

Detecting Cooling Coil Fouling Automatically— Part 1: A Novel Concept

Daniel A. Veronica, PhD, PE
Member ASHRAE

Received July 29, 2009; accepted March 5, 2010

Fouling of surfaces within the heat exchangers of heating, ventilating, and air-conditioning (HVAC) systems of buildings is an equipment fault that wastes appreciable amounts of energy; however, it escapes detection under current building automation technology. A novel concept is introduced to automatically detect this fouling on water-side and air-side surfaces of water-cooled HVAC air coils. The concept incorporates a model-driven component contributing a table of human expert information and an embedded data-driven component assimilating real-time data sampled from HVAC plant instrumentation. Supervisory programming (the “agent”) conducts real-time surveillance for coil fouling using the data-driven component, which is a model replicating the current dynamic thermal behavior of the coil. The surveillance is a specific characterizing transient (a “query”) exercised periodically on the data-driven dynamic coil model as a surrogate for exciting the real coil. When a query returns a suspect result, the agent determines if coil fouling or some other change caused that result by using the tabulated expert information. The concept makes use of all data sampled from the plant, reflecting transient and steady behavior. Fouling can be discerned from other developments, such as instrument drift, and the agent can distinguish air-side fouling from water-side fouling, estimate the severity of fouling, and estimate an uncertainty for its classification. The values tabulated describe three-dimensional surfaces characterizing the varied impact fouling generically has on coil thermal effectiveness when considered over the state space of coil operation. Categorical use of that information by the concept is justified by analysis in the effectiveness-NTU state plane.

A companion paper (Veronica 2011) gives results of exercising a crucial task within the concept on simulated data by using one form of a dynamic data-driven model; a multilayer perceptron.

INTRODUCTION

Fouling of heat transfer surfaces reduces the thermal conductance and thus the thermal effectiveness of water-cooled air coils in HVAC systems. A coil fouled on its water side (e.g., by slime, scale, or corrosion products) or air side (e.g., by particulate or fibrous buildup on fins) requires more flow of chilled water than when clean to maintain a regulated leaving air temperature at a given airflow rate. A penalty imposed by both water-side and air-side fouling is the added pump power demanded by this higher water flow requirement. Higher chilled-water flow delivered to a given coil heat load diminishes the temperature of water leaving the coil and thus diminishes the temperature difference the coil water side presents to the chiller plant (generally written in HVAC literature as “delta-T”). Another penalty occurs then for chiller plants whose control strategy has the side-effect of lost efficiency under reduced delta-T, as described by Taylor (2002). In the case of air-side fouling in coils subject to closed-loop airflow regulation,

Daniel A. Veronica is a mechanical engineer at the National Institute of Standards and Technology (NIST), Gaithersburg, MD.

added resistance to airflow imposes a third penalty: increased fan power to maintain the regulated airflow rate. The actual magnitude of these energy penalties depends on factors particular to each system installation and operation. Normal automatic regulation of air-handling unit (AHU) leaving air temperature and airflow rate compensates, and thus unavoidably masks, the ongoing loss of performance and efficiency. Fouling then escapes casual detection by building occupants or maintenance staff until it worsens to the extent that conditions in occupied spaces can no longer be kept comfortable. At that point, the penalties of fouling may have accrued undetected for months. A concept to detect coil fouling automatically at an earlier stage is thus warranted. The concept should be able to keep its surveillance using only information obtained readily from a typical building automation system (BAS).

Concepts to automatically detect faults (unwanted conditions) in a generalized process are distinguished most fundamentally by what a concept claims *a priori* as fact when new data, just acquired from process instrumentation, are first analyzed. The claimed knowledge is primarily either (1) a deductive model expressing human expert knowledge of physical principles motivating the process or (2) data sampled earlier from the process, inviting some inferred, data-driven model(s) of its behavior. Two distinct approaches to fault detection and diagnosis (FDD) thus result, being respectively (1) model driven and (2) data driven. A more detailed taxonomy of FDD is provided by Katipamula and Brambley (2005a, 2005b). The simpler distinction made in this paper clarifies the points important here yet is consistent with their more exhaustive treatment.

Data-driven methods of fault detection can assimilate complexity in a process autonomously, whether that complexity is due to nonlinearity of components or system effects on components due to their installation. A model-driven approach can explicitly dispatch well-known physical principles to detect faults in the process, avoiding the need to infer essentially the same physical principles stochastically—and thus incompletely or incorrectly perhaps—from large amounts of data.

For clarity, the term *model* in this paper refers only to a data-driven mapping of process inputs to output obtained autonomously by a program, typically embedded in microcontroller hardware, assimilating samples from instruments. *Simulation* and *simulator* refer to numerical algorithms written deductively from first principles and run independently of sampled data on, for example, a desktop computer. An unavoidable exception is the term *model-driven*, which already has wide acceptance in the literature as including desktop simulations.

A CONCEPT TO DETECT COIL FOULING AUTOMATICALLY

Figure 1 depicts a new concept to detect cooling coil fouling automatically by combining the strengths of data-driven and model-driven approaches. The decision that classifies the coil at any particular time as being clean or fouled is made by an autonomous supervisory program called the “agent,” which employs internally as needed (1) a real-time, data-driven component and (2) a lookup table embedded as the product of many simulations previously run off-line in an independent, model-driven component. Comprehensive texts on artificial intelligence, such as that by Luger (2002), describe agents having capabilities far broader than the term is meant to suggest here. The simpler agent considered in this paper is one meeting Luger’s more basic description of agents as “...mechanisms supporting decision making in the context of limited processing resources” (p. 16). The agent, data-driven component, and lookup table are presumably programmed into a microcontroller embedded at the AHU with access to its data.

As seen in Figure 1, the data-driven component used by the agent is an input-output model of the coil process dynamics. The agent maintains that data-driven dynamic model on-line to assimilate periodic samples of input data (entering air and water temperature, entering air humidity, airflow, and water flow or control valve signal) and output data (leaving air temperature) arriving

