Outwardly Pointing Cameras

Vic Nalwa, PhD (EE, Stanford), ***Fellow of the IEEE***
President, FullView, Inc.

Abstract

Panoramic and spherical-view cameras today are by and large one of three types:

1. *A single camera with wide-angle optics, as provided by a fisheye lens or curved mirror;*
2. *Multiple cameras looking out directly in different directions; or*
3. *Multiple cameras looking out off flat mirrors, which is FullView's patented approach.*

Whereas multiple cameras offer much higher and more uniform resolution than single cameras, outwardly pointing multiple cameras are in general incapable of producing composite images that are seamless, artifact-free and blur-free no matter what, due to parallax. FullView evades parallax through its patented approach in which multiple cameras look out off flat mirrors such that each camera is effectively looking out in a different direction, but from the same single viewpoint. As a result, FullView's composite images — whether video or still, and irrespective of their resolution — not only are always seamless, artifact-free and blur-free, but also always provide much higher and more uniform resolution than outwardly pointing cameras.

1. Parallax: What ails outwardly pointing cameras

Parallax is the disparity between views of an object due to a difference in viewpoints.

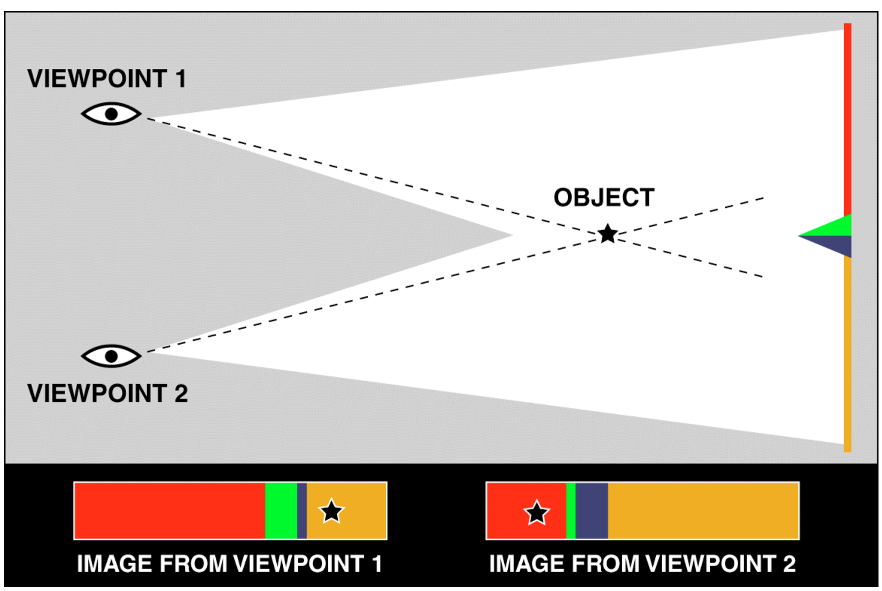

It is parallax — manifested in the different relative positions of the same object point in images from different viewpoints — that is the hurdle to combining images from outwardly pointing cameras into seamless, artifact-free, blur-free composite images. Parallax causes perspectives of the same object from different viewpoints to be different, as of the peak in the background above, and it also causes objects to occlude others differently in images from different viewpoints, as does the star in the foreground above. **Importantly,**

1.1 **The parallax of an object decreases as its distance from the viewpoints increases.** But the farther away an object, the higher the image resolution we need to see it clearly, and the higher this resolution, the more significant is previously insignificant parallax!

