Virtual Closed Loop Control for Adaptive Energy Management in Retail Buildings

Paul Campbell
Vice President of AI & ML Solutions, facil.ai

In order to produce a more precise operation of HVAC assets in retail buildings, we need to emerge from the limitations of the Open Loop control environment of commonly available Building Automation Systems. However these Open Loop control assets can still be leveraged and precision achieved through an external translation of Open Loop to Virtual Closed Loop control executed locally using Advanced Supervisory Control and by doing so elevate precision as the foundation upon which the strategic domain of Adaptive Energy Management is built.

TABLE OF CONTENTS

Paul Campbell
Vice President of AI & ML Solutions, facil.ai

Open v Closed Loop Control: The Role of Set Point in Control Theory

Control Theory applies only to closed loop systems. The negative feedback loop is the critical element to controlling the output that is absent in Open Loop Control. Without the negative feedback loop the controller is independent from the output, which makes the word "control" a misnomer if one infers that precision equates to control. Open Loop Control with all its severe limitations is common among commercially available Building Automation Systems and represents an enormous opportunity for enterprising engineers and programmers and a fertile ground for data science.

Understanding the proper role of Set Point in Control Theory is critical. In Control Theory, Set Point is just a point of reference or parameter as opposed to data. To oversimplify, data coming into the controller of a Closed Loop Control system is compared to its Set Point parameter with the objective of determining if the output will be precise. If so, no changes will occur; if not, the controller will call for an adjustment to the parameter to create a precise output.

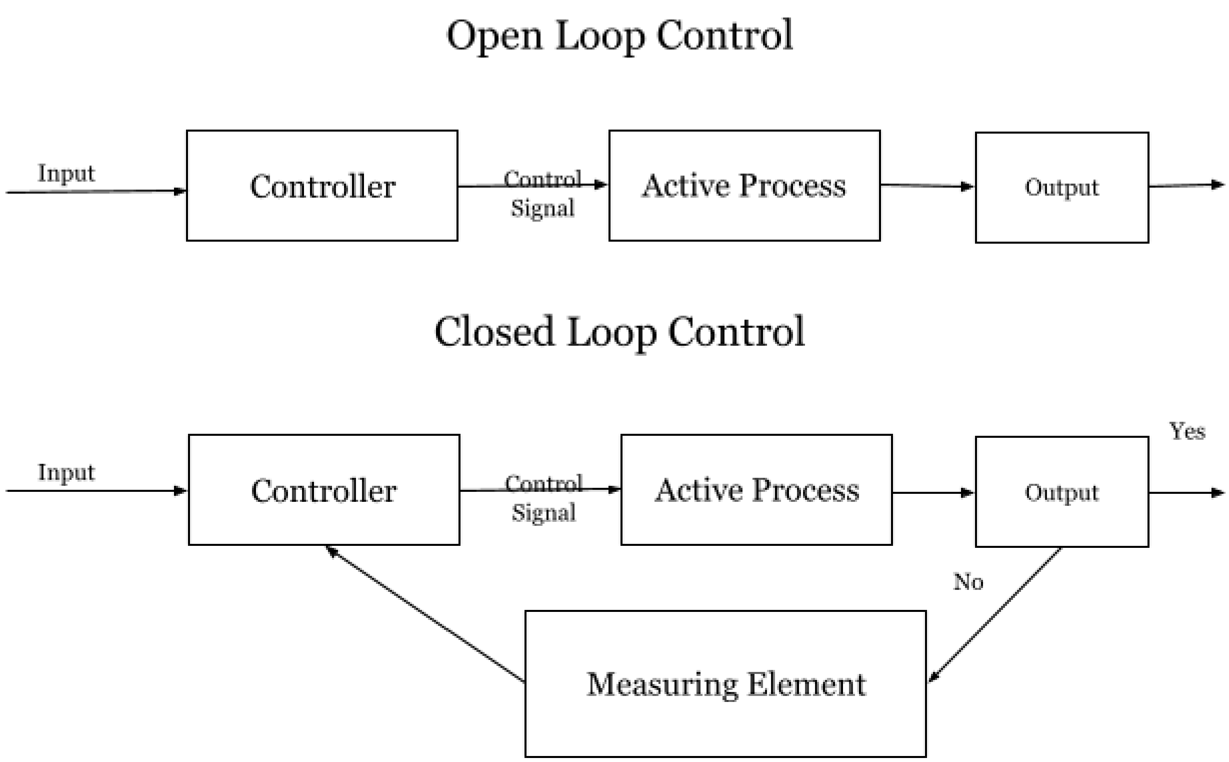

Paul Campbell
Vice President of AI & ML Solutions, facil.ai

Limitations of PID loops native to Closed Loop controllers

The standard vehicle for negative feedback native to Closed Loop control systems has been the PID (Proportional - Integral - Derivative) Loop. PID Loops reside in the controllers and need to be tuned to properly overcome bias (referred to as error in the controller industry). Without getting too much detail, PID Loops manage error by summating the corrective gains of the amplitude (Proportional) of error (KP), the cumulative error over time (Integral) or "area under the curve" (KI) and the rate (Derivative) of change (KD) as it drives error toward zero.

