WIND AND WAVE MODELLING IN AN OPERATIONAL
WARNING SYSTEM AGAINST FLOODING

Gerbrant van Vledder, Lykle Voort, Gijs van Banning

Alkyon Hydraulic Consultancy & Research, PO Box 248, 8300AE Emmeloord, The Netherlands
Email: vledder@alkyon.nl, voort@alkyon.nl and banning@alkyon.nl

Elizabeth (Ellen) J. Claessens and Yede Bruinsma
Rijkswaterstaat, Institute for Inland Water Management and Waste Water Treatment RWS RIZA, PO Box 17, NL- 8200 AA, Lelystad, The Netherlands.
Email: E.J.Claessens@riza.rws.minvenw.nl, and Y.Bruinsma@riza.rws.minvenw.nl

1. INTRODUCTION

Dikes protect large parts of the Netherlands. Under normal circumstances these dikes are high enough to protect the low-lying land against flooding. During storms the dikes are attacked by a combination of high water levels and waves leading to dangerous wave run-up levels. All dikes are designed to withstand extreme conditions with return periods up to 2,500, 4,000 and 10,000 years (depending on the location). Fortunately, such extreme conditions have not yet occurred. In the event that a storm approaches the Netherlands, several warning systems are used to warn against possible dangerous situations. In this paper, a prototype of a recently developed warning system for the inland lakes in the former Zuiderzee is described.

In 1932 a 30 km long closure dike was build forming the IJsselmeer. In this lake various dikes and polders were created forming several lakes of which the largest are the IJsselmeer and Markermeer. Several rivers flow into these lakes of which the river IJssel (a branch of the river Rhine) is the largest. The IJsselmeer and the North Sea are connected to each other with discharge sluices enabling the regulation of the water level in this lake. An outline of these lakes and polders is shown in Figure 1. Several local authorities manage the dikes around these lakes.

Figure 1: Main inland water lakes in the IJsselmeer area.

In 1986 a special service of the Dutch Ministry of Transport and Public Works (Rijkswaterstaat) was created to issue warnings for dangerous storm situations in the IJsselmeer area. This service, the Warning service Dikes IJsselmeer area (WDIJ), issues warnings for the dike authorities when a risk of flooding is expected to occur. The current warning system is a database containing results of pre-computed storms. These storms cover a large number of wind and water level combinations for which storm levels (i.e. the sum of the water level and wave run-up) have been computed. In the event of a storm this system is fed with the still water level of the various lakes, the river discharges and the expected maximum wind speed and direction. Subsequently, a search is made in the database to find the most appropriate standard storm, followed by retrieving the pre-computed storm levels. The result may consist of a number of dike sections where the storm level exceeds a preset alarm level. If necessary, the dike authorities are warned to take appropriate safety measures.

The advantage of the existing system is that it is quick and robust. The disadvantage of this system is that it is based on highly schematised stationary storms using homogeneous wind fields. In practice both waves and water levels are rather sensitive to changes in wind speed and direction. Therefore, this simple approach may lead to inaccurate and consequently unreliable predictions of dangerous situations along the dikes.

To improve this situation a prototype of a new operational warning system has been developed. It is scheduled that this new system will become operational in the storm season 2006/2007. This new system uses predicted inhomogeneous and instationary wind fields, real-time measured water levels and river discharges to simulate storm levels along the dikes. The measured data are used to optimise the hindcast part of the flow model to create an optimal initial starting condition for the predictive part. The new system consists of a coupled set of simulation models for the calculation of the wind fields, currents and water levels, wind-generated waves and wave run-up. Wind fields are obtained from the Royal Netherlands Meteorological Institute (KNMI) and processed in a high-resolution downscaling module (Verkaik *et al.*, 2006) to obtain a fine grid spatial variation of wind speeds and directions. Water levels are simulated with the flow-model WAQUA (Stelling, 1984) using the downscaled winds as input. The wind-generated waves are simulated with the third-generation wave prediction model SWAN (Booij et al., 1999). The downscaled winds and computed water levels are input for SWAN. Finally, the wave run-up is computed with a dedicated program using the water level, the incident wave conditions and the dike profile and roughness as input.

A crucial requirement of any operational warning system is the speed with which warnings can be issued. For the present system it is required that every 3 hours updated information on future storm levels (up to 48 hours) should be known within 1.5 hours after the start of the compound simulation system. This strict requirement forced the development team to focus on an efficient implementation of the downscaling module, WAQUA and SWAN. A PBS job scheduler is used to distribute the tasks over a cluster of PC's. MPI is used for parallelisation, meeting the CPU demanding requirements of the wind, flow and wave modules.

The development of the new operational system is the main focus of this paper. Section 2 describes the elements of the new system. The principles of the downscaling module are presented in section 3. The implementation of the SWAN model in this system is described in section 4. Section 5 describes the operational use of the new system, the quality of the warning system, and some future developments of the new system. Conclusions are presented in Section 6.

2. STRUCTURE OF THE OPERATIONAL WARNING SYSTEM

The new warning system consists of 4 simulation programs that basically convert wind fields into storm levels along the dikes in the IJsselmeer area. The general structure of this system is shown in Figure 2. The basic input of the system consists of the weather prediction from the HIRLAM model of the KNMI together with the measured water levels in the lakes and the measured river discharges in the rivers flowing into the IJsselmeer. The key elements of the prediction system are:

- the program for the downscaling of the winds,
- the flow model WAQUA for the simulation of currents and water levels,
- the SWAN wave model for the simulation of wind-generated waves, and
- a module for the computation of the wave run-up.

