DETAIL-ORIENTED:

National Weather Service Introduces the HRRR

by Greg Waxberg

ISTOCK/CHEREZOFF

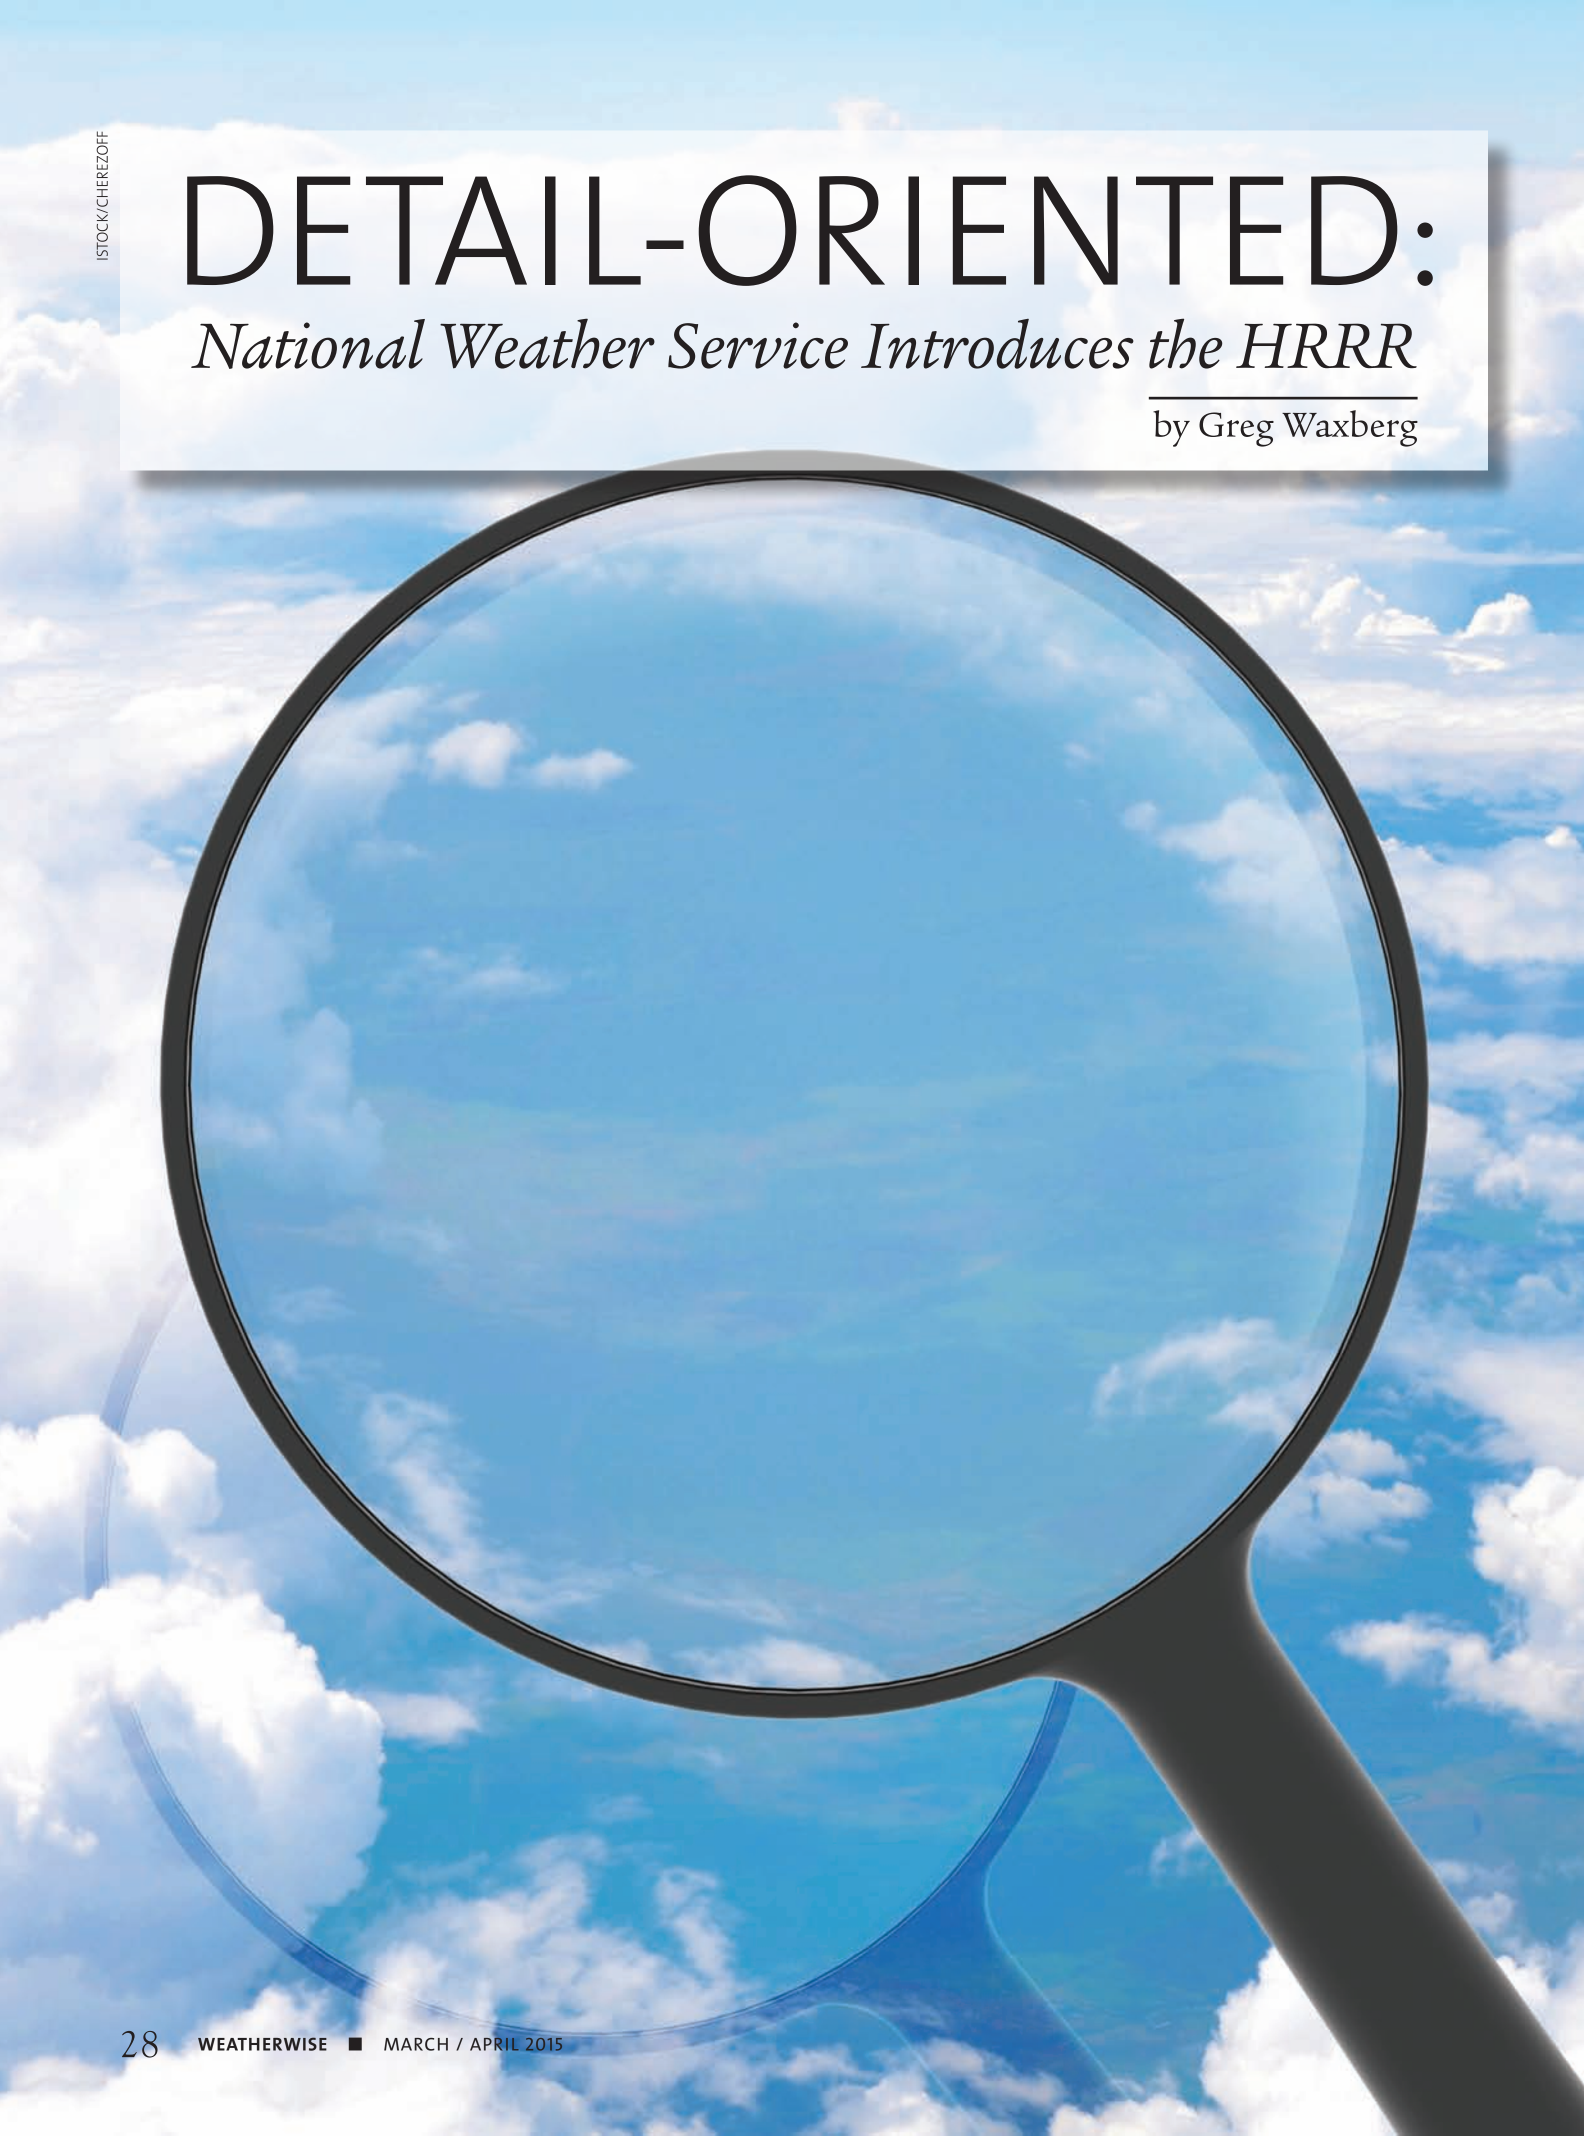

Think about the amount of detail you see when you zoom in on an online map, and the importance of that additional detail. Information such as highway exit numbers, one-way streets, business names, and the stores inside malls are visible if you zoom in far enough. Now, imagine that level of detail in weather forecasting, showing exactly what type of severe weather is moving across the map, how large and intense it is, and when it will pass over a particular city or neighborhood.

Many forecasters have been able to see those specifics since September 2014, when the High-Resolution Rapid Refresh (HRRR) forecasting model became operational at the NWS. It is complementing the Rapid Refresh model (RAP, in use since May 2012) and may eventually replace it. The difference