# COMPUTER SOFTWARE ARCHITECTURE TO SUPPORT AUTOMATED BUILDING DIAGNOSTICS

Whole building diagnostician software architecture

D. P. CHASSIN
Pacific Northwest National Laboratory[†], Richland, Washington USA

Durability of Building Materials and Components 8. (1999) *Edited by M.A. Lacasse and D.J. Vanier*. Institute for Research in Construction, Ottawa ON, K1A 0R6, Canada, pp. 2665-2676.


## Abstract

Developed by Pacific Northwest National Laboratory's Building Sciences Group, Honeywell Technology Center, and the University of Colorado Joint Center for Energy Management with support from the U.S. Department of Energy, the Whole Building Diagnostician is a Windows-based application that provides building operators with easy access to system diagnostic information. The architecture of the software infrastructure presented provides essential data collection, validation, integration, analysis, and management functions for the large sets of discontinuous asynchronous time-series data used by all the modules in the application. The proposed architecture uses a central database to store both the data and the diagnostic results from the various modules. Although the use of a centralized database has many advantages, it has several shortcomings. This paper will discuss the advantages and shortcomings of such an approach when deployed on a large scale.

Keywords: building diagnostics, computer software, ventilation diagnostics

## 1 Introduction

The Whole Building Diagnostician (WBD) is a system created to provide building operators access to information about the performance of the buildings for which they are responsible. The first operational version of the WBD provides failure diagnostics for air handlers, and performance tracking at the whole building and the aggregate plant level, (Brambley et al. 1998).

---

[†] Pacific Northwest National Laboratory is operated by Battelle Memorial Institute for the US Department of Energy under contract DE-AC06-76RLO 1830.

The WBD infrastructure provides the systems in support of the interaction between the various WBD modules. In general, it provides the following services as shown in Fig. 1:

- data transport and storage
- module execution and control
- installation and setup

## 2 System requirements, design, and implementation

### 2.1 Data collection

The WBD collects data from a variety of Building Automation Systems (BAS) using DDE, trend logs, or custom interfaces. (Some interfaces such as those for BACnet are planned but have not been implemented to date because of a lack of demand.) For each BAS, a driver must be designed and implemented. Data is then automatically collected from the BAS, processed and delivered to the diagnostician's data sources.

### 2.2 Data access

The user views and edits configurations, schedules (see Fig. 2), and views data and diagnoses during WBD session. Data pertaining to the following categories of information is delivered to all WBD modules from the data sources:

- building component hierarchy and connectivity,
- sensor channels, connectivity, and conditions (locations, errors, tolerances and biases),
- building, plant, air handler, zone, and BAS configurations,
- building, plant, air handler, zone, and data collection schedules,
- sensor readings
- diagnostic results and intermediate calculations
- energy rates and costs

Database transaction control and disaster recovery services are used to maintain database consistency in the event of a variety of system failures (e.g., power outage, hardware or software failure).

### 2.3 Database schema maintenance

The infrastructure provides database services to perform the following tasks:

- debug the implementation of all the diagnostic modules and the infrastructure itself,
- automatically upgrade the database schemas to the most recent design,
- audit every data table for consistency with other tables, and
- repair damage resulting from crashes, power outages, and hardware failures.