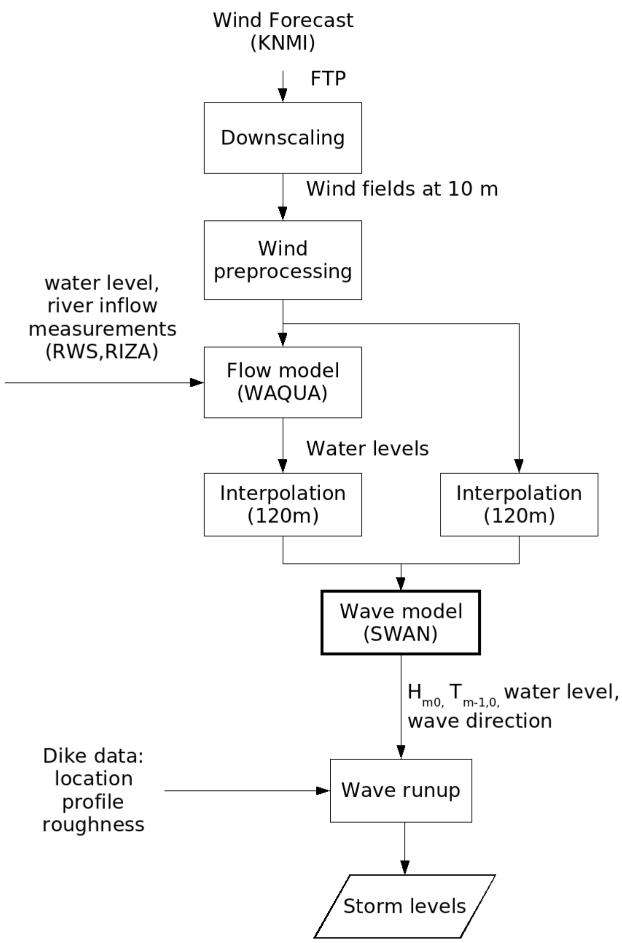

Figure 2: Structure of the prediction system (boxes indicate modules, arrows indicate data flows).

In the first step wind information is obtained from the KNMI via an automated FTP-link. Predictions up to +24 hours on an 11 km resolution are obtained from the HIRLAM 11 km model every 3 hours at a hourly interval. Wind predictions up to +48 hours on a 22 km resolution are obtained from the HIRLAM 22 km model every 6 hours at a 3 hourly interval. These wind fields cover an area much larger than the Netherlands.

In the second step the available HIRLAM wind fields are processed in the downscaling module to obtain fine-resolution wind fields on a user-defined resolution. In the present system, downscaled wind fields are generated for every hour up to +24 hours on a spatial resolution of 500 m. Linear interpolation in time is used to convert the wind to the required time steps of the flow model. For the period +24 hours to +48 hours, downscaled wind fields are produced every 3 hours. Detailed information on the downscaling is presented in section 3 of this manuscript.

In the third step the hydrodynamic flow model WAQUA (Stelling, 1984) is used to simulate the time- and space-varying currents and water levels in all lakes in the IJsselmeer area. Due to the presence of a dike between the IJsselmeer and the Markermeer two separate flow models have been made. One model for the IJsselmeer, Ketelmeer and Vossemeer and one model for the Markermeer, Gooimeer and Eemmeer (See Figure 4). The spatial resolution of the curvilinear WAQUA models varies from 10 m in the IJsseldelta to 1000 m in the middle of the IJsselmeer. The internal time step of the WAQUA model is 20 seconds. The IJsselmeer model is the most critical application of the flow model, as it has to deal with the discharge sluices in the Afsluitdijk and with the IJssel discharge. Especially for this model an integrated model data approach has been chosen. The 3 day simulations exist of a hindcast part, covering the past 24 hours and a forecast part, covering the coming 2 days. Essential is that the starting point for the forecast is as good as possible representing the actual situation. Measurements for the last 24 hours at a number of 4 water level stations around the IJsselmeer are used to give a balanced correction for the average water level at every time step. Test simulations show that this approach gives a more reliable prediction of the water levels for the coming two days. Water level fields are generated every hour on a square grid with a spatial resolution of 100 m for the larger lakes and on a resolution of 50 m for the smaller lakes.

In the fourth step the SWAN wave model is run in stationary mode to simulate the wave conditions at a number of selected prediction times. Due to the geometry of the various lakes and the presence of dams, four SWAN models have been made. One model for the IJsselmeer and one model for the Markermeer. The third model for the Ketelmeer and Vossemeer is nested in the model for the IJsselmeer. The fourth model for the Gooimeer and Eemmeer is nested in the Markermeer model. In the present set-up wave conditions are computed for +3, +6, +12, +18, +24, +36 and +48 hours. In addition, the wave conditions are computed at the (hourly) moment of maximum wind speed (determined by a small dedicated program) in the first 24 predictive hours. The wind fields obtained from the downscaling and the water levels from the WAQUA model drive the SWAN model.. The basic output of the SWAN model consists of the significant wave height H_{m0}, the spectral wave period $T_{m-1,0}$, the mean wave direction θ and the water level at a selected number of output points. To allow checking the system, additional output can be generated of the wind speed and direction at these output points, and fields (in the form of SWAN block files) of the significant wave height H_{m0}, the spectral wave period $T_{m-1,0}$, the mean wave direction θ and the water level. A detailed description of the setup of the SWAN models for the different lakes is presented in section 4 of this manuscript.

