

(51) International Patent Classification:
H05B 37/02 (2006.01)(21) International Application Number:
PCT/EP2016/080097(22) International Filing Date:
7 December 2016 (07.12.2016)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
15202004.6 22 December 2015 (22.12.2015) EP(71) Applicant: PHILIPS LIGHTING HOLDING B.V.
[NL/NL]; High Tech Campus 45, 5656 AE Eindhoven
(NL).(72) Inventors: CAICEDO FERNANDEZ, David, Ricardo;
c/o High Tech Campus 45, 5656 AE Eindhoven (NL).
PANDHARIPANDE, Ashish, Vijay; c/o High Tech Cam-
pus 45, 5656 AE Eindhoven (NL).(74) Agents: VAN EEUWIJK, Alexander, Henricus, Walter-
us et al.; Philips Lighting B.V. - Intellectual Property High
Tech Campus 45, 5656 AE Eindhoven (NL).(81) Designated States (unless otherwise indicated, for every
kind of national protection available): AE, AG, AL, AM,AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY,
BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM,
DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT,
HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR,
KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME,
MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ,
OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA,
SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM,
TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM,
ZW.(84) Designated States (unless otherwise indicated, for every
kind of regional protection available): ARIPO (BW, GH,
GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ,
TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU,
TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE,
DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU,
LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK,
SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ,
GW, KM, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

— as to applicant's entitlement to apply for and be granted a
patent (Rule 4.17(ii))

Published:

— with international search report (Art. 21(3))

(54) Title: SENSOR SYSTEM.




Fig 9

(57) Abstract: A person detection sensor unit comprises: a communications interface; a sensor device configured to capture over time sensor data from an area covered by the sensor device; a processor configured to implement the following steps: detecting in a first portion of the sensor data captured at a first time a predetermined synchronization code, and measuring the first time based on a clock signal available at the sensor unit; detecting in a second portion of the sensor data captured at a second time, at least one person present in the area, and measuring the second time based on the clock signal; based on said detection at the sensor unit of the at least one person, generating from the second portion of the sensor data presence data for the second time pertaining to the detected at least one person; and outputting via the communications interface the presence data for the second time and associated timing data, which conveys the second time as measured at the sensor unit relative to the first time as measured at the sensor unit. A correspond method and computer program are disclosed, as it a sensor system that includes the sensor unit.

Sensor system

TECHNICAL FIELD

The present invention relates to synchronization of a sensor system.

BACKGROUND

5 A lighting system for illuminating an environment may comprise one or more luminaires, each of which comprises one or more lamps that emit illumination into the environment, plus any associated socket, housing or support. Each lamp may take any suitable form, for example an LED-based lamp comprising one or more LEDs, or a filament bulb, gas discharge lamp, etc.

10 Such luminaires may be inter-connected so as to form a lighting network. For example, in order to control the illumination, a gateway, such as a lighting bridge, may be connected to the network. The gateway can be used to communicate control signals via the network to each of the luminaires, for example from a general-purpose computer device such as a smartphone, tablet or laptop connected to the gateway.

15 The lighting network may have a mesh topology, whereby the luminaires themselves act as relays within the lighting network, relaying control signals between the gateway and other luminaires in the network. Alternatively, the network may have a star topology, whereby luminaires communicate with the gateway “directly” i.e. without relying on other luminaires to relay the control signals (though possibly via other dedicated network components). Generally, the network can have any suitable network topology, e.g. based on a
20 combination of star-like and mesh-like connections. In one example, the lighting network may for example operate in accordance with one of the ZigBee protocols, while the computer device connects to the gateway via another protocol such as Wi-Fi.

 The luminaires or the lighting system may also be equipped with sensor
25 mechanisms. Historically, such sensor mechanisms have been relatively unsophisticated. For example, combinations of timers and motion sensors have been used to selectively active luminaires in response to recently sensed movement in the environment. An example of such a motion sensor is a passive infra-red (“PIR”) motion sensor, which uses infrared radiation emitted from moving bodies to detect their motion.

More modern lighting systems can incorporate sensors into the lighting network, so as to allow the aggregation of sensor data from multiple sensors in the environment. Using suitable sensors, this allows the luminaires to share information on, say, occupancy, activity patterns, changes in temperature or humidity, daylight levels, etc. These sensor signals may be communicated via the lighting network to the gateway, thereby making them available to the (or a) computer device connected to the gateway.

Such sensors have also been used in a lighting system to extract information relating to people in the area covered by the lighting system. For example, people counting techniques have been utilised to generate a count of people in the area based on the aggregation of sensor data from individual image capture devices. The ability to detect a count of people over a particular area may have a number of applications, such as space optimization, planning and maintenance, HVAC control, and data analytics driven marketing. For example, in marketing analysis, people count is needed as one of the input data for analysis. E.g. for space optimization, a count of people in (pseudo) real time may be desired to identify temporal and spatial usage patterns.

SUMMARY

Where sensors in a sensor system are not properly synchronized, this can result in errors in higher-level information derived from their outputs, such errors in a people count. The present invention allows the outputs of the sensors to be synchronized, thereby ensuring that accurate information can be derived from those outputs.

A first aspect of the present invention is directed to a method of synchronizing first and second sensor units of a sensor system, the method comprising implementing by a synchronisation system of the sensor system and external to the sensor units the following steps: receiving from the first sensor unit: at least one first measurement generated at the first sensor unit, and a timestamp of that measurement generated at the first sensor unit based on a first clock signal available thereat; receiving from the second sensor unit: a plurality of second measurements generated at the second sensor unit, and for each of the second measurements a timestamp of that measurement generated at the second sensor unit based on a second clock signal available thereat; comparing the first measurement with each of the plurality of second measurements to identify which of the second measurements has a maximum correlation with the first measurement; determining a timing offset between the first and second clock signals by determining a difference between the timestamp of the first measurement generated at the first sensor unit and the timestamp of the identified second

measurement having the maximum correlation with the first measurement generated at the second sensor unit; receiving from the first sensor unit at least a subsequent first measurement, and a timestamp of the subsequent first measurement generated at the first sensor unit based on the first clock signal; receiving from the second sensor unit a plurality of subsequent second measurements, and a timestamp of each of the subsequent second measurements generated at the second sensor unit based on the second clock signal; and using the determined timing offset to determine which of the subsequent second measurements was performed substantially simultaneously with the subsequent first measurement, by determining that the timestamp of that subsequent second measurement differs from that of the subsequent first measurement by substantially the determined timing offset.