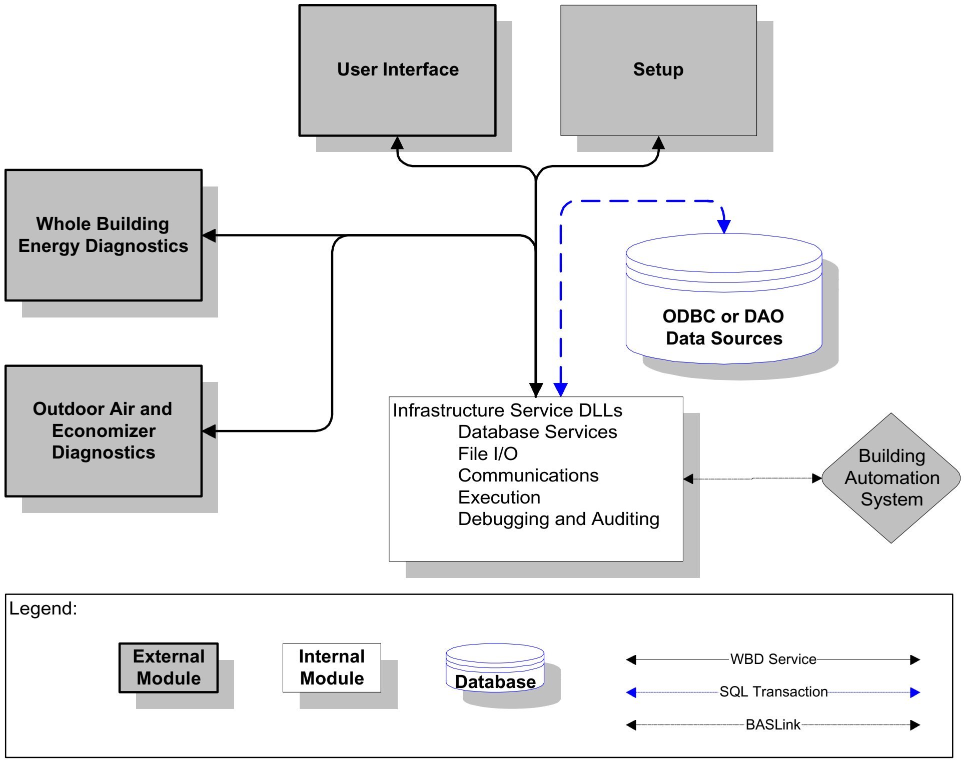


**Fig. 1: Whole building diagnostician system structure**

## 2.4 Process scheduling

The WBD controls execution of diagnostic modules and specifies the analysis desired by the user. Users may instruct each diagnostician individually to start either when the system boots up, when user logs in, when a WBD session is opened, or never start at all. In addition, users may individually instruct each diagnostician to shut down either when the system shuts down, when the user logs out, or when a WBD session ends.

The processing schedule is determined based on which building components the user elects to monitor. The schedule is then dispatched to the appropriate diagnosticians which will perform the required operations to ensure the user receives the desired results.

When a fatal processing exception is encountered during analysis, the subject analysis is abandoned while other analyses are completed. After all the other analyses have been given an opportunity to complete, the subject analysis is restarted in an attempt to overcome the problem. If the processing problem persist, it may require user intervention. The user is signaled in the appropriate WBD session by indicating an exclamation point beside the appropriate component's icon.

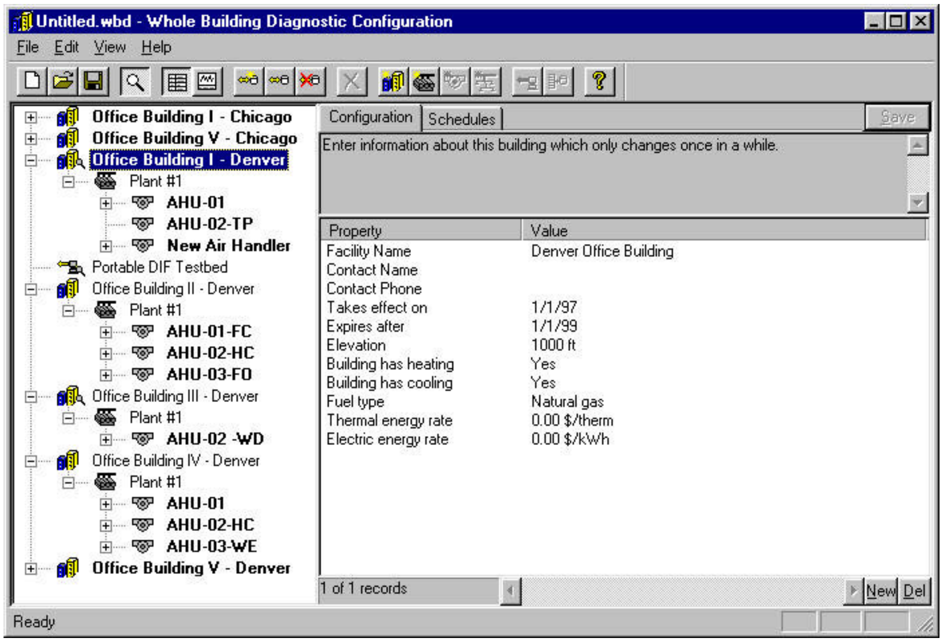


**Fig. 2: WBD building configuration screen**

## 2.5 Building automatic system links

The WBD has been successfully linked to several buildings having Johnson Control's (JCI) Metasys system and one having a CSI automation system. At the time of this writing, construction of links to a number of the other systems is in progress or anticipated, but results from these are currently unavailable.

