control loops, it can provide elaborate temperature profiles, imbedded analyzer values, valve positions, and a host of other seemingly important and relevant data. And, as you suspect, the operator has too much data and perhaps too little information. User-centered design principles suggest that much less data are needed; much more context and interpretation is required. Moreover, the method of presentation itself can either obscure understanding or provide just the necessary context. A simple way to do this would be to provide graphs with goal marks located at appropriate levels on the same scale instead of tabular data with numerical targets. Recall the directness and clarity of the deviation diagram.

Implementability

You might suppose that this part of the user-centered design would be the least offended requirement. Operator workstations have been designed and built around providing all of the necessary "handles" for operators to manipulate the things operators do during their shift. Let me suggest that such is not always the case. All too frequently an operator has to "hunt around" for the right valve to check or observe one process value while he is required to manipulate another process value—but neither can be found on the same display. And there are the situations where the operator is required to shut down part of a process by gradually reducing the setpoints for a number of controllers in strict unison to avoid thermal or other operational stresses. Not only are the controllers on different displays, but also many of the variables needed to check the progress of the work are not colocated. The operator has to do a lot of shifting around and remembering.

The message for achieving good design is to ensure that routine operations and strategically important operations can be carried out in ways that ensure success and provide for a minimum of operational risk.

Unified Feel

When we ask for equipment to possess a unified feel, we are suggesting that once operators learn how to interact with some of the equipment, they will know how to interact with all the equipment. A very good example of this requirement can be found by approaching any conventional personal computer. Whether it be run under a Windows system, a UNIX system, an Apple Mac system, or any of a number of others, it would be obviously clear to the user how the keyboard is used and what should happen when a mouse or other pointing device is moved. We know what should happen when one "right clicks" the pointing device. This is also to say that the activities that are used by employing the keyboard and pointing devices are also going to be interacted with according to an expected feel.

12.5 OUR BIOLOGICAL CLOCK

We humans are designed to be generally compatible to our surroundings.

