It is a Matter of Balance: New Understandings of Open Plan Acoustics

It's been thirty years now. Thirty years since the first open-plan work spaces started replacing individual enclosed offices as a standard in the American workplace. Thirty years of progressively more compact workstations in progressively more densely populated work areas. Thirty years of office workers (as many as one in four, according to a recent British survey) 1 variously complaining about the acoustics in their workstations.

As the head of the ASTM (American Society of Testing and Materials) Subcommittee on Open Plan Spaces points out, in exchange for the improved communication afforded by these open-plan offices, organizations have had to face new challenges in acoustical privacy: "No one likes to sit in an office, trying to concentrate on critical words or numbers, with a racket coming from adjacent areas." 2 So acoustical experts have spent the last thirty years trying to figure out how to minimize the effects of any "adjacent racket" in the open-plan office--in short, how to achieve good speech privacy.

The more the experts learn, the more confusing it gets for everyone else--mostly because of all that technical jargon they use: "interzone attenuation," "sound transmission class," "noise reduction coefficient." Once that language barrier is surmounted, however, it's easy to understand why controlling the acoustics of the open-plan office has been so much of a challenge for so many years.

Rule of threes
Controlling open-plan acoustics is difficult, the experts have discovered, because they're never dealing with just one thing. As Michael Wodka, a product designer and consultant who's been involved in open-plan acoustics since the late 1960s, puts it, "There isn't just one problem, but three; and there isn't just one tool to use against those problems, but three (or four, depending on how you count); there isn't just one optimal solution, but at least three, involving combinations of at least three design elements. It takes a systems approach to sort it all out." 3

Sorting it all out requires starting with the problem--or rather the problems. Regardless of the particulars, the experts have determined that controlling open-plan acoustics always comes back to controlling the same three acoustical problems: sound level, speech intelligibility, and sound paths.

Sound level--it's not just loudness
If the acoustical problem in an open-plan office were simply a function of loudness, that would be one thing. But it's not. In determining where the sound level in an average office falls on the decibel (dB) chart (remembering, of course, that each additional 10 dBs doubles the loudness), something very interesting becomes obvious.

At around 60 dBs, the noise level in the average office is probably four times quieter than that in the crowded Paris bistro (80 dBs) where Hemingway says he wrote his best stuff, "one true simple sentence" at a time. 4 And who has not tried to work in a library (at 20 dBs, 16 times quieter than the average office), only to be rudely disturbed by someone whispering at the next table?

No, it's not strictly a question of loudness; it's a question of distraction. And the main distracters are "people noises"--two-way conversations, bits and pieces of phone conversations, throat clearings, squeaky new shoes walking by--anything that "piques the listener's curiosity." 5

Speech intelligibility--how much a casual eavesdropper can understand at what distance
For environmental psychologist Franklin Becker, it becomes a matter of "control of information flow" --an input problem of distractions and interruptions and an output problem of confidentiality. 6 Even though a normal conversation (65 dB) is only half as loud as a ringing telephone (73 dB) or a quarter as loud as a copy machine (81 dB), the "information content" 7 of normal conversation makes it much more distracting than the much louder equipment noise in an open-plan office. That information content also carries a long, long way: Even with carpet and a very sound-absorbent ceiling, a conversation held in an open space can still be easily understood at least 32 feet away. 8 So, to ensure that people can focus on doing

their jobs in an open-plan office, one has to find a way to control all the spreading sound waves from other people doing their jobs.

Workstation panels are an obvious solution.

Sound paths--over, under, around, and through whatever is put in the way
Because sound waves act like light waves--i.e., they spread out spherically, in all directions, on all planes at once--the most workstation panels can ever do is act something like a lampshade. No matter how large a portion of a sound wave a set of panels can corral, there will always be portions left free to leap over, creep under, bend around, or seep through those panels--continuing on their way in search of someone to disturb.

Applying different tools to each sound path
To keep as few "someones" as possible from being disturbed, acoustical experts have learned to attack sound waves on each of the paths those waves follow to get into adjacent workstations. (Leaving any one of the paths uncontrolled, they've found, negates what's been achieved with the others.) 9 These paths are as geometrically precise as the angles followed by a billiard ball: There are two bending (diffracting) paths--around the sides and over the top of a panel. In addition, there are two reflecting paths--off the ceiling and off vertical surfaces like walls or other "flanking" panels. There is also one direct path (other than the obvious line-of-sight path from open doorway to open doorway)--through a panel.

