David Sellers

What to Consider When Designing for *N* + 1

BY DAVID SELLERS, P.E., MEMBER ASHRAE

The September 2020 *ASHRAE Journal* column by Nathan Ho, "Performance-Based Approach to Laboratory Exhaust Systems," mentioned $N + 1$ redundancy for the exhaust fans. In reviewing it, I recalled some lessons I learned regarding what that truly means during my tenure as a system owner at a silicon wafer fabrication facility (aka, a "fab").

In my experience the definition of $N + 1$ is "it depends." Generally, the relationship is an expression of redundancy. For Nathan's system, the $N + 1$ requirement was expressed in the context of fan redundancy. It was met by designing a system that required two identically sized fans to handle the design condition and providing a third fan with the same capacity as the other two. If one of the operating fans failed, the reserve fan could be started to restore full capacity while repairs were made.

The lab was not provided with a redundant distribution system. If one of the ducts imploded, the system would be out of service until repairs could be made, even though there was an extra fan. There are other *N* relationships that address this. For example, 2*N* implies that there is a totally independent, fully redundant distribution system. And the concept is not just about having redundancy; it's about seamless transition to the redundant systems upon failure.

That means it is possible to "mathematically" address $N + 1$ in the design solution only to discover your system is not as redundant as the relationship implies when it comes online due to system dynamics, and perhaps due to the details of the definition of a "failure."

Prior to joining the team at the fab, I had limited exposure to industrial process sites. $N + 1$ had been about having an extra chiller in the central plant that was as large as the largest chiller. In a health-care environment, it might have been about complying with a code requirement to have an extra boiler and related auxiliaries to allow the facility to continue to be heated on the design day despite the failure of a prime mover. In both scenarios, an operator had time to facilitate the transition to the redundant equipment before things went "amok."

Within hours of arriving at the wafer fab, I learned there may be little or no time for such an operator intervention in an industrial process environment. Designs needed to address failure modes in the context of process requirements, safety, equipment performance characteristics and system dynamics. This needed to be complemented by an automation strategy that could

David Sellers, P.E., is senior engineer at Facility Dynamics Engineering's office in Portland, Ore.

detect a failure and seamlessly transition to the backup machinery. Lacking that, things could quickly "unwind."

Setting the Scene

The fab was built on an accelerated timeline to catch an economic wave sweeping through the industry. Thus, it was done as a design-build project with a very basic owners project requirement (OPR), which was specific about the configuration of the cleanrooms and systems serving them, but left many details up to the contractor.

Figure 1 illustrates the cleanroom. This is a complex figure, and that is part of my point; the *N* + 1 system design—and in this case the entire cleanroom was the system—*must* address the complexities of the system to be successful. If it focuses on specific equipment or subsystems rather than the whole system, problems like the ones I will describe can occur.

The OPR specified *N* + 1 for the cleanroom airside with the intent being that the failure of a fan in the makeup, recirculation or process exhaust subsystems would not take the cleanroom out of production. This column will focus on the *N* + 1 issues in the process exhaust subsystem. *Figure 1 Notes* provides additional information on how *N* + 1 was achieved (or not) for the MAU and recirculation subsystems.

Ultimately, the contractor's interpretation of "not taking the cleanroom out of production" and "failure" ended up being different than what the owner had in mind. Specifically, the contractor viewed a brief outage, for instance, a process exhaust system being offline for three or four minutes during a transition, as not being out of production. But for us, there were production loss implications beyond the three to four minutes of downtime. And the contractors' perspective considered a failure in the context of the fans, while in our perspective it was crashing the cleanroom.

N + 1 and Conservation of Mass and Energy

When I arrived on site at the fab, a good two-thirds of the fab was coming out of the ground (*Photo 1*). But the epitaxial (EPI) cleanroom was online in qualification runs (demonstrating that we could consistently make good product). With part of the site under construction, the operating team faced a number of challenges. These included dealing with problems created by systems that we didn't own and couldn't touch that were not fully commissioned—but that had a direct impact on operations in the EPI cleanroom.