Advantageously, the synchronization of the first aspect is entirely passive, in the sense that it is based entirely on measurements from the sensor units without requiring any special communications between the sensor units and the synchronization system outside of the normal operation of the sensor units. That is, the first aspect can be implemented without any additional signalling overhead to the sensors within the sensor system. In this respect, it is noted that the synchronization of the subsequent measurements is not achieved by adjusting the first or second sensors units (which would require communication between the external synchronization system, and thus additional signalling overhead), and in particular is not achieved by adjusting how they apply their respective timestamps – rather, the first and second units continue to output timestamps that are “inconsistent” (in the sense that the timing offset persists, such that the timestamps of the first measurement differs from that of the substantially simultaneous second measurement by substantially the timing offset), and this inconsistency is accounted for externally at the external synchronization system based on the earlier determination of the timing offset at the external synchronization system.

A second aspect of the present invention is directed to a person detection sensor unit comprising: a communications interface; a sensor device configured to capture over time sensor data from an area covered by the sensor device; a processor configured to implement the following steps: detecting in a first portion of the sensor data captured at a first time a predetermined synchronization code, and measuring the first time based on a clock signal available at the sensor unit; detecting in a second portion of the sensor data captured at a second time, at least one person present in the area, and measuring the second time based on the clock signal; based on said detection at the sensor unit of the at least one person, generating from the second portion of the sensor data presence data pertaining to the detected

at least one person; and outputting via the communications interface the presence data for the second time and associated timing data, which conveys the second time as measured at the sensor unit relative to the first time as measured at the sensor unit.

The second aspect can provide highly accurate synchronization and, when sensor units according to the second aspect are connected in a sensor network, allows time differences due to the transport of the message over the sensor network to be accounted for. Moreover, although a dedicated synchronization code is used, it is communicated in manner that is detectable by the sensor device of the sensor units, and thus does not create any additional signalling overhead within the sensor system.

In embodiments, the first and second aspect can be combined such that synchronization of timestamps is performed locally at the sensor units (according to the second aspect) and measurements are additionally synchronized externally (according to the first aspect).

In embodiments of the first aspect, a plurality of first measurements may be received from the first sensor unit at the synchronization system, and for each of the first measurements a timestamp of that measurement generated at the first sensor unit based on the first clock signal.

A correlation may be determined for each of a plurality of time difference values, by applying a correlation function to the first and second measurements for that time difference value, the determined time offset between the clock signals corresponding to the difference value for which the determined correlation is maximised.

Note that the term “maximized” in relation to a correlation means most correlated. Depending on how the correlation function is defined this may (for example) correspond to a maximum value of the correlation function, but in other cases may correspond to a minimum value of the correlation function depending on how it is defined.

The first measurement and the second measurements may pertain to an area of overlapping sensor coverage between the first and second sensor units.

Each of the first and second measurements may comprise a respective measured location of a person detected in an area covered by the first and second sensor units respectively.

The locations measured by both sensor units may be in the area of overlapping sensor coverage.

Each of the first and second measurements may comprise a respective light level measured over all or part of an area covered by the first and second sensor units respectively.

The light levels may be measured by both sensor units across the area of overlapping sensor coverage.

The first measurement may pertain to only a part of the area covered by the first sensor unit, and each of the second measurements pertains to only a part of the area covered by the second sensor unit.

The first measurement may be compared with each of the plurality of second measurements by multiplying the first measurement with that second measurement.

A plurality of first measurements may be received from the first sensor unit, each with a timestamp of that measurement generated at the first sensor unit based on the first clock signal, and the comparing step may comprise determining a correlation for each of a plurality of difference values by: for each of the first measurements, multiplying that first measurement with the second measurement whose timestamp corresponds to the timestamp of that second measurement offset by that difference value, the determined time offset between the clock signals corresponding to the difference value for which the determined correlation is maximised.

The comparing step may comprise determining a correlation for each of a plurality of difference values by: multiplying the first measurement with the second measurement whose timestamp corresponds to the timestamp of the first measurement offset by that difference value, the determined time offset between the clock signals corresponding to the difference value for which the determined correlation is maximised.

For example, a plurality of first measurements may be received from the first sensor unit, each with a timestamp of that measurement generated at the first sensor unit based on the first clock signal, wherein the correlation for each of the candidate timing offsets may be determined by: for each of the second measurements, comparing that second measurement with the first measurement whose timestamp corresponds to the timestamp of that second measurement offset by that difference value.

A plurality of first measurements may be received from the first sensor unit at the synchronization system, each with a timestamp of that measurement generated at the first sensor unit based on the first clock signal, wherein the comparing step may comprise determining a correlation for each of a plurality of difference values by: determining a sum of differences between each of the first measurement and the second measurement whose

timestamp corresponds to the timestamp of the first measurement offset by that difference value.

For example the sum of differences may be a sum of absolute or squared differences.

5 The method may comprise using the determined timing offset to account for inconsistencies in timestamps pertaining to subsequent measurements by the first and second sensor units when the sensor system is in use.

10 The method may comprise estimating a people count for a desired area, which comprises a total area covered by the first and second sensor units, based on the subsequent measurements, their timestamps and the determined timing offset.

Each of the sensor units may be a person detection sensor unit configured to detect locally at that sensor unit any person or people present in an area covered by that sensor unit, and to output to the synchronization system presence data pertaining to the person or people detected locally at that sensor unit.

15 According to another aspect of the present invention a sensor system comprises: a first sensor unit configured to generate: at least one first measurement, at least a subsequent first measurement, and a timestamp of each of those first measurements generated based on a first clock signal available at the first sensor unit; a second sensor unit (3b) configured to generate: a plurality of second measurements a plurality of subsequent second
20 measurements, and a timestamp of each of those second measurements based on a second clock signal available at the second sensor unit; a synchronisation system external to and connected to the sensor units and configured to: compare the first measurement with each of the plurality of second measurements to identify which of the second measurements has a maximum correlation with the first measurement; determine a timing offset between the first
25 and second clock signals by determining a difference between the timestamp of the first measurement and the timestamp of the identified second measurement having the maximum correlation with the first measurement; and use the determined timing offset to determine which of the subsequent second measurements was performed substantially simultaneously with the subsequent first measurement, by determining that the timestamp of that subsequent
30 second measurement differs from that of the subsequent first measurement by substantially the determined timing offset.

In embodiments of the second aspect, the presence data for the second time may comprise:

a presence count indicating a number of people detected by the sensor unit in the covered area at the second time; and/or

a presence score indicating a likelihood that there is a person or people in the covered area at the second time; and/or

5 a number of person location identifiers, each identifying a location of a person detected in the covered area at the second time.

Each person location identifier may be a two or three dimensional location vector.

10 The sensor device may be a photo sensor device configured to sense visible and/or non-visible radiation.

For example, the photo sensor device is an image capture device and the sensor data is image data, the first portion of the image data being one or more first images captured by the image capture device of the area and the second portion being one or more second image captured by the image capture device of the area.

15 The second images may not be outputted by the sensor unit.

The processor may be configured to detect the synchronization code embedded in the radiation as amplitude and/or phase modulations.

