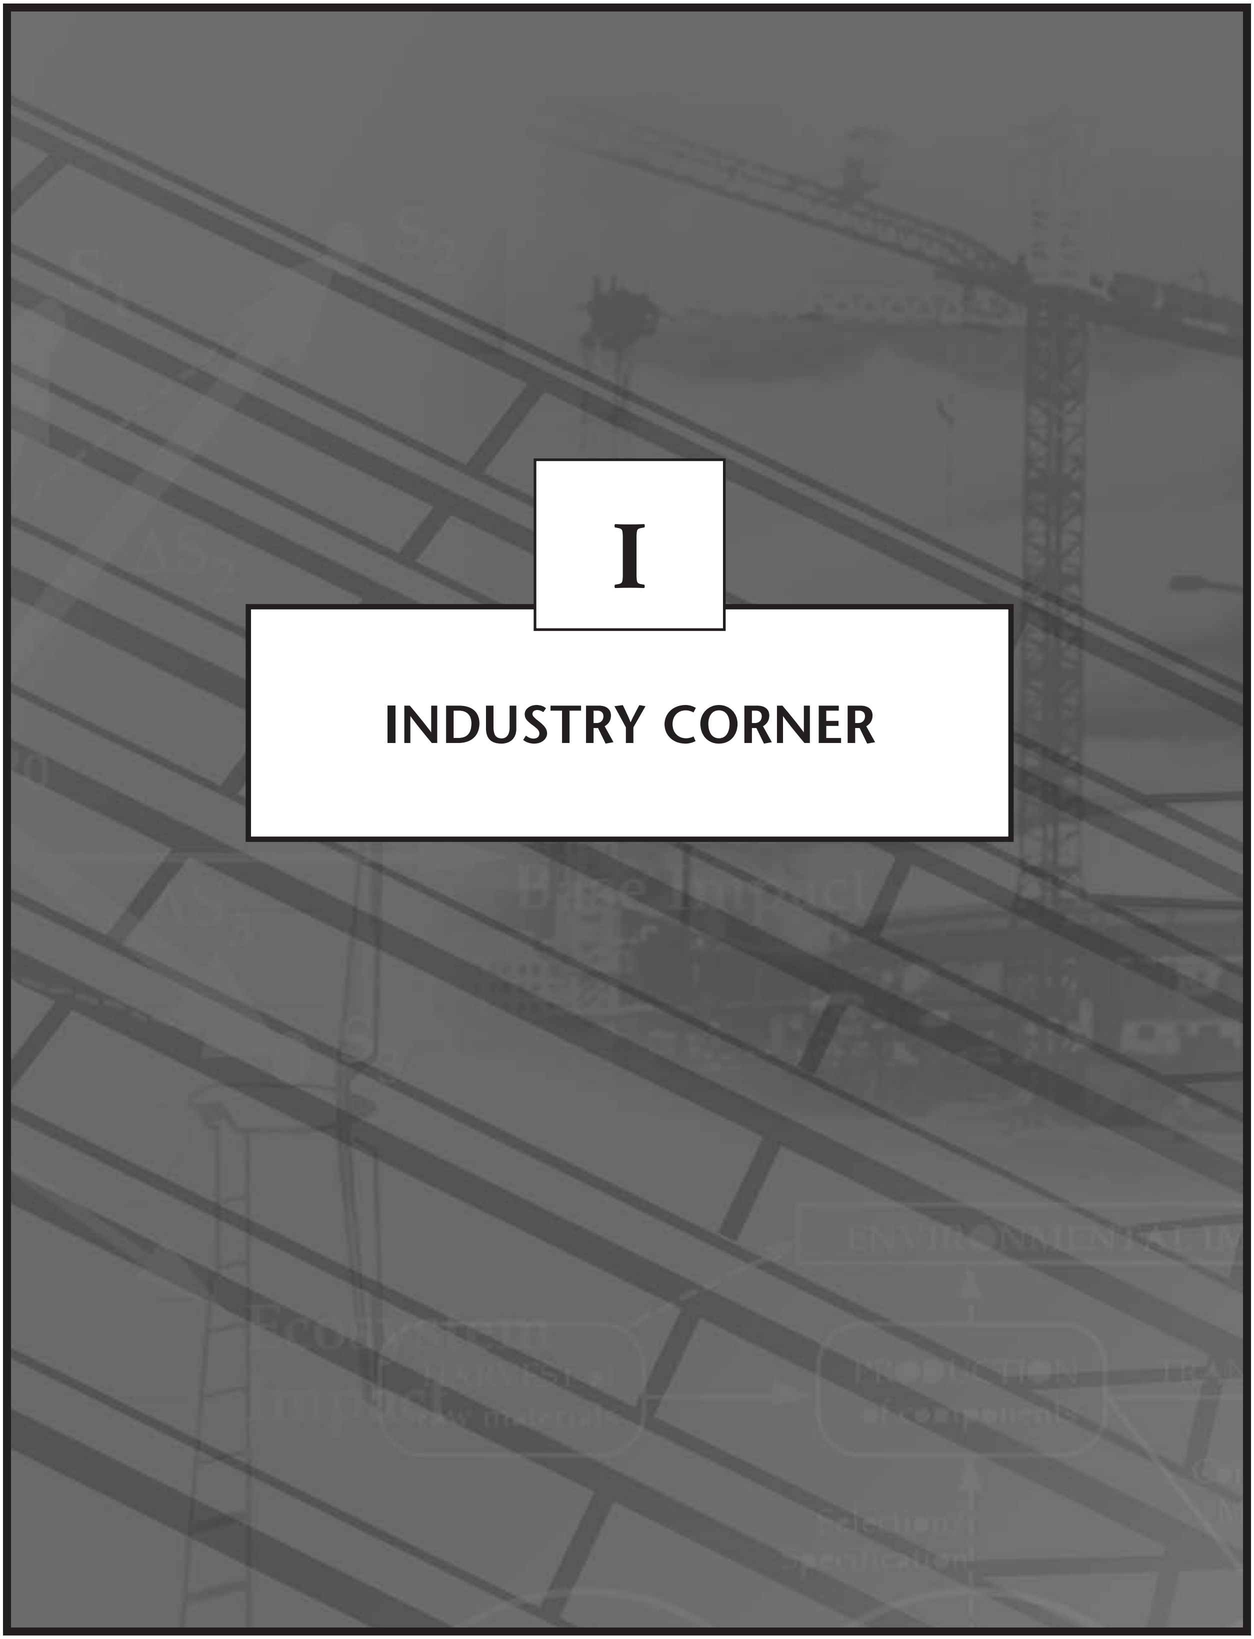

I

INDUSTRY CORNER

INTRODUCTION TO GREEN WINDOW DESIGN AND PERFORMANCE

Alex McGowan[1]

INTRODUCTION

Green building design can be a complex process, and windows are an important part of that process. A wide range of performance parameters can make selecting the right options difficult, and some parameters provide conflicting performance trade-offs. Good information about how to design green windows is not readily available, as most people look through *windows, but not many people look* at *them. Designed correctly, windows can greatly enhance the occupants' enjoyment of the space and enhance the energy performance of the entire building. Designed* incorrectly, *windows can be a source of irreparable failing of the building to meet the occupants' requirements for comfort, energy efficiency, and long-term durability of the building envelope.*

This article discusses the physics behind window performance, describes some of the design parameters, and explains how to determine whether you can get the performance you want on site. This is not meant to be an exhaustive reference on the subject, as that would require several books. Rather, the intent is to provide an overview of the physical behavior of window components, so that appropriate design choices are based on sound knowledge and not guesswork. This article also provides some fluency with window-performance terminology for use in discussing design options with specialists.

PHYSICS OF GLAZINGS

To discuss these aspects of window performance, and the parameters used to measure them, let's start with the source of all light and heat (Figure 1). The sun emits a spectrum of electromagnetic radiation comprising a wide range of wavelengths (Figure 2). These wavelengths are measured in nanometers (nm), or billionths of meters. If we graph the full range of wavelengths that