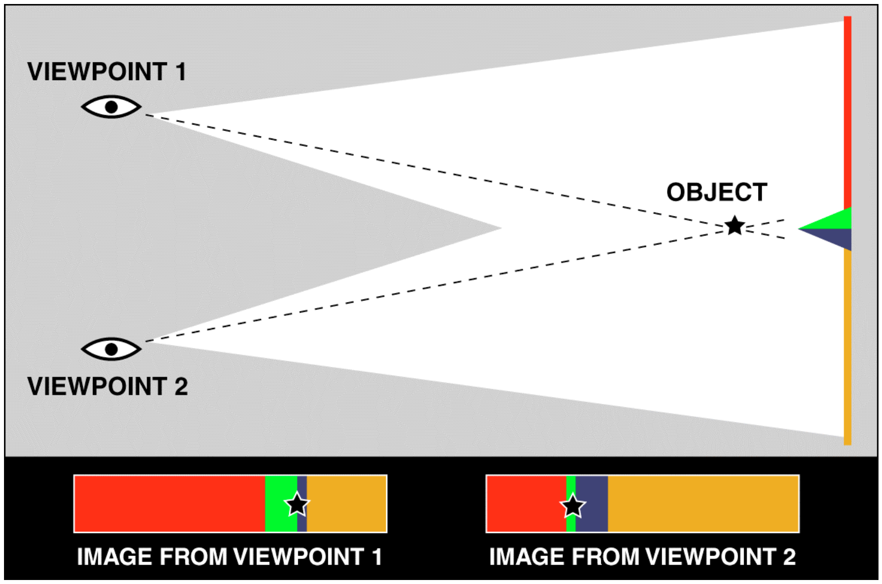

1.2 **The parallax of every object decreases as the distance between the viewpoints decreases.** Every object's parallax approaches zero as this distance approaches zero.

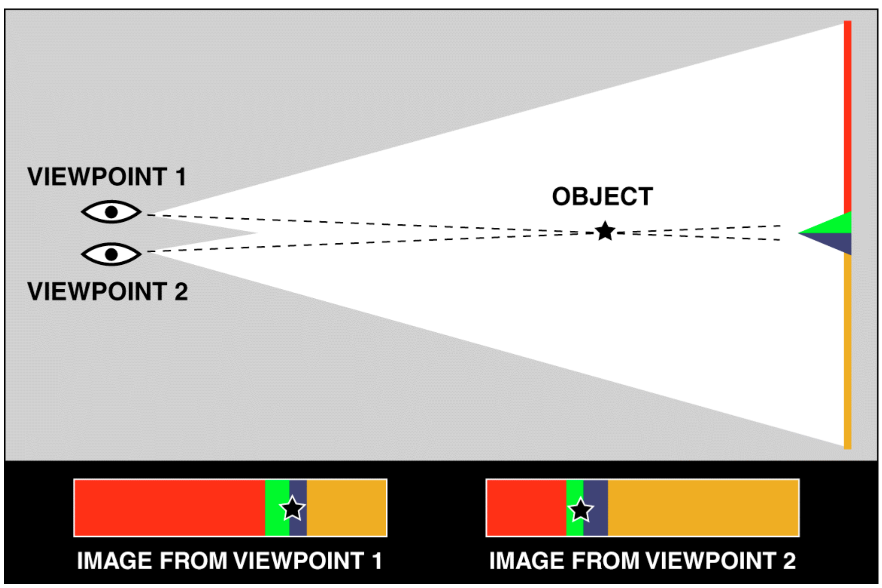

Parallax hinders combining images from different viewpoints into seamless, artifact-free, blur-free composite images for precisely the same reason that it facilitates stereoscopic vision: Because the parallax of an object point encodes its three-dimensional (3D) position. As clear from the figures above, if the 3D position of any object point is known relative to the viewpoints, we can deduce its differing image positions, and hence its parallax, and vice versa. Then, reconciling parallax between images from different viewpoints to create a composite image requires that we implicitly or explicitly decode the 3D shapes and positions of all those portions of objects that are visible in more than one image. There is no way around this. But this is the stereo problem of computer vision, a problem well known to be without any robust and general computational solution to date despite decades of research.

2. Image Matching: How parallax is determined, often erroneously

Parallax between images is always determined — whether directly or indirectly — by image matching: By finding the best match to an image region encompassing an object point in one image, in the other image. Image matching assumes that every image region sought to be matched will be replicated — relatively uniquely — in the other image. But this, as clear from Section 1, in general requires that all objects in every matched image region have the same parallax, which in turn in general requires that all objects be extremely distant.

As image matching, in general, can reliably determine the parallax of only those objects at best that are extremely distant from the viewpoints relative to the distance between the viewpoints, composite images of objects that are not so distant typically exhibit artifacts, irrespective of whether any image matching is performed. Such artifacts are red boxed in the 3-image panorama below from Point Grey's LadyBug2 spherical camera's website in 2008.

