Editor: Miguel Encarnação

The Digital Chameleon Principle: Computing Invisibility by Rendering Transparency

Frank Nielsen
Sony Computer Science Laboratories

How many of us have had a dream of being invisible? But is that really that far-fetched, impossible, or improbable? As a child I used to daydream of magical powers that let me pursue exciting adventures in magical wonderlands. These great magical powers exhibited technology far beyond the current state of the art. I still clearly remember today those fantastic dreams. Indeed, dreaming that kind of fantasy let me feel good and relaxed, and they certainly colored my childhood. I personally think that daydreaming is also a kind of healing therapy of our day-to-day busy life to satisfy our curiosity. However, nowadays I also feel somewhat more conservative (and do not dream unfortunately anymore of Harry Potter worlds), partly because of my science background gained over the last 20 years or so of school education. School taught me the way to reasonably and rationally think without leaving much space for extravagance. That phenomenon is what Sir Arthur C. Clarke called the failure of the nerves in his book *Profiles of the Future: An Inquiry into the Limits of the Possible*.[1] A good example cited by Clarke are excerpts of scholarly articles written by experts claiming the then impossibility of heavier-than-air flying machines. Nowadays, we all take airplanes, so why do children still dream of magical worlds? Maybe, a simple answer is that children use and look at technology without fully understanding their underlying mechanisms making it possible. It's easier for them to naturally extend the scope of those technologies according to their will. In this article, I will examine the concept of invisibility in detail with a good flavor of computer graphics. But first, let me quickly review its historical background.

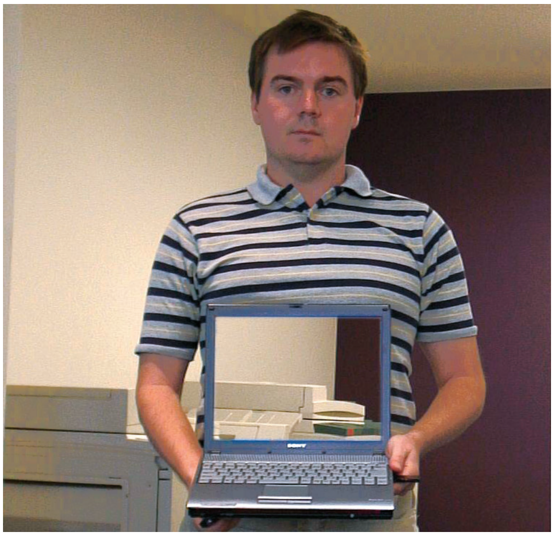

1 A glimpse at tomorrow's invisibility using today's digital forgery tools.

The Invisible Man and the quest for cloaking devices

The quest for invisibility was certainly spurred by the science fiction novel by H.G. Wells, *The Invisible Man* (1897), that was later turned into a movie directed by James Whale (1933). The myth and quest of the invisible man had then entered households. Eventually, the seminal movie spawned a number of sequels such as the recent *Hollow Man* (2000). Although the original 1933 movie was poor in visual effects (remember the "invisible" Griffin shown with his face completely hidden by bandages, wearing goggles and a hat?), the *Hollow Man* movie was, in contrast, so rich in computer graphics effects, that it was nominated for an Academy Award—not for the best actor performance but rather for its visual effects. In *Hollow Man*, the invisibility power is reached by a team of scientists injecting some substance into Kevin Bacon's body. Flesh would become quickly but progressively transparent.

As H.G. Wells asked, what is the purpose of being invisible? After all, what motivates us that much? Beyond the obvious goals of peeping, robbing, or fighting, it's not so clear. Ideally we would like to have the invisibility power reversible. But this turns out to be difficult, at least in the movies, where film directors following H.G. Wells emphasize the dangers of invisibility rather than its merits. Becoming irreversibly invisible turns out to yield an unstable mental state that reaches various paroxysms beyond the point of no return. This explains why most of the science fiction movies such as the *Star Trek*, *Predator*, or *Harry Potter* ones consider reversible invisibility as a magical gift or better as a technological tool. In the *Star Trek* television series' "Balance of Terror" (1966) episode, Romulans hide their spaceships simply by pushing a button that activates the invisible shield. That is, invisibility is obtained after activating an advanced stealth gadget-cloaking device, which causes someone to be extremely difficult to detect with conventional sensors. Can we legitimately ask

Published by the IEEE Computer Society 0272-1716/07/$20.00 © 2007 IEEE

about the physical rationale of invisibility? A trivial but naive way would be to consider changing the body refractive index to that of the surrounding ambient air, ensuring that our body absorbs no light. We would then become almost as air is—that is, invisible or scientifically speaking, transparent (see Figure 1 for a fake preview created by image compositing). But is this that simple?