are emitted from the sun, and that appear to an observer in the International Space Station above the atmosphere, we would obtain the upper curve in Figure 2. As the sunlight passes through the atmosphere, some of the energy is absorbed and the energy that reaches the earth's surface is shown by the middle curve in Figure 2.

This curve comprises high-energy ultraviolet light, which has very short wavelengths, visible light (electromagnetic radiation in the range of 300–700 nm, which is what our eye recognizes), and long-wave infra-red radiation, which we sense as heat energy in the longer wavelengths.

WINDOW PERFORMANCE PARAMETERS

For daylighting, you want windows to transmit as much visible light as possible. The property that characterizes this performance is visible transmittance (VT or τ_v), a decimal quantity (i.e., between 0.0 and 1.0), representing the fraction of visible light incident on the outer surface of the glass that is transmitted into the building. Ideally, visible transmittance should be as high as possible, but all materials reduce light transmission. Even a single layer of "clear" glass has a visible transmittance of approximately 0.90, and a double-glazed window is approximately 0.81.

To provide linkage between the occupants and the outdoor environment, you would like the visible light to be transmitted with the truest possible color rendition, so that the outside view looks like it should, and interior colors take on their proper hues in natural light. This is actually built into the way in which VT is determined. In the NFRC 200 standard that defines VT, specific values are measured at each

[1]P.Eng., LEED® AP, Levelton Consultants, Building Science Division, 760 Enterprise Crescent, Victoria, British Columbia Canada V8Z 6R4, amcgowan@levelton.com.

FIGURE 1. Source of solar gains and visible light.

FIGURE 2. Solar and visible spectrum.

