

SIMULATING THE FIRES IN THE WORLD TRADE CENTER

Kevin McGrattan & Charles Bouldin
National Institute of Standards and Technology, USA

INTRODUCTION

In the months following the attacks on the World Trade Center and the Pentagon, there was an active debate in the fire protection engineering community about the fires that erupted following the impact of the aircraft on the buildings. Because fires of this magnitude in these types of buildings are rare, there is a wide spectrum of opinion about the fire temperatures and their effect on the structural steel. Much of the fire literature consists of empirical correlations derived from experiments ranging from bench scale to room scale. Extrapolating these well-known correlations to the WTC requires a re-examination of the underlying assumptions. Many of these correlations are appropriate for a narrow range of fire sizes and building geometries, and can not be directly applied to the WTC fire scenarios. As a result, computer fire models that have been developed over the past decade are being applied to the analysis.

The National Institute of Standards and Technology (NIST), an agency of the United States Department of Commerce, is conducting an investigation of the collapse of WTC 1, 2 and 7. WTC 1 and 2 were the 110 floor towers, and WTC 7 was the 47 floor office building north of WTC 1 that collapsed at 5:20 in the afternoon of September 11, 2001. As part of the investigation, NIST has conducted simulations of the fires in each building using a computational fluid dynamics (CFD) model known as the Fire Dynamics Simulator (FDS). This paper will describe the experiments conducted at NIST to calibrate and validate the FDS model for use in the WTC project, and it will describe the techniques developed to simulate the very extensive fires that spread over 6 to 12 floors in the different buildings. This paper does not present any of the subsequent work to analyze the thermal response of the steel structure, nor does it present any collapse hypotheses. These efforts will be described elsewhere.

Experimental Program

Two large-scale test series were conducted to provide validation for the FDS, plus various small-scale experiments were conducted to provide the model with input data for different materials. The large scale test programs are referred to as Phase 1 and Phase 2. Both test series involved fires in compartments with the same ceiling height as a floor in WTC 1 or 2. Phase 1 was a series of fire tests with a liquid fuel spray burner generating a fixed amount of energy in a compartment with various targets and obstructions, like columns, trusses and other steel objects. These tests were designed to test the accuracy of the model, and its sensitivity to changes in various input parameters. Phase 2 was a series of fire tests in which office workstations similar to those used in WTC 1, 2 and 7 were burned in a compartment with limited openings to simulate the under-ventilated conditions of the WTC fires. These tests were designed to test the model's ability to characterize the burning behavior of real furnishings under conditions typical of the WTC fires. Only the Phase 2 work will be discussed in this paper.

NIST Fire Dynamics Simulator

Fire Dynamics Simulator (FDS) is a computational fluid dynamics (CFD) model of fire-driven fluid flow. It solves numerically a form of the Navier-Stokes equations appropriate for low-speed, thermally-driven flow with an emphasis on smoke and heat transport from fires (McGrattan *et al.* 2003). Version 1 was publicly

released in February 2000. The core algorithm is an explicit predictor-corrector scheme, second order accurate in space and time. Turbulence is treated by means of the Smagorinsky form of Large Eddy Simulation (LES). For most applications, FDS uses a mixture fraction combustion model. The mixture fraction is a conserved scalar quantity that is defined as the fraction of gas at a given point in the flow field that originated as fuel. The model assumes that combustion is mixing-controlled, and that the reaction of fuel and oxygen is infinitely fast. The mass fractions of all of the major reactants and products can be derived from the mixture fraction by means of "state relations," empirical expressions arrived at by a combination of simplified analysis and measurement.

Radiative heat transfer is included in the model via the solution of the radiation transport equation for a non-scattering gray gas, and in some limited cases using a wide band model. The equation is solved using a technique similar to finite volume methods for convective transport; thus the name given to it is the Finite Volume Method (FVM). Using approximately 100 discrete angles, the finite volume solver requires about 15 % of the total CPU time of a calculation, a modest cost given the complexity of radiation heat transfer. FDS approximates the governing equations on a rectilinear grid. The user prescribes rectangular obstructions that are forced to conform with the underlying grid.

All solid surfaces are assigned thermal boundary conditions, plus information about the burning behavior of the material. Usually, material properties are stored in a database and invoked by name by the user. An extensive effort was undertaken to characterize the thermal properties of common items found in an office setting, like privacy panels, stacks of paper, computer monitors, office chairs, pressboard tables, desks, and carpeting. These materials will be described next.