is in the resolution: 13 kilometers/eight miles for RAP versus three kilometers/two miles for HRRR. Those numbers mean that the RAP produces a forecast imposed on a horizontal map grid with 13-kilometer squares, while the HRRR produces a forecast imposed over a three km-per-square grid. There are four times as many grid points in each direction, so 16 times as many grid points cover the same area.

"The smaller grid can resolve more features—on the ground, where a ridge or valley that may be entirely within a 13-kilometer grid space will likely spill over the edges of a three-km grid, and in the output, where a thunderstorm complex would likely occur entirely within the 13-kilometer grid space, but, again, would spill over the edges of the three-km grid," said Nezette Rydell, meteorologist-in-charge at the NWS forecast office in Boulder.

What the HRRR has in common with RAP is that it generates a new forecast every hour (the NWS has been producing hourly updated forecasts since 1998), which is a sizable difference from the next-most-frequent updates of *six* hours from global forecasting models that are standards in forecast offices. One such model is the Global Forecast System (GFS, produced by the NWS's National Centers for Environmental Prediction [NCEP]), which produces forecasts 240 hours into the future with standard resolution, or up to 16 days with coarser resolution, based on 28 km between grid points. It uses an atmosphere model, a land/soil model, and a sea ice model. Another six-hour model is the North American Mesoscale Forecast System (NAM, also produced by the NCEP), that produces forecasts 72 hours into the future and generates multiple grids at various resolutions over North America. Models are also produced in Europe and Canada—again, all operated on global scales every six hours.