The EPI fab manager was having a recurring problem in the EPI quality control (QC) area (see *Figure 1 Notes,* Note 7). At random times, the doors would slam against their frames, and the force across them made it difficult if not impossible for someone in the room to force them open and get out. Adding to the excitement, floor tiles had also blown out. Given that the floor tiles were made of ±40 lb (18 kg) cast-aluminum, this was somewhat of a concern.

Forty pound (18 kg) projectiles and impossible-to-open emergency exits aside, the events contaminated the cleanroom, shutting down production. The root cause of the problem turned out to be the makeup air unit (MAU) controller rebooting at random. When that happened, it de-energized all its outputs and went through an orderly restart. Even though we were *N* + 1 regarding fans on paper, we were not in the context of the subsystem because the controller failure was shutting both fans down concurrently.

Because there were no interlocks between the MAU and process exhaust control system, when the MAU went offline, the process exhaust system continued to operate. Thus, the process exhaust fan did whatever it took to achieve design conditions. Since the fan was attempting to provide the QC area with 6,000 cfm (2832 L/s) of exhaust, with no source of makeup air, it was pushed up its curve. Since the discharge was referenced to atmosphere, the inlet pressure and cleanroom went extremely negative. This is why the doors slammed against their frames and became difficult to open.

The floor tiles in most of the cleanroom were perforated to provide a recirculation path to the subfloor and maintain the required air change rates. For the QC area, the required air change rates were achieved without recirculation due to the high makeup and exhaust flow rates, thus the tiles in that area were solid.

During the MAU failure, the pressure differential created by the process exhaust fan was imposed across the floor tiles. Since the recirculation systems had no direct access to makeup air, the subfloor went negative as the process exhaust fans tried to pull air directly from the cleanroom, causing the subfloor plenum divider to collapse, blowing out the floor tiles.

This opened a path to other areas of the cleanroom, but without makeup air, the entire cleanroom was