CALIBRATION AND VALIDATION EXPERIMENTS

The experimental program concentrated on the thermal properties of the office furnishings that constituted the bulk of the combustible fuel within the WTC buildings under study. Several types of office workstations typical of those used in WTC 1 and 2 were purchased at area office supply stores. The thermal properties of the major materials making up the workstations were derived from cone calorimeter experiments. These properties were input into FDS, which was used to simulate the burning behavior of a single workstation burning under a 2.5 m ceiling with baffles to contain a hot layer of smoke above the burning workstation. Other than the baffled ceiling, no walls surrounded the workstation other than its own privacy panels. The thermal properties of the workstation components were adjusted slightly so that the FDS prediction of the heat release rate would match the experiment. Then the model was used to predict the heat release rate of 3 workstations burning within a large enclosure. The purpose of this exercise was to determine if FDS could simulate the dynamics of a fire in a setting similar to WTC 1, 2 and 7.

Description of the Workstation Components

Cone calorimeter experiments at three different heat fluxes were performed for the carpet, desk (wood), computer monitor, chair, privacy panel, and stacked paper. For the simulations of the WTC fires, only the carpet, desk and privacy panel data were used directly. The carpet and privacy panel were modeled as thermoplastics, that is, the burning rate is assumed to be proportional to the heat flux from the surrounding gases. The desk was modeled as a charring solid, in which a pyrolysis front propagates through the material leaving a layer of char behind that insulates the material and reduces the burning rate. Details of the pyrolysis models can be found in the FDS Technical Reference Guide (McGrattan *et al.* 2003).

The desk was modeled as a charring solid. The thermal properties of the wood and its char were taken from both the calorimeter experiments and the work of Ritchie *et al.* (1997). It is 2.8 cm thick with density 450 kg/m³, specific heat 1.2 kJ/kg/K at 20 °C and 1.6 kJ/kg/K at 900 °C, conductivity 0.13 W/m/K at 20 °C and 0.16 W/m/K at 900 °C. The ignition temperature is 360 °C and the heat of combustion is 14,000 kJ/kg ± 800 kJ/kg. Its total available energy content is 210 MJ/m² ± 50 MJ/m².

The carpet was modeled as a thermoplastic with density 750 kg/m^3 , specific heat 4.5 kJ/kg/K , conductivity 0.16 W/m/K , ignition temperature 290°C , thickness 6 mm , heat of vaporization $2,000 \text{ kJ/kg}$, heat of combustion $22,300 \text{ kJ/kg} \pm 600 \text{ kJ/kg}$, and total available energy content $61 \text{ MJ/m}^2 \pm 2 \text{ MJ/m}^2$.

The privacy panel was modeled as a thermally-thin thermoplastic. The product of specific heat, thickness and density is $0.73 \text{ kJ/m}^2/\text{K}$. Its surface density is 0.25 kg/m^2 , ignition temperature 380°C , heat of vaporization $6,000 \text{ kJ/kg}$, heat of combustion $30,000 \text{ kJ/kg} \pm 500 \text{ kJ/kg}$. Its total available energy content is $6.0 \text{ MJ/m}^2 \pm 1.3 \text{ MJ/m}^2$.

The test compartment walls and ceiling were made of three layers of 1.27 cm (0.5 in) thick Marinite I, a product of BNZ Materials, Inc. (<http://www.bnzmaterails.com>)*. The manufacturer provided the thermal properties of the material used in the calculation. The density is 737 kg/m^3 , conductivity 0.12 W/m/K . The specific heat ranged from 1.2 kJ/kg/K at 93°C to 1.4 kJ/kg/K at 425°C .

In the simulations of the fires within the WTC, the chair, computer, paper, and other miscellaneous items within the workstation were modeled as a single item by lumping the mass together into large “boxes” and distributing them throughout the workstation. It is common practice in fire protection engineering to use surrogate materials for fire experiments, and this practice has been extended to numerical modeling. Over the years, various items have been developed that are representative of various types of commodities. For example, wood cribs are often used to represent ordinary combustibles found in residential or light industrial settings. Paper cartons with various amounts of plastic within are also used as surrogates for a wide range of retail commodities. One in particular is called the FMRC (Factory Mutual Research Corp.) Standard Plastic Commodity, or more commonly, Group A Plastic. This test fuel is often used in sprinkler approval testing at Factory Mutual and Underwriters Laboratories in the US, and similar test fuels have been developed in Europe. In the late 1990s, FDS was used to simulate large scale rack storage fires to determine the effectiveness of the combined use of sprinklers, roof vents and draft curtains (curtain boards). As part of this effort, a considerable amount of work was done to characterize the thermal properties of Group A Plastic (Hamins and McGrattan 2003). Because Group A Plastic has been shown to be fairly representative of fires fueled by a mixture of paper (cellulosic materials) and plastic, and because it has been used in numerous FDS simulations, it was decided to model the contents of the office workstations with a fuel similar to Group A Plastic. Blind predictions of the single open workstation burns were made using the material properties obtained during the sprinkler/roof vent study, and then these properties were adjusted to match the results of the experiments. Thus, the single workstation burns served to calibrate the model. They were not intended to be validation experiments. The validation experiments consisted of burning 3 workstations at a time in an under-ventilated compartment.