The JCI linkage was established using the Microsoft Window's Dynamic Data Exchange (DDE) protocols support by the JCI Metalink module. The WBD implementation of DDE-based data collection allows the WBD system installer to link to most DDE-capable BASs, using an approach similar to that used when linking from Microsoft Excel, and without implementation of custom software.

The CSI linkage is done using the WBD's trend-log loader. Many legacy BASs do not support data interchange protocols and can only output trend logs or database files. The WBD is capable of reading Comma Separated Value (CSV) files, Data Interchange Format (DIF) files, and databases supporting Microsoft ODBC. These control systems typically output a summary of short-period data and the WBD collects and integrates that data whenever necessary.

The implementation of BACnet support in the WBD is not currently available but it is anticipated within the foreseeable future.

## 2.6 Time integration services

The WBD supports two types of data: hourly and irregular data. Hourly data is stored as an integrated value at the end of the hour (i.e., data stored at 2pm represents an integration of values from 1pm to 2pm). Irregular data is stored at the time the value becomes valid and remains in force until another value is obtained and stored.

Hourly data must be integrated by the data collection modules. Since most of the WBD's diagnostic analysis is performed on hourly data this presents a convenient efficiency for data storage and processing. In the cases where daily summaries are required, hourly values are integrated.

However, in some cases values changing in a sub-hourly time scale or at irregular intervals must be available to diagnostic module. In such cases integration is performed on arbitrary period data. The integration method and interval may be specified at run time. Summary information supported and statistics delivered are:

- integrated, averaged, accumulated, minimum, or maximum values
- count of samples used
- first and last time stamps of samples used
- root mean squared error and mean bias error (for the indicated degrees of freedom)

## 2.7 Units, errors and reliability, and psychrometrics

A number of statistical and thermodynamic utilities are provided to simplify implementation and ensure consistent usage across the diversity modules. These include unit conversion services, error and reliability propagation calculations, and psychrometric routines.

Unit conversion is based on universal constants. Any physical quantity is expressed as a function of six constants. They are:

| | |
|---|---|
| c | speed of light (2.99792.5 x $10^8$ m/s) |
| e | charge of an electron (1.602189246 x $10^{-19}$ Coulombs) |
| h | Planck's constant (6.62617636 x $10^{-24}$ J/Hz) |
| k | Boltzman's constant (1.38066244 x $10^{-13}$ J/K) |
| m | rest mass of an electron (9.10953447 x $10^{-31}$ kg) |
| | (All found in CRC Handbook of Chemistry and Physics 1980-81) |
| s | value of 1 kilogram of gold (1.233270 $10^4$ 1990$/kg), (GIAGP 1997) |

After converting these quantities to SI, any physical quantity can be computed using the equation:

$$Q = c^{\alpha} e^{\beta} h^{\gamma} k^{\delta} m^{\varepsilon} s^{\phi} \rho + \sigma$$

where:

$\alpha$ is the exponent to c,
$\beta$ is the exponent to e,
$\gamma$ is the exponent to h,
$\delta$ is the exponent to k,
$\varepsilon$ is the exponent to m,
$\phi$ is the exponent to s,
$\rho$ is a scalar, and
$\sigma$ is a constant offset.

The scalar is used to adjust to specify non-SI units and to translate from one unit to another. The basic SI and currency units are calculated as follows.

