HI-02-18-1

Modular HVAC Simulation and the Future Integration of Alternative Cooling Systems in a New Building Energy Simulation Program

Richard K. Strand, Ph.D.
Member ASHRAE

Daniel E. Fisher, Ph.D., P.E.
Member ASHRAE

Richard J. Liesen, Ph.D.
Associate Member ASHRAE

Curtis O. Pedersen, Ph.D.
Fellow ASHRAE

ABSTRACT

The interest in both modular simulation and alternative cooling systems continues to rise both in the United States and in other countries, particularly those nations where concerns for climate and the global environment are high. Modular simulations allow program users to test out configurations that are different from the standard systems and may sometimes lead to innovative design solutions. At the same time, the U.S. Department of Energy has released its new building energy simulation program, EnergyPlus. Its integration of a modular HVAC simulation within the framework of a comprehensive building thermal simulation has resulted in a program with significant initial capabilities and flexibility. This relative flexibility in comparison with its parent programs and the feedback from the integrated building, system, and plant simulation modules allows users to investigate the use of nonstandard systems and will eventually allow the analysis of more complex alternative systems. This paper first provides an overview of the HVAC simulation method of the new program as background and then discusses some of the alternative cooling system models that are made possible by the simulation environment and will hopefully be implemented in a more integral way within the program in the near future.

INTRODUCTION

One of the largest problems with the major hourly building energy simulation programs developed in the United States in the past, BLAST (Building Systems Laboratory 1997) and DOE-2 (Winkelmann et al. 1993), has been the lack of ability to accurately model HVAC systems that are different from the "standard" systems that one tends to encounter in most buildings. Both programs have relied on template systems that were intended to cover most typical building installations but were not extremely flexible in modeling various options or improvements to the systems that are continuously being developed. Each system type was formed by various pieces of equipment in a set order that could not be altered. While it is true that standard systems are used quite often and, in general, adequately meet the temperature and humidity requirements of the spaces that they are serving, the lack of flexible models in BLAST or DOE-2 resulted in a significant lag between the time when new technology is developed and when it could be simulated in these programs.

Once outside the capabilities of the template systems of BLAST or DOE-2, the user typically had to either find a way to manipulate output data to mimic the performance of the altered system or request that changes or enhancements be made to the existing model. The process of using data produced by a standard system and attempting to alter that data outside the program can be quite a challenge. One must know enough about the model to understand any detailed output (if such output was even available) to be able to create a second model to read that data and produce a second set of output. Once outside the program, the user no longer has the same confidence factor that the program itself has and may need to prove that the data produced outside the program are accurate.

Requests for changes or additions to the programs were also fraught with difficulties. While, in theory, some of the major simulation codes produced in the United States over the last several decades were "open" and available to the public, one might argue that the code was really only open in theory but not in practice. Large programs of the past were difficult to understand and required a significant investment of time and expertise in order to understand the program. Unfortunately, understanding the program or even just a limited

Richard K. Strand is an assistant professor in the School of Architecture at the University of Illinois at Urbana-Champaign. **Daniel E. Fisher** is an assistant professor in the School of Mechanical and Aerospace Engineering at Oklahoma State University. **Richard J. Liesen** is a senior research engineer and **Curtis O. Pedersen** is a professor emeritus in the Department of Mechanical and Industrial Engineering at the University of Illinois at Urbana-Champaign.

portion of the program was not enough because of the interconnection of algorithms and data within the codes. Basically, modifying the legacy codes required an expert in the particular program and a financial commitment of some kind to fund the work. New technology or systems do not always have enough interest at first to fund such work. In addition, the number of experts for the programs was relatively small (on the order of a handful per program), meaning that there was also a lack of experts to keep the programs up-to-date. This led to a significant divide in the 1980s and 1990s between what was “common” in industry and what could be simulated by the hourly building energy simulation programs.

The new program (Crawley et al. 1997) tries to address all of these concerns and solve the major obstacles to keeping simulation programs from being used and being useful for a much wider segment of the architectural and engineering communities. A major effort was made to create and enforce strict programming standards. These rules help to unify the program with a common appearance and organization. The program was also restructured using a modern programming language (FORTRAN90) to be as modular as possible. Modularity addresses the issues of data sharing and the need to understand a large portion of the program to make small changes. With the modularity of the new program, it is much easier to modify the program and takes less time for those unfamiliar with the code to make changes because only a small portion of the code needs to be understood in order to make changes. This significantly increases the number of people who will be able to upgrade the program and also cuts down considerably on the financial investment required to produce a new model.

Moreover, the designers of the HVAC section of the program strove to make the simulation flexible with some defined limits. The goal was to be a compromise between the template-based programs such as BLAST and DOE-2 and the completely modular programs such as TRNSYS (Solar Energy Laboratory 1990), HVACSIM+ (Clark and May 1985), SPARK (Buhl et al. 1993), etc. The intention was to gain some of the benefits of each simulation strategy while avoiding some of the potential drawbacks. This compromise allows the user to map components that are to be analyzed in whatever order they would appear in the building being investigated. The benefit to the strategy chosen for the new program is that there is a greater ability to switch component order or types without having to program a new model. Thus, many of the common modifications to the standard systems can be modeled through input file changes rather than code changes. This flexibility is simplified from more detailed simulation approaches that require users to connect equipment and flow variables. In the new program, the user must simply map the components and their interconnections at the macro level (component to component) rather than the micro level (variable to variable).