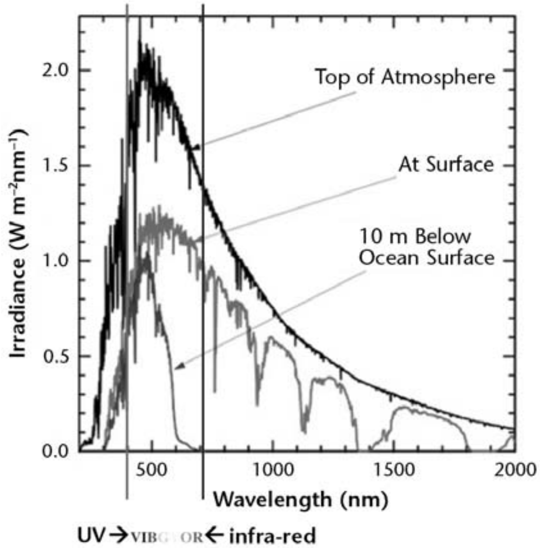

wavelength of light and combined into a weighted average value. The weighting of individual values (to develop the average VT) is based on the visual response of the human eye, as defined by a "Standard" observer, so it favors the blue-green part of the spectrum. Thus, the single averaged value is generated in a way that represents true color rendition.

The solar heat gain coefficient, or SHGC, is another decimal quantity (i.e., between 0.0 and 1.0) representing the fraction of sunlight incident on the outer surface of the glass that is transmitted into the building. A high SHGC will provide free solar energy to a building to reduce the need for heating load, whereas a low SHGC will reduce solar gains in a building that already has a high cooling load. Designers used to rely on a related quantity called the shading coefficient (SC), which is the SHGC of a window relative to the SHGC of a single layer of glass. Eventually, researchers and designers decided to take the single layer of reference glass out of the picture and just use the SHGC (not many designers use single glazing any more).

The U-factor or U-value is the heat loss coefficient, which represents a rate of heat transfer per unit area, per degree of temperature difference. Units are Watts per square metre per Celsius degree, or BTU per hour per square foot, per Fahrenheit degree, depending on the favorite system of units. U-factor is the inverse of R-value (U = 1 ÷ R), so higher U-factor implies a lower R-value (i.e., less resistance) and therefore higher heat transfer. Heat transfer happens by *conduction* (in solid materials), *convection* (via movement of a heated or cooled liquid or gas coming into contact with a surface), and *radiation* (radiant energy transfer between surfaces through a gas or vacuum). The path to high-performance windows lies in reducing these three modes of heat transfer.

Higher U (or lower R) means higher heat loss, which means auxiliary heat (what we pay for, with its associated greenhouse gas and carbon costs) is required to maintain the room air temperature within a comfort zone for the occupants. But there is a secondary effect: higher heat loss means that the room-side surface of the window is colder during winter conditions. This can cause condensation-related problems, as the water vapor in the room air condenses on the colder surfaces. This is an important consideration: condensation is bad because the water that condenses out of the room-side air can ruin interior finishes, support mold growth (which may affect immuno-compromised or allergic occupants), and generally impair indoor air quality.

If that were not enough, there is a tertiary effect of cold windows. The cold window surface impairs occupant thermal comfort, and it does that in three ways:

- First, the air next to the cold window is cooled via conduction. As the air is cooled, it becomes denser as it falls down the surface of the glass, and getting cooled more as it travels along the colder surface. By the time it gets to the bottom of the window, the air can have achieved an appreciable velocity, so that anyone sitting next to the window would feel the air movement and interpret it as a drafty window. And they would be right; the window does not need to be open, or to allow air to leak into the room for a draft to have occurred in this way. The phenomenon is called *induced draft*, and the result is that the occupants feel cold—so they will go turn up the thermostat, as long as they have the ability to do so.

- Second, the cooled window presents a cold surface to the room. The occupants of the room (most of whom should be considered warm-blooded mammals) are all radiating energy at a rate characteristic of their surface temperature, which is typically around 22°C. The cold window surface will result in a higher radiant exchange between the occupants and the window, which the occupants (most of whom should be considered sentient beings) will interpret as a chill—so they will go turn up the thermostat, as long as they have the ability to do so.
- Third, the geometry of most rooms is such that windows are predominantly on one wall, so the radiant loss to the cold window surface is not balanced by radiant losses to other similar surfaces. The occupants will interpret this asymmetric radiant heat loss as another type of thermal discomfort—so they will go turn up the thermostat, as long as they have the ability to do so.

The effect on thermal comfort related to cold windows, as noted above, is the main reason that heating systems often deliver heat to the rooms (using registers for forced-air systems, or radiators in hydronic systems) below the windows. In this way, warm air rises from the register to meet the cool air falling off the window. The air streams mix, and tempered air is delivered to the room. This mixing also helps to reduce the potential for condensation to form at the bottom of the window.