But the HRRR covers the Lower 48 United States, as well as adjacent areas in Canada and Mexico, and generates short-term forecasts. According to Rydell, the model needs 30–35 minutes to assimilate data and then delivers output 40–45 minutes after the top of the hour:

> *Every hour, the HRRR gives you an updated forecast for the next 15 hours [RAP forecasts go out 18 hours], and we get new data every 15 minutes. If a global model has a bad run—maybe it got bad data—it's six hours before we get another chance. With the HRRR, we get new data very quickly, and we can look at the last six hours together as "ensemble predictions" and see how the forecasts changed, or didn't change, or trended, all in real time.*

ISTOCK/06PHOTO

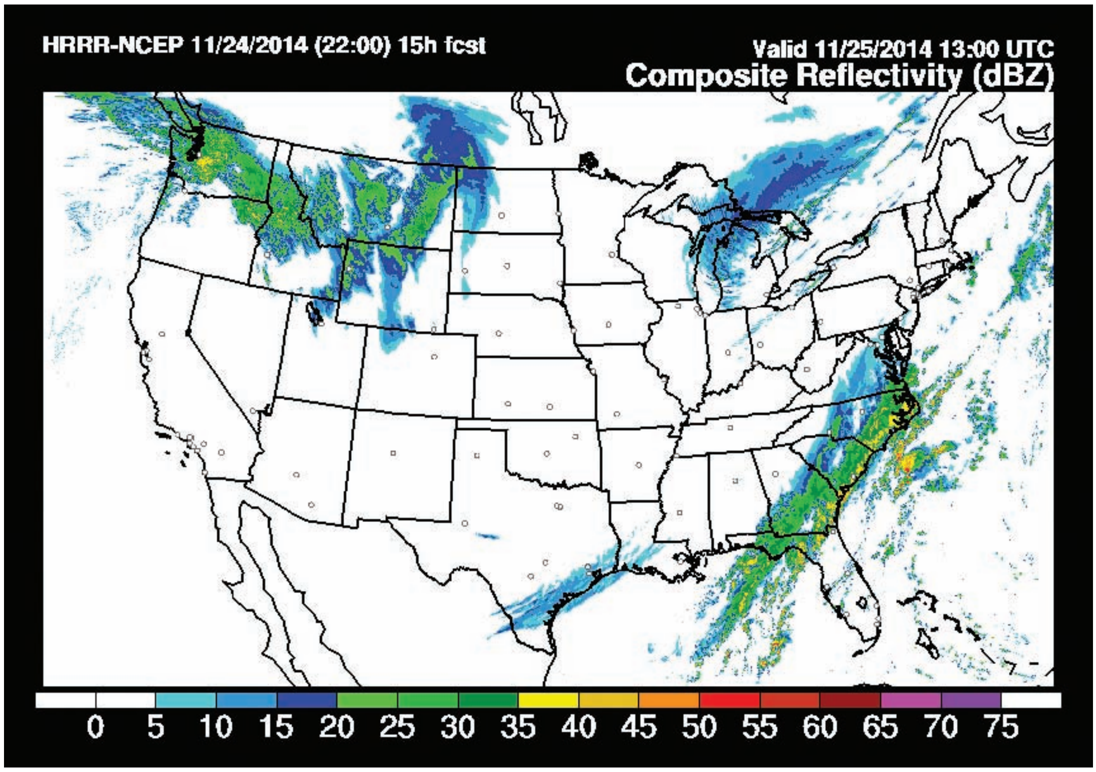

GSD

HRRR coverage in the United States.

That kind of ensemble look at the weather can increase our confidence in the forecast.

The big news, then, is that this combination of hourly updates and quadrupled resolution of the RAP is helping forecasters pinpoint severe weather for specific counties, cities, and neighborhoods with even more detail and accuracy than in the past, allowing for increased warnings for targeted areas, thereby keeping the public safer from tornadoes, hurricanes, hail, thunderstorms, snow bands in winter storms, and other severe weather—with the same level of accuracy for all types of conditions. Visually, the difference on a weather map is mainly a matter of colors: before, you might have seen large patches of green, yellow, and red for a thunderstorm. Now, you might see the same thunderstorm divided into intensity scales of five or six colors, and the storm itself might appear to be more segmented, rather than one large, uninterrupted system.