The experts have found two ways to keep sound traveling on these paths from disturbing people in adjacent workstations: 1) change the angles of the two diffracting paths (so that they miss the people in adjacent workstations), or 2) lower the sound level and speech intelligibility of the sound traveling on reflecting and direct paths (so the sound is no longer distracting).

The first is fairly easy to accomplish; the second is much more difficult. If one expects to achieve any level of speech privacy at all, both must be accomplished. The only way to do that is by applying several different tools.

Increased acoustical shadow--for diffracting paths
By increasing the height and width of dividing panels or by adding hang-on storage components, the angles of diffracted sound waves (i.e., those bending over and around panels) can be flattened. More of the sound traveling on those flattened angles now misses the heads (and therefore the ears) of people seated in adjacent workstations, increasing the "acoustical shadow" 10 in those workstations.

The effect of these acoustical shadows can be multiplied by moving people closer to their panels; however, this will increase the distraction from sound waves coming directly through the panel. 11 It also will do nothing to protect those people from sound waves bouncing around on reflecting paths (which is what any kind of sound wave does once it encounters a hard surface like a floor, a wall, or a hang-on component on the other side of the workstation).

Direct and reflecting sound waves, therefore, must be handled using the second method--lowering their sound level and speech intelligibility. To do that, the acoustical experts have come up with three tools: sound absorption, sound blocking, and sound masking.

Sound absorption--for reflecting and direct paths in larger workstations
Everyone knows how rugs and pillows, thick curtains, and overstuffed chairs can take the echo out of a big, empty room. Like a thick paper towel soaking up water, these porous/fibrous, sound-absorbing materials soak up sound waves by trapping them in an internal maze of air pockets. As sound waves work their way through that maze, the sound energy they carry dissipates--turning into microscopic heat. 12 When they emerge from that maze, sound waves are weaker, and therefore less loud.

Sound absorption was all anyone knew to use in the first open-plan offices. By adding sound-absorbent materials (e.g., fiberglass, shredded wood fiber, paper honeycomb) to the structure of ceilings and walls, the experts learned early on how to weaken reflecting sound waves.

Sound absorption also worked, they learned, to weaken sound waves following a direct path through panels--as long as people in adjacent workstations were more than 12 feet apart. At that distance, direct sound waves are already sufficiently weakened by the distance traveled between workstations to make sound absorption effective. 13

Sound blocking--for direct paths in smaller workstations
As workstations have become more compact, however, people are often working (and talking) within only a couple of feet of dividing panels. Under these conditions, sound waves traveling on direct paths are too strong when they reach the panel to be significantly weakened by passing through its sound-absorbent maze. Such direct sound waves penetrate the panel, bringing easily understood conversations into adjacent workstations. 14

The only way to prevent this direct penetration, the experts have learned, is to add sound-blocking materials to the dividing panels. 15

Whereas the best sound absorption relies on lightweight, porous materials, the best sound-blocking materials (solid masonite, metal, or hardboard) are dense and heavy. With no air spaces for sound waves to slip into and through, panels containing these tightly packed interiors essentially cut off the direct path into adjacent workstations.

Sound masking--to cover up whatever sound is left over
Even with the best sound absorption and sound blocking (in ceilings, walls, and panels), there will always be some sound left bouncing around the open-plan space. By adding a sound-masking system, the experts discovered they could cover up any distracting noises that aren't absorbed or blocked.

It's similar to turning on a water faucet when someone is speaking in another room. Before the faucet is turned on, a person can hear and respond to the words just fine. That's because there's a big difference between the sound level of the "intrusive speech" and the ambient backgroundsound of the room. When the faucet is turned on, however, the ambient background sound level increases, reducing the ear's sensitivity to the intrusive speech. One may actually still hear the speaker's statement, yet not consciously acknowledge it. If the water is turned up even more, one may remain aware that someone is talking, but won't be able to understand what's being said, no matter how hard one tries. 16

The continuous, specifically tuned sound signal used in today's dedicated electronic sound-masking systems is carefully designed to render speech unintelligible, without being distracting in itself. 17

Balancing the four tools in compact workstations
The experts have known since the early '70s that it takes all four of these tools (increased acoustical shadow, sound absorption, sound blocking, and sound masking) to control the acoustics in the open-plan office. As workstations have shrunk, it has become more and more important (and difficult) to balance the four of them.

