The Future of Fire Simulation

By Kevin McGrattan, Ph.D., Howard Baum, Ph.D., Ron Rehm, Ph.D., Glenn Forney, Ph.D., and Kuldeep Prasad, Ph.D.

INTRODUCTION

Scientists and engineers are often asked to make predictions of the state of technology in the future and are usually laughably wrong. The best prognosticators get the trends right, but cannot possibly fill in the details. Think of Jules Verne predicting a trip to the moon, albeit in a projectile decked out in lavish red velvet, manned by champagne-sipping adventurers, and shot out of a giant cannon. Unlike Jules Verne, we dare only look 10 years into the future, rather than 100. Also, we focus on the work going on in the Building and Fire Research Lab at NIST because, even though most agree that modeling will play an increasing role in fire research, the nature of the models is a subject of intense debate. We do not presume to speak for the entire community, and we welcome the opinions of other researchers as to the direction of modeling in the future.

In the Spring 2000 issue of *Fire Protection Engineering*,[1] Howard Baum wrote a brief history of fire simulation in which he listed three major challenges facing fire modelers:

First, there are an enormous number of possible fire scenarios to consider. Second, we do not have either the physical insight or the computing power (even if we had the insight) to perform all the necessary calculations for most fire scenarios. Finally, since the "fuel" in most fires was never intended as such, the data needed to characterize both the fuel and the fire environment may not be available.

Ten years from now, these issues will remain. Certainly the wide range of fire scenarios will persist, even widening due to the constant emergence of new materials and new architectural forms. Computing power will certainly increase, but not to the point of allowing for direct numerical solutions of the governing equations. As models focus in on the small-scale combustion processes in a fire, ever-more complex challenges will emerge that are, for now, neglected. Fortunately, there is hope. The reason is that models based on fundamental principles will improve automatically as computers get faster and the temporal and spatial resolution improves. In looking towards the future, we need to adopt fundamentally sound physical mechanisms that will retain an elegance and simplicity over time, that will shift us from empirical to deterministic descriptions of fire behavior, and that will be useful to fire protection engineers and researchers alike.

BLOWING SMOKE

Even as we develop more sophisticated numerical algorithms to describe the growth and suppression of fires, the majority of design calculations will continue to address a subject for which zone and field fire models were first developed – smoke movement. Because smoke inhalation and carbon monoxide poisoning are and will remain the most dangerous actors in a fire, code officials will continue to enforce regulations designed to ensure safe evacuation of a building. Originally, two-zone fire models were developed to predict the descent of the smoke layer in a fairly simple building, but as building geometries become more complex, fire protection engineers are turning to field models to track the smoke in open-plan, multilevel buildings.

Ten years from now, engineers will still be interested in smoke movement from fires whose *size* and *growth* rate will be predefined. Current field models can handle these problems in theory, but computation times are often too long or the grid resolution is too coarse to capture important features of the flow. The solution to this problem is faster computers, better allocation of grid cells and parallelization, all of which

at NIST

are subjects of active research by computer scientists because the application of these ideas goes way beyond fire.

What we can do now is adopt models that improve automatically as numerical grids become more refined. The best example of this idea is Large Eddy Simulation (LES). We have found over roughly twenty years that good simulations result from solving the Navier-Stokes equations with as few empirical parameters as possible on grids with as many cells as possible. While we always want more, we have found that very good results are obtainable with modest calculations, allowing engineers running our models to investigate a wide variety of problems without having to worry about numerical parameters for which they have little training.

It is our belief that the LES concept will emerge as the prevailing methodology for smoke and heat transport because of its ability to render realistic, time-resolved animations of the flow of gases throughout a building. It is inevitable that as computers get faster, users of CFD models will demand more lifelike simulations rather than time-averaged or steady-state images. This is particularly true of the fire community since the audience for many of the simulations are the authorities having jurisdiction who often have little training in fire modeling. The design engineer must demonstrate to the official what is being calculated with something more than static images or time-temperature plots. Animations of smoke flow provide a visual check of the building geometry, grid resolution, and other features of the calculation that are difficult to convey any other way.

WHAT ABOUT THE FIRE?

Skeptics of fire models have complained from the outset that the fire is not really modeled in a fire model. To a large extent, this criticism is valid, and here are two reasons why. First, most engineers are usually interested in smoke movement, so there's no reason to model the fire other than as a point source of smoke and heat. Second, the combustion processes occur at length and time scales below the resolution limits of most practical calculations, so much so that information obtained from the resolvable scale. like an average cell temperature, is useless in even the most simplistic of combustion models.