In the fifth step the wave run-up is computed with a dedicated program. This program contains wave run-up equations based on physical model investigations by Van der Meer (1997). The run-up module is run for a selected number of locations and for the same moments of time as the wave model SWAN. The input for the wave run-up module consists of a dike profile and the incident significant wave height, spectral wave period, incident wave direction and the water level. It is noted that the input water levels are obtained via the SWAN model and not from the WAQUA model. The reason is the simplification and unification of data interaction. The output of this module consists of the wave run-up level at the selected output locations.

In the sixth and last step the simulated storm levels are compared with pre-set threshold values for the storm levels. These thresholds are the alarm levels. If one or more of these alarm levels are exceeded, a warning is issued to the local dike authorities.

3. PRINCIPLE OF THE DOWNSCALING MODULE

Wind speed forecasts by numerical weather prediction(NWP) models in heterogeneous terrain lack locally correct representations as they are derived using grid-box averaged roughness lengths. For example, the most detailed HIRLAM model used in our warning system has a resolution of 11 km. The local accuracy of NWP-wind speed forecasts is improved using a downscaling method developed at the Royal Netherlands Meteorological Institute (Verkaik, 2006, Verkaik et al. 2005). The text below is largely taken from Verkaik et al. (2006).

The method includes a simple two-layer model of the atmospheric boundary layer, used in combination with a high-resolution surface roughness map. The basis of the two-layer (2L) model is the assumption that sub-grid wind speed variations are caused mainly by surface roughness changes. The 2L model is used to post-process direct NWP-model output. The model comprises of a surface layer and an Ekman-layer. In the surface layer vertical wind speed transformations are performed using the logarithmic wind speed profile. In the Ekman-layer geostrophic resistance laws are applied.

The downscaling of NWP winds using the 2L-model consists of the following steps. First, the 10 m wind of the NWP is transformed to a wind at the top of the Ekman layer using the roughness from the NWP model and geostrophic drag laws. In this step neutral stability of the Atmospheric Boundary Layer (ABL) is assumed. Next, the wind speed at the top of the Ekman layer is interpolated bi-linearly to the target location. Then, at the target location, the wind speed at the top of the surface layer is computed, followed by a further downward transformation using the local sub-grid roughness. Experience shows that errors in the upward transformation due to the assumption of neutral stability are usually counter-balanced in the downward transformation (De Rooy and Kok, 2004).

The roughness map is derived from a land-use map and a simple footprint model. The roughness lengths are wind direction dependent. For the Netherlands this information is available with a directional resolution of 5° and a spatial resolution of 500 m. The footprint area of the Ekman-layer extends further upstream than that of the surface layer. The roughness lengths compare well to those derived from gustiness analysis for station locations. The adjustment of the surface wind after a roughness transition as modelled by the two-layer is similar to that of internal boundary layer models.

The NWP-model wind and the downscaled wind have been evaluated at a number stations in the Netherlands. Verification shows that the downscaling reduces the NWP surface wind speed error significantly, largely in terms of bias. The quality of the downscaled wind, however, depends highly on the quality of the high-resolution roughness map.

An example of the downscaling is presented in Figure 3. The left panel shows the spatial variation of the wind field in the IJsselmeer area, obtained with the HIRLAM 11 km model. The right panel shows the wind field based after the downscaling. The arrows indicate the wind direction.

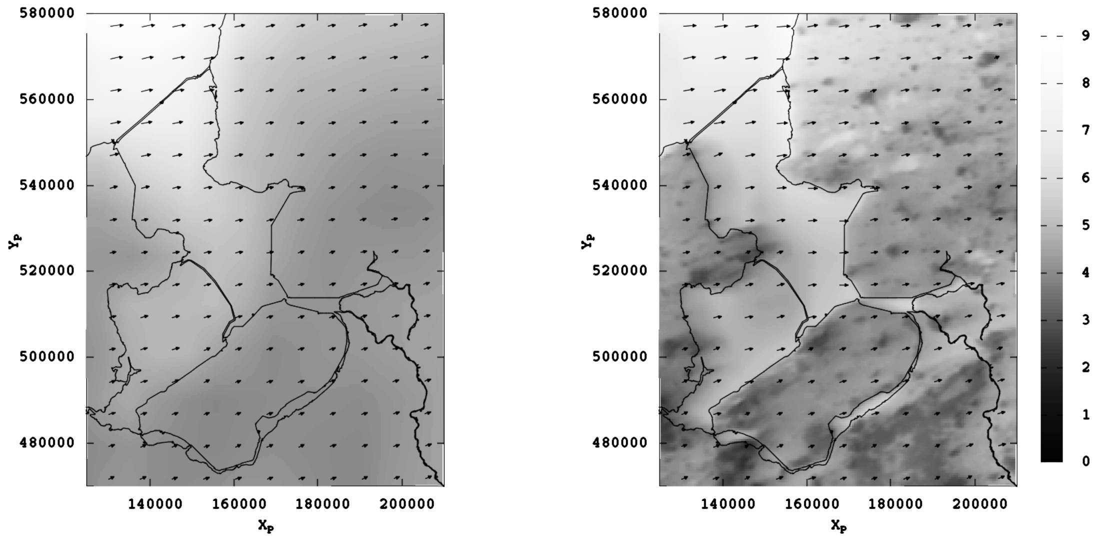

Figure 3: Spatial variation of the wind speed and direction based on the HIRLAM 11 model (left panel) and based on the downscaling (right panel). July 10, 2006, 5:00 hours.