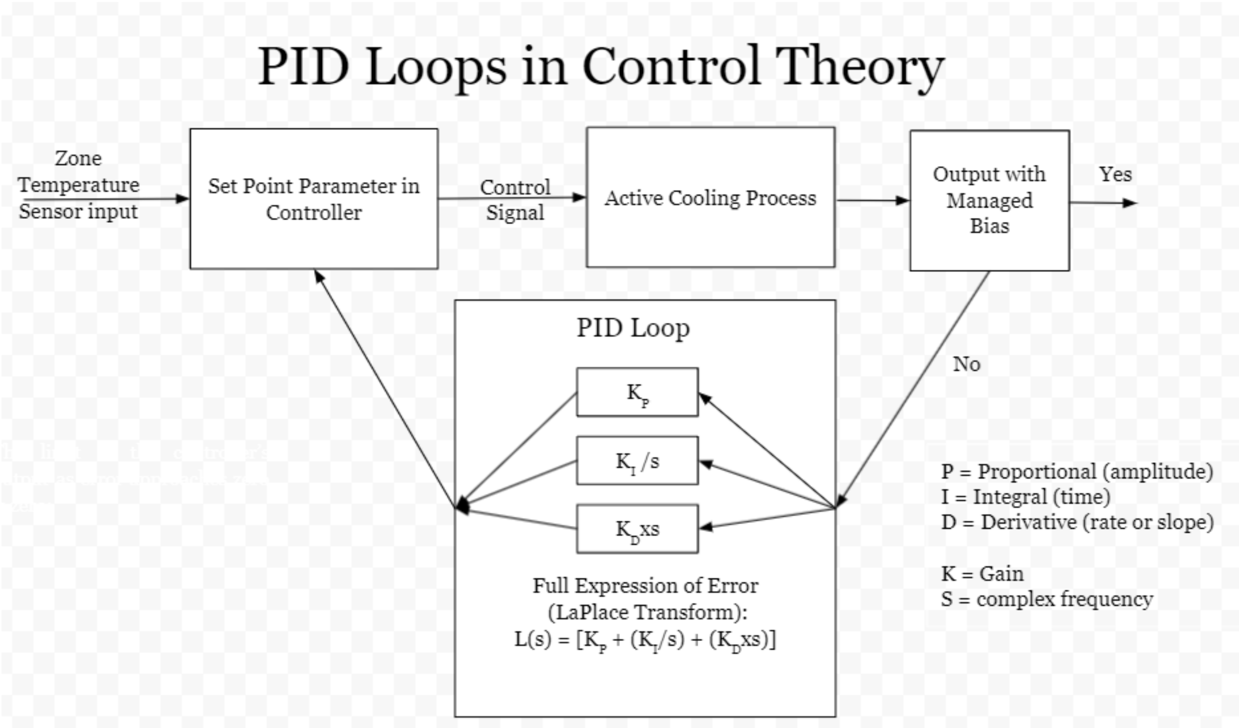

The Proportional, the Integral and the Derivative complete a multi-dimensional expression of error and that kind of exactitude in controlling extremely complex, mission-critical and potentially dangerous applications (such as nuclear power plant control rods, oil platforms, etc.) where error cannot be tolerated is imperative.

At first glance, it would seem a straightforward exercise to tune a PID Loop. However in practice, it is a highly technical function that does not lend itself to ease of use. Much has been written on this subject but even a quick search reveals how complex tuning can be:

Paul Campbell
Vice President of AI & ML Solutions, facil.ai

"PID tuning is a difficult problem, even though there are only (up to) three parameters and in principle is simple to describe, because it must satisfy complex criteria within the limitations of PID control. There are accordingly various methods for loop tuning, and more sophisticated techniques are the subject of patents; this section describes some traditional manual methods for loop tuning."

Source: https://en.wikipedia.org/wiki/PID_controller

In contrast, Closed Loop control of Building Automation Systems in retail buildings sector only focuses on the Proportional or amplitude of the error. In the retail building sector the Integral and the Derivative are set to zero for ease of use and the controller attends only to amplitude as the "total" expression of error. Therein lies the vulnerability to precision with this limited kind of PID Loop. Instantaneous amplitude without a wider time dimension or rate leaves only a snapshot upon which to tune, assuming the effort to tune is made.