Much of the future research in fire modeling will focus on improvements in the way small "subgrid-scale" physical processes are modeled. Examples of such processes include (but are not limited to) soot formation and growth, combustion in vitiated atmospheres, heat transfer to pyrolyzing surfaces, and radiation from gaseous combustion products. All of these phenomena occur in both laboratory-scale experiments and material test apparatus. These processes will serve as the start-

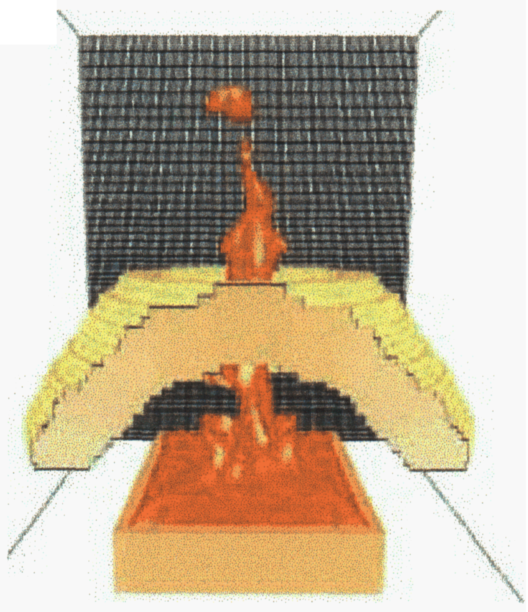

FIGURE 1. **Simulation of a sample of wood burning in the cone calorimeter performed with FDS.** Courtesy Simo Hostikka, VTT Building and Transport, Finland.

ing point for developing better fire submodels because they are well-controlled, relatively simple, and, most importantly, small. Because they are *small,* the calculations can be performed at sufficiently high resolution to capture the important phenomena directly, and then the same calculations *can* be performed at lower resolution to see how well the new algorithms perform on larger-scale simulations. The objective of this effort is not to produce the most detailed description of the phenomena. Very detailed submodels of most fire phenomena exist now; the challenge is to design an overall fire model that balances the accuracy of each submodel. Balance means that the level of detail incorporated into each is roughly the same. Everyone learns in high school that adding a measurement accurate to the nearest millimeter and one accurate to the nearest centimeter yields a result that is only accurate to the nearest centimeter. Similarly, a fire model will only be as accurate as the least accurate of its components.

A good starting point for a better fire model is a well-controlled test apparatus, like the cone calorimeter (Figure 1). A set of solid and gas phase models should be developed that would hopefully provide a reasonable, balanced description of the interaction of the fire with the test apparatus. In essence, this is the procedure that was followed in the development of the Fire Dynamics Simulator (FDS), a general-purpose fire field model released into the public domain in the year 2000. The approach had been to model the large-scale gas phase transport as faithfully as possible for a given numerical grid, and then introduce extra features that were consistent with the detail (or lack thereof) afforded by the smoke transport algorithm. In some sense, the fire itself was just another one of these extra features. At first, the fire was a Gaussian-distributed blob of heat superimposed on the numerical grid, then the fire was a set of hot particles ejected from the burning object, and for the time being the fire is a surface on which fuel and oxygen meet and react infinitely fast. The emerging fire model may move beyond this, but at the moment it should be possible to work with this description of the combustion while the solid phase mechanisms are brought up to par.