It is important to remember that the NWS continues to use about 20 forecasting models, one of which is the HRRR, and it can be used in combination with those other models to produce the desired forecast (for example, a seven-day forecast can be obtained by merging HRRR with GFS and NAM). And, depending on where you live and what television station you watch or radio station you listen to, you might be receiving a forecast from a private company, such as AccuWeather. Different weather companies use different models that satisfy various requirements for forecasts, but the HRRR is likely being incorporated into those models in some fashion.

Researchers in the NOAA's Earth System Research Laboratory (ESRL) developed the HRRR over a five-year timespan. Dr. Stan Benjamin, chief of the Earth Modeling Branch in the Global Systems Division (GSD), led the team. "There was a *tremendous* need from the public for weather models to provide better storm prediction. Many in our group have worked in the prediction of severe weather in the Oklahoma and Texas areas, and several of us have been storm-chasers, so there is motivation to address this key area. The timing came together with the technical side," he said. That "technical side" included data assimilation (complicated mathematical equations and interpolation problems to incorporate observations, including radar, into a computer model), beginning with the Rapid Update Cycle, an hourly updated model that NOAA started to operate in 1994 at a 60-kilometer scale, preceding the 13-kilometer RAP. He explained that "the RAP model responded well to a novel method for assimilation of radar data locating individual storms, so we were able to then downscale to a three-kilometer model."

Four cumulative factors allowed NOAA to develop the new storm-scale model. First, since 2007, all of the three-dimensional radar data across the country have been available at central locations, such as at the NCEP, making it easier to incorporate the data into forecasts. Second, and simultaneously, GSD developed a technique for using that radar data to start the HRRR with knowledge of existing storms. Third, since GSD uses the Weather Research and Forecasting Model (WRF), which has over 25,000 registered users in over 130 countries, GSD developed an optimal version of that model with revisions and new codes—to better capture clouds, and even heat and moisture from the land surface—that would give the best possible wind and storm forecasts. Fourth, the NWS bought processors to upgrade its computers to be fast enough to run the HRRR with its other models.

During those five years, ESRL took steps to fine-tune the model, beginning in September 2009, when it ran the HRRR on research computers, which are less expensive and reliable than those used for operational models, but perfectly sufficient for testing. Benjamin said:

> *In the period from 2010–2011, we discovered weaknesses in the system. The HRRR was not particularly good at forecasting* groups *of thunderstorms. It wouldn't form them, or it wouldn't sustain them. In the same period, there was a "moist bias," meaning that we had some false-alarm storms that weren't occurring at all. The average number of storms being reported by the HRRR was more than what was happening in reality.*

In addition to those problems, the forecasts for the first two hours after the model started did not look realistic—for example, the HRRR might have placed a storm in the wrong county, or forecasters could not easily see detailed wind structure—and this situation happened repeatedly for about four years. The problem was solved in 2013 when researchers added three-km radar assimilation every 15 minutes to the 13-kilometer assimilation of the previous hour. (Prior to 2013, the HRRR depended on radar data assimilated only at 13 kilometers into the RAP.)

To improve the HRRR, the development team refined its methods to make more accurate use of radar data and surface observations, such as those recorded at Surface Weather Observation Stations, described by Benjamin as another pillar of the observing network: "Radar and satellite images are critical, but observations at the surface and from aircraft add a 3-D picture of temperature, moisture, and winds, so we can better estimate conditions from near the ground up to jet level and above. Those details are important to predict how storms might evolve." Testing the HRRR involved objective comparisons to other models; comparisons to observations from weather balloons, aircraft, radar, and satellite; and feedback from up to 60 NWS offices, often with two or three forecasters representing each office. Those comments ranged from "outstanding job" and "too many storms" to both "good job for snow forecasting" and "bad snow forecasts."

Benjamin noted that, even with the challenges and feedback inherent in testing, the HRRR could have become operational in 2011 if computer speed had not been a concern: "*All* weather models work that way, with constant refinements. The HRRR's skill was sufficiently accurate in 2011 to be useful. Three years later, investments in the size and speed of operational computers allowed the NWS to take advantage of GSD's improvements to the model."