The results in Figure 3 clearly show that the downscaling reveals many (physically plausible) local variations of the wind speed. As expected the relatively low wind speeds occur in forested and urban areas. Also the gradual increase in wind speed after roughness transitions, especially after land-water boundaries, can be seen. A striking feature is that the wind speed on the small lakes (e.g. the Ketelmeer) becomes much higher than in the HIRLAM 11 model winds.

In general the quality of the downscaled winds is much higher than the relatively crude HIRLAM winds. However, experience shows that the downscaled winds are underestimated in storm situations. This needs to be remedied since accurate wind speed predictions are essential for the prediction of water levels and waves. Ongoing verification and calibration studies are performed to improve the downscaling.

4. SWAN MODEL IMPLEMENTATION

For the IJsselmeer area four SWAN models have been developed. Each of these models covers one or more lakes. The largest model covers the IJsselmeer. Nested in this model is a model for the Ketelmeer and Vossemeer. The third model covers the Markermeer. Nested in this model is the fourth model for the Gooimeer and Eemmeer.

For each SWAN model attention was given to the following aspects:

- stationary or instationary mode;
- bottom topography;
- computational grids and nesting;
- physical settings;
- numerical settings;
- correction factors;
- coupling with downscaling module;
- coupling with flow model;
- coupling with the wave run-up module.

4.1 Stationary versus instationary

By definition, natural storm systems change in space and time. This implies that the wind fields are instationary and inhomogeneous. These features are accounted for in the wind model HIRLAM and in the flow model WAQUA. For consistency reasons also the SWAN model should be run in instationary mode. However, Claessens *et al.* (2002) conclude that it is not (yet) feasible to use SWAN instationary in the warning system. The computational effort required in the non-stationary mode is too large and the quality of the results is too poor. Additional investigations revealed that for the IJsselmeer the evolution of the wave field in a storm could also be simulated using a sequence of stationary SWAN simulations. This is acceptable when the time step between the simulation times of succeeding SWAN simulations is equal to or larger than 30 minutes.

4.2 Bottom topography

For each lake detailed digital bottom topographies are available. For the IJsselmeer and Markermeer this information is available on a resolution of 40 m, for the smaller lakes this information was available on a 10 m grid. Dams are schematised as obstacles in the SWAN model. Examples are the dam towards the former island Marken and the dams in the new city area IJburg in the southwest corner of the Markermeer.

4.3 Computational grids and nesting

The main challenge in developing the SWAN models is to find a balance between accuracy and computational requirements. Therefore, various sensitivity studies have been performed to find an acceptable spatial resolution, required number of iterations and efficient nesting.

The purpose of the operational warning system is to provide storm levels near the dikes of the inland lakes. Many of these dikes have a shallow foreland where the waves are affected by shallow water effects. An accurate prediction of the wave conditions in these areas would require a high spatial resolution in the order of 10 m to 40 m. Such a resolution is not needed for the upwind and central area of the IJsselmeer and the Markermeer. Here, a resolution of about 200 m would be sufficient.

Choosing an overall resolution of, say 40 m, for the IJsselmeer is possible but this leads to unacceptable memory requirements (4 Gbyte) and unfeasible long simulation times (order of days). Another option is to apply high-resolution nested grids near the boundaries of the dikes. Such nested grids need only be activated for downwind situations. Using nested grids enhances the complexity of the operational warning system since many intermediate data files need to be stored and managed. Therefore, this option was not (yet) considered. Using a curvi-linear grid with a spatially varying resolution was considered, but it is topologically impossible to create such a grid with a finer resolution along the boundaries than in the central part.

For simplicity regular rectangular grids were created for the different lakes and the smaller lakes are nested in the larger lakes. The IJsselmeer, Ketelmeer and Vossemeer are connected to each other and they could be combined in one grid. Since the smaller lakes have different length scales than the larger lakes it is more efficient (at least in SWAN) to use a separate grid for the Ketelmeer and Vossemeer. This nested grid has a finer resolution than the IJsselmeer. Similarly, the grid for the Gooimeer and Eemmeer is nested in the larger Markermeer. The outline of these four grids is shown in Figure 4. This figure also shows the (present) output points of the warning system in the IJsselmeer.

