Ultra-high Resolution Simulation of a Downburst-producing Thunderstorm

Robert Sisneros
National Center for Supercomputing Applications
Urbana, IL, USA
sisneros@illinois.edu

L.G. Orf
Central Michigan University
Mt Pleasant, MI, USA
leigh.orf@gmail.com

G.H. Bryan
National Center for Atmospheric Research
Boulder, CO, USA
gbryan@ucar.edu

***Abstract*—In this work we investigate simulation data from the CM1 cloud model. Specifically, typical thunderstorms are simulated under environmental conditions known to produce downbursts, such as those in the High Plains. A downburst is created by a column of sinking air that, upon impinging the ground, can create an outburst of damaging winds. Understanding the specific capabilities of downbursts is critically important to the maintenance of structures susceptible to these intense downdrafts, such as power transmission lines. In this video we first provide a cloud- and precipitation-based overview of a downburst-producing thunderstorm. We then provide the context to transition from the more recognizable and general features of the thunderstorm to those pertaining to downbursts. In high-detail we illustrate the significant small-scale characteristics of the downburst. Video link: http://www.ncsa.illinois.edu/~sisneros/sc_showcase.html.**

I. INTRODUCTION

An advance in any number of fields under the umbrella term "atmospheric science" has real potential to be both far-reaching and general. This is certainly the case for the problem of understanding the formation of destructive atmospheric phenomena such as tornadoes or downbursts. The use of simulations is invaluable in studying complex interactions among the contributing factors of these events. However, understanding such small scale features requires a robust, high-resolution code coupled with a capable, cutting edge resource.

This work is direct output from a multi-institutional collaboration spanning multiple steps typical of large-scale science. In the remainder of this paper we will begin with a more in-depth description of our target problem in Section II. We follow outlining processes involved in the generation of the data in Section III as well as the visualization of the data in Section IV.

II. DOWNBURSTS

Downbursts are intense downdrafts that are spawned by thunderstorms. In some cases, the downburst-producing thunderstorms are weak and short-lived, but occur in environments which are ideal for the formation of strong downdrafts. Downbursts can cause tornado-strength winds at the earth's surface, resulting in significant damage to property. The simulation we present occurs in such an environment, one that is typical for the High Plains of the United States. A thunderstorm grows and produces snow and graupel within the cloud which, as it falls downward into warm air, melts to form rain which then evaporates. The melting and evaporation cool the air, which is free to accelerate downward in this type of environment.

The main research objective of this work is to understand the specific wind loads that downbursts present to structures such as power transmission lines, which are vulnerable to downburst winds [5]. Up until recently, downbursts were typically simulated using simple models that force the downdraft using either a momentum or density source without including the thunderstorm cloud itself. Only recently has it been possible to simulate the lifecycle of an entire downburst-producing thunderstorm at resolutions sufficient to properly capture the flow near the ground where the strong winds do their damage. In addition to providing data valuable to wind engineers interested in designing structures that can withstand downburst winds, thunderstorm simulations at resolutions only possible using supercomputing facilities can also be used to aid meteorologists in better understanding both the physical processes causing the downbursts as well as the surface damage patterns caused by downbursts.

The imagery of this simulation highlights the origin and motion of the fastest-moving air. We chose to focus on potential temperature perturbation as a volume-rendered field in order to trace the motion of the air, exploiting the fact that in this type of thunderstorm, which produces a low-precipitation "dry downburst," the negative buoyancy that forces the downdraft is primarily thermodynamic in origin, resulting from evaporation, sublimation, and melting of hydrometeors. In so-called "wet" downbursts, which typically occur in more humid environments, negative buoyancy results primarily by the drag induced by heavy precipitation. Negative potential temperature perturbation is proportional to positive density perturbation, and is the source of the term term "density current" which describes phenomena such as downbursts.

III. GENERATING THE DATA

The CM1 [1] cloud model is spatially three-dimensional, non-hydrostatic, non-linear, and time-dependent. With CM1 initialized to match environments known to be conducive to downburst-producing thunderstorms, our data was generated on a run on the kraken XSEDE supercomputer at NICS. It utilized 18,000 cores and produced 25 TB of data in four wallclock hours. Although CM1 was designed specifically for parallel architectures, at our current, and especially future scales, there were certain obstacles we had to overcome to efficiently use our allocated resources.

The CM1 model came with three output format options, all of which either result in one file per MPI process per output time or one file per output time. Both of these options present problems for large simulations; one file per process will result in millions of small files, making post-analysis unwieldy. This approach is inefficient, suffering from latency, metadata server overhead, and general overhead associated with concurrently writing such a large number of files to disk on a shared resource. However, writing one file per time becomes prohibitively inefficient for a large number of MPI processes. Blue Waters exhibits the best aggregate I/O throughput when many (but not too many) large files are written concurrently.

