TROUBLESHOOTING CHILLED WATER DISTRIBUTION PROBLEMS AT THE NASA JOHNSON SPACE CENTER

Various modifications over the years had skirted a problem with low chilled water ΔT without solving it. Now, discovering the cause could help improve the system's present level of efficiency.

By WAYNE KIRSNER, PE,
Kirsner, Pullin & Associates, Marietta, Ga.

The NASA Johnson Space Center had long had a problem with chilled water (CHW) distribution to the 40 buildings served by the central CHW plant. Originally designed in 1964 for a 16 F temperature drop across the central plant chillers, the chilled water system could only attain a 7 F Δ*T* on average and a 10 F Δ*T* at best. This meant not only that twice as much CHW as originally intended had to be pumped around the 5 mile campus loop to satisfy the cooling load but also that the seven 2000-ton chillers in the central plant couldn't be loaded much beyond half of their capacity.[1] Thus, operators were forced to run twice as many chillers to meet the campus load, and the frictional loss in the mains due to the excessive flow made it tough to deliver sufficient CHW flow to hydraulically distant buildings.

Beginning in 1978, NASA tried to remedy this problem by replacing the CHW pumps serving the chillers with new 250-hp pumps that would pump twice as much flow. To mitigate the roughly fourfold increase in head pressure drop across the chillers due to increased flow, they retrofitted the evaporator shells on all seven chillers from three-pass to two-pass. This limited the pressure rise in the chiller evaporator bundles to only about a 10 percent increase but still left problems in the rest of the system. Piping in the CHW plant and buried mains serving the 265-acre central campus could not accommodate the twofold increase in CHW flow without excessive velocity and

[1] *Chillers were designed to cool 56 F CHR to 40 F CHS. If CHW returns at an average of less than 48 F, without an increase in flow through the chillers, chillers cannot load beyond 50 percent.*

2000-ton chillers in the original chilled water plant.

Photos courtesy of National Aeronautics and Space Administration, Johnson Space Center concomitant head loss in critical legs of the system. The most hydraulically distant buildings continued to be starved for CHW flow periodically.

In 1990, a second 24-in. chilled water main was buried parallel to a segment of the existing underground loop to shunt CHW to certain critical buildings. The new main helped free up pipe volume in the existing CHW main segment serving the western portion of the campus, but buildings on the other side of the campus still suffered.

In 1991, a second 4000-ton chilled water plant was constructed on the opposite side of the campus from the existing plant. Its function, in addition to serving some new buildings' loads, was to feed the campus chilled water system from the "backside" so as to provide plenty of chilled water supply pressure to force water to the most hydraulically distant buildings from the original CHW plant. Fig. 1 shows the chilled water distribution loop at the Johnson Space Center.

This "nuke it" approach to solving the problem did, in the end, make sufficient CHW available to all buildings, but it left the Johnson Space Center (JSC) with a hydraulically complex CHW pumping and piping system that used a lot of pumping energy. In the course of an energy study we performed on the JSC campus, we set out to determine the cause of the original problem of low CHW ΔT with the ultimate aim of improving the efficiency of the system.

The original design

At least three times, we were informed that the problem originated with the initial design. We were told it called for a 16 F ΔT at the