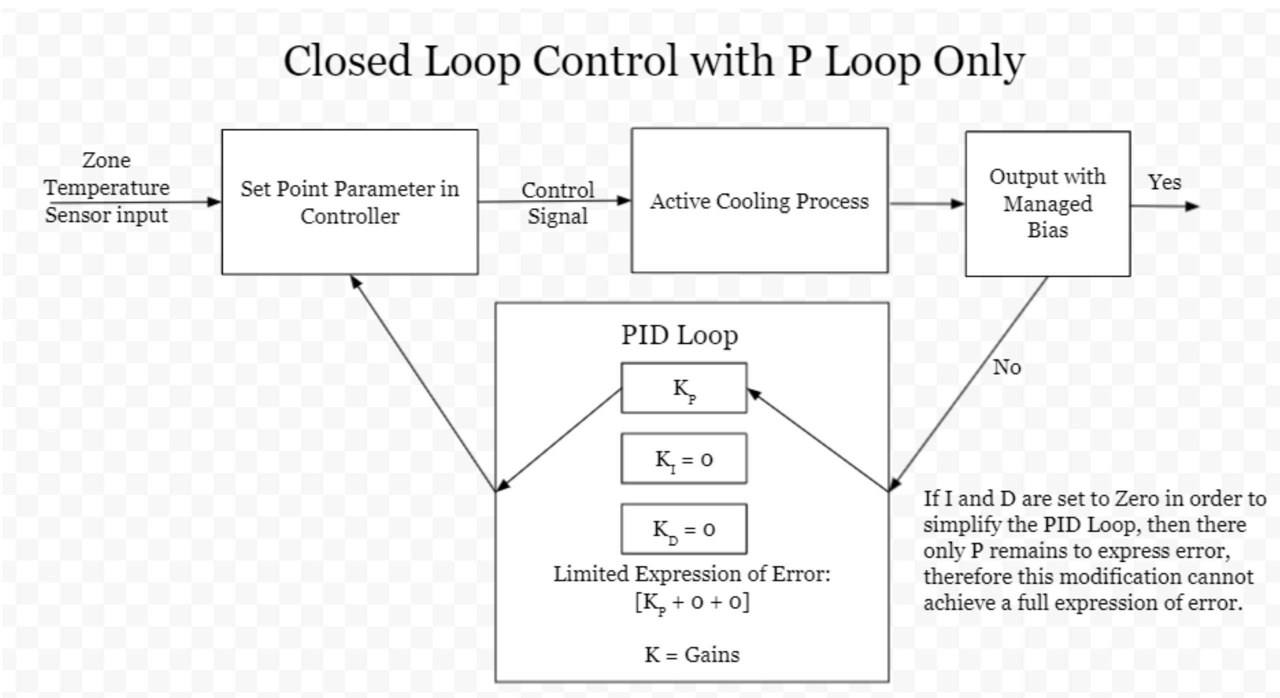

Paul Campbell
Vice President of AI & ML Solutions, facil.ai

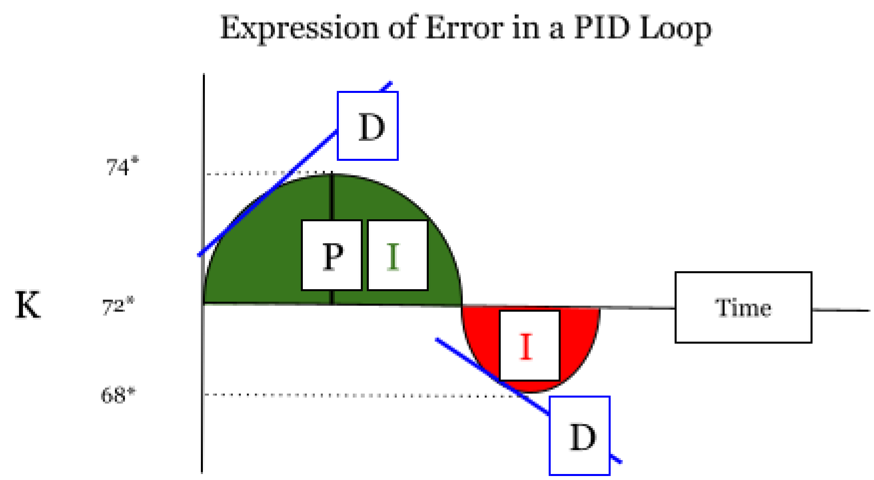

However, just having a snapshot in hand to tune amplitude with is only the beginning of the difficulty that lay ahead. As an example, if a 40 ton HVAC unit where only amplitude of error is corrected comes on at 74*F (2* F "deadband", set-point is 72), it blasts the zone with 55*F air and the zone temperature plummets down to 68*F almost instantaneously. The unit overshot its set point by 4* and demonstrated the classic signature of oscillating control.

There is a whole set of dependencies in this scenario that are external to the reach of the PID loop:

1. The HVAC unit must have been properly sized for the space during the design phase.
2. The Building Automation System controls application Engineer must have designed the system and programmed the controller properly.
3. The Building Automation System installation technician must have installed the controller properly.
4. The HVAC contractor must have balanced the room's air-flow properly.
5. After installation and air-balancing, the commissioning agent must have commissioned (tuned) the system properly.
6. The customer must operate and maintain proper configuration of the system properly.
7. The space requirements for cooling (heat load) must remain fixed and constant for the next 40-50 years (not likely) and must not ever deviate. This would prohibit LED lighting upgrades (which alters the heat load) and any tenant improvements, such as demising or re-demising walls, moving thermostats, etc.
8. Over its life-cycle, the system must be serviced and maintained properly, with the understanding that implementing adjustments must be within the confines of the "design intent".

Paul Campbell
Vice President of AI & ML Solutions, facil.ai

Distinction between data and parameters

Understanding the bias between the set point and data it pertains to is the first step in achieving adaptivity in energy management. The word bias is more misunderstood than perhaps it should be. Statisticians use this term routinely to describe differences between expectation and performance whereas accounting professionals implicitly understand bias but refer to it as variance. Since most folks are less exposed to statistics than accounting, the term bias becomes a bit intimidating. To eliminate confusion, we'll maintain the term bias.

In practice it is a simple exercise to understand bias and there are parallels all around us. In the realm of general accounting for instance one would need to understand why bias (or variance) between expense and budget is out of tolerance and take corrective measures. Occasionally, the budget model is not up to the task. Every model including the best models has error built into it. Bad models are poorly structured but otherwise good models can be compromised by either the use of poor variables or the misuse of good variables as inputs. The latter circumstance is the essence of this examination.

It takes an alert individual when looking at a dense graphic trend to discern data from parameters within a visualization and we would be well-served to make this vital distinction:

1. Data is empirical, meaning it is metered or measured and is continuous in nature, subject to changing value over time.
2. Parameters are artificial thresholds laid upon the data as one sheet of overhead projector transparency film would overlay another in order to gauge if performance conforms to a desired outcome. Parameters do not reside on the same plane as data.

We will simplify our demonstration by trending just two things - HVAC zone temperature (data) and cooling set point (parameter) - and study the bias with the objective of transforming that bias to an operational asset.

Transitioning from an Open Loop to a "Virtual Closed Loop" control

We propose in this examination, rather than procuring a portfolio of Closed Loop Control Building Automation Systems for retail building applications to replace disappointing Open Loop Control systems, an option to strongly consider is to create and install a Virtual Negative Feedback Loop. Through programming, control engineering knowledge and data science a negative feedback can be inserted into an Open Loop using appropriate communication protocols. Doing so requires strength of will as well as knowledge native to several technical disciplines but the payoff and the return on investment for the effort may be well worth it and more achievable than a capital-intensive procurement.

What follows are the challenges we encountered, the discoveries made and the results achieved.

Zone Temperature Bias Management: Managing to Zone Temperature rather than Set Point

In the trend below, cooling set point in the left hand side of the graph (pink box) had been set at 74*F during occupancy yet the HVAC unit had extreme difficulty reaching a 74*F zone temperature as a lower limit and hit an upper limit of 77*F routinely in mid occupancy. A simple manual adjustment to a 70*F cooling set point allowed the unit's zone temperature to stay within the bounds of a lower limit of 71*F and an upper limit of 74*F in the first 2.5 days after the adjustment was made (green box). This is the most fundamental step in managing to zone temperature.

