Open-Area Smoke Imaging Detection (OSID).
F.P.R.F. SUPDET 2010 Full Paper

Ron Knox
Xtralis Pty Ltd.

BACKGROUND

A number of large-open spaces in the built environment present unique challenges to fire detection systems. Buildings such as stadiums, large atria, airports and rail stations, hotels and convention centers and warehouses demand a fire detection solution that is sensitive to diluted smoke but that is non-intrusive on the space.

In the SUPDET Conference of 2008 Xtralis presented a novel technique for open-area smoke detection called OLSD (Open-area Light-scattering Smoke Detection). This technology works by measuring the light scattered off-axis from a laser light beam by smoke particles in that beam, thus revealing both the concentration and location of the smoke.

Although this technology works extremely well, especially for Very Early Warning detection, it was realized there are a number of non-technical barriers to transforming this into a marketable product, not least of which is that the very novelty of the principle of operation makes it difficult to get type approval against existing codes and standards. It remains the intention to commercialize OLSD as a Very Early Warning system, but in addition a new approach known as Open-area Smoke-imaging Detection (OSID) has been invented and developed for use in normal sensitivity applications.

The new OSID technology re-uses many of the techniques developed in the OLSD program for image capture and processing. However, the principle of operation of this technology differs from OLSD in a key area. Essentially, rather than using light-scattering, OSID measures the extinction along a light beam caused by smoke particles along the direct path of the light beam, just like a traditional projected beam smoke detector, except that by using dual wavelengths and image processing many limitations are overcome and benefits added. Moreover, it can be approved against existing standards without requiring severe re-interpretations.

To understand the operation of the new system we must first look at how existing projected beam products work, and at their benefits and limitations.

Principle of Operation of Traditional Projected Beam Detectors

The projected beam smoke detector is arguably the easiest to understand of all the smoke sensors available, since it relates well to the simple human observation that a light gets dimmer when smoke obscures the view. Inherently, measuring light attenuation (aka extinction, or obscuration) cannot reach the level of stability and hence sensitivity of a light scattering detector, particularly over short distances. Hence beam detectors are not generally considered as capable of very early warning performance as are light scattering instruments. Put simply, this is because a scattering detector is measuring a large increase in a near-zero signal whereas an extinction detector needs to resolve a small decrease in a big signal. This gives an inherently lower stability, higher noise reading.

Nonetheless, beam detectors when applied correctly can be surprisingly effective in many circumstances and can surpass spot-detector performance [Ref 1]. But they do have some fundamental issues that have often caused them to be a “grudge buy” and to be considered as a low-end, cost-driven choice, suitable only when nothing else can be made to fit.

A classical beam detector uses two units, called a transmitter and a receiver. Inside the transmitter a light source (typically an infra-red LED) flashes periodically. Light from the LED is focused into a tight beam by a lens, and a finely adjustable mechanism is provided to allow the beam to be directed from where the transmitter is mounted at one end of a room towards the receiver at the other end, which could be 100m away. The receiver also has an alignment mechanism and a lens which focuses the beam onto a light sensor, typically a silicon photodiode. In highly simplified form the key optical elements are represented in this diagram.

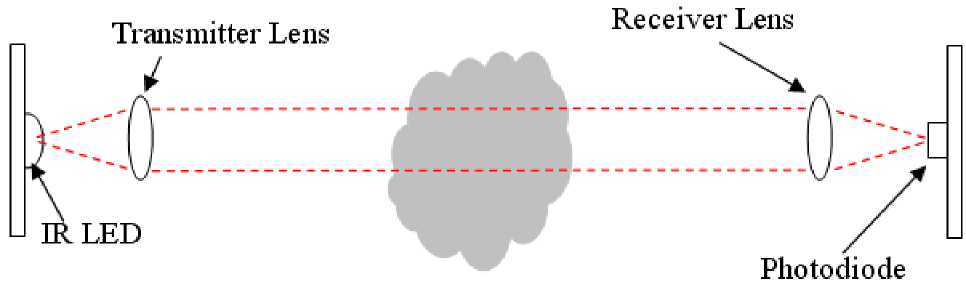

The electrical output from this photodiode is amplified and measured so that the signal reduction due to smoke present between the transmitter and receiver can be determined. Usually, the transmitter and receiver are wired together so that the light pulse is synchronized with the receiver. In alternative designs the transmitter and the receiver are housed together in a single enclosure, and are aligned onto a remote reflector. The reflector is not a flat mirror (as this would need to be exactly aligned) but is made up of corner reflector elements which reflect the light strongly back towards the source.