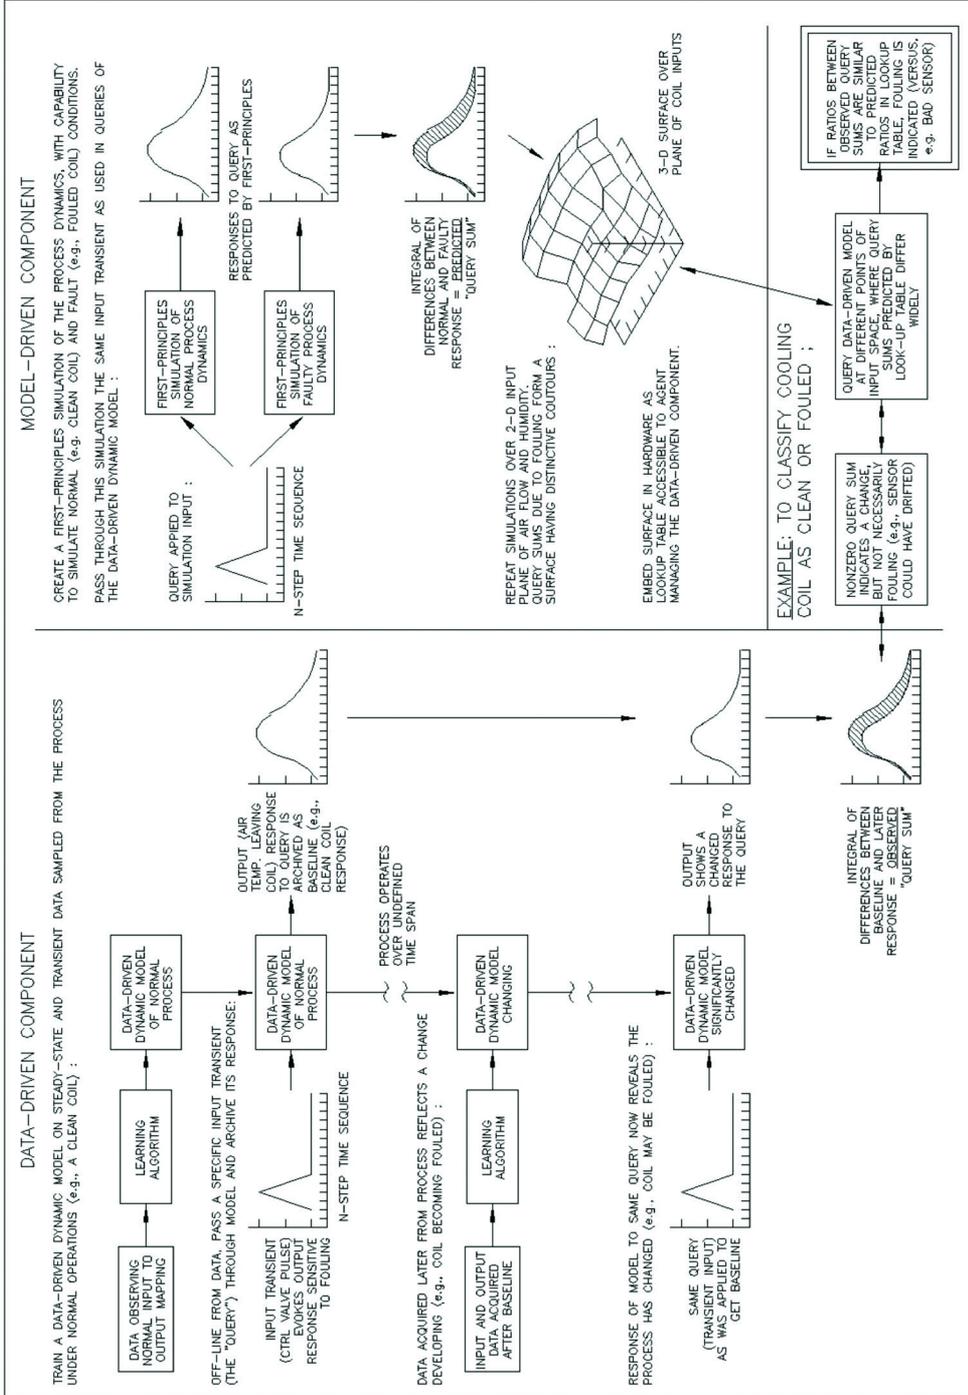


Figure 1. A novel concept to detect process faults automatically.

from AHU instrumentation. A classical technique to identify the unknown properties of a real input-output system is to pass an input transient through the system and record the dynamic output response (e.g., “reaction curves” in control tuning). The same idea is employed here as the fouling detection “query.” In this case it is not the real coil but a data-driven dynamic model of the real coil that is tested (“queried”) by the agent to determine if the real coil has become fouled.

The model-driven component contributes to the agent an embedded lookup table containing predicted values of a quantity called the “query sum.” As seen in Figure 1, a query sum integrates the disjoint area under two sequences responding to a query differently. A nonzero query sum obtained between successive queries of the data-driven model indicates a change developed in the model between the two queries, with coil fouling being only one of multiple causes possible. Query sums resulting from fouling, though, exhibit a distinctive shape when plotted over the state space of variables input to the coil process. This fact is depicted in Figure 1 as a surface established *a priori* by simulating many queries of clean and fouled coils off-line in the model-driven component. The surface shape results consistently from generic thermal effects that fouling alone exerts on cooling coils. The shape is expressed adequately by a modest number of query sum values stored in the lookup table. The agent exploits characteristic features of the shape in its decision making, discerning fouling on the coil air side from that on the water side and discerning fouling from other faults such as sensor drift.

Early in operation of the coil, when it is known to be clean, the data-driven model assimilates data sampled throughout the various regimes and excursions of normal AHU regulation. From that data, the model is converged numerically to (it “learns”) a mapping found satisfactory over a range of normal coil operation. The query specifically is a triangular pulse in the position of the coil water control valve. The four remaining input quantities (entering air and water temperatures, and entering air humidity and flow rate) hold constant values, locating the query at a specific point of input space. The coil model can be queried at points in the input space of the coil regardless of where in that space the real coil happens to be operating. This capability is used to build a library of baseline (“coil clean”) query responses soon after the model has learned normal coil operation adequately. Building the baseline means similar queries are exercised on the model at different points of input space—points that later will be used to determine if the coil has become fouled.

A crucial feature of the concept is illustrated by the example in the lower right corner of Figure 1. It is seen that detection of fouling does not rely on comparing data (e.g., leaving air temperatures) sampled from the real coil directly to corresponding values generated by first-principles simulations. Rather, surveillance queries of the data-driven model are made during AHU operation periodically at a well-learned point of input space known (from the model-driven lookup table) to yield a query sum significantly different from its baseline if fouling is present. A significantly nonzero query sum resulting from a single surveillance query indicates the model has changed since the baseline query but does not alone conclude fouling in the real coil as the cause. Complementary queries of the data-driven model are then exercised at other points of input space. If ratios between multiple query sums from the model are reasonably similar to ratios between query sums predicted by first principles at the corresponding points in the table, then the evidence supports fouling in the real coil as causing the change in the model.

The concept in Figure 1 resulted from answers to questions addressed by the research here. How can a data-driven model of a real coil process be employed to tell when the coil has become fouled? What distinctive evidence does fouling write into sampled data that is not found given other faults such as instrument drift? Is that evidence generic to all cooling coils? Can fouling on the coil air side be distinguished from fouling on the water side? Simulating the behavior of clean and fouled cooling coils numerically, a model-driven exercise of human expert knowledge answered those questions to motivate the concept.

SIMULATIONS ESTABLISH THE BASIS OF THE CONCEPT

Many numerical simulations of a cooling coil having various configurations (four-row and eight-row, and within each of those, full-circuited and half-circuited) were conducted, using two water-side fluids (pure water and water-glycol), over a comprehensive itinerary of steady-state and transient coil operations (Veronica 2008). Versions of the steady-state and dynamic cooling coil simulators developed by Zhou and Braun (2007) were modified by Veronica to incorporate water-side and air-side fouling as described in the appendix to this paper.

Three mechanisms are postulated through which surface fouling, whether water side or air side, affects coil performance. First, fouling covers convective metal surfaces of the coil with a largely stationary layer having, conservatively considered, conductivity equal to the adjacent moving fluid. Second, fouling reduces the area of convective surfaces exposed to the moving water and air. Third, given a water pump with closed-loop control sensing differential pressure across the coil pipe branch, fouling reduces the flow rate of water passing convective surfaces. The reduction of airflow rate through the coil due to air-side fouling, although significant to residential systems as shown by Siegel and Carey (2001), is not simulated here since large chilled-water (i.e., commercial) systems typically maintain airflows by closed-loop regulation of the fan pressure.

A nominal severity of water-side and air-side fouling is defined. *Nominal* is used here in its classic sense (“named”), fouling of a severity named only to enable different simulation scenarios to be compared equitably. The nominal fouling was established by trial simulations as that thickness of fouling diminishing peak heat transfer through the coil by 20%. Simulating fouling as uniformly having this nominal thickness does not necessarily imply the fouling material itself is deposited over convective surfaces uniformly. With conductivity of the fouling layer simulated as equaling that of the adjacent fluid, the implication is only that deposits of fouling material adhering to a convective surface impede motion of the layer of fluid next to the surface. This allows that a less conductive foreign fouling material might itself be distributed over the convective surface unevenly, but its thermal effect is represented conservatively by a relatively even, stilled layer in the more conductive fluid. In the full-circuited, four-row, water-cooled coil, the 20% performance loss occurs when 7% of the clean tube water flow radius is stilled or 35% of the clean fin air channel width is stilled.

Water Flow Transients Afford an Opportunity to Detect Fouling

Classical system identification includes applying an impulsive input. The output responding to an impulse reveals information on the “memory” of the system—its signal (i.e., energy) storage capacity and diffusion resistance. Fouling adds thermal capacity and resistance to a coil, changing its memory characteristics as evidenced by changed response to an impulse of water or airflow. A data-driven coil model cannot be expected to assimilate an impulse response, since impulses are not normally in data from real plants. A triangle-shaped pulse of flow approximates an impulse as a practical behavior learnable from real coil operation.

Figure 2 presents three plot panels. The top two collectively show that an upward triangular pulse in water flow produces a downward excursion in air temperature leaving a cooling coil. The top panel shows water flow ranging from 20% of full flow to full flow and back to 20% over a period of 96 seconds, a realistic rate of change. Values of the four input quantities remaining—entering air temperature ($T_{a,ent}$), entering air humidity ($W_{a,ent}$), air velocity (v_a), and entering water temperature ($T_{w,ent}$)—are listed. These remain constant to fix the query at a specific point of input state space.