Is corporate policy intended to maintain a 74*F occupied cooling setpoint or is corporate policy intended to maintain a zone temperature of 74*F? Too often policy - perhaps because of the data v. parameter confusion - enforces the parameter rather than the data.

You will note that a static manual adjustment of the occupied set point initially achieved an upper limit of 74*F zone temperature in the 2.5 days defined by the green box and effectively demonstrated it can overcome bias. However, the static manual adjustment of the occupied cooling setpoint also demonstrated its limitations in the 1 day of occupancy defined by the purple box. On that day, the unit achieved a zone temperature with upper limit of 72*F during occupancy, 2*F below our target zone temperature for occupancy (red arrow) so the manual adjustment to set point actually over-cooled the zone slightly. This trend signals an opportunity to develop a dynamic approach for adjusting the occupied cooling setpoint as conditions require to overcome the bias as well as maintain a tight tolerance to the 74*F target zone temperature.

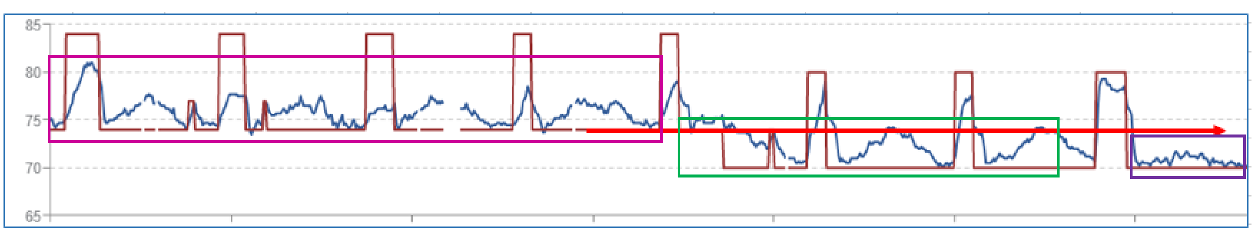

Eliminating arbitrariness in parameters - the gateway to a virtual closed loop

In the field of data science, an arbitrary parameter is commonly accepted as a reasonable starting place in the early stages of data examination as demonstrated by our manual occupied set point adjustment above.

Paul Campbell
Vice President of AI & ML Solutions, facil.ai

It is not however accepted as a destination for data examination. Both a "one-size-fits-all" corporate set point policy approach or a more variegated set point that relies on arbitrariness will fail, also as demonstrated by our manual occupied set point adjustment above. Arbitrariness over time and over large data sets in actively operating models within negative feedback loops is simply not sufficient and will prove unsatisfactory to our objective.

Through Machine Learning bias can be overcome and a precise zone temperature can be maintained. Moreover Machine Learning allows us to perpetually commission the set point to any value required to achieve a precise zone temperature.

$$\lim_{S \to Z} \mathbf{f(X) = 0}$$

This mathematical expression of the limit of a function illustrates nicely what we are trying to achieve and it reads: the limit of the controller's output as the zone temperature bias approaches zero (S-->Z in the expression) is zero.

In order to drive our bias toward zero and sustain it we must as a matter of course eliminate arbitrariness. We can never achieve precision and accept arbitrariness in a parameter. A parameter absolutely must be informed by the data. Though preached and practiced in data science, the human quality of imposing one's will on parameters signals a competing vision of control and it is an arduous task to relieve someone of the emotional need to register their cognitive bias on a parameter even if an important strategic objective weighs in the balance. Unyielding persistence in this endeavor is vital.

Reverting the roles: Zone Temperature as control input and Set Point as Control Output

"An 'essential variable' is defined as "a variable that has to be kept within assigned limits to achieve a particular goal"

> Jan Achterbergh, Dirk Vriens (2010). "§2.3 Cybernetics: Effective methods for the control of complex systems". Organizations: Social Systems Conducting Experiment

The set point is the control output (or essential variable per Achterbergh & Vriens) and the zone temperature is the control input. At first glance this may seem to be reversed but it is actually the intended roles for zone temperature and set point in control theory. In closed loop control, the set point needs to be set at such values (which are subject to change based on conditions like weather or seasons) as to overcome the bias specific to that zone in order to meet the target or objective zone temperature. That requires the will to study the unit's empirical input data against the output parameters and to overcome

Paul Campbell
Vice President of AI & ML Solutions, facil.ai

the thoughtlessness of holding set point at a specific value regardless of the bias, assuming that bias is even acknowledged.

There are common precedents elsewhere. Cruise Control in road vehicles for instance will accelerate or decelerate a vehicle in order to maintain a 40 mph target speed depending on which side of the hill the vehicle is on. Its 40 mph target speed is the control input and the rate of acceleration is the control output (or essential variable). As kids we rode our 10-speed bicycles all over and intuitively maintained speed by shifting to a lower gear going uphill or a higher gear going downhill. Nobody argues about which is the control input or control output in these two examples but strangely maintaining an optimal HVAC zone temperature is not so easily understood.