Problems with traditional beam detectors

The perceived issues with traditional beam detectors are primarily difficulty in alignment and proneness to false alarms. Such false alarms may be triggered by, for example, objects such as banners, balloons or even birds entering the beam path, dust in the air or insects such as moths crawling on the optical surfaces of the transmitter, receiver or reflector. Normal building movement caused by temperature changes etc will also affect alignment.

Typical beam detectors require the initial alignment to be accurate to about 0.1 degrees of movement; tricky but achievable. Some designs of detector use software-controlled motor driven mechanisms to adjust the fine alignment automatically to obtain and subsequently maintain the strongest signal available.

OSID PRINCIPLE OF OPERATION

The improvements that OSID offers against traditional beam detectors stem from three core design ideas [Ref 3, 4, 5]:

- Two wavelengths of light are used:
  - Use of ultra-violet (UV) and infra-red (IR) wavelengths outside the human visible range assists the identification of real smoke compared to larger objects such as fork-lift trucks, insects and dust; thus reducing opportunities for false alarms.

- A CMOS imaging chip with many pixels (just as used in a digital camera) is used rather than a single photo-diode, providing:
  - Multiple source capability (i.e. several beams into a single receiver)
  - Automatic alignment and movement tracking by software only
  - Location of the smoke in a large space

- A unique method for aligning

A core objective remains that the reading of smoke concentration is ***quantifiable*** and is ultimately traceable back to an industry/scientific standard. This requires that the illumination is "active", i.e. provided with intensity and timing under the direct control of the system. Variation in lighting conditions from total darkness to full sunlight should have no effect on the measurement system; neither should low contrast conditions like white smoke on a white background.

APPLICATION & BENEFITS

Simple linear configuration

In its simplest configuration, a system consists of one *Emitter* and one *Imager* placed on opposite walls, and roughly aligned with one another.

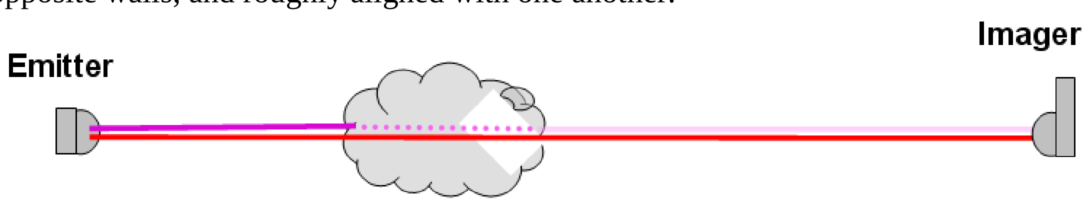

Figure 1 Linear OSID layout

Again, this is a simplified diagram for clarity; in reality the beams from the Emitters are conical with a width of +/-5º. Roughly pre-aligning the Emitter to 0.5º is easy using an alignment tool described in more detail later. For a single beam path, an Imager fitted with a "telescopic" lens giving +/-5º field-of-view is used and this too can be easily pre-aligned with the same tool. With 10º total field of view lens the maximum range for the system is ~150m (500 feet). Note that the field of view angle given is for the horizontal plane; since the Imager chip has a conventional 14:9 aspect ratio the vertical axis is reduced accordingly.

One significant benefit is that, rather than having to mechanically align the optical system with great precision, the exact location of the Emitter in the Imager's field of view is determined automatically by the Imager software. This software identifies the location of the image of the Emitter which may be anywhere on the active surface of the Imager chip. To visualize this, one might simply think of the active part of the imaging chip as creating the picture that would be seen on a normal TV or computer monitor screen. In a single frame taken when the transmitter is blinking on, it appears as a bright spot in the picture. Any future re-positioning of the image caused by building movement is also tracked by the software, eliminating false alarms due to movement without needing any motorized mechanical parts.