Research has been done to suggest that high-performance windows, and their resulting warmer window temperatures, would allow designers to move the heat-delivery devices away from the exterior walls. This would keep more of the heat in the room, thus increasing the overall efficacy of the system, and would reduce the amount of piping or ductwork. Apart from reducing construction costs, less piping or ducting means less likelihood of leaks, lower maintenance costs, and a greener design.

WINDOW DESIGN OPTIONS

Now that we know about U-factors and their units, and we know about SHGC and visible transmittance (which have no units), and the various modes of heat transfer and what they mean to occupant comfort, we have some of the tools to examine window options and how they affect window performance.

Reducing *conductive* heat transfer in the solid components comprising the frame of a window or curtain-wall assembly is done by using materials with lower thermal conductivity (e.g., vinyl or wood window frames instead of metal) or by reducing the surface area through which the conductive heat transfer occurs. Wood or vinyl windows are typical alternatives for metal-framed windows in residential design, especially single-family or low-rise multi-family dwellings.

All of the materials used for framing have advantages and disadvantages, when viewed from a green perspective. Aluminum has a very high embodied energy: the cost of mining, smelting, and extruding make this an energy-intensive material. On the other hand, it is easy to recycle and has a very high strength/weight ratio, so less material is needed to achieve a desired structural capacity. Vinyl is somewhat less energy-intensive, but it is derived from a petroleum base, so vinyl-framed windows are not exactly carbon-neutral. On the other hand, vinyl produces a reasonably well-insulating frame, as it encapsulates air cavities, and so can provide lower long-term energy consumption (and operating costs). Also, vinyl frames are very easy to recycle, although they typically require virgin material at manufacture, because any recycled-material content would make it difficult to produce the characteristic white frame. Pultruded fiberglass frames are also very energy-efficient, but cannot be recycled and contain petroleum-based polymers in the resin matrix.

Wood framing is a clear winner as a green material. The embodied energy is relatively low (mostly the cost of transportation, although the finishing work involves milling, which is moderately energy-intensive (although obviously not as much as, say, aluminum). Wood represents locked-in carbon, and as long as the window and its installation are properly designed to prevent decay, the carbon can stay locked in for a long time—on the order of 100 years, with proper maintenance. On the other hand, wood is not a suitable material for high-rise or non-residential applications where structural loads can be significant and non-combustible construction is required by building codes.

If structural, fire-safety, or aesthetic considerations dictate that the window frame must be metal, we can still replace part of the metal with a low-conductivity material such as nylon, urethane, or vinyl; this interruption in the conductive metal frame is called a *thermal break*. The more metal (that is, metal that forms part of the conductive path) that is replaced by low-conductivity material, the better the thermal performance of the framing system. A common curtain-wall assembly features a simple thermal break, formed by the vinyl gasket between the body of the frame and the pressure plate that holds the glazing in place (Figure 3). This gasket is only 3 mm thick, and so does not displace very much metal. An improvement on this design is the "*pour-and-debridge*" thermal break (Figure 4), so-called because liquid urethane is poured into a pocket in the metal frame, then the urethane is allowed to cure, and then the back of the metal pocket is milled out, or *debridged*. This replaces more of the metal (typically 6–8 mm) than is the case for the simple vinyl-gasket design. Even better is the "dog-bone" thermal break, a separate nylon or vinyl component that can be 25, 50, or even 75 mm long.

It might appear that we are chasing small details in this discussion of thermal breaks, but consider this: the frame of a typical curtain-wall assembly comprises only about 10–15% of the total area, but accounts for 50% of the heat loss through the assembly.

Reducing conductive heat transfer in the edge-of-glass region of the window involves replacing the metal in a typical spacer—the device that keeps the layers of glass separated—with a lower-conductivity material such as silicone foam, vinyl, or butyl (see Figure 5).

These so-called warm-edge spacers slightly reduce the overall heat loss of the window assembly. In other

FIGURE 3. Simple thermal break.

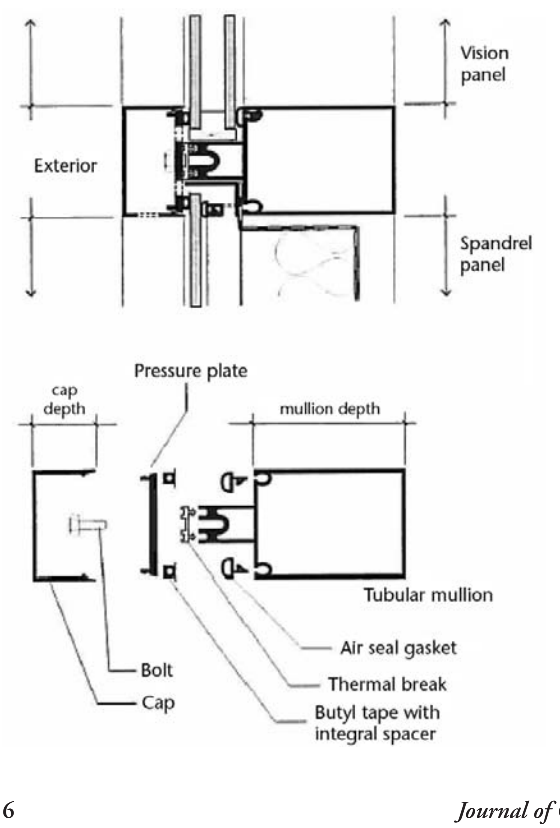

FIGURE 4. Pour-and-debridge thermal break.