Open Loop v Virtual Closed Loop Control

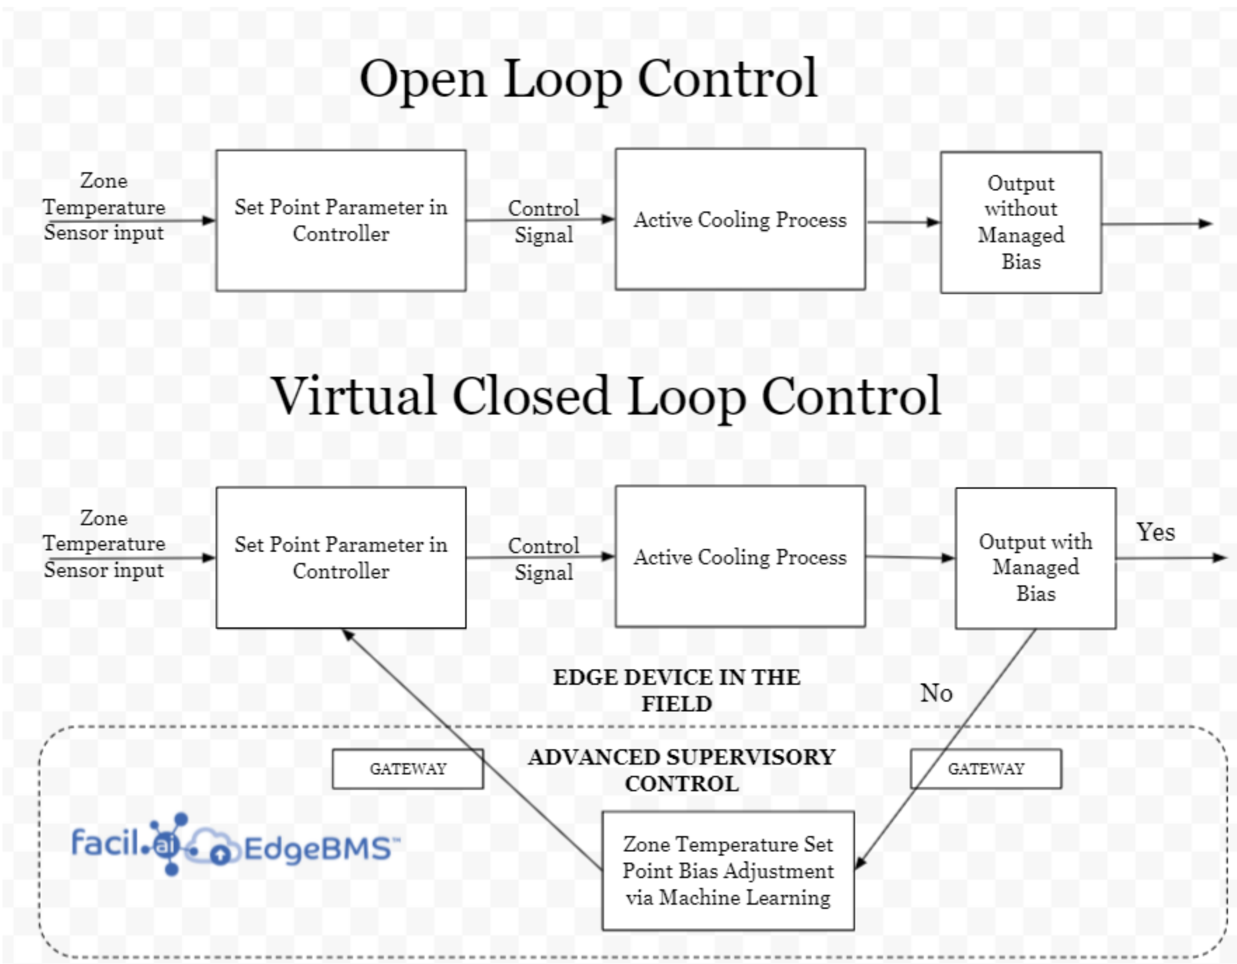

Paul Campbell
Vice President of AI & ML Solutions, facil.ai

"Give me a lever long enough and a fulcrum on which to place it, and I shall move the world," per Archimedes. In this instance of a Virtual Closed Loop, the lever would be the zone temperature's bias v set point, the fulcrum would be machine learning and the world would be the set point. The zone temperature as well as the zone's set point bias inform the set point continually to produce a precise output and drive bias toward zero in perpetuity.

Granted our objective here is a bit less ambitious than Archimedes' but his point is well taken. In fact our example here is univariate but only for ease of demonstration. Imagine a controller exercising multivariate levers to produce a precisely optimized total environment of temperature, humidity, pressure, airflow, CO2 and light. The lever in our examination would not have to be "that long " since the data to optimize the parameters are already available in most open loop control Building Automation Systems. All that remains is to understand each datum's bias against its set point parameter. Achieving that understanding is a matter of will and to invoke the Archimedes analogy one final time, one's will represents the length of the lever. Understanding bias can be done if it is willed to be done.

Perpetual Commissioning of HVAC assets through Machine Learning

Once biases are looped back into the controller, outputs become more precise. However, precision is maintained only if biases are continually managed and through the advent of Machine Learning precision can be the perpetual result rather than an unattainable ideal. The inherent velocity of Machine Learning as the controller's negative feedback vehicle elevates our biases from a simple Closed Loop control input to immediately actionable business intelligence and can in a manner of speaking be monetized. Taken that far these biases don't just inform the controller to be more precise but become the currency of untolled millions of automated decisions annually. As such all HVAC units in an enterprise with a Virtual Closed Loop control using machine learning negative feedback to manage biases represent a fleet of decision engines.

These virtual decision engines are strategic in nature since we are managing these biases instantaneously across the enterprise and not through some audit regime, which can only be tactical. Because the negative feedback loop is external to the controller, this gives Virtual Closed Loop control a decided advantage over native Closed Loop control in that any from among a suite of strategic decisions can be applied instantaneously as well as predictively. Automating decisions that are strategic, wide ranging, low amplitude, high volume and predictive are not manageable in any other way.

Psychometrics of Comfort

There are all sorts of biases and data affecting HVAC performance. Our ultimate objective is not to drive performance into the comfort zone illustrated in the psychometrics chart below as that would be reactive. Rather our goal is pre-emptively keeping performance bound there efficiently through bias management by interposing our Virtual Closed Loop control. Without the negative feedback based on Machine Learning we have no effective means to keep performance bound within the comfort zone and will not escape being reactive.

