WHITE PAPER

# How the Latest Generation of Supercomputers Speeds Global Climate Research

Silicon Graphics International Corp.
900 North McCarthy Blvd.
Milpitas, CA 95035

TABLE OF CONTENTS

## 1.0 Introduction

Not so long ago, weather forecasters endured a reputation for being wrong more often than right about the need to carry an umbrella. However, meteorologists have recently gained vastly improved capabilities for developing their projections. The big change came with advances in high-performance computing (HPC).

Because of the complexity involved, the length of the simulation period and the amounts of data generated, weather prediction locally or on a global basis requires some of the most powerful computers in the world.

Predicting the weather next week or in the next hour involves collecting data observed by satellites, land-based stations, and other sources.

The computational requirements are driven by the need for higher resolution models for more accurate and extended predictions of future weather patterns. In addition, more physics and chemistry processes are included in the models so scientists can observe the very fine features of weather behavior. These models operate on 3D grids that encompass the globe. The closer the points on the grid are to each other, the more accurate the results.

Meteorologists combine this observed data with models that incorporate a scientific understanding of atmospheric processes, and then use large computing environments—often clusters with hundreds or thousands of nodes that add up to large petaflops-scale systems—to simulate a range of weather phenomena. These HPC-powered simulations allow forecasters to predict everything from the path of a tornado on the Great Plains of the U.S., which can change from minute to minute, to weather in the Swiss Alps, where it can be cool and rainy in one valley and warm and sunny in the next.

## 2.0 The Huge Challenge of Seeing Over the Horizon

Compared with forecasting short-term weather events, predicting global climate change is a bit of a different game. Climatologists take a longer view, and they rely on multiple models that simulate many years of climate change—typically 10 to 100 or more. While using similar or the same models as used for weather forecasting, climatologists compute quantities that represent averages over large areas and over longer time periods of time. These long-term climate simulations rank among the most complex and computationally demanding problems in science, right up there with predicting earthquakes or research into astrophysics, aeronautics, fusion power, and exotic materials.

The amount of data required to simulate long-range climate patterns is staggering. Simulations begin with collecting an initial set of data from hundreds of thousands of surface stations, ships and buoys, aircraft, airborne radiosondes (tiny airborne weather stations that radio their findings back to Earth) and dozens of weather satellites that stream terabytes of information every day. Climate researchers also need to augment this gigantic set of observed data with vast amounts of historical data. Comparing model predictions for past years with actual observations for the same period in the past enables them to verify their model. This historical data also provides a basis for predicting trends among climate-related factors.

Additionally, the governing equations of numerical weather prediction are what mathematicians call non-linear. This means that small changes in input data can result in big changes in the resulting prediction. One way to compensate for the relatively low precision of observed data and to “smooth” over the non-linearity fluctuations is by using the method of “ensembles”—i.e., making multiple runs that differ slightly in their input data, and then weight-averaging the results. In addition, scientists use multiple climate models—which may differ in how they simulate some phenomena or how their grids are constructed—to get an “average” prediction.

This explains why tomorrow's forecast can be way off, yet the climate prediction—the average for the season—can be spot on.

Not only are the computational requirements of climatology severe, but the capricious nature of climate itself can present enormous challenges for researchers. Local weather may obscure global trends and climate changes are not uniform across the globe—on any given day, there can be a drought in one region and floods in another. Additionally, what seem like very small changes in global average temperatures have a significant cumulative effect when the trend is persistent.

## 3.0 Added Complexity: Our Changing Oceans and Atmosphere

Climate researchers must also factor in the complex role played by oceans, which cover approximately 70 percent of the Earth's surface. "The world's oceans have a two-way relationship with weather and climate," according to the U.S. Environmental Protection Agency. "The oceans influence the weather on local to global scales, while changes in climate can fundamentally alter many properties of the oceans."[1]

For example, since the turn of the last century, human activity has caused an astonishing increase in the release of CO2 and other greenhouse gases (see Figure 1). As these gases accumulate in the atmosphere, they trap more energy from the sun and the oceans absorb more heat, which results in rises both in sea surface temperatures and sea levels as polar ice melts. Ocean heat content not only determines sea surface temperature but also affects currents, which contribute to climate patterns.[2] Ultimately these changes in temperatures, due to growing emissions of greenhouse gases from various sources, are leading to alterations in climate patterns around the world.