Resistance to false alarms

If any smoke enters the beam the small particles in the smoke will reduce the UV light transmission more significantly than the IR light transmission, whereas dust & objects affect both equally. Software can examine the strengths of these signals, and how they change over time, and make a determination of whether to raise an alarm or flag a trouble condition. This appropriate use of UV as well as IR light both reduces the probability of a false alarm and enhances the sensitivity to small particle smokes, which has often been relatively low on optical detectors.

Additionally, clear materials loaded with dyes which absorb UV and IR to differing extents can be used as a convenient "smokeless" test sheet for commissioning and maintenance tests in the field. Conventional beam detector filters obscure both wavelengths and are simply reported as a trouble condition, not as a fire.

Multiple Emitters for area protection

For protecting a room, up to 7 Emitters can be deployed around the walls.

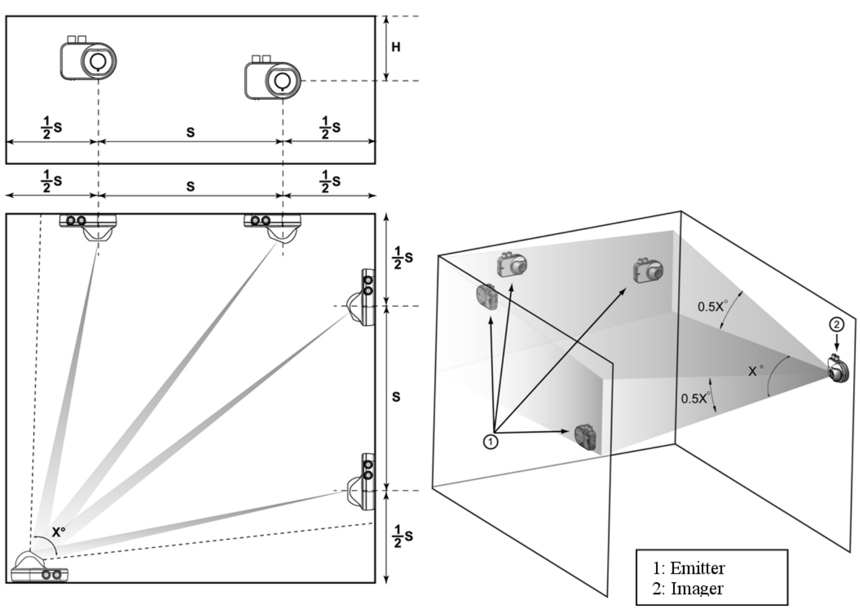

Figure 2 Multi-Emitter area OSID layout

In this case a 90º lens (+/- 45º field of view horizontally) can be fitted to the Imager. With 90º field of view lens the maximum range for the system is 34m (110 feet).

Alternative lens options include a 45º field of view giving a range of 70m (220 feet).

An important target of the design is that the overall power consumption of the emitter is very low, enabling it to run for many years on an internal battery (although a wired version is to be made available for those who wish it). By using modern Lithium battery, up to 10 years of life can be anticipated. This reduces installed cost by removing the need for cabling to the Emitters. Installation of multi-beam applications needs wiring only between the Imager and a fire panel.

Range could be further increased by increasing the Emitter power. However, since four times the power is needed to obtain twice the range, the effect on battery life is significant, so such systems would usually be wired.

TECHNOLOGY DETAILS

Pre-alignment

To operate correctly it is only necessary that the Imager and Emitters are very roughly pre-aligned by hand to ensure that all Emitters are comfortably contained within the field of view, and that the Imager falls within the wide beams of the Emitters.

This could be achieved in many ways, but the design choice for the current product design is to mount both the Emitter and the Imager optical assemblies inside a “ball and socket” housing which allows a range of movement of +/- 60º in the horizontal and +/-15º in the vertical. These may be swapped around if needed, eg to look down a staircase, by simply mounting the housing on the wall in a 90º rotated orientation.