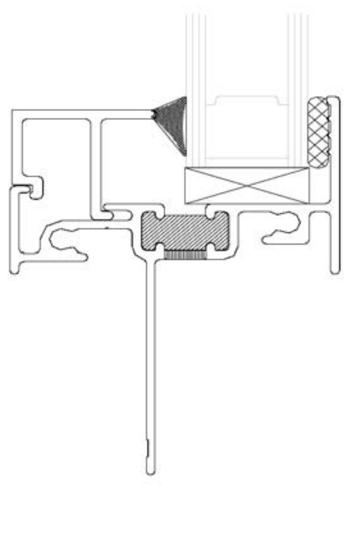

FIGURE 5. Various types of edge-glass spacers.

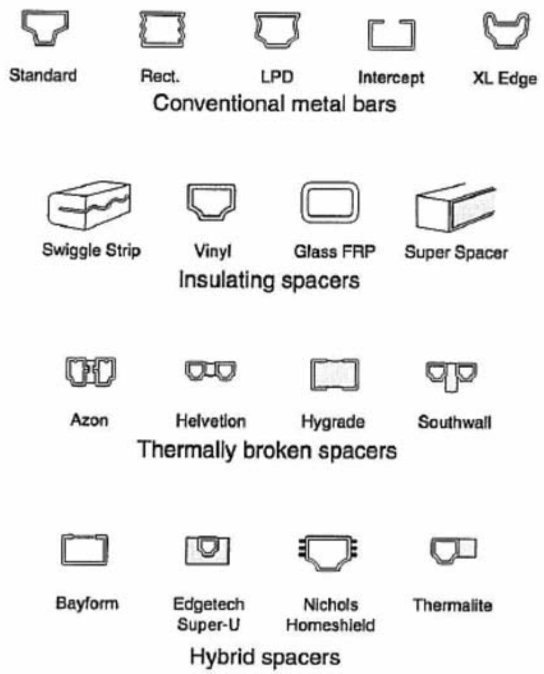

words, they reduce the total-window U-factor, but not by much, because the spacer comprises a small percentage of the total area of the assembly. They do a good job of reducing the local heat transfer in that small area, however, so the local surface temperature at the edge of the glazing is warmer in windows with warm-edge spacers, so that the potential for condensation at this location is reduced.

Other warm-edge spacers designs insert a thermal break between metal components or increase the conductive path length, but all are attempting the same trick: reducing conductive heat loss around the perimeter of the glazed area of a window.

Reducing conductive heat loss in the glazed area is typically done by replacing the air in the glazing cavity (i.e., between the layers of glass) with a low-conductivity gas such as argon or krypton. This also reduces *convective* heat loss, as these gases are more viscous than air.

The use of argon-gas fill is debated in some circles: it is a low-conductivity gas, and it does reduce conductive and convective heat loss, but there is some concern that the argon molecules can more easily leak out of the sealed unit (because argon molecules are smaller than nitrogen or oxygen, the principal constituents of air). Over time, the argon concentration in a sealed unit can be depleted to the point that it is no longer providing its intended function.

As fill gases are invisible, it is difficult to know when the argon has leaked out of the assembly. Typical estimates are from 5 to 7 years for an appreciable concentration of argon, but not enough research has been done to verify this estimate.

It is possible that the argon might not leak out of a particular sealed unit for quite some time (and maybe not at all), so that the presence of the gas can help save energy for several years. Eventually one should expect that the sealed unit will fail—the seal is only warranted for 10 years—which will result in loss of the inert-gas fill, so the increased cost of the gas-fill option should be weighed against the energy savings over the expected service life of the sealed unit. When considering the overall green benefits of an inert-gas fill, one should also factor in the significant amount of energy required to separate the fill gas out of air in the first place. The life-cycle cost of this measure should be weighed against the energy savings for the specific application, as a result of reduced conductive and convective heat loss.

Convective heat transfer is also used to beneficial effect. As Figure 3 shows, the surface area on the interior of a typical curtain-wall assembly is much larger than the surface area on the exterior. This is intentional: the increased exposure on the interior will increase convective heat transfer between the warmer room air and the metal surface, while the reduced area on the exterior controls the amount of heat loss from the assembly to the exterior. This allows the room-side air to warm the curtain-wall assembly, and reduces the potential for condensation to form. The penalty is increased heat loss at the curtain-wall frame, but the chance to reduce condensation in a passive way (i.e., without adding auxiliary heating) is generally considered to be worth that price.