The surrogate fuel is modeled as a homogenous solid whose density is 172 kg/m^3 . The paper carton is treated as a thermally-thin material whose density \times specific heat \times thickness is $1.0 \text{ kJ/m}^2/\text{K}$. Its ignition temperature is 370°C and the heat of combustion is $30,000 \text{ kJ/kg}$. The heat release rate of the boxes ramps up to 450 kW/m^2 in about 1 min . Note that this fuel package is similar, but not the same, as Group A Plastic. The density has been increased to account for all the miscellaneous items within the workstation. Also note that unlike the desk, partition and carpet, the boxes are simply given a burning rate rather than a heat of vaporization, meaning that the boxes will burn at the given rate regardless of the heat flux upon them as long as the surface temperature remains above its ignition temperature. The reason for this is that it is not possible to predict the burning rate using the heat feedback approach because the geometry of the scattered fuel is too complex to directly predict the response of the material to the thermal insult. By collecting all the scattered items into boxes, the geometry of the combustibles is greatly simplified, and as a result the burning behavior must be simplified as well.

Description of the Simulations

*Certain commercial equipment, instruments, or materials are identified in this document. Such identification does not imply recommendation or endorsement by the National Institute of Standards and Technology, nor does it imply that the products identified are necessarily the best available for the purpose.

The geometry of the compartment is relatively simple. The overall enclosure is rectangular, as are the vents and most of the obstructions. Numerical grids of 20 and 40 cm were used to model the fires. The purpose of the grid variation was to ensure that the model was not sensitive to the change in grid cell size. Typically, enclosures of this size are modeled using 10 cm grid cells. However, for the simulations of WTC 1, 2 and 7, a 40 cm grid was used. By simulating the experiments at 20 and 40 cm, we can test if the model produces significantly different results with grid cells of different sizes. Figure 1 is a snapshot of a simulation showing the fire and the major geometric features of the compartment for the simulations. Note that the surrogate fuel packages are placed roughly where the computer monitor, chair and paper were located. Six tests were performed, with various ignitor locations and fuel arrangements. A 2 MW burner was placed either near the windows of the compartment overlooking the workstation nearest the openings in Tests 1, 2 and 3. The burner was placed towards the rear of the compartment overlooking the workstation in the rear of the compartment in Tests 4, 5 and 6. In Tests 3, 5 and 6 Jet A fuel was poured over the workstations and surrounding carpet. To simulate this in the model, spray nozzles were positioned over the center of each workstation, 2 m above the floor. These nozzles are normally used to simulate water sprinklers, but in this case, the water was replaced by a liquid having similar properties to Jet A. The nozzles were activated for 2 s, in which time the equivalent amount of liquid as in the tests was ejected and spread over the furnishings.

In Test 5, Workstations 1 and 2 were disassembled prior to the burn and the contents were piled on top of the respective load cells. To model this scenario, the burning rate of the collective fuel packages was reduced by one half to account for the decrease in burning area of the fuel pile. The choice of one half was somewhat ad hoc. No free burns of workstation parts had been performed. This was the only test in which the simulated fuel packages had to be modified from their free-burn values. In this regard, Test 5 was used to calibrate, not validate, the model.

From a modeling perspective, the objective of the simulations of the Phase 2 experiments was to demonstrate that a simplified model of an office workstation can be used to predict the burning behavior of a group of workstations in an enclosure with features similar to WTC 1, 2 and 7. Because of the magnitude of the simulations of the building fires, the model of the workstation had to be fairly crude. However, because of the many uncertainties in the initial conditions of the fire simulations, the lack of detail in the model is not considered to be a problem. The model fires had similar growth patterns, peak heat release rates, decay patterns, and compartment temperatures.

The model also captured the major features of the individual tests. For example, Tests 1 and 4 were similar in design except for the burner location. In Test 1 the burner was near the windows; in Test 4 it was near the rear of the compartment. The peak heat release rate was reached in about 15 min in Test 1, whereas it was reached in about 10 min in Test 4. The model shows a similar trend. The faster growth of Test 4 is probably due to the fact that the compartment heated up more quickly with the fire deep inside rather than near the windows, leading to more rapid spread of the fire across the pre-heated furnishings. Even though ceiling tiles were distributed over the desk and carpet in Test 4, this did not seem to have a noticeable effect on the growth, or at the very least the burner position seemed to have a far greater role in explaining the difference between Tests 1 and 4. The comparison of HRR between model and experiment is shown for Test 1 in Figure 2. The upper layer temperature in the rear of the compartment for this same test is shown in Figure 3. The results for the other tests are comparable. The peak HRR and temperature are predicted well, as well as the duration of the fires. Both the peak values and the duration of the burning are important for the WTC simulations because it is not only important to predict the temperatures that the structural steel was exposed to, but also the duration of the exposure.