20 The processor may be configured to determine a difference between the first time as measured at the sensor unit and the second time as measured at the sensor unit, wherein the associated timing data may comprise the determined difference.

Alternatively the associated timing data may comprise a first timestamp of the first time and a second stamp of the second time, and thereby conveys the first time relative to the second time. That is the first and second timestamps may be outputted separately, at the same or at different times.

25 In another aspect, a sensor system comprises: a plurality of person detection sensors units; a transmitting unit configured to emit at a first time a synchronization code detectable by the sensor units and embodying the synchronization code; wherein each of the sensor units is configured according to any embodiment of the second aspect.

30 The sensor system may further comprise a people counting apparatus; wherein each of the sensor units may be configured to output respective presence data for the second time and associated timing data which conveys, to the people counting apparatus the second time relative to the first time; and wherein the people counting apparatus may be configured to use the respective presence data to estimate a people count for a total area covered by the sensor units.

The transmitting unit may be a luminaire configured to emit illumination at the first time in which the synchronization code is embedded, the sensor device of each sensor unit being a photo sensor device.

The synchronization code may be embedded using visible light communication, whereby it is imperceptible to a human eye.

Another aspect of the present invention is directed to a people detection method implemented by a person detection sensor unit of a sensor system, the method comprising: capturing over time sensor data from an area covered by the sensor unit; detecting in a first portion of the sensor data captured at a first time a predetermined synchronization code; measuring the first time based on a clock signal available at the sensor unit; detecting in a second portion of the sensor data captured at a second time at least one person present in the area; measuring the second time based on the clock signal; based on said detection at the sensor unit of the at least one person, generating from the second portion of the sensor data presence data for the second time pertaining to the detected at least one person; and outputting, to a processing apparatus external to the sensor unit, the presence data for the second time and associated timing data, which conveys to the external processing apparatus the second time as measured at the sensor unit relative to the first time as measured at the sensor unit.

Any embodiments of any of the above aspects may implement features of any of the other aspects, or embodiments therefor. For example, embodiments of the first aspect may implement any feature of the second aspect or any embodiment thereof and *vice versa*.

According to a yet further aspect of the present invention, a computer program code comprises executable code stored on a computer readable storage medium and configured when executed to implement any of the methods, sensor system functionality or sensor unit functionality disclosed herein.

BRIEF DESCRIPTION OF THE DRAWINGS

For a better understanding of the present invention, and to show how embodiments of the same may be carried into effect, reference is made to the following figures, in which:

Figure 1 is a schematic illustration of a lighting system,

Figure 2 is a schematic block diagram of a sensor unit,

Figure 2A is a schematic block diagram of a luminaire with embedded sensor unit,

Figure 2B is a schematic block diagram of a luminaire,

Figure 3 is a perspective view of a pair of adjacent luminaires,

Figure 3A is a plan view of part of a lighting system,

Figure 4 is a schematic block diagram of a central processing apparatus for
5 operating a lighting system,

Figure 4A is a schematic block diagram illustrating an exemplary control
architecture of a lighting system,

Figures 5 illustrates how local image processors cooperate with a central
processing apparatus to provide a people counting function,

10 Figure 6 illustrates how a correctly synchronized sensor system may be used
to implement people counting,

Figure 6A is a block diagram showing how sensors may communicate
timestamped measurements to a central processing apparatus;

Figure 7 is a flowchart for a sensor synchronization method;

15 Figure 8 shows a first example of how sensors may operate in accordance with
the synchronization method of figure 7;

Figure 9 shows a second example of how sensors may operate in accordance
with the synchronization method of figure 7;

Figure 10 shows a flow chart for another synchronization method and

20 Figure 11 shows a third example of how sensors may operate according to the
other synchronization method of figure 10.

DETAILED DESCRIPTION OF EMBODIMENTS

“Vision sensors” comprising visible light cameras are useful for people
25 counting. Consider a system wherein for privacy reasons each vision sensor does not provide
entire images to a central processing device when operating in real time, but only presence
decisions based on performing the image recognition locally at the sensor (where the
presence decision could be a soft or hard decision). A vision sensor configured in this manner
is an example of a person detection sensor unit. Alternative person detection sensor units, for
30 example which do not operate on captured images but collect and analyze other type(s) of
sensor data are also within the scope of the present disclosure. Each vision sensor covers (i.e.
provides sensor coverage of) a respective area, defined by its field of view from which it is
able to capture sensor data.

The vision sensors may have overlapping field-of-views (FoVs) and sensing areas. In a vision sensor system, it may happen that one or more issues arise, which in turn may lead to application errors, e.g. errors in an estimated number of people over a particular space, at a higher system level that are not immediately noticeable. The vision sensors are
5 connected, so as to form a sensor network having any suitable topology, for example of the kinds described above, and which operated according to any suitable protocol (e.g. ZigBee, Wi-Fi, Ethernet etc.).

Examples are described below in the context of a sensor system with multiple vision sensors in communication with an external processing apparatus, such as a people
10 counting apparatus to offer data-enabled applications based on people counting. In order to perform further processing (for example fusion, analytic etc.), reasonable synchronization of measurements generated by the vision sensors is required.

Each of the vision sensors has available to it a respective clock signal, which it uses to apply timestamps to its outputted measurements. However, due to differences in
15 respective clock signals available at the vision sensors, a shift in the local time of each vision sensor can occur, leading to inconsistencies between the timestamps generated by different sensors. Over long periods of time, this shift can become significant and lead to increase errors in an estimated people count over the total area covered by the vision sensors as a whole.

20 Traditional methods of sensor synchronization are based on exchanging messages via the sensor network itself and hence result in additional network load within the sensor network, which is a problem particularly (though not exclusively) in wireless networks.

In the described examples, the outputs of the vision sensors are synchronized
25 without requiring any additional transmission overhead for synchronization within the sensor network, i.e. without any additional network traffic via the sensor network.

The vision sensors only send a limited amount of information (e.g. location, illumination changes) to an external processing apparatus via a bandwidth-limited communication channel. The information received from the vision sensors at the central
30 processing apparatus is synchronized within a reasonable value by exploiting embedded patterns in the same information; this is achieved without increasing bandwidth requirements and ensuring that the communication channel is used for data communication only.

In a first and a second of the described examples, each vision sensor communicates to the external processing apparatus, at each of a plurality of times, the following:

a measurement generated by the vision sensor at that time and pertaining to the area covered by the vision sensor;

an associated timestamp denoting that time, as measured at the vision sensor based on a respective clock signal available as that vision sensor; and

a vision sensor identifier ("ID") unique to the vision sensor within the system.

The external processing apparatus determined, by correlating the reported measurements of adjacent sensors, and compensate for any time shift within a given time window. That is, to determine account for any timing offset(s) between the respective clock signals available to the different vision sensors.