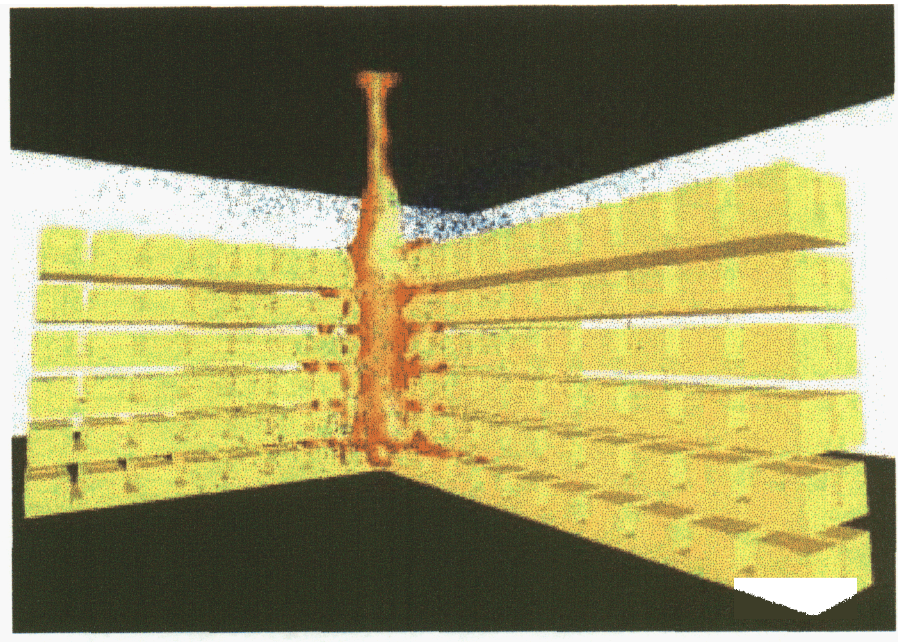

FIGURE 2. Simulation of a rack *storage* commodity fire. The **simulation predicts** the ***growth* and suppression of a fire that originates at the** floor. **These types of simulations are by far the most challenging attempted to date, and it remains to be *seen* how much** the **relatively simple solid** phase **pyrolysis and suppression algorithms can be improved.**

There are two advantages to this evolutionary strategy. First, the various submodels, even in their primitive states, have been useful to FPEs for smoke movement and simple heat transfer calculations, and to introduce the next generation to the technology. Second, all aspects of the simulation improve at the same pace – sprinklers, radiation, burntng objects – so that no part of the calculation looks out of place. A good analogy is classical sculpture. The artist transforms a block of marble into a human form by painstakingly chipping away stone first to reveal the gross outlines of the head, arms. torso, etc.. and only then focuses in on finer details. Consider that the most beautifully sculpted hand would look ridiculous if one arm were longer than the other.

INDUSTRIAL-STRENGTH FIRES

A few years ago. in parallel with large-scale tests, the development of FDS turned towards the problem of fire suppression in large warehouses and warehouse retail stores[2] (Figure 2). As simplistic as the combustion and heat transfer algorithms were at the time, they were not nearly as primitive as the sprinkler spray and suppression models. With the exception of the thermal activation equation, which by that time had become widely adopted, the water droplets emerging from the pipe, landing on the commodity, and eventually interacting with the fire were by far the greatest source of uncertainty in the simulation.

A series of bench-scale experiments was conducted at NIST to develop necessary input data for the model. These experiments generated data describing the burning rate and flame spread behavior of the cartoned plastic commodity, thermal response parameters and spray pattern of the sprinkler, and the effect of the water spray on the commodity selected for the tests. The missing link in the analysis was the spray characterization of the sprinkler itself; that is, the water was assumed to leave the sprinkler in a simple umbrella pattern quantified only by visual observation. What made the model work reasonably well was the fact that the spray parameters were tweaked until a match between computed and observed water density patterns on the floor was obtained. Hundreds of hours were needed to roughly characterize one fuel and one sprinkler because the characterization

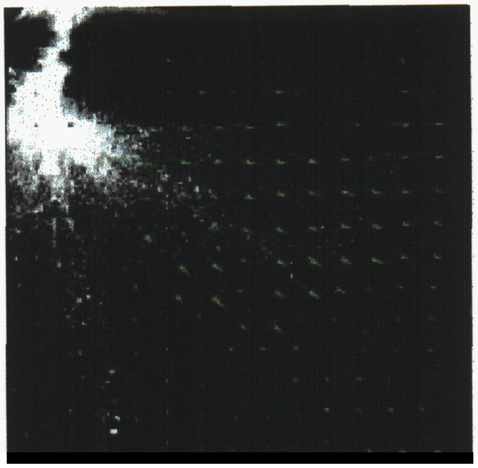

FIGURE 3. Image of a sprinkler spray created with Particle Image Velocimetry (PIV). The green arrows represent the velocity vectors of water droplets leaving the sprinkler orifice. The technique involves taking two photographs of the spray in rapid succession and backing out the velocity from the displacement of the droplets. Courtesy Dave Sheppard, Northwestern University and Underwriters Laboratories.

was almost all empirical – little of it was based on fundamental physical models because the phenomena was so very complex. As a result, users of the FDS model were not able to apply it easily to other commodities and sprinklers; a problem that persists to this day.

Sprinkler spray characterization will remain largely empirically based because each sprinkler has its own unique design that makes predicting which way the water will go difficult. To simulate the sprinkler spray, we need to know the initial distribution of the droplet size and velodty. Measuring these quantities has proven to be very difficult and still very expensive. The most promising technique for measuring droplet size is through Phase Doppler Interferometry (PDI) and droplet velodty through Particle Image Velocimetry (PIV) (see Figure 3). Both are nonintrusive, laser-based techniques that require very expensive equipment and skilled technicians with a high level of training in laser diagnostics. This is worrisome because calculations should be cheaper than experiments, or else what's the point? If high-level modeling of challenging industrial fire scenarios becomes more routine and starts to show potential benefits to sprinkler manufacturers and building owners, there ought to be more investment in the measurement techniques required for input data. The catch-22 is that it's hard to show benefits with little data.