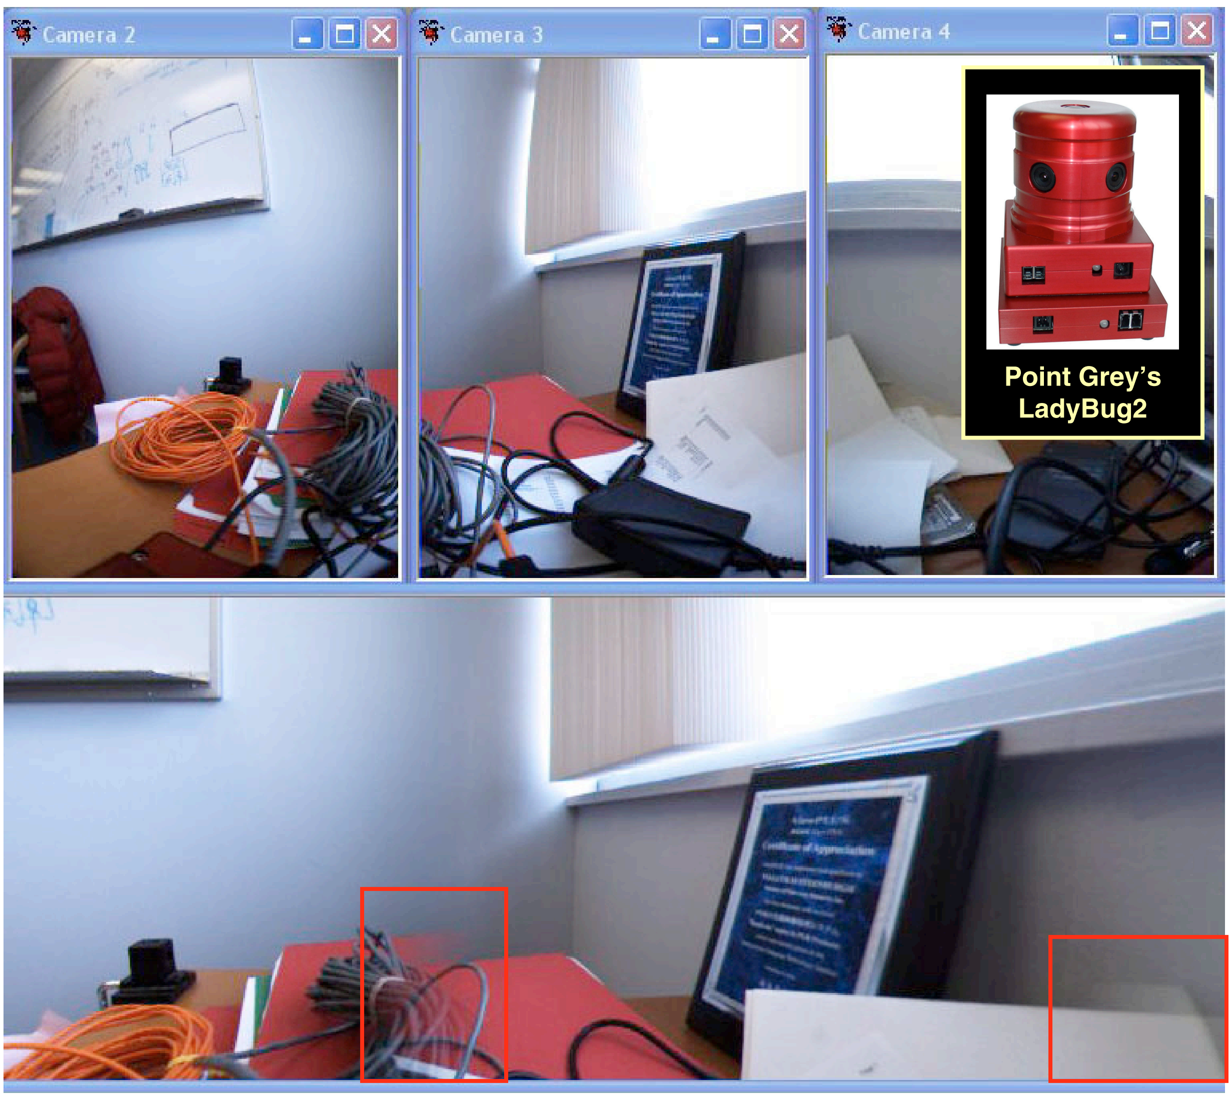

Even when objects are all extremely distant from their viewpoints, composite images from outwardly pointing cameras often exhibit glaring artifacts. For instance,

2.1 When images have periodic patterns, several equally good matches might exist between regions in one image and regions in another, and then which match is

chosen becomes a matter of chance. This could be the case in the following screen grab from [Google's Street View] in 2008, in which not only is the periodic railing split, but apparently as a result, so are the road and the rampart in the distant background.