The flexibility of the program is an asset that has paid dividends in the development process and will eventually benefit the simulation community through an ever-increasing list of capabilities as models of new HVAC technology are integrated with the new program. One area where this flexibility benefit is already showing its usefulness is in the area of radiant systems (Strand 2001; Strand and Pedersen 2002). Many of the modular aspects of the program made the installation of the initial radiant system model simpler than with previous research programs. In addition, the flexibility of the program provided the opportunity to quickly investigate the performance of a modified version of a radiant system. The combined modularity and flexibility made it possible to model a scenario that could not have been investigated in either BLAST or DOE-2 and provides evidence that these two features will be valuable to both users and model developers alike.

The remainder of this paper provides a discussion of the HVAC simulation technique that has been implemented within the new program, focusing mainly on the primary system equipment. In addition, it provides an overview of the radiant system model as well as a look into what types of alternative systems can be implemented in the future and the relative difficulty of such integrations.

CURRENT PRIMARY SYSTEM MODELING METHODOLOGY

In general, the HVAC specification and solution scheme can be seen as various groups collected into hierarchies. For example, there are three main loops within the HVAC simulation in the new program: an air loop, a plant loop, and a condenser loop. The air loop is assumed to use air as the transport medium as part of an air-handling system, while the plant and condenser loops may use a fluid of the user’s choosing (typically water). A user may have any number of each type of loop in a particular input file. There are no explicit limits on the number of loops within the program—the user is only limited by computer hardware. Execution speed will naturally vary with the complexity of the input file.

Main loops are further divided into “subloops” or “semiloops” for organizational clarity and simulation logistics (see Figure 1). These subloops are matched pairs that consist of half of a main loop. For example, the air loop is split into “air loop” and “zone equipment” halves. Each half of the loop has a distinct function: the air loop contains centralized equipment, while the zone equipment portion of the loop contains

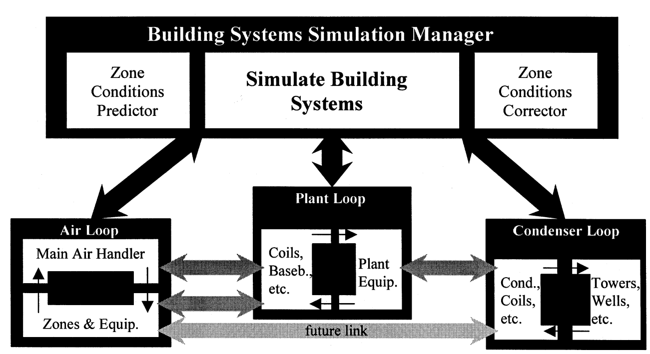

Figure 1 Connections between the main HVAC simulation loops and subloops.

zones, terminal units, zone-specific equipment, etc. Plant and condenser loops are broken into supply and demand sides. The plant demand loop contains equipment that places a load on the primary equipment. This might include coils, baseboards, radiant systems, etc. The load is met by primary equipment, such as chillers or boilers on the plant supply loop. Each plant supply loop must be connected to a plant demand loop and vice-versa. A similar breakdown is present on condenser loops where the demand side includes the water side of condensers, while the supply side includes condenser equipment such as cooling towers.

The breakdown into subloops allows for better handling and control of information and simulation flow throughout the program. Figure 1 shows an overview of the HVAC system groupings within the new program. Direct connections between the subloops of the air, plant, and condenser loops are enhanced by indirect connections between the various main loop types. For example, coils (heating or cooling) are in reality heat exchangers with an air and a water or refrigerant side. The air side of the coil is handled within the air loop where the control of the device is also maintained. The fluid side of the coil is handled within the plant demand side, which passes the energy requirements of the coil onto the plant supply side. All loops are simulated simultaneously, though subiteration loops are maintained between the two sides of any loop to speed convergence. Overall iterations ensure that the results for the current time step are balanced and updated information has been passed to both sides of the subloops as well as across to the other side of indirect connections such as coils.

Branches further divide the subloops into groups as they would appear within any HVAC system. Elements can be defined in series, in parallel, or both with some restrictions. Figure 2 provides an overview of a generic subloop representation. Branches are defined as individual legs within the loop structure. Thus, the segment between point A and point B is defined as a branch, as is the section between points E and F. There may be multiple sections (C1 to D1 through Cn to Dn) in between the splitter and mixer. Each subloop may only have one splitter and one mixer. Thus, equipment may be in parallel between the mixer and splitter; however, within any branch, there can only be elements in series and not in parallel. The topology rules for individual subloops allow a reasonable amount of flexibility without requiring a complicated solver routine to determine the actual flow and temperature conditions. Note that since plant supply and demand are broken up into two separate subloops, chillers or boilers may be in parallel to each other in the supply side and coils may be in parallel to each other on the demand side. Thus, the restriction of only a single splitter and mixer on a particular subloop does not limit the normal configurations. Also, a subloop does not require a splitter or mixer if all equipment on the subloop is simply in series—this would correspond to a single branch that would define the entire subloop.