These issues motivated the development of new I/O code for writing out model data for near real-time and later use [4]. An approach was developed that meets the following objectives: reduces the number of files written, minimizes the number of times actual I/O to the filesystem is performed, and creates larger files. This is accomplished by writing only one file per node (a typical simulation has 16 MPI ranks per node) and by buffering data from multiple times in memory before writing an HDF5 file to disk. In addition to the obvious reduction of files on disk, the reduction in I/O requests (typically by a factor of 50) results in better overall performance of the simulation.

IV. THE VISUALIZATION

The simulation was rendered on the Blue Waters supercomputer utilizing VisIt [6], [2]. In addition to VisIt's proven track record of running at large scales on architectures similar to Blue Waters [3], the level of support for VisIt by the Blue Waters staff makes the package a natural choice. For all rendered frames, VisIt's ray caster was used with the engine executed on anywhere from approximately 1K to 8K processors (required memory may change on a per-frame basis from even slight changes in lighting, view points, sample rates, etc.). The nearly 600 timesteps of the simulation were rendered for each section of the video (see below). However, there was a final step to ensure efficient interfacing of all pieces of the pipeline.

A. *Parallel Processing*

The library responsible for the significant increase in CM1's I/O throughput (see Section III) accomplishes the task through I/O aggregation. That is, multiple time steps of the full spatial resolution are combined and written into a single file. While VisIt has support for many data types, including HDF5, a typical reader equates a single file with a single data block. This provides a high level of support for typical file-per- process simulation output, but would eliminate the possibility of parallel analysis in our case.

First we added a routine to the aforementioned I/O library that reads metadata from the CM1 output files and creates a new header file (a .cm1visit file). This file contains all necessary information for creating the data structures VisIt needs: grid size and spacing, variable names, number of time steps, etc. Second, a data reader VisIt plugin was created that is associated with this header file. Using this plugin VisIt has each of its current processors request a block of data. The reader, based on the number of processors allocated in a VisIt session, automatically decomposes the data to feed these processes.

B. *The Dissipation Phase of the Storm*

This part of the video portrays the typical features of a thunderstorm from typical angles. This section begins with a description of the variables rendered (Figure 1(a)). For the majority of this video, the camera is located above the storm looking down. The video ends after a transition of the camera around to under the storm looking up (Figure 1(b). The transfer function shown in Figure 1(a) is a direct screen capture of VisIt's transfer function widget with the used gaussian selections and values that are unchanged across the set of timesteps.

The effect of rendering two variables in one volume rendering was achieved through a derived variable in VisIt's expression control window. This expression looked at both the cloud and precipitation variables (which were also derived quantities from a combined total of five CM1 variables); the variable with the largest value was selected for a location, the direct value if precipitation, the negative if cloud. This is possible because there are few overlapping structures between these two variables. A simple sum of these variables leads to visually similar structures, but it is not possible to color each separately.

C. *Focus on the Downburst*

After a transition effect (Photoshop tween, see Figure 1(e)) the remainder of the video keeps the camera directly on the very small section of the dataset containing the downbursts. The first and last timesteps of this video are shown in Figures 1(c) and (d). The rendered variable is potential temperature perturbation, and the selected low values correspond to the sinking cold air associated with downbursts. The simulation's high resolution allows for very close inspection of small features, some of which we stop to annotate within the video (Figure 1(f)).

