

Introduction

There are now several air sampling smoke detection systems on the world Fire Protection market that operate at 'High Sensitivity' levels as defined by the tests described in BS6266 1992. The AirSense Technology Ltd **Stratos-HSSD**[®] detector is at time of writing, the latest entry in this field, and is designed to overcome many of the shortcomings found in other systems. In order to understand Stratos it is useful to have a cursory knowledge of the fire detection systems that preceded it, whether point detectors, large area detectors or other aspirating smoke detection systems.

It is a truism that early detection of fires is the key factor in preventing them from developing. The earlier they are detected and extinguished, the less damage will be caused. It is not until one looks at the evolution of detection systems that one can see where such a simple truism can lead. The major requirement of any system is, obviously, to detect fire, but, not quite so obviously, to **not** give a warning when **no** fire occurs! The problem of unwanted alarms is a major one; for instance frequent call outs to deal with non-existent fires can lead to a tendency for alarms to be ignored. Unwanted alarms, from whatever cause, can largely negate the usefulness of any alarm system and for this reason the ability to not give unwanted alarms must be rated as high in importance as the ability to signal a true alarm.

About the author

The Kit Girling has been involved for many years with the design of various successful fire detection products. For many years he was also Electronic Engineering Manager of the Company that first imported High Sensitivity aspirating Smoke Detection products to Europe from Australia. He was subsequently involved in making this earlier generation system suitable for application in the European fire market and developing differing variants of this system. Since February '93, Kit Girling has been responsible for leading the design team that has developed the Stratos-HSSD[®] system.

Background

Many principles have been used to detect fire and are still in use, although some are only used in special applications. The following comments describe some of the more commonly used principles:

Heat Detectors

All fires give off heat, and sensing the temperature of the air at a given point in a given zone can indicate the presence of a fire occurring in it. The method is simply to sense if the temperature rises above a pre-determined trigger level and give a warning. If the trigger temperature is low, relatively early warning can be given. Unfortunately an unwanted alarm could be generated by heating systems being switched on or a heat wave occurring, taking it above the trigger temperature. If the trigger temperature is high, the fire may be well alight before warning is given. Such a system is unlikely to be ideal in many applications.

In order to improve on this situation, the 'Rate of Rise' temperature detector is more often used. These sensors do not measure the actual temperature, but the rate at which it increases. In this detector, the background temperature is largely ignored, but if the rate at which the temperature rises is greater than a given amount, it is assumed that this is due to a fire. Unfortunately, for general use, the unwanted alarm conditions still occur even if not so readily as with the simple heat detector; also the fire must be well alight to generate the rates of rise of temperature required to trigger them.

Heat detectors must either be strategically placed with respect to potential fire sources or have a very high placement density in order to be effective. The heat of the fire must also be transmitted to them rapidly.

Ionisation Detectors

Carbon products are given off in the form of smoke and gasses when a fire occurs. If a small sample of air is ionised by means of a radioactive substance, then the ions will allow a small electrical current to flow between two electrodes placed in the sample. If the combustion products mix with the air in the sample then they will inhibit the movement of the ions that then as a result de-ionises, and the electrical current will decrease. This is the operating principle of the Ionisation detector. A low concentration in the air of the products of combustion is required to effect a detectable change in the current, and this can make the detector quite sensitive. Unfortunately the current is also effected by draughts in the air and to a lesser degree by humidity, atmospheric pressure and temperature. Although these unwanted effects can be minimised, the combination of them, and the very small nature of the ionisation current (10 - 50 picoamps), limit the sensitivity and usefulness of ionisation detectors. It should also be noted that the use of radioactive elements is generally frowned upon by today's society because of the difficulty in safely disposing of them. The main advantage of this type of detector is its simplicity of manufacture and low cost.

General Optical Detectors

The standard used for measuring the concentration of smoke in air is the amount of attenuation it will cause in the intensity of a light beam shining through it, usually expressed as 'percent obscuration per metre'. When the smoke is invisible, such as that which can effect an ionisation detector, then a correlation is derived between the different types of smoke, visible and invisible. The measurement units for optical detection systems always end up as 'percent attenuation per unit length' because they can be translated into an accurate laboratory measurements.

The different principles of Optical detection are described below.