The middle panel of Figure 2 shows the downward excursion in the temperature of air leaving the coil ($T_{a,lvg}$) as it results coincidentally from the upward triangle pulse in water flow. Two

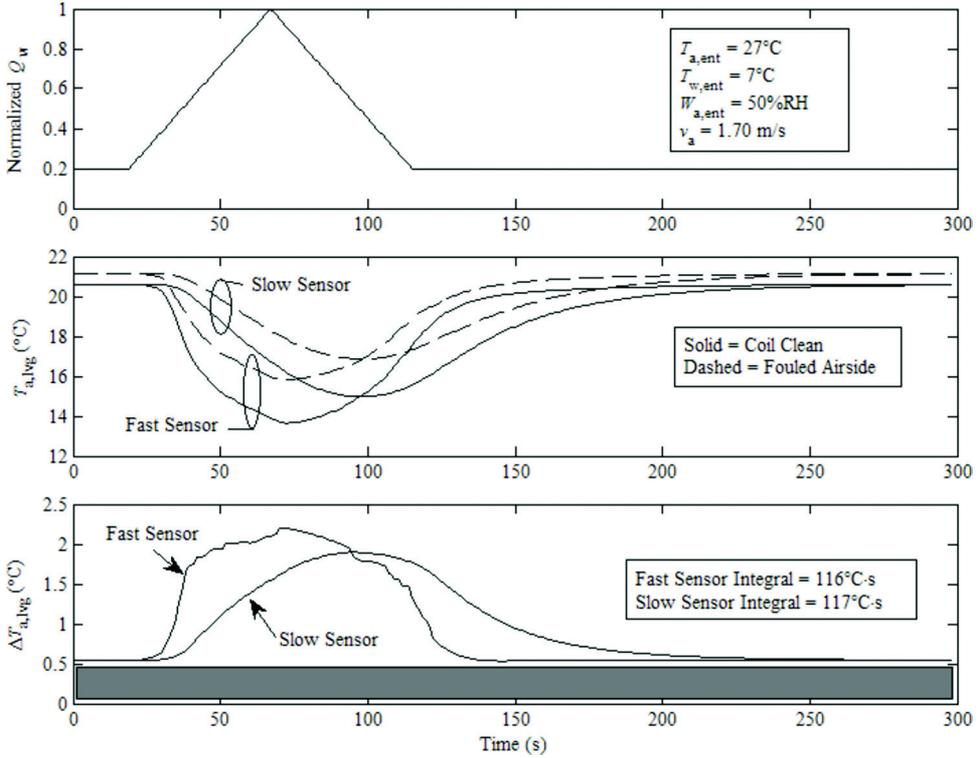


Figure 2. Response of clean and air-side fouled cooling coil to pulse in water flow.

pairs of temperature lines are plotted, each pair having a solid (coil clean) and dashed (coil fouled) line for two cases of responsiveness in the air temperature instrument (sensor).

In one case (“fast”), the sensor writes temperature into the data with no time lag, representing an unshielded device. Leaving air temperature calculated by the simulation is taken directly as the value measured. In a second case (“slow”), the sensor is shielded for mechanical protection. The shield mass and the larger film of relatively stagnant air surrounding it incur a significant delay between the time bulk airflow leaving the coil and passing the sensor factually has a particular temperature value and the time the state being sampled within the sensor (e.g., electrical resistance) reflects that fact. The temperature value sampled from the sensor at any instant is thus not necessarily the current ambient air temperature (T_{amb}) but a potentially different value (T_{lag}) driven by earlier conditions, the two quantities being equal only when conditions are unchanging. An ordinary differential equation represents the situation and has the following solution for the temperature sampled:

$$T_{lag}(t) = T_{amb} + (T_{lag}(t = 0) - T_{amb})e^{-t/\tau} \quad (1)$$

Given heat conduction within the sensor is far better than heat conduction to the surrounding air, it is appropriate to model the distributed physical properties actually involved more simply by a lumped-parameter time constant, τ :

$$\tau = \frac{mc}{\bar{h}_{film}A} \quad (2)$$

where m , c , and A represent respectively the mass, specific heat, and surface area of the shielded sensor. An averaged coefficient of heat convection, \bar{h}_{film} , across the thermal boundary layer surrounding the shield is the only variable involved, being dependent upon air properties and velocity as well as the shape of the shield. The coefficient is calculated using well-established relations for cylinders in cross flows (e.g., see Holman [1976]):

$$\bar{h}_{film} = \frac{Nu_{film}k_a}{d_{RTD}} \quad (3)$$

where

$$Nu_{film} = P(Re^b)(Pr^{1/3}) \quad (4)$$

Values for P and b are obtained from empirically determined tables found in most heat transfer texts. The Reynolds number (Re) and Prandtl number (Pr) are derived from the flow regime. Air thermal conductivity is k_{air} , and shield diameter is d_{RTD} . The parameter values used are those representing a commercially available averaging resistance temperature detector (RTD). A numerical simulation requires a solution defined not in continuous time from zero but as an ongoing sequence of discrete values employing finite differences. The finite-difference approximation to Equation 1 is

$$T_{lag}[n] = T_{lag}[n-1] + (T_{lag}[n-1] - T_{amb}[n-1])\left(\frac{-\Delta t}{\tau}\right) \quad (5)$$

A slow-responding, shielded air temperature sensor is thus simulated by sending raw output from the coil simulation through the low-pass digital filter defined by Equation 5.

The middle panel of Figure 2 shows that air-side fouling causes the excursion of leaving air temperature resulting from the water flow pulse to be shallower. The bottom panel of Figure 2 shows the temperature difference ($\Delta T_{a,lvg}$) defined by subtracting the air temperature leaving the clean coil from the air temperature leaving the coil when fouled. Summing values of each difference quantity at a one-second interval integrates approximately the area under each $\Delta T_{a,lvg}$ plot. Units of each summation are $^{\circ}C \cdot s$ ($^{\circ}F \cdot s$). The shaded rectangle under both the slow and fast sensed temperature differences ($\Delta T_{a,lvg}$) indicates an area removed from each summation, the area contributed by steady coil operation regardless of any transient. By subtracting that area out of the summation, the resulting integral reflects only effects on the coil dynamics. Any steady offset, whether due to fouling or other difference between the two cases (e.g., sensor drift), is purposely excluded. Given the sum is taken over a time span long enough to not significantly truncate integrating slow sensors, practically the same sum is obtained from fast or slow sensors. At the point of input space defined in the top panel, the bottom panel shows that nominal air-side fouling yields a summation of about $117^{\circ}C \cdot s$ ($210.6^{\circ}F \cdot s$) whether the sensor is fast or slow. This agreement is expected, as the water pulse ultimately transfers the same amount of thermal energy from the air regardless of the sensor response. A benefit to using a pulse as the excitation is that the integrating summation then becomes self-closing, not requiring imposed limits. Temperature differences responding to the triangular pulse typically follow a roughly trapezoidal path, as seen in the bottom panel of Figure 2, so their integrated area (the summation) is roughly linear with the height (peak $\Delta T_{a,lvg}$).

Figure 3 considers fouling only on the coil water side. As with air-side fouling, water-side fouling makes a shallower transient in leaving air temperature result from the water flow pulse. The bottom panel of Figure 3 shows a summation of about 113°C·s (203.4°F·s), nearly the same value as yielded by air-side fouling. No correlation is involved. The top panel of Figure 3 shows this 113°C·s (203.4°F·s) value occurs at a different point of input space, where airflow is one-fourth that considered in Figure 2.

The shallower excursion in leaving air temperature given fouling clearly shows that fouling reduces the heat exchange effectiveness of the coil. But only in particular regions of input space, specific to each fouling mode (air side or water side), is the effect as significant as seen in Figures 2 and 3. In other regions of input space (i.e., other conditions of the coil air and water), fouling of the same severity exerts very little impact upon the excursion of leaving air temperature. In those regions, the thermal effectiveness of any cooling coil is practically unaffected by fouling. There are yet other regions where fouling diminishes the excursion of leaving air temperature even more significantly than seen in Figures 2 and 3.

Multiple Query Sums Conforming to a Known Surface Indicate Fouling

Figure 4 shows the summation value predicted from the query at each of 49 points over the flow rate–humidity plane. The values are those predicted by the simulations based on air-side fouling of nominal severity. It is not arbitrary that there are 49 points (i.e., 7 × 7). More points do not reveal further useful surface features, and fewer points mask useful features. The shape of

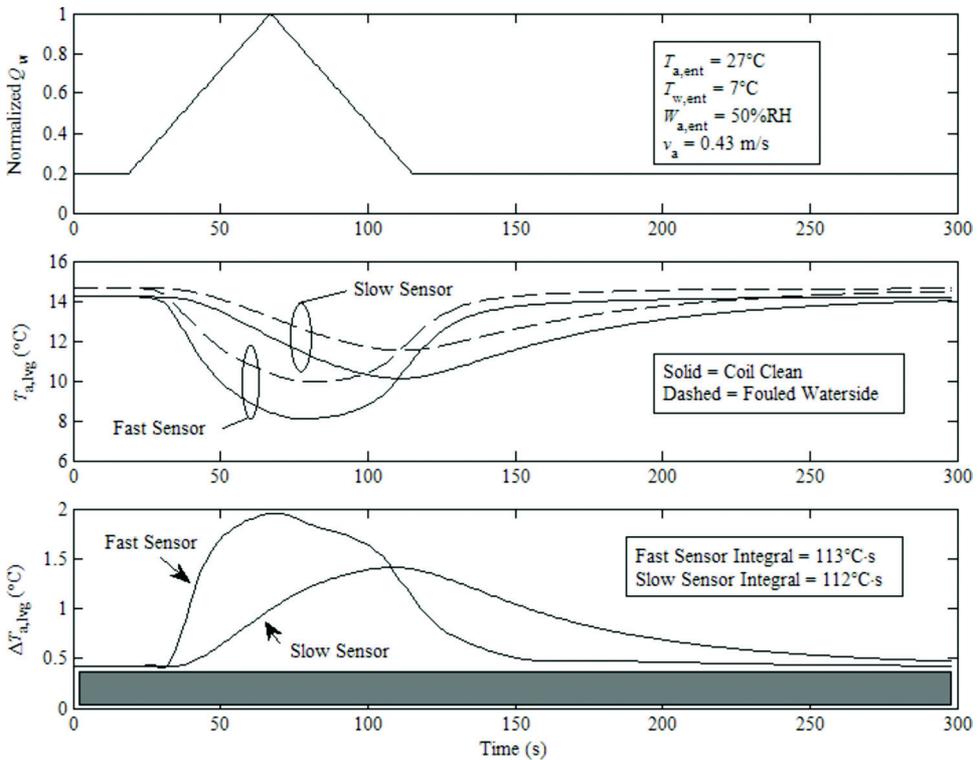


Figure 3. Response of clean and water-side fouled cooling coil to pulse in water flow.

the summation (or, “query sum”) surface over the airflow-humidity plane is very contorted, despite the fouling having a constant severity.