In the first example, for each of the plurality of times, the measurement identifies a location of any person whose presence has been detected in the area covered by the sensor at that time. The external processing apparatus uses knowledge of any sensing region overlap of adjacent sensors to determine the timing offset(s) by correlating the locations reported by adjacent sensors.

In the second example, the measurement comprises one or more light levels measured over all or of the covered area, or part of the area covered by the sensor. For example, at least two light levels may be measured simultaneously at the sensor unit each over a respective predefined region of the area covered by the sensor. Each of the predefined regions is large compared with a pixel size of the vision sensor, such that the measurements in this second example are effectively a heavily quantized image, having a significantly lower resolution. The external processing apparatus uses knowledge of any sensing region overlap of adjacent sensors to determine the timing offset(s) by correlating the light levels reported by adjacent sensors.

In a third of the describe examples, the vision sensor are arranged such that there is a luminaire within the field of view of at least two adjacent vision sensors. During a synchronization phase, the luminaire is dimmed according to a predetermined sequence a sequence, so as to emit from the luminaire a visible light synchronization code. Each vision sensor is configured to recognize the dimming sequence, thereby allowing it to detect a starting time of the dimming sequence from the luminaire. Each vision sensor communicates to the external processing apparatus, at each of a plurality of times:

a measurement for that time, which is presence data for that time, which conveys information about the presence of any detected person or people in the area covered by the sensor at that time;

an associated timestamp denoting that time relative to the starting time of the dimming sequence, such that all vision sensors output timestamps in the same temporal frame of reference; and

its vision sensor ID.

The external processing apparatus arranges the measurements received from the sensors in chronological time with respect to the start of the dimming sequences for each vision sensor.

The first and second examples pertain to external synchronization, performed by the central processing apparatus. The third example pertains to internal synchronization, performed within each of the vision sensors. The internal synchronization techniques of the third example can be combined with the external synchronization techniques of the first or second examples, such that both types of synchronization are performed in parallel, which can allow greater accuracy and robustness.

The presence data for a given time may for example comprise a people counting metric for that time, such as a people count for the area covered by the sensor or one or more probability scores for estimating a people count for that area.

Each measurement generated by each vision sensor is based on a respective portion of the sensor data, such as image data, captured by that vision sensor. Any person in the area covered by that vision sensor is not identifiable from the measurement itself, even if they are identifiable in that portion of sensor data (e.g. image data) from which that measurement is generated. The measurement has a lower information content, and thus a smaller size than, than the portion of sensor data from which it is generated, which reduces signaling overhead within the sensor network.

In the first and second examples, timestamps outputted by different vision sensors may be expressed in different temporal frames of references due to the clock signals available at the different vision sensors being out of sync – at the very least, the system does not assume that the timestamps outputted by the vision sensor are in the same frame of reference, and takes steps to identify and account for any time-base or time-stamp correction externally to the vision sensors, using measurements outputted by the vision sensors as part

of their normal function within the system so that no additional signaling overhead is required.

In the third example, a synchronization code is embedded in visible light that is detectable by the vision sensors as part of their normal function, or more generally in a manner that is detectable by the vision sensors as part of their normal sensor function. That is, such that the synchronization code is received by the vision sensors in sensor data collected by them, and not via the sensor network. This synchronization code is used by the vision sensors to correct their timestamps locally, such that the timestamp outputted by the different vision sensors are all in substantially the same temporal frame of reference. Again, this does not require any additional signaling overhead, as the synchronization code is not sent via the sensor network.

The respective clock signal available to each vision sensor is preferably a locally generated clock signal generated by a respective local clock of that sensor, for example a crystal oscillator (e.g. quartz) clock, as this requires no signalling overhead. When using a stable clock the sensors can report the number of “ticks” (or a derivative thereof) since the last synchronization code. The advantage of using a stable clock is that it has limited sensitivity to drift and thus the frequency of re-synchronization can be limited.

When clocks are more susceptible to synchronization drift (such that their timing offsets are variable) the present techniques may be performed repeatedly at suitable intervals to account for the variable offsets. Alternatively the vision sensor may receive its local clock signal from the sensor network, this could be an existing clock, e.g. the form of the clock of the communication network (e.g. the TSF in 802.11) or a dedicated clock distributed purposefully though this is less preferred due to the signalling overhead it requires. Clock signals received via the network are still prone to synchronization errors, for example due to different clock signal transmission times and/or changing network conditions.

Figure 1 illustrates an exemplary lighting system 1 in which the technique disclosed herein may be employed. The system 1 comprises a plurality of luminaires 4 installed in an environment 2, arranged to emit illumination in order to illuminate that environment 2. In embodiments, the system may further comprise a gateway 10 to which each of the luminaires 4 is connected via a first wired or wireless networking technology such as ZigBee. The gateway 10, sometimes referred to as a lighting bridge, connects to a computing apparatus 20 (which may or may not be physically present in the environment 2) via a second wired or wireless networking technology such as Wi-Fi or Ethernet. The computing apparatus 20 may for example take the form of a server (comprising one or more

server units at one or more sites), or a user terminal such as a smartphone, tablet, laptop or desktop computer, or a combination of any such device. It is able to control the luminaires 4 by sending control commands to the luminaires 4 via the gateway 10, and/or is able to receive status reports from the luminaires 4 via the gateway 10. Alternatively in embodiments the gateway 10 may not be required and the computing apparatus 20 and luminaires 4 may be equipped with the same wired or wireless networking technology, by which they may be connected directly into the same network in order for the computing apparatus 20 to control the luminaires 4 and/or receive the status reports from the luminaires 4.

In the illustrated example, the environment 2 is an indoor space within a building, such as one or more rooms and/or corridors (or part thereof). The luminaires 4 are ceiling-mounted, so as to be able to illuminate a surface below them (e.g. the ground or floor, or a work surface). They are arranged in a grid along two mutually perpendicular directions in the plane of the ceiling, so as to form two substantially parallel rows of luminaires 4, each row being formed by multiple luminaires 4. The rows have an approximately equal spacing, as do the individual luminaires 4 within each row. However it will be appreciated that this is not the only possible arrangement. E.g. in other arrangements one or more of the luminaires 4 could be mounted on the wall, or embedded in the floor or items of furniture; and/or the luminaires 4 need not be arranged in a regular grid; and/or the environment 2 may comprise an outdoor space such as a garden or park, or a partially-covered space such as a stadium or gazebo (or part thereof), or a combination of such spaces.

Multiple people 8 may occupy the environment, standing on the floor below the luminaires 4. The environment 2 is also installed with one or more “vision sensor” units 3, each of which being a visible-light based imaging unit for detecting the presence of people 8 in the environment based on the images it captures. E.g. these may also be mounted on the ceiling in a regular pattern amongst the luminaires 4, and may be arranged to face downwards towards the illuminated surface beneath (e.g. the ground or floor, or a work surface). Alternatively the sensor units 3 may be mounted in other places such as the wall, facing in other directions than downwards; and/or they need not be installed in a regular pattern.