1 meter = $4.121487 \times 10^1$ h / mc, with a precision of 7 significant digits,
1 kilogram = $1.09775094 \times 10^{30}$ m, with a precision of 10 significant digits,
1 second = $1.235591 \times 10^{10}$ h / $mc^2$, with a precision of 7 significant digits,
1 Ampere = $5.051397 \times 10^8$ $emc^2/h$, with a precision of 7 significant digits,
1 K = 1.686358 $mc^2/k$, with a precision of 7 significant digits,
1 candela = 1.447328/683.0 $m^2c^4/h$, with a precision of 7 significant digits, and
1 US$ = $1.097751 \times 10^{30}$ sm, in 1990 with a precision of 7 significant digits.

Mathematical computations are then greatly simplified since conversion between like units (i.e., the exponents to the constants are the same) is performed automatically by adjusting the scalar and constant offset and deriving new units by calculating new exponents from those of the old units.

Finally, all sensor readings are assigned limits, error boundaries, and reliability rates. If a value is outside the limits, it becomes invalid. Valid values have a lower and upper error boundary, and a reliability rate which indicates what percentage of the samples can be expected to have true values that fall between the lower and upper boundaries. In addition, all mathematical and boolean operations propagate the error boundaries and reliability rates. As a result, each diagnosis is given a reliability based on the reliability of the original sensor readings and the amount of arithmetic performed on those readings.

## 2.8 Text compression

All text is stored in the database using a tokenized compression algorithm. The system maintains a list of keywords used to compress commonly used phrases. As a result a message such as "The mixed air temperature is less than the return air temperature" is stored as "Tmix < Tret." This system is also used to maintain consistent terminology in all the many diagnostic messages, and to support rapid revisions to them. In addition, dynamic data retrieval is supported in order to allow the current values to be retrieved for a given message. Thus the message "ActualTmix is GetActualTret and ExpectedTmix is GetExpectedTmix." is translated to "The actual mixed air temperature is 68ºF and the expected mixed air temperature is 72ºF."

## 2.9 Outdoor air and economizer diagnostician

The Outdoor Air and Economizer (OAE) diagnostic module includes an instance of the system scheduler, the diagnostic processor, and the OAE user interface component. It provides state analysis and diagnostics for air handlers and economizers.

The scheduler reads the processing schedule dispatched by the user interface to the OAE module. Each building component that is to be monitored by the OAE is checked to see if any new data are available. As new data from the BAS are introduced in the database, the diagnostic processor handles these data.