The ball housing is supplied free to move, but can be rigidly locked into place using a tool with a hex-key end. This tool engages with a steel-lined aperture at the front of the ball which is precisely aligned to the optical centre of the Imager or Emitter. The tool is equipped with a pre-aligned laser pointer, the design of which is similar to devices used to pre-set telescopic rifle sights. The installation procedure is simply to position the laser spot as required (details below) and to rotate the tool by one quarter-turn to both lock the ball in place, and in the case of the Emitter, to switch it on. A reed switch activated by a magnet on the locking mechanism switches on the Emitter only when mounted and locked. This prevents the battery discharging during shipping and storage, and also confirms that the unit has been aligned and locked in place. This alignment and fastening tool has been dubbed a “laser screwdriver”.

When the system is using a single Emitter at a long range the field of view of the lens may be 10º width in total; ie +/- 5 degrees. Preferably, the Emitter image appears in roughly the middle of the picture. Achieving 0.5 degrees is easy (this represents a target about 1 foot across 60 feet away) in fact, people find it harder to resist spending unnecessary time by getting the spot exactly onto the 2 inch window.

When a 90º field of view lens is used with multiple Emitters distributed around a room, perhaps at different heights, the laser is firstly directed onto the estimated mid-point of the Emitters. It is then also necessary to verify that all of the Emitters are comfortably captured in the field of view. For this, a further tool is inserted into the aperture in the ball. This tool has a joint designed to allow the laser spot to trace a rectangle around the room matching the field of view of the lens.

Emitter Operation

The OSID Emitter contains a number of key elements, illustrated in this diagram.

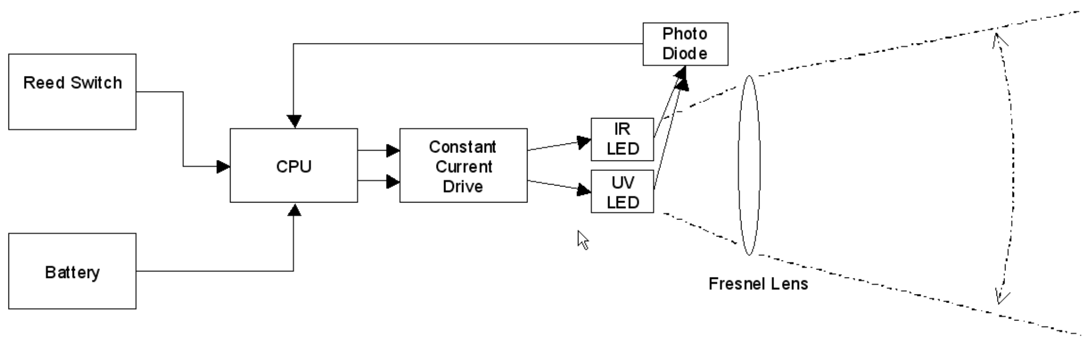

Figure 3 Emitter Unit Physical Architecture

Light from the two (or more) LEDs is focused into a projected beam by a Fresnel lens/diffuser into a diverging beam of about +/-5º width. It is desirable that the LED chips are physically close together to minimize divergence between the UV and IR beam patterns and to prevent small objects like insects from blocking one wavelength more than the other, and so the LED dies are mounted side-by-side within a custom made package.

A low power micro-controller (Texas Instruments MSP430 variant) delivers a carefully defined sequence of pulses to the LEDs, which is unique to each Emitter made, as is expanded on below.

The intensity of each pulse is measured by the photodiode and the CPU's internal A/D converter. This measurement is used to provide compensation for LED temperature and ageing effects. Varying the drive pulse compensates the intensity as required. The effects of battery voltage and LED forward voltage variations are eliminated by the use of a constant current drive circuit.