Now that they are using the HRRR, Rydell and her staff are especially happy about the

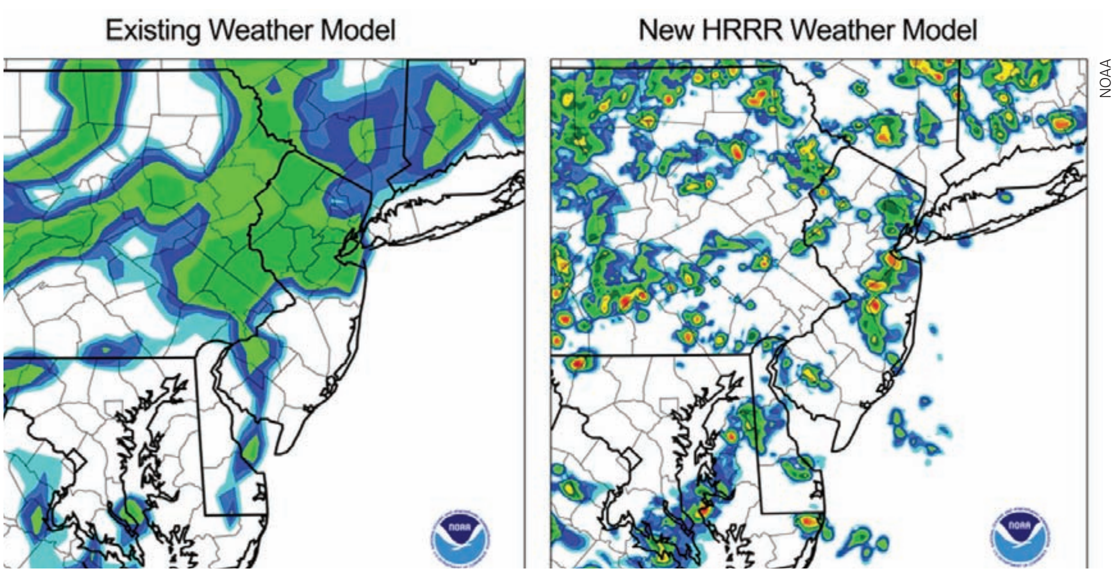

A comparison between the existing Rapid Refresh weather forecasting model and HRRR.

amount of observational data that the model uses to initialize, including data from satellites, radar, and ham radio. "The HRRR knows more about what's happening *now* than any model ever has before," she declared. Benjamin elaborated: "You don't just see smears of storms. You see individual thunderstorms, wind storms … lots of detail, and *realistic* details."

Over 100 fields of data are output from each HRRR forecast. A few of them are surface visibility, air pressure, VIL (vertically integrated liquid, the total mass of precipitation in the clouds; the maximum VIL of a storm can help determine its severity), velocity of updrafts and downdrafts, helicity (the amount of rotation in a storm's updraft air; significant rotation means that the predicted storm is likely a supercell and could even produce a tornado), graupel (sleet), lightning threats, heating rates, moistening rates, wind speed, soil temperature, soil moisture, dewpoint, relative humidity, precipitation, heat flux, and cloud cover.

Also included in the HRRR are advanced representations of clouds and wind. "Especially for short-term forecasts, it is important for weather models to have knowledge of existing clouds because clouds shield sunlight from reaching the surface, and clouds affect radiation and temperatures at all times of day," Benjamin said. Across the United States, laser beam ceilometers measure the distances from the ground to the tops and bottoms of clouds. Ceilometers are part of the Automated Surface Observing System (ASOS), which are sites owned by the NWS, the Federal Aviation Administration (FAA), and, in some cases, the Department of Transportation. As far as wind is concerned, the HRRR shows accurate local variations in wind speeds, particularly around mountains.

On the subject of the FAA, the HRRR is also supplementing forecasts for the aviation industry, so that airline dispatchers can direct pilots in avoiding hazards such as turbulence and thunderstorms. "Dispatchers file flight plans—the intended routes for aircraft—so that Air Traffic Control will know where each plane is going and when it will pass certain points. Air Traffic Control's main responsibility is to keep the proper amount of space between planes, so pilots need to use their onboard radar to make sure they're deviating safely around thunderstorms," explained Clinton Wallace, Deputy Director of the NWS Aviation Weather Center. Wallace and his staff work with the FAA System Command Center and with airlines to provide weather information for aircraft.