Figure 1. Geometry of the Phase 2 simulations.

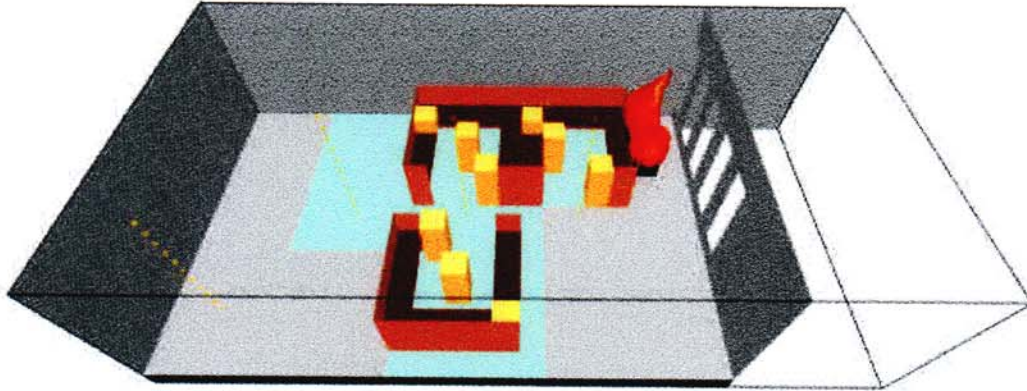


Figure 2. Comparison of HRR for Phase 2, Test 1.

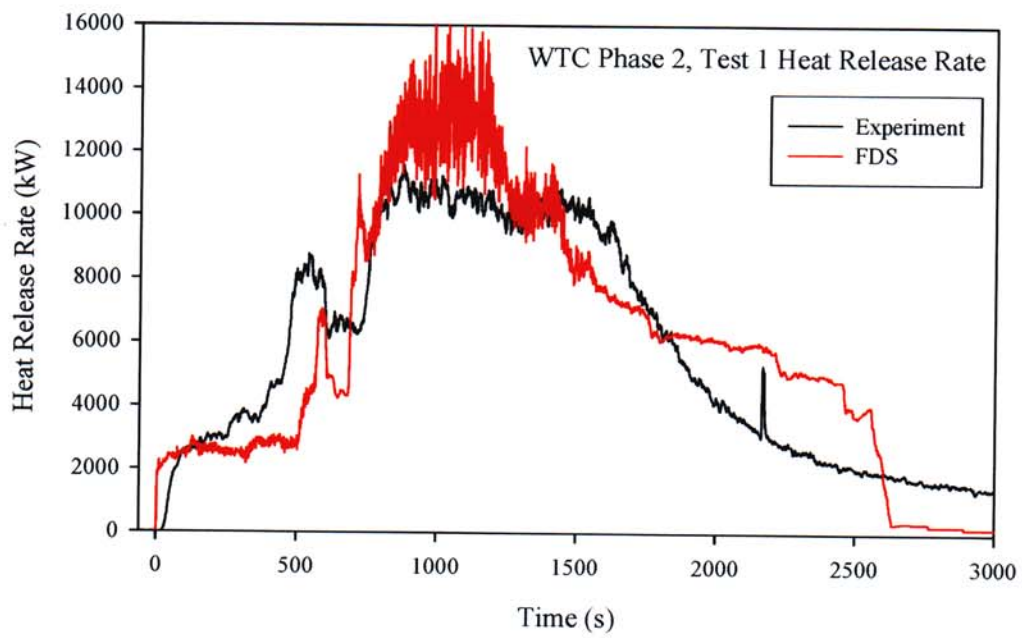
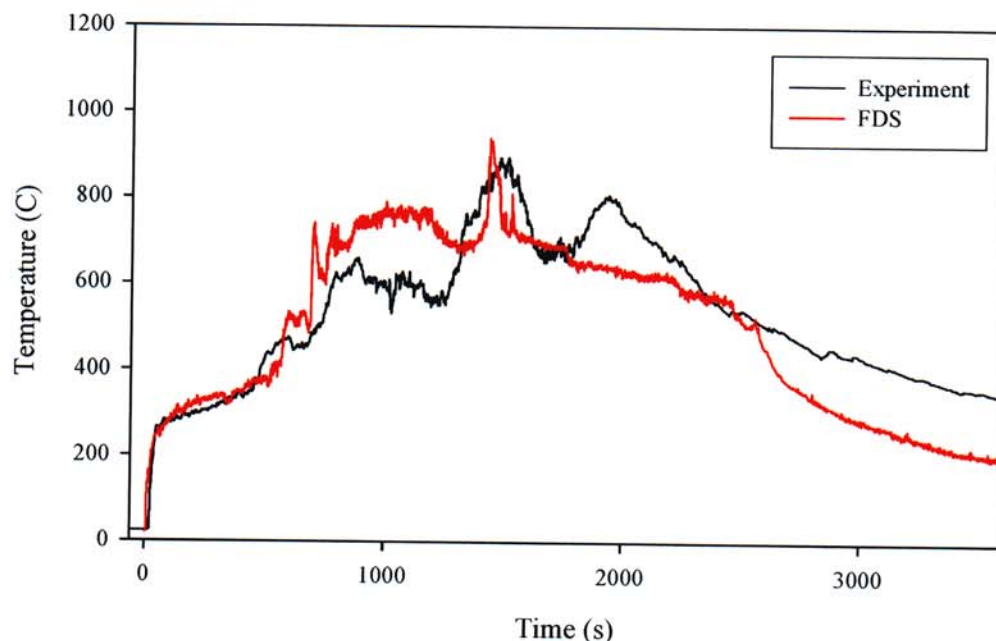