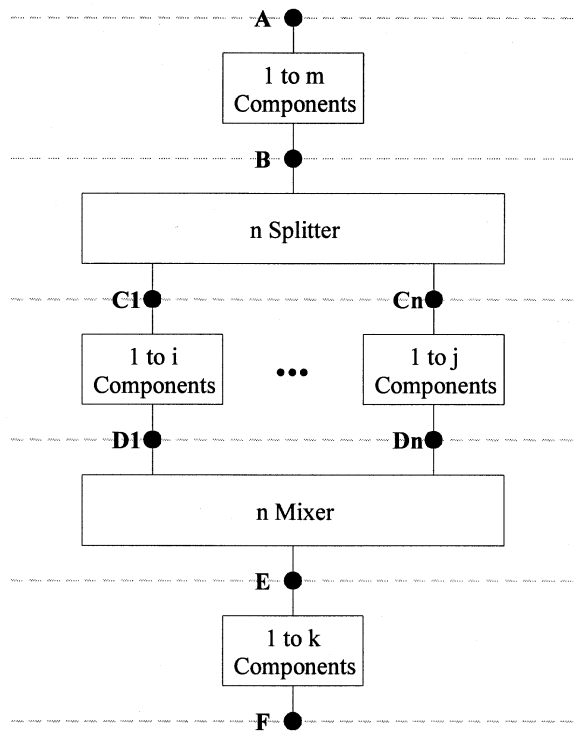

Figure 2 *Branch layout for individual HVAC subloops.*

A basic description of how branches are constructed of components and nodes is available in the literature (Fisher and Taylor 1999). Essentially, each branch is made up of one or more components linked together in series. The branch has an information node containing properties of the loop (temperature, enthalpy, flow rate, etc.) at the beginning and end of the branch as well as between components. Components on the branch take the conditions of the node at their inlet and use that information as well as overall control information to simulate the component and write the outlet data to the node following the component. This information is then used either by the next component on the branch or establishes the outlet conditions for the branch.

Resolution of Flow Conditions

One of the most important aspects of the solution procedure within the plant and condenser loops of the new program is the method used to solve the various subloops. This involves making the supply side meet a particular load based on the simulation of the demand-side loops. Load distribution is an issue that must be addressed as well as how flow rates are adjusted and temperatures are updated. These issues are discussed in the next several subsections, and the algorithms described are important to how the HVAC simulation functions. These features play a role in how any primary system would be modeled and, thus, also are things to consider when developing alternative cooling models.

Pump Control for Plant and Condenser Loops. The pump is quite simply the component that drives the flow. How it reacts depends on several different conditions. In total, there are three different decision variables, two of which are defined by user input. These three deciding factors are whether the pump is constant or variable-speed, whether the pump operation is continuous or intermittent, and whether or not there is a load on the loop. The pump is simulated first on the supply-side loop after the demand-side loop has determined what the demand on the loop will be. The load is simply calculated by multiplying the requested flow rate from the demand side by the difference between the enthalpy at the supply-side inlet and the enthalpy that corresponds to the current loop setpoint temperature. This setpoint temperature is the fluid temperature that one is attempting to maintain at the outlet of the supply side and can be scheduled to different values on an hourly basis.

The operation of a constant speed pump is fairly straightforward. If the user designates a constant speed pump that is operating continuously, the pump will run regardless of whether or not there is a load. This may have the net effect of adding heat to the loop if no equipment is turned on. If the pump is constant speed and operates intermittently, the pump will run at its capacity if a load is sensed and will shut off if there is no load on the loop.

A variable speed pump is defined with maximum and minimum flow rates that are the physical limits of the device. If there is no load on the loop and the pump is operating intermittently, then the pump can shut down. For any other condition, such as the loop having a load and the pump operating intermittently or the pump continuously operating (regardless of the loading condition), the pump will operate and select a flow somewhere between the minimum and maximum limits. In these cases where the pump is running, it will try to meet the flow request made by demand-side components.

In many cases, the first estimate of flow requested by the demand side tends to be fairly accurate and the flow rate does not vary in subsequent iterations. However, because there is the possibility that the coils or some other component might request more flow in future iterations during the same time step, the program must not only set flow rates but also maintain a record of the current maximum and minimum flow rate limits. This information is important, not only to the pump itself, but also to other pieces of equipment that may control their flow rates and, thus, require knowledge of the limits within which they may work. In general, the decisions on what to set the maximum and minimum flow rates is directly related to the type of pump (constant or variable speed). For constant speed pumps, the maximum and minimum flow rate values are the same and, thus, if the flow requested does not match this, the other components must either deal with the flow or a bypass branch must be available to handle the excess flow. For variable speed pumps, the maximum and minimum flow rates are set by the user-defined limits.

Plant/Condenser Supply Side. Component models, such as boilers, chillers, condensers, and cooling towers, are simulated on the supply-side of the plant and condenser loops. In order to allow specification of realistic configurations, the plant and condenser supply-side loop managers were designed to support parallel-serial connection of component models on the loop. In addition, loop managers were designed to support both semideterministic models (e.g., the parameter estimation models of the ASHRAE Primary Toolkit [Pedersen et al. 2001]) and "demand-based" models (e.g., the performance map models of BLAST and DOE2.1E). As a result, the loop manager must be able to simulate models that require the mass flow rate as an input and models that calculate the mass flow rate as an output—sometimes in the context of a single loop configuration.

In order to achieve these design criteria without resorting to a pressure-based flow network solver in the HVAC portion of the code, a rules-based "flow resolver" was developed for the EnergyPlus plant and condenser supply-side managers. The flow resolver is based on the following assumptions and limitations:

- Each loop is only allowed to have a single splitter and a single mixer.
- Due to the fact that there can only be one splitter and one mixer on a given loop, it follows logically that there can be, at most, one bypass on each loop.
- No other components may be in series with a bypass, i.e., a branch that contains a bypass may have no other equipment on that branch.
- Equipment may be in parallel only between the splitter and mixer component of a loop or between one of those types of equipment and the loop inlet/outlet nodes.
- Equipment may be hooked together in series in each branch of the loop.
- Flow rates on individual branches will be controlled using maximum and minimum available flow rate limits.

The flow resolver employs a simple predictor-corrector algorithm to enforce mass continuity across the plant loop splitter, as shown in Figure 3.