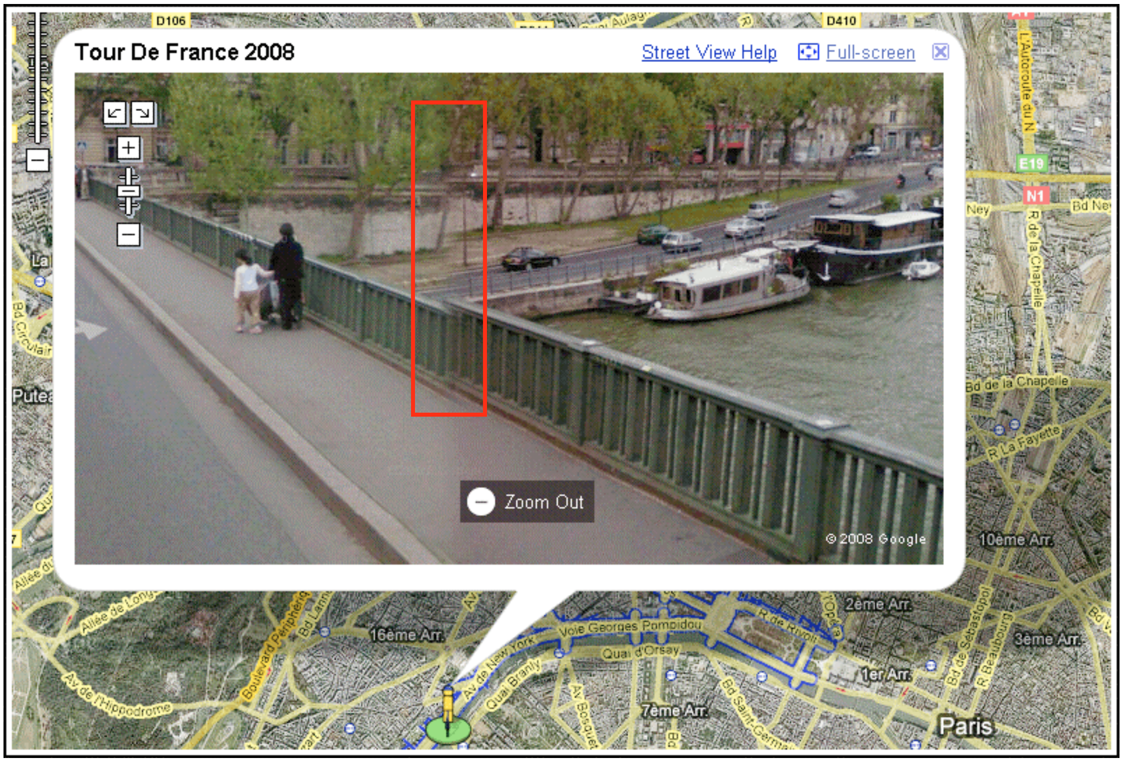

2.2 When there is a dearth of image texture, there is little to match, and then the result of image matching is unpredictable. This could be the case in the following screen grab from [Google's Street View] in 2007, in which the textureless station wagon is split in two, apparently causing the railing behind it and the tire marks in front of it to also split.