Understanding how various standard commodities burn and how they respond to water ought to be less empirically based than sprinkler sprays, assuming the necessary solid phase models are developed that retain enough of the fundamental physics to accommodate a better description of suppression, yet simple enough to be used in large-scale simulations. We discussed above the need for more fundamentally based models of pyrolysis, starting with relatively small-scale calculations of standard test apparatus and eventually moving to large-scale. It is unclear how to describe the burning of real commodities, which are mixtures of cardboard, plastics, woods, etc., other than with the simple lumped parameter models developed to date. It is hoped that at a minimum, we will have a way of relating the burning rate of the fuel to the heat feedback to the surface based on the thermophysical properties of the fuel rather than simply an exhaustive series of experiments that are often too expensive to perform given the wide variety of fuels burning in a single fire. This is possible now with a limited number of pure fuels, liquids especially, but hopefully this list can be extended in the future.

THE FIRE HAS LEFT THE BUILDING

The fire models with which we are familiar were developed to describe residential and commercial building fires. However, there is a different class of models developed by the forest management and agricultural communities designed to predict the spread of wildland and forest fires. These models are semiempirical and are built upon very different assump-

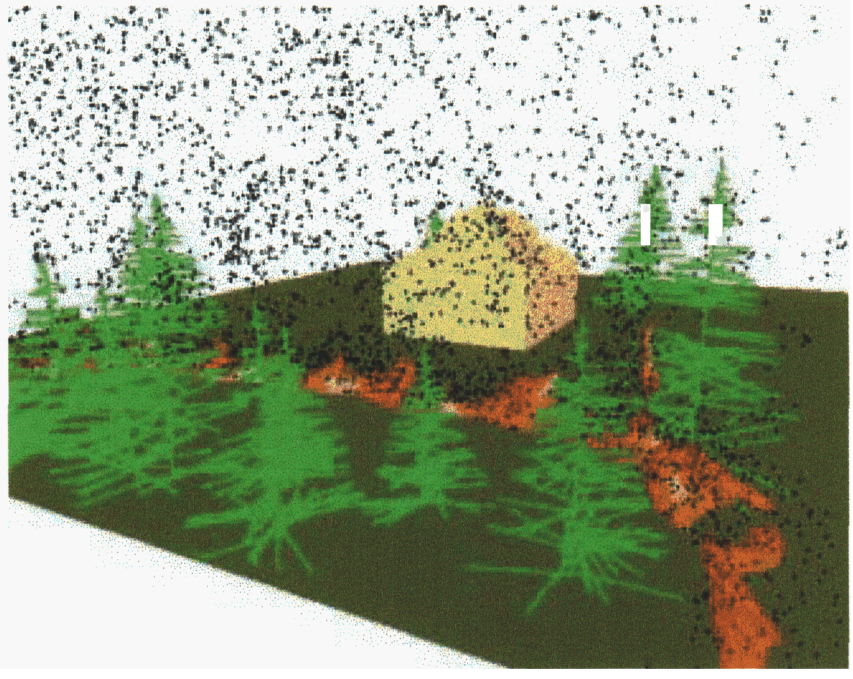

FIGURE 4. Simulation of a brush fire advancing on a house. Preliminary calculations such as these are now being performed to assess the feasibility of simulating community-scale fire spread. Here, the domain is a few hundred meters on a side, the grid cells about 1 m. The trees serve as a drag on the oncoming wind.

tions than building fire models. They are closer to zone models in philosophy since they are designed to run faster than real time so that they can be used the day of the actual fire. It would be hard to believe that field models for community-scale fire spread could be developed in ten years' time to be run in real time, but field models are being developed now, both at NIST and elsewhere, to study the behavior of large outdoor fires to aid in planning efforts.

Last year, wildland fires cost an estimated one billion dollars for just the firefighting. However, fires in the built environment are generally even more costly. The famous 1991 Oakland Hills fire in Berkeley, CA, alone did an estimated $1.5 billion in damage while killing 25 people and injuring 150 others. While large-scale models of fire propagation are useful in wildland settings, corresponding models for community-scale (rural, suburban, or urban) fire spread, i.e., fires spreading between structures and natural fuel, are still in their infancy. Development of such models suffer from the following Catch-22 – validated models of community-scale fire spread are needed because experiments on that scale are almost impossible to carry out: but without experiments, how do we validate the models?