GLAZING SELECTION

The main factor defining the energy performance of a window or curtain wall assembly is, of course, the selection of the glazing. The glazing affects radiant heat transfer, and therefore defines the SHGC and

VT of the assembly, and has a dominant effect on the U-factor as well. Let's look at how that works, as illustrated in Figures 6 and 7.

First, energy from the sun strikes the outer layer of glass at surface #1 (surfaces are always numbered from the outside inward). Depending on the solar-optical properties of the outer layer of glass, some percentage of that energy is reflected off the outer surface, and the rest is transmitted inward. Some of this energy is absorbed in the outer layer of glass (more, if the glass is specifically formulated to be heat-absorbing), but most of the energy is transmitted inward. This inward-flowing energy next encounters surface #2 (the inner surface of the outer layer of glass).

Once again, some of the energy crossing this surface is reflected back (this time, back into the glass), and the rest is transmitted into the glazing cavity. The gas in the cavity is typically not dense enough to absorb an appreciable amount of energy, so all the energy that leaves surface #2 arrives at surface #3, the outer surface of the inner layer of glass (let's just assume a double-glazed window for now, to keep the explanation simpler, but the principle would hold true for multiple glazings). The same interplay of reflection and transmission occurs, resulting in some fraction of the radiation that was incident on the exterior surface being transmitted to the interior space; this fraction, expressed as a decimal, is the SHGC. The SHGC of a clear single glazing is approximately 0.86, while that of a clear double glazing is approximately 0.76 and a clear triple glazing is 0.69.

The amount of solar energy that does wind up in the interior space can therefore be computed by multiplying the solar energy on the outer surface by the SHGC of the glazing, if this parameter is known. That energy is absorbed by the surfaces in the interior space, raising their temperature according to their surface properties and thermal mass. These surfaces will then re-radiate energy, at a longer wavelength and greatly reduced intensity. A portion of the re-radiated energy will reach the surface of the window; again, some of this energy will be reflected into the room, but some will be transmitted outward, and the radiant heat-transfer process proceeds as described above, except that it occurs in the longwave part of the spectrum and in the outward direction.

So that's how radiant heat transfer works in the context of window performance. How, then, do we use this knowledge to reduce heat transfer?

The answer is different depending on the type of building and the predominant loads expected in each case. Non-residential buildings tend to have large internal heat sources such as lights and office equipment, meaning they are dominated by cooling loads. It would be better in these cases to select glass that will reduce the amount of solar energy transmitted into the building, so one would select a glazing with a low SHGC. The best candidates would

FIGURE 6. Radiant transfer in outer glazing (#1).

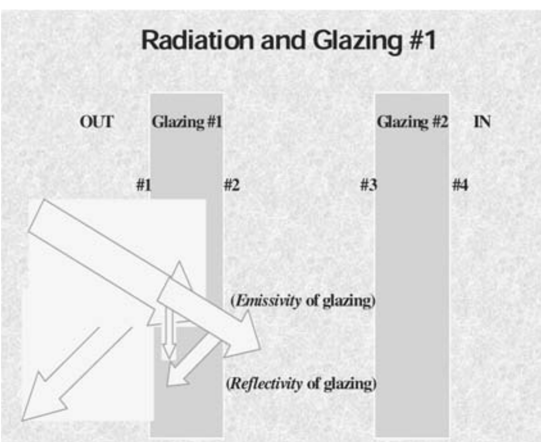

FIGURE 7. Radiant transfer in inner glazing (#2).

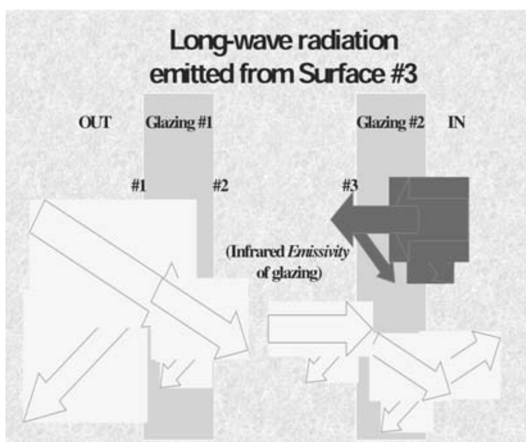

be reflective or tinted glass. The SHGC of reflective double glazing is typically 0.15–0.20, whereas double glazing with a tinted exterior lite would be more in the range of SHGC = 0.40–0.60, depending on the depth of the tint.

Reflective glass has a layer of metal or metal-oxide ions embedded in the outer layer and, as the name implies, reflects more of the incident energy back outside before it can contribute to the overheating of the building. Typical colors are silver, gold, copper, and bronze, as these reflect (pun intended) the metals used to develop the reflective property. Such properties can also be used to achieve a dramatic building façade. Figure 8 shows a bank tower that uses reflective gold glass to achieve a literal and metaphorical tower of gold, besides reducing the cooling load on the building.