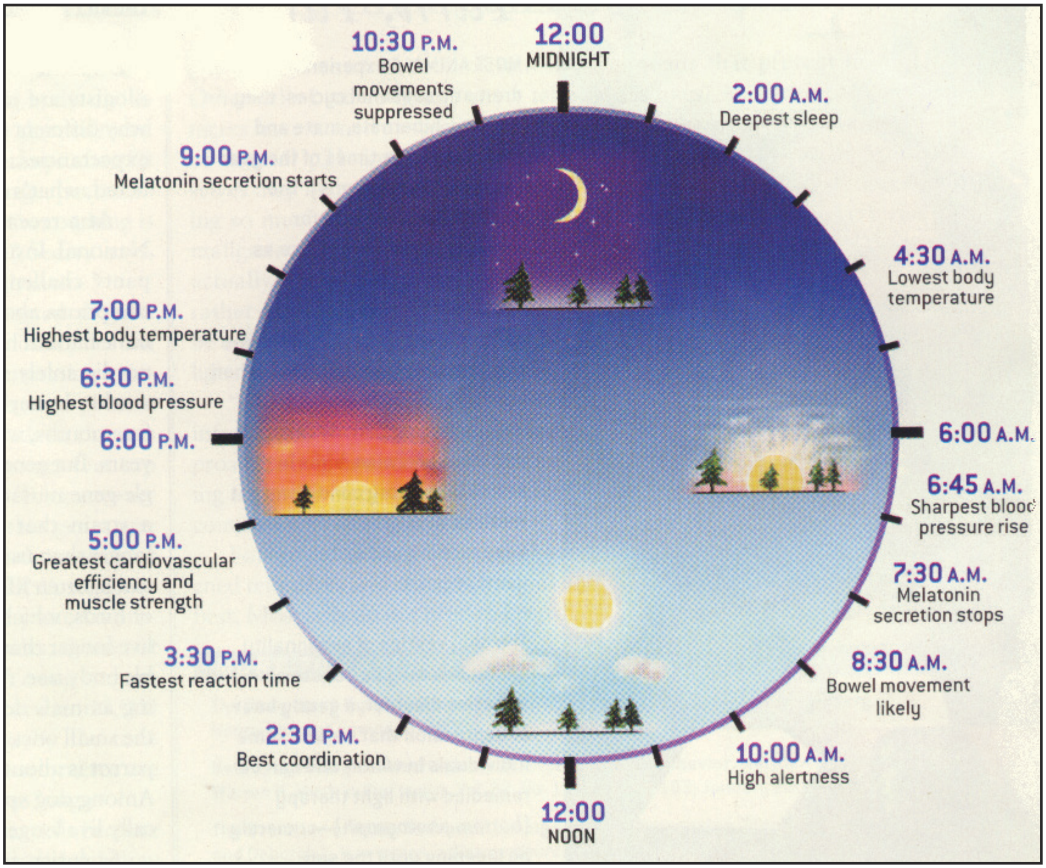

Figure 12.5.1. Circadian rhythms; our internal biological clock

If we trust Darwin, we believe that those of us who adapted best to our environment have survived. Adapting means being in harmony. We get enough sleep, have the ability to reason quickly and efficiently, and expect the dependability of our bodily engine to process foods and rid of waste. One of the most observable adaptations is the synchrony of our bodies to the daily cycle of life. For all but the briefest of moments in human history, the sun has governed our activities. It is little wonder that our inner clocks are so linked to this powerful diurnal rhythm.[6]

Alas, the demands of a modern manufacturing society, facilitated by the invention of ways to produce artificial light, have made us necessary and able to shift our personal clocks different from the natural sun clock. Doing so has exacted a sometimes-tragic price for such rashness. Yes, we can move our inner clocks. Certain parts of the world have needed to do so because of the very large variability of this natural clock—witness what goes on in the highest latitudes. And take note of their accommodation difficulties.

These natural variations, however, occur very slowly when compared to the artificial, forced shift-time changes of the modern workforce.

It has been observed that a healthy adult person takes from 1 to 3 days to accommodate to each hour time shift of their inner clock. And this accommodation is often aided by the shifting of the sun clock at the same time—as in jet plane travel. On the other hand, a totally artificial shifting of the inner clock has no such lucky assistance. Little wonder that those shifting inner clocks have difficulty performing for days and days. And little wonder that such difficulty may affect their stress performance response. Accidents result. Situation awareness must be able to reach even those individuals.

12.6 OTHER OPERATOR SUPPORT ISSUES

Understanding the limitations of operator awareness is another vital resource in our kit bag of operator support. What is it that we see? What we sometimes see might not always be what is there.

Intent Recognition

Let's suppose for a moment that you were looking over the shoulder of an operator at his station. And let's suppose that the operator was working on an important task, one that you understood very well. And let's further suppose that the operator was less familiar than you and somehow faltered in making his way down an action path toward managing a situation. Assuming that your presence was for assisting and not for testing, you would likely ask the operator what he was trying to do. And once you assured yourself that it was the task you understood, you might suggest that he do "thus and so" to get the job done. We call this help.

Continuing with this line of thought, suppose that instead of you, there was an "agent" (software and such) watching the operator at his duties. This agent determined, by whatever means, that it strongly suspected the operator was trying to perform a specific task. It would be extremely helpful if the agent could query the operator to determine if the operator was, in fact, trying to perform the detected task. Having affirmative confirmation, the agent might also suggest ways for the operator to do the intended task or, in the alternative, and with permission, do the task for the operator. This is intent recognition.

Are you convinced that you know how it works but just as sure that such assistance is not very useful? Let me assure you that this is far from the case. Developing hundreds of intent recognition scenarios for routine tasks would be hardly beneficial. But developing a few scenarios for tasks or situations that the operator must do well but is often called on to do under high stress might just be the ticket to better and safer operations. Part of situation awareness technology would be to identify just those situations where intent recognition and assistance could be critical. The other part would be to provide this assistance when needed. Implementation of intent recognition and assistance normally follows the operator consent construction introduced in chapter 8.

Operator Vigilance

Operators have relatively long periods of time on duty. Eight hours is usual. Twelve hours is increasingly popular. To make matters worse, those 8 hours do not occur during the same clock time of the day for each day of the working month. We know from the study of the effects of attempts at changes in human circadian rhythms (section 12.5) that what our bodies are trying to do is strongly associated with where our inner clock is at the moment, regardless of the actual daily clock time. If they are mismatched, we are not at our best. However, clocks being out of match are not the only reasons for operators not being at their best. Personal health, home matters, and exciting plans for the future all play a part in occupying and distracting our minds. Distracted minds pay little attention to tasks, important or otherwise.