Invisibility: science facts or urban legends?

Several times in the 20th century, different sources reported that invisibility was allegedly successfully obtained. Perhaps the most famous but arguable story is the Philadelphia experiment. The Philadelphia experiment myth appeared in 1955 with the publication of *The Case for UFO's* (Bantam Books) by Morris K. Jessup. After the publication of his book, Jessup received a letter from C.M. Allende that detailed a claimed secret experiment carried out in 1943 by the US Navy in Philadelphia. Allende said that the USS *Eldridge* ship was rendered invisible and transported from Philadelphia to Norfolk, Virginia, and back in a mere few minutes. This was achieved by applying the principles of Einstein's unified field theory. This is all the more puzzling as Einstein indeed worked for the navy during World War II. The experiment supposedly had some terrible aftereffects for all the ship crew members who became badly ill and developed schizophrenia, among other ailments. Although Allende claimed in his correspondence to have witnessed this experiment from another ship, he could not be later identified, nor could authorities trace alleged reports of this experiment in Philadelphia newspapers. The Philadelphia experiment is nowadays considered a hoax. But we cannot rule out for sure the technological possibilities of Einstein's space–time theory.

Invisibility and camouflage in nature

Coming back to the scientific rationale of invisibility, we can start by having a look at Mother Nature. We find in nature a beautiful source of inspiration for reaching invisibility. The first example that comes to mind is jellyfish. Jellyfish are well-known sea creatures that are translucent and look transparent so that they are almost invisible in seawater. But their functionalities are unfortunately limited. Let us look at other species. Consider flat fishes such as winter flounders that hide themselves from predators by adapting their skin pigments according to the background environment. These fish can mimic the color of their background environment by controlling pigment cells, or chromatophores, located in the dermis and epidermis. Chromatophores are irregularly shaped cells containing chemical pigments that are also sensitive to temperature changes and stress condition. Chromatophores are also found in crustaceans, lizards, and amphibians. Color changes take place in chromatophores by hormonal control (a slow process depending on the body's emotional state) that yields either dispersion or gathering of granules within the cells. This is precisely that density change that lets these cells control either the darkening or lightening effect of their appearance. Camouflage is a must value to survive in a world populated by predators. Unfortunately, predators also use camouflage tricks themselves to lure prey for lazy but efficient foraging—think of a chameleon or a stonefish, for example. In the long run, chemical or molecular switching might offer the best method for a photochemical cloaking device. For now, they exhibit far too slow a latency. Let us look now at the feasibility of this transparency principle by rendering background imageries, using off-the-shelf, low-latency cameras and projectors commonly found in computer graphics and vision laboratories.


© Paul Debevec, Univ. of Southern California

2 **A photorealistic computer graphics rendering obtained by lighting the 3D scene geometry with a light probe (shown in the inset).**

Plenoptic functions and light fields

Let me start the principle of rendering transparency by introducing the notion of a plenoptic function for modeling a scene's visual appearance. The word *plenoptic* stems from Latin roots *plenus* (meaning full or complete) and *optics* (meaning see). The plenoptic function is a 7D function that records for every spatial position (X, Y, Z) the intensity of the light of the ray anchored at (X, Y, Z) for every orientation (θ, ϕ) at every time t (six parameters). Because light such as visible light is in general nonspectral, it's best described by its spectrum decomposition into pure (ideal) monochromatic lights, thus adding the extra wavelength dimension λ to the plenoptic function. Such a decomposition and analysis of light is done using a spectrogram. The visible spectrum is a subband of the electromagnetic field with wavelengths ranging in the 400 to 700 nanometers. In summary, the 7D plenoptic function $P(X, Y, Z, \theta, \phi, \lambda, t)$ describes the full environment light intensity contribution of the monochromatic lights of wavelength λ at any location in space at any time. For a given position at a given time, the plenoptic function is simply called a (panoramic) light probe. Light probes are nowadays often captured for photorealistic image-based lighting of 3D computer graphics. Figure 2 shows a synthetic rendering of a 3D scene using a high dynamic range light probe.