Tinted glass can be achieved by mixing pigment into the glass, or by placing a reflective layer on surface #2. Both options will result in more energy being absorbed in the glass, which may lead to breakage if the glass is not tempered or heat-strengthened. As with any glass where breakage is a concern, more attention needs to be paid to the edge cuts on the glass. Nicks and notches at the edge of glass, formed when cutters are dull or used improperly, will produce stress concentrations where cracks will initiate and propagate.

Glass with a lower SHGC (typical values can range from 0.20 to 0.60) will prevent overheating in the occupant space, but will also tend to have greatly reduced visible transmittance. While the exterior view is not necessarily as important in a non-residential building, daylighting is an important source of light in new green designs. Therefore, some trade-off is required, as increasing the visible transmittance also increases SHGC. Daylighting can reduce internal heat sources by reducing the need for artificial lighting, depending on the availability of natural light. So part of the trade-off in non-residential window design involves determining whether the increase in visible transmittance (and consequent reduction in cooling load due to artificial lighting) can balance the increase in SHGC (which will therefore increase cooling load due to solar gains). Computer-based whole-building energy models are the best way to examine these trade-offs, because the answer depends on a variety of factors including: building configuration and orientation, the amount of thermal mass (both inside the

FIGURE 8. Example of reflective glass.

FIGURE 9. Radiant transfer in inner glazing (#2).

building and as part of the envelope), operating schedules, type of heating and cooling systems used, and building usage).

Residential buildings in southern climates (e.g., Miami or southern California) typically require cooling as well, but in this case the view is usually more important than in non-residential buildings. Therefore, the ideal glazing would have a low SHGC but a high VT, and good color rendition (which is typically a feature of higher VT). For this application, you need a *spectrally selective* glazing, which uses a combination of metallic oxide coatings to shape the transmissive properties of the glazing. A wide range of options are available, but generally the design intent is to meet the requirements of low SHGC and high VT.

Residential buildings in northern climates (such as Minneapolis or Toronto) are easier to deal with in design. Generally speaking, the heating load dominates the energy profile of the building for most of the year, even though the occupants might wish for more cooling for one or two weeks in August. The glazing option that best complements this load profile admits the most solar gain while minimizing heat loss. As we have seen by looking at the glazing physics, low-emissivity coatings provide the required performance.

LOW-EMISSIVITY COATING OPTIONS

In fact, there are two kinds of low-emissivity coatings. *Sputter coats*, or low-emissivity coatings developed by vapor deposition, are produced by the bombardment of metal or metal-oxide atoms on finished glass product (i.e., after the float-glass process). The specific metal used determines the solar-optical properties, but most sputter coats actually have several layers of oxide coatings. Sputter coats can be applied to any glass substrate, including tempered or laminated glass. In fact, a sputter-coated low-emissivity coating on tempered glass traditionally had to be tempered *before* the sputter-coat application. Newer sputter-coat low-emissivity coatings are being developed to be post-temperable.

When sputter-coat low-emissivity coatings are used in a sealed glazing unit, the coating must first be removed so that the spacer assembly can bond to the glazing (it would not provide a durable seal for the spacer assembly to adhere to the low-emissivity coating, as the coating is not that strongly bonded to the glass). This process of removing the low-emissivity coating around the perimeter of the glazing is called "edge deletion."

Sputter coats are also called "soft coats," and are not as durable as pyrolitic coating (see below). The first generation of soft coats featured emissivities in the range of 0.35–0.4, but second- and third-generation soft coats are much lower. Typical products include AFG Comfort Ti™ ($\varepsilon = 0.03–0.04$), PPG Solarban 60 ($\varepsilon = 0.035$) and PPG Sungate 100 ($\varepsilon = 0.09$), and Cardinal Lo-e2 ($\varepsilon = 0.04$) or Lo-e3 ($\varepsilon = 0.05$). With emissivities this low, the U-factors of soft-coat glazings are quite low, but the SHGC is also low. While these products are good choices, if the desire is exclusively to maximize thermal resistance (i.e., low U-factor), they do not provide the best choice for passive-solar design, because they tend to block more solar gain. The SHGC for a double glazing with soft-coat low-e on surface #2 is approximately 0.42, whereas that of a hard-coat on surface #2 is 0.65. Interestingly, a hard-coat low-e on surface #3 has an SHGC of 0.71 (higher than the SHGC for the same coating on surface #2, because the latter reflects more of the sunlight before it reaches the room).

Pyrolitic coatings are layers of metallic oxides applied to the molten glass during the float-glass production process. As this produces a durable coating that is actually part of the glass, pyrolitics are commonly called "hard coats." The first generation of hard coats produced emissivities in the range of 0.60, but second-generation pyrolitics have much lower emissivities, with consequent reduction in radiative heat loss. Typical products include AFG Comfort E2™ ($\varepsilon = 0.20$) and Stopsol™, Pilkington Energy Advantage ($\varepsilon = 0.16$), and PPG Sungate 500 ($\varepsilon = 0.21$).