Beam Detectors

A simplified adaptation of the laboratory meter described above immediately becomes interesting as a smoke detector. The laboratory instrument itself is costly to manufacture, and is impractical as a commercial smoke detector due to calibration problems in an industrial situation. In order to make it commercially viable, the specification must be changed so that dust build-up etc. does not cause unwanted alarms. A large attenuation of the beam is required (typically 30 - 50%) and consequently it needs to cover large areas. The attenuation in light is achieved by several means, but a commonly used method is to shine a collimated light beam through many metres to a reflector from which the light is reflected toward a light sensor near the beam source. Smoke occurring in the beam will attenuate its intensity, and this reduction in intensity will be used to trigger an alarm. Such a system is also good at detecting mist, steam or dust, also anything that may physically interrupt or reduce the beam intensity. Dirt build-up on lenses and reflectors have been known to cause serious problems, as have normal movement in buildings caused by temperature changes and wind.

Light Scattering Detectors

Another type of optical detector is immediately suggested by the beam detector, when one considers what happens to the light in the beam that is **not** returned to the sensor. Very little of this light will be 'absorbed' (e.g. translated to another energy form). Most of it will be scattered out of the light beam, and hence not returned to the sensor. The suggestion is, that if the scattered light itself is sensed, its intensity would be a direct measure of the scattering medium or smoke. This principle is used very successfully in many detectors although it is more expensive to realise than an ionisation type detector, but gives a more reliable and potentially more sensitive detector.

A second type of optical light scattering detector uses a tightly focused laser beam to produce pulses of light which are scattered from individual particles in an air stream passing through the point of focus of the laser. This is the 'particle counter' detector, which was originally designed for monitoring the quality of Clean Rooms (microelectronics and pharmaceutical manufacture and research). For smoke detection it is only viable when used in an aspirating system.

Light Sensing Detectors

Most fires will generate light of one sort or another, infrared light being the most prominent. Extremely sensitive infrared cameras are available, which could be used in an alarm system. Such cameras are not used due to their high cost and the problem of identifying the position of a fire but they could provide a means of identifying very small temperature rises. The source of heat must be visible to the detector, either directly or by reflection. This is not the case though with most other types of detector. A simplified version that did not produce an image could be acceptable for fire detection although it would not pinpoint the source of the heat in the same way that a camera would.

A flaming fire will produce a flickering ultraviolet light, and an ultraviolet sensor can be used to detect it while also discriminating against other sources of ultraviolet, arc welding for example, by means of the flicker rate. When materials burn they produce light of different colours depending upon the elements that constitute the material. These colours are produced in certain proportions and sensors can be made selectively responsive to these colours by means of colour filters. By sensing the colours of light in the required proportions an alarm can be triggered. The response time of such a detector can be extremely rapid.

Light sensing fire detectors tend to have specialised uses, furthermore, although the response time to fire may be rapid enough to detect flame within a few milliseconds, most fire situations will progress through a relatively long smouldering period, and generally light sensing detectors are either not sensitive to this stage or are inefficient at detecting it.

There are many variations on these principles, they are also used with varying degrees of technical expertise in order to suit different markets and applications. In general the sensitivity and reliability of a detector, of whatever principle, can be improved with an increase in cost. The large majority of detector designs are 'point' type, for either consumer or commercial use, where, due to the highly competitive nature of the market, costs have to be kept to the minimum. In the great majority of cases, detectors are operated from a remote control unit, and one control unit will usually have many detectors connected to it. The control unit provides the power necessary for them to operate and itself operates fire sounders, automatic signalling to a remote centre (e.g., the Fire Brigade), fire extinguishing systems, etc. It normally indicates the signals it is receiving from the detectors by means of lamps but the use of special PC programmes to display more comprehensive information, is becoming more common.

Point Detectors

This is the name given to the type of detector normally found in most buildings distributed on the ceilings. A point detector depends upon normal air currents and convection to carry the smoke, products of combustion or heat to it. For this reason, the volume covered by the detector is limited. Where large volumes need to be covered, more detectors are used and in very large volumes a great many detectors may be required to provide complete coverage. There needs to be more than one tier of detectors for very high volumes.

In the earlier days of such detection, all these detectors were wired in parallel on a single pair of wires from a control panel. The wires provided electrical power to the detector and also carried the alarm signal. The cost of the wires makes a significant contribution to the total cost of the installation. Since the system must continue running during a mains power failure, back-up batteries are required to take over at such times. The capacity of the batteries will depend upon how much power each detector needs to

run it. In order to keep the size and cost of batteries low, point detectors are designed to draw very low currents.

Any detector on the pair of wires which signals an alarm is signalled back to the control panel. This type of control panel is unable to say which of the many detectors had operated, so the signalling is only 'zonal'. Later developments use an addressable system whereby the detector is given a coded unique identification which forms part of the signal back to the control panel. The control panel can identify which of the detectors (or other alarm device) on a given pair of wires is signalling. The detector can also signal back more information than a simple alarm state, for instance; it can signal back that it has a fault or pre-alarm condition. This has the effect of allowing the 'pinpointing' of which detector has given the alarm and it also allows rapid identification of certain faults in a detector.

Since a large number of detectors is often needed in a point detection system, the cost of each unit must be kept within tolerable limits. In order to do this the design is kept reasonably simple and the performance is consequently comparatively modest compared to that possible with modern technology. In addition to this, National Standards always demand a limit to maximum sensitivity of these devices in an effort to reduce unwanted alarms. Point detectors are usually set to give an alarm at a smoke density corresponding to between 2% and 4% obscuration per metre. Although they are capable of signalling a fault condition, they will not signal the occurrence of every possible type of fault.

Although these detectors often provide quite adequate protection, the probability of them giving an unwanted alarm could, in some cases, be lessened by a more costly design, and the sensitivity could be increased if cost were no object. Their general performance reliability, although good, may not be as high as that required by a person wishing to protect a very valuable property, and a higher sensitivity would certainly give an earlier warning. Where very valuable property is to be protected by point detectors, or where certain fire extinguishing systems are to be automatically released, two different types of point detector are commonly used in a 'coincidence' (or 'double knock') configuration. This configuration requires an alarm to be indicated by both of the detectors before it is accepted. Since two detectors of the same principle would be equally susceptible to the same source of unwanted alarm, usually they are of different operating principles, e.g. a heat detector and an ionisation detector. The probability that such a configuration will give an unwanted alarm is very much less than with a single detector. Normally detectors are specially selected to suit the hazards in the volume they are protecting; to give a rapid and sensitive alarm. The price paid for this degree of certainty is the time delay required for both detectors to trigger an alarm and the effective sensitivity are reduced to that of the least suitable detector of the pair.

High Sensitivity Aspirating Systems *(a.k.a. Air Sampling Systems)*

Where a large volume is to be covered or an area is to be covered at a high sensitivity; instead of using many low cost, inherently low sensitivity point detectors, an Aspirating system may be used. Such a system consists of a single, more sensitive and highly developed detector and an air sampling system to cover the required volume. The air sampling system consists of a length or lengths of pipe with strategically placed sampling holes along its length. An air impeller is used to draw air along the pipe from the sampling holes and through the detector measuring chamber. The transit time for air to travel from the sampling holes to the detector head can be significant, so there must be significant advantages to this system to offset this delay. The detector is designed to be stable, reliable and highly sensitive (approximately 10 to 200 times the sensitivity of a point detector). Such additional sensitivity allows the detection of smoke at a very much earlier stage in the development of a fire, and as a consequence, this may reduce damage levels considerably. This saving of time more than offsets the delay due to the transit time.

Aspirating smoke detectors are susceptible to considerable dilution of smoke in the air sample, and it is important that this effect is understood. For example, if it is assumed that a detector chamber has a typical sensitivity of 0.1%/metre obscuration, and it is drawing its sample through ten holes in the length of pipe, and that each sampling hole contributes equally to the volume through the detector measuring chamber, then if smoke only effects one of the sampling holes, then, due to dilution from the other 9 holes, the actual detector responsiveness will only be equivalent to a 1% detector. If however, say 5 of the 10 holes are effected by smoke, the apparent sensitivity will be 0.5%/metre. The more widespread the distribution of smoke among the sampling holes, then the higher will be the density of smoke appearing in the chamber, and the *apparent* sensitivity will be greater. Sampling pipe design should take this effect into consideration.

The action of pulling air into the air sampling system gives a better smoke collection method than the naturally occurring air convection relied upon by point detectors, and the sampling pipe systems frequently allow the sampling holes to be strategically better placed, for example in an air extraction system or within smoke strata which may be anticipated.

A high sensitivity aspirating detector can cover the same volume as many point detectors, but with added advantages. Economically speaking this allows more money for the cost of the detector, which, in turn allows it to be designed to be considerably more sensitive, but with the added advantages of being considerably more stable. The removal of stringent economic factors also allow a higher degree of repeatability in performance with a similar reliability to a great quantity of point detectors.

The added advantages of an aspirating system make it the natural choice where very reliable performance is required. The ability to use a very sensitive system is an advantage for protecting very clean environments such as computer rooms or microelectronics or pharmaceutical manufacturing clean rooms. The air sampling inlets can be very nearly invisible which makes them ideal for use in historic buildings, where point smoke detectors would detract from the appeal of the building. Other important applications are found where their specific properties include; museums, art galleries, warehouses, telephone exchanges, penal establishments, dusty or dirty areas, unusually hot or cold areas, areas with high levels of radio energy, atrium buildings, etc.,

The detectors used in successful aspirating systems in recent times, have nearly all been designed to work on the optical, light scattering principle. There is still one produced in the UK that uses an Ionisation chamber and another is available from USA that uses a variant of the Optical principle called a Wilson Cloud Chamber.

The improvements in sensitivity make the optical detector capable of detecting invisible quantities of smoke. One of the advantages of their very high sensitivity is that indication occurs so early that it is of a *potential* rather than an *actual* fire. This allows corrective action to be taken long before an extinguishing action is needed, or indeed, could be effective. This fact alone translates into savings in the cost of extinguishing agents and the possible damage its discharge could incur. Also, some of the most common extinguishing agents are Halons (halogenated hydrocarbons). The manufacture of these materials is being curtailed because of its damaging effect on the Earth's ozone layer.

Thus the air sampling system can overcome most of the disadvantages of other detection methods, but tend to be economically designed to the limit of the market. (i.e. The design maximises the possible performance advantages on the assumption that the customer will pay more for them.)

Specific types of Aspirating Systems.

There are many *aspirating systems* available around the globe, although there are only three generally available *High Sensitivity* aspirating Detection systems in common usage. It is known that there are a few products under development at this time which may become available in this rapidly expanding market sector. In definition of the term High Sensitivity, this is taken to mean that the system is capable of detecting the electronic fire simulations described in the appendix of BS 6266 1992. All three detectors have essentially similar levels of sensitivity, with no single product having any great advantage over the other in this aspect. Interestingly, all three detectors use the Light Scattering principle of operation, but each in a different manner; One uses a Xenon flash bulb as its light source, whilst two take advantage of the very reliable semiconductor laser (in this application, Mean Time Between Failure can be in excess of 1000 years). Of the two laser based products, one is a particle counter, whilst the other is responsive to the Mass of airborne material over a wide size range rather than the number (count) of particles.

A potential problem with High Sensitivity Smoke Detectors is the possibility of alarm signals being given by dust particles. This problem is dealt with in various ways depending upon the actual system and its detection principle, but it is usually at least partly dealt by an air filter of some description on the assumption that dust particles are physically larger than smoke particles. The two laser based systems also use dust discrimination to enhance this effect, although both do so in differing manners. It is also useful to understand that there is an unfortunate effect whereby the more a filter becomes clogged with dust, the smaller the particles are which it will remove out of the air-stream. This can be a very serious shortcoming, because air will still flow easily through the filter when the size of particle being stopped is the size of smoke particles.

Particle counters

Particle counting works on the principle of counting the number of particles in a given volume of air sampled. In order to do this, the rate of air-flow through the chamber must be maintained within limits, or the effect of increased counting due to increasing the air-flow must be allowed for.

The method used is to focus a laser beam into a very small volume in the air-stream and measure the light scattered by a particle entering this volume. This will give a light scatter pulse of a duration and size proportional to the size of particle. By its very nature it is an absolute measuring instrument, i.e. its reading is the number of particles per unit volume; it is not relative to anything. At first sight this appears to be a point in its favour, and for measuring the impurity of air in a clean room it certainly is. It should be taken into consideration though that the smoke level that is indicative of fire is not absolute, it may be any level depending upon the normal amount of smoke present. The increase in the count rate above normal is an indication of unusually large numbers of particles being present. The size of the particle being counted is not a prime factor in the count, since either a very large particle, or a very small particle will increase the count by one unit.

It is assumed that dust particles are larger than smoke particles, and the particle counter can give a signal indicative of particle size and the system may be set to ignore those above a certain size. In relatively clean environments this may entirely obviate the need for a filter. Unfortunately, the size of a fire can only be accurately measured either by the weight or volume (mass) of smoke particles emitted and **not** the *number* of particles. A particle counter will reach saturation when the rate at which it is required to count particles is beyond its capability. This can be near the normal amount of smoke present which includes dust particles when they are present.

Light scatter (mass detection) systems.

These systems look for the amount of light scattered by a stream of air through an optical chamber. With perfectly clean air there is a very small amount of scatter. As the volume of particles below a given size increases, the amount of light scatter increases. Measuring the amount of scatter thereby gives a measure of the volume of impurity in the air. This has a direct relationship to the size of the source fire. The ultimate limit to the sensitivity of such a detector is set by the amount of electrical noise being generated by the sensor, the light source and their associated circuitry. This is because the system must be able to differentiate between the true signal given by smoke, and the noise signals with sufficient reliability to prevent an alarm being triggered by the electrical noise. The amount of light scattered by a given volume of particles is highly dependent upon the wavelength of the light, with the amount of scatter increasing rapidly with a decrease in wavelength. The **direction** of scatter is determined by the relationship between the wavelength of the light and the particle size and shape. This is particularly the case with large particles. It is however of great importance to note that all shapes and sizes of particle give a large amount of scatter in the **forward direction**, that is, scatter that diverts light only a few degrees from its original path. Detectors that rely primarily on small angle scatter will be sensitive to a wide range of particle sizes, almost irrespective of the wavelength of light.

Light sources

One type of aspirating detector, developed in the late 1960's (although reaching the market in the '70's) uses a Xenon flash bulb as the light source. This gives a very wide spectrum of light output extending into very short wavelengths, well below the visible spectrum. The intensity of scattered light at these wavelengths is high but is only of use when a light sensor is used which is sensitive to them. The best response of normal silicon light sensors, as is used in this detector though, only covers the visible spectrum and the light at longer wavelengths. The advantage promised by the broad spectral light output is lost.

Other, more recently developed types use semiconductor lasers giving 5 to 100 mW. of light power. The types used are monochromatic coherent light sources at the red end of the visible light spectrum. The coherent wave front of a laser beam allows excellent manipulation of it by lenses. This does not, at first sight, appear to be a promising light source, but, as pointed out below, can be entirely satisfactory.

Light emitting diodes can give similar light powers to a semiconductor laser and are also nearly monochromatic. They are very much cheaper than lasers, however they do not provide a coherent light source and are usually only (able to be) used in point detectors.

An unfortunate characteristic of Xenon flash tubes is a deterioration of light output over time. Very expensive tubes can be obtained where this effect is reduced to a minimum, but they are neither commercially viable or used in this application. This deterioration provides a very difficult problem to overcome where a detector is intended to behave as an absolute measuring instrument. Due to the deterioration in light output, the signal will deteriorate to a commensurate degree, requiring the detector to be frequently re calibrated. The high electrical noise generated by a Xenon tube will offset the large signals it generates, and the promise of high sensitivity is not fully realisable. The light from a Xenon flash tube is impossible to focus into a tight beam and it is not possible to detect much light scattered at small angles from a wide angled beam. However, complex (dust accumulating) iris systems in the light sensors viewing path can allow a reasonable signal to be obtained from light scattered at larger angles. The Xenon flash tube has a high energy requirement for its drive, which is a disadvantage when intrinsically safe systems are required, or when a system runs from its stand-by battery source. As a consequence, although at first sight the Xenon lamp looks ideal as a light source, it has problems that effectively rule it out as a technically viable one with the advent of

semiconductor lasers. Although future developments could conceivably change this situation, it is difficult to envisage that the effort required would be practical or cost effective when the alternative of a laser exists.

The semiconductor laser is at present the subject of many semiconductor manufacturers development programmes because of its application in printers, compact disc players photo-copiers and fibre optic systems. These are expanding markets and auger well for future improvements in cost and performance. Using a laser as the light source in a smoke detection chamber has immediate advantages and disadvantages when compared with a Xenon flash tube. The main disadvantage is that it gives light at a single wavelength and this is at the long wavelength end of the visible spectrum, where the amount of scatter signal is low. However, the single wavelength, coherent light source that it provides, can be focused to a very tight beam allowing light scattered at very small angles to be sensed without complex iris arrangements. This may be seen to more than offset the disadvantage of the low intensity of scatter at these wavelengths, since the intensity of scatter at small angles is consistently comparatively high, regardless of particle size. Practical signals can be obtained, as with the Xenon flash tube, for particles down to large molecule (gas) size. The semiconductor laser gives a consistent intensity of light output with time and, when it is correctly driven, has an extraordinary length of life. These advantages score heavily in a competition against the Xenon flash tube.

Absolute v. relative scaling

Absolute Scaling

The output of the smoke detector in the conventional aspirating smoke detector system is assumed to be (essentially) zero for a clean air sample and a fixed known level for a given amount of smoke pollution. It is, in effect, a meter giving an absolute measurement of 'smoke-like' pollution. The full scale reading may be set to a minimum of 0.05% obscuration per metre which makes it very sensitive. A usual full scale reading is between 0.1 - 0.2% obscuration per metre. The zero of the scale corresponds to zero pollution.

There are only a few applications where the air is totally free of pollution. There is normally a constantly varying amount of background pollution. With the '**Absolute**' type of detector the alarm levels must be individually set to trigger at the desired level. The correct level is one that will be reached by as low an amount of pollution as is possible without being reached by the normally occurring background variations. Alarms triggered by the normal background amounts of pollution will be unwanted alarms which, as previously explained, must be avoided. The background levels will vary according to the time of day and the particular activities in the protected area. The requirement then, is to find what are the appropriate levels to set the pre-alarm and alarm triggers points. In order to do this the detector output is used in the following way: Smoke is sampled during the normal usage of the system for a period of 10 to 14 days, and a chart recording of the levels against time is generated. This chart is assumed to represent reasonably the future variations that will occur. The chart is examined for high smoke levels during both the day and the night periods, and the pre-alarm and alarm levels for day and night usage are both set to be higher than these respective peak levels. The importance of avoiding unwanted alarms is so great that the trigger levels are routinely set, as a matter of judgement on the part of the installer, to be *well* above the peak levels, in case the chart recording is not truly representative of the future readings. The aim of the exercise is to set the trigger levels at a point where the *probability* of an unwanted alarm is very small and where there is still a high probability of detecting unusually large signals.

There are a number of problems with the absolute scaling system;

- The background level does not always vary as it did during the first two weeks after installation.
- The peaks indicated on the chart recorder may correspond to levels that should have triggered an alarm.
- The two weeks required for setting it up are inconvenient and are hence frequently omitted or reduced, and the alarm trigger levels set at a 'safe' level (i.e., 'safe' from unwanted alarms where the system is comparatively insensitive).

As with all very sensitive equipment, variations in the detector zero point and sensitivity will occur, and these variations must be taken into account in the installer's judgement, when setting the trigger levels. An arbitrary element is contained in this procedure, and it is unlikely that two different installers will set the same trigger levels.

Relative Scaling

An alternative method of setting the detector scaling is to determine that zero on the measurement scale will correspond to the mean (average) of the background level of pollution, and the full scale point set a given amount above the alarm trigger level. The alarm trigger level with this method is set by performing a continual statistical analysis of the background pollution and continuously calculating the level at which a given probability of a signal occurring is acceptably low. Such measurements and calculations are well within the capability of the modern micro-processor, and all that is required is a suitable method. Since the zero and full scale points are set relative to the background pollution levels, this makes the detector inherently more suited to its specific environment. It also makes the movements of the detector's output for zero pollution immaterial since it is constantly defining its own scale. However, since the scale is set according to a history of constantly changing background environment, the method of implementation must place a limit on the duration of the history. For example, the background level readings taken thirty days ago are likely to be considerably less significant than the readings taken thirty minutes ago. This method is basically an on-going 'learning method', and to be successful it must 'learn' logically. There are a number of minor problems with this method but all of them may be successfully overcome by using a micro-processor, in fact the micro-processor becomes the centre of the detector, governing all its procedures and actions.

Stratos™ - HSSD®

Stratos is a relatively scaled detector designed around a modern microprocessor of the same family as that used in Personal Computers. The detector chamber is based on the light scattering principle with a high power semiconductor laser used as the light source. The detector operates on the forward scattering detection principle and is responsive to a wide range of particle sizes. Both dust discrimination and filtration techniques are employed, which makes the system virtually immune to problems caused by dust pollution. The problems encountered in its design, their solutions and the general method of implementation, are listed below.

Description

Histograms and Learning

Readings from the detector chamber(s) are obtained by the micro-processor at the average rate of once per second. A continually updated histogram is generated from the detector(s) output, in which the histogram classes contain the percentage of readings taken between different signal levels. The classes cover all possible readings that may be obtained, and it shows the distribution of signal levels. From the histogram the Standard Deviation and Mean of the distribution are continually calculated. A decaying factor is used in the build up of the histogram. The effect of the decaying factor is to limit the past time period or historical period that is represented by the histogram, it also sets the time which the histogram requires to become truly representative of the distribution. The method for doing this is an AirSense Technology Limited patented process called '**ClassiFire**'. If the ClassiFire factor is set so that the a fully maintained histogram always represents a history of the last ten hours, then it will take ten hours running from the initial set up to build up a representative histogram. Similarly, a factor used to make the histogram represent a ten minute history will require a ten minute period from the initial set up to become fully representative. This is called the Learning Time. Since the ideal should be a seven to eight hour history, this is a problem for initial installation. The Learning Time is also a problem for a change in use of the protected area such as a day-to-night or a night-to-day change. Both of these problems are overcome by maintaining **two** histograms; one with a seven hour learning time (the standard histogram) and the other with a fifteen minute learning time (the fast histogram). On initial power up the alarm is disabled for fifteen minutes while the fast histogram becomes representative of the signal distribution during this period. The histogram so obtained is then used as the standard histogram which then has the decay factor set to cover a multi-hour historical period. Since the historical period is initially only fifteen minutes, the spread of the distribution will be smaller than if the historical period had been several hours and it will take several hours before the distribution shows the true normal spread. The alarm trigger level is set by the spread, in terms of the standard deviation. In order to compensate for the initial FastLearn, the figure for this is increased by a factor that is gradually decreased to unity over the initial period. This gives an artificially slightly high alarm level for the first few hours after installation but needs only fifteen minutes before the detector becomes operative. As the detector has to discern between day and night periods, then the detector takes an initiation period of 24 hours as part of the initial self set up.

The fast histogram is constantly maintained as well as the standard histogram, and is used to compare with it in order to show short term changes or trends. Short term changes will occur with changes in use of the protected area. The problem with changes in use is that, with, say, a seven hour learning period, the detector will have only just learned about one period after a change, when another change occurs and it begins to learn about that. The detector never has time to settle fully to its current use. For all practical purposes, changes of use cover only two different levels of use. These

are usually day time and night time, but for the sake of generality may be referred to as normal time and quiet time.

The problem of maintaining correct alarm trigger levels, in spite of changes in use at occupied and unoccupied times, is solved by having two standard histograms, one for the normal time and the other for the quiet time. When the detector is first installed the approximate times of day for the changes in use are programmed into the micro-processor. Within a 140 minute period of these set day and night start times (70 minutes either side of the actual time set), the processor looks for a change of use by comparing the fast histogram with the present standard histogram. When it is satisfied that it has detected a change within this 140 minute period, it switches to the alternative standard histogram and notes the time it occurs. The alternative standard histogram will contain all the learning from the previous time it was used and consequently needs no time to re-learn it. The time that was originally programmed for the change to occur is compared with the time it actually occurred and if there is a difference, the programmed time is modified to be nearer to the actual time. After some days, the programmed time will be showing the average of the actual times the changes occur and will have the ability to capture the next time a change occurs if it is within an hour of it. This is useful when the clocks change from summer to winter time and activities change according to the clock.

If the comparison of fast histogram and the current standard histogram shows no change of use within an hour of the programmed time, then no switch to the alternative standard histogram is made. This is particularly beneficial if the detector remains in a quiet period for a day or more, such as for a weekend or holiday period, when it will remain with the quiet period histogram for the whole period, as, after all, this type of system is usually used where the optimum in sensitivity is required, and this feature provides a unique level of sensitivity with this **artificial intelligence** sensitivity switching.

Dust

An area that requires special care with the light scattering principle, is the exclusion of abnormal levels of dust from the detector chamber. It is as well to mention that perfect protection against dust particles is an ideal which is unlikely to ever be fully achieved although a method of discriminating between dust particles and smoke particles of the same size is always being sought. It is difficult to obtain, because many dust particles exist in the same size region as smoke particles. As such, the combination of fully monitored and compensated (see below) dust particle separation and dust discrimination used in Stratos provides a very adequate solution.

Dust Compensation

For reasons of power economy and reliability a low power air impeller is required to draw air down the sampling pipes. This leads to the problem that the bore of the sampling pipes must be large in order to obtain the high flow rates required (usually 15 to 25 mm.). High flow rates are required to obtain short transit times between the sampling holes and the detector head. Even with a well designed system these may often be as great as 60 seconds for the furthest holes. The result of fast air flow down large bore pipes is a large volume of air being moved. In order to overcome long term problems, the air which passes through the detector must be filtered to some degree to remove the majority of dust particles. The nature of this filter is critical though, because the matter it filters from the air remains in the filter medium, and as it builds up will increase the efficiency of the filter. This manifests itself by the filter removing progressively smaller and smaller particles the more it becomes loaded with dust, until it is removing the majority of the smoke particles. At this time it will still be passing air satisfactorily and there will, unless compensated for, be no external indication that it is blocking smoke until such time as it fails to detect a fire. One indication that occurs for a *relatively scaled* detector (Stratos) is that the spread and level of readings will gradually decrease.

Stratos will automatically decrease the scale range to compensate for this, but the indication will consist of a marked decrease in the spread of readings. The essence of the problem is not so much the filtering out of dust, which is a straight forward design problem, but the life of the filter, and giving a reliable indication that it is becoming fully loaded. The Stratos detector tackles the problem of dust in four ways;

- Only a small part of the total air sample is taken through the detector head. While a large air flow down the sampling pipes is an advantage, the detector head itself can measure the pollution density equally well for a small sample as a big sample. As a consequence, only the small part of the total air sample flowing through the detector head requires filtering and the filter and detector chamber life is consequently greatly extended. Additionally the frictional resistance that would otherwise restrict air flow is greatly reduced.
- Progressive decreases in the variance and mean of readings in the histogram give a warning that the filter is becoming fully loaded. This progressive reduction in signal spread causes the detector to compensate, thereby automatically maintaining the sensitivity of the Stratos detector.
- The flow rate through the detector sampling chamber is relatively fast and an individual dust particle which does get into the chamber will give just one abnormally high reading from paired light pulses from the laser. Stratos looks for such individual, abnormally high readings and ignores them, only accepting a signal from the lower of the paired pulses. This process is known as Dust Discrimination.
- The filter itself is specifically designed to have a very high capacity for holding dust, which further extends the filter life.

Bargraph Scale

The Stratos scale is represented on the front panel of the Master Stratos detector by a ten segment bargraph display, with the segments numbered 1 to 10. The bargraph illuminates continuously from 1 upwards, to the measured amount of pollution. The alarm trigger level is always pre-set at bargraph level eight. **If** there were a segment numbered zero, then this is where the measured mean level of background signal would be indicated. The main purpose of the bargraph is to indicate the pollution level being accepted by the micro-processor as actual smoke pollution. Unlike other High Sensitivity Smoke Detectors, it does *not* indicate the level of background pollution. The normal indication on the bargraph will be no segments illuminated, although in environments that are normally very smoky, bargraph 1 may be periodically illuminated. The trigger levels are set according to the bargraph. If bargraph number 8 is illuminated an alarm condition is indicated, although not necessarily accepted. If it is not illuminated, an alarm condition is not indicated. A Pre-Alarm trigger level may also be set to a bargraph level below number 8 and an Auxiliary trigger level may be set at any bargraph. Once a trigger level is illuminated the actual function will not be accepted (i.e. signalled to the control unit) until the end of a pre-set delay time. The appropriate segment must be illuminated continuously for the whole duration of the delay time in order to be signalled unless the signal has risen very rapidly to beyond full scale. If this is the case then the time delays may be programmed to be ignored. The bargraph has other uses during the initial fifteen minute installation period and the Test mode. These are detailed in the Stratos-HSSD Installers Handbook.