Operator vigilance is a process by which a specialized "agent" (software or otherwise) monitors operators to determine the degree of attention they are able to pay or appear to be paying to their jobs at hand. How this monitoring is done is not a subject for this book. Where diligent operator attention is required, the plant environment must ensure that it happens. This might be done without any specialized monitoring through the use of a program of careful attention to varying operator duties. It might be done by varying control room ambience or by construction of efficient reminders of routine tasks or by other comfortable and effective means. However it is done, it must be done. It is unfortunate that most control room protocols do this only as a by-product of coffee, operator jokes on one another, and other idle distractions attempting to break the heavy mantle of monotony.

To Push or to Pull

In information processing language, "push" means that a possessor of information actively provides that information to a party of expected need. In effect, they push the information to them. "Pull" means that the party that thinks itself in need must find the possessor of that information and grab it. Again, in effect, pull it (by sharing, not taking possession) from where it is and toward oneself. Both modes have their uses and problems. For example, alarms are always a "push." Operators are immediately notified—they never have to request. State changes for equipment do not push. Yes, certainly, if the operator was looking at information that normally includes the current state, then the current state would update to reflect that change. But otherwise, state changes are generally not "push" events. This section will not cover the push/pull subject in much more detail.

Before we leave the push/pull discussion, there is one misuse that has been suggested by responsible operator station designers. They design so that operators get a chance to see everything they should see by pushing it all to them. The normal way it is pushed is to design a series of displays that rotate on a time basis between several screens. That way, based on a clock timer, each screen appears according to schedule so the operator can view it. A common use comes to mind: a security guard stationed in front of several monitors; each monitor is split between perhaps four cameras. For high-security operations, no rotating screens are used or operators are changed very frequently—say every half hour. The normal approach, however, is to switch the four camera views every 5, 10, or 15 seconds in a continuous rotational loop. In short order, all the cameras are displayed for the guard to see. It is well known that in such situations, the security guard rapidly succumbs to the events of screen changes but does not pay much attention to their content. From a human factors or user-centered design point of view, what is wrong with this situation? In short, everything. In long, there are several interesting aspects to point out:

- Display movement distracts from all other tasks and generally reduces the attention span for everything.
- Forced rotation of screen images actively prevents operator focus on anything but the grossest of detail, or they blank entirely and see nothing at all.
- Display movement (the changing of camera images) reduces the ability to discern image movement within a given screen, save the grossest.
- Forced rotation requires actively interfering in order to return to a previous mode or even stay focused on a desired one. At the same time, other views might contain important information, but since viewing is the only way to know that, it will be missed. Nefarious individuals often take advantage of this phenomenon by creating a distraction to draw attention away while the bad stuff takes place somewhere else, off view.

12.7 OPERATOR DISPLAYS

The video display unit (VDU) would seem to represent a significant step forward in human-machine interaction. Actually, it was a step backward. But not in the way you might think. Only recently have humans been faced with interacting with "things" that were not real and seen with their own eyes. Herds of animals did not need a video display to be seen. Crops during a drought did not require them either; neither did sailing ships, managing armies, nor cooking soup in a pot over the fire. It is when the stuff that needed to be managed was composed of things, real or virtual, that cannot be seen with eyes directly that we are required to employ surrogate machines to visualize them.

Unfortunately, it was the evolution of displays, not their intrinsic faults or limitations, that led us down the wrong early paths. The earliest displays were able to show lots of (weakly formatted) text, and later on, limited graphics. These early video displays for process control started from where the predecessor hardware equipment left off. They first mimicked control stations as faceplates on the VDUs. They were arranged in rows with most of the display agents and handles that were found on their physical analog counterpart it so quickly replaced. What was gained, of course, was the ability to place many times more displays easier and cheaper on video units than on metal control walls.

Once the faceplate barrier was broken, so to speak, the world of graphic design opened up. As soon as color made the scene, all the primary PCS manufacturers started a race to see who could use the more appealing and flashy colors to preen in front of prospective buyers. I do not actually know who won that race, but I do know that the purchasers were the collective losers. What was lost was the objective of what the video display should do and how best to do it. Now, many decades later, we know a bit more. This chapter will help you to understand.