That balance is delicate; no gains come in one area without trade-offs in another.

Compromising sound absorption with hang-on components
While hang-on components like flipper door units can increase the acoustical shadow of panels, they also cut down on the sound absorbency of those panels. Clearly, panels with no hang-on components have a lot more surface area available to absorb sound waves than do those that carry flipper door units, work surfaces, work organizers, tackboards, and filing drawers or cabinets. 18

As workstations have become more compact, these hang-on components now cover a larger proportion of a shrinking amount of panel surface, undercutting the absorption capabilities of the panels that carry them. (Work surfaces, in particular--positioned as they are between the speaker's face and the lower portion of panels--render those lower portions totally ineffective as acoustical elements.) 19

Sound-masking system undercut by increasing sound absorption
One might think, as the experts originally did, that just beefing up the sound absorption in panels could counter these "counter-absorbency" effects of hang-on components. An odd thing happens, however: Sound waves from the sound-masking system get soaked up too fast. This leaves the space without sufficient ambient noise to cover up the speech that is still penetrating the panel. That speech becomes even more noticeable (and disturbing) to people in adjacent workstations. 20

Sound blocking more critical in compact workstations
The surprising thing the experts have learned about sound absorption in panels is this: Rather than having a critical effect on speech privacy in more densely populated spaces, highly absorbent panels mainly affect the sound-masking system and reverberations (from reflecting sound paths). To achieve good speech privacy in compact workstations, the experts have learned, panels must function primarily as barriers 21 --that is, they must have increased sound-blocking capabilities.

This is not to say, of course, that panels need no sound absorption at all. Because sound-blocking materials simply send direct waves back where they came from, using sound-blocking materials alone creates further acoustical challenges (in the form of reverberations 22) back in the original workstation. The more tightly packed the barrier, the more sound waves will be reflected back into the original work space or go careening off into others. 23

To counter this phenomenon, acoustic experts have learned to design panel cores out of a combination of sound-blocking and sound-absorbing materials. This mix of sound absorption and sound blocking helps keep people on both sides of the panel more acoustically comfortable.

No "one-size-fits-all" solution
Even as the acoustical experts were figuring out the appropriate balance among acoustical shadows, sound absorption, sound blocking, and sound masking for a

hypothetical open-plan office, they were starting to recognize the futility of trying to come up with a solution that would work in every open-plan office.

That's because open-plan spaces can vary so much in architectural and design elements--among them, room dimensions; ceiling height and structure; placement and structure of lighting fixtures; floor and wall coverings; and number, size, and placement of windows and doors. 24 Most of these elements are usually set by the time anyone starts thinking about acoustics. And yet, changing any one of these variables changes the mix of potential diffracting, reflecting, absorbing, and blocking surfaces. Changing that mix changes what's ultimately feasible in terms of good speech privacy.

With so much to consider and so little to control, is it any wonder that as late as the mid-1980s acoustical experts were recommending four-sided workstations with the tallest panels, in a valiant effort to squeeze the most speech privacy out of otherwise "acoustically challenged" spaces; or that they went a little overboard warning facility managers not to let people hang calendars, posters, photographs, or other reflecting surfaces on panels, "as this may lead to a lack of coordination in the designed acoustical . . . elements." 25

Fitting the solution to the acoustical need
By the time the '90s rolled around, with workstation configurations becoming more and more open, the experts finally conceded that it is often not feasible (either economically or physically) to give every area in an open-plan space the maximum level of speech privacy. This concession has led to a three-tiered definition of speech privacy 26 --one based on how well a given space is able to control speech intelligibility.

Experts concede that interzonal speech privacy is extremely difficult to achieve in the open-plan office. In addition to a uniform sound-masking system that perfectly covers speech frequencies and a ceiling that absorbs reflected sound waves "as effectively as the sky," this level of privacy requires panels that are high enough (60 to 80 inches 27) to create deep acoustical shadows, and panels that absorb nearly as well as the ceiling, and block so effectively that what little sound does leak through can be easily covered by the masking system. 28

Zonal and regional privacy are clearly more within reach.