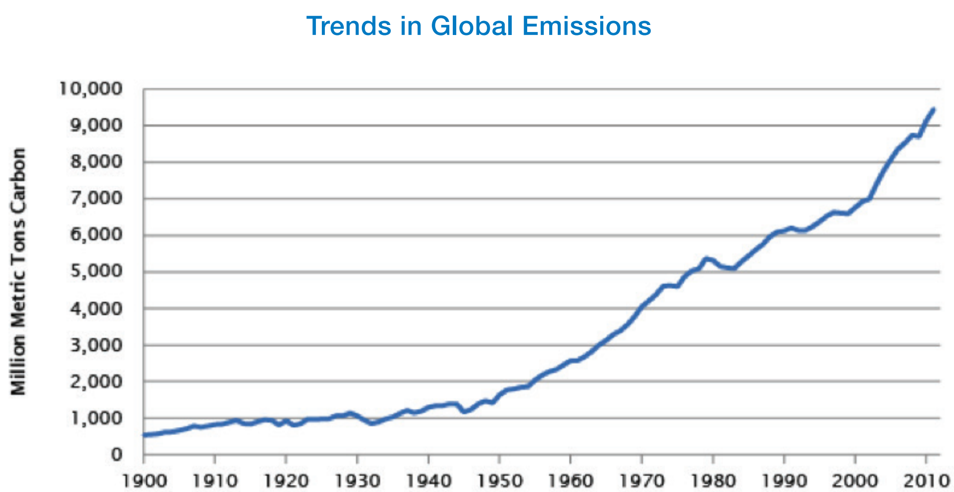


*Figure 1. Global carbon emissions from fossil fuels have significantly increased since 1900. Since 1970, CO2 emissions have increased by about 90%, with emissions from fossil fuel combustion and industrial processes contributing about 78% of the total greenhouse gas emission increase from 1970 to 2011. Agriculture, deforestation, and other land use changes have been the second largest contributors.*[3]

[1] United States Environmental Protection Agency (2016): Climate Change Indicators in the United States.

[2] https://www3.epa.gov/climatechange/science/indicators/oceans/.

[3] Source: Boden, T.A., Marland, G., and Andres R.J. (2015). Global, Regional, and National Fossil-Fuel CO2 Emissions. Carbon Dioxide Information Analysis Center, Oak Ridge National Laboratory, U.S. Department of Energy,doi 10.3334/CDIAC/00001V2015.

Over the long-term, daily forecasts have little dependence on what happens under the ocean surface, but the long-term climate is hugely impacted by the maritime world. Warmer southern oceans may promote the development of stronger storms in the tropics, which as we have seen in recent years can cause enormous property damage and loss of life. For example, in 2010, almost all of Pakistan was affected when massive flooding caused by record-breaking rains hit the region. The recent drought in California that led to enormously destructive wildfires (which burned more than 300,000 acres) can be traced back partially to the climatic influence of the warming oceans and changing currents. In just the last decade alone the U.S. experienced the two costliest storms in its history: first Hurricane Katrina virtually submerged the city of New Orleans, and then Superstorm Sandy slammed into the nation's densely populated east coast and financial center, leaving $65 billion of damage in its wake. Coastal communities are impacted by storms like these along with rising sea levels.

While any single extreme weather event may seem anecdotal, climatologists observe that there is an increased pattern in all manner of extreme events. The accuracy of climate models projections is measured, among other things, by predictions of frequency and severity of storms and precipitation.

*In 2010, almost all of Pakistan was affected when massive flooding caused by record-breaking rains hit Khyber-Pakhtunkhwa and Punjab.*

*2015 California wildfires as seen from space.*

Predicting climate change—and preparing for and possibly even remediating it—requires more powerful high-performance computing systems and climate modeling to improve the resolution and precision by orders of magnitude. These will allow scientists to provide more actionable projections about the impact of climate change for specific regions and assist agencies throughout the world to develop more accurate weather predictions on a local and global scale.