We've focused upon volume and velocity but we only briefly touched upon variety. Our examination has been limited to only zone temperature's bias v its set point. However just by introducing outside air temperature and relative humidity along with the various set points that attend to them as additional data and parameters to consider modeling, a full suite of bias management Machine Learning algorithms in a feedback loop affecting other functions within a HVAC unit's order of operations (such as economizer control) quickly intensifies the demands placed upon Virtual Closed Loop control. The implications of bias management across an enterprise is one of geometric progression that is orders of magnitude larger than bias management via audit regime. The primary implication is that Virtual Closed Loop control requires a strategic vision to project it as a force multiplier.

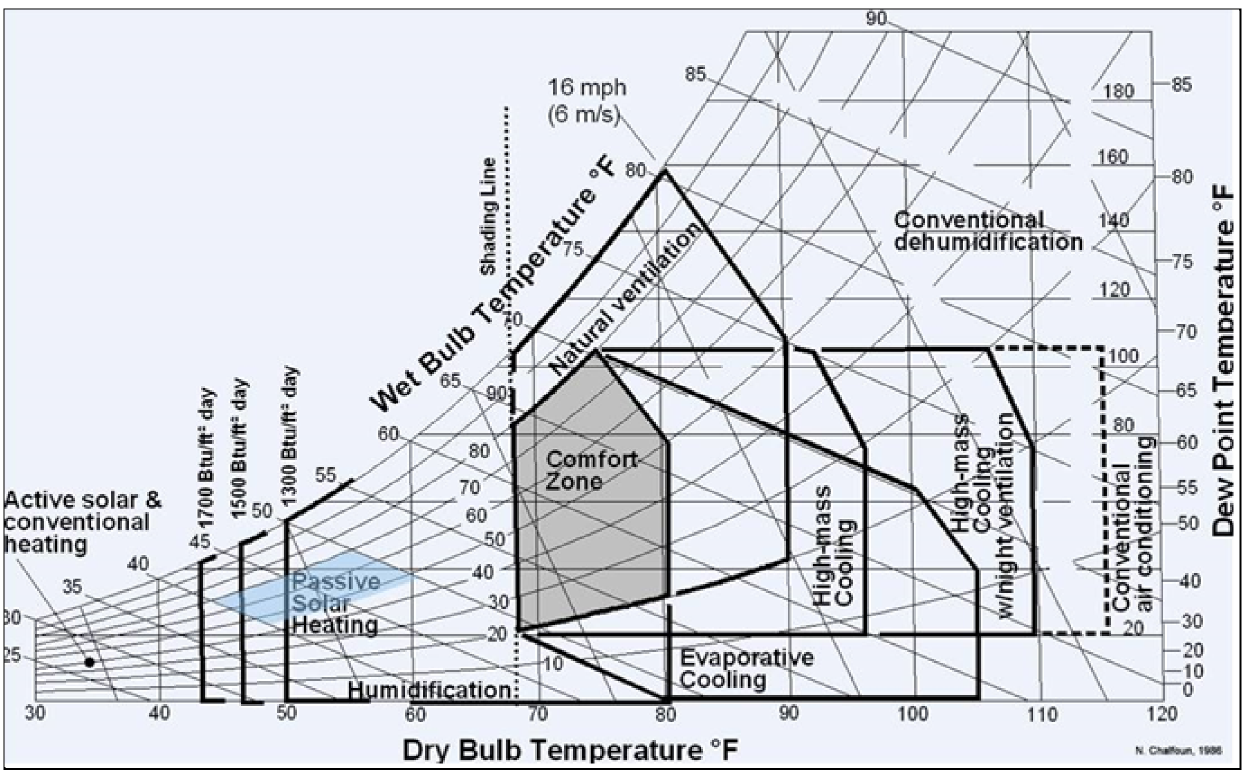

Horizontal v Vertical Energy Management

Once the elimination of bias and arbitrariness has been achieved and precise control established across an enterprise's HVAC fleet, the use of energy can then be optimized. What follows is an overview of Adaptive

Energy Management as a strategic framework. There is a distinction that needs to be understood once in the orbit of demand side management, defined here as Vertical Energy Management and Horizontal Energy Management.

Demand Response is often the first thing that comes to mind when the active control of energy use is introduced in conversation. It is not a program that is designed with customer energy efficiency in mind but rather designed by the utility or Independent System Operator (ISO) to relieve stress and preserve the operational integrity of the electrical distribution system. The utility or ISO incentivizes the customer to reduce load measurably and will offer rebates to enable that capacity reduction, expressed in kW. The shorter the notification window the greater the incentive. In order to be effective in a demand response scenario, the signal from the utility or ISO would have to be recognized and the eligible sites batched and the chosen tactic deployed within the response window. These features – primarily the tactical enable status default position being set to OFF until the ISO signal is received – comprise the essence of Vertical Energy Management where loads are taken down hard and noticeably.

Contrast this with an enterprise's optimization program where the program's enable status default position is set to ON. Such a program is always enabled and goes in and out of active status as local zone or site criteria are reached then satisfied. The optimization program's cumulative active time is the element that defines this as an example of Horizontal Vector Energy Management and aligns more closely with permanent load reduction. This program would be measured in both in terms of kW and kWh.

Tactics v Strategy in Adaptive Energy Management

Too often it is said that "We have a Demand Response strategy", which is a miscast of the term strategy. It would be better said that "We have a Demand Response tactic. It is one of the various tactics in our Adaptive Energy Management Strategy." A Demand Response tactic is just one arrow in the quiver.