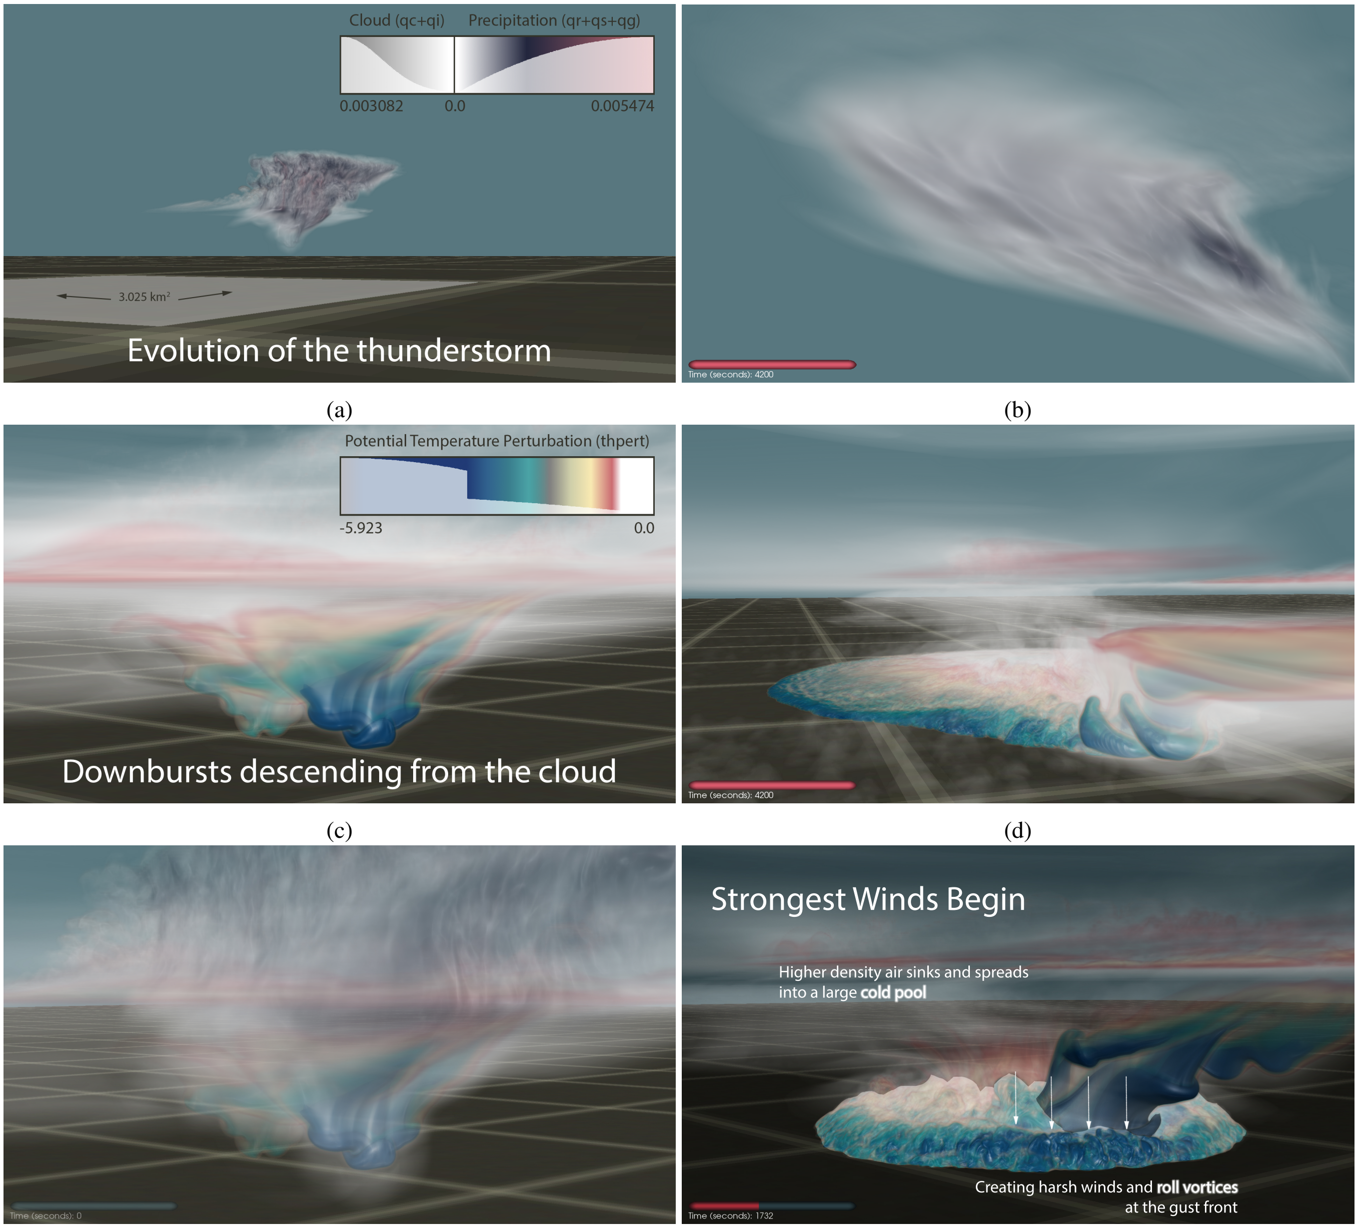

(a) (b)

(c) (d)

(e) (f)

Figure 1: Sample frames from the video highlighting the annotated first frames and the last frames from the dissipation phase (a-b) as well as section focused on the downburst area (c-d). A frame from the tween transition between sections of the video (e), and an example of one extensive mid-video annotation (f).

REFERENCES

[1] George H. Bryan and J. Michael Fritsch. A Benchmark Simulation for Moist Nonhydrostatic Numerical Models. *Monthly Weather Review*, 130(12):2917–2928, 2002.

[2] Hank Childs, Eric S. Brugger, Kathleen S. Bonnell, Jeremy S. Meredith, Mark Miller, Brad J. Whitlock, and Nelson Max. A contract-based system for large data visualization. In *Proceedings of IEEE Visualization 2005*, pages 190–198, 2005.

[3] Hank Childs, David Pugmire, Sean Ahern, Brad Whitlock, Mark Howison, Prabhat, Gunther H. Weber, and E. Wes Bethel. Extreme scaling of production visualization software on diverse architectures. *IEEE Computer Graphics and Applications*, 30(3):22–31, 2010.

[4] Leigh Orf. An I/O strategy for CM1 on Blue Waters. Technical report, Central Michigan University, 11 2012.

[5] Leigh Orf, Erica Kantor, and Eric Savory. Simulation of a downburst-producing thunderstorm using a very high-

resolution three-dimensional cloud model. *Journal of Wind Engineering and Industrial Aerodynamics*, 104106(0):547 – 557, 2012.

[6] VisIt users portal. http://visitusers.org/, 2011.