Imager Operation

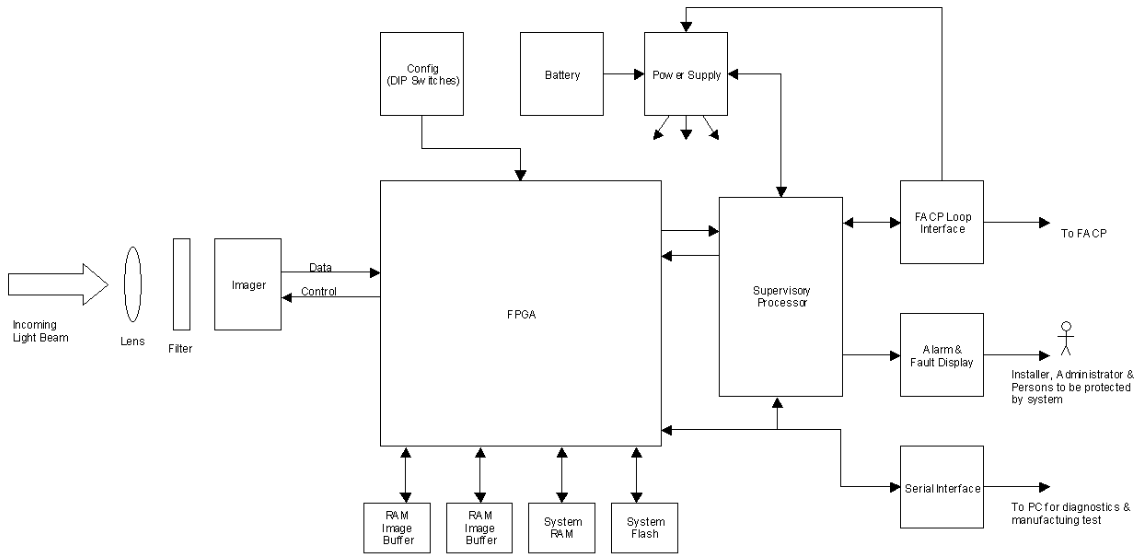

Figure 4 Imaging Unit Physical Architecture

With the exception of the field wiring termination card, the OSID Imager optics and electronics are fully housed within a ball that can be moved to pre-align the device. A flat flex cable links the electronics to the termination card. The optics is internally sealed in a moisture proof tube, equipped with an optional heater to prevent condensation forming on the outer surface in humid / cold condensing environments.

The Imager is fitted with a CCTV type lens that has been selected for minimum dispersion for UV and IR wavelengths (i.e. the focal length of the lens for both wavelengths are nearly identical) and for good temperature response characteristics.

The Imager is also fitted with a dyed glass filter, designed to be almost opaque to all but the 2 wavelengths of interest. This contributes to the systems ability to work in a wide range of lighting conditions, including a full sunlit scene, and in strong artificial lighting including flickering sources like mercury vapor lamps.

However, most of the sensitivity and tolerance to bright lighting comes from the technique of "background subtraction" previously described in the OLSD SUPDET paper in 2008 [Ref 2]. This uses the very fast capture speed of the imaging chip to measure the light level around the Emitter image immediately before and after the wanted flashes, and then subtracts them so that the uncorrelated background contribution disappears entirely. This technique requires fast processing, which is readily available if one is prepared to use one of the many energy-hungry microprocessors on the market. As the desire for this product is to operate from a fire panel's limited power supply a Field Programmable Gate Array (FPGA) which can be programmed for high performance at low consumption was selected. Programming such a device is much more laborious, but ultimately gives excellent performance. The FPGA selected also permits a sophisticated power saving strategy, as detailed later.

CHALLENGES IN THE TECHNOLOGY DEVELOPMENT

Using a video imaging chip rather than a simple photo-diode provides benefits for a beam detector, but in practice these techniques are challenging to apply, especially when tight budgets apply to both cost and electrical power consumption (to permit powering directly from a fire panel). For one thing, affordable video imaging devices are not designed with calibrated measurement of UV and IR light levels in mind, and some ingenious methods were required to be developed.

The smaller particles that are generated in most fires at an early stage, and by all fires once they transition to the increased threat stage of flaming, interact more strongly with short wavelength light (UV of ~400nm) than they do with long wavelength light (IR of ~850nm). This is a consequence of the Mie theory of light scattering. Incidentally, this *scattering* theory is still applicable to an *extinction* measuring device, since much of the light beam is not actually absorbed by the smoke so much as scattered away from the receiver. Having said that, black smokes do absorb a larger proportion of the light, which is exactly why they look black, and is also why light-scattering detectors and light-extinction detectors cannot be calibrated to give the same readings for both white smoke and black smoke.