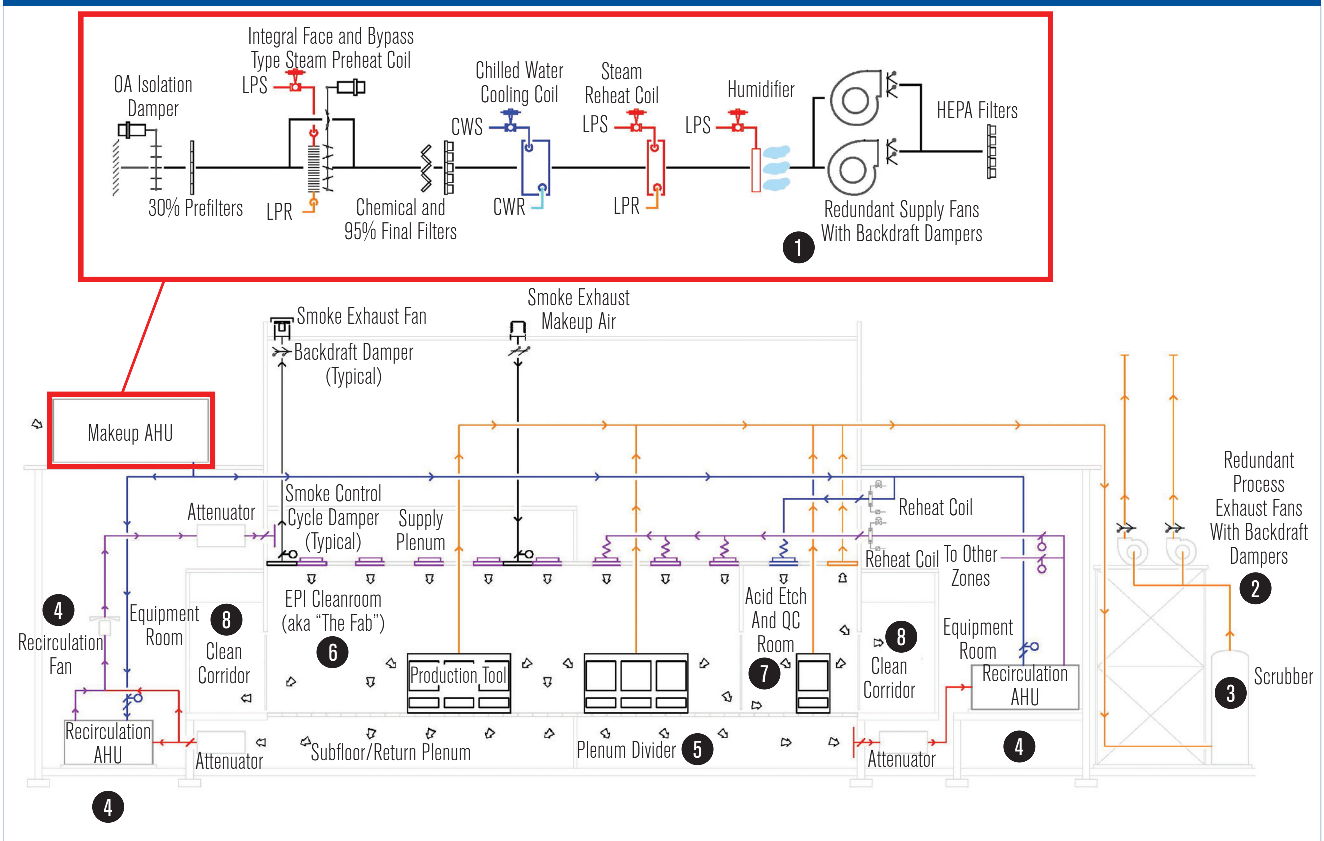

FIGURE 1 The epitaxial (EPI) cleanroom system. See *Figure 1 Notes* on facing page for details.

pulled negative, This reversed the flow through all the door cracks and other leakage paths and contaminated the entire cleanroom in addition to the QC area. But it did provide enough flow to allow the process exhaust fan to achieve near design flow.

Bottom-line, the *N* + 1 MAU and process exhaust fans did not protect the fab from going out of production in this failure mode because the failure was not associated with the fans; it was with the systems that controlled the fans.

N + 1 and Fab Crashes

As I settled into my new position, a bigger problem emerged. The cleanroom was running under a temporary certificate of occupancy, allowing us to begin qualification work while the final details of the fab construction were worked out. This included demonstrating the response of the cleanroom to various failure modes, including the failure of a process exhaust fan.

Since the owner's personnel believed the contract was for a system in which the failure of a process exhaust fan would not take the cleanroom out of production, once the various safety devices were verified (hazardous gas alarms, flow switches, etc.), they saw no reason to not work in the fab while testing was going on.

The design build contractor's approach to operating the *N* + 1 process exhaust fans was: *If the fan that was online failed, start the lag fan.*

PHOTO 1 The fab under construction.

FIGURE 1 NOTES Additional information on how *N* + 1 was achieved (or not) for the MAU and recirculation subsystems shown in *Figure 1*.

1 This illustrates the details of the makeup air unit (MAU). The design intent for the MAU fans was also for full redundancy (*N* + 1), and each fan had been sized to provide all the air that was removed by the process exhaust system, plus an additional 6,000 cfm (2832 L/s) allowance for envelope leakage. The fan speed was modulated to maintain the cleanroom pressure, allowing the fans to compensate for changes like doors opening and closing or filters loading over time.

The cleanroom envelope leaked more than twice the design allowance, and this was after significant time and money were spent trying to find and seal leaks. As a result, we had to run two fans under normal operation to maintain the required positive pressure cascade (see Note 7), and the *N* + 1 design target was not achieved due to a failure in the envelope, a subsystem served by the MAU.

Even without the envelope issues, the *N* + 1 integrity of the MAU was compromised. While backdraft dampers allowed the lag fan to come online and maintain flow to the cleanroom when the lead fan failed, there was no way to isolate the failed fan from the system to make repairs while keeping the system online. Pressure differentials were too high to allow opening the fan compartment to enter it while the unit was operating.

2 Unlike the MAU fans, the process exhaust fans were truly 100% redundant (*N* + 1) in terms of flow and capacity once the fab was online. They operated at constant speed, but were equipped with variable speed drives to facilitate balancing.