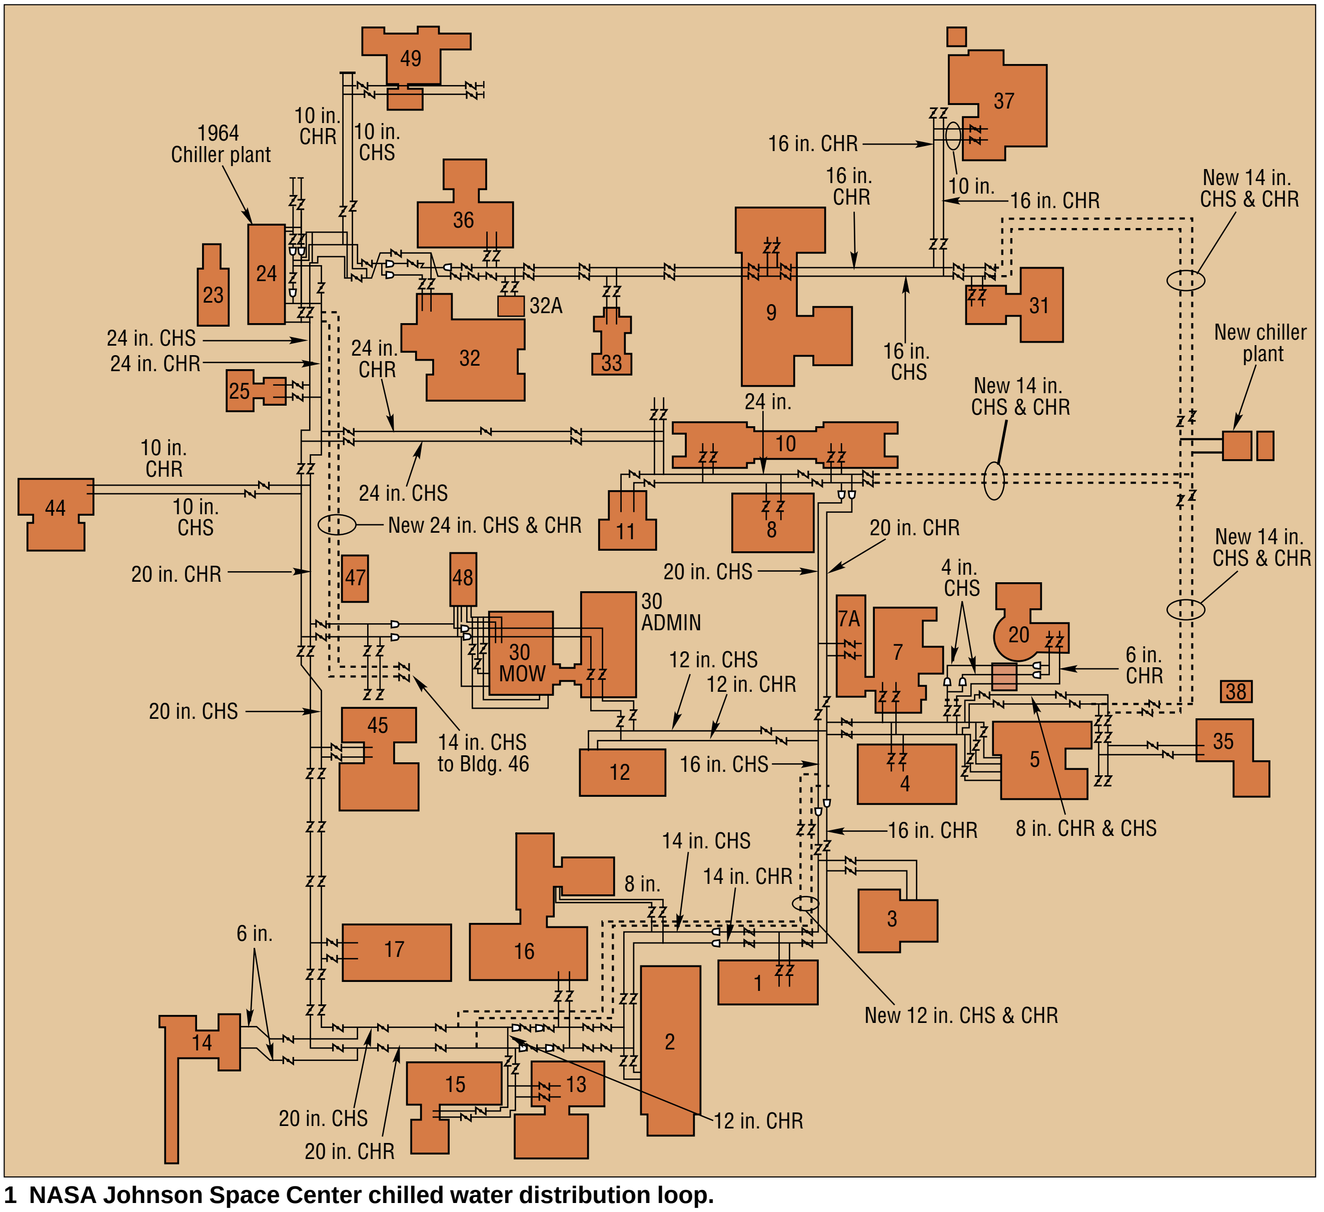

1 NASA Johnson Space Center chilled water distribution loop.

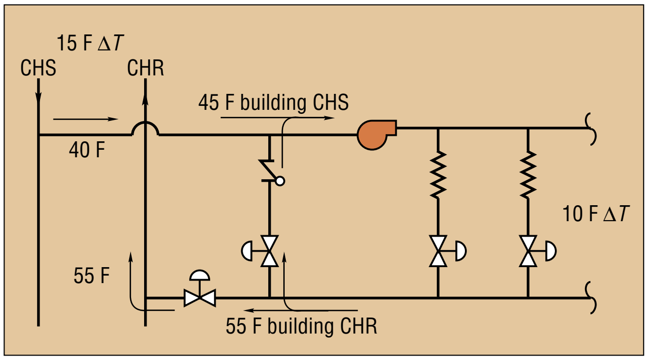

2 Chilled water blending station.

TABLE 1—Increase in CHW flow with warmer CHS temperature. (Based on four-row, 103 fins per ft, 500 fpm air velocity, 75/62 F EDB/EWB, maintaining 55 F.)

CHS temperature, F	Required gpm	Percent gpm	ΔT, F
40	51.8	100 (base)	10
42	67.6	130	7.7
43	81.3	157	6.6
44	exceeds 8 fps in coil	—	—

chilled water plant but only a 10 F ΔT across the buildings' systems.

This was true. Buildings were provided with blending stations such as that shown in Fig. 2. Campus CHW was to be supplied at 40 F and blended at individual buildings with chilled water return (CHR) from the building to make 45 F chilled water supply (CHS). A building pump then circulated the 45 F CHS through the building's air handling unit cooling coils, which were typically selected for a 10 F rise from 45 to 55 F at design conditions. Some building CHR was recirculated back through the building for blending, and the rest returned to the central plant through the campus CHR main. Thus, the building's AHU coils experienced a 10 F rise in CHW temperature, and the campus saw a 15 F rise—from 40 to 55 F—at design load conditions.

This was the *intended* design; it was not the source of NASA's problem. So what was the origin of the problem? We investigated several possibilities.

CHS temperature