As previously discussed, the pump establishes the total loop mass flow rate by setting the flow in the first supply-side branch. In the second step, a predictor algorithm polls each piece of equipment on the loop and "predicts" branch mass flow rates based on the requested flow rate for each. The loop manager calls the appropriate module to simulate (in order) all of the components on each branch of the loop except for splitters and mixers. In this step, each component sets the conditions at its outlet node, including temperature, flow rate, maximum allowed (design) flow rate, minimum allowed (design) flow rate, maximum available flow rate, and minimum available flow rate. These predicted values are based purely on the component's own control scheme and, thus, each component is free to request as much (or as little) flow as desired.

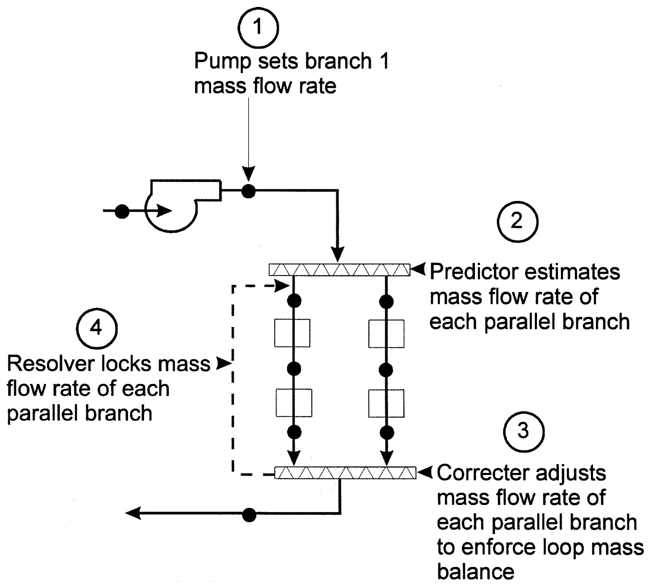

Figure 3 *Plant/condenser supply-side solution scheme.*

Each component is tagged in the user input file as an ACTIVE, PASSIVE, or BYPASS type of model. An ACTIVE type describes a demand-based plant model that calculates mass flow rate as an output. A PASSIVE type describes a semideterministic model that is simulated with the mass flow rate as an input. The BYPASS type designates a loop bypass.

The predictor algorithm first establishes the desired flow rate of each branch by searching for ACTIVE components on the branch. The first ACTIVE component in simulation order sets the desired branch flow. Branches with only PASSIVE components require a flow rate between the minimum and maximum allowable branch flow. Branches with a BYPASS component have a branch flow only when all other branches combined cannot handle the entire loop flow.

In the third step, the loop manager makes any necessary "corrections" to the requested branch flows in order to enforce overall continuity on the loop. If mass conservation allows all ACTIVE branches to be satisfied, then the remaining flow is divided between the PASSIVE branches and, as a last resort, the BYPASS. If there is insufficient flow to meet the branch demand, ACTIVE branch requests are met first in the order that the branches appear in the branch list in the input file.

Plant/Condenser Demand Side. The plant and condenser demand side are simulated in a different manner than the supply sides because, in reality, there are no components to simulate or control. On the supply sides, there is a load management scheme and other constraints that must be resolved. On the demand sides, all of the components have already been simulated and controlled by the air loop, the zone equipment, or the plant supply side. Thus, the demand-side management module only needs to resolve the actual flow rate through each section or branch of the subloop and also monitor the maximum and minimum flow rates that are available.

The flow rate is resolved first for each individual branch. For every branch, the program cycles through each node on the branch and determines what the flow requests and flow limits are. The most restrictive flow constraints are assumed to be valid for the entire branch regardless of component type. Since there may be several components in series on a particular branch, there is also a defined scheme for assigning priority to components that will have the ability to control the flow. The user may specify individual components as either active or passive. Active components are given highest priority for requesting a particular flow rate. If there is more than one active component on a particular branch, then it is assumed that the first active component on the branch is the highest priority and dictates the flow request.

Once all of the branches have set their flow rates and constraints, the splitter and mixer must resolve the various flow requests. In the demand-side scheme, the mixer and any branch following the mixer is completely passive. Thus, all of the control happens at the splitter. The splitter first attempts to sum the maximum and minimum constraints from all of the branches coming out of the device and compare those to the constraints are that valid for the branch leading into the splitter. When there is a mismatch between the outlet constraints and the inlet constraints, the simulation will defer to the inlet constraints due to the fact that the pump is, in reality, controlling flow on the loop. Since the constraints of the pump would be passed across to the demand side from the supply side, an assumption is made that the coils or other demand-side components must live within the bounds of the pump.

Once the flow has been resolved at the splitter, the branch flow rates and constraints between the splitter and mixer can be adjusted, if necessary. In some cases, this will be mandatory to maintain a mass balance at the splitter. When the flow rate coming out of the splitter does not match the branch requests, individual branch flow rates must be adjusted to provide for the extra flow or the "flow deficit." When there is extra flow, flow is sent through any bypass branch first and is then sent to passive branches in reverse order of their appearance in the splitter outlet list. When all of these branches have been exhausted, flow will be increased to the active branches, also in reverse order. The reverse order guarantees that the branch appearing first has the highest priority to receive the flow rate it has requested. If there is not enough flow for all of the requests, flow rates will be decreased in a similar order: passive branches first in reverse order, followed by active branches in reverse order. Flow rates are increased or decreased until a mass balance at the splitter exists.