Pyrolitic coatings are more durable during assembly and storage, and can be heat-treated and laminated after the glass is produced. Also, there is no need for edge deletion with hard coats, as the low-emissivity coating is actually part of the glass, so the bond between the spacer and the coating is quite strong.

SOURCES OF DESIGN VALUES

Understanding what affects window performance is vital to producing good designs, but you need data to evaluate the performance of the window in the context of the overall building envelope. Values for VT are required for input to daylighting designs, and SHGC and U-factor data help to define the heating and cooling loads for the mechanical systems. Unfortunately, there is no single “correct” answer, as we have already noted—each situation requires a different solution, and will result in a different selection of glass options that produce characteristic values of VT, SHGC, and U-factor. The question is: where can we get these values?

The ASHRAE Handbook of Fundamentals contains a chapter on fenestration, with values for SHGC, VT, and U-factor for a range of typical window framing systems and glazing options. These were developed to provide representative values for fenestration products at standard sizes, and are a good first approximation for rough design. More accurate values should be used as the design becomes more refined. Many individual window products have been rated in accordance with national standards (e.g., NFRC 100 for U-factor, NFRC 500 for SHGC), but these are often only for residential applications. Also, these products are rated at standard sizes, and SHGC, VT, and U are all affected by size. If the fenestration panels in your application are within 10% of the size of the standard products, the standardized performance parameters will be a reasonable approximation. Otherwise, you should determine a size-specific value for your case. The ASHRAE Handbook of Fundamentals or the NFRC standards provide guidance in determining total-product values at various sizes.

You can also use software designed for this purpose. The WINDOW and THERM programs, developed at Lawrence Berkeley National Laboratories, allow users to compile various layers of glass into window framing systems of any description. The software will produce total-product values for any size when used in accordance with NFRC standards and guidelines. If software is not desirable, there are NFRC-accredited laboratories that can perform the analysis. For a list of these experts, see www.nfrc.org.

FIELD VERIFICATION

So, having designed your green windows, how do you make sure you get what you wanted at the construction stage? The difficulty in determining the presence of a gas fill or low-emissivity coating is that these are generally invisible. There are some options for determining the presence of a low-emissivity coating in the field, however. Figure 10 shows an electronic device that can determine whether a low-emissivity coating is present on the glass (although it does not identify the nature of the coating or its emissivity).

Figure 11 shows a low-tech method of verifying the presence of a low-emissivity coating in the field. With a light source held at an angle to the glass, several reflected images of the light source can be observed (one for each glass surface in the assembly, so that Figure 11 illustrates the example of a double-glazed window with four glass surfaces). If one of the surfaces has a low-e coating, the image

FIGURE 10. Electronic detection of low-e coating

FIGURE 11. Using a light source to check for low-e

of the light source reflected from that surface will be redder than the other images. This is because the light reflected from that surface is shifted more toward the red end of the spectrum by the presence of a low-emissivity coating. The coating prevents more of the radiant energy from being emitted from the surface, and reflects it back toward the light source. The fact that the image appears redder means that more of what is being reflected is in the red end of the spectrum, because the low-emissivity coating operates more in the infra-red range (see Figure 2).

It is difficult to verify the presence of argon gas, as it is invisible, tasteless, and odorless. Some researchers are attempting to develop field-testing equipment that uses a high-voltage discharge to measure argon concentration, but it has proven to be difficult to make a portable and inexpensive tester. One possible way to check for the presence (or absence) of fill gas would be to use an infrared thermometer to measure the room-side surface temperature at the center of the glazing unit. The center-of-glass simulation software (i.e., WINDOW or VISION) can be used to simulate the window assembly using boundary conditions that match the field conditions (warm- and cold-side temperatures and an approximation of surface film coefficients). The fill-gas concentration can be varied in the software until the predicted surface temperature matches the measured value, which should provide a reasonable estimate of the presence of fill gas. Of course, if you don't specify fill gas, you don't have to check for it.

Poor installation techniques can compromise the energy savings of even the highest-performance windows. Proper installation can be verified *in situ* by using a field test for air leakage around the perimeter of the window (following ASTM E783), or for water leakage in the presence of a simulated wind-driven rain (following the ASTM E1105 protocols). These tests should be performed by qualified testing agencies. A rigorous maintenance program is also important in improving the long-term durability of seals and weatherproofing components.

CONCLUSION

We have seen how green window design starts with an understanding of the basic physical principles behind window performance. Proper field review is also necessary to ensure that the intended design actually gets installed. Taken together, these measures will result in improved window performance, which ultimately improves the thermal comfort of the occupants and contributes to greener buildings.