As seen in Figure 5, query sum values from water-side fouling differ markedly from those due to air-side fouling. This affords an opportunity to distinguish the two fouling modes automatically, based upon relative changes in the query sums obtained at different points of input space. A query is made of the model at point A in the coil input space, labeled in Figures 4 and 5. A nonzero query sum indicates the model has assimilated some change to the coil process since the baseline query at point A. There can be no conclusion yet as to what the change may be. A query sum near the $125^{\circ}\text{C}\cdot\text{s}$ ($225^{\circ}\text{F}\cdot\text{s}$) value corresponding to point A in Figure 4 suggests the coil is fouled on its air side at the nominal severity. Further queries at other points could build further evidence toward that classification and away from other possible reasons for the change in the model.

For example, if the coil is fouled only on its air side, a complementary query at point B must yield a sum 52% lower than that obtained at point A (e.g., $60^{\circ}\text{C}\cdot\text{s}$ [$108^{\circ}\text{F}\cdot\text{s}$] versus $125^{\circ}\text{C}\cdot\text{s}$ [$225^{\circ}\text{F}\cdot\text{s}$]). The thermal principles behind air-side fouling dictate that the sums obtained conform to the distinctive surface seen in Figure 4. Thus, the 52% lower sum at point B builds evidence that air-side fouling, not some other problem, is causing nonzero sums, since sums at points A and B relate by a ratio known to be consistent with air-side fouling. Conversely, a query at point B returning a sum about 285% larger than that at point A (e.g., $231^{\circ}\text{C}\cdot\text{s}$ [$415.8^{\circ}\text{F}\cdot\text{s}$] at point B versus $60^{\circ}\text{C}\cdot\text{s}$ [$108^{\circ}\text{F}\cdot\text{s}$] at point A) builds evidence for water-side fouling, since the ratio of the query sums conforms to the surface map of Figure 5.

The query sum surfaces of Figures 4 and 5 are generated for entering air and water temperatures of 32°C and 7°C (89.6°F and 44.6°F), respectively. These particular values are chosen because they offer a large thermal driving potential that exhibits fouling well. Also, the influence entering air and water temperatures exert on the summation value obtained is practically linear.

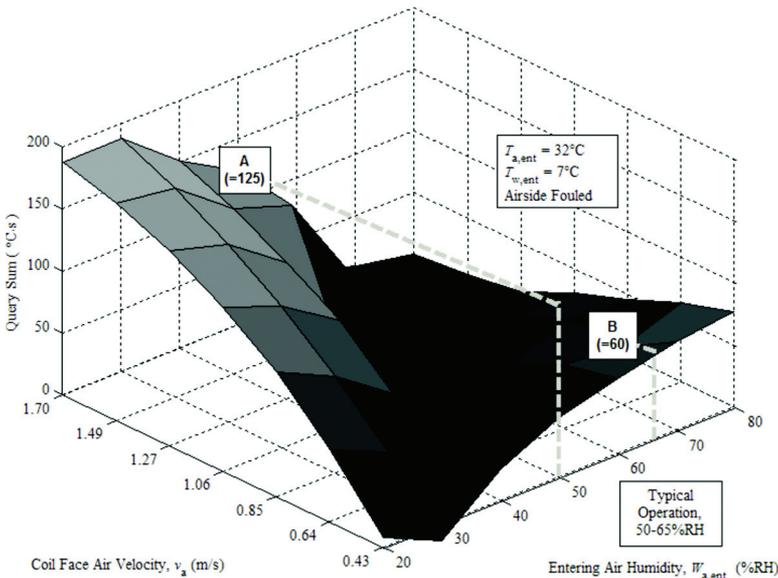


Figure 4. A summation surface for air-side fouling.

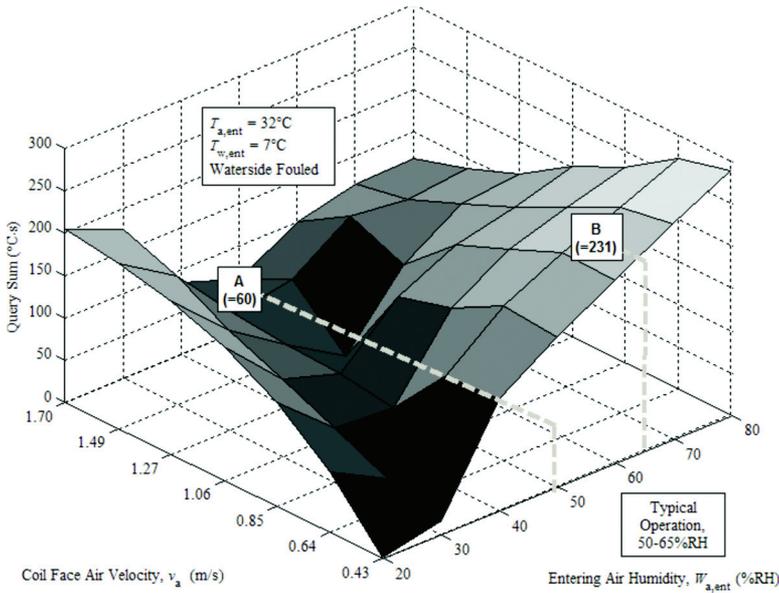


Figure 5. A summation surface for water-side fouling.

As a result, query sums for each mode of fouling are characterized adequately across the entire input domain by an ensemble having as few as four distinct surfaces—one surface for each permutation of high and low values for entering air and water temperatures (32°C and 23°C [89.6°F and 73.4°F] for air and 13°C and 7°C [55.4°F and 44.6°F] for water). The ensemble of surfaces essentially amounts to a lookup table having as few as 196 numbers (i.e., four surfaces comprised of 49 query sum values each). At any point of the input space, the value of the sum expected due to nominal air-side or water-side fouling is thus obtained by linear interpolation within the table.

CONCLUSION

Merits of the Concept

The concept illustrated in Figure 1 departs radically from prior approaches to FDD that also reference real-time data-driven models to simulation results. The departure is radical in the type of information referenced and in the nature of the reference. Dexter and Benouarets (1996) simulated a model-driven scheme to detect fouling in cooling coils by suggesting a qualitative interpretation of fouling. They modeled the coil in real time as mapping input measurements (e.g., airflow rate, valve position) to output measurements (approach) through a qualitative transform space of fuzzy partitions. Deciding that the coil is either clean or fouled falls to the graduated “degrees of similarity” by which the observed mapping agrees with each member in a reference library of compatible fuzzy mappings generated from simulations of known clean and fouled coils. While the researchers address the potential for ambiguity in that approach, a fuzzy transform of quantitative measurements into qualitative partitions is regarded necessary because “It is difficult to obtain adequate representations of the complex, ill-defined, and often highly nonlinear behavior of a faulty plant for use in [FDD]” (p. 550).

The concept in this paper suggests that the complex quantitative behavior of the coil is exactly what should be exploited for FDD, given the information referenced and the manner of reference are sufficiently robust. The query sum is potentially a more robust extraction of a fault feature (the disjoint area) than obtained by individual, isolated comparisons of sampled and predicted values. Eliminating all steady-state offset from the query sum furthers its robustness to drift of instruments and bias in the data-driven model.

Current approaches to coil fouling detection employing a quantitative first-principles simulation as the reference include Haves et al. (1996) and Xu et al. (2005). The former use on-line estimation of current coil state by a simple intermediary data-driven model, used in turn to estimate physically meaningful, fouling-sensitive parameters (e.g., conductance UA) in a parallel-ing first-principles simulation. The appendix to this paper shows that in some regimes of coil operation the much-reduced UA conductance in a fouled cooling coil is not reflected well by any change ($\Delta T_{a,avg}$) in its leaving air temperature. Unless an FDD scheme specifically exploits that fact as *a priori* knowledge, which the concept in this study does, it emerges on its own to complicate coil classification. FDD relying solely on real-time parameter estimation could be blind to a diminished UA conductance that is not reflected well in the arriving data.