It is also necessary to monitor the flow constraints at the branches and components since once the flow rates are changed, the components must be resimulated by the controlling loop (air loop, zone equipment, or plant supply side). The controllers for these components must know if the constraints have been modified so that the simulation does not toggle between a component requesting a flow that the pump cannot meet and the pump then resetting the flow to what it can provide. Note that once a flow rate for any component has changed, this signals the need to resimulate any subloop to which it might have an indirect connection. Currently, this means that if a flow rate on the plant demand side changes, the simulation must recalculate the conditions on both the air loop and zone equipment subloops since coils and other equipment could be on either side of the main air loop. Similarly, if the condenser demand-side simulation results in a change in flow rate through a chiller condenser, then the plant supply side must be triggered to perform its calculations again. Care has been taken to avoid cases where the various subloops might simply keep triggering the resimulation of their indirect connections in an infinite loop.

Temperature Resolution. The transition from load or energy-based plant models to a loop-based arrangement makes variables of both the flow rate and the fluid temperature. This means there are more degrees of freedom that must be controlled. The flow resolver concept discussed previously controls the flow rates through the components and maintains an overall mass flow balance through the loop. However, the temperatures still need to be controlled. A purely iterative procedure can be expected to converge to the appropriate loop temperatures, but the procedure can become slow to converge under conditions where the demand changes rapidly. This situation is somewhat analogous to that existing in the link between the zone and the air system. In that case, the convergence and stability of the iterative solution were greatly improved by adding the thermal capacitance of the zone air and other fast-responding mass within the zone. Based on that experience, it was decided to add thermal capacitance to the plant loop and benefit from the added stability. Because the thermal capacitance in the zone/system interaction is relatively small, it was necessary to use a third order numerical solution there. Since the plant loop thermal capacitance is higher, a simple first order solution has been found to be satisfactory.

To implement the capacitance, each loop is assigned a fluid volume as user input. This is used to determine a capacitance concentrated in the supply-side outlet node. If the loop setpoint cannot be maintained, this node becomes an energy storage location and its temperature reflects the current capability of the supply side. The size of the thermal capacitance affects the speed of recovery from situations where the setpoint was not maintained. The user must estimate a fluid volume based on the size of the pipes in the loop. However, rough estimates seem to be sufficient. The supply-side outlet node temperature and the demand-side inlet temperature proceed through smooth paths from one time step to the next. No energy is lost or gained because of storage in the loop capacitance. Once setpoint temperature is reached, the storage effects are not involved.

INTEGRATION OF ALTERNATIVE COOLING SYSTEMS

In the U.S., when cooling is required to maintain thermally comfortable conditions within a building, architects and engineers turn almost exclusively to forced air systems. By providing cold air to a space, the system is able to extract excess heat provided by internal gains, solar or transmission gains, etc., to maintain thermal equilibrium at a desired setpoint temperature. This is the preferred technique for a variety of reasons, including a vast amount of experience with such systems, the ability to model and predict the effect and energy costs of these systems, and a current lack of interest from various groups in providing other alternatives. In this way, the HVAC community in this country lags somewhat the industry within other parts of the world, such as Europe. In Europe, where energy concerns seem to play a greater role, alternative cooling strategies such as radiant cooling and slab cooling have been successfully implemented with significant energy and cost savings.

Some of the arguments against the implementation of such energy-saving cooling strategies in the U.S. have centered on the important issue of climate. Europe tends to have a much cooler and drier environment in general than the U.S. The lower sensible loads encountered in Europe mean that much larger percentages of the cooling load could be met by the alternative strategy, making it much easier to justify any added financial investment. Also, the lower humidity levels in Europe mean lower latent loads from ventilation and decreased likelihood of condensation in a radiant cooling system.

The purpose of this next section is to provide a brief overview of existing alternative cooling models in the new program, such as the low-temperature radiant system, and to discuss the possibility of expanding the capabilities of the radiant system model in the future. The added alternative systems that could be simulated in the future will, in large part, take advantage of the existing model and the modularity and flexibility that the new program offers. This will improve the success of the integration of the alternative cooling models as well as significantly decrease the development time.

Low-Temperature Radiant Cooling System Overview

Low-temperature radiant cooling systems appear, on the surface, to be relatively simple systems. The system circulates cold fluid through tubes embedded in a wall, ceiling, floor, or panel. Energy is, thus, removed from the space, and zone occupants are conditioned by both radiation exchange with the system and convection from the surrounding air.

Despite the relative simplicity of the low-temperature radiant systems, the integration of such a system within an energy analysis program requires one to overcome several challenges. First, for systems with significant thermal mass, the conduction transfer function method for modeling transient conduction must be extended to include embedded heat

TABLE 1
Major Features of the Low-Temperature Radiant Cooling System Model

Features	Notes
Rigorous Model Foundation	• Integrated with a heat balance approach • Takes advantage of existing and tested loads calculation • Improvements to heat balance are immediately available to radiant system model
Transient Conduction	• Accounted for using a modification of standard conduction transfer functions • System can be defined with any material or insulation level
Controls	• Model varies flow rate to meet cooling loads • User specifies maximum water flow rate (can be different for heating and cooling) or maximum electric power of system • User specifies setpoint temperature for system (can vary on an hourly basis) • User specifies throttling range for controls • User specifies temperature to which the setpoint is compared (can be MAT, MRT, or operative temperature) • User specifies water temperatures (can vary on an hourly basis)
Simulation Flexibility	• Model can adjust time steps to accurately account for rapidly changing conditions • Zone time step integration of heat source/sink seeks to guarantee that energy is not created or lost • Nearly all of the limits (such as number of zones or number of surfaces) that are found in programs of this type have been eliminated

sources or sinks. Strand (1994, 1995) showed that this was possible and that the low-temperature radiant system could be handled like any other surface within the heat balance framework.