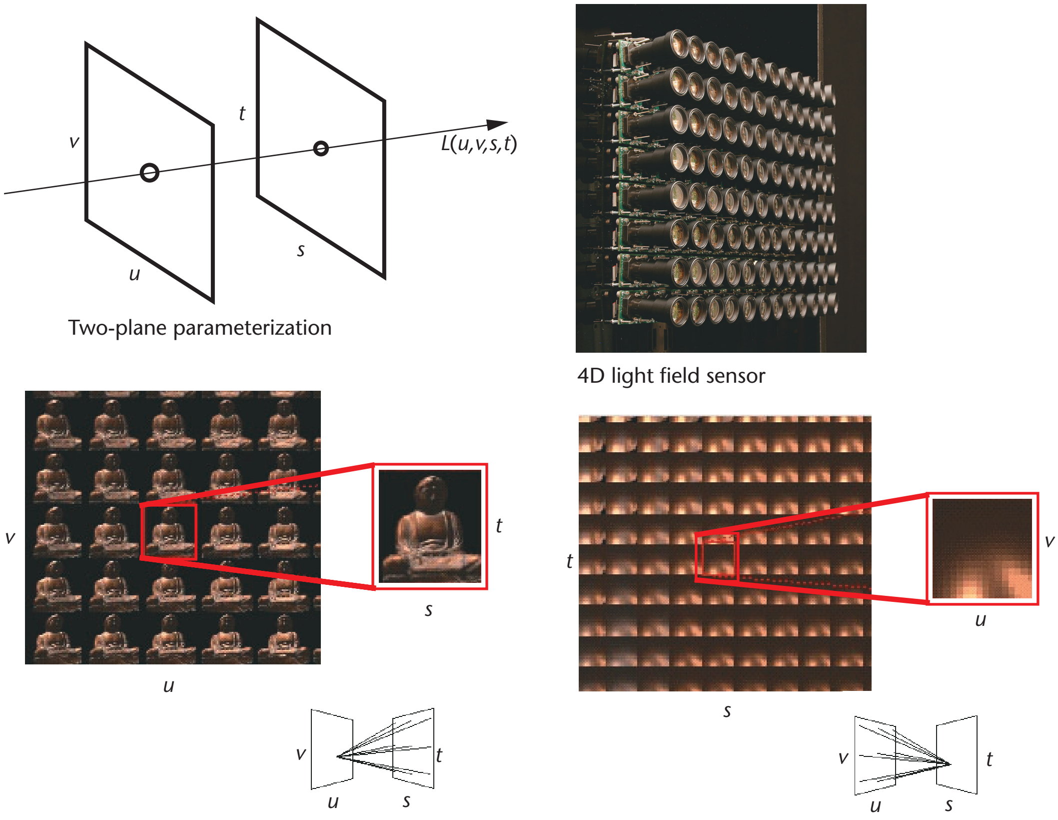

3 A 4D light field reduces the 7D plenoptic function by considering a static scene from its exterior for a given color channel.

Courtesy Stanford Computer Graphics Laboratory

Researchers have extensively studied plenoptic modeling, which has evolved into a subfield of computer graphics called image-based rendering. One such key pioneer image-based rendering work is the light field project carried out in 1996 by Levoy and Hanrahan (see http://graphics.stanford.edu/projects/lightfield/ for more information).[2] Levoy and Hanrahan captured and rendered particular cases of plenoptic functions: 4D light fields. A 4D light field simplifies the plenoptic function to practical cases. A light field is only described by four parameters (s, t, u, v) that parameterizes 3D lines using two pairs of 2D coordinates (s, t) (the focal plane) and (u, v) (the nodal plane) of two disjoint parallel planes (see Figure 3). A light field $L(s, t, u, v)$ considers static scenes and a given color channel (red, green, or blue). Light fields are also called radiance fields because they are related to the space distribution of light energy. To capture a static light field using a pinhole camera, we can select a few (u, v) positions for the camera's nodal point and acquire the corresponding bundle of (s, t) rays (see Figure 3). To capture time-varying dynamic light fields, we need a 2D camera array that acquires synchronously these bundles of rays, temporal frame by temporal frame. Once we acquire a light field, we can generate new synthetic pinhole camera views on the fly using pixel rebinning of source images.