The approach of Xu et al. (2005) is to establish well-calibrated, mutual agreement between real plant instrumentation and paralleled first-principles simulations. Some iteration of mutual corrections and retesting is likely. Fault surveillance commences once that is done, and future residuals between sampled data and simulation predictions are taken as suspect. It is the absolute nature of the comparisons made that demands such rigorous calibration. The concept in this paper fundamentally makes relative comparisons of data to first principles, by way of ratios between query sums, so calibration may be less rigorous. Steady disagreement in bias, for example, is cancelled both by calculating the sums and by comparing ratios.

Simulations show query sum values rise nonlinearly over a very broad range of fouling severity. The rise is practically linear, however, over the initial low range of less severe fouling more likely of concern to field surveillance (e.g., up to the 20% performance loss made nominal here). Estimation of fouling severity from query sum values is possible. The query sum surfaces obtained for all combinations of coil configuration simulated (by row quantity, circuitry, and coolant) shared a characteristic appearance having features similar to what is seen in Figures 4 and 5. The scale of those features differs sufficiently to warrant different tables for each configuration. Size and details such as tube and fin spacing might exert minor impacts, but they are not considered pivotal to concept proof and thus are not pursued here.

Lastly, prior work such as Dexter and Benouarets (1996), Haves et al. (1996), and Xu et al. (2005) were all considered with data sampled from the coil only at thermally steady (i.e., static equilibrium) states. The extensive study of HVAC commercial practice by Roth et al. (2002) ascribes a poor state of tune to many closed-loop controllers in the field. So, it is possible many HVAC systems do not sustain steady states for long enough periods to make it realistically practical to automatically detect coil fouling by relying solely upon steady-state models and equilibrium data. The concept in this paper assimilates data from transient and steady states and in fact benefits from more transients since the data-driven model sought must be dynamical.

Challenges for the Concept

Using the contours of query sum surfaces to ideal advantage requires that the data-driven model respond to queries at multiple points of input space with nearly equal validity so that significantly different query sums predicted by the model-driven component are available to help isolate fouling as a cause. If the model “knows” the current process dynamics of the real coil well in only a limited region of input space—the space visited by only the most recent AHU operations, for example—queries made on the model also must be confined to that region. In

that case some, but not all, functionality of the concept is lost. A nonzero query sum after the baseline would still indicate a change has developed in the model. But, without valid complementary queries, fouling cannot be automatically isolated as the cause. A principal merit of the concept thus relies on capabilities within the data-driven component. The model must learn coil dynamics quickly enough that even relatively short exposures to data from the regions of input space vital to FDD functionality are assimilated well and then retained well even after operation departs to another region. It is a challenging application of machine learning, and very specific development is required. Multiple paralleled models, each adept in a particular region of input space, may be needed. The agent, referring to the lookup table, could score relative experience, proficiencies, and uncertainties for each model and manage their use appropriately. That this division of labor might be necessary is why the functions of model and agent are differentiated.

Implementation Issues

The data-driven coil model requires an input related to water flow. At most HVAC sites, water flow will need to be mapped from the control valve command, an acceptable condition that at worst requires adding a water pressure input. The need to sense coil leaving air temperature (LAT) presents a harder issue. Air handlers typically sense that temperature only downstream of the fan. Modeling further assumptions or inputs to back out deduced LAT values might fold excessive uncertainty into the data. Given recent advances in networked semiconductor sensors, an inexpensive retrofit of a sensor array to obtain a real LAT input seems plausible. Because the capability of the data-driven model is crucial to viability of the concept, a companion paper (Veronica 2011) addresses that issue separately.

ACKNOWLEDGMENT

A fundamental idea in the concept presented here, that changes in a process can be detected automatically by periodic queries of an adaptive agent modeling the process, was contributed by Dr. Peter Curtiss of Curtiss Engineering, Boulder, Colorado. The appendix is based on guidance from Dr. Mike Brandemuehl, University of Colorado-Boulder. I thank Karl and Madeleine Larson, principal benefactors of the Larson Building Systems Laboratory at the University of Colorado-Boulder, where this research was performed, for their generous support.

NOMENCLATURE

A	=	area	n	=	index, discrete sample
b	=	exponent from reference data	P	=	coefficient from reference data
c	=	specific heat	Q	=	flow rate, volumetric
C	=	heat capacity	R	=	resistance, thermal
d	=	diameter	t	=	time
h	=	coefficient of convection	T	=	temperature
H	=	enthalpy	U	=	heat conductance (per unit fluxed area)
k	=	conductivity, thermal	v	=	velocity (at coil face)
m	=	mass	W	=	humidity

Subscripts

a	=	air or at air-side interface	lag	=	lagging (the true value)
amb	=	ambient	lvg	=	leaving (the coil)
ent	=	entering (the coil)	min	=	minimum
fa	=	of the material fouling air side	p	=	at constant pressure
fw	=	of the material fouling water side	sat	=	at moisture saturation
fin	=	fin (extended heat transfer surface)	tot	=	total
H	=	enthalpy-based (vs. sensibly-based)	w	=	water or at water-side interface

Superscripts

* = for fully wetted air-side control volume

Greek Symbols

Δ = difference (finite)

η = efficiency

ε = effectiveness, thermal

τ = time constant

REFERENCES

- Dexter, A.L., and M. Benouarets. 1996. A generic approach to identifying faults in HVAC plants. *ASHRAE Transactions* 102(1):550–56.
- Katipamula, S., and M. Brambley. 2005a. Methods for fault detection, diagnostics and prognostics for building systems—A review part I. *HVAC&R Research* 11(1):3–25.
- Katipamula, S., and M. Brambley. 2005b. Methods for fault detection, diagnostics and prognostics for building systems—A review part II. *HVAC&R Research* 11(2):169–87.
- Kuehn, T.H., J.W. Ramsey, and J.L. Threlkeld. 1998. *Thermal Environmental Engineering*, 3d ed. Upper Saddle River, NJ: Prentice-Hall.
- Haves, P., T.I. Salsbury, and J.A. Wright. 1996. Condition monitoring in HVAC subsystems using first principles models. *ASHRAE Transactions* 102(1):519–27.
- Holman, J.P. 1976. *Heat Transfer*, 4th ed. New York: McGraw-Hill.
- Luger, G.F. 2002. *Artificial Intelligence*, 4th ed. Harlow, Essex, UK: Pearson Education Ltd.
- Roth, K., D. Westphalen, J. Diekmann, S. Hamilton, and W. Goetzler. 2002. *Energy Consumption Characteristics of Commercial HVAC Systems Volume III: Energy Savings Potential*. Cambridge, MA: TIAX LLC.
- Siegel, J., and V. Carey. 2001. Fouling of HVAC fin and tube heat exchangers. Publication LBNL-47668. Berkeley, CA: Lawrence Berkeley National Laboratory.
- Taylor, S. 2002. Degrading chilled water plant delta-T: Causes and mitigation. *ASHRAE Transactions* 108(1):641–53.
- Veronica, D.A. 2008. Detecting heat exchanger fouling automatically with an embedded data-driven agent using expert signature maps. PhD dissertation, Civil, Environmental, and Architectural Engineering Department, University of Colorado at Boulder.
- Veronica, D.A. 2011. Detecting cooling coil fouling automatically—Part 2: Results using a multilayer perceptron. *HVAC&R Research*. Accepted for publication.
- Xu, P., P. Haves, and M. Kim. 2005. Model-based automated functional testing - methodology and application to air-handling units. *ASHRAE Transactions* 111(1):979–89.
- Zhou, X., and J. Braun. 2007. A simplified dynamic model for chilled-water cooling and dehumidifying coils—Part 1: Development (RP-1194). *HVAC&R Research* 13(5):785–804.

APPENDIX

A principal merit of the concept presented in this paper is that it distinguishes fouling generically from any other cause of nonzero query sums by reference to the characteristic surfaces' query sums from fouling present over the input space, exemplified by Figures 4 and 5. This appendix shows that reference is based factually upon a thermal relationship generic to fouling in cooling coils, and thus the surfaces inform the FDD agent consistently and reliably. The relationship exists not in the observable state space of the coil instrument data but in multiple, unobservable state planes of overall coil effectiveness and overall coil number of transfer units (NTU). The effectiveness-NTU state plane is a tool of steady-state analysis, however, not directly applied to transients. To explain how input state affects query sums despite a fixed fouling severity, two steady “end” states are defined from the query transient as seen in Figures 2 and 3. One steady end state is at the initial low water flow of the triangular pulse, and the other is defined by its momentary high water flow. These two end states are employed in a series of steady-state simulations bracketing an assumed operating region of coil input space, as seen in Table A1.