Once the transient nature of the system is accounted for, one must then turn to the next difficult issue: controls. Controls are problematic for almost any simulation program. The problem is not whether something can be simulated because typically a simulation program offers the ability to experiment with many different control strategies. Rather, the problem is typically the diversity of controls that are implemented and keeping the controls that can be simulated up to date. In this area, the new program should be seen as a first attempt at modeling basic low-temperature radiant systems and not as the definition of all radiant systems. Plans call for the addition of other control strategies in future versions of the program.

As a result, controls for low-temperature radiant systems within the new program are fairly simple, though there is some flexibility through the use of schedules. The program user is allowed to define a setpoint temperature as well as a throttling range through which the system varies the flow rate of water (or current) to the system from zero to the user-defined maximum flow rate. The flow rate is varied linearly with the flow, reaching 50% of the maximum when the controlling temperature reaches the setpoint temperature. Setpoint temperatures can be varied on an hourly basis throughout the year if desired. The controlling temperature can be the mean air temperature, the mean radiant temperature, or the operative temperature of the zone, and this choice is also left to the user's discretion. Since flow rate is varied, there is neither explicit control on the inlet water temperature nor mixing to achieve some inlet water temperature in a hydronic system. However, the user does have the ability to specify on an hourly basis through a schedule the temperature of the water that would be supplied to the radiant system.

One remaining challenge is the merging of the low-temperature radiant system model with an integrated building simulation program. In the past, most simulation programs have simulated the building envelope, the space conditioning systems, and the central plant equipment in three separate steps. While this had some advantages and was partly due to a lack of computing capacity, the large drawback for this arrangement is that there is no feedback from the space conditioning system or central plant to the building conditions. Thus, if the system or plant was undersized, it was reported as an "unmet load" and did not affect the temperatures experienced within the building. A predecessor (Taylor et al. 1991) to the new program resolved this issue by integrating all three major components of a building simulation and, thus, allowing feedback between the equipment and the building envelope. The radiant system model conforms with this solution technique.

Further details on the low-temperature radiant cooling model that was part of the initial release of the new program can be found in the literature (Strand and Pedersen 2002). Major features of the model as implemented in the new program are summarized in Table 1.

Potential Expansion of Modeling Options

There are nearly limitless different schemes for providing alternative cooling for a building. Many of these rely heavily on climate or at least have only been tried in climates that are thought to be conducive to such alternate schemes. This paper focuses in part on combining a radiant slab with alternate ways of obtaining chilled water for circulation through the system. There are several possibilities for improving the current model. One improvement would be to utilize “free cooling” from some form of condenser, such as a ground loop or a cooling tower. Another improvement would be to combine the radiant system with an air loop to avoid the prospective problems of condensation on the radiant slab. Each of these technologies and the potential ways they could be integrated with the current radiant system model within the program are discussed below.

One study has already been done linking a radiant slab to a cooling tower with the new program. This was, in reality, a “virtual link,” in that the coupling was made possible by input parameters. While the results of the study (Strand 2001) are in need of some validation, the fact that the study could be done without changes to the program is a testament to the modularity and flexibility of the new program. In the study of linking a radiant slab to a cooling tower, the link was approximated by scheduling the plant loop temperature on an hourly basis. This can be done in input data and does not require modification of the source code. The loop temperature was scheduled to equal a constant 3°C above the outdoor wet-bulb temperature. This was assumed to approximate the performance of a cooling tower. The study showed that by running the system at various times during the night, the total cooling load on a residential building could be cut in half for a fairly massive slab, even for a fairly humid climate.

While this study showed what was possible, it was obvious that modification of input to mimic a slab connected to a cooling tower was not the same as having a model that integrated these two features. Cooling loop setpoint temperatures had to be manipulated by hand and then entered as a schedule into the input file. In addition, the determination of when to run the water through the slab had to be determined manually and then also scheduled for “optimal” performance. Clearly, such exercises are not tolerable for long time frame simulations or by frequent users of energy simulation programs. In addition, a user would require flexibility in the type of condenser to which the system would connect.

Integrating a radiant slab with a condenser will require two issues: connections and controls. The most convenient form of the enhanced model would allow the user to connect the water side of the radiant system to the condenser demand side rather than the plant demand side. This would allow a direct connection between the system and the condenser. In order to accomplish this, the program would need to be aware of the possibility of connecting a condenser demand-side subloop directly to the zone equipment subloop (the “future link” shown in Figure 1). An alternate means of approximating this direct connection that would avoid the addition of an indirect link would be the creation of a special “perfect heat exchanger” equipment type for use in the plant supply side. This could grab the conditions from the condenser demand side and invert them over on the plant supply side. Unfortunately, this does not seem to be a viable alternative because there is a clear disconnect between what is happening in the condenser equipment and the radiant system, and this would not fit in with the load management scheme used in the plant supply sides.

Even with the addition of an indirect link between the condenser demand side and the zone equipment subloops, the issue of controls and when the system should be run is not completely solved. Currently, the radiant system does not look at the temperature of the fluid in the loop at all. It assumes that this temperature is adequate to produce a cooling effect on the system (or a heating effect in heating season). There are no checks made to decide whether the system should be shutdown because the fluid temperature will have the opposite of the intended effect on the space.

The need for additional controls to shut the system off based on the fluid temperature coming back from the cooling tower, ground loop, or other condenser type is not insurmountable. The current control algorithm for radiant systems is contained within a single subroutine and could be amplified to include a fluid-side temperature check. In fact, this would probably be a reasonable safety check that should be added regardless of how the system is linked to a fluid loop.