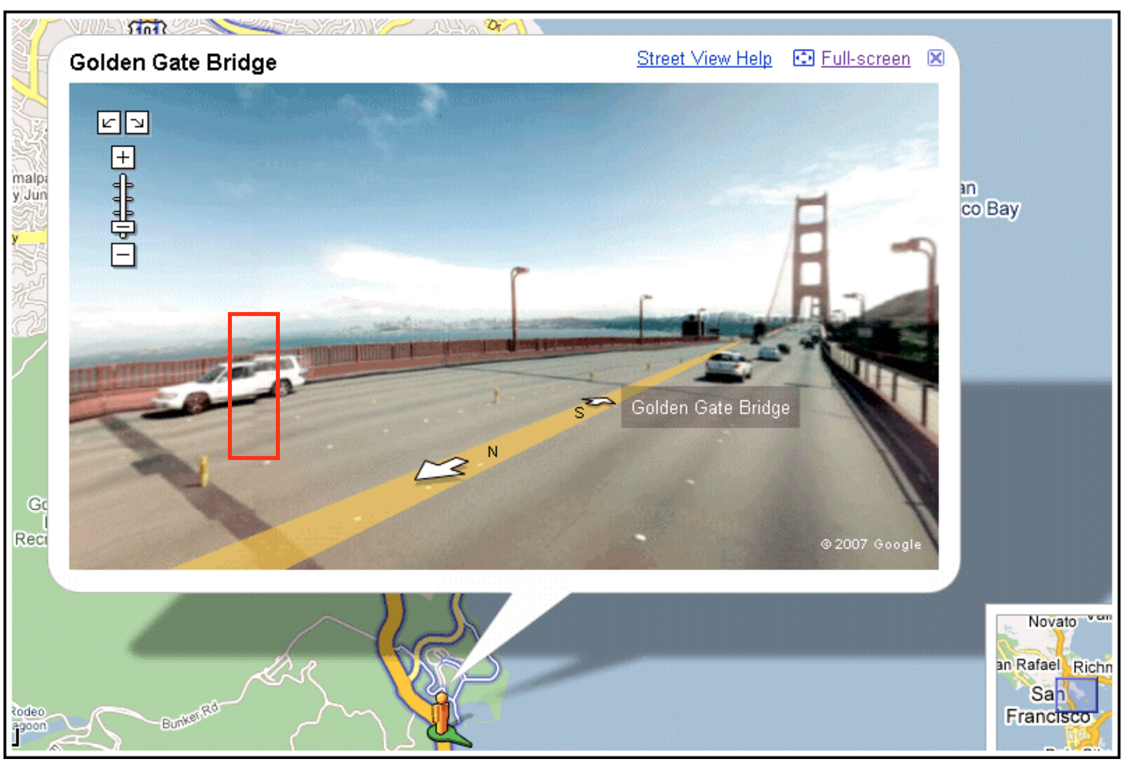

3. Image Blending: How residual parallax is camouflaged, at a cost

Even when image matching is at its best, or all objects distant, there will always remains some residual disparity between images, no matter how small this disparity. The reason for such disparity even after image matching is that, excluding the singular case of exactly constant parallax over a matched image region, image matching can at best provide only an "average" parallax over an image region, and not the exact parallax at any specific point in the region.

It is to camouflage residual parallax — and for no other reason — that independent of whether any image matching is performed, component images from outwardly pointing cameras are unfailingly blended together to create composite images, such blending causing blurring and ghosting in the composite image and thus lowering its resolution relative to its component images. These phenomena are all visible upon zooming into the red boxes in the 3-image panorama below, this panorama taken from Point Grey LadyBug2 camera's website in 2008, the same source as for the earlier indoors image exhibiting ghosting.