We knew part of the present-day problem with low ΔT and excess CHW flow was an increase in campus CHS temperature from the original 40 F design up to 42 F. This increase was instituted by NASA engineers to save chiller compressor energy. The downside of this energy-conservation measure was that instead of blending two parts 40 F campus CHS to one part 55 F building CHR to make 45 F CHS, just over three parts of 42 F campus CHS to one part of 55 F building CHR was required to blend to 45 F. This change caused campus CHW ΔT to shrink by 2 F and increased the overall campus flow requirement by 15 percent.

In addition to increasing the requirement for campus CHW at blending stations, the increase in CHS temperature contributed to the problem of low ΔT in another way. Raising campus CHS temperature above the 40 F design adversely impacted the CHW flow requirements at those AHUs scheduled to receive unblended 40 F CHS. Some 4096 of roughly 28,000 gpm of chilled water supplied to the campus is scheduled to be at 40 F. Buildings receiving unblended CHS contain computers, mission support equipment, and other specialized services. Increasing CHS temperature to cooling coils decreases the ability of these coils to transfer heat. To compensate, CHW flow must be increased to maintain heat transfer.

Table 1 shows, for a typical coil selection, CHW flow required as a function of increased CHS temperature with the constraint that coil heat transfer is maintained constant. According to the table, resetting CHS from 40 up to 42 F requires 130 percent of the CHW flow in the base selected coil to achieve the same leaving air temperature and reduces ΔT across the coil from 10 to 7.7 F at full load.