3 Scrubbers are used to neutralize hazardous exhaust prior to discharge to atmosphere. The tool sets in the EPI cleanroom used very strong acids and alkalis to etch the wafers for various production and quality control reasons. Thus, the process exhaust had to be "scrubbed" to remove the acid and alkali vapors from the airstream. Failure of the process exhaust system would expose the cleanroom staff to some very hazardous vapors.

4 The recirculation systems recirculated air to maintain the targeted air change rates, which was the primary mechanism used to manage particle counts. Generally speaking, higher air change rates equated to lower particle counts. The three recirculation systems shown represent multiple fans and units. Generally speaking, we could have one fan or unit for each area offline and still maintain particle counts, achieving the *N* + 1 target. But we usually ran all of them all the time. The units contained a fan, filter racks to support 95% filters that were used to clean up the space after construction and a cooling coil. Much of the load on the cooling coil was the fan heat associated with the recirculation fans.

5 Our cleanrooms included a subfloor that contained a maze of piping supporting the process tools and served as the return air path for the recirculation systems. Since the cleanroom itself was divided into different circulation zones, the subfloor was also divided by plenum walls, arranged to match the circulation zones in the room above.

6 The EPI cleanroom was one large open cleanroom, but there were different levels of quality created by the way air change rates were managed, the ceiling HEPA and ULPA filter layouts in different areas and plenum separations above the ceiling and in recirculation space in the subfloor.

7 The quality control (QC) room contained hoods with different acid baths used to etch wafers as a part of the quality control process. It was supplied with 5,000 cfm (2360 L/s) of air directly by the MAU. The process exhaust system removed 6,000 cfm (2832 L/s) directly via the hoods. The 1,000 cfm (472 L/s) difference was supplied by infiltration through three double-wide automatic doors from adjacent spaces. The pressure in the main cleanroom was controlled by varying MAU fan speeds and some manual balancing adjustments to create a "pressure cascade" from the cleanest areas to the dirtiest areas. Since the QC room was a dirty area, it operated at a slightly lower pressure than the surrounding areas, thus the infiltration through the entry door cracks.

8 The clean corridor surrounding the cleanroom along with the adjacent office area was the dirtiest place in the fab other than the mechanical rooms. If the cleanroom went negative, the air came from the clean corridor and any points of entry to the fab and leaks in the envelope.

Unfortunately, it took about 30 to 90 seconds to detect the failure and bring the lag fan online. That meant that for 30 to 90 seconds after the event, there would be no flow in the process exhaust system. That was more than enough time for various flow sensors and hazardous gas sensors to trip and "crash the fab."

In a crash, operating rules and safety protocols required that personnel in the fab evacuate immediately via the closest available emergency exit. Thus, 100 or so people in their cleanroom "bunny suits" would end up standing out in the parking lot waiting for someone to tell them it was safe to go back to work again.

To return to work, the team had to:

- Verify that the cleanroom was safe;
- Remove and discard the contaminated cleanroom attire;
- Don fresh cleanroom attire;
- Enter the cleanroom through the air lock, two people maximum, 15-second cycle time;
- Clean up the contaminated areas near the crash doors;

FIGURE 2 Recreation of the sketch I drew to help solve the fab crashes.

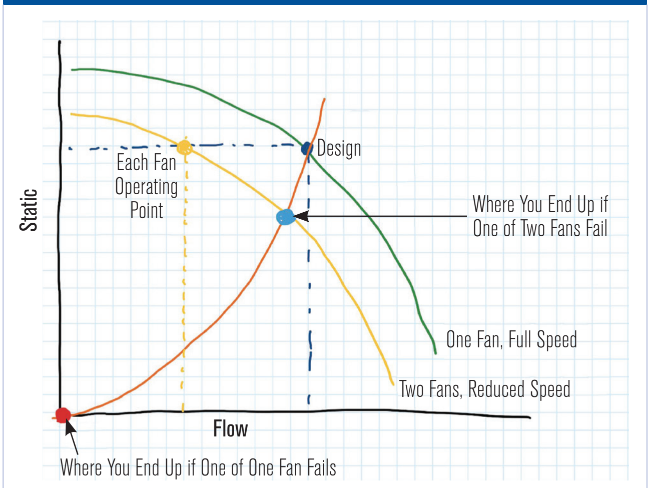

- Discard or recycle product that was in the contaminated areas; and
- Monitor and document cleanroom quality to verify that the space was back to production standards.