The luminaires 4 have known identifiers (“IDs”), unique within the system in question, and are installed at known locations. The vision sensor units 3 also have known IDs, and are also installed at (at least what are initially believed to be) known locations in the environment 2. The sensor units 3 are not necessarily co-located with the luminaires 4. The locations of the luminaires 4 are determined during a commissioning phase of the luminaires 4, i.e. before the luminaires 4 and sensor units 3 are actually put into operation for their

purpose of illuminating the environment 2 and detecting presence of people 8 respectively. Typically commissioning is performed at the time of installation or shortly afterwards.

During commissioning of the luminaires 4, a commissioning technician determines the location of each of the luminaires 4, either manually or using automated means such as GPS or another such satellite based positioning system. This may be the location on any suitable reference frame, e.g. coordinates on a floorplan, map of the area, or global coordinates. By whatever means and in whatever terms determined, the commissioning technician then records the location of each luminaire 4 in a commissioning database 21 mapped to its respective luminaire ID. The commissioning technician also performs a similar

commissioning process for the sensors 3 during a commissioning phase of the sensor units 3, i.e. prior to the actual operational phase, before the vision sensors 3 are actually put into operation for their purpose of detecting presence of people 8. The sensor commissioning phase comprises storing the (believed) location of each in the commissioning database 21 mapped to its respective sensor ID.

Note that the commissioning database 21 could be anything from a large database down to a small look-up table. It could be implemented on a single device or multiple devices (e.g. computing apparatus 20 represents a distributed server, or a combination of server and user terminal). E.g. the table mapping the vision sensor locations to the vision sensor IDs could be implemented separately from the table mapping the luminaire locations to the luminaire IDs. Of course it will also be appreciated that the commissioning could be performed over different occasions, and/or by more than one technician. E.g. the commissioning of the vision sensors 3 could be performed by a different commissioning technician on a later occasion than the commissioning of the luminaires 4.

Knowing the locations of the luminaires and the sensors 3 allows the position of the luminaires 4 relative to the sensor units 3 to be known. According to the present disclosure, this is advantageously exploited in order to check for commissioning errors or other problems with the sensor units 3. In fact, for the purposes of the present disclosure, only the relative locations of the luminaires 4 relative to the sensor units 3 need be known (e.g. stored in terms of a vector in the commissioning database 21). However, for other purposes it may alternatively or additionally be desired to store the absolute locations of the sensor units 3 and/or luminaires 4 in the commissioning database 21, such as to enable the absolute location of a person (e.g. on a floorplan or map) to be determined based on the sensor units 3, or to allow the luminaires 4 to be used as a reference for indoor navigation, etc.

In other embodiments, each of one, some or all of the sensor units 3 may be incorporated into the housing of a respective one of the luminaires 4. In this case the locations of the luminaires 4 are known relative to the sensor units 3 implicitly, i.e. can be assumed to be co-located. For such sensor units 3 the commissioning database 21 is not necessarily required for the purpose of checking the sensor units 3, though may optionally be included anyway for other purposes (e.g. again to enable detection of the location of a person 8, or for indoor navigation).

The use of the luminaires 4 to check for problems with the sensor units 3 will be discussed in more detail shortly. However, first the context of a people counting system will be described with reference to Figures 2 to 5.

Figure 2 shows a block diagram of a vision sensor unit 3, representing the individual configuration of each sensor unit 3 in the lighting system 1. The sensor unit 3 comprises: an image sensor 6 in the form of a visible light camera, a local processing module 11, a network interface 7, a local memory 13 connected to the local processing module 11, and a local clock 18 connected to provide a local clock signal 19 to the local processing module 11. The camera 6 is able to detect radiation from the luminaires 4 when illuminating the environment, and is preferably a visible light camera. However, the use of a thermal camera is not excluded. The local processing module 11 is formed of one or more processing units, e.g. CPUs, GPUs etc.; and the local memory 13 is formed of one or more memory units, such as one or more volatile or non-volatile memory units, e.g. one or more RAMs, EEPROMs ("flash" memory), magnetic memory units (such as a hard disk), or optical memory units. By whatever means implemented, the local memory 13 stores code 12a arranged to run (e.g. execute or be interpreted) on the local processing module 11, the processing module 11 thereby being configured to perform operations of the sensor unit 3 in accordance with the following disclosure. Alternatively the processing module 11 could be implemented in dedicated hardware circuitry, or configurable or reconfigurable hardware circuitry such as a PGA or FPGA.

By whatever means implemented, the local processing module 11 is operatively coupled to its respective camera 6 in order to receive images captured by the camera 6, and is also operatively coupled to the network interface 7 in order to be able to communicate with the processing apparatus 20. The processing apparatus 20 is external to each of the sensor units 3 and luminaires 4, but arranged to be able to communicate with the sensor units via the respective interfaces 7, and to communicate with the luminaires 4 via a similar interface in each luminaire 4 (not shown).

The local clock signal 19 is a periodic, regular (i.e. having a fixed or approximately fixed period) signal, which the processing module 11 can use to generate a timestamp of an event denoting a current, local time, i.e. measured locally at the sensor unit 3, of the event.

5 The timestamp can have any suitable format, and the term “timestamp” herein generally refers to any data that conveys a time of an event in any temporal frame of reference, generated based on a clock signal. In one of the simplest cases, a timestamp may be counter value e.g. expressing the time as a single integer value (or a set of integer values). As another example, a timestamp may express a time in any combination of hours, minutes,
10 seconds, ms etc., as appropriate to the individual circumstances; or more generally as a floating point or set of floating point values. A timestamp may express a time to any degree of accuracy and precision that is appropriate to the individual circumstances.

 The local clock 18 comprises a crystal oscillator crystal clock, and the clock signal 19 is derived from by applying a current to a crystal oscillator e.g. quartz crystal of the
15 clock 18. The local clock signal 19 denotes a current time relative to a local reference time, for example a current time expressed as an integer count (where the reference time is e.g. a count of zero). Alternatively the sensor unit 3 may derive its local clock signal 19 by some other means, for example from based on a locally available AC (alternating current) e.g. from a power supply that is powering the sensor, or it clock signal may be a locally received
20 version of a clock signal broadcast through the sensor network (though this is less preferred, due to the additional signalling overhead it requires).

 In any event, the local clock signal 19 expresses the current time relative to the local reference time of the sensor unit 3, and has a frequency at which the current time is updated.

25 The local clock 18 and processing module 11 are shown as separate components for the sake of illustration. However, part of the functionality of the local clock 18 may be implemented by the local processing module 11 itself, for example the local clock may provide a periodic input to the processing module 11 from which the local processing module generates the clock signal 19 itself, i.e. such that the local processing module 11
30 computes the current time relative to the local reference time.