Physical Display Architecture

Ironically, it was the very flexibility of the early VDUs that contributed to our failure to appreciate the need for a physical architectural arrangement that incorporated multiple units. Add the high cost of the proprietary VDUs to the equation and rarely would one find more than two or three units per operator station. Early on, there was an inherent conflict between flexibility and navigational difficulty. Operators became overburdened with their primary task: keeping an eye on things. In those days, keeping an eye on things required constantly shifting displays to find all the relevant states and situations. Only now do we understand how difficult that burden was. Being difficult, it was often done incompletely or ineffectively. It should not be surprising that alarms were (over) used to be watchful servants for the operator. It was a natural way to automatically be able to keep an eye on his plant. Everything that might be abnormal would cause an appropriate alarm on operational displays.

We know that more screens are as necessary as more displays. Figure 12.7.1 shows a recommended architecture. This architecture, first proposed in the late 1980s, has very important structural aspects. First, we have an expected location for all necessary

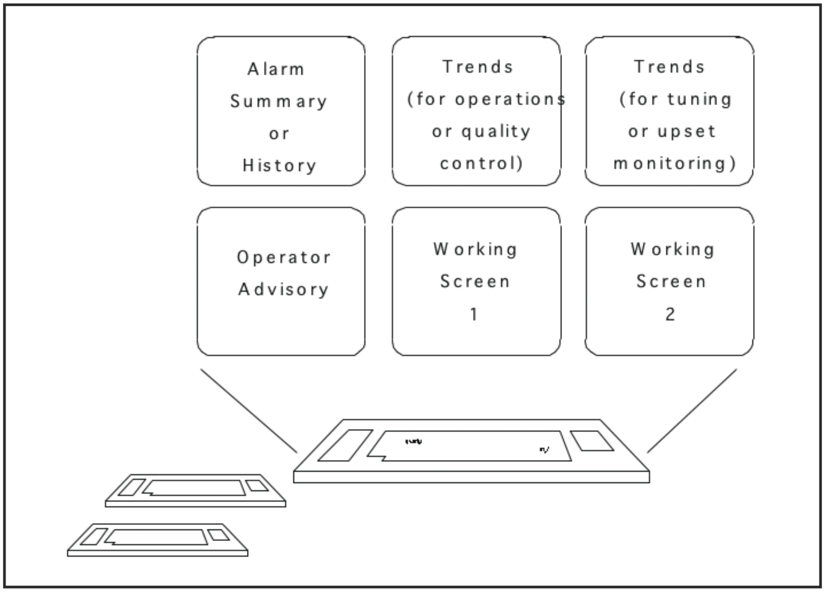

Figure 12.7.1. Suggested physical display arrangement

information needed for general operations. Second, there are enough screens so that most tasks, including monitoring, can be viewed at the same time without requiring switching displays on a screen. The locations are arranged so that related information is naturally located.

The choice of which resident displays to locate on which screens is made so that those requiring close interactions are center and lower. Note the two working screens (1 and 2) are at the bottom center and right. Here is where the operator would be monitoring specific control points and related variables. Or he may be intervening to manage an abnormal situation or other event, for example, by altering a controller setpoint or moving a valve. Close at hand (bottom row, to the left) would be other advisories to assist him. For example, the screen could provide assistance, as needed to augment the operator's current activities. This assistance would show procedures, provide relevant background analytical data, alarm diagnostic assistance data, and the like.

The displays that provide more global information are located above. These displays, as a package, complement the operator's role of observing and managing. To the top left would be the screen dedicated to the alarm system. The screen above center provides overview information on how well the process is working. This aids the operator working to ensure the plant does not go astray. He may also be working on process improvements. The displays that support improvements are located on the upper-right screen. Again, we have a hierarchy of detailed displays:

- Task displays within easy reach in the front and sides
- Overview information above, within easy sight

Modern Displays

Six screens that were so radical and extravagant 20 years ago are now rather commonplace. Local control rooms, once located in satellite areas with each area separate from the other, have been replaced with a single centrally located one in a protected area remote from the actual production equipment. Figures 12.7.2 and 12.7.3 show photographs of a typical central control room at a modern petroleum refinery. There are six operator areas within view and a seventh (out of view) for engineering and maintenance.

There is also a growing tendency to group closely linked plants with separate operators together into a structured super operator station. Figure 12.7.4 illustrates one. Notice that this arrangement also includes the increasingly common addition of large overview screens located high above the operator areas so as to be visible to others. This permits important information to be shared among operators, so vital for crisis management.

Also present are optional special-purpose hard-wired alarms and the ever-present collection of emergency shutdown switches and interlock management hardware.