The main numerical problem to community-scale fire prediction is grid resolution. Consider a square kilometer of terrain containing both structures and dry vegetation. Any field model tracking the progression of a brush fire through the area would require several million grid cells, which, even if cleverly distributed, would provide spatial resolution of at best a meter. Existing large-scale models of wildland fires regard the fuel (vegetation) as continuous and assume the fire to propagate as a line. Resolvable-length scales for these models are tens to hundreds of meters. The technical challenge for the community-scale fire model is to develop a mathematical description for the ignition and burning of individual trees and shrubs, and to determine fire spread between wildland elements and structures. Such a mathematical description must include fire spread by brand generation, transport, and subsequent ignitions for both wildland fuels and structures. As with any useful model, these descriptions must be validated using experimental data and must then be integrated into a CFD flow solver generalized to account for an atmospheric boundary layer flow conditioned by natural topography, upwind structures, and trees.

In addition to the numerical challenges posed by community-scale fire prediction, it is often difficult or impossible to obtain meteorological and topological information in a form that can be used in the calculation. The meteorological conditions driving the fire have to be postulated or derived from a mesoscale weather model with a minimum resolution measured in kilometers and the terrain features obtained from a database that may or may not exist at the required resolution for that particular patch of the earth. Fortunately, there are now efforts within the meteorological and geographical research communities to develop and maintain models and databases that would be useful to the fire community. For example, digital elevation data from LIDAR overflight measurements is being made increasingly available and cost-effective.

PROOF BY PRETTY PICTURE

Modelers are looked upon with skepticism by the rest of the fire community because of the perception that they all too often hide behind an eye-catching image or animation without understanding the physics underpinning the model. In fact, some have started to refer to CFD as "colorful fluid dynamics." This is often a fair assessment, but it is short-sighted. While the rapid improvement in visualization techniques has been a boon to many in the field who use field models on a regular basis, within the next 10 years what is now, gee-whiz will become ho-hum. This is good because as field modeling becomes routine, the discussion will be raised beyond the superficial level we are at now to a point where the quality of a simulation will be judged by the spatial and temporal fidelity of its images and animations. With any field model, the user chooses a numerical grid on which to discretize the governing equations. The more grid cells, the better but more time-consuming the simulation. The payoff for investing in faster computers and running bigger calculations is the proportional gain in realism manifested by the images.

As the community at large becomes accustomed to looking at various pictures and animations, model developers will find new ways to dazzle. Up to now, most visualization techniques have provided useful ways of analyzing the output of a calculation, like contour and streamline plots, without much concern for realism. A rainbow-colored contour map slicing down through the middle of a room is fine for researchers, but for those who are only accustomed to looking at real smoke-filled rooms, it may not have as much meaning. Visualization in the next 10 years will turn towards providing as much information as the rainbow contour map but in a way that speaks to modelers and nonmodelers alike. Take, for example, Figure 5. Presented are two ways of visualizing the same calculation, each figure made for a different audience. The trend in scientific visualization is to combine the features of each into one to reach both groups of people.

A good example is smoke visibility. Unlike temperature or species concentration, smoke visibility is not a local quantity but rather depends on the viewpoint of the eye and the depth of field. Advanced simulators and games create the illusion of smoke or fog in ways that are not unlike the techniques employed by fire models to handle thermal radiation (Figure 6). It is envisioned that eventually graphics hardware and software will play a role in actually computing results rather than just drawing pretty pictures. AHJs often ask whether or not building occupants will be able to see exit signs at various stages of a fire. The fire model can predict the amount of soot in each grid cell, but that doesn't answer the question. The harder task is to compute on the fly within the visualization program what the occupant would see and not see.

CAN'T YOU MAKE IT GO FASTER?

Computational Fluid Dynamics was built around weather prediction and aerospace design. A quick browse