In NASA's case, a 30 percent increase in the total amount of unblended CHW flow supplied was only an extra 4 percent of the total campus flow. This 4 percent added to the 15 percent additional flow previously mentioned was not insignificant but certainly could not account for the roughly twofold in-

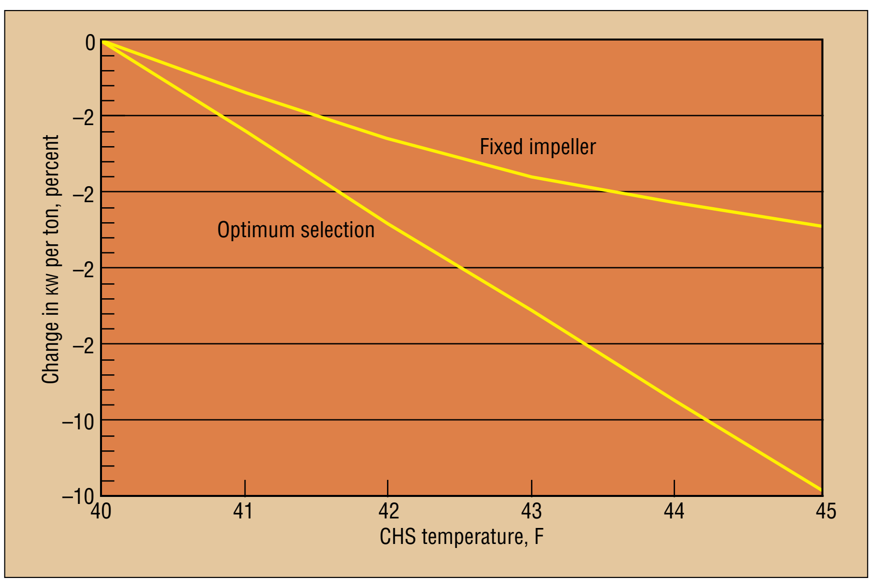

3 Change in chiller efficiency with CHS reset, based on 500-ton three-stage chiller with fixed impeller selected to provide 40 F CHS. Optimum selection allows impeller and motor to vary with CHS to provide 500 tons. (Data courtesy of Terry Dugan, The Trane Company.)

CHW ΔT at Part Load

At say 50 percent part load across a cooling coil, if CHW flow fell off to 50 percent also while maintaining a constant ΔT, the heat balance between air and water would be neatly maintained. This would be simple. Unfortunately, this is not how things work. The rate of heat transfer, Q, between air-to-copper and copper-to-water governs how much CHW flow will be required at part loads:

$$Q = U\,A\,LMTD$$

As air flow and CHW flow vary so does the composite heat-transfer coefficient, U. Its primary components are the convection heat-transfer coefficients, h, of the airside and waterside of the coil:

$$\frac{1}{U} = \frac{1}{h_a} + \frac{1}{h_w}$$

As air and water flow velocities decrease, the respective h values, and hence U, decrease. For water flow in pipes:

$$h = 158(1 + 0.11T)\frac{V^{0.8}}{d^{0.2}}$$

for turbulent flow, heat in.

For air flow across tubes:

$$\frac{hD_0}{k} = C(N_{\mathrm{Re}})^{0.6}(N_{\mathrm{Pr}})^{0.33}$$

where

$$N_{\mathrm{Re}} = \frac{V_c D_0}{v}$$

$$N_{\mathrm{Pr}} = 0.72$$

From the equations above, one can see that U, the heat-transfer coefficient, falls off with flow—*i.e.*, velocity—less rapidly than Q, the load. Whereas Q falls off linearly (*i.e.*, to the first power), h for water falls off by a 0.8 power, and h for air falls off at a 0.6 power.

Hence, at say 50 percent of the air flow across a coil, the cooling coil load will be 50 percent, but *UA* will have decreased by something less than 50 percent—maybe about 40 percent. Thus, *LMTD* must decrease to balance the heat-transfer equation above.

Of the four temperatures that determine the *LMTD*, three stay essentially constant; the entering and leaving air temperatures (EAT and LAT, respectively) don't change much, and likewise the CHS temperature is constant. Therefore, it is the CHR temperature that will act to increase (by means of reducing CHW flow) to reduce the *LMTD* in response to the reduced load.

The point is: **CHW ΔT *increases* at coil part load**, or alternately, *CHW flow is reduced to a greater degree than load as load falls off.*

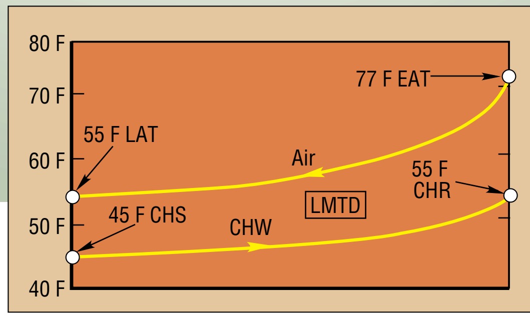

flation in design CHW flow that the JSC was experiencing.