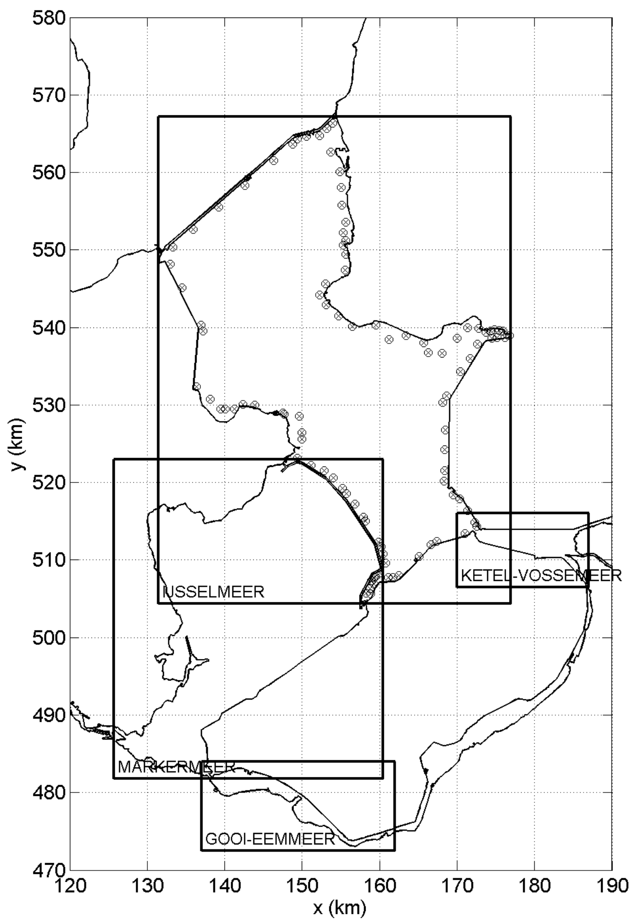

Figure 4: Overview of the SWAN model grids and output points in the IJsselmeer

The spectral resolution was equal for all SWAN models. In line with the SWAN guidelines a directional resolution of 10° was used and a geometric frequency spacing from 0.08 Hz – 1.5 Hz with a 10% increase between succeeding frequencies.

The optimal spatial resolution was determined by performing a series of simulations with different spatial resolutions and a fixed number of 50 iterations. Computed values of the significant wave height H_{m0} and spectral period $T_{m-1,0}$ on the output points for the warning system were exported to data files. Subsequently, for each output point the variation of a wave parameter as a function of grid resolution was determined relative to the value with the highest resolution. In this way the effect of the grid resolution on the computed wave parameters could be determined. Next, an optimal grid resolution was chosen, taking into account the required computational effort. For the IJsselmeer simulations were performed on a resolution of 120 m, 160 m, 200 m and 240 m. The wind speed was 25 m/s and 4 wind directions were considered (60°, 150°, 240° and 330°).

Figure 5 shows an example of the relative variation of the significant wave height H_{m0} and spectral period $T_{m-1,0}$ as a function of grid resolution. Each line corresponds to one output point. The thick lines refer to downwind output points, whereas the dashed output points refer to upwind output points. The results in this figure clearly show that the relative variation of both integral wave parameters is less than 2% for the downwind points, whereas for upwind points the relative variation may exceed 10%. These findings demonstrate that grid resolution mainly affects the results for short fetches. Since the warning system is meant for high wave conditions that only occur on the downwind side of the lakes we may only consider the solid lines to choose an optimal spatial grid resolution. Similar figures (not shown here) were made for the other wind directions in the IJsselmeer, the Markermeer and for the Ketelmeer and Vossemeer, and in the Gooimeer and Eemmeer.

For the IJsselmeer the spatial resolution was set to 200 m, for the Markermeer to 120 m and for the Ketelmeer-Vossemeer and Gooimeer-Eemmeer to 50 m.

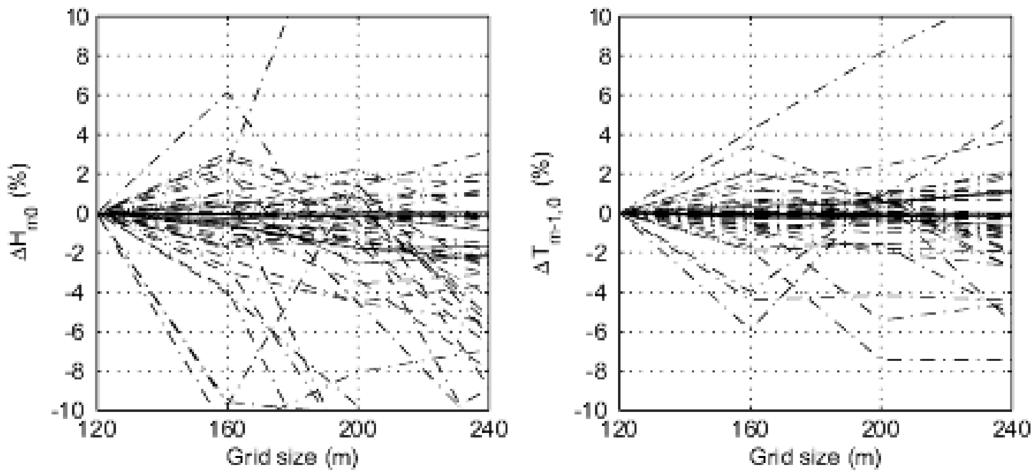

Figure 5: Relative variation of the significant wave height H_{m0} and spectral period $T_{m-1,0}$ as a function of grid resolution for all 105 output points. Solid lines for downwind output points, dashed lines for upwind output points. Wind speed 25 m/s and wind direction 240°. Each line represents one output point.

4.4 Physical settings

The inland water lakes have typical depths in the range of 5 to 10 m. This implies that during storm conditions shallow water effects will play a role in the evolution of the wave field. In some locations very shallow foreshores exists where waves may break due to depth limitations. Another feature is that dikes, thus preventing the inflow of swell waves from the Waddenzee, surround all the lakes. This implies that initial wave growth occurs along all upwind boundaries.

