This case study was created as part of
the free and voluntary UberCloud Experiment

Case Study Author – Henrik Nordborg, Adrian Rohner, HSR

CFD Simulation of Vertical Axis Wind Turbines

Summary

- The Cloud is ideally suited for Computationally Demanding Parameter Studies. Since each wind speed and tip-speed ratio runs independently, progress is limited only by core-count and licenses.
- CD-adapco's flexible Power on Demand license is ideal for on-demand work in the cloud.
- UberCloud's customized cloud cluster solution comes with all software packages pre-installed. There is no investment in HPC equipment or expertise needed.
- With most software vendors catching up to the idea of software as a service, the authors expect a large part of engineering simulations to run in the Cloud in the near future.

"Cloud based computing extends our potential to design/develop better products. By utilizing this potential the products can be optimized with a much faster pace and higher quality."

MEET THE TEAM

- ***End user and Team Expert -*** Henrik Nordborg, Adrian Rohner, Hochschule fur Technik in Rapperswil (HSR), Switzerland
- ***Software Provider -*** CD-Adapco providing STAR-CCM+
- ***Resource Provider -*** Microsoft Azure with UberCloud STAR-CCM+ software container
- ***Technology Experts -*** Fethican Coskuner, Hilal Zitouni, and Baris Inaloz, The UberCloud Inc.

USE CASE

Vertical Axis Wind Turbines (VAWT) represent an interesting alternative to the more conventional horizontal axis design, promising a simpler and more robust design with lower costs. Furthermore, the fact that the turbine does not have to be turned to face the wind makes it ideally suited for difficult locations with constantly changing wind direction.

Figure 1: Large scale vertical axis wind turbine (VAWT) (courtesy of Agile Wind Power, Switzerland).

Aerodynamically, the VAWT is more difficult to model than the horizontal axis design. The problem is that the blades see very different wind velocities and flow patterns at the upwind and downwind positions. A semi-analytical approach to estimate the performance of a VAWT, the Double Multiple Stream Tube method (DMST), is therefore not very accurate can only be used for rough estimates. A more promising approach is to use CFD simulations, which fully resolve the flow pattern inside and around the turbine. Obviously the best results are obtained from a fully transient simulation of the entire structure, including the all the struts in additions to the blades. As this is very time consuming, it makes sense to optimize the blades using two-dimensional simulations of a cut through the turbine.

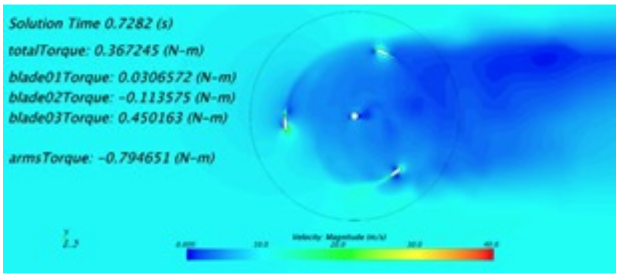

Figure 2: Two-dimensional cut through the simulation of a VAWT.

A more difficult problem is the proper handling of stall. It is not very difficult to compute lift and drag of an airfoil for small angles of attack. For larger angles, however, the flow will detach from the blade, leading to stall and a significant reduction of the lifting force. It is well known that standard turbulence models have problems capturing this transition. On the other hand, the effect on the total torque might not be dramatic, as the main torque is produced from blades with a fairly small angle of attack.

In order to validate the simulations, a small VAWT and a simple wind tunnel have been developed at the HSR. The blades can be manufactured using 3D printing, making it easy to test new designs. The goal is to compare measurements with accurate transient 3D simulations in order to determine the accuracy which can be obtained from numerical simulations.

Figure 3: The small VAWT used for the validation studies.

The power generated by a wind turbine can be written

$$P = C_p(Re,\lambda)\frac{1}{2}pv^3A$$

where ***p*** is the density of air, ***v*** is the wind speed, and ***A*** is the swept area. The efficiency is given by the dimensionless power coefficient, $c_p(Re,\lambda)$ which depends on the design of the wind turbine, the Reynolds number, and the tip speed ratio (TSR)

$$\lambda = \frac{wR}{v}$$

In the case of a VAWT, the power generated by the turbine is not a constant but depends on the momentary position of the blades. This makes it necessary to simulate a number of complete revolutions of the turbine in order to obtain the average power. Furthermore, in order to determine the dependence on the Reynolds number and the TSR, we need to run the simulations at different wind speeds and different angular velocities. This means that a large number of simulations have to be computed during a short period, making the application ideal for Cloud computing.

TECHNOLOGY: Running STAR-CCM+ in the Cloud

This STAR-CCM+ wind turbine simulation with about 20 million cells first ran on the in-house 3-year old 96-core computing cluster of HSR, the Hochschule fur Technik in Rapperswil, Switzerland, and we decided to repeat the same simulation as an UberCloud Experiment on 96 cores of Microsoft Azure A8/A9 compute instances interconnected with Infiniband with UberCloud's STAR-CCM+ software container. This case study presents a quite common situation: in-house hardware ages quickly, while on Azure you always get the latest and greatest hardware. In fact, the simulation of the wind turbine on the Azure A8/A9 cluster ran 22% faster.