Figure 3. Upper layer gas temperature in the rear of the compartment, Phase 2, Test 1.



SIMULATIONS OF THE FIRES IN WTC 1 AND 2

This section describes how the physical geometry of the buildings was described in the numerical model. Information about the layout of the relevant floors was obtained from architectural drawings provided by the occupants. For floors where information was not available, the geometry of a nearby floor or a floor of similar use was substituted. Information about exterior damage and window breakage was obtained by studying thousands of photographs and videos. There was no attempt made to predict the window breakage in the simulations. This information was provided as a boundary condition.

Numerical Grid

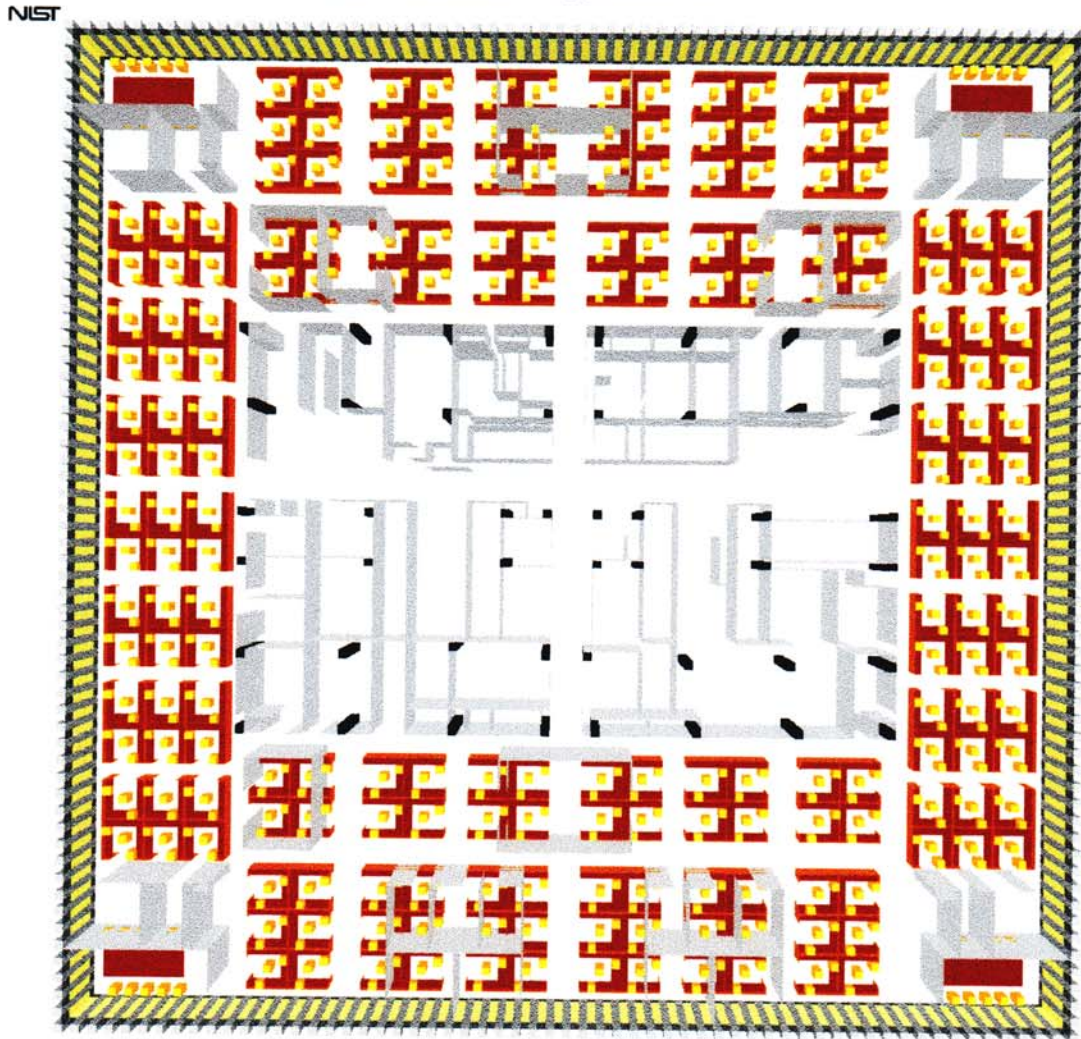
The windows in WTC 1 and 2 were nominally spaced 1 m apart. In addition, the external columns plus their aluminum cladding were assumed to be 0.5 m wide. The slab-to-slab floor spacing was assumed to be 3.6 m. Because of these approximations, a uniform numerical mesh consisting of cells whose dimensions were 0.5 m by 0.5 m by 0.4 m was used. In the model, each tower face consisted of 58 windows, 61 columns, and two 0.5 m spacers next to each corner column. In the real tower geometry, these spacers formed the bevel. Figure 4 shows a single floor of the WTC 1 as it is approximated by the numerical model.