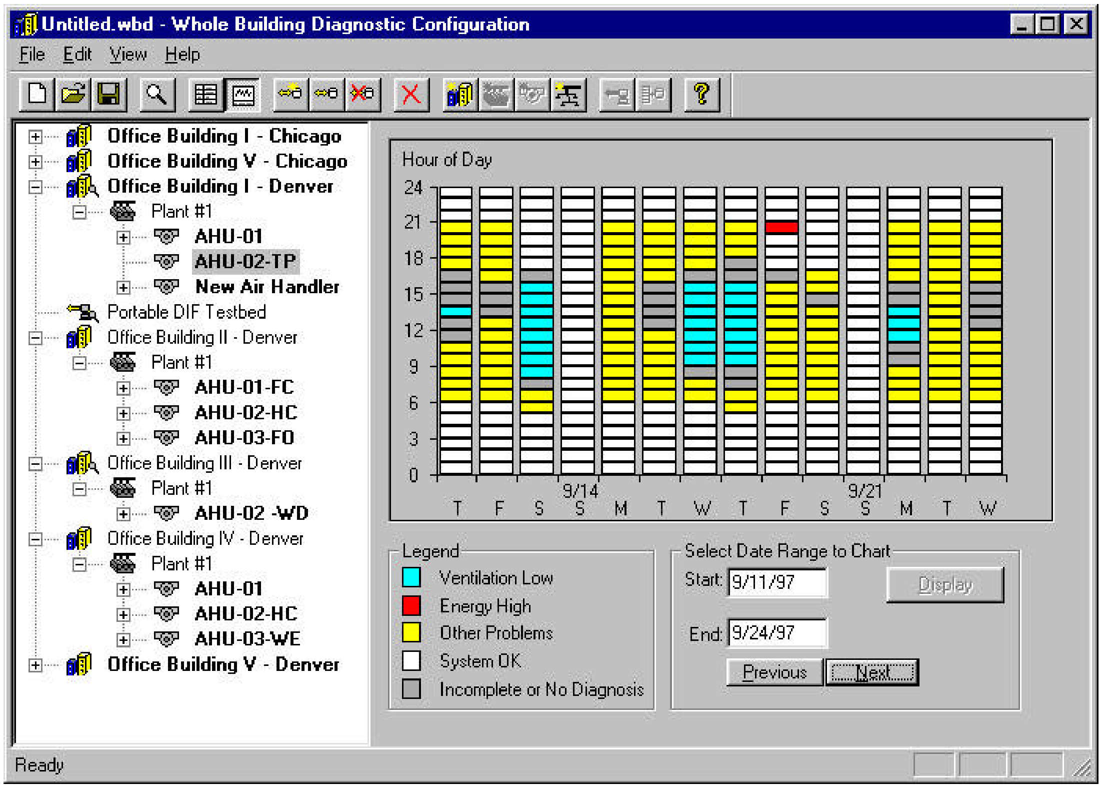


**Fig. 3: OAE state display**

All new data are processed through the OAE module's state detection decision tree. As possible states are excluded using individual tests and calculation, the module records any impossible failures that can be inferred at that point. When an air handler's state is eventually determined, the OAE module records that state and any possible failures that can be inferred from that state. The module then processes the impossible and possible failures accumulated over recent hours and days in order to determine the most likely failure(s) that are consistent with that majority of reliable observations.

The OAE user interface (see Fig. 3) displays the diagnosis concluded for each hour based on the data obtained from the BAS. Users can drill down on each hour to obtain details on the system state and the diagnosis (see Fig. 4).