Bottom line: a crash cost about $10,000 when you took lost work, lost product, laundry and cleanup supplies into account.

N + 1 and Spinning Reserve

About one month into my tenure, after a week with one or more crashes per day, I found myself in a meeting with about forty people in attendance, including the person from corporate headquarters responsible for the project.

There was a lot of discussion about increasing the time delays on the safeties that triggered a fab evacuation and reducing the transition time required to bring on the backup fan. All of them had implications in terms of cost and schedule, and none of them ensured success.

As I listened to the discussion, it reminded me of a problem (of my own making) I had experienced with a hospital chilled water system. I solved it by creating spinning reserve (see "Spinning Reserve" sidebar). Recalling that experience, I started to wonder what would happen if we created spinning reserve in the process exhaust system. Thinking it through, I made a sketch similar to what is shown in *Figure 2* and showed it to my supervisor.

Right about then, the person from corporate headquarters, very diplomatically, asked if there might be an alternative control strategy that would address the problem. After encouragement from my supervisor, I screwed up my courage and showed the people the sketch. After a bit of discussion, it was suggested we go out and try it.

It was late in the day, and we were only running one shift, so the risk level was low. And if it worked, once the control algorithms were modified we would have a solution to the problem.

So off most of us went to the scrubber pad (*Photo 2*) and up the ladders to the platform where the exhaust fans and their VFDs were located. I was extremely nervous.

Spinning Reserve

Early in my career, a newly installed chilled water-cooled MRI machine would trip out when the lead distribution pump in the central plant I had designed and programmed was changed. That happened because I had not considered the dynamics of the system when I punched the lead-lag sequence into the control system.

The sequence looked at the run time on the lead pump, and if it had reached an even multiple of 1,000, shut that pump down and started the other pump. As a result, there was a brief loss of flow in the distribution system. Prior to the installation of the MRI equipment, the thermal inertia of the system easily masked the short loss of flow. But not so for the hypersensitive MRI flow switch.

The solution was to think about it like an experienced operator charged with making the change. In doing that, I realized I should start the lag pump and ramp it up while slowing down the lead pump. And, I should not shut down the lead pump until the lag pump had the situation under control (a lesson passed on to me by an experienced operator).

That process created what I now know is called "spinning reserve" and protected the system from a total loss of flow during the transition. Learning to think like an experienced operator can be a very valuable asset when writing control sequences.

PHOTO 2 The EPI scrubbers. The tall cylinder just right of center is the process exhaust scrubber. One of the two process exhaust fans is visible to its right.

After all, this was based on a sketch. I had not seen the fan curves, had no idea where the fans were operating on their curves, no idea where the surge line was and no idea if we would cross it when we went into the two fan operating mode.

I will never forget working with the operators to manually set up the scenario; bringing a second fan online and making educated guesses at the settings we needed Then, we tripped a fan off and held our breath. There were no alarms. Granted, we ended up at less than design flow with one fan running. But that flow was sufficient to prevent a crash, allowing the cleanroom to remain in production. Subsequently, we demonstrated that manually speeding the fan up restored design flow, all without crashing the fab. Adding a sensor and some control code allowed the process we had demonstrated manually to be automated.

N + 1 and Energy Efficiency

At the time, I thought this strategy would save energy; after all, the affinity laws said that running two fans to serve a given load at reduced speed will save energy compared to what it will take if you run one fan at full speed to do the same thing. Unfortunately, my thinking was abusing the affinity laws.

I was "clued in" when the kW readings from the VFDs showed higher consumption for the new operating mode. Contemplating what we had done in the context of *Equation 1* revealed my oversight.

$$kW = \left(\frac{Flow \times Static}{6,356 \times \eta_{Fan} \times \eta_{Belts} \times \eta_{Motor} \times \eta_{VSD}} \right) \times 0.746 \qquad (1)$$

where

kW = Input to the system to produce the flow and static pressure

Flow = Flow rate, cfm

Static = Fan static pressure, in. w.c.

