wind and waves always acting in one direction.

The starting point criteria for comparing fully coupled simulations with any simpler model should be a comparison of fatigue damage at a number of sections and azimuth positions. The first criterion for comparison of a full and simplified model in the article is tower top deflection in the downwind direction. Although some earlier work is cited, this seems too abrupt, keeping in mind the tower failures in the early days of onshore wind turbines, due to lack of consideration of shear stresses due to tower torsion. There are nonlinearities involved from tower top deflection components to total stress, and from effective stress to fatigue damage that should be kept in mind when defining simplified criteria for comparisons of models with different levels of complexity.

My comments will go in the opposite order of the paper, from complex to simple models. If this makes sense, please consider the sequence in the paper as well.

6 Decoupled models with full rotor loads

On page 17, line 12 and through section 6, the fully integrated model is compared with a decoupled model, applying time series of all 6 degrees-of-freedom (DOF) forces/moments from an isolated rotor computation to the tower top. Aerodynamic damping is applied through linear viscous dampers in all 6 DOF.

The optimization problem is here defined to adjust the six damping coefficients to match the fully integrated model by minimizing eq. 9. In eq. 9, the variance residuals of the 6 tower top deflections have equal weight, although they have different influence on the von Mises stress on the different parts of the structure. As mentioned earlier in the article, this is OK if a perfect match is found for all DOFs, but this is not always possible. Also, two time series of deflections may be very different, but have the same variance. Please comment if other optimization criteria have been tested as well, such as fatigue life due to all loads, in in selected hot spots, or important stress components at selected hot spots.

The comparisons of PSD plots of forces in the article is a very powerful tool for explaining the differences in the results, although the logarithmic scale mask some of the differences. The way the study is set up, also gives the opportunity to look at comparison examples in the time domain, since a perfect decoupled model and detailed rotor model in theory should give the same time history of loads. Please show examples of comparison. I would like to see 100s of each of the six load components at the tower root comparing 1) the integrated simulation, 2) the decoupled model with full rotor loads and six optimized dampers, and if possible, 3) the decoupled model with full rotor loads and no dampers, at 8 m/s, rated wind speed and 20 m/s.

In the EFL comparisons of figures 15 and 16, please include a comparison of “total” EFL, based on von Mises stress, on the upwind side of the tower at the tower root and just below the nacelle.

In the conclusions, it is stated that “the method to obtain the damping coefficients for all six viscous dampers used in this paper is computationally demanding and impractical for industrial applications. Hence, a more efficient method or formula to obtain optimal damping coefficients for viscous dampers is desirable”. I think one reason the optimization algorithm is demanding, is that we are trying to optimize one “convoluted” objective (eq. 9) by varying six coefficients simultaneously. Maybe the starting point for the optimization could be improved by tuning one and one or two and two coefficients by looking at outputs more directly linked to those coefficients. For example, select the damping coefficient for x motion, Cx to match variance of My at the tower bottom, Cy to match the variance of Mx at the tower bottom, set Cz equal to Cy (should be the same, except for the influence of shear and rotor tilt), and Crz to match the variance of torsion shear stress at the tower top. If approaches like this are already used, or tested and discarded, please comment on this.

How different are Cx found in DC 1 through DC 4?

On page 15, line 8, the full and decoupled models are the same as above, except that only one damper is applied in the x direction at the tower top for the decoupled model. This coefficient is tuned to match 1) tower top variance of displacement, 2) the variance of the overturning moment, and 3) the equivalent fatigue load for the overturning moment (My) at tower bottom. In figure 10, the damping coefficient optimized for matching EFL is very different from the damping coefficient optimized for matching tower bottom bending moment variance. Please examine why the results differ so much, because this would help us understand more of the challenges using simplified criteria such as tower top deflections. One point is that even if two models have the same variance of tower bottom bending moments, the EFL may be different due to the strong nonlinearity from loads to EFL (eq. 8). Then it is surprising that we seem to get identical results in the figures 11 and 12, for the three optimization methods, since the decoupled rotor model has been run with very different damping. Please double-check that the figures 11 and 12 are correct and comment on this. In the beginning of the article it is stated that the aerodynamic damping is very important for fatigue, but many of the comparisons show very small differences for significant differences of damping levels.

On page 17, line 3, it is stated that "The results of the decoupled model (DC 3) are not conservative anymore. In fact, the ratio for displacements for bending moments around the y-axis are up to four times higher for a wind speed of 12 m/s (not shown) and still almost three times higher for a wind speed of 20 m/s." I think it should be bending moments around the z-axis. Please clarify the use of "conservative" or simply state whether the loads for model DC3 are higher or lower than the integrated model. When I hear that a model is "conservative", I assume it has higher loads than the reference.

On the same page, it is stated that "It seems that especially the torsion (rotational movement around the vertical axis) of the OWT has to be damped." I think this is a very important comment. In my opinion, the next logical step is a model DC3b with dampers in tower top x translation, and tower top z rotation. It is even possible that the optimization could be carried out first as in DC3 (tuning only coefficient Cx), and then, keeping Cx constant, tune Crz. Even if the two coefficients after being initialized separately have to be fine-tuned simultaneously, should this give a significant speedup compared to model DC4. Would model DC3b then perform almost as good as model DC4? What would be the increase in computational cost over model DC3 (due to optimizing two damping coefficients vs. one)? How sensitive is the life-time fatigue at the tower root and tower top due to the torsion shear stress? Please do this, or comment if this has been tested and discarded already.

5 Decoupled models with reduced rotor loads

I think this section should examine or give more details of the cited articles on the importance of the different force and moment components on the life time fatigue at the hot spots in the tower root and top. This can be achieved in several ways, e.g.:

1. Re-run the integrated model through the fatigue load cases, and remove one force/moment component from the rotor at the time (except rotor thrust), or
2. Remove one stress component at the time at selected hot spots at the tower root and tower top.
3. Supplement the figures 3, 4, 6, 7, 11, 12 , 15 and 16 with “total” EFL based on all stress components, at the upwind side of the tower (since all load cases here have wind and waves aligned).

Model DC 1 applies only rotor thrust to the tower top. This model is attractive for users without access to a detailed rotor model, since overall CT curves are often available, and the time series can be pre-computed from eq. 1.

Model DC2, including tower top overturning moment, requires a detailed rotor model. This step up in complexity relative to model DC1 is maybe a bit understated. Once we have the six force/moment components available from a detailed rotor model, the extra effort over model DC 2 is to optimize the additional damping coefficients, for each windspeed bin.

A separate file contains the article, supplemented with highlights and sticky notes that provide some further (and overlapping) comments.

Technical corrections

A separate file contains the article, with highlights in yellow for checking of language and spelling, and comments in sticky notes.