A more interesting question for NASA relating to the CHS temperature reset was: *Should chillers be operated at 42 F instead of the original design 40 F CHS temperature?* In other words, was NASA wise in sacrificing pumping energy to save compressor energy at the chillers. The following thumbnail calculation compares the pumping power saving with that of the chillers.

If campus CHS temperature was returned to 40 F, we calculate that actual campus CHW flow would fall by 3850 gpm. The power to be saved by not having to pump this extra flow through roughly 139 ft of head (60 psi) experienced by the campus CHW pumps is:

$$\frac{(3850\text{ gpm})(139\text{ ft})}{(3960)(0.83\text{ eff.})} \times \frac{0.746\text{ kW/hp}}{0.95\text{ eff.}}$$

$$= 88\text{ kW}$$

In contrast, Fig. 3 shows that a chiller (with fixed impeller) would use about 2.6 percent less power to produce 42 F CHW rather than 40 F.[2]

The JSC's average load over the year was over 6000 tons. Even if all chillers were new, fairly efficient, electric centrifugal chillers operating at 0.6 kW per ton (which they were not), the average year-round compressor power saving would be, roughly, 0.026 × 6000 tons × 0.6 kW per ton = 93.6 kW—virtually a wash with the pump saving of 88 kW.

At the cooling peak load of about 10,000 tons, however, when the electric billing demand for the year is established, the chiller compressor power saving would far outshine the 88 kW of pumping saving: 0.026 × 10,000 tons ×0.6 = 156 kW.

Therefore, we recommended that NASA stick with the 42 F CHS temperature. The chillers were saving more energy and power operating at 42 F than the pumps were sacrificing in additional pumping energy.

Blending stations

A second place to look for the source of the low CHW ΔT problem was the building blending stations. It seemed reasonable that if CHR from the buildings was to reach 55 F at design conditions, the blending stations had to be operative to blend 42 F campus CHS up to 45 F at the buildings. In fact, we found from examining computer monitoring records that 27 percent of the blending stations that should have been working were not. Building air conditioning systems served by these inoperative blending stations were receiving unblended 42 F water even though they were designed to receive 45 F CHS.

One might think that this was the source of the problem. But consider this—a coil supplied

continued on page 57

[2]*Allowing the impeller selection to vary, as would be the case with a new chiller, doubles the energy savings.*

continued from page 54

with 42 F CHS achieves a larger ΔT and requires less flow to satisfy its leaving air temperature control thermostat than if supplied with 45 F CHS. Coil simulation runs confirm that flow requirement falls to between 50 and 80 percent of the design flow if CHS temperature is reduced from 45 to 42 F. Since an operational blending station requires 77 percent campus CHS to mix with building CHR (3.3 parts 42 F CHS to 1 part 55 CHR), it's arguable that buildings with inoperative blending stations may actually require less CHW flow. (In fact, it's arguable whether blending stations are such a hot idea to limit campus CHW flow in the first place.)

The degree of CHW flow reduction that can be achieved depends on the initial CHW flow velocity within the coil's tubes. If it's high (above 4 fps), a large reduction can occur; if initial design flow is low (say below 3 fps), however, the heat-transfer coefficient suffers as decreasing flow velocity approaches nonturbulent flow. Below about 1 fps in a 5/8-in. tube, flow is no longer turbulent. This effectively limits flow reduction since the leaving air temperature sensor will sense a loss of heat transfer and act to restore flow. Of course, if AHU cooling coil controls are not, for whatever reason, throttling flow, CHW flow will not be reduced at all. So which is it? At buildings with inoperative blending stations was flow reduced or excessive?

We examined the CHW flow at the buildings with inoperative blending stations to ascertain if flow was in excess of design. Measurements of campus flow recorded by the central control and monitoring computer verified that buildings with inoperative blending stations were, in every case, calling for excess flow and contributing to the low ΔT problem.