**Table A1. Using Steady-State Simulations
 to Derive a Data Signature for Fouling in a Cooling Coil**

Input Code	Steady-State Valuation of Coil Input Space					$\Delta(T_{a,avg})$, (Fouled – Clean), °C (°F)	Change in $\Delta(T_{a,avg})$, Between High and Low Water Flow, °C (°F)		
	$T_{a,ent}$ °C (°F)	$W_{a,ent}$ % RH	v_w m/s (fpm)	$T_{w,ent}$ °C (°F)	Q_w L/s (gpm)	Water Side Fouled	Air Side Fouled	Water Side Fouled	Air Side Fouled
0	23 (73.4)	20	0.43 (85)	7 (44.6)	0.5 (7.9)	0.65 (1.17)	0.48 (0.86)		
1	23 (73.4)	20	0.43 (85)	7 (44.6)	2.4 (38)	0.48 (0.86)	0.34 (0.61)		
2	23 (73.4)	20	0.43 (85)	13 (55.4)	0.5 (7.9)	0.41 (0.74)	0.30 (0.54)		
3	23 (73.4)	20	0.43 (85)	13 (55.4)	2.4 (38)	0.30 (0.54)	0.21 (0.38)		
4	23 (73.4)	20	1.70 (335)	7 (44.6)	0.5 (7.9)	0.46 (0.83)	0.43 (0.77)		
5	23 (73.4)	20	1.70 (335)	7 (44.6)	2.4 (38)	2.46 (4.43)	2.35 (4.23)		
6	23 (73.4)	20	1.70 (335)	13 (55.4)	0.5 (7.9)	0.29 (0.52)	0.27 (0.49)		
7	23 (73.4)	20	1.70 (335)	13 (55.4)	2.4 (38)	1.54 (2.77)	1.47 (2.65)		
8	23 (73.4)	80	0.43 (85)	7 (44.6)	0.5 (7.9)	0.46 (0.83)	0.05 (0.09)		
9	23 (73.4)	80	0.43 (85)	7 (44.6)	2.4 (38)	2.12 (3.82)	0.68 (1.22)		
10	23 (73.4)	80	0.43 (85)	13 (55.4)	0.5 (7.9)	0.23 (0.41)	0.07 (0.13)		
11	23 (73.4)	80	0.43 (85)	13 (55.4)	2.4 (38)	1.26 (2.27)	0.38 (0.68)		
12	23 (73.4)	80	1.70 (335)	7 (44.6)	0.5 (7.9)	0.13 (0.23)	0.48 (0.86)		
13	23 (73.4)	80	1.70 (335)	7 (44.6)	2.4 (38)	2.24 (4.03)	1.15 (2.07)		
14	23 (73.4)	80	1.70 (335)	13 (55.4)	0.5 (7.9)	0.12 (0.22)	0.45 (0.81)		
15	23 (73.4)	80	1.70 (335)	13 (55.4)	2.4 (38)	1.14 (2.05)	0.82 (1.48)		
16	32 (89.6)	20	0.43 (85)	7 (44.6)	0.5 (7.9)	1.01 (1.82)	0.73 (1.31)	-0.31	-0.25
17	32 (89.6)	20	0.43 (85)	7 (44.6)	2.4 (38)	0.70 (1.26)	0.48 (0.86)	(-0.56)	(-0.45)
18	32 (89.6)	20	0.43 (85)	13 (55.4)	0.5 (7.9)	0.77 (1.39)	0.55 (0.99)		
19	32 (89.6)	20	0.43 (85)	13 (55.4)	2.4 (38)	0.53 (0.95)	0.37 (0.67)		
20	32 (89.6)	20	1.70 (335)	7 (44.6)	0.5 (7.9)	0.74 (1.33)	0.69 (1.24)	3.06	2.94
21	32 (89.6)	20	1.70 (335)	7 (44.6)	2.4 (38)	3.80 (6.84)	3.63 (6.53)	(5.51)	(5.29)
22	32 (89.6)	20	1.70 (335)	13 (55.4)	0.5 (7.9)	0.56 (1.01)	0.52 (0.94)		
23	32 (89.6)	20	1.70 (335)	13 (55.4)	2.4 (38)	2.89 (5.20)	2.76 (4.97)		
24	32 (89.6)	80	0.43 (85)	7 (44.6)	0.5 (7.9)	0.61 (1.10)	-0.03 (-0.05)	3.87	1.22
25	32 (89.6)	80	0.43 (85)	7 (44.6)	2.4 (38)	4.48 (8.06)	1.19 (2.14)	(6.97)	(2.20)
26	32 (89.6)	80	0.43 (85)	13 (55.4)	0.5 (7.9)	0.41 (0.74)	-0.01 (-0.02)		
27	32 (89.6)	80	0.43 (85)	13 (55.4)	2.4 (38)	3.42 (6.16)	0.82 (1.48)		
28	32 (89.6)	80	1.70 (335)	7 (44.6)	0.5 (7.9)	0.14 (0.25)	0.50 (0.90)	3.00	0.45
29	32 (89.6)	80	1.70 (335)	7 (44.6)	2.4 (38)	3.14 (5.65)	0.95 (1.71)	(5.40)	(0.81)
30	32 (89.6)	80	1.70 (335)	13 (55.4)	0.5 (7.9)	0.12 (0.22)	0.54 (0.97)		
31	32 (89.6)	80	1.70 (335)	13 (55.4)	2.4 (38)	2.11 (3.80)	0.83 (1.49)		

Characterizing the Steady-State Impacts of Coil Fouling with Signature Surfaces

Table A1 shows the 32 permutations (“input codes”) resulting when the five variables input to the coil simulations are each given high and low values defining extremes of steady operation. The simulations yield an intuitive outcome, that fouling of a cooling coil is exhibited best in data sampled when the driving potential for heat is greatest—that is, when the source fluid (air) is hottest and the sink fluid (water) is coldest. This alone explains the relatively low values of $\Delta T_{a,lvg}$ in 24 of the 32 possible input valuations (the unshaded rows in Table A1). Each of those 24 yields smaller $\Delta T_{a,lvg}$ values when compared to an input valuation that is similar except for entering air temperature being higher, entering water temperature being lower, or both. A convenience of the way the inputs in Table A1 are ordered is that the remaining eight (shaded) valuations of the input space lie as four complementary pairs, differing only in water flow rate. Not so intuitive is how the other three inputs (entering air humidity, airflow rate, and water flow rate) influence the effect fouling has upon leaving air temperature.

When water flow rate (Q_w) is high and airflow rate (Q_a) low, water-side fouling is most evident in data sampled when air humidity is high. That is, the largest $\Delta T_{a,lvg}$ value between a water-side fouled coil and a clean coil is yielded at the point of input space labeled “Input Code 25” in Table A1. But, if both water flow and airflow rates are high, then water-side fouling is more evident in data sampled at low humidity. That is, water-side fouling yields a larger $\Delta T_{a,lvg}$ at Input 21 than at Input 29. The consistent thermal principle behind such seemingly inconsistent observations is revealed by reducing the input dimensions remaining free (i.e., having a yet unexplained influence).

The free inputs are reduced to two by defining, in terms of either water flow or airflow, a quantity serving as the signature fouling writes upon data sampled from a coil. The quantity is a fouling “signature” because it is significant in data from fouled coils but insignificant in data from clean coils. The signature quantity most sensitive to both water-side and air-side fouling is defined in terms of differing steady water flows. Consider two arbitrary rates of water flow, low and high. The quantity $\Delta(\Delta T_{a,lvg})$ then expresses the difference between the resulting two values obtained for the change ($\Delta T_{a,lvg}$) fouling exerts on air temperature leaving the coil ($T_{a,lvg}$) at each water flow rate. Though it seems abstract, subsuming water flow into the $\Delta(\Delta T_{a,lvg})$ signature quantity allows the impact fouling exerts on coil performance to be visualized easily over the input space remaining free as a three-dimensional surface. This surface lies over a plane having airflow rate and humidity as axes.

The two rightmost columns of Table A1 each show four values of the fouling signature, $\Delta(\Delta T_{a,lvg})$, which become “corner” values of two such surfaces, one due to water-side fouling and one due to air-side fouling. Figure A1 shows the surface produced by water-side fouling given the high entering air temperature and low entering water temperature used in Table A1, 32°C and 7°C (89.6°F and 44.6°F), respectively. Importantly, trial simulations confirm entering air and water temperatures each exert a largely linear influence upon leaving air temperature and thus the fouling signature value. This is expected, as entering air and water temperatures are used linearly in the finite-difference equations of the simulations. As a result, the signature is reasonably well characterized over the entire coil input space by only three more instances of this surface, one at each remaining permutation of high and low values for entering air and water temperature. More usefully, the relatively linear influence of entering air and water temperatures means the query sum expected from a data-driven model of a fouled coil is also similarly characterized by four three-dimensional surfaces, one at each permutation of the high and low values of entering air and water temperatures.