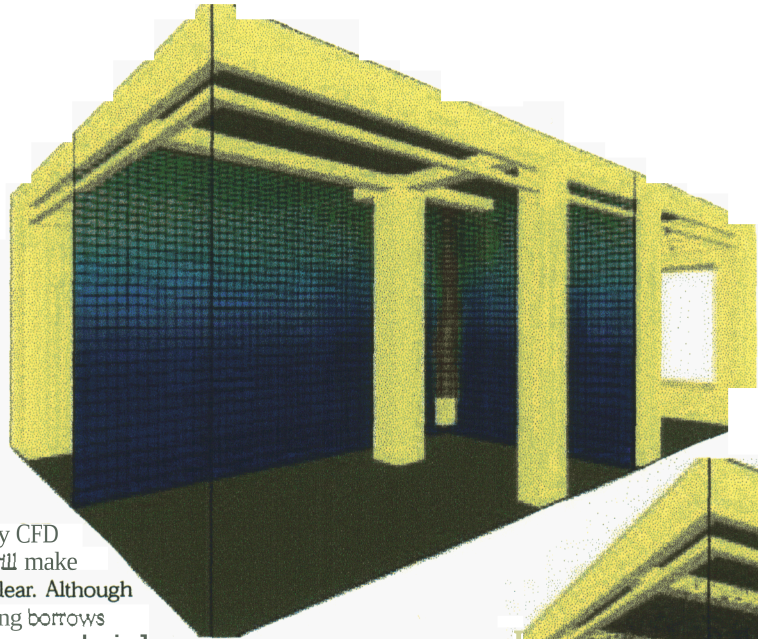

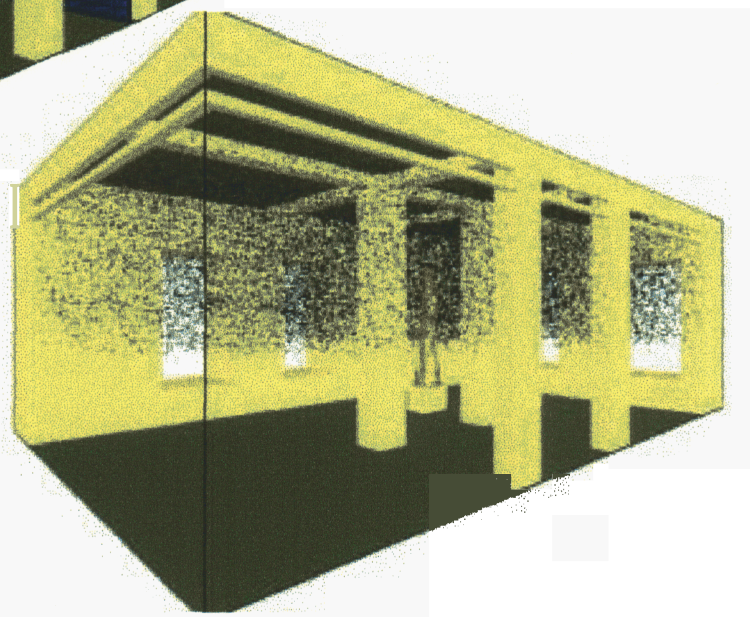

FIGURE 5. Two different ways of visualizing a fire simulation. On the left, contours of gas temperature are shown superimposed on the numerical grid. On the right, the fire and smoke are shown as an orange-colored sheet and black dots.

through any CFD textbook will make this point clear. Although fire modeling borrows many of the same physical assumptions from weather modelers and numerical algorithms from the aerodynamics community, it is different in one important way – it is practiced by relatively small organizations whose engineers have limited backgrounds in numerical methods and computing. Although many small FPE firms have been absorbed by larger, more diversified design and architecture firms, the typical fire modeling effort within one of these organizations is modest – a few engineers working for a few weeks on a given design problem with computers not much more powerful than those found in the home. The reason for this is that fire protection is but one of many features of an overall building design, and one which is typically squeezed when other items in the budget run in the red.

Because of how it is practiced, fire modeling has always emphasized simplicity and efficiency. One of the first questions that we are asked whenever we demonstrate the latest simulation is how long did it take to set up the case and how long did it take to run. The answer to both of these questions needs to be on the order of a day or less (and don't tell me I can't run it on my laptop). If it's more, then we've lost 90% of our audience. This presents us with a bit of a dilemma – how do we stay at the forefront of CFD but still serve the community of practitioners? One way is to design the model so that one can easily start doing simple calculations with simple geometries and then systematically work up towards more elaborate appli-

FIGURE 6. This figure is an example of how fog is used to bring more realism into the scene. Shown is a simulated smoke plume made with the ALOFT (A Large Outdoor Fire plume Trajectory) model. The plume is embedded within a fog-shrouded oil field visualized with the commercial software package SGI Performer.

cations. The trouble with many types of engineering software is that it is packed with so many features that the learning curve for even the simplest of problems is too steep. Fire modeling, especially field modeling, will advance only if there is a large enough core of users to justify the time and expense of developing and maintaining a very complex computer code. Sustaining that core of users means making the software accessible to a wide audience.