Ventilated slabs where cooler nighttime air is circulated through the core of a slab floor/ceiling cannot currently be modeled in the new program. In fact, because the fluid is air and the air is connected to the outside environment rather than a fluid loop, the modeling of such systems will require the creation of a separate model. These systems are actually easier to simulate from the standpoint that they do not require any supplemental equipment connections except perhaps the specification of a fan that would be used to circulate air through the slab core. The control of these systems would again require modified control algorithms that must compare the temperature of the outside air with the temperature within the slab. However, there is overlap between the control of a ventilated slab and a radiant system connected directly to a condenser, and lessons learned from the algorithm of one system will most likely benefit the other. Of the two alternative systems just discussed, it appears that the ventilated slab will require less time to implement.

One other alternative system type that cannot yet be modeled within the new program is what some segments of the engineering community term a “hybrid system.” This system is the linkage of an air-handling system with a radiant system and is intended to avoid condensation problems that might result when a radiant cooling system is specified for a humid climate. In this case, the surface temperature of the radiant system may drop below the dew-point temperature of the air in the zone, resulting in condensation and, in some cases, “indoor rain.” To avoid this, an air loop is specified to provide dehumidification, and the radiant system attempts to meet most of the sensible load. Some of these systems will run the cold fluid through the cooling coil of the air loop first and then circulate it through the radiant system. This most likely would prevent condensation.

The new program is well positioned to model such a system. One reason is that the plant and condenser loops are both specifically set up to handle such series organization of components. In addition, the new program allows multiple systems to serve a single zone so that an air system and a radiant system can already serve a particular zone. However, the sequencing of these systems is a scheduling and priority scheme that is based on availability, load, and temperature—not on humidity. A hybrid system needs to be able to sense the humidity level within a space and determine whether it will need to run the air loop to dehumidify the zone air and avoid condensation. Since the radiant system currently does not control on humidity and the sequencing of system operation does not depend on humidity, it seems highly likely that such a hybrid system will require a separate system model that links an air loop and a radiant system together. The disadvantage with defining this as a separate system is that it might limit the types of systems that can be linked together with a radiant system to form a hybrid system. This will certainly require further investigation before the development of a model begins, but the prospects appear favorable for integrating a hybrid system model as well as ventilated slabs and condensers linked with radiant systems in the new program in the near future.

CONCLUDING REMARKS

This paper has described the portions of a new energy simulation program (EnergyPlus) related to modeling HVAC systems and the outlook for using this program to model alternative cooling systems. The more apparent advantages of the new program over the programs from which it descended are its flexibility and modularity in the HVAC section in particular and, in general, the entire code and in the integration of the solution of the three main elements of the building description: the thermal loads on the envelope and the primary and secondary systems. While the program does not have a model for every available technology, the program was intentionally designed to allow a large number of model developers to gain access to and make relatively fast changes to the simulation code. This will reap significant benefits in the long term as more developers are able to contribute and as more users begin relying on the program as their simulation tool.

The program has already shown that its flexibility can already allow the simulation of some situations for which there are not yet official models. This was seen by using the options within an input file to mimic a radiant cooling system that is connected to a cooling tower. This successful test that did not require changes to the program itself is both an encouraging demonstration of the program's flexibility as well as a first step in defining models that will handle such cases in a more sophisticated fashion. Potential future models that were discussed included the radiant system linked to a cooling tower or some other condenser type, such as a ground loop, a ventilated slab, and also a hybrid system that combines an air loop and a radiant system. The outlook for integration of such models appears favorable, in part due to the modularity of the new code.

Clearly, these are not the only examples of new models that could be integrated with the new program. Other researchers will have other ideas that they wish to investigate. The strength of the new code is that this is more possible than it ever was before. The flexibility of the program allows users to model in a limited way potential new techniques with changes only to the input file, while the modularity of the code will speed the eventual integration of new models within the program. It is hoped that this will benefit both the development of new technology and its implementation as the nation and the world continue to strive for more energy efficient buildings.

ACKNOWLEDGMENTS

The authors gratefully acknowledge the leadership of Drury Crawley and the support of the United States Department of Energy through various contracts for the research described in this paper. The authors also wish to note the significant contributions of Linda Lawrie of the United States Army Corps of Engineers Construction Engineering Research Laboratory, Fred Buhl of Lawrence Berkeley National Laboratory, and Don Shirey of the Florida Solar Energy Center to the HVAC and supporting portions of the new program.

REFERENCES

Buhl, W.F., A.E. Erdem, F.C. Winkelmann, and E.F. Sowell. 1993. Recent improvements in SPARK: Strong-component decomposition, multivalued objects and graphical editors. In *Proceedings of Building Simulation '93*, pp. 283-289, August 1993, Adelaide, South Australia Australia, IBPSA.

Building Systems Laboratory. 1997. *BLAST user's manual*. University of Illinois at Urbana-Champaign.

Clark, D.R., and W.B. May, Jr. 1985. HVACSIM+ building systems and equipment simulation program. National Institute of Standards and Technology, Washington, D.C.

Crawley, D.B., L.K. Lawrie, F.C. Winkelmann, W.F. Buhl, A.E. Erdem, C.O. Pedersen, R.J. Liesen, and D.E. Fisher. 1997. The next-generation in building energy simulation—A glimpse of the future. In *Proceedings of Building Simulation '97*, Volume II, pp. 395-402, September 1997, Prague, Czech Republic, IBPSA.

Fisher, D.E., and R.D. Taylor. 1999. A modular, loop-based approach to HVAC energy simulation and its implementation in EnergyPlus. *Conference Proceedings of Building Simulation '99, Kyoto, Japan*. International Building Performance Simulation Association (IBPSA).

Pedersen, C.O., R.J. Liesen, R.K. Strand, D.E. Fisher, L. Dong, and P.G. Ellis. 2001. Loads toolkit RP-987. Atlanta: American Society of Heating, Refrigerating and Air-Conditioning Engineers, Inc.