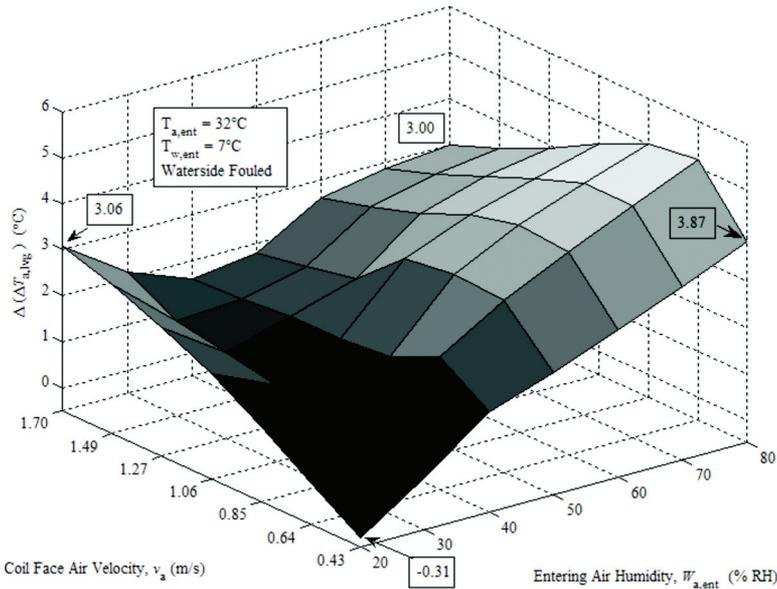


Figure A1. A signature surface for water-side fouling.

Simulating Coil Fouling Numerically

The Figure A1 surface is shaped by the differing rates of heat the coil is able to transfer in different regions of the airflow-humidity state plane, despite the source and sink temperatures remaining constant. This behavior is consistent only when expressed in terms of the thermal effectiveness of the coil and is explained best by first reviewing the thermal modeling employed in the simulations created by Zhou and Braun (2007).

The coil is divided into control volumes, each representing one water tube row (i.e., transverse plane of air cooling) in the finned core. When the humidity of entering air is sufficiently low (e.g., 20% RH), only sensible energy is removed from the air and all air-side surfaces (and thus, all control volumes) are dry. A simplifying feature of this thermal model regards the heat transfer between coil metal and each fluid (i.e., water side and air side) as a separate relation of thermal effectiveness (ε) and NTU for each control volume. A constant metal temperature is presumed across each control volume (tube row). Though in reality it is a counter-flowing array of cross-flowing tube rows, the coil is thus modeled as a lineal chain of control volumes, each representing one water tube row, each containing purely counter-flowing fluids. This simplification allows on each side of the coil metal the special case of thermal effectiveness being regarded as a rising exponential asymptotic to unity:

$$\varepsilon = 1 - e^{-NTU} \tag{A1}$$

Separate NTU values are defined for each convective interface to the metal. A contribution of this paper is to incorporate surface fouling at this juncture, adding the reciprocals of “fouling factors” (R_{fa} for air side and R_{fw} for water side) in series with the coefficients (h) of clean surface convection. Units of fouling factors are thermal resistance over a unit fluxed area ($[\text{m}^2 \cdot ^\circ\text{C}]/\text{W}$ [$\text{ft}^2 \cdot \text{h} \cdot ^\circ\text{F}/\text{Btu}$]), the reciprocal of units in convection coefficients. The fouling factors are calcu-

lated from the thickness and conductivity of stilled fluid layers yielding a nominal performance loss as described earlier. In each control volume, the NTU for the water side is then

$$NTU_w = \frac{(h_w + 1/R_{fw})A_{w,tot}}{\dot{C}_w} \quad (A2)$$

The NTU for the air side is then

$$NTU_a = \frac{\eta_{fin}(h_a + 1/R_{fa})A_{a,tot}}{\dot{C}_a} \quad (\text{dry}) \quad (A3)$$

$$NTU_a^* = \frac{\eta_{fin}^*(h_a^* + 1/R_{fa})A_{a,tot}}{\dot{C}_a} \quad (\text{wet}) \quad (A4)$$

While the NTU denominator is conventionally the minimum of the two fluid heat capacity rates involved, here each fluid interface is addressed separately using its own capacity rate. Zhou and Braun (2007) detailed origins of the convection coefficients (h_w , h_a , and h_a^*) and fin efficiencies. It suffices here to see the NTU values map through Equation A1 to result in effectiveness values used to calculate thermal resistances. These thermal resistances are exercised linearly in the finite-difference equations of the numerical simulations. They cover all energy transfers possible within the control volume: (1) the coil metal to water (R_w) and (2) the air to dry coil metal (R_a) or condensing air to wet coil metal (R_a^*):

$$R_w = \frac{1}{\varepsilon_w \dot{C}_w} \quad (A5)$$

$$R_a = \frac{1}{\varepsilon_a \dot{C}_a} \quad (A6)$$

$$R_a^* = \frac{1}{\varepsilon_a^* \dot{M}_a} \quad (A7)$$

It would have been incorrect to incorporate fouling simply as additions to the resistances here rather than as additions inside the NTU formulations above. A constitutive characteristic of indirect heat exchange between moving fluids is that the impact from a given change in NTU diminishes as NTU grows large and grows large as NTU is diminished. That character is constituted in the simulations by the exponential mapping of Equation A1, so any developments in thermal conductance also must pass through that mapping.

Factors that Shape the Impact Fouling has on Leaving Air Temperature

Figure A2 compares clean and water-side fouled cases of overall thermal resistance for an entire coil at various steady-state water flow rates, given two different entering air humidity levels. Here, "overall" thermal resistance sums in parallel the individual total thermal resistances of all control volumes in the coil. Each "total" thermal resistance sums in series all elemental film and metal resistances lying between bulk air and water in a control volume. Fouling on either air-side or water-side convection surfaces adds another resistance element into the series between air and water in each control volume. Lines plotted in Figure A2 end at states corresponding to input conditions listed in Table A1, so the figure essentially maps all the steady

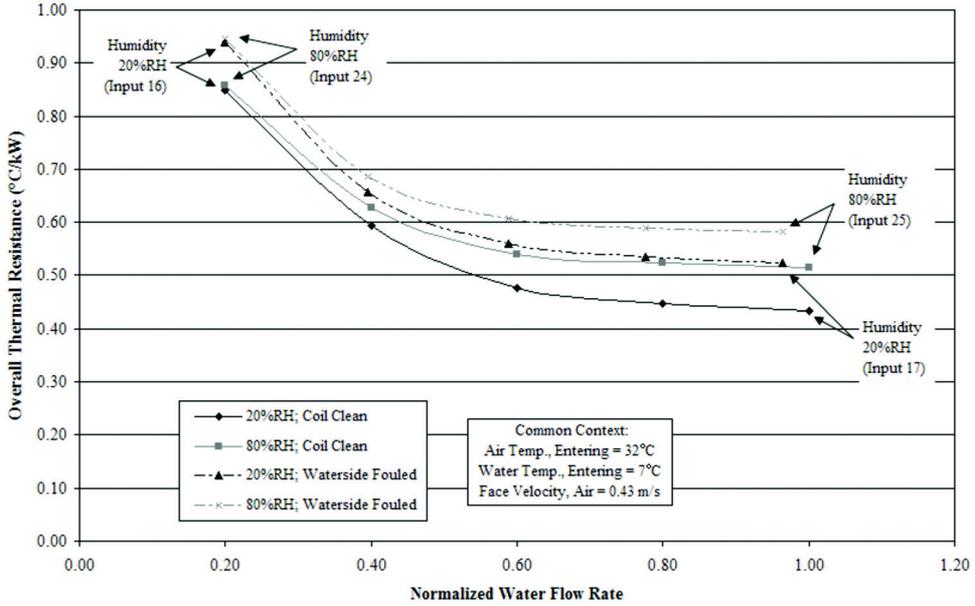


Figure A2. Impact of water-side fouling on coil overall thermal resistance.

states of water flow lying between Inputs 16 and 17 (low humidity) and between Inputs 24 and 25 (high humidity). Overall thermal resistance is seen falling asymptotically with higher water flow due to better heat convection across the turbulent boundary film. The fixed amount of resistance added by a nominal layer of water-side fouling shifts total resistance upward nearly symmetrically across the water flow range at both low and high humidity. This nearly equal shift of resistance is proportionally a greater share of total resistance at high water flow than at low water flow. Further, the resistance proportion due to fouling grows more steeply, and ends up higher, when humidity is low. Based on thermal resistance effects alone, it then seems fouling should exhibit a larger signature (the $\Delta(\Delta T_{a,lvg})$ in Table A1) at low humidity (between Inputs 16 and 17) than at high humidity (between Inputs 24 and 25). But, Table A1 shows the opposite is true. In fact, the -0.31°C (-0.56°F) value at low humidity is not only insignificant compared to the 3.87°C (6.97°F) change at high humidity, it is slightly negative. The impact of fouling on thermal resistance alone does not correctly explain why some regions of coil input space are more conducive to fouling detection than others.