Simply having the Building Automation System to execute a tactic is not being strategic. Under Adaptive Energy Management, tactics are matched, sequenced and executed dynamically through machine learning via Advanced Supervisory Control just as bias management is. Among other things, orders of precedence must be established to manage tactical traffic. For example if we have five tactics enabled and conditions are met that would activate all of them at once, the order of precedence needs to govern which of the five tactics gets approved to proceed and which remain in standby and by extension, which positions (X thru 5, assuming that more than one tactic can be active so long as there is no conflict) the remaining tactics occupy in the queue. Tactics working at cross purposes or competing for resource support would be an example of a failed strategy. Tactics working in concert - possibly concurrently either within the same subset or in parallel subsets, possibly working consecutively within a subset - is an example of a successful strategy.

Paul Campbell
Vice President of AI & ML Solutions, facil.ai

The graph below illustrates the differences between Vertical Energy Management and Horizontal Energy Management as well as demonstrates the space tactics occupy within an Adaptive Energy Management strategy. The hour of the day is on the X-axis and average hourly demand expressed in kW is on the Y-axis and the data are a representation of a cumulative monthly data segment for one retail site.

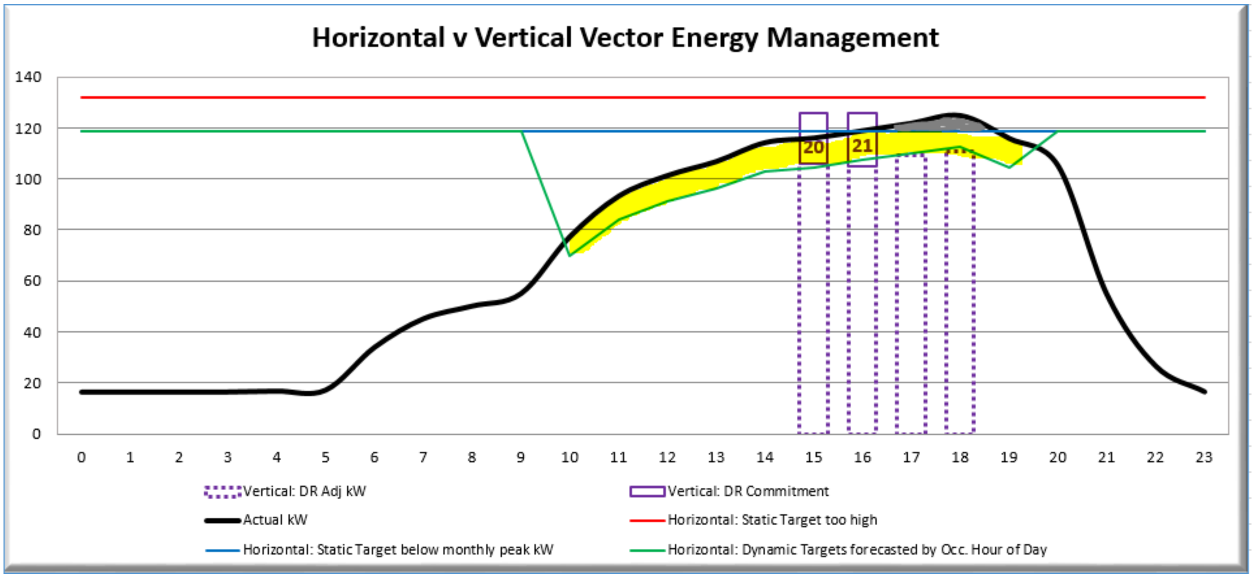

Highlights of note:

1. The black line represents the average hourly actual recorded demand.
2. The red line is a representation of static demand target that some Building Automation Systems might use as a demand limiting set point. This is an example of a static target that was set too high to be useful to limit demand.
3. The blue line is a representation of a static target set to 90% of the peak monthly demand to show that limited avoidance can be achieved, but only on days that approach the static target. The demand limiting is inhibited by this low rung approach.
4. The green line is a representation of a forecasted hourly target produced by machine learning that undercuts the actual demand during specified hours or under specified conditions. In this demonstration the forecasted target was active between 10:00 AM and 8:00 PM. The dynamic target demand here shows how such a tactic might express itself but from day to day may have a different signature.
5. The purple outlined stacked bars represent the adjusted demand (dotted outline) and the demand reduction commitment (solid outline) in a fictitious Demand Response event covering a 3:00 PM to 6:00 PM timeframe. Commitments that are not achieved do not receive the incentive and in fact may be subject to penalties in some programs. Achieved and failed commitments are demonstrated here.

Paul Campbell
Vice President of AI & ML Solutions, facil.ai

The intent of this visual is to emphasize that we need not settle for the choice of trimming off only the very top of the monthly peak on the one hand or to wait idly by for a Demand Response event to be called nor do we have to manually re-commission demand limiting targets as buildings' mechanical systems composition evolves over time. Even as buildings become totally 100% optimized there will always be a peak demand to be managed. The emphasis of Adaptive Energy Management is decidedly on the volume ultimately driven by time, i.e. our X-axis.

Notwithstanding, the true graphic impact that an Adaptive Energy Management strategy produces does have a vertical as well as a horizontal signature and the two-dimensional area such as the yellow shaded arc in the graph demonstrates the impact of just one tactic. The aggregate of a series of dimensional areas like this for all tactics that have been activated demonstrates the full impact of a vibrant Adaptive Energy Management strategy.

Building as a Battery

A closer examination of the Horizontal v Vertical Energy Management graph above shows that while it represents a 24 hour period the tactics within it congregate only on the occupied part of the day as demand rises. This demonstrates misguided thinking and it was done quite deliberately to emphasize a point, namely that more progressive thinking would be to manage energy use horizontally over the entire 24 hour range. It may surprise some that the trimming of the afternoon hour peaks will occur organically as a matter of course if we apply the proper tactical approach during the overnight hours within our overall Adaptive Energy Management strategy.

Too often the instantaneous use of electricity is mislaid as the template for thermal capacity where the use of energy and the thermal benefit are exchanged only in the moment as though they transact at a cash register. Heat does not dissipate in synchrony with the use of the specific volume of electricity that supports mechanical cooling operations nor does heat dissipate at a constant rate over a 24 hour period. Accordingly, the timestamp range when one meters electricity for mechanical cooling purposes may be hours prior to the timestamp range when the thermal output provided by that volume of electricity is being drawn down.