Light field rendering is a brute-force approach in computer graphics; however, it yields spectacular photorealistic results where traditional scene modeling has failed. Now that I have finished describing the main technical concepts, it's time to see how to implement the low-latency optical active camouflage following the chameleon principle.

The digital chameleon principle

The main difference with the approach I take and, for example, a simple snapshot of a translucent jellyfish is that I am going to shade real objects (that is, physical objects) transparent. These dynamically rendered objects will appear almost invisible. For simplicity, consider a 3D box as depicted in Figure 4 with an observer standing in front of the box. My goal is to deliver in the viewer's exploring area a plenoptic field that matches as closely as possible the plenoptic field that would exist if the box were not there. The result is an exact replica of the image appearing behind the box from any viewpoint in front of the box. In doing so, I can conceal inside the chameleon box arbitrary objects (a small world) to hide from the outer world. Let us proceed step by step by considering several cases.

First, the simplest case—a fixed viewpoint (see Figure 4a). Assume that a pinhole camera correctly approximates the geometric vision of our observer. In that case, the set of rays emanating from the observer's pinhole eye has a fixed geometry and can be acquired by light sensors located on the other side of the chameleon box. If we further assume that the scene is located far away with respect to the viewer's distance to the box, then we can acquire with a good approximation this set of rays using another similar pinhole camera located on the chameleon box's back side. Evaluating the transparency resolution obtained is a delicate matter as many factors intervene in the final transparency rendering. However, we might easily guess (by extrapolation) the result of the digital chameleon box experiment by referring to a similar project: a team of researchers led by Tachi and Inami of the University of Tokyo uses a retroreflector material called X'tal vision to project background images onto the masked object that appears transparent (see Figure 5).[3] The method is similar in spirit to that of Figure 4a. The retroreflector cloak consists of a coating of small beads that reflects the light exactly in the opposite incoming direction without being too scattered or absorbed. Retroreflector materials are already massively used for traffic road signs, allowing drivers to read faraway road signs with their vehicle lights. Such retroreflector cloak systems, however, provide only a good illusion of invisibility for a viewer's gaze direction, matching the surface normal of the object needing to be concealed. To acquire the optical camouflage background, the University of Tokyo researchers used a video camera located on the back of the standing human or object. This limited the freedom of motion of the quasi-invisible actor. The chameleon box experiment, however, gives a reasonable glimpse of the box's quality, and incites us to expand venues for an active optical camouflage technique.

Next, consider a fixed dot on the front side of the chameleon box (see Figure 4b). The geometric position of the dot and the observer's viewpoint defines a geometric ray that a projector can render as the same color of the ray passing entirely through the box. As the viewer moves freely around, the ray's geometry changes, but another closely matching projector ray can still render it with the correct color.

Finally, let us consider a general case (see Figure 4c). Namely, viewing the front side of the chameleon box from an arbitrary viewpoint. It's easy to extend the two previous cases to this general case. By using a camera array on the backside of the chameleon box, you can acquire the 4D light field incident onto the wall and reproduce it on the box's front side using a projector array. This process guarantees that the viewer will receive the same set of rays as if the box were not there. That is, the box is rendered transparent (shaded transparent), and I can compute the invisibility illusion by digitally remapping 3D light rays acquired on the backside camera panel to rays projected on the front side projector panel. The chameleon box in its simplest form is a naive, brute-force, digital-ray-remapping engine. Cameras play the role of light detectors while projectors play the role of light emitters. In other fancy words, the two opposite box sides are light portals of the plenoptic function. Of course, in practice, we would rather like a free-form cloak (the invisible suit) that perfectly fits the body or object to conceal: a cloaking device as shown in *Predator* or *Harry Potter* movies.