In Figure A3, overall NTU of the coil is plotted for each case in Figure A2. For the 20% humidity case plotted in Figure A3, all the control volumes in the coil have dry air sides, so only sensible energy is transferred from the air. Overall coil NTU is thus obtained using the conventional formula dividing overall conductance (the reciprocal of the resistance in Figure A2) by the minimum fluid heat capacity. When all control volumes have wet air sides, as in the 80% humidity case of Figure A2, both latent and sensible energy are transferred. An enthalpy-based extension to the conventional effectiveness-NTU model is then required. The extension is due to Threlkeld (cited in Kuehn et al. [1998]). Water and air enthalpies in a wet control volume are considered potentials driving energy across an “enthalpy conductance,” an analogy to the sensible heat conductance (UA product) of dry control volumes. In the wet control volume, enthalpy conductance is

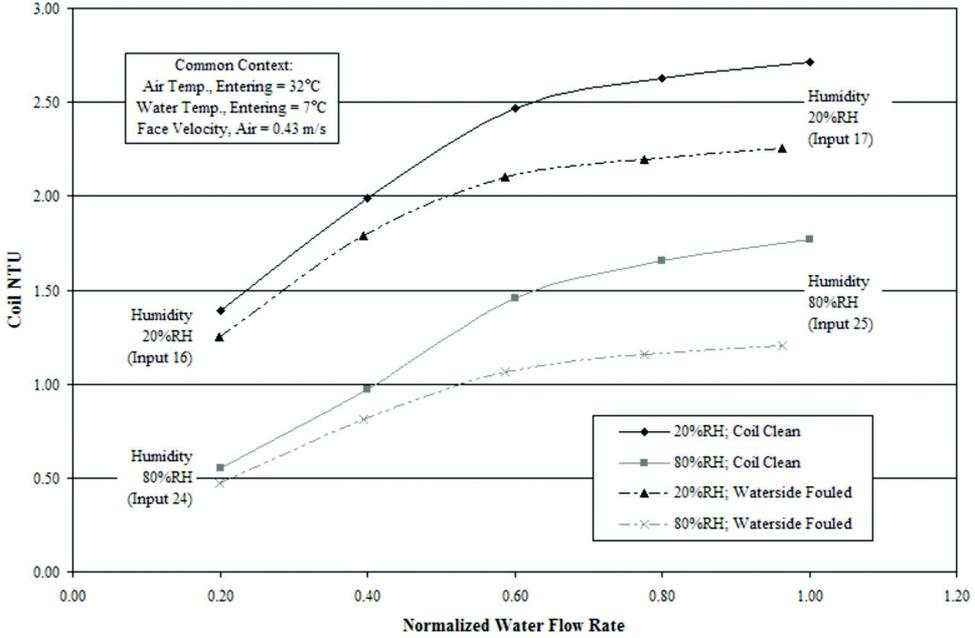


Figure A3. Impact of water-side fouling on coil NTU.

$$(UA)_H = \left(\frac{c_{p,a}}{(UA)_a} + \frac{c_{p,sat}}{(UA)_w} \right)^{-1} = (c_{p,a}R_a^* + c_{p,sat}R_w)^{-1} . \quad (A8)$$

Units of enthalpy conductance are those of a mass flow rate (kg/s [lb_m/s]), less intuitive than units of sensible conductance, such as kW/°C (Btu/h·°F). From enthalpy conductance, a dimensionless, enthalpy-based NTU is defined as follows:

$$NTU_H = \frac{(UA)_H}{\dot{m}_a} \quad (A9)$$

In Figure A3, NTU_H from the enthalpy-based model is plotted for the 80% humidity case and NTU from the conventional model is plotted for the 20% humidity case. In both cases of humidity, the NTU lost to fouling widens markedly at higher water flow. This contrasts with Figure A2, where the thermal resistance gained from fouling changes little with water flow. The widening NTU loss is one of two key features in Figure A3, so its origin must be emphasized. While fouling adds a physically serial, fixed increment to the drop in resistance across the fluid interface given a higher flow, its corresponding decrement to the greater interface conductance involves parallel addition, a nonlinear operation. So the NTU lost to fouling must widen at higher NTU. The widening is not contingent on anything else. The two key features in Figure A3 are (1) that the NTU lost to fouling widens with higher water flow by similar amounts at both high and low humidity and (2) that NTU is uniformly larger at low humidity than at high humidity.

NTU maps to overall thermal effectiveness of the coil (ϵ_{coil}) in Figure A4. This is not effectiveness at individual convection interfaces, as used in Equations A5, A6, and A7, but a quantity calculated by psychrometric analyses of the coil inlet and outlet states:

$$\epsilon_{coil} = \frac{H_{a,ent} - H_{a,lvg}}{H_{a,ent} - H_{min}} \quad (A10)$$

Here, H_{min} is the minimum value of enthalpy the air-cooling process can attain theoretically under the second law of thermodynamics. This is not the apparatus dew point but rather a less energetic state intersecting the coil entering water temperature with the leaving air humidity. In a coil having dry air sides, the air temperature observed to leave the coil is linear with effectiveness since air specific heat is constant over the temperatures applying. If the air side is wetted, leaving air temperature is nearly linear with effectiveness, curved only slightly by the specific heat of air changing slightly as it dries. So, the transform from the impact fouling has on overall thermal effectiveness of the coil in Figure A4 to the impact it has on air temperature leaving the coil is practically linear for all cases.

The two features of Figure A3 now play distinct roles that ultimately determine the change in overall thermal effectiveness of the coil, which results from nominal water-side fouling. Given higher values of steady water flow, overall coil effectiveness maps as states rising in an arc asymptotic to unity. The particular path of the arc depends on the remaining coil inputs, such as humidity and airflow. The arc reflects the saturating character inherent to energy exchange between moving fluids, modeled in the control volumes by Equation A1. At the high 2.4 L/s (38.0 gal/min) water flow, nearly equal NTU is lost to fouling at low humidity as at high humidity (i.e., the similar shifts from solid lines to dashed lines along the plot abscissa in Figure A4).

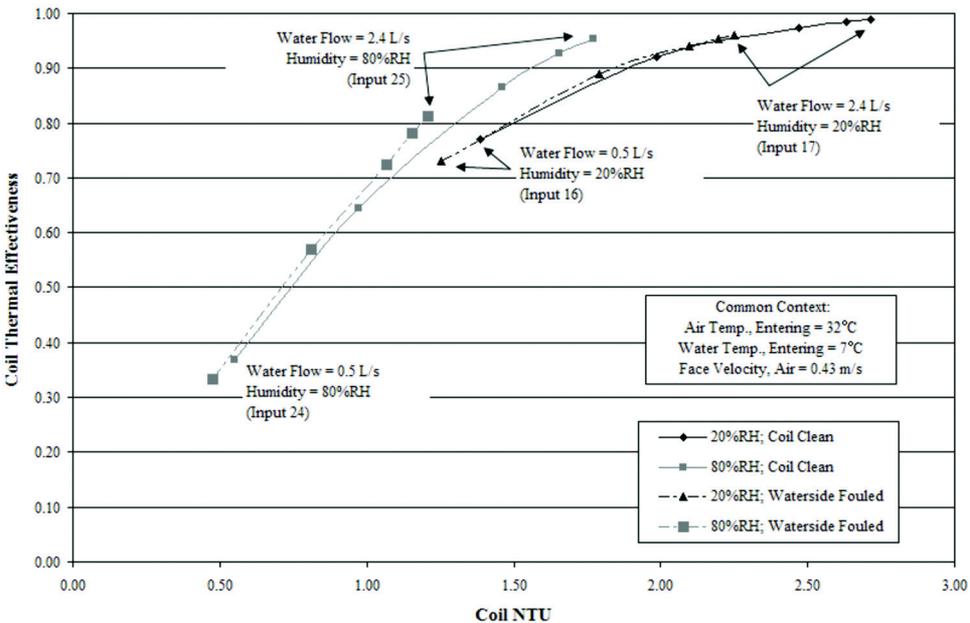


Figure A4. Impact of water-side fouling on coil thermal effectiveness.

Since the low humidity case lies globally at higher NTU, where the effectiveness curve is much flatter, the equal NTU lost there projects to the ordinate axis as a much smaller loss of coil effectiveness. Fouling shifts the arc of high-humidity states leftward on the Figure A4 abscissa, to lower NTU. Since NTU is globally lower at high humidity, the shift occurs on a much steeper segment of the effectiveness curve. A shift of NTU equal to that in the low-humidity case now results in a more significant change in coil effectiveness and, hence, in leaving air temperature.

NTU and effectiveness of the coil were examined comprehensively over the coil input space, as expressed by surfaces such as those in Figures 4, 5, and A1. Regions conducive to easy detection of fouling—that is, having higher query sums from fouling—are generally characterized by (1) a relatively large shift of NTU due to fouling, combined with (2) a relatively low NTU value (e.g., below 2.0, roughly). That combination puts a large NTU shift on the steeper part of the effectiveness curve, where it translates into a larger effectiveness change and thus a greater change in leaving air temperature.