Not only must the software be easy to use, but the calculations must run as fast as possible. Veteran CFD practitioners do not find week-long calculations unusual, but fire protection engineers who only have experience with zone models find it intolerable. Faster computers have soothed some, but the demands for moredetailed calculations often negate gains made in computer speed. To keep up with demand, the flre models will need to exploit advances in computer science and numerical methods that go beyond just faster chips. Parallel processing is becoming more of a reality in certain fields, but still is a few years away for those using the current generation of personal computers. However. in the not-too-distant future, relatively inexpensive desktop computers will come with 2, 4, or 8 processors, plus the necessary hardware and software to make these chips work together effectively. Also, techniques to better distribute the grid cells will allow for greater flexibility in the design of simulations. One technique that is used by many CFD packages (but not yet FDS) is called multiblocking. An example of how this would work is a house in which every room has its own numerical grid. Those rooms requiring more spatial resolution could have finer grids, those that don't need it could remain coarsely gridded. The numerical algorithms presently used in single-block codes will not change except there needs to be extra logic built into the algorithms so that information is properly communicated across block interfaces. Such a technique is perfect for flre models because most simulations investigate buildings with relatively simple. rectangular geometries. Contrast this with the aerospace industry where simulations are performed on very complicated body shapes. These models utilize numerical grids that are far more sophisticated and difficult to construct than those needed for fire models.