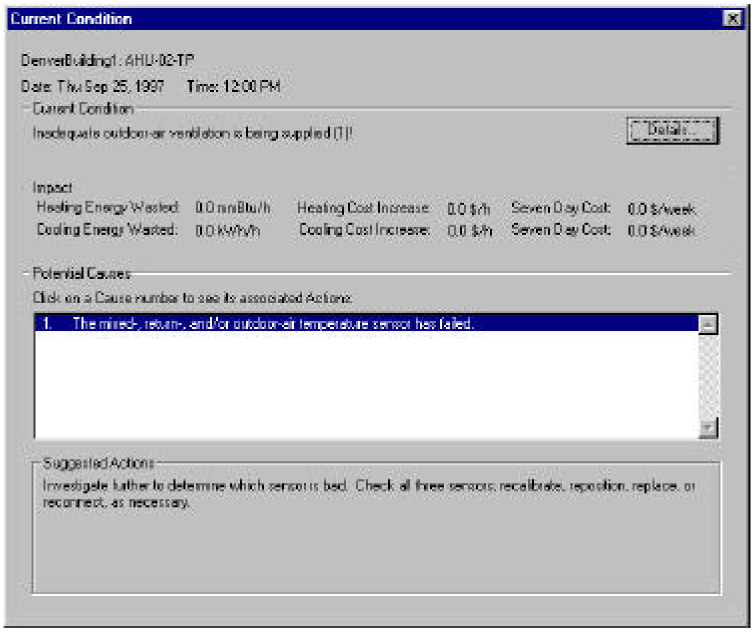


**Fig. 4: OAE state details and diagnosis**

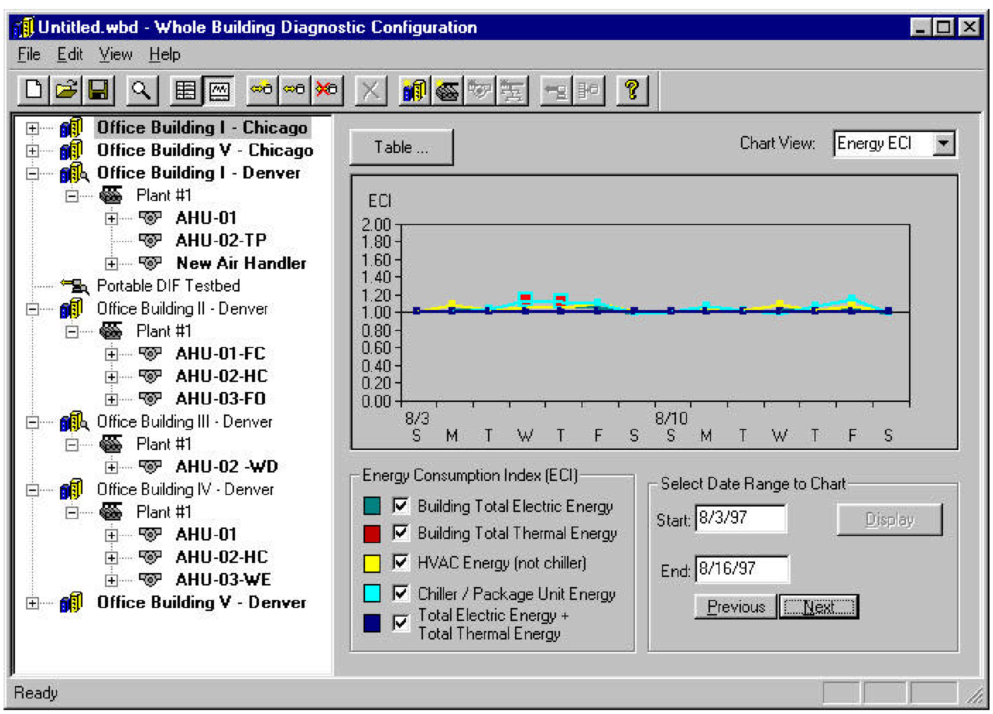


**Fig. 5: WBE tracking screen**

### 2.10 Whole building efficiency diagnostician

The Whole Building Efficiency (WBE) diagnostic module includes an instance of the same scheduler as the OAE module, but has its own efficiency tracking module and WBE user interface. The WBE provide Energy Consumption Index (ECI) calculations for whole building energy, whole building electric, whole building thermal, chiller and non-chiller energy usage. An ECI is the ratio of actual usage to expected usage. In addition the WBE is capable of detecting some power, gas, temperature and dew-point sensor failures.

The data processed through the WBE are used to calculate daily actual and expected ECIs. Any ECI with a statistically significant difference between the actual and expected value is flagged as a potential problem (see Fig. 5). The user can drill down on these days to obtain details on the severity of the problem and the anticipated cost impact.

### 2.11 Common user interface controls

The user interface module uses two common controls: the building component tree and the component configuration dialog. Both these controls are linked to the data sources and automatically maintain their synchronization to the data sources. In addition, the building component tree maintains the processing schedule information by displaying icons on the components to indicate active monitoring (the magnifying glass), a problem (an exclamation point--not shown), and the bolded component name (data and/or diagnosis available for viewing).

## 3 System performance evaluation

Database size is an important consideration when installing the WBD. In particular, large facilities or campuses must ensure that adequate storage is provided. The WBD collects and processes extremely large amounts of data. As the reader can observe in Table 1, building components with problems can be particularly burdensome to the storage facilities of the host computer. In general, the more complex or difficult the problem is to discern, the more storage is required to successfully diagnose the situation. This is especially true if multiple complex problems are present which are characterized by contradictory but simpler diagnoses. As a result, it is generally advisable to segregate a single large database into groups of smaller databases each with only a few components.

We have observed that processing speed is generally not a consideration when installing the WBD for continuous on-line operation. However, when batch operation is employed on a large number of buildings for long time periods, the delays in obtaining results can be considerable. For this reason, and for the sake of completeness, processing speed analysis was performed on a Pentium 166 MMX type computer running Windows 95. The results of this analysis are summarize in Table2.