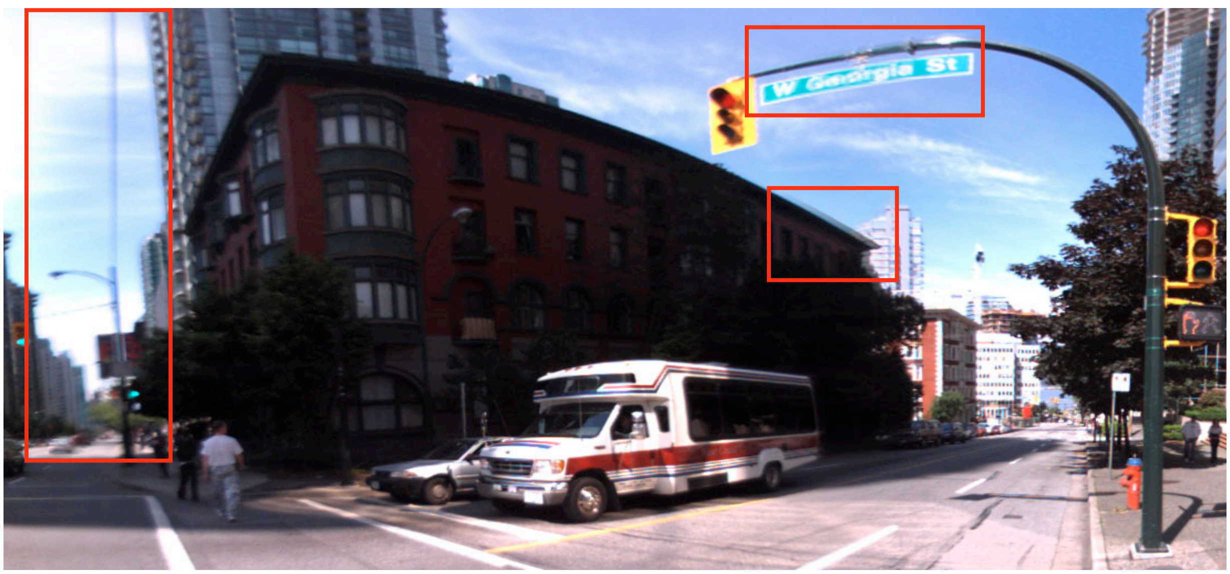

4. What Distinguishes Fullview: Lack of parallax

FullView's patented designs evade parallax by having their multiple cameras each, except at most one, look out off a flat mirror such that each camera is effectively looking out in a different direction but from the same single viewpoint. Below is a pedantic FullView design. Whereas 4 cameras are shown below, the number and type of cameras and lenses, and camera orientations and mirror inclinations, are all design parameters.

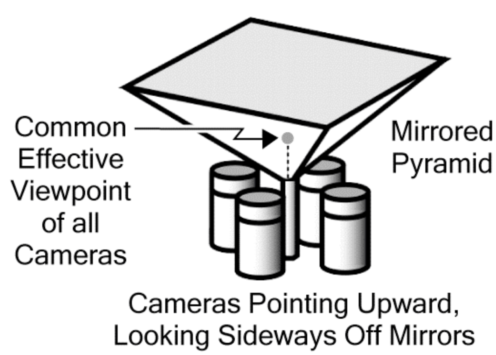

In contrast to imagery from outwardly pointing cameras, FullView's wide-angle imagery, whether video or still, and whatever its resolution, is always created in the same scene-independent way: By placing parallax-free images from multiple cameras side-by-side after warping each image in a predetermined fashion. As a result, FullView's imagery is always seamless, artifact-free and blur-free, irrespective of the scene and how far objects are in it.

This is illustrated below by a 3-image panorama taken inside Bell Labs in 1999 by an early FullView prototype. Notice how well adjacent images line up along each of the two red-boxed vertical seams despite some objects being relatively close to the camera, something outwardly pointing cameras cannot cope with. Please zoom in to see the seams, noting in particular how well the wrinkles in the blue shirt remain intact across one of the two seams.

5. What FullView Offers: Unrivaled wide-angle imagery

A single image sensor using wide-angle optics — e.g., a fisheye lens or a curved mirror — typically produces spherical-view images whose resolution is too low and too uneven for most applications. Whereas using higher-resolution sensors can improve image resolution, both here and of multi-camera designs, there is a price to pay: Camera sensitivity then drops, as each pixel (i.e., sensing element) must then be smaller, and the frame rate drops too, as image data must then be clocked (i.e., moved step-by-step) across a larger array of pixels.

Multi-camera designs in general offer higher and more uniform resolution than any single-camera design. However, all multi-camera designs are not equal: **FullView offers much higher and more uniform resolution than outwardly pointing cameras that use the same number and type of component cameras because, as FullView avoids parallax,**

5.1 FullView has no need to overlap its component fields of view — such overlap necessary for image blending — which allows each of FullView's component fields of view to be smaller than otherwise and thus provide higher resolution.

5.2 FullView has no need to blend its component images, which allows it to avoid the uneven blurring of composite images that always accompanies image blending, such uneven blurring unevenly lowering the effective resolution of the composite image.

In addition to its panoramic-view design of Section 4, **FullView also offers the following spherical-view design**. Note that in it, if 5 cameras were to look off mirrors instead of 4, every camera here would have a one-to-one correspondence with a camera of the LadyBug2.