AHU cooling coils, controls

We came to believe that the source of low campus ΔT must lie at the individual building AHUs. Our first thought was to ask if normal part-load performance of the AHU's cooling coils could explain the low ΔT. Should not CHW ΔT across a coil fall at part load? No—the opposite is true. As load decreases, CHW ΔT across a cooling coil will *increase*. This point is explained in the accompanying sidebar.

Low CHW ΔT in buildings' cooling coils, then, could only stem from two causes:

- An inability to transfer heat by the coils.
- Problems with CHW throttling valve control.

Dirty coils would have provided a convenient explanation for poor heat transfer. But if the coils had an inability to transfer heat, spaces served would be undercooled, and they weren't. Even on 100 F or hotter days, NASA kept the building lettuce-crisper cool. A further check of coil faces confirmed that they were clean. This was to be expected as the contract maintenance at JSC was very good.

Our focus shifted to controls. Mechanical drawings showed three-way CHW valves at AHUs, but NASA engineers told us valves were two-way throttling type. Converting the operation of three-way valves to two-way by closing off a valve in the bypass is not an uncommon strategy for converting a constant flow system to variable flow. The problem with this retrofit, however, is that three-way valves do not generally have the spring strength to close against any significant pump pressure. Could pump pressure be lifting the CHW valves off their seats?

Nope. NASA had indeed changed all the three-way valves to two-way in an earlier effort to reduce campus CHW flow. Flow was controlled by two-way throttling valves, which in turn were controlled to maintain a constant leaving cold deck temperature by a local pneumatic receiver-controller with a remote set point control adjustment, or RCPA. The RCPA was adjustable from the central monitoring and control computer located in the Central Plant.

This seemed to be a neat feature of the system. If an occupant was uncomfortable, he could notify the central computer operator by telephone, and the operator could remotely adjust the set point of the cooling coil throttling valve controller up or down via the RCPA. This seemingly innocuous feature of the central control and monitoring system, however, coupled with human nature, turned out to be a primary contributor to the problem of low campus CHW ΔT.

On August 27, 1991 at 10:35 AM, we used the central computer to record the cold deck temperatures and RCPA values for the seven AHUs serving Building 45—a 112,000 sq ft office building occupied by the NASA engineers who oversaw operation of the campus. The computer readouts, as well as the CHW temperature measurements, we recorded at the AHUs that morning are shown in Table 2. The outdoor temperature at the time of the recordings was a relatively mild 82 F.

Note that all chilled water throttling valves except that for AHU-1

TABLE 2—Building 45 cold decks on August 27, 1991.

AHU	Cold deck temp., F	RCPA, percent	CHS temp., F	CHR temp., F	ΔT, F	CHW valve percent open
1	60	5	44.0	63.0	19.0	40
2	53	4	44.3	52.9	8.6	100
3	53	5	44.1	54.1	10.0	100
4	53	50	44.3	52.9	8.6	100
5	54	30	44.4	54.1	9.7	100
6	54	4	44.3	54.0	9.7	100
7	53	5	44.1	52.7	8.6	100
9	51	50	—	—	—	—

were 100 percent open. (AHU-1 served the common spaces in the building on the first floor and thus had no "constituency." The other AHUs served office spaces.) Note also that according to the computer, all cold deck temperatures except that for AHU-1 were maintaining leaving air temperatures below the design set point value of 55 F. Further note that, as we would expect, the CHW ΔT for AHU-1 exceeded the design value of 10 F as it should at part load. Why were CHW throttling valves 100 percent open at all other AHU cooling coils even though their cold decks were below design set point? Obviously, the set points had been adjusted downward.

Check out the RCPA values in the table; 50 percent is neutral. AHU cold deck set points were being reset downward by well-meaning operators at the central plant. We theorize that every time someone important complained, the RCPA was set down a little bit, but apparently there was no procedure to reset the RCPA back up (nobody ever complained that it was too cool). Essentially, there was a one-way ratchet on cooling set points.

This situation wasn't an isolated case in Building 45. Building 1, where the top brass resided, was worse. The average CHW ΔT in this building was only 6 F.