The numerical grid for each floor of WTC 1 and 2 was of dimension 128 by 128 by 9 cells. The 128 cells in the horizontal directions allow for several meters of simulation outside of the external walls. The calculations were run in parallel, thus each floor was assigned to a different processor. The floor slabs, core walls, and workstations were approximated as thin obstructions. As described in the previous section, the contents of each workstation were collected into boxes and distributed throughout.

Penetrations in the floor slabs representing elevator shafts and HVAC ducts were created in the model by defining rectangular plates on top of the floor slab that were removed at the start of the calculation. This served to carve out holes in the floor. Window breakage was modeled by removing thin obstructions serving as windows at times obtained from the analysis of photographs and videos. Broken external columns were


removed the same way.

Figure 4. Plan view of a typical floor in WTC 1.



Parallel Processing

Modeling the fires on multiple floors of WTC 1 and 2 is computationally intensive, both in terms of CPU time and memory. Up to this point in its development, FDS has been limited to calculations small enough to run on a single CPU and fit into the memory of a desktop personal computer. The WTC study is an example of a large-scale fire modeling problem that is impossible to analyze without the use of parallel processing. In terms of parallelization, the exact details of FDS are not important. The approach taken to run the code on a cluster of machines can be applied to virtually any CFD code, in particular those that involve three spatial dimensions and time. In such cases, the computational demand is fairly well represented by the product of the number of computational grid cells and the number of time steps taken to advance the solution of the governing equations in time. For example, if the computational grid consists of one million cells and the simulation requires ten thousand time steps, we say the demand is 10^{10} cell-cycles. We can break down the overall demand into memory requirement and CPU time. The memory requirement is a function of the number



of grid cells; the CPU time is a function of the number of time steps.

Roughly speaking, state of the art 32 bit processors can complete roughly 100,000 FDS cell-cycles per second. Realistic simulations of fires such as those in the WTC require on the order of 500,000,000,000 cell-cycles, or about two months of calculation on a single 2 GHz processor. Plus, the calculation would require 6 to 12 gigabytes of memory (RAM), well over the 4 gigabyte address space of 32 bit processors. Because of this, the WTC calculations are not only impractical on single processor systems, they are impossible on any 32 bit processor. A 64 bit processor system may theoretically handle the static memory requirements of a large simulation time, but the run times for large calculations remain prohibitive.

Because of the computational and memory issues of large fire simulations, a parallel version of the fire model must meet the two fundamental requirements discussed above, as well as satisfy a number of practical implementation concerns. We must distribute both the computational and the memory requirements across multiple processors. The simulation must be done so that each processor uses less than 4 gigabytes of memory, while enough processors must be used to reduce the simulation to a practical length of time, of the order of one week.

Because the computational load is distributed throughout much of the source code, we have chosen to break up the calculation into multiple spatial blocks, with each block essentially doing the same type of calculation. A feature common to most CFD codes is multi-block or multi-mesh structure in which more than one structured grid is used in the calculation. We exploit this feature by simply putting the data and computation for each block on a different processor. This has advantages and some limitations. The advantages are (1) a natural and scalable extension of the existing code, (2) the amount of data communication will be kept to a minimum, since we need only communicate overlap information, rather than the data for full blocks, (3) source code changes are localized in small communication routines, (4) development is fairly fast. The disadvantages to the multi-block approach are (1) equal distribution of work across processors (load balancing) depends on spatial symmetry in the simulation, such as the translationally symmetric geometry of the WTC floors, (2) the level of parallelism and the speed up of the calculation is limited to the number of spatial blocks that can be used in the calculation. These limitations are not severe in many cases, including the WTC.