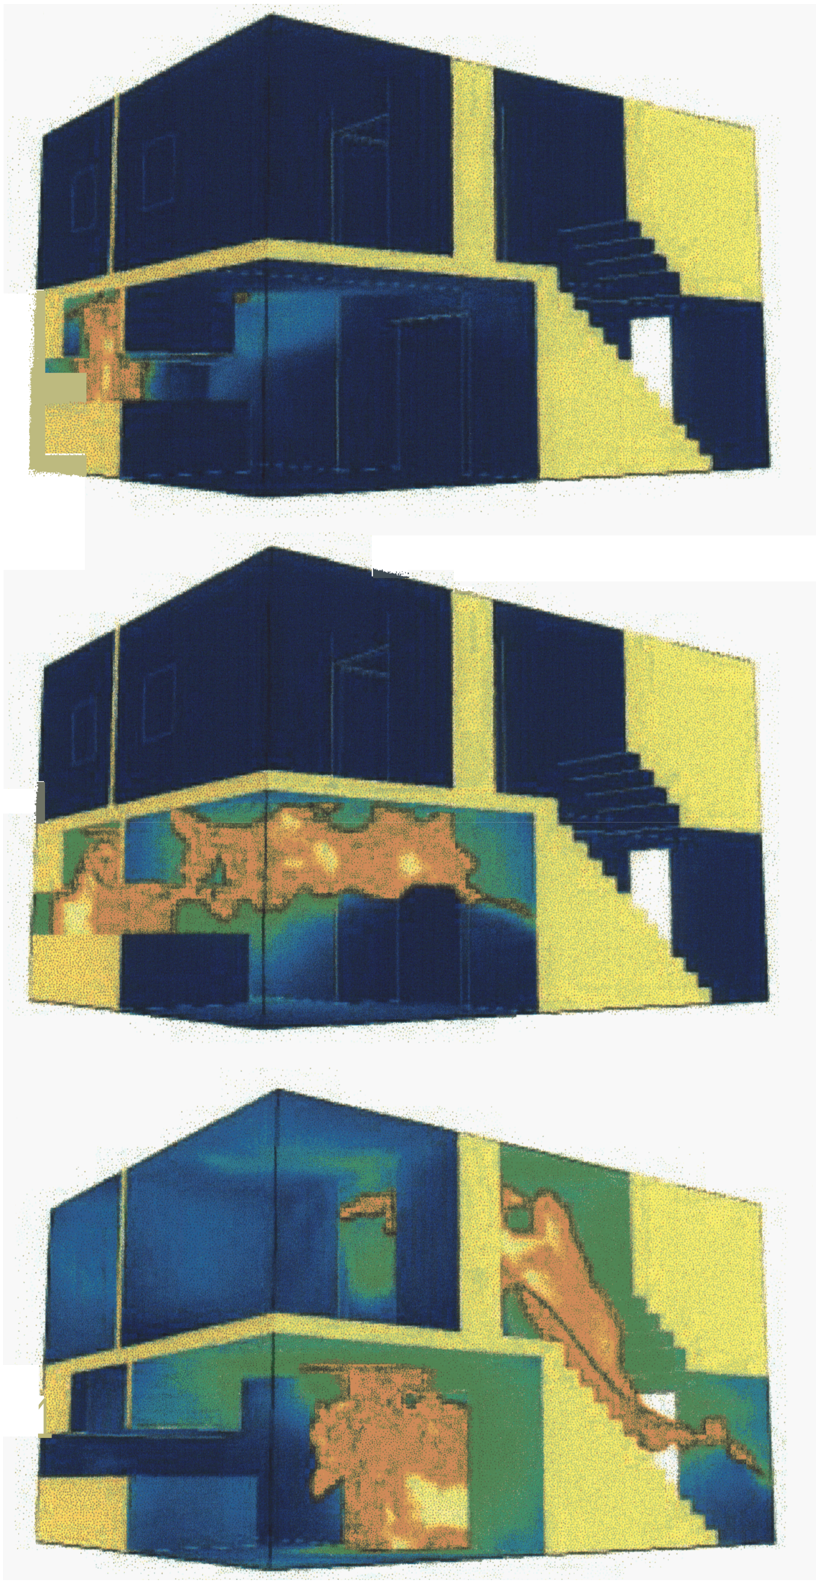

FIGURE 7. Several snapshots of a fire spreading through a townhouse. The tire originates in the Wtchen area (lower left), and eventually spreads throughout the house. The front door (right) is assumed to be open, as are the windows on the second level.

WHO'S GOING TO USE FIRE MODELS?

The discussion thus far has focussed mainly on the application of fire models to design problems. This is not surprising, since fire models have traditionally worked best when the fire itself was considered merely a model input rather than a model output. However, fire models have been used as forensic tools in the past, and their use as such will accelerate in the future. In fact, much of the work to improve fire models past the point of just smoke movement has benefited the fire investigators moreso than the designers who are most often content to dial in a design fire rather than try to predict its growth and suppression.

The fire service in particular has traditionally been skeptical of any type of model, usually preferring full-scale experimental data over computer simulations. However, recent work[3] with fire models to reconstruct several fire losses has moved some in the fire service to consider the use of fire models as training tools for fire fighters. If the present interest in simulation by the fire service continues, a great deal of effort will be placed on understanding the spread of fire through an entire house, not just a single room. Reconstructing a raging house fire goes way beyond simple smoke and heat transport because, as the fire spreads from its point of ignition to envelop entire rooms in flame, the response of the wall materials becomes tightly coupled to the progression of the fire in a way that up to now fire models have largely neglected. Presently simulations of entire house fires are being performed to roughly scope out the grid resolution and physical mechanisms necessary to at least capture qualitatively the sequence of events from primary ignition to second-object ignition to room flashover to room-to-room fire spread (Figure 7). It is difficult to validate such calculations, but at least we are starting to understand what we're up against. Validation will come from more controlled single-room experiments, like the ISO 9705 room corner test, and from simulations of test apparatus.

Ultimately, the users of fire models, whether they be researchers, design engineers, firefighters, or litigators, will drive the development of new features. The challenge to developers is to create fire models that can be used and understood by all of these groups. Even if some do not need the entire set of model features, their use of the models will speed the development and acceptance of them because much of the skepticism associated with modeling will diminish as more people grow comfortable with the capabilities and limitations.

ACKNOWLEDGMENTS

The work described in this article is the contribution of many people at NIST and beyond. We are indebted to a number of guest researchers and post-doctoral associates who have contributed to the modeling effort over the past 10 years, including Dike Ezekoye, Ruddy Mell, Javier Trelles, Francine Battaglia, Jason Floyd, and Simo Hostikka; in addition, our colleagues at NIST for input data, validation experiments, and software testing: Jason Averill, Dave Evans, Paul Fuss, Bill Grosshandler, Anthony Hamins, Dan Madrzykowski, Rick Peacock, Bob Vettori, Doug Walton, and John Widmann. Finally, thanks to the many users of our various fire models, for useful feedback, suggestions, and critiques.

The authors are with the National Institute of Standards and Technology.

REFERENCES

1 Baum, H.R., "Large Eddy Simulations of Fire." *Fire Protection Engineering*, 6:36-42, Spring 2000.

2 McGrattan, K.B., Hamins, A., and Stroup, D., "Sprinkler, Smoke & Heat Vent, Draft Curtain Interaction – Large Scale Experiments and Model Development." Technical Report NISTIR 6196-1, National Institute of Standards and Technology, Gaithersburg, MD, September 1998.

3 Madrzykowski, D., and Vettori, R.L., "Simulation of the Dynamics of the Fire at 3146 Cherry Road NE, Washington, D.C., May 30, 1999." Technical Report NISTIR 6510. National Institute of Standards and Technology, Gaithersburg, MD, April 2000.