Lights, Insect Eyes and Fruit

1. Introduction

In the last decade, solid-state LED light sources have reached levels of efficiency, power, durability and reliability necessary to allow them to replace fluorescent tubes and incandescent bulbs in many lighting applications. A secondary advantage of such technology is the ability to modulate the light, at high frequency, to transmit data to remote devices.

During this same period, imaging technology has also progressed in leaps and bounds. The development of CMOS image sensors has led to a very low cost imaging solution, with these sensors being deployed in a large variety of products.

The problem of estimating our location in an outdoors environment can be easily solved using GPS, although the cost is still an issue for some potential applications. The problem is harder in an internal space, such as an office. Although many solutions have been proposed over the years, they all tend to be either inaccurate, or rely on a lot of infrastructure to work. In this working note we propose a location sensor based on a combination of LED lighting and CMOS-based apposition eyes.

2. Office lighting with UIDs

Consider a typical open-plan office, with lights embedded in the ceiling every couple of meters. Now imagine replacing these lights by a Lumileds Luxeon panel. At first glance a user may see little difference. However, suppose now we modulate the light emitted from each panel with a unique signature or ID. If the lights are dimmable, using pulse width modulation, then the additional cost of adding a modulated UID might be negligible. The user would be unaware of this modulation, as it would occur at a rate far above human perception. Nevertheless, a sensor could decode this signal, potentially yielding a low-cost position sensing mechanism for the building. Of course, in practice, there will be many lights within a room. If we only require a crude estimate of position then we could opt for only having one lighting panel in the room emitting a unique ID. A more interesting alternative would be for all panels to have their own ID, and to then develop low-cost sensors that could cope with the potential confusion that results.

3. Apposition eyes

A simplistic approach to sampling the light at different angles is illustrated in Figure 1, where a compound lens is used in conjunction with an LCD shutter to sample the light at different angles of incidence.

A more advanced approach might be based on artificial apposition eyes, for example those developed by J. Duparré et al. from the Fraunhofer-Institut fur Angewandte Optik und Feinmechanik. The basis of their approach is illustrated in

Fig. 1 Position sensing using modulated LEDs

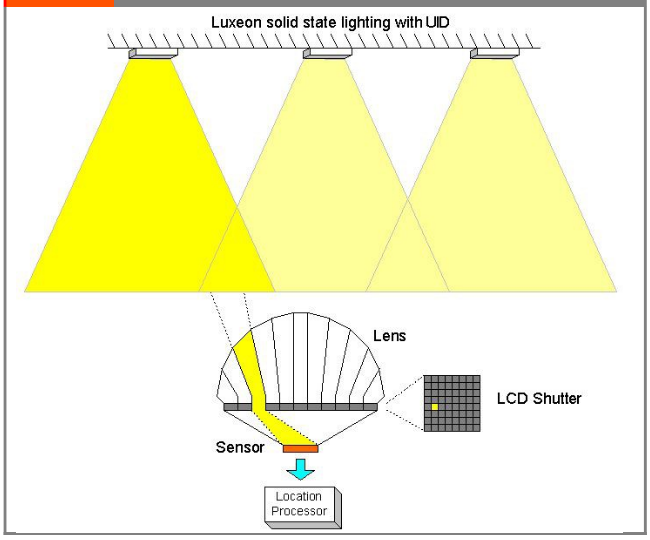

Figure 2. Could Agilent's current crop of CMOS sensors be adapted to use such a technique? And if so, could this form the basis of a low-cost location sensor that exploited UID lighting?

Fig. 2 Artificial apposition eye

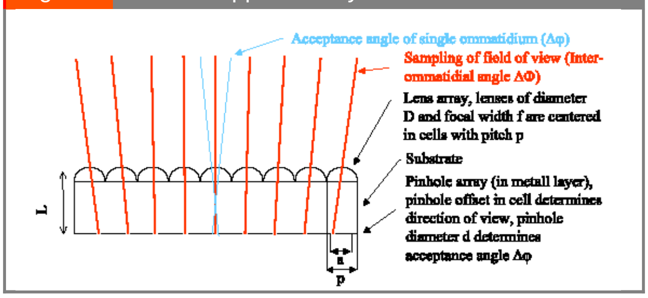

3. Apples and lemons

Consider a 51x51 cell apposition eye pointing at an office ceiling. Figure 3 shows an example of what might be observed by such a sensor. Here we have used the Radiance package to construct a model of the office, using realistic representations of the lights, and then calculated the light intensity in each direction the sensor is looking in. Each square in the display represents 2≅ .

Clearly the first step in estimating our position is to detect the UIDs being emitted by the lights so we can annotate the picture with light identities. We are currently exploring how best to solve this part of the problem in a cheap fashion.

Knowing the identity of the lights, and their angle from the surface normal of the sensor, can we calculate an estimated

Fig. 3 View of ceiling from an apposition eye

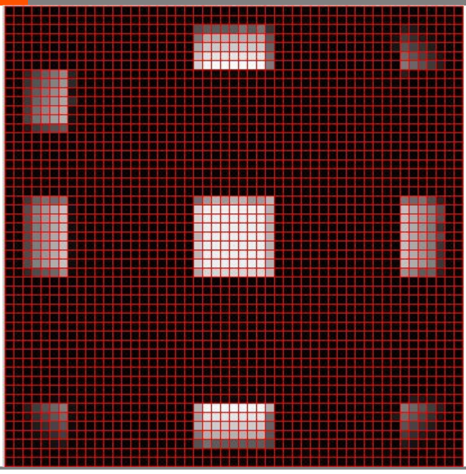

position for the sensor? Given a bearing between two known points, in two dimensions our position will be on a circular arc, as illustrated in Figure 4. The extension to three dimensions involves rotating this arc around the axis formed by the known points. This produces a surface called a spindle torus. In the common case of a sensor being some distance from the ceiling, and where the angle is therefore < 90≅ , the surface forms an "apple", as shown in Figure 5. Where the angle is > 90≅ then the surface forms a lemon shape. With just one pair of observations, and no additional information, we could be anywhere on this surface. If more lights are visible we construct additional apples and lemons, one for each pair in the cross-product, and then attempt to find a point where they all intersect. Unfortunately, there are many potential sources of inaccuracy. For example, the apposition eye will accept a range of angles for each cell, and our estimate of the centre of each light will also be imprecise. We therefore seek to find a position that minimizes the distance to each surface, rather than trying to find a common intersection that is unlikely to exist. For the example shown in Figure 3 the calculations produce an estimated position of (3.30062, 1.000340, 2.70073). The real position is (3.3, 1.0, 2.7), showing that in the best case we can produce very accurate results. However, moving the surface normal away from the vertical, whilst keeping our position fixed, introduces more error, up to 15cm in our example. This is to be expected, as this moves some of the lights towards the edge of the sensor, where it becomes more difficult to estimate the angle to the centre of the light. There are various strategies we could pursue to improve such estimates. For example, weighting the observations based on some estimate of confidence in the reading may produce more accurate results in these cases, as will a more refined algorithm for estimating the cell containing the centre of each light.

The accuracy will also decrease as we reach the edge of the room, where less lights will be visible. In some cases additional constraints may be known. For example, in the case of a sensor attached to a robot we may know the height above the ground, or the surface normal. This will allow us to either create a more accurate estimate of our position, particularly where we only have a few observations.

In summary, we have presented an approach that might have the potential for forming the basis of a low-cost position location system. We are currently exploring how best to encode the UIDs to enable fast and cheap detection. We are also investigation various heuristics to improve the location estimates, as described earlier.

Fig. 4 2D positions given two points and angle

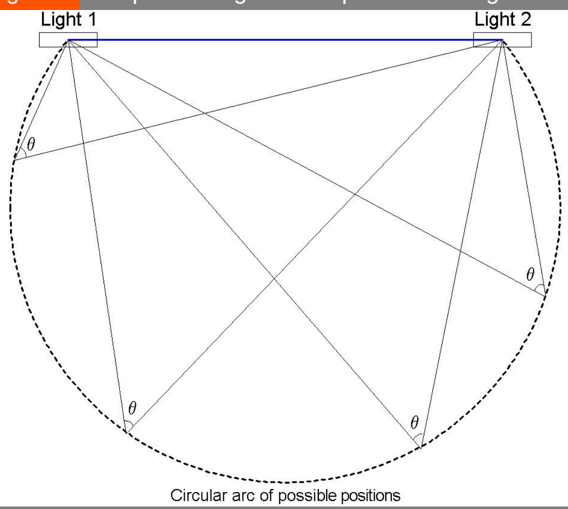

Fig. 5 Apple torus