Because we are interested in a scalable, portable code, we use Message Passing Interface (MPI). This is a standard, well-documented system of implementing parallel processing, that can work with shared memory, distributed memory, or combinations of those architectures (Gropp, 1999). Our goal in using MPI was to produce a code that, except for the requirement of the MPI library, would be as portable and standardized as the sequential version. The parallel code runs on most computer platforms, including networked Windows-based PCs. At NIST, we opted for a cluster of commodity personal computers running Linux, connected by a gigabit ethernet network. The individual processors are in the range of 2.0 to 2.8 GHz and we chose dual processor machines to save space and to allow the addition of OpenMP code as a future extension to the MPI-based code. For production work in our lab, we use two clusters: a smaller, development cluster that was used to develop and debug the code, and a larger cluster with 128 processors. Using both clusters, we have the capability to run six to eight large parallel processing jobs simultaneously.

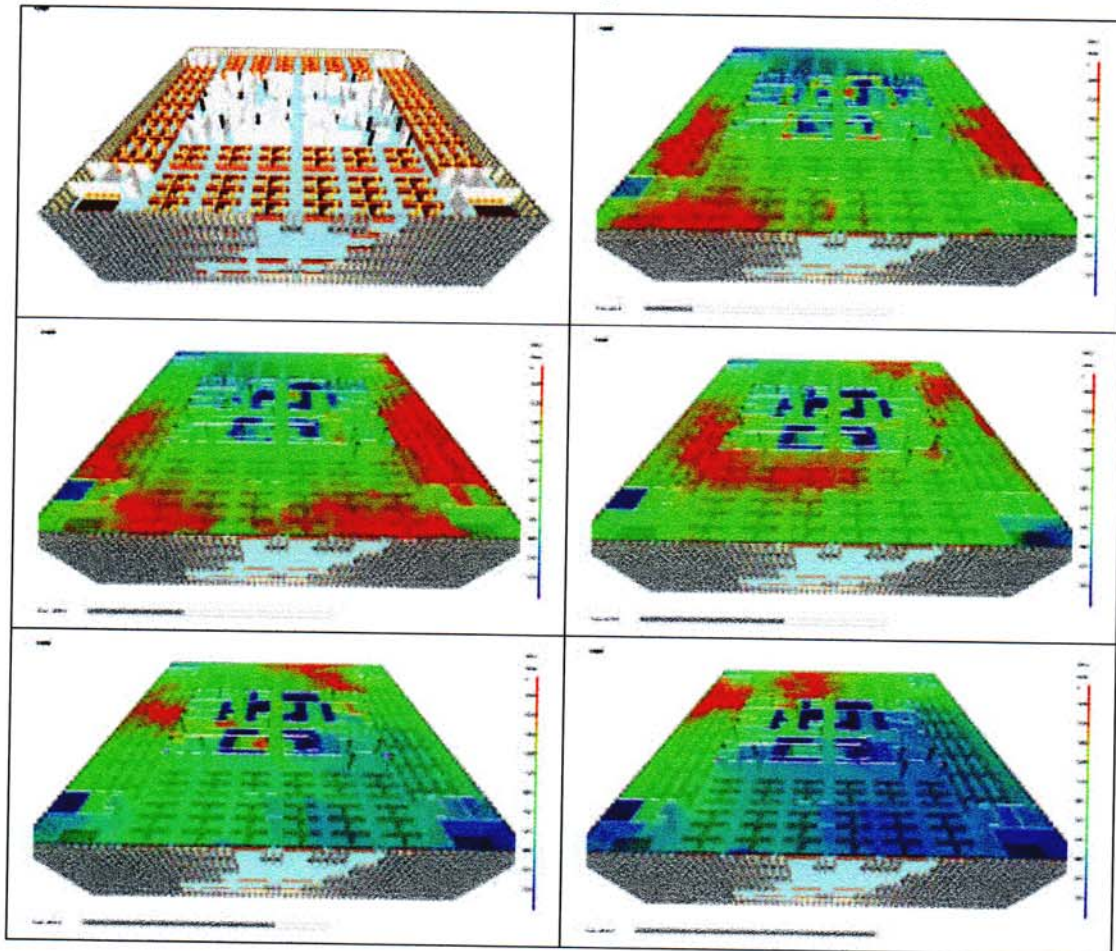
Sample Simulation

Shown in Figure 5 is a sequence of snapshots showing the predicted upper layer temperatures on a floor of WTC 1 at time increments of 15 min. The first image is a cut-away showing the damage to the north face of the tower and the layout of the walls and furnishings. The subsequent images are color coded by temperature, with the red (or dark) patches representing temperatures in the vicinity of 1000 °C. Initially, these hot areas of active burning are near the impact zone at the north end (foreground of picture), but migrate towards the south as the combustible furnishings are exhausted. Driving the progress of the fires is the breaking of windows that provide air to the oxygen starved fire. The window breakage is not predicted by the model; it is an imposed boundary condition resulting from the analysis of thousands of photographs and videos recorded