What happens when the cooling coil leaving air temperature is reset downward? Table 3 demonstrates the increase in CHW flow for which a throttling valve controller will call to achieve a colder leaving air temperature (LAT). At just 2 F reset below design, the controller would call for well over double the chilled water flow in an attempt to achieve the lower LAT. At and below a cold deck set point of 51 F, no amount of CHW flow will permit the coil to achieve this LAT, so a throttling valve would go wide open in a futile attempt to reach an unattainable condition.

TABLE 3—Increase in CHW flow to achieve colder LAT. (Based on four-row, 95 fins per ft, 500 fpm air velocity, 75/62 F EDB/EWB, 45 F EWT.)

Cold deck LAT, F	Required gpm	Building ΔT, F	Percent gpm
56	43.7	10	base
55	67.9	7.2	155
54	115.0	4.7	263
53.75	134.7	4.1	308
51	cannot be achieved	4.1	—

So this was the final piece of the puzzle. Campus CHW ΔT was low and flow high for three reasons:

- Building AHU cooling coil control set points were ratcheted downward over time, causing many CHW control valves to run wide open, admitting as much CHW as the pumps would push through the system.
- Primary CHW temperature was delivered at 42 F versus the 40 F design.
- Approximately 27 percent of the blending stations were inoperative, and in buildings where CHW controls allowed cooling coils to "run wild," this especially caused excess campus CHW use.

How much could NASA save by remedying the low ΔT problem? A 16 F ΔT was no longer achievable at JSC due to modifications over the years. Our building-by-building tabulation of existing conditions showed that better than a 12.2 F ΔT was achievable at full design load at each building (and remember that at part load, ΔTs should be even greater). Achieving a 12.2 F ΔT would reduce average chilled water flow to 62 percent of the current requirement and save more than *50 percent of the annual chilled water pumping cost.*

The work to achieve this saving could be done in house:

- Reset all cold deck set points at building AHUs to control at *design* leaving air temperature.
- Ensure buildings achieve design ΔT or better at part load by monitoring *existing* computer monitoring points reporting CHS and CHR temperatures at each building and by surveying buildings for cause when low ΔT is recorded.
- Activate blending stations and set them to control for design building CHS as scheduled on drawings.

Eliminate throttling valves

There was another ripe opportunity to reduce CHW pumping cost by yet another 50 percent: eliminate the use of throttling valves to control CHW flow through the chillers and out to the campus. Located at the discharge of each chiller was a 16-in. pneumatic throttling valve. In a "snapshot" of pressure readings recorded at four of five chillers operating at part load on an August evening, the discharge throttling valves were found to be devouring 44 of 69 psi produced by the 250-hp CHW pumps. The effect of the valves was to throttle pump flow through the chillers from 6000 gpm down to 4300 gpm. The plant operators used the throttling valves to maintain a more or less constant CHS pressure out into the system. In the evening, when load would subside and building pumps were deenergized, system flow resistance would increase and plant CHW pumps would back up on their curves. Operators, noticing a rise in CHS pressure at the exit from the plant, would respond by adjusting the throttling valves to consume more pump pressure so as to reduce the CHS pressure experienced out in the system. This was vaguely thought to save energy.[3]

Obviously, this was an extremely inefficient way to control CHW flow, but it also had another consequence—sometimes CHW flow tried to reverse itself at hydraulically distant buildings. In other words, pressure in the campus CHR main exceeded that in the campus CHS main so that CHR flow attempted to enter the

[3]*In fact, it did save some energy in buildings where cooling coils were running wild and increased CHS pressure would cause them to run wilder, so to speak.*

distant building. This had long been a baffling and intermittent problem at the JSC campus. We found it to be caused by the "non-decoupled" nature of the individual building CHW pumps and the central plant pumps and exacerbated by the use of the throttling valves at the central plant. (For more information, see the accompanying sidebar.)

By replacing the function of the throttling valves with a variable-speed drive on each pump, we showed that pumping energy could be reduced another 50 percent on top of the 50 percent saving already predicted for reducing the campus CHW flow. Controlling the drives to maintain positive CHS pressure with respect to CHR pressure at the most hydraulically distant building would eliminate the tendency to reverse flow.