6,356 = A units conversion constant that is good for air at approximately 0 ft to 2,000 ft (mean sea level) and between –40°F and 120°F (–40°C and 49°C)

η_{Fan} = Fan static efficiency

η_{Belts} = Belt efficiency; well-adjusted V-belts typically have an efficiency of 97% to 98%

η_{Motor} = Motor efficiency

η_{VSD} = Variable speed drive efficiency.

0.746 = Horsepower to kW conversion constant

If you consider the terms in the numerator, the flow term was unchanged from what it was for single-fan operation. And because both fans were connected to a common duct system (other than for some minor dynamic loss differences associated with splitting the flow at the fans vs. having all of the flow go to one fan or the other, and reducing the flow by 50% in the fan discharge stack) the static term was not significantly changed. Thus, energy savings would be solely due to improvements in the efficiency terms in the denominator.

For these fans (and many fans and pumps I have looked at where one machine was selected to serve the full load), the overall efficiency did not improve. *Equation 2* (which is based on *Equation 1*) and *Figure 3* illustrate this for both operating modes.

Fan Efficiency. The affinity laws assume that when you reduce the fan speed, you are moving down a system curve that runs through the selection point. But, in this case, the operating point shifts horizontally across the curve to a significantly less efficient location (*Figure 3a*).

Motor Efficiency. Motor efficiency improved slightly due to the slight rise in the motor efficiency curve. However, this could have gone the other way depending on where you started from (*Figure 3b*).

Before:

$$kW = \left[\left(\frac{28,000 \times 6.0}{6,356 \times 0.802_{Fan} \times 0.9800_{Belts} \times 0.9252_{Motor} \times 0.9600_{VSD}} \right) \times 0.746 \right] \times 1_{Fan} = 28.25 \text{ kW} \qquad (2a)$$

After:

$$kW = \left[\left(\frac{14,000 \times 6.0}{6,356 \times 0.669_{Fan} \times 0.9800_{Belts} \times 0.9350_{Motor} \times 0.9440_{VSD}} \right) \times 0.746 \right] \times 2_{Fans} = 34.07 \text{ kW} \qquad (2b)$$

Added Cost:

(34.07 kW – 28.25 kW) = 5.82 kW: For 8,760 hours per year = 50,983 kWh at 3.8 cents per kWh=$1,937

Drive Efficiency. Belt efficiency will be the same for either case. However, VSD efficiency will tend to drop as the load drops (*Figure 3c*).

N + 1 Bottom Lines

If the system had been 2*N*, then we would have saved energy. But, if the two-fan operating mode avoided one fab crash per year, we would see a net savings of about $8,000 a year, despite the spinning reserve energy penalty. Even at current utility rates, there would be significant savings. And, while modest, the service life of the motors, bearings and drives were probably increased due to the lower loads they experienced.

On the other hand, we would be wearing out the fans at the same rate vs. having the option of keeping one in reserve with less total run-time on it. But our operators were very rigorous in doing preventive maintenance. And, they could accomplish it by doing work on one fan and allowing the other fan to handle the load. Scheduling the work for periods of time when the fab was out of production minimized financial risk of a fab crash during maintenance.

Bottom line, I learned via experience that there is more to the various *N*+ relationships than math; system dynamics and definitions were critical, and nonenergy benefits may prevail.

That may sound like heresy from a tree-hugging, salmon-loving Oregonian. But our systems were there to provide a clean, safe, comfortable, productive environment in the facility. The energy consumption needed to be considered in the context of that broader mission. And the fans were not the only things consuming resources. The $10,000 per crash price tag ultimately represented an environmental hit in terms of wasted resources, human energy and embedded energy.

So, I would like to think that in addition to solving a financial problem, our solution was a holistic solution. If nothing else, the experience taught me to think in broader terms about how our systems use resources. Hopefully, this column has given you similar food for thought. ■

https://bit.ly/309ZpBE
Rate this Column

FIGURE 3 Efficiency penalty for running two fans vs. one fan.