A view of the expanded operator area is shown in Figure 12.7.5. It is easy to appreciate how the screen layouts and presence of large overhead screens can facilitate operators viewing and sharing information and collaboratively working problems.

Figure 12.7.2. Typical modern operator area

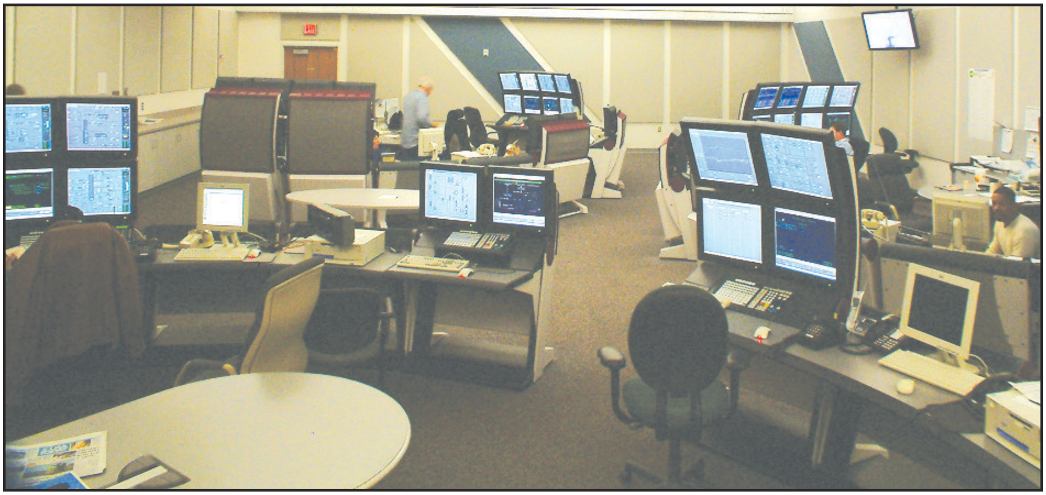

Figure 12.7.3. Typical operator areas within a control room

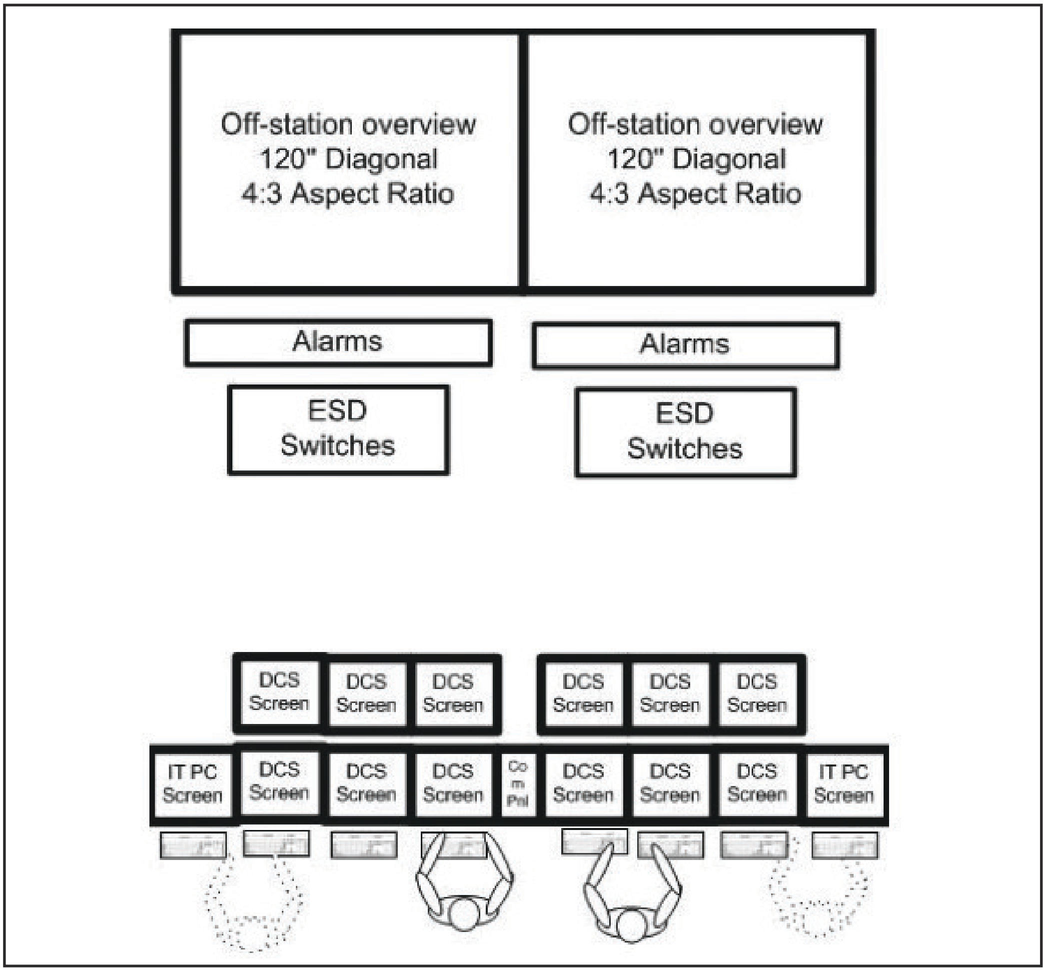

Figure 12.7.4. Expanded screen layout architecture

We now turn our attention away from the number and arrangement of screens and toward the content and purpose of the displays placed on those screens. This is actually where the awareness of situation is to be found.

Hierarchical Display Architecture

A three-level display hierarchy (Figure 12.7.6) delivers a robust structure that encourages ready access to information while at the same time keeping important situation context and promoting efficient navigation to go deeper. The first two levels follow an expected progression from the general to the more detailed. But the third one is a

Figure 12.7.5. Expanded screen control room

FIRST
Produce an overview
(provides a one-view glimpse of everything that might go wrong)

SECOND
Provide ready access to important details
(second line of defense—confirm or deny)

THIRD
Provide detailed guidance
(if problem requires support—here is where it is) .

Figure 12.7.6. Display hierarchy levels

departure. It does not provide a more detailed view. Rather, it is presumed that if the first and second levels are not capable of providing the needed detail, another paradigm is needed. That paradigm is support, not more detail. This structure is a very clear departure from the widely accepted norm for navigation and structure: "clicking to oblivion" that characterizes most personal computer operator interface exchanges.

Clicking one's way down a "tree" of choices to locate information is a tedious, time-consuming activity. If the tree arrangement logic differs from the mental expectations of the user, it can be distracting. Even if there is a good match, it takes time and requires a constant stream of locating the click target area, using pull-down menus to provide the accepted choices, deciding which choice matches the current need (and often finding no good match, one must click down and back to attempt to find the closest or the least not-closest match), selecting the choice, and going down to the next level. As you might imagine, this approach is always frustrating and often becomes nonproductive during times of operator stress. At the very time it needs to be at its best, it fails. This is why efficient operator-machine interaction uses as little clicking down as possible.