Parasitic pump heat

Reducing the pumping energy imparted to the CHW stream garners yet another saving. Pump energy equal to the pump brake horsepower is transferred into the CHW stream in the form of flow and head to eventually, via friction, become heat. Reducing pumping energy reduces this parasitic heat gain. Each CHW pump, at 216 bhp, imposed almost 50 tons of refrigeration load on the chilled water system. That's almost 300 tons of extra cooling when six pumps are operated during the summer. Reducing pump brake horsepower to 25 percent, as made possible by the modifications described above, would reduce the parasitic cooling by roughly 225 tons in summer.

The bottom line

The overall dollar saving to the NASA Johnson Space Center for all the fixes described above was calculated to be $231,400 per year. Implementation cost for the project was $349,300. No cost was included for in-house recalibration of building cold deck temperatures and blending stations. The projected simple payback for the project is 1.5 years. Ω

Reverse Pressurization of CHW Mains

On a pump curve (and in real life), system flow is established at the equilibrium point where system resistance equals total pump head at that flow. At the NASA JSC, plant CHW pumps act, at least partially, in series with building CHW pumps because building and campus CHW loops are not decoupled (see Fig. 2 in the text). Hence, total pump head is the sum of the head contributed by the plant and building pumps. Because the two sets of pumps "share" the total head requirement to move CHW through the entire system, there is a point in the system at which the plant pumps hand over, so to speak, the job of moving the CHW to the building pumps. Ideally, this point should occur at the entrance to the farthest building's CHW system.

In the diagram of the JSC CHW system below, however, this point occurs out in the campus loop just beyond the third building from the central plant. At this point, the 70 psi generated by the central plant's pumps (of which only one is shown) has been expended. Beyond this point, building pumps are using a portion of the head they generate to reach back into the campus CHS main and "pull" CHW into the pumps. The remaining head is consumed in pushing the CHW through the building and then pushing it out into the campus CHR main. The result of the building pumps pushing (*i.e.*, adding pressure) to the campus CHR while pulling (sucking) from the campus CHS is to create an inverse pressure gradient near the end of the campus CHW mains. This is the case at the last two building takeoffs in the example below.

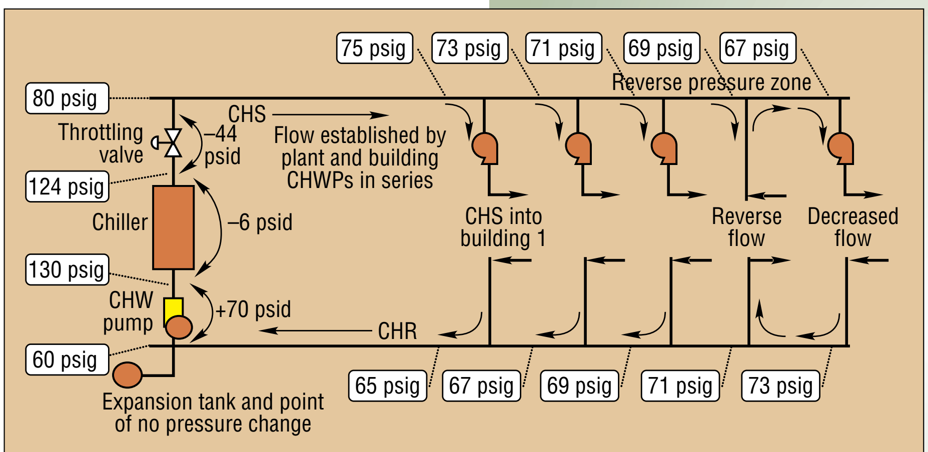

Schematic of central plant serving campus buildings with CHW pressure gradient reversal.

The consequences of an inverse pressure gradient are:

- Buildings without a pump located in the reverse pressurization zone will experience either no CHW flow (if there's a check valve) or reverse flow.
- Buildings downstream of any building undergoing reverse flow will receive a mixture of campus CHS blended with recirculated CHR from the upstream building.
- Building pumps drawing flow from the zone of reverse pressurization will not produce design flow because of the extra head they must generate to overcome the inverse pressure gradient.

The reversal of pressure gradient is avoided when plant pumps output sufficient pressure to overcome all resistance in the campus CHW system out to the entrance of, hydraulically speaking, the last building. The plant pumps at NASA JSC easily have sufficient capacity to do this if the throttling valves are not used to emasculate them by eating up a large portion of the head they produce.