A systems approach
Applying this three-tiered definition, acoustical experts have learned how to create different sections of an open-plan space to accommodate different groups' differing needs for speech privacy. They've also learned how to give individual people the differing levels of speech privacy they may need for differing activities in different spaces--as exemplified by Fritz Steele's proposed "cave" (private) and "court" (public) areas. 29

Both of these activities require a systems approach: involving end users, facility managers, designers, specifiers, contractors, and manufacturers up front in choosing the appropriate ceiling, masking system, and panels. Without such an approach, the experts have learned, you'll end up with "an eavesdropper's paradise." 30

As they worked out the trade-offs involved in these different levels of speech privacy, the acoustical experts have learned something else: how little acoustical difference panels make if the other required elements are not present.

Panels—not the most critical factor
Although the experts have determined that it's doubtful someone can get good speech privacy without a properly designed panel, 31 they also recognize that panels are not the most important factor: Ceiling height, ceiling absorbency, and the sound-masking system are all more critical to speech privacy than are panels. 32

As product designer Wodka puts it: "If you haven't got a good acoustical ceiling and sound-masking system, there's no point in putting a lot of money into acoustical panels at the end. The most you can expect to do with panels is to fine-tune what you've already accomplished with the ceiling and background sound system." 33

This research summary has already shown how little fine-tuning the absorbency of the panel can affect speech privacy. Now that the general trend in open-plan offices is toward lower panel heights, there are also limits to how much one can fine-tune the sound blocking in a panel. The ASTM committee on open-plan spaces has determined, in fact, that anything lower than 60 inches is "not effective" as a sound barrier. 34

The chart on this page shows how little can be achieved from adding a good acoustic panel to a space that lacks a good ceiling and walls (shown in the two columns on the left). In fact, the effect is almost nil--reducingthe original 58 dB sound level by an imperceptible 2 dBs and leaving the occupant "extremely dissatisfied." The columns on the far right, however, show a big payoff for using masking, plus adding good panels to good walls and ceiling. The original sound level is halved (with more than a 13 dB reduction), bringing the level down to 42 dBs and moving the occupant into the "satisfied" range. 35

The key to choosing the correct panels for a particular open-plan space, then, is to be very clear about which level(s) of speech privacy one is trying to achieve. Equally important is to remember that, although panel absorbency will help with reverberations in the space as a whole (and is, therefore, a necessary part of the whole acoustic equation), the sound-blocking ability of panels is more critical to speech privacy and should, therefore, be the first priority. 36

In the meantime, if the workstation configuration needs other attributes from panels (e.g., tackability, visual access via window frames, work-organizing components), there's no harm in addressing them. In most cases, they won't make that big a difference in the overall acoustics.

The confusion about panel rating systems
So, how does one know how much sound absorption or sound blocking a particular panel has? Usually by checking out its ratings.

Here's where things can get confusing. There are several different rating systems, all measuring different things, in different ways, on different scales, and all with names that are so arcane everyone refers to them by their acronyms.

Two rating systems have been around since the 1940s: for sound absorbency, NRC (Noise Reduction Coefficient); and for sound blocking, STC (Sound Transmission Class).

These two systems were originally designed to rate architectural building materials during World War II. For lack of anything better, the furniture industry adopted them in the '70s to compare the acoustical performance of materials used in ceilings, architectural walls, and panels.

Both have their drawbacks.

The misleading nature of NRC ratings
Perhaps because of its more intuitive scale (it runs from 0 to 1.0) the NRC (absorbency) rating predominated as an industry favorite far into the 1980s. This predominance led manufacturers to overbuild panels for high absorbency, at the expense of sound blocking--long after the experts had determined that the STC (blocking) rating was a much better predictor of speech privacy. 37 While the NRC rating is fine for ceilings, it can be misleading for panels, for several reasons.

First, because it tests panels with no hang-on components, the NRC rating indicates more sound absorption than can actually be delivered by the panel once components are added. 38 Second, it's impossible to tell from the NRC rating how well a panel will absorb higher (speech) frequencies. Therefore, a person could get better speech privacy performance from a panel with a lower NRC rating than from one with a higher rating, simply because the first panel more effectively soaks up higher frequencies. 39 Third, a panel with an NRC rating of .80 absorbs about 5 to 6 dBs more than one with an NRC rating of .50. For most ears, that's an imperceptible difference. 40 Fourth, and perhaps most important, the panel height, the distance between speaker and listener, the sound-masking system, and the panel's sound-blocking ability are all more important to speech privacy than is panel absorbency 41 --the only thing the NRC test measures.