Above considerations imply that the so-named deep-water source terms for wind, whitecapping dissipation and non-linear four-wave interactions should be activated. In addition, the so-named shallow water source terms for bottom friction, surf breaking and non-linear triad interactions should be activated

4.5 Numerical settings

An important feature of the SWAN model is its numerical scheme to reach a solution of the wave energy equation by iteration. The required number of iterations depends on the user specified convergence criteria. These criteria specify the number of points where certain relative differences between integral wave parameters should be obtained. Detailed information about the convergence criteria in SWAN can be found in Zijlema and Van der Westhuysen (2005). In addition, the maximum number of iterations can be specified. Especially, this latter option is needed in an operational system where an answer must be provided within a fixed time window.

The convergence behaviour of SWAN was investigated by performing simulations with a large fixed number of iterations (typically 50) and outputting the values of the integral wave parameters in a number of test points per iteration. Figure 6 shows an example of the iteration behaviour of SWAN for an IJsselmeer simulation in a downwind grid point (x=150.6 km, y=544.5 km) of the significant wave height H_{m0}, the spectral periods T_{m01} and $T_{m-1,0}$ and the mean wave direction θ. The non-directional parameters are normalized with the value after 50 iterations. Similar simulations and figures were made for the other models and used to determine the optimal number of iterations.

Figure 6 shows the typical convergence behaviour of SWAN in the IJsselmeer. The significant wave height reaches convergence (difference less than 1% with value after 50 iterations) after 30 iterations, whereas the period measures typically require 50 iterations. The mean wave direction varies a few degrees in the first 30 iterations. To ensure sufficiently accurate results of SWAN in the IJsselmeer typically 50 iterations are needed.

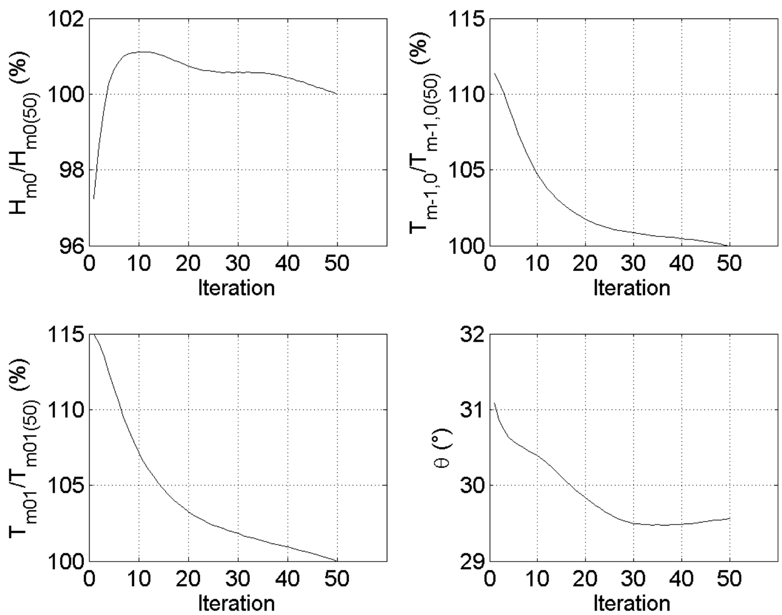

Figure 6: Convergence behaviour of four integral wave parameters for a point in the IJsselmeer. Location x=150.6 km, y=544.5 km.

3.6 Coupling with downscaling module

The wind fields from the downscaling module are available on a spatial resolution of 500 m. The area of these wind fields covers all computational grids of the various SWAN models. For consistency reasons these wind fields are also used by the flow model WAQUA.

3.7 Coupling with flow model WAQUA

In theory currents and waves interact in two-ways. The flow model provides per time step a two-dimensional current and water level field. These quantities can be accounted for in the SWAN model using linear wave theory. On the other hand in areas with strong wave dissipation wave-driven currents and set-up may be generated. Both types of interactions may be accounted for in the operational system.

For efficiency reasons, the present operational systems only accounts for a simple one-way coupling in which the spatial variation of water levels is accounted for in the SWAN model. To that end the curvi-linear WAQUA model provides water level fields on a regular rectangular grid suitable for the SWAN model. Currents are not (yet) included in the coupling since we expect that they have a much smaller effect on the wave conditions than spatial variations of the water level.

3.8 Coupling with the wave run-up module

The SWAN model is coupled with the wave run-up module via an output data file containing the computed significant wave height H_{m0}, spectral period $T_{m-1,0}$, mean wave direction θ and the water level h in all output points of the warning system. It is noted that the water level h can also be retrieved directly from the WAQUA model, but this would create an additional unnecessary data link in the whole system.

3.9 Correction factors

The present SWAN model implementations use default physical settings (as recommended in the SWAN manual) and a sufficiently large number of iterations. In addition, the spatial resolution was chosen in such a way that discretisation errors in the predicted significant wave height H_{m0} and spectral wave period $T_{m-1,0}$ are less than 2%. However, these choices are no guarantee that the SWAN models produce accurate results.

The quality of the wave model results can be improved by calibration of the optimal parameter settings for the physical processes. In such a procedure the bulk difference between predicted and observed wave parameters is minimized. Calibrating the various SWAN models is rather complex. Instead correction factors were determined that will be applied on predicted wave quantities.

Correction factors for the significant wave height H_{m0} and spectral wave period $T_{m-1,0}$ were determined on the basis of a comparison of routinely observed and simulated parameter values. Only simultaneously available wind and wave data could be used. The maximum observed wind speed was 22 m/s. A statistical analysis was performed to determine systematic differences in these parameters. Figures 7 and 8 show the comparison between observed and simulated values and the derived correction factors for the IJsselmeer. The results showed a weak distinction in correction factors for waves with significant heights smaller than 1 m and those with heights larger than 1 m. This distinction may be related to different model behaviour per wind direction such that the geometry of the IJsselmeer plays a role. A further (wind direction and/ or fetch length related) refinement of these correction factors is scheduled.