At least every two hours, the FAA issues a national Operations Plan for routing of air traffic, including wind warnings at airports. Wallace said:

> *The FAA needs to balance the load of aircraft flying through particular sectors of airspace, so the agency sets these different routes in partnership with airlines. Well, once you throw a thunderstorm into a sector, you've just reduced the amount of airspace that's available for aircraft to fly safely. So, the FAA wants to balance that impact on airlines, and the HRRR is a tool for the FAA to generate that strategic plan.*

One problem with which air travelers are all too familiar is the "ground stop," when Air Traffic Control stops aircraft that are on the ground so that airplanes in the sky have the airspace to land at an airport that is impacted—or that will be impacted—by thunderstorms. The ground stop can prevent a plane from taking off until its flight time positions it correctly with other airplanes en route to an airport. Also, the ground stop may delay a certain number of airplanes to reduce the demand on an airport. Under these circumstances, Wallace praised the HRRR as "the ideal

ISTOCK/ANDREWSOUNDARAJAN

model for a lot of flight planning because of updated conditions based on real observations."

His praise does not end there, since a major plus of the HRRR is that forecasters can see data at the same scale as the actual thunderstorms:

> *In the past, we saw precipitation amounts for a certain time period. Now, since the HRRR looks like radar, we can see a visual representation of thunderstorm spacing, size, and intensity, very important for routing aircraft around weather. Several small, widely scattered cells or one large cell are relatively easy to route around, but it may not be possible to route a plane around a line of storms. With lower-resolution models, we weren't able to pick out those details of how thunderstorms were developing.*

The Massachusetts Institute of Technology (MIT)/Lincoln Labs and the National Center for Atmospheric Research (NCAR) used the experimental HRRR as an input to the Consolidated Storm Prediction for Aviation's (CoSPA) blending algorithm, under funding from the FAA's Aviation Weather Research Program. Essentially, CoSPA combines the best capabilities of a group of forecast systems into one high-quality system; it resulted from federal mandates to reduce the number of competing and redundant aviation weather forecasts that, themselves, resulted from research funded by the FAA and NWS and then caused confusion in air traffic management.

Along with passenger flights, those for shipping are benefiting from the HRRR. Thus far in his department's daily experience with the model, FedEx's manager of weather services Kory Gempler said the new model is the "perfect fit" for FedEx's continental United States Terminal Aerodrome Forecast (TAF). TAFs, usually issued every six hours and valid for 24 hours within five statute miles of the center of an airport's runway complex, provide cloud heights and ceilings, wind direction and speed, wind shear, visibility, and other conditions important for aviation. "Although it's not perfect—no model is—we feel that the HRRR is usually the first choice due to its update cycle and its radar assimilation capability. Proof of this can be noted in the performance of CoSPA, which still struggles with certain storm structures, but, overall, is one of the better convective forecast products due to its update cycle," Gempler said.

Gempler's staff also uses the HRRR to keep an eye on weather for airports that FedEx serves, but for which the company does not issue a TAF, and to pass along changing conditions to flight dispatchers. "We focus on FedEx hubs, major markets, and other weather-sensitive airports for our TAFs. For example, Phoenix does not see much bad weather, so we use the NWS TAF, but we have the ability to write a TAF for any city in the United States if needed," he said.


ISTOCK/JMICHL

So, since we know that weather models are constantly fine-tuned based on technology, faster computers, and new or improved scientific ideas, the question is how the HRRR might evolve. Benjamin listed several possibilities: improving the starting points for the model by incorporating better observations of radar and satellite data; running multiple models each hour with slightly different starting points to use those ensemble forecasts to provide better probability information for decision-making; extending the HRRR forecast beyond 15 hours; and covering the globe, perhaps within the next 10 years. Benjamin reports:

> *There is a lot of motivation to improve the HRRR. People make decisions every day based on weather forecasts. Sometimes, it's a safety issue, like getting home from work before a snowstorm hits. Or, in different industries, critical weather-dependent decisions are made every hour by energy managers or aviation planners, all needing better forecasts. What do they really need? We still have only one atmosphere, but we can make a single model* better *to help with those decisions.* W

GREG WAXBERG, a writer and magazine editor for The Pingry School, is an award-winning freelance writer.