**Table 1: Database size statistics**

| Database Contents | Database Size (Megabytes) |
|---|---|
| Configuration data only | 0.32 |
| Whole building energy data only | 1.5 |
| Properly operating air handler data only | 3.6 |
| Improperly operating air handler data only | 10.4 |

**Table 2: Diagnostic Processing Speed Statistics**

| Processing Scope | Processing Time (seconds) |
|---|---|
| Diagnose 1 day of air handler data | 7 |
| Diagnose 1 day of whole building energy data with 3 months of baseline data | 11 |
| Diagnose 1 day of whole building energy data with 9 months of baseline data | 31 |

## 4 The future of building diagnostics

An essential feature that we anticipate needing in future versions of the WBD is a linkage to simulations of physical components to provide a reference basis for model-based diagnosticians and for training baseline data needed by diagnosticians based on statistical or artificial intelligence methods. We anticipate that such simulation engines must meet a number of requirements that, at the current time, appear to be unsatisfied by the current generation of building simulation engines or those under development. First, simulations must be cost-effectively linked to diagnostic modules (either by integration or interoperability). Second, the simulation engine's data modeling requirements for diagnosticians are somewhat incompatible with the data that can be collected from building automation systems. Third, the simulation engines must be able to simulate a wide variety of realistic (i.e., commonly seen in the field) faulty behavior. As a result it is likely that the simulations must separate the component control scheme from physical behavior. Fourth, the simulations must be scalable. Thus, simulation engines that are given less data should provide more rough estimates of results, and the same ones when given more data provide more accurate estimates of results. Fifth, the simulation must provide comprehensive time-series reports of the interaction between the control system and the building components, reporting all sensor and actuator signals. Finally, these simulation engines must be capable of considering the uncertainty of some of the input data (e.g., occupancy, weather, sensor noise) when computing outputs such as energy use.

The International Alliance for Interoperability has been working on interoperability models called Industry Foundation Classes that will enable diagnosticians to obtain performance metrics from simulation tools in a more practical way than currently possible. However, the models to support the requirements suggested above have yet to be defined.

Tools such as the Facility Resource Energy Data (FRED) developed at Pacific Northwest National Laboratory collects, manages, and reports on large amounts of building energy usage data. The current installation of FRED at the DOE's Hanford Operations in Richland, Washington collects energy end-use data from about 90 buildings. Efforts are currently under way to integrate the WBD into this tool in order to enable tracking of building energy usage from within the FRED environment.

The data collection tools used by the WBD are not only capable of reading building automation system data, but they can send control overrides as well. While no current tools have been developed that exploit this capability, we anticipate developing diagnosticians which are capable of overriding control signals for short time periods in order to confirm diagnoses more quickly than is possible using the reasoning techniques already implemented in the WBD.

Continuous commissioning is another consequence of the ability to send control overrides from tools such as the WBD to the building automation system. In this approach, diagnosticians would periodically run building components through a series of exercises designed to demonstrate proper operations. In this way, a diagnostician would be capable of detecting a problem with a component (e.g., an economizer) before the component is needed.

## 5 Conclusions

The current field trials of the WBD have shown that system architecture is flexible enough to support a variety of building types and automation systems. These trials have also shown that the diagnosticians are effective in identifying installation and operation problems in air handling units in particular. Furthermore, the number of faulted air handling units found seems to support the conventional wisdom that the majority of economizers are experiencing problems due to installation, control, or maintenance failures.

The system architecture is designed to permit controls manufacturers to embed the diagnosticians within their control products. However, to date, no manufacturer has attempted this. We argue that by monitoring component performance proactively and continually, control systems will be able to ensure that equipment is more cost-effectively maintained, properly operated and with due diligence. In particular, we contend that the advent of tools such as the WBD will eliminate the excuse that is often used for failures to meet performance standards such as ASRHAE Standard 62-1989, *Ventilation for Acceptable Indoor Air Quality*: that the tools to practically and continuously verify compliance are not available. While it is certain that the WBD does not assure indoor air quality any better than it might assure any other standard of performance is met, it does provide a first line of defense in doing so: that an adequate proportion of outdoor air is indeed being drawn in. Given a family of diagnostic tools developed in the coming years, we expect that more and more buildings will achieve the performance standards to which they were designed.

## 6 References


Brambley M. R., R. G. Pratt, D. P. Chassin, S. Katipamula, D. Hatley (1998) "Diagnostics for Outdoor Air Ventilation and Economizers", ASHRAE Journal, Vol. 40, No. 10, October 1998.

CRC Handbook of Chemistry and Physics (1980/81), p. F246

Gold Institute Average Gold Prices (1997), Web Site, http://www.goldinstitute.com/average.html