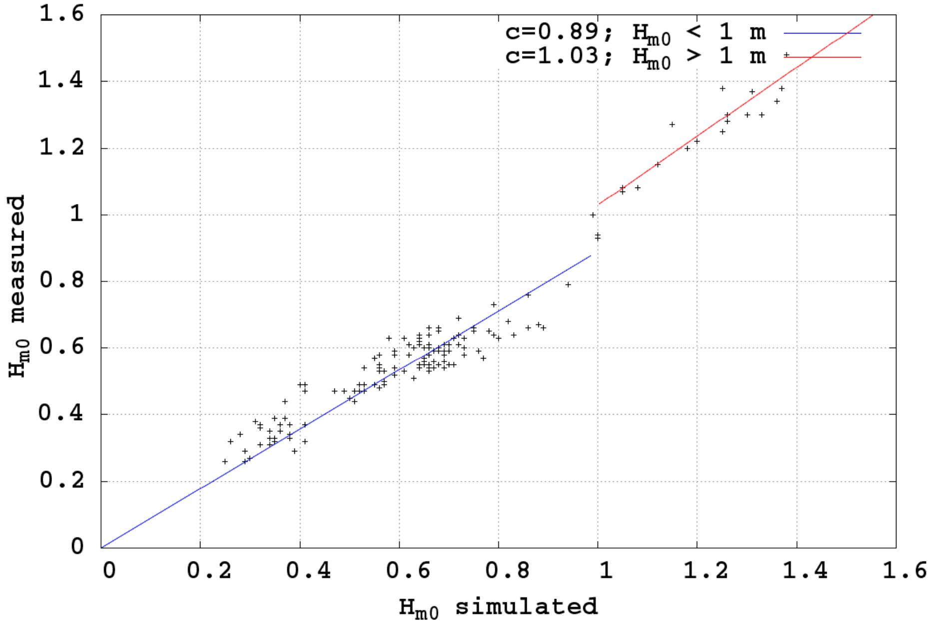

Figure 7: Comparison of the simulated significant wave height and the measured significant wave height; the measurements were performed during three storm situations. Correction factors based on a piece-wise least squares analysis.

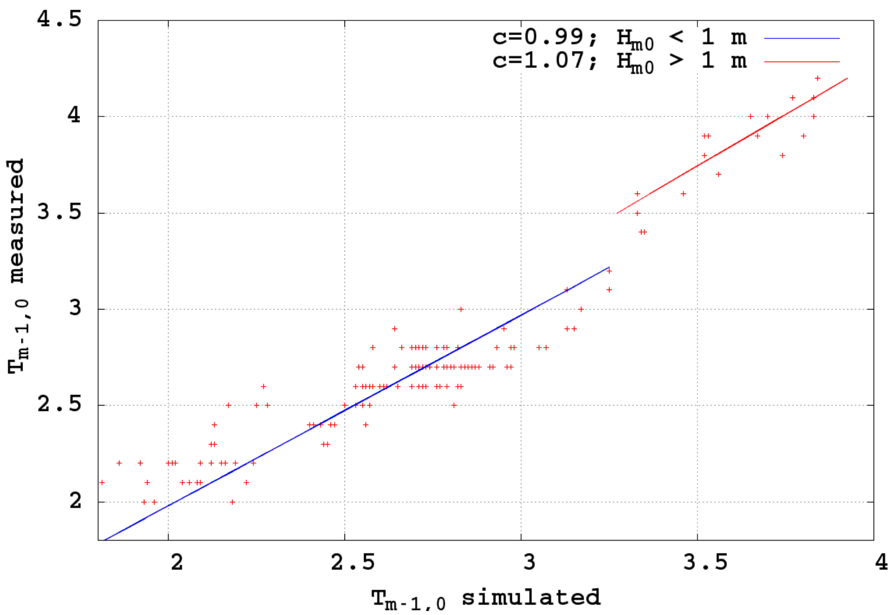

Figure 8:Comparison of the simulated spectral wave period and the measured spectral period; the measurements were performed during three storm situations. Correction factors based on a piece-wise least squares analysis.

5. OPERATIONAL ASPECTS

The different modules that the operational warning system was composed of were not designed to be used in an operational system. The turn around time of one complete forecast and the quality control in particular were important considerations. The quality and reliability of the system required special software design and hardware considerations.

5.1 Reduction of turnaround time

A crucial requirement of the new warning system is that all computations should be made within 1.5 hours such that the local dike authorities get a timely warning. The total cpu-time of one compound forecast is approximately 12 hours. To generate a forecast within an acceptable time, a computational cluster with 8 nodes was employed. The downscaling and wave calculation are split in separate, independent tasks, which are distributed over the different nodes. In particular, the downscaling module is time-consuming, but the different tasks are independent. This reduces the turnaround time of the individual models. A good reduction in turn around time was achieved by distributing the individual tasks over multiple processors.

However, the minimum required run time of each stage is still determined by the run time of the individual tasks. Both the individual IJsselmeer flow and wave models require the use of multiple processors in order to further reduce the turn around time. For each forecast one (time-dependent) IJsselmeer flow simulation and 9 (stationary) wave simulations are performed. With this setup, using 4 processors for the flow simulation and 2 processors for each wave simulations leads to the minimum turnaround time. The Markermeer flow and wave simulations are run on one processor.