Solar Energy Laboratory. 1990. TRNSYS—A transient simulation program. University of Wisconsin-Madison.

Strand, R.K., and C.O. Pedersen. 1994. Analytical verification of heat source transfer functions. First Joint Conference of International Simulation Societies, Zurich, Switzerland.

Strand, R.K. 1995. Heat source transfer functions and their application to low temperature radiant heating systems. Ph.D. dissertation, University of Illinois at Urbana-Champaign, Department of Mechanical and Industrial Engineering.

Strand, R.K., and C.O. Pedersen. 1997. Implementation of a radiant heating and cooling model into an integrated building energy analysis program. *ASHRAE Transactions* 103(1).

Strand, R.K. 2001. Modeling a radiant slab coupled to a cooling tower using the new national energy analysis program. Proceedings of the Cooling Frontiers Symposium, Arizona State University, October 4-7, 2001.

Strand, R.K., and C.O. Pedersen. 2002. Modeling radiant systems in an integrated heat balance based energy simulation program. *ASHRAE Transactions* 108(2).

Taylor, R.D., C.O. Pedersen, D. Fisher, R. Liesen, and L. Lawrie. 1991. Impact of simultaneous simulation of building and mechanical systems in heat balance based energy analysis programs on system response and control. Conference *Proceedings of Building Simulation '91, Nice, France*. International Building Performance Simulation Association (IBPSA).

Winkelmann, F.C., B.E. Birdsall, W.F. Buhl, K.L. Ellington, A.E. Erdem, J.J. Hirsch, and S. Gates. 1993. DOE–2 Supplement, Version 2.1E, LBL-34947, November 1993, Lawrence Berkeley National Laboratory. Springfield, Va.: National Technical Information Service.

DISCUSSION

Edward F. Sowell, Calif. State Univ. Fullerton: I notice that thermal capacitance is included the plant loop energy balance.

As noted in the paper, this capacitance was not in the original formulation of the model, presumably because the engineering judgment was that the dynamic effects of it are unimportant with regard to the intended purpose of the model. This technique, i.e., introducing false (or, at least physically insignificant) dynamics purely to resolve numerical stability problems, has a long history dating back to the times of analog computation. However, it is of dubious merit. First, if the dynamic phenomena are indeed unimportant, finding accurate data needed to represent them is, at best, unnecessary aggravation for the user. Moreover, if care is not taken in setting realistic capacitance values the false dynamics introduced into the results will be confusing or, in the extreme, misleading. On the other hand, if the correct capacitance is indeed provided the dynamic equations may be so fast that the solution time has to be reduced below that which would otherwise be required, leading back to where the exercise began, i.e., numerical difficulties.

My question is therefore, why do you not use well known methods, e.g., Newton-Raphson, to solve the algebraic part of the model?

Although not discussed in the paper itself, another issue was raised during presentation of the paper. The slides suggested that the dynamic terms were introduced in the loop energy balance equation itself. It would seem that a modular simulation software system should incorporate something like this as a separate module. For one thing, this approach would make the implementation of the actual physical model of the plant loop more transparent, i.e., an energy and mass balance. Also, if done in this manner the dynamics could be optionally omitted.

When these questions were posed at the session there were basically two responses.

For the benefit of the Transactions reader I would like to paraphrase these responses and my counters. The first response was that "the time step does not really get shortened because the basic architecture of the program sets the time step in a top down manner." Without knowing exactly what that means, I would venture to say that this policy might not always work, since there is nothing in the physics that ensures phenomena at a lower level cannot have faster response times than those at higher levels. In other words, without using some kind of nonphysical heuristic in the code, I do not see how you can allow arbitrary setting of the plant loop capacitance and simultaneously ensure that a time step set at a "higher level" will be short enough for proper handling of the plant loop dynamics. The second comment was basically that since "all systems are really dynamic" we are justified in introducing dynamics where ever necessary to achieve dynamic stability.

While I agree that nothing happens instantaneously, I would not want to be so constrained in modeling as to use only differential equations. In my view, a well defined simulation model aims to answer a specific range of performance questions. Typically, these questions will have a dynamic threshold, so to speak, in that events and processes with faster response can meaningfully be treated as instantaneous. If one insists on modeling such events and processes with differential equations there will be significant increase in runtime with no corresponding benefit to the user. Thus introducing dynamics below this threshold merely to avoid solution of algebraic equation systems seems hard to justify.

Rick Strand: In EnergyPlus, the dependent state variables in the algebraic part of the model are fluid temperature and fluid mass flow rate. The mass flow rates are resolved without a pressure based flow simulation by means of a “rule-based” flow resolver. The “rule based” flow resolver requires a successive substitution scheme, but it also enforces stability by excluding unstable configurations. Thermal mass was not added to allow simulation of unstable configurations, it was added to improve the computation time of stable configurations.

In order to use a Newton method to resolve the algebraic part of the model, either pressure must be included as a state variable or the mass flow rate must be specified for each branch (or mass flows must be specified-the trivial solution). A Newton solver would require a full pressure based flow simulation. The overhead in terms of both input and computation time for detailed piping configurations (hundreds of branches and dozens of components) would be prohibitively large, and balancing and controlling the flow would prove as difficult and expensive in the simulation as it is in practice. With the "rule based" flow resolver, relatively complex systems can be specified quickly and reliably.

There is, however, a place for detailed simulation of HVAC systems, and EnergyPlus was in fact designed to accommodate both "rule-based" (successive substitution) and pressure based (Newton) solvers. A predecessor of the EnergyPlus heat balance (IBLAST) demonstrated the feasibility of this concept by supporting a seamless link to HVACSim+.