 Without synchronization, the local clock signals 19 available at any two different sensors 3, by whatever means they are generated, may have a timing offset (i.e. be out of sync), for example, because they are based on different reference times – causing in a substantially constant time offset between the clock signals – and/or because they have

slightly different frequencies – causing a time offset between the clock signals that increases over time.

Figure 2B shows an example of a luminaire 4 in embodiments where the luminaires 4 are separate to the sensor units 3. Here, each luminaire 4 may comprise one or more lamps 5, a respective interface 7', a local memory 13' and a local processing module 11'. The local processing module 11' is operatively coupled to the lamp(s) and the interface 7'. Each lamp 5 may comprise an LED-based lamp (comprising one or more LEDs), a filament bulb, a gas-discharge lamp or any other type of light source. The memory 13' comprises one or more memory units and the processing module 11' comprising one or more processing units. The local memory 13' stores code 12b arranged to run (e.g. execute or be interpreted) on the local processing module 11', the processing module 11' thereby being configured to perform operations of a luminaire 4 in accordance with the present disclosure. Alternatively the processing module 11' of the luminaire 4 could be implemented in dedicated hardware circuitry, or configurable or reconfigurable hardware circuitry such as a PGA or FPGA.

In general each of the above-mentioned interfaces 7, 7' could be a wired or wireless interface, but is preferably wireless. For example in embodiments the interface 7 of each of the sensor units 3, and the interface 7' of each of the luminaires 4, may be a ZigBee interface arranged to connect to the gateway 10 using a first wireless networking protocol such as one of the ZigBee standards, e.g. ZigBee Light Link; while the processing apparatus 20 (e.g. a server, or a desktop computer, laptop, tablet or smartphone running a suitable application) connects to the gateway 10 via a second wireless networking protocol such as Wi-Fi or Bluetooth. The gateway 10 then converts between the protocols to allow the external processing apparatus 20 to communicate in one or both directions with the sensor units 3 and luminaires 4. Alternatively, the interface 7 in each of the sensor units 3, and the interface 7' in each of the luminaires 4, may comprise an interface of a type (e.g. Wi-Fi or Bluetooth) directly compatible with that of the external processing apparatus 20, thus allowing the communication to occur directly between the processing apparatus 20 and the sensor units 3 and luminaires 4 without the need for a gateway 10. Generally the network can have any suitable network topology, for example a mesh topology, star topology or any other suitable topology that allows signals to be transmitted and received between each luminaire 4 and the gateway 10 and/or processing apparatus 20.

Whatever the network topology, the external processing apparatus 20 is configured to send control commands to the sensor units 3 and luminaires 4 and to receive

information back from the sensor units 3 and luminaires 4, via the relevant interfaces 7, 7'. This includes receiving soft or hard presence decisions from the sensor units 3, and in some cases receiving measured light levels (as in the second example below). The various communications disclosed herein between components 3, 4, 20 may be implemented by any of the above-described means or others, and for conciseness will not be repeated each time.

Figure 2A shows a variant of the arrangement shown in Figures 1 and 2, wherein the sensor unit 3 is integrated into the same housing as one of the luminaires 4, and therefore the sensor unit 3 is substantially collocated with the respective luminaire 4. In this case, the combined luminaire and sensor 3, 4 unit further comprises (in addition to the components described above in relation to Figure 2) at least one lamp 5 such as an LED-based lamp (comprising one or more LEDs), gas-discharge lamp or filament bulb. The communication with the combined sensor unit and luminaire 3, 4 may both implemented via a shared interface 7 of the unit, and/or any control, processing or reporting associated with the sensing and or luminaire functionality may be implemented by a shared local processing module 11. Alternatively separate interface 7' and/or separate local processing module 11' could be provided for each of the sensor and luminaire functions, but in the same housing.

By way of illustration, embodiments below may be described in terms of separate sensor units 3 and luminaries 4, as shown in Figures 1, 2 and 2B, but it will be appreciated that the various teachings disclosed herein can also apply in relation to the integrated arrangement of Figure 2A.

The local processor 11' of the luminaire 4 (or the local processor 11 of the combined unit 3, 4) is connected to the lamp(s) 5, to allow local lighting control code 12b executed on the local processor 11' (or 11) to control the dimming level of the illumination emitted by the lamp(s) 5, and or to switch the emitted illumination on and off. Other illumination characteristic(s) such as colour may also be controllable. Where the luminaire 4 comprises multiple lamps 5, these may be individually controllable by the local processor 11' (or 11), at least to some extent. For example, different coloured lamps 5 or elements of a lamp 5 may be provided, so that the overall colour balance can be controlled by separately controlling their individual illumination levels.

The local controller 11' of the luminaire 4 may be configured to control one or more such properties of the emitted illumination based on lighting control commands received via the interface 7' from the external processing apparatus 20. E.g. the processing apparatus 20 may comprise a server arranged to receive presence metrics from the sensor units 3 indicative of where people are present in the environment 2, and make decisions as to

which luminaires 4 to turn on and off, or which to dim up and down and to what extent, based on an overview of the presence detected by the different sensor units 3. And/or, the processing apparatus 20 may comprise a user terminal such as a smartphone, tablet or laptop running a lighting control application (or “app”), through which the user can select a desired adjustment to the emitted illumination, or select a desired lighting effect or scene to be created using the illumination. In this case the application sends lighting control commands to the relevant luminaires 4 to enact the desired adjustment or effect. In further alternative or additional arrangements, the local controller 11' of the luminaire 4 may be configured to control any one or more of the above properties of the illumination based on signals received from one or more other sources, such as one or more of the sensor units 3. E.g. if a sensor unit 3 detects occupancy then it may send a signal to a neighbouring luminaire 4 to trigger that luminaire to turn on or dim up.

In each sensor unit 3 (or combined unit 3, 4) the respective image sensor 6 is connected to supply, to its local processor 11, raw image data captured by the image sensor 6, to which a local person detection algorithm is applied by local image processing code 12a executed on the local processor 11. The local person detection algorithm can operate in a number of ways based on any suitable image recognition techniques (e.g. facial recognition and/or body recognition). Based on this, the local person detection algorithm generates one or more “presence metrics” indicative of whether a person 8 is detected to be present in a still image or moving image (video) captured by the image sensor 6, and or how many people 8 are detected to be so present. For example the one or more presence metrics may comprise: a hard indication of whether or not a person 8 is detected to be present in the image (yes/no), a soft indication of whether or not a person 8 is detected to be present in the image (an indication of a degree of certainty such as a percentage), or a momentary count of people 8 simultaneously present in the image, a count of the number of people appearing in the image over a certain window of time, and/or a rate at which people appear in the image. The code 12a running on the local processing module 11 reports this information to the external processing apparatus 20, for use in determining a person count centrally.