At present, the number of wave calculations has been reduced in order to reduce the turnaround time of the system. Also, the numerical settings are a compromise between accuracy and turnaround time. As more models might be added and accuracy requirements might change, more computational nodes might be needed. The system is designed to easily include more nodes.

5.2 The quality of the warning system

Continuous monitoring of the system is necessary to control the proper functioning and to detect eventual failures. This requires extensive logging of all activities and abnormal situations. All input and output files are kept for some time for later inspection. As the number of files is large, a consistent naming convention is used, requiring a strict internal administration.

The operational warning system depends on wind forecasts, water level measurements and current measurements, which are all acquired from external sources using potentially media. This also increases the risk of missing required input. When required input is missing, the system must still generate a forecast based on estimated input (‘fall-back’). In this system, missing water levels and currents are substituted with the last available measurements. In the case of a missing wind forecast, the last available wind forecast is used again.

At present, the overall quality of the warning system cannot yet be judged. Comparing the results of the simulations with measurements in a daily routine will check the quality of the operational system. Distinctions are made for the first 24 hours of the prediction and the second 24 hours of the prediction. Statistical differences between predicted values and measured values the day after should proof the reliability of the system under any condition. It is in fact also the way to verify future improvements of the system.

5.3 Software design and hardware considerations

The operational warning system was composed of pre-existing modules. Rather than modifying the existing modules to cooperate with each other, the modules were left unchanged. Instead, the input for each module is generated from the output of other modules, either by simply reformatting data or by interpolation. The data input is formatted into the required format using input templates. The required data is extracted from the output of other modules. Using templates makes it easier to update a module with a newer version.

In order to guarantee consistency of information, information is kept in unique locations. If more than one module uses this information, the necessary input is generated from one source. For example, both the wave module and the run-up module require the locations where a storm level prediction is to be computed. Both input files are generated from one file containing unique information needed by both modules.

Generating input files, starting the different models (in parallel), collecting the output and maintaining an administration is all performed by a Perl script, linking the different modules together. This system is implemented on a Linux cluster with 8 PC’s (Pentium 3.0 GHz, 1Gb RAM). The operating system is Linux distribution Slackware 10.1, kernel 2.4.31. In addition, SWAN version 40.51 is used and WAQUA/Kalmina release of June 2006. Both the flow model (WAQUA) and the wave model (SWAN) utilize the ‘Message Passing Interface’ (MPI) to distribute the workload over the computational nodes.

5.4 Improvements

Experience with the system will show where efficient improvements can be made. These further improvements may concern:

- inclusion of wave-current interaction;
- reducing the number of iterations by improving the first guess of SWAN;
- extending the number of processors;
- grid resolution;
- refinements of the correction factors;
- using fine grid nested areas along the boundaries;
- verification of each module and of the total system;
- implementation of SWAN instationary mode;
- fail safe options (double ups processors, dual data links and hard disks RAID 5);
- fall back options (what to do if data is not delivered and yet a prediction has to be given).

6. CONCLUSIONS

A new warning system against flooding has been developed for the lakes in the IJsselmeer area. To determine storms levels against the dikes in the IJsselmeer area this system automatically collects:

- wind fields from the national weather service;
- performs a downscaling of the wind to resolve local wind variations;
- runs a flow model;
- runs a wave model, and;
- runs wave-runup module to determine storm level along the dikes in the IJsselmeer area.

Special attention was given to an efficient implementation of the downscaling module and the SWAN model.

A crucial requirement of the new warning system is that all computations should be made within 1.5 hours such that the local dike authorities get a timely warning. This requirement was achieved by performing the simulations of the various models in parallel mode on a multi-processor system.

An overall assessment of the quality of the warning system cannot yet be made. This requires monitoring the system output of winds, water levels and wave run-up for a long period and performing a statistical analysis on the results.

REFERENCES

Booij, N., R.C. Ris, and L.H. Holthuijsen, 1998: A third generation wave model for coastal regions, part I, J. Geophys. Res. 104, C4, 7694-7666.

Claessens, E.J., J.P. De Waal, H.C. Van Twuiver, and M. Bottema, 2002: The application of a coupled wave, waterlevel and wind model in a warning system against flooding. 7th Int. Workshop on Wave Hindcasting and Forecasting, Banff, Canada, 410-421.

De Rooy, W.C. and Kok, 2004: A combined physical-statistical approach for the downscaling of model wind speed. Weather Forecasting, 19, 485-495.

Stelling, G.S., 1983: On the construction of computational methods for shallow water flow problems. Ph.D. Thesis Delft University of Technology. The Netherlands.

Van der Meer, J.W., 1997: Wave run-up and overtopping on dikes. Delft Hydraulics Report H2458/H3051.

Verkaik, J.W., 2006: On wind and roughness over land. Ph.D. Thesis University of Wageningen, the Netherlands.

Verkaik, A.W., A.J.M. Jacobs, A.B.C. Tijm, and J.R.A. Onvlee, 2006: Local windspeed estimation by physical downscaling of weather model forecasts. Submitted to J. Wind Engineering and Industrial Aerodynamics.

Zijlema, M., and A.J. Van der Westhuysen, 2005: On convergence behaviour and numerical accuracy in stationary SWAN simulations of nearshore wind wave spectra. Coastal Engineering. Vol. 52, 237-256.