## 4.0 Super Speed and Scale to Create Simulations in Practical Time Frames

Running simulations on the spatial and temporal scale needed for climate change research requires maximum computational speed and scalability in order to execute in practical time frames. It turns out that in today's HPC cluster technology, moving data in and out of the processing units is often the limiting factor in minimizing time-to-solution. To be effective, systems working with weather forecasting and climate modeling require high memory bandwidth and fast interconnect across the system, as well as a robust parallel file system.

Extremely high capacity storage and file systems are essential to dealing with two key data problems of weather forecasting and climate modeling. The first problem involves the relatively short-lived—but time-critical—nature of the data used as input into the models. This data comes from the Earth's surface through many different instruments, from under the ocean surface, and from sensors and satellites. It corresponds to numerous physical, chemical and biological properties, and it is not all collected at the same instance. Before running a model, the data must be "assimilated"—i.e., extrapolated in time and space to match the mathematical grid points used to approximate the area over which the simulation is run.

The second data source is the detailed numerical output resulting from the runs of many models and/or from a model that simulates hundreds of years of evolving climate.

Moreover, it is not only the final state that is captured. For most studies, data is saved for many (if not all) "time steps" of the simulation so that the evolution leading to the final state can be analyzed. Climate research requires following a measurable pattern (temperature, precipitation, etc.) over the years. This simulated data amounts to the majority of data involved in climate research.

Previously, crunching the huge volumes of data required to create high-resolution climate simulations could take days to months to complete even on the most powerful computing clusters. Today, climate researchers are welcoming the new generation of HPC systems with speed and addressable memory measured in petaflops and hundreds of terabytes, respectively. These supercomputers, such as the SGI® ICE™ XA system, are uniquely suited to support large-scale needs in weather and climate simulations.

A machine such as the petaflops-scale supercomputer recently purchased by the U.S. government for the National Center for Atmospheric Research (NCAR) is one example. Called "Cheyenne," this new system built by SGI is capable of 5.3 quadrillion calculations per second, and will have an astonishing 313 terabytes of storage memory. It can make quick work of interpolating terabytes of weather data to fit a three-dimensional grid that approximates the globe, and then run simulations over this grid to produce long-range forecasts. Cheyenne can also handle the huge datasets that incorporate topography, winds, temperatures, radiation, gas emission, cloud forming, land and sea ice, vegetation and more. Additionally, this new generation of petascale supercomputers has the power to model massive phenomena's that affect climate.


An SGI ICE XA system, like the one installed at NCAR, is capable of incredibly complex, data-intensive calculations to dramatically improve the resolution and precision of climate change predictions. As a result, governments, industry and local communities can be better prepared for the future. It will also be used to predict climate patterns over the next ten years or further into the future to assess drought risk or the extent of melting sea ice in the Arctic.

The importance of this information cannot be understated, especially when it comes to anticipating, understanding and coping with weather and climate disasters. For example, according to the National Centers for Environmental Informatics, there were eight U.S. weather and climate events in 2014 with losses exceeding $1 billion each, including droughts, tornados, flooding and severe storms such as a major winter storm event.

## 5.0 Conclusion

Data accumulation is just one of the challenges facing today's weather and climatology researchers and scientists. They must deal with a lack of data regarding the future as well as complicated interrelationships, such as between oceans and the atmosphere. To understand and predict Earth's weather and climate, they rely on increasingly complex computer models and simulations based on a constantly growing body of data from around the globe.

Some of today's largest and most sophisticated computer hardware and software have been used to predict our weather and investigate climate changes. Using these HPC technologies, scientists have done an exceptionally better job of predicting weather and climate than was possible only a few years ago. Now, a new generation of petascale supercomputers is enabling scientists to not only refine their predictions by orders of magnitude but also better address the impacts of climate change and assist society in confronting adverse weather phenomena.

**Global Sales and Support**: sgi.com