The Overview Level

An overview must provide a way to take in the entire operator's area of responsibility. It must show summarized alarm statistics, grouped status information, and important conditions of both upstream and downstream units and provide ready navigation to everywhere. It must communicate whether or not the production unit is operating well. If well, how well? If not, then how not and where? And it must do this without requiring the operator to search it out.

Let's take a look at an example of how individual displays might be organized to provide both the continuous ability to maintain an overview of the process as well as the necessary information and handles required to manage problems. We start with an unusual form of overview, shown by Figure 12.7.7. It is called a taxonomy view. This example permits a very clear illustration of the essential elements of an overview, without the distraction of the familiar. But please do not go out and build one just yet. The illustrated process is a refinery boiler unit. Colocated there are the plant steam and plant air systems. Taxonomy views most often are used when the particular plant has a number of very similarly functioning units, regardless of how different their physical appearances might be. This arrangement has the benefit of providing information in the same format for each element. Normal means the same thing for each, and so for the abnormal. Knowing what to do with one during an abnormal situation will serve as a close guide for what to do for the others. Shown in the figure are places for four boilers: COB stands for a CO boiler, PB is a Power Boiler, RB is a Riley Boiler, and FUTURE is for a future one. Each boiler has six constituent components: BFW is the boiler feedwater, STM is the produced steam, B/D is the blow-down and condensate, FG is the combustion fuel gas, AIR is the combustion air, and FLUE is the burned gas exhaust. The columns for BFW and FG are for the shared boiler feedwater system and fuel gas supply. The BFW component includes the following: DEA for deaerator, COND for condensate system,