Taking all this into consideration, one can see that an NRC rating may not really tell much about how well a panel will perform in any one particular open-plan space. Further, the difference in sound reduction between an NRC rating of .50 and one of .80--at 5 to 6 dB--is barely a 25 to 30% reduction in sound level. And, the most sound reduction one can expect from sound absorption in the open-plan office is 20 dBs. That leaves a sound level of 45 dBs that can be controlled only by a good sound-masking system. 42

STC rating: better, but still not "real-world"
Although the STC rating is better than the NRC at predicting a panel's contribution to speech privacy, it still doesn't tell the whole story.

Like the NRC system, the STC system tests panels in isolation in reverberating chambers, using a much broader spectrum of sound frequencies than that encompassed by human speech. 43 Because it measures how much sound passes through a full-height panel (i.e., a floor-to-ceiling wall) dividing the test chamber in two, the STC test does not tell much about how well a shorter panel (with gaps underneath and at joints) will block sound. 44 Because the test chamber has nonabsorbent walls and ceiling, the STC test (like the NRC test) assumes the panel will have multiple chances to interact with bouncing sound waves. In a real open-plan space, however, "spreading sound must be controlled immediately, since the critical relationship is between two adjoining work zones." 45

For these reasons, the experts criticize both the NRC and the STC systems for "not measuring the real world," 46 and have, therefore, turned to interzonal rating systems as a more relevant predictor of speech privacy in the open-plan office.

Interzonal tests to simulate the open-plan office
As has already been discussed, one cannot expect good speech privacy in an open-plan office without coordinating several things: ceiling and wall absorption, the sound-masking system, and sound absorption/blocking in panels. This systems approach is carried over into the interzonal rating systems, which use human speech frequencies to test the speech privacy capabilities of panels--in simulated workstation configurations--in conjunction with a sound-masking system and absorbent walls and ceiling.

As product designer Wodka points out, because these interzonal tests consider the effects of sound absorbency and sound blocking in determining a speech privacy rating, what an interzonal rating provides, in effect, is a combined NRC and STC rating. And, because it tests how all the components of a particular open-plan space function together, the interzonal system allows acoustical designers to observe the effects of tradeoffs--e.g., how much can be gained by cranking up the sound-masking system (say, to 47 dB) to compensate for not-so-terrific blocking capabilities in a panel. (In reality, it may be able to slash what a person in an adjacent workstation can understand from 34 percent to 20 percent.) 47

NIC', AC, and AI ratings to measure speech privacy
There are three basic interzonal rating systems. The first one-NIC' (Noise Isolation Class, with the ' pronounced "prime")--was developed in the late '70s by the Public Building Services (PBS) branch of the government. This interzonal rating system has three subtests: NIC'b rates panels in a barrier mode (perpendicular to the sound path); NIC'f rates panels in a flanking mode (at acute angles to the sound path); a third test rates panels in a mock-up of a pair of typical workstations (including hang-on components). Together, these three NIC' tests measure the speech privacy afforded by panels in a typical workstation configuration. 48

The second interzonal rating system--AC (Articulation Class)--was developed in the late '80s by the ASTM as a modification of the NIC' rating system and was finally adopted in 1991 as an industry standard. AC test setups are essentially the same as NIC' setups. Only the rating scale differs. 49

Different manufacturers may choose to use either NIC' or AC tests on their panels. Because the scales used by these two rating systems are anything but intuitive, however, the experts have figured out a way to correlate them to a third rating system--the AI (Articulation Index)--which is a little easier to interpret.

Developed in World War II as a way to define how well headphones and intercoms articulated speech on ships and planes, the AI test originally consisted of one person reading a list of words to another to see what percentage of words could be understood. Whereas it's ideal for headphones and intercoms to have a high AI (indicating that most of what's said can be understood), an open-plan office should have a low AI (indicating that not much can be understood). 50 The ASTM adopted an electronic simulation of the original AI test in 1991 as a way to verify the speech privacy of a particular open-plan space--either in the field or in a mock-up. 51