In an ideal world, more than 2 wavelengths at even wider separations could be used to give further particle size characterization. However, the cost/benefit tradeoffs give rapidly diminishing returns.

Practical limitations to the wavelengths chosen are driven by design choices and constraints including:

- Affordable light sources are required. Today, that means reasonably low-cost high intensity LEDs.
- The light from the source should be as near to invisible as possible; to avoid annoying visibility in deliberately darkened environments such as entertainment venues.
- While affordable conventional imaging chips manufactured in volume for CCTV or machine-vision applications do provide response to wavelengths outside the human-visible spectrum, they are beginning to roll-off in sensitivity, as shown in Figure 5 - Imager Quantum Efficiency vs. Wavelength Curve. Fortunately, the signal to noise ratio is still perfectly sufficient for good performance. Another aspect is that part to part variation in both the LED’s wavelength and the Imager’s response curve must be accommodated in the design. This is managed at manufacturing by the addition of a test and calibration step that verifies that the wavelengths are in specification, and programs into the Emitter compensation parameters which are required. These are later communicated to the Imager in the field as part of a data transmission encoded in the flash sequence.

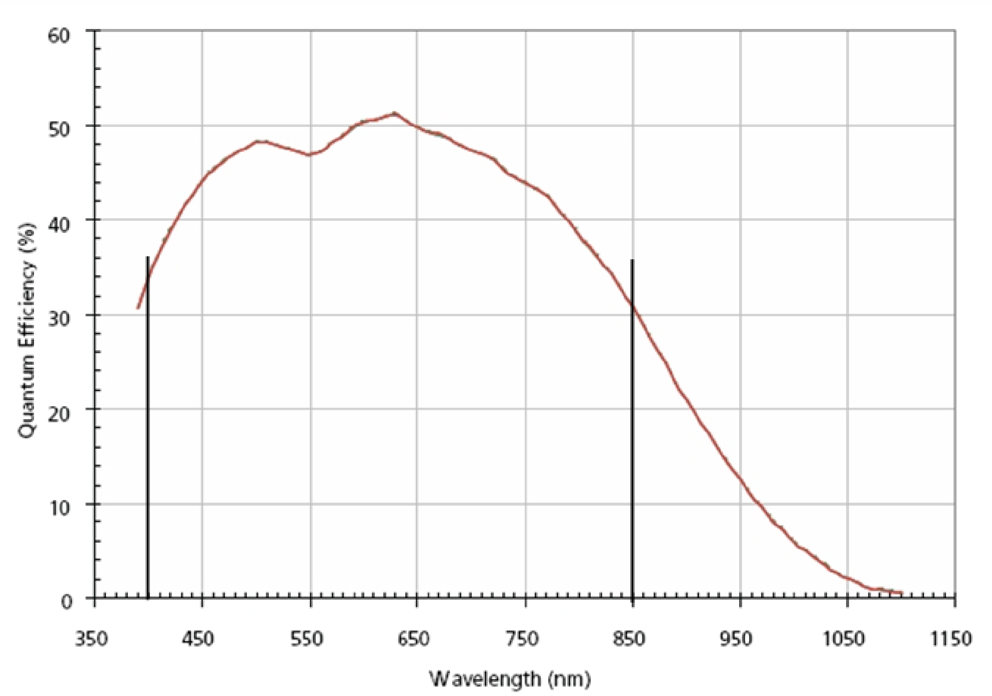

Figure 5 - Imager Quantum Efficiency vs. Wavelength Curve

Conceptually, a monochrome video imaging chip is just a grid of many light sensitive elements, each one acting like a photodiode to form a single pixel in a picture. It also contains circuitry and micro-code to capture individual picture frames and to transfer each pixel output to memory for processing. After being converted to a digital value, the signal from each pixel is expressed in “grey levels” going from black to full brightness. Ideally, each pixel should range from 0 grey levels, meaning complete darkness, up to about a typical maximum of about 1000, meaning full measurable intensity. Above this the pixel is said to be saturated; which is to be avoided since any stronger light level cannot be measured.