that day by eye witnesses. The uncertainty in the window break times is on the order of 5 min in areas not obscured by smoke.

The burning behavior shown by the simulation is similar to that of the fires in experiments conducted by Ian Thomas and Ian Bennetts (1999). They looked at fire spread in long and wide enclosures with a single ventilation opening, where the fires were ignited at various points deep within the bench-scale compartments used. The fires would rapidly spread across the liquid or solid fuels covering the floor without consuming much of the fuel. The fires would then surround the compartment opening and burn back into the compartment as the fuel near the opening was exhausted. In the WTC simulations, fires are ignited over a wide area by simulated spray nozzles ejecting a liquid with properties of aircraft fuel. Much of the available oxygen is consumed rapidly, driving the fires to the openings created by the aircraft. The fires move away from the initial impact area as the nearby furnishings are exhausted, and as windows are broken out away in other parts of the building.

Figure 5. Predicted upper layer temperatures of a floor of WTC 1.



SUMMARY

The investigation into the cause of the collapse of WTC 1, 2 and 7 by NIST will not be completed until the fall of 2004. Work is on-going to simulate the weakening of the structural steel due to the aircraft impacts and the fires. Nevertheless, the fire experiments and simulations performed to date have improved our ability to analyze the response of any large building or structure to fire. In the years ahead, these techniques will become increasingly widespread due to faster computers and the ability to harness an entire set of off-the-shelf personal computers to perform very large calculations. Effective modeling is a combination of fast computers, efficient algorithms, and well-planned small and large scale experiments to provide both input to the model and a validation of results. Projects as complicated as the WTC study are rarely conducted using modeling alone. There is and will always be a need to coordinate computation and experiment to reconstruct the dynamics of large fires.

ACKNOWLEDGMENTS

The authors wish to thank all those who participated in the validation experiments for the WTC investigation, including Anthony Hamins, Richard Gann, Alex Maranghides, Tom Ohlemiller, George Mulholland, Rik Johnsson and Michele Donnelly. Bill Pitts provided the time history of window breakage in WTC 1 and 2. Elisa Baker and Johnathan Demarest converted architectural drawings of WTC 1 and 2 into electronic form for input into FDS. Harold (Bud) Nelson aided in defining initial conditions for the simulations. Glenn Forney helped to convert a group of personal computers into a cluster of parallel processors to run the large simulations of the fires in WTC 1 and 2, and also made necessary modifications to his visualization program Smokeview to handle the extremely large data sets generated by the calculations. Kuldeep Prasad developed the necessary data structures and multi-blocking techniques within FDS to enable parallel processing. He also developed the necessary interfaces to convert output from FDS fire calculations into input for the thermal and structural analysis of the steel joists and columns.

REFERENCES

- Gropp, W., E. Lusk, and A. Skjellum. 1999. *Using MPI*, MIT Press, Cambridge, Massachusetts, USA.
- Hamins, A., and K.B. McGrattan. 2003. Reduced-Scale Experiments on the Water Suppression of a Rack-Storage Commodity Fire for Calibration of a CFD Fire Model. Fire Safety Science: Proceedings of the Seventh International Symposium. International Association for Fire Safety Science.
- McGrattan, K.B., H.R. Baum, R.G. Rehm, G.P. Forney, J.E. Floyd, K. Prasad and S. Hostikka. 2002. Fire Dynamics Simulator (Version 3), Technical Reference Guide. NISTIR 6783. National Institute of Standards and Technology, Gaithersburg, MD.
- Ritchie, S.J., K.D. Steckler, A. Hamins, T.G. Cleary, J.C. Yang, and T. Kashiwagi. 1997. The Effect of Sample Size on the Heat Release Rate of Charring Materials. Fire Safety Science: Proceedings of the Fifth International Symposium. International Association for Fire Safety Science.
- Thomas, I.R. and I.D. Bennetts. 1999. Fires in Enclosures with Single Ventilation Openings—Comparison of Long and Wide Enclosures. Fire Safety Science: Proceedings of the Sixth International Symposium. International Association for Fire Safety Science.