How far are we from building a real prototype of such a digital chameleon box? Actually, not as far as you might think. Indeed, by noticing that the light field captured by the 2D camera array on one side of the box can also be entirely computed by simulation using computer graphics, you can remove at first the camera array (and reduce messy synchronization engineering details). In that case, observers see a 3D virtual world projected in front of the box. This is the basic principle of some families of 3D displays that work by reproducing virtual synthetic scene rays. Several industrial systems are currently built and are shown all over the

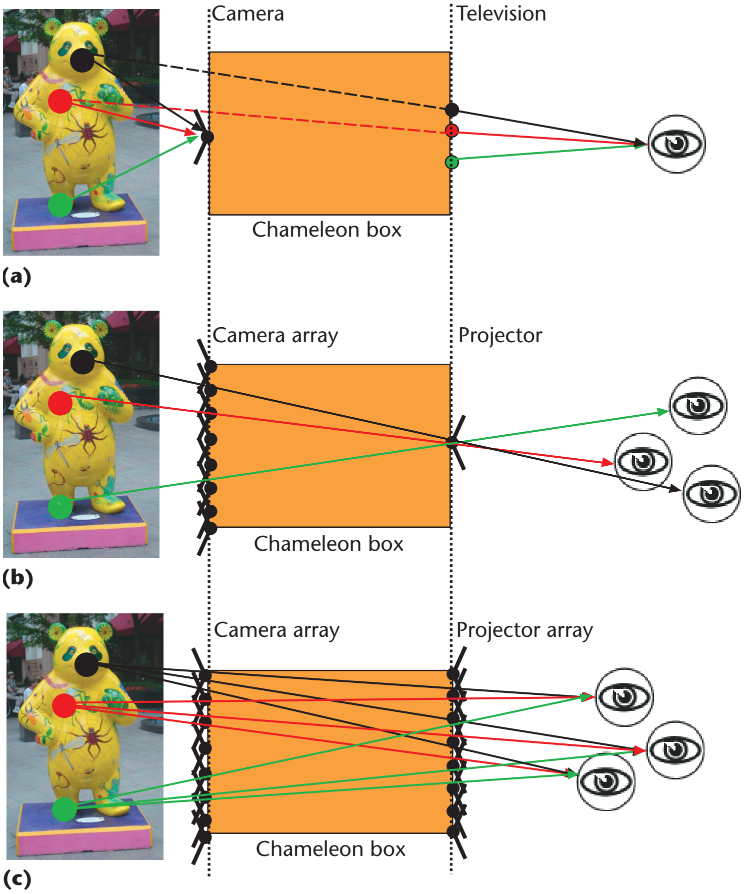

4 The digital chameleon principle relies on a camera array to acquire a light field on one side of the box and uses a projector array to reproduce that light field on the other side, thus visually concealing the box to an observer located in front of the ray display. (a) One camera and one screen give a correct image from a fixed viewpoint. (b) Multiple cameras and one projector give a single dot view, correct from any direction. (c) Multiple cameras and multiple projectors give a correct image from anywhere.

Courtesy Tachi Laboratory, Univ. of Tokyo

5 The invisible cloak experiment carried out by researchers at the Tachi laboratory.

world. See, for example, the Holografika display shown at the 2006 Siggraph Emerging Technologies exhibit (see http://www.siggraph.org/s2006/main.php?fconference&p=etech&s=holographic and http://www.holografika.com/).

Thus to build a digital chameleon box, you would need to combine such a 3D ray display with a 2D array of cameras. This is, technologically speaking, feasible since nowadays camera arrays are mainstream in computer graphics and vision laboratories. This tight coupling of camera arrays and projector arrays will likely soon strengthen as 3D cinematography and next-generation videophones challenge these technologies.