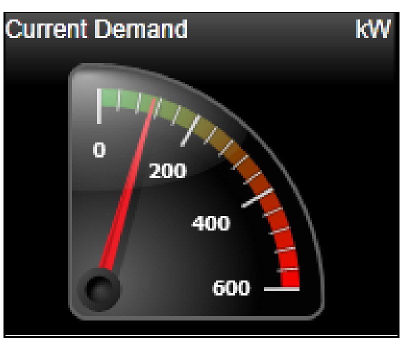

Paul Campbell
Vice President of AI & ML Solutions, facil.ai

The graphic above is a nice illustration of the asynchrony of the use of electricity and heat dissipation. If we saturate a retail space overnight, the thermal output will be stored until the media storing the cooled air (everything possible from retail product, to internal fixturing to structural elements of the building) seeks osmosis and draws in warmer air as the thermal heat gain from the sun and occupancy intensify until the cooled air is depleted. Hence the building acts like a battery: it gets recharged during the night and it is depleted during the day. Much has been published on this subject over the years. Using Thermal Mass and Night Purging techniques are well established but perhaps underappreciated.

“Opportunities for reducing operating costs through use of building thermal mass for cooling are due to four effects: reduction in demand costs, use of low cost off-peak electrical energy, reduced mechanical cooling resulting from the use of cool nighttime air for ventilation precooling, and improved mechanical cooling efficiency due to increased operation at more favorable part-load and ambient conditions. However, these benefits must be balanced with the increase in total cooling requirement that occurs with precooling of the thermal mass. Therefore, the savings associated with load shifting and demand reductions depend upon both the method of control and the specific application.”

> Source: Load Control Using Building Thermal Mass by James E. Braun,
> Ray W. Herrick Laboratories, School of Mechanical Engineering, Purdue University,
> ASME Journal of Solar Engineering, Volume 125 August 2003

It should be noted that overnight saturation does not stand in isolation. Dr. Braun’s ASME article from 2003 also emphasizes the need to shift load as required and seek a balance and to make this 24-hour approach viable.

The graphic below illustrates this 24-hour nature of Building as a Battery. As before, time occupies the X-axis and demand (kW) occupies the Y-axis. Along the X-axis are labels describing the tactical phases within our Adaptive Energy Management strategy: Saturate, Coast, Shift and Recover. The blue trend line shows a typical daily demand curve that peaks in the afternoon hours and mirrors the grey trend line plotting the Outside Air Temperature (OA_T). The objective of Building as a Battery is to re-distribute this pattern strategically. The gold trend line shows the type of demand curve that we intend to achieve where demand increases to support saturation then drops off significantly as HVAC coasts while heat dissipates.

Paul Campbell
Vice President of AI & ML Solutions, facil.ai

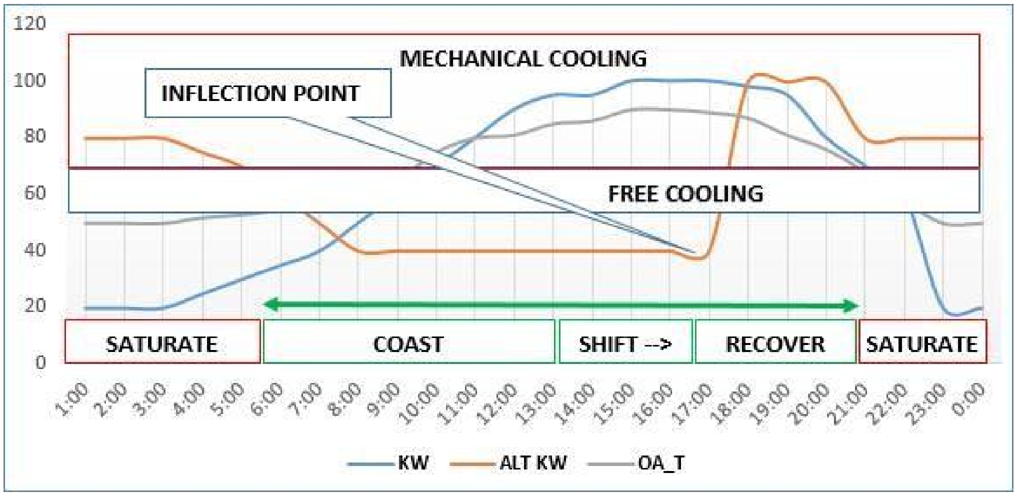

Critically, the need to shift load -i.e. limit demand - in many cases would need to be deployed as a tactic of the overall strategy. The rate of heat dissipation will vary widely based on geographic location and season so the fleet of "batteries" of an enterprise may approach depletion at different timestamps. It screams for an analytics solution to accurately predict when - or if - to begin shifting load to suppress the demand curve as far into the late afternoon - early evening hours as possible until released to go into the recovery phase, most probably determined by an upper limit for zone temperature.

It is important to note one should not infer that saturation means 100% mechanical cooling as when outside air temperature and relative humidity are favorable fan operation alone can inject a significant quantity of the saturation by bringing in outside air. Also the presence of a solar array or onsite physical battery storage assets can be brought to bear on the establishing the boundaries of the coast, shift & recovery phases of Adaptive Energy Management.

The validation of the Building as a Battery is in determining the energy reduction by calculating the difference between the areas under the curve of the blue and gold demand trends below. If the area under the blue trend is greater than the area under the gold trend, energy reduction has been achieved in this 24-hour period. The validation of Adaptive Energy Management as a strategy is in configuring the most effective sequence, selection and duration of the suite of tactics behind the saturation, coast, shift and recovery phases, thereby maximizing the impact of Building as a Battery across an enterprise.