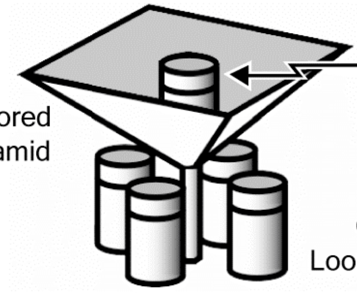

This design, however, has the disadvantage that it provides much lower resolution through its central camera than through its other cameras owing to its central camera's relatively much larger field of view, exactly as in LadyBug2. Then, when seeing the ceiling or the sky is unimportant, **FullView also offers the following panoramic-view design whose vertical field of view is twice that of FullView's pedantic design, but whose resolution is much higher than and more uniform than that of FullView's preceding spherical-view design.**

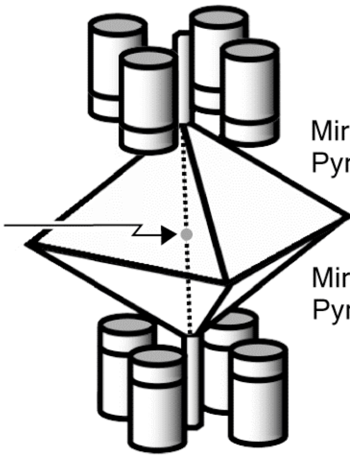

6. Conclusion: Why Settle for Less?®

FullView® cameras offer panoramic and spherical imagery without equal. Technology for these cameras was invented in 1995 at Bell Labs Research, the legendary birthplace of not only the transistor, but also *The Sampling Theorem*, which is the basis of our Digital Age.

FullView alone — through its patented devices, methods and products of methods — can offer you multiple cameras looking out into the world from effectively a single viewpoint, without parallax, to provide you seamless, artifact-free, blur-free composite images. Whereas all FullView's cameras to date have been real-time panoramic video cameras, FullView's parallax-free approach clearly applies to still cameras as well — and to spherical-view cameras — and is independent of the resolution of the component cameras used.

FullView's 10-Camera Unit for U.S. Navy

Microsoft's Licensed 5-Camera RoundTable

What is important here is that FullView's technology is mature and tested, and has been deployed successfully in the most demanding of circumstances: To provide live video in applications with neither an appetite for defects, nor the time to camouflage them. At one end of the spectrum, where uncompromising performance is key, the U.S. Navy has chosen FullView cameras for its next generation of aircraft carriers. At the other end of the spectrum, where dependable, low-cost, artifact-free and blur-free performance is key, Microsoft has chosen to license FullView's technology for its videoconferencing product called the Microsoft RoundTable — now sold as Polycom CX5000 — this license sought by Microsoft only after Microsoft had first spent years pursuing outwardly pointing cameras.

To conclude, FullView can reconfigure any set of outwardly pointing cameras to remove parallax between them, which allows it to provide you not only seamless, artifact-free and blur-free imagery, but also much higher and more uniform resolution than others.

ABOUT THE AUTHOR

Vic (Vishvjit S.) Nalwa is President of FullView, which he cofounded with Lucent Technologies in 2000. He invented the original FullView camera at Bell Labs in 1995, in recognition of which he was elected a *Fellow of the IEEE* in 2004.

After his Junior Year at St. Columba's High School, New Delhi, India, he attended the Indian Institute of Technology (IIT), Kanpur, India, from where he graduated with the B.Tech. Degree in Electrical Engineering in 1983. At IIT Kanpur, he won *The First Prize for Academic Excellence in the Core Curriculum* in 1981, and he was subsequently *The Best Graduating Student in Electrical Engineering* in 1983, both jointly. He then received Stanford University's inaugural *Information Systems Laboratory Fellowship*, graduating from Stanford with the M.S. and Ph.D. Degrees in Electrical Engineering in 1985 and 1987, respectively.

From 1987 to 2000, he was a Principal Investigator at Bell Labs Research, where over the summer of 1993, he pioneered Automatic On-Line Signature Verification — creating a system whose equal-error rate was less than a tenth that of three different competing multi-year group efforts by the renowned Statistics, Speech and Neural Networks Departments of Bell Labs Research. For this, he won a Bell-Labs-wide competition in 1994 and he was thereon afforded unfettered freedom by the President of Bell Labs, freedom that led to FullView in 2000. In 1989, he was concurrently on the faculty of Electrical Engineering at Princeton University.

He is the author of the text A Guided Tour of Computer Vision, Addison-Wesley, MA, 1993. From 1994 to 1998, he was on the Editorial Board of the *IEEE Transactions on Pattern Analysis and Machine Intelligence*. He has published several well-known academic papers, received numerous well-known patents and academic honors, and given invited talks at numerous well-known academic institutions, among them, in alphabetical order, CMU, Dartmouth, Delft IT, Harvard, IIT Delhi, INRIA, MIT, NYU, Princeton, Stanford, Technion, U British Columbia, UC Berkeley, U Illinois, U Maryland, U Penn, USC and Yale.