When the unit is first started up the software has no knowledge of exactly where in the picture frame any Emitters might appear; so it begins a search. The CMOS video imaging chip used has a few hundreds of thousands of pixels to search, and each Emitter flashes for less than 1/1,000th of the time, so finding an Emitter (and there may up to 7 in the view) is challenging. This is made more difficult by the fact that there may well be many other bright and varying light sources in the picture. In one example demonstration we have conducted the room was lit by chandeliers each holding dozens of fluorescent light bulbs, each flickering on and off faster than the human eye can see. Fortunately, imaging chips are available that can operate at high speed, and they can be re-programmed in real time to capture only partial frames even faster. In that way, the system can first identify candidate light sources that *might* be Emitters and then examine them closely to determine if they have the right timing characteristics to be definitely identified as the wanted sources.

The Emitter design used may be battery powered and so are free-running. This creates the issue that any two Emitters might happen to flash at close to the same time, causing a timing collision. This may seem unlikely at first consideration since the duty cycle of less than 1 in 1000 is so low, but in practice 7 independent oscillators will drift in and out of phase and inevitably cause unacceptable periods of overlap.

While in principle the Imager can see all Emitters at once, in practice it can only capture a partial frame quickly, and so any timing collisions must be brief events that preferably cause only one or two flashes to be skipped and so do not impact system performance. To achieve this every single Emitter has a unique code identifier which is communicated to the Imager via a data pulse attached to each flash sequence which is used to 'jitter' the flash sequence timing in a unique but predictable way so that no 2 units will ever stay in lock-step.

A further complication arises from the desire to have very low electrical power consumption. Although the imaging system can see more than one Emitter in the picture there is insufficient electrical power available to leave the camera and processor running continuously, therefore both need to be put into a suspended (or sleep) state in between the expected flashes from the Emitters. Before the imaging chip and processor are suspended, a small independent timer (acting like an alarm clock) is set to wake them up just before the next Emitter flash is due. On awakening, the next frame is captured and processed; any alarm and trouble conditions are analyzed; I/O functions are performed; the pseudo-random timing of the next Emitter flash is calculated and the alarm clock is accordingly re-adjusted and the system goes back to sleep for another few milli-seconds.

Research

Apart from theoretical response curves based on light scattering and absorption theory, it is obviously essential to perform real smoke tests to validate a system's performance. In Xtralis's laboratory facilities the technology has been tested with a wide range of smoke types, including the standard UL and EN test fires; and with a wide range of nuisance sources including the ISO standard dusts, steam and talcum powder. OSID shows a good and consistent responsiveness to all of the common smoke types. These test results are proprietary, and may be published separately in the future.

Conclusion

In fire detection the most important trade-off is the reliable detection of actual fire threats while minimizing the cost, disruption and perhaps most importantly the loss of credibility caused by false alarms. The industry strives to improve methods to reliably and economically identify false stimuli such as dust, steam and macroscopic objects while ensuring a safe response to real threats. Dual wavelength measurement alone is not a complete panacea, but used intelligently and in combination with both a careful signal analysis approach and an imaging-based automatic alignment scheme it can substantially improve a fundamentally promising technology that has developed a negative reputation. As familiarity and confidence grows we expect that the new technology will expand into many new application areas.

REFERENCES

1. Projected-Beam Smoke Detectors - More Than Just a Substitute for Spot Detectors. Fire Protection Engineering; Journal. NEMA, Summer 2004. http://www.fpemag.com/archives/article.asp?issue_id=21&i=179&p=1
2. Open Area Light-Scattering Smoke Detection, Xtralis Pty Ltd.
Fire Protection Research Foundation – Proceedings SUPDET 2008
3. Knox et al; International Patent Number WO2004/102498 “Particle Detector” Nov 2004.
4. Knox, Boettger, Ajay; International Patent Number WO2006/050570 “Particle Detector, System and Method” 2006
5. PCT/AU2009/000727, PCT/AU2009/901922..27 “Particle detection” 2008/2009