Note, detecting whether a person appears in an image may comprise detecting whether a whole person appears in the image, or detecting whether at least a part of a person appears in the image, or detecting whether at least a specific part or part of a person appears in the image. The detection could also comprise whether a specific person appears in the image, or detecting whether a specific category of person appears in the image, or detecting whether any person appears in the image.

Figure 3 shows a perspective view of a first and a second of the sensor units 3a, 3b, as described above. The first and second sensor units 3a capture images from a respective sensor area 30a, 30b, which experience light from one or more of the luminaires 4a, 4b. By way of example, each sensor unit 3a, 3b may be associated with or incorporated into a different respective one of the luminaires 4a, 4b adjacent one another in a grid, or each sensor unit 3 could be associated with a different respective group of the luminaires 4 (e.g. placed at the centre of the group).

The respective lamp 5 of each of the luminaires 4a, 4b is arranged to emit illumination towards a surface 29 (e.g. the floor, or a workspace plane such as a desk), thereby illuminating the surface 29 below the luminaires 4. As well as illuminating the environment 2, the illumination provided by the luminaires 4 renders the people 8 detectable by the sensor units 3.

The respective image sensor 6 of each sensor unit 3a, 4b has a limited field of view. The field of view defines a volume of space, marked by dotted lines in Figure 3, within which visible structure is detectable by that sensor unit 3a, 3b. Each sensor unit 3a, 3b is positioned to capture images of the respective portion (i.e. area) 30a, 30b of the surface 29 that is within its field of view ("sensing area") below. As can be seen in Figure 3, the fields of view of the first and second sensor units 3a, 3b overlap in the sense that there is a region of space within which structure is detectable by both sensor units 3a, 3b. As a result, one of the borders 30R of the sensing area 30a of the first sensor unit 3a is within the sensor area 32b of the second sensor unit 3b ("second sensing area"). Likewise, one of the borders 30L of the sensor area 32b of the second sensor unit 3b is within the sensor area 30a of the first sensor unit 3a ("first sensing area"). An area A is shown, which is the intersection of the first and second sensor areas 30a, 30b. The area A is the part of the surface 29 that is visible to both of the first and second sensor units 3a, 3b ("sensor overlap").

Figure 3A shows a plan view of a part of the lighting system 1, in which a 3x3 grid of nine sensor units 3a,...,3h is shown, each having a respective sensor area 30a,...,30h, which is the sensor area of its respective image sensor 6 as described above. The sensing area 30 of each sensor unit 3 overlaps with that of each of its neighbouring sensor units 3, in both directions along the grid and both directions diagonal to the grid, as shown. Thus every pair of neighbouring sensor units (3a, 3b), (3a, 3c), (3a, 3d), (3b, 3c), ... has an overlapping sensor area (or field of view, FoV). The overlapping sensing areas of the vision sensors ensure that there are no dead sensing regions.

Although nine luminaires are shown in Figure 3A, the present techniques can be applied to lighting systems with fewer or more sensor units 3. It will also be appreciated that the grid arrangement of Figure 3A is just one example for achieving a desired overlap.

Figure 4 shows a block diagram of the processing apparatus 20. The processing apparatus comprises at least one computer device for operating the lighting system 1. For example the computer device may take the form of a server, or a static user terminal such as a desktop computer, or a mobile user terminal such as a laptop, tablet, smartphone or smart watch. Whatever form it takes, the computer device 20 comprises a processor 27 formed of one or more processing units, and a network interface 23. The network interface 23 is connected to the processor 27. The processor 27 has access to a memory 22, formed of one or more memory devices, such as one or more RAMs, EEPROMs, magnetic memories or optical memories. The memory 22 may be external or internal to the computer device 20, or a combination of both (i.e. the memory 22 can, in some cases, denote a combination of internal and external memory devices), and in the latter case may be local or remote (i.e. accessed via a network). The processor 27 is also connected to a display 25, which may for example be integrated in the computer device 20 or an external display.

The processor 27 is shown executing people counting code 24, from the memory 22. Among other things, the people counting code 27 applies an aggregation algorithm, to aggregate multiple local presence metrics received from different ones of the sensor units 3 so as to generate an estimate of the number of people 8 in the environment 2. In this way the processor 27 implements a processing module connected to receive data relating to the captured images of the image capturing device, and to thereby determine a count of the number or rate if people being found in the environment 2.

The network interface 23 can be a wired interface (e.g. Ethernet, USB, FireWire) or a wireless interface (e.g. Wi-Fi, Bluetooth, ZigBee), and allows the computer device 20 to connect to the gateway 10 of the lighting system 1. The gateway 10 operates as an interface between the computer device 20 and the lighting network, and thus allows the central processing apparatus 20 to communication with each of the luminaires 4 and sensor units 3 via the lighting network. The gateway 10 provides any necessary protocol conversion to allow communication between the computer device 20 and the lighting network. Alternatively the interface 23 may enable the computer device 20 to connect directly to the luminaires 4 and sensor units 3. Either way, this allows the computer device 20 to transmit

control signals to each of the luminaires 4 and receive measurements from each of the sensors 3.

Note that the figures herein, including Figures 2 and 4, are highly schematic. In particular, the arrows denote high-level interactions between components of the luminaires 4, sensor units 3 and central computer 20 and do not denote any specific configuration of local or physical connections.

Note that the computer device 20 may be local to the environment 2 (e.g. present in the environment 2 or in the same building) or may be remote from it (at a remote geographic site), or the processing apparatus 20 may even comprise a combination of local and remote computer devices. Further, it may connect to the gateway 10 via a single connection or via another network other than the lighting network.

Figure 4A shows an exemplary lighting system control architecture for implementing a remote or networked connection between the computer device 20 and the gateway. Here, the computer device 20 is connected to the gateway 10 via a packet basic network 42, which is a TCP/IP network in this example. The computer device 20 communicates with the gateway 10 via the packet based network 42 using TCP/IP protocols, which may for example be effected at the link layer using Ethernet protocols, Wi-Fi protocols, or a combination of both. The network 42 may for example be a local area network (business or home network), the Internet, or simply a direct wired (e.g. Ethernet) or wireless (e.g. Wi-Fi) connection between the computer device 20 and the gateway 10. The lighting network 44 is a ZigBee network in this example, in which the luminaires 4a, 4b, 4c,... communicate with the gateway 10 using ZigBee protocols. The gateway 10 performs protocol conversion between TCP/IP and ZigBee protocols, so that the central computer 20 can communicate with the luminaires 4 and sensor units 3 via the packet based network 32, the gateway 10 and the lighting network 44.

Anywhere herein where there is referred to a processing apparatus 20 external to the sensor units 3 and luminaires 4, this may comprise any one or more computer devices arranged according to any of the possibilities discussed above or others. Note also that “external” or “externally” means the processing apparatus 20 is not housed within any shared housing (casing) of any of the sensor units 3, and in embodiments nor in any housing of the luminaires 4. Further, this means the processing apparatus communicates with all of the involved sensor units 3 (and in embodiments luminaires 4) only using an external connection via a networked and/or wireless connection, e.g. via the gateway 10, or via a direct wireless connection.