One of the novelties in this experiment was to use UberCloud's new STAR-CCM+ software containers, https://www.TheUberCloud.com/containers/, which are ready-to-execute packages of HPC software. These packages are designed to deliver the tools that an engineer needs to complete his task in hand. The ISV or Open Source tools are pre-installed, configured, and tested, and run on bare metal without loss of performance. They are ready to execute, literally in an instant, with no need to install software, deal with complex OS commands, or configure. This UberCloud Container technology allows wide variety and selection for engineers because the containers are portable from server to server, Cloud to Cloud. The Cloud operators or IT departments no longer need to limit the variety, since they no longer have to install, tune and maintain the underlying software. They can rely on the UberCloud Containers to cut through this complexity. This technology also provides hardware abstraction, where the container is not tightly coupled with the server (the container and the software inside isn't installed on the server in the traditional sense). Abstraction between the hardware and software stacks provides ease of use and agility that bare metal environments lack.

CHALLENGES

The main challenge in this application is running a large number of very time-consuming simulations in order to perform parameter studies. For each set of parameters, we ran 10 full revolutions of the turbine, requiring 3600 time steps per simulation. The geometry was meshed using the built-in polyhedral mesher of STAR-CCM+, resulting in meshes with 4.7 million elements. The boundary layer around the blades was resolved using inflation layers, and the SST k-w model with wall functions was used to model turbulence.

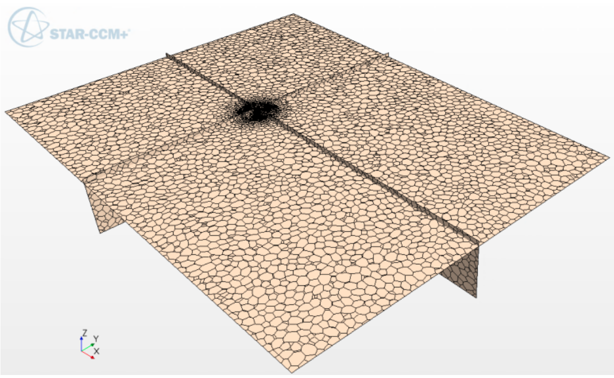

Figure 4: A polyhedral mesh used in the simulation.

PROCESS AND BENCHMARK RESULTS

The computations were performed on the in-house Windows Cluster at the HSR and the in the Cloud using Microsoft Azure. In both cases, 96 cores and 40 Gbit/s Infiniband networking were used. The compute nodes at the HSR are equipped with dual Intel Xeon X5645 processors running at 2.40 GHz and using Windows Server 2012 with the Microsoft HPC Pack 2012 R4. The hardware in the Cloud is provided by Microsoft Azure and used A8/A9 nodes. Given the large number of the elements and the highly efficient CFD software, the scaling up to 96 cores was very good in both cases. However, it turned out that the Cloud was significantly faster, requiring only 25 hours per simulation (10 revolutions) as compared to 32 hours with the in-house cluster. This represents a speedup of 22%.

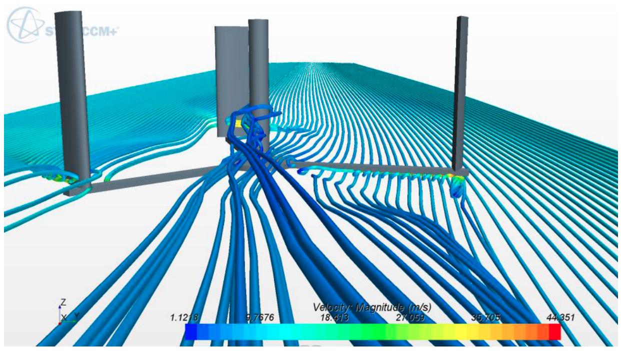

Figure 5: Streamlines illustrating the influence of the blades on the flow.

BENEFITS

The Cloud is ideally suited for this kind of computationally demanding parameter studies. Since each simulation with a particular wind speed and tip-speed ratio can be run independently, progress is only limited by the number of available cores and licenses. However, the CD-adapco licensing model using Power on Demand license tokens is very flexible and well suited for on-demand work in the cloud.

CONCLUSIONS

- We showed that the Microsoft Azure based UberCloud cloud solution is a beneficial solution for CD-adapco users who have the need to deliver their simulation results in a much faster time manner.
- To use cloud based cluster computing, there is no investment in in-house HPC equipment or expertise needed, since UberCloud offers customized and handy cloud cluster solutions with all requisite software packages pre-installed.
- With most software vendors catching up to the idea of software as a service, we expect a large part of engineering simulations to run in the Cloud in the near future.

ABOUT UBERCLOUD

UberCloud makes it easy to run your simulations on powerful cloud infrastructure.

No more compromises on mesh quality or model fidelity because of hardware limitations. With UberCloud's flexible software platform and network of cloud partners, you get on-demand access to major providers such as Microsoft Azure, HPE and others. Choose from a variety of secure data centers, and hardware options such as InfiniBand, GPUs etc.

Unleash the full power of your analysis software and boost confidence in your results.

With over 200 technical-computing-as-a-service case studies, UberCloud has the experience, software platform and partnerships required for your success.

Engineers and scientists rely on UberCloud to manage the complexity of cloud and software operations, so they can focus on their analysis.

FOR MORE INFORMATION

2310 Homestead Rd. Suite:C1-301 Los Altos CA 94024

help@theubercloud.com
www.theubercloud.com/help