Further, to reduce the complexity of using 2D camera arrays and/or 2D projector arrays, we can consider using monolithic, high-resolution acquisition and rendering devices—that is, a single, high-resolution camera/projector. For a century now, the integral photography field has investigated this architecture. Instead of using camera/projector block arrays, integral photography uses a sheet of tiny hemispherical lenslets, usually arranged in a honeycomb pattern, manufactured using well-mastered photolithographic processes. Each tiny lens captures the light rays coming into the center of the lens with a given incidence angle. A high-resolution area sensor captures the sheet, yielding a sheer amount of ray information. Ideally, the smaller the lenslet the better the resolution. There are, however, some physical constraints, and you will typically face the geometric optics diffraction problems encountered at small scales.

Let us envision a bit more of the technology of the optically invisible man, beyond the digital chameleon box. In practice, we would prefer wearing a suit equipped with interspaced, dense, tiny button-size camera/projectors that dynamically autocalibrate themselves on the fly rather than using the massive equipment mentioned previously. This is just a matter of time as miniaturization of these key camera/projector light-sensing and emitting devices progress drastically, and novel manufacturing techniques—such as silicon on glass, which hides the large-scale-integration circuitry in translucent material—become possible.

To wrap-up, perfect invisibility seems quite far away from today's practicalities. Yet, computational invisibility, an emerging topic of computational photography, has numerous applications before reaching the consumer level. Indeed, how useful would it be to have the floor of an airplane cockpit rendered transparent! No more visual dead zones. Also think of how easy it would be to park buses or simply your car with the back rendered transparent, replacing our current loud warning buzzers used for safety.

Invisibility and shading

Invisibility is the task of masking physical objects so that they are no longer visible. In contrast, shading is the task of making up visual appearances of objects. Visual properties of object appearances include reflectance characteristics and numerous phenomena such as subscattering or diffusion effects. The appearance of an object is simply expressed as a function of the illumination (given by the surrounding environment), surface reflectance properties (depending on materials), and viewer location. Looking back at this column's title, I chose to write "The Digital Chameleon Principle: Computing Invisibility by Rendering Transparency," but I could also have replaced the subtitle using more technical jargon like "Computing Invisibility by Shading Transparency," since invisibility can also be interpreted as the process of creating a transparent shader. Shading is nowadays an essential technology for achieving photorealism in computer graphics. Let us address a computer graphics–related project called the shader lamps that also deals with physical objects.[4]

The main idea of a shader lamp is to replace our usual light bulbs by tiny computer-controlled calibrated projectors that emit per-pixel controlled light color on objects to produce realistic object appearances, as depicted in Figure 6. This method first molds physical objects from white clay or covers them with a material sheet that has similar diffuse matte finish characteristics. It then interactively shades the objects by rendering computer graphics from projector viewpoints and projects these synthetic images onto the diffuse matte objects that play the role of geometric proxies. Once again, let me emphasize that these special visual effects take place in the real world and not on the computer screen as usual, thus making it easy to design tangible user interfaces for interactive shading. The number of shader lamps (projectors) depends on the intrinsic complexity of the geometric object since the system needs to account for potential self-shadowing of objects and must fully cover the scene surfaces. Shader lamps not only visually rewrite the object's texture (the diffuse component) but also reproduces any bidirectional reflectance appearance (whether view-dependent, using a bidirectional reflectance distribution function, BRDF, or not). Moreover, not only do shader lamps enable us to project virtual reflectance on registered real-world geometry, they also animate physical static objects as well by projecting dynamic image-based illumination movies. The visual appearances of projector-shaded objects are thus lifted from the real to the mixed-reality world of computer-controlled projectors. The shader lamp approach is currently limited to physical, diffuse material objects, and requires tracking the single viewer's location if view-dependent effects such as specular highlights are required.

Modalities of invisibility

Even if the active optical camouflage looks good, meaning human eyes can't detect it, an infrared or thermal camera might easily show someone hiding inside the chameleon box. That is, we have to take into account the modality of invisibility. Indeed, to give yet another example, consider a silent, black-painted airplane traveling high enough at night in a cloudless sky: no one on the ground would be able to notice it. Yet, this would be trivial for radars to detect that airplane using the reflections of emitted radar signals. Aircrafts thus become vulnerable. To overcome that issue and hence bypass radar technologies, the American SR-71 Blackbird airplane is using a special radar-absorbing paint to trap those signals and limit as much as possible their reflections, making them

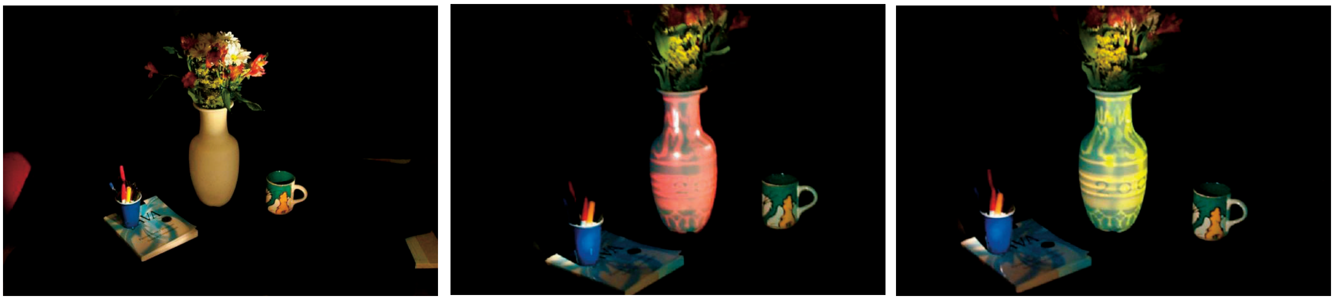

Courtesy Ramesh Raskar, Mitsubishi Electric Research Laboratory

6 **Virtual reflectance of a white diffuse physical vase. Shader lamps provide a virtual, bidirectional reflectance distribution function appearance to diffuse material, physical objects.**

radar invisible. However, we could possibly come up with a new sensitive detection scheme using another modality. Modality emphasizes that invisibility is not ultimate but rather limited to various sensing methods.

Stealth and antistealth technologies

For an observer, we might reverse the challenge of invisibility: can we see invisible things? The question might sound absurd at first since we have concentrated so far on making objects optically invisible. But remembering that invisibility and visibility is simply a matter of imagery techniques, we might create novel modalities for seeing so-far invisible objects. Let us take a typical example encountered in daily life at the airport security check screening. Guards try to detect efficiently invisible objects hidden in clothes and bags. The usual bag screening procedure is the x-ray conveyor belt. Yet, nowadays with the high level of terrorism alerts, more and more substances and objects need to be detected for preventing their access on board. For that crucial purpose, there have been novel imaging modalities that have been quickly commercialized to automatically check whether someone is wearing undesirable substances, such as suspicious liquids or traces of explosives. The T-ray camera is one such amazing technique. It's sensitive to terahertz frequencies and literally lets us see through clothes or other packaging (see Figure 7). T-rays also allow cameras located in space to see through clouds. T-ray cameras are certainly one of the first great technological breakthroughs of the 21st century.

My approach to invisibility based on the plenoptic-modeling paradigm is scalable and should yield a certain amount of invisibility. It's clearly not a pure approach such as physics-based ones, but rather an engineering illusion. One drawback is that the digital chameleon principle does not prevent the box (or rendered objects) to have reflected or scattered light due to its physical properties. Thus, ultimately we need a direct, physics-based, self-contained structure method that does not violate the basic laws of physics.

Microscopic-level invisibility

Recently, researchers in physics pointed out a new family of materials, metamaterials (also called superlens), exhibiting the property of negative refractive indices. By using metamaterial, light can theoretically bend around objects. This property yields a potentially pure physics-based invisibility technique. Actually, historically, warping light was the first proposed approach to invisibility, reportedly studied in the 1930s using an electromagnetic field. How does this metamaterial work? Alu and Engheta of the University of Pennsylvania use plasmon waves to reduce the light scattering effect of incoming light, thus making the object appear less visible.[5] Indeed, we see objects because light bounces off them. The plasmonic cover prevents light scattering by resonating at the same frequency of the incoming light. That is, plasmons are electrons of a metallic surface moving in rhythm defining a shell. That shell reduces light scattering providing that the light frequency is close to the resonant frequency of the plasmons. Gold and silver are two suitable materials for building plasmonic visible-light shields. Although the idea of a plasmonic shield is theoretically sound, its current practicalities are limited to microscopic size objects under controlled lighting. To be practical, the method needs to take into account real sunlight (rich blend of wavelengths) and scale objects to a reasonable human size. That is, this does not yet offer a self-contained magic cloak. But with more

Courtesy QinetiQ

7 **Example of a T-ray image.**

research, plasmonic screening could theoretically hide large objects from telescopes whose imageries are based on long wavelengths instead of short visible waves.

Conclusion

Today, invisibility does not seem out of reach to the realm of science. At least, it's more probable than it was a century ago. There are yet many technology gaps to bridge to reach true invisibility. I hope that I've convinced the reader that transparency could be in principle computed and implemented by a special shader using both camera and projector arrays, such as with the digital chameleon box. This basic idea has already yielded various patent applications, and my guess is that this active optical camouflage technique will soon be successfully demonstrated.

I should emphasize that we are seeing with our brain. Our eyes are only the raw photo sensors that deliver basic electrochemical signals to our brain, which then processes these low-level cues into higher cognition notions. Thus, it might be possible to think of invisibility at the human brain level. What does being invisible mean then? In cognitive science, this invisibility phenomenon is called cognitive blindness. A typical example is the case of somebody leaving a meeting room without being noticed by an audience that is deeply engrossed in a conversation. Such a cognitive invisibility could be individually selective compared with real-world, physics-based absolute invisibility. Yet another potential investigation topic linking physics-based and human-eye invisibility is vibration. Indeed, as stated previously, we see using the persistence of vision property. That is, light is first accumulated in retinal photosensors (cones and rods) before propagating the impulses into electrochemical reactions. Thus, we average light, and this causes various scene aliasing effects such as the car wheels that look like they are spinning in reverse. Looking at an air fan, for example, we literally see through it by virtue of the persistence of vision. Thus vibration and light averaging might also be a future direction for finding other invisibility tricks.

The approaches to invisibility presented here are within the probable reaches of today's science. Yet the door is widely open to our imagination, and we can find other ingenious schemes. As the pace of technology keeps ineluctably increasing, recent achievements and progress suggest that the active camouflage paradigm of the digital chameleon box is not just a utopian dream, but is rather becoming closer to reality. For now, let me conclude by citing Alfred A. Montapert: "Nature's laws are the invisible government of the Earth," and these are the hard ones to spot! ■

Acknowledgments

I thank Paul Debevec (University of Southern California), Marc Levoy (Stanford University), and Ramesh Raskar (Mitsubishi Electric Research Laboratory) for their kind image contributions. I am indebted to Shigeru Owada (Sony Computer Science Laboratories) for his valuable advice in drawing Figure 4. This article was reformatted, edited, and inspired by the original article appearing in the December 2006 edition of *Janus* magazine (see http://www.janusonline.net/, "The Digital Chameleon Principle: Computing Invisibility by Rendering Transparency") that runs a full-length issue on the many fascinating facets of invisibility.

References

1. A.C. Clarke, *Profiles of the Future: An Inquiry into the Limits of the Possible*, Indigo, 2000.
2. M. Levoy and P. Hanrahan, "Light Field Rendering," *Proc. Siggraph*, ACM Press, 1996, pp. 31-42.
3. S. Tachi, "Telexistence and Retro-Reflective Projection Technology (RPT)", *Proc. 5th Virtual Reality Int'l Conf.* (VRIC), Istia Innovation, 2003, pp. 69/1-69/9; http://projects.star.t.u-tokyo.ac.jp/projects/MEDIA/xv/VRIC2003.pdf.
4. R. Raskar et al., "Shader Lamps: Animating Real Objects with Image-Based Illumination," *Proc. 12th Eurographics Workshop Rendering Techniques*, Springer, 2001, pp. 89-102; http://www.cs.unc.edu/~raskar/Shaderlamps/.
5. A. Alu and N. Engheta, "Achieving Transparency with Plasmonic Coatings," *Physical Rev. E*, vol. 72, no. 1, 2005.

Contact author Frank Nielsen at Frank.Nielsen@acm.org.

Contact department editor Miguel Encarnação at me@imedia-labs.com.