The memory 22 of the external processing apparatus 20 stores a database 21. This database 21 contains a respective identifier (ID) of each sensor unit 3 and each luminaire 4 in the lighting system 1 (or just IDs of the luminaires 4 when the sensor units 3 are integrated into luminaires 4). These uniquely identify the sensor units 3 and luminaires 4 within the system 1. Further, the database 21 also contains an associated location identifier 71 of each sensor unit 3 and luminaire (of again just the location identifiers of the luminaires 4 if the sensor units are integrated into luminaires). For example, each location identifier 71 may be a two dimensional identifier (x,y) or three dimensional location identifier (x,y,z) (e.g. if the sensor units 3 are installed at different heights). The location identifier 71 may convey only relatively basic location information, such as a grid reference denoting the position of the corresponding luminaire 4 or sensor unit in a grid – e.g. (m,n) for the mth column and nth row - or it may convey a more accurate location on a floor plan or map, e.g. meters, feet or arbitrary units, to any desired accuracy. The IDs of the luminaires 4 and sensor units 3, and their locations, are thus known to the processing apparatus 20. The memory 22 may also store additional metadata 26, such as an indication of the sensor overlap A, and any other sensor overlaps in the system.

Figure 5 illustrates how the processing apparatus 20 and the sensor units 4 cooperate within the system 1. First, second and third sensor units 3a, 3b, 3c are shown, though this is purely exemplary.

The image sensor 6 of each sensor unit 3a, 3b, 3c captures at least one respective image 60a, 60b, 60c of its respective sensing area (each of which could be a still image or a video). The local processing module 11a, 11b, 11c of that sensor unit applies the local person detection algorithm to the respective image(s). That is, the local person detection algorithm is applied separately at each of the sensor units 3a, 3b, 3c, in parallel to generate a respective local presence metric 62a, 62b, 62c at each, also referred to equivalently as a people counting metric herein. Each of the local presence metrics 62a, 62b, 62c is transmitted to the processing apparatus 20, e.g. via the networks 42, 44 and gateway 10. To preserve privacy, the images 60a, 60b, 60c themselves however are not transmitted to the central processing apparatus 20 (or at least not in a high enough resolution form for people to be recognizable or a least identifiable).

The external processing apparatus 20 applies the aggregation algorithm to the presence metrics 62a, 62b, 62c in order to estimate the number of people 8 in the environment. The aggregation algorithm generates an indicator of this number (people count)

64, which may be outputted on the display 25 to user of the processing apparatus 20 and/or stored in the memory 22 for later use.

The process may be real-time, in the sense that each local processing module 11a, 11b, 11c repeatedly generates and transmits local presence metrics as new images are captured. The people count 64 is updated as the new presence metrics are received, for example one every few (e.g. ten or fewer) seconds. Alternatively, the process may be pseudo-real-time, e.g. such that the people count 64 is updated every minute or every few minutes, or every hour (for example), or it may be pseudo-static e.g. a “one-time” people count may be obtained in response to a count instruction from the user of the external processing apparatus 20, to obtain a snapshot of current occupancy levels manually. That is, each count may be instructed manually.

Each presence metric 62 may be generated over a time window i.e. based on multiple images within that time window. This allows movements above a certain speed to be filtered out. I.e. objects moving fast enough to not appear in all of those images may be filtered out so that they do not affect the people count 64.

As discussed, in embodiments the sensor unit 3a captures images of the part of the surface 29 directly below it. This means the image 60a is a top-down view of the person 61, whereby the top of their head and shoulders are visible. Note that, in the case that the person 61 is in the sensor overlap area A, they would be similarly detectable in an image captured by the second sensor unit 3b. That is the same person 61 would be simultaneously visible in images from both the first and second sensor units 3a, 3b, at different respective locations in those images. A similar scenario can also occur even if the sensor units 3 do not face directly down, e.g. are at an angle in a corner of a room, or face sideways from the wall. It will be appreciated that the present disclosure is not limited to a top-down arrangement.

In embodiments, each sensor unit 3 (or rather its local image processor 11) communicates a respective one or more presence metrics, along with its ID and a timestamp, to the external processing apparatus 20 (e.g. a centralized people counting computer device). The timestamp is generated based on that sensors local clock signal 19.

The presence metric(s) reported by each sensor unit 3 comprise at least an indication of whether a person 8 is detected, or likely to have been detected, by the sensor unit 3. For example, this may comprise a yes/no flag indicating whether a person was detected. Alternatively or additionally, it may comprise a block-pixel-by-block-pixel score matrix, e.g. a 10 by 10 matrix of binary values e.g. with each element a “1” or “0”, indicative of presence or no presence - this choice ensures that the communication from the sensor units

3 to the external processing apparatus 20 maintains privacy, and is also low rate. Another alternative or additional possibility is to report a score a probability score indicative of the probability that a person is present. The probability score may be computed over a time window, thus filtering out movements above a certain speed. These may be estimated using known statistical methods, e.g. maximum *a posteriori* (MAP). Further, in embodiments, the reported presence metrics may comprise a location vector denoting the location of the detected person 61, e.g. which may be expressed relative to the sensor unit 3 that captures the image 60, or as a position within the image.

The external processing apparatus 20 collects such metrics from all the sensor units 3 associated with a region over which a people count is of interest (e.g. all or part of the surface 29). Additionally, the external processing apparatus 20 has knowledge of sensing region overlap of the sensor units 3, from the metadata 26. It aggregates the individual vision sensor counts while avoiding double-counts over overlapping regions within a given time window. For example, if the reported presence metric(s) from each sensor unit 3 comprise a yes/no indication of whether or not they detected a person 8, plus an associated location vector indicating where the person was detected, then the processing apparatus 20 can determine when a result from two different sensor unit 3 is within an overlap region (e.g. A in the example of Figure 3) and occurs at approximately the same location. This may be counted as one result instead of two to avoid double counting. As another example, if the reported presence metric(s) comprise a block pixel matrix, then the processing apparatus 20 can again determine when a result from two different sensor unit 3 is within an overlap region and occurs at approximately the same location in order to avoid double counting. The regions of overlap A can be determined at the commissioning stage and pre-stored in the memory 22 as part of the metadata 26, or alternatively can be determined automatically by the processing apparatus 20 automatically based on outputs from the sensors.

Alternatively or in addition, at least part of the metadata 26 may be available to the sensors 3 themselves, such that the sensors themselves have knowledge of sensing region overlap.

As a further alternative it is not necessary to have an overlap between the sensing regions 30 of different sensor units 3. While an overlap is preferred to avoid blind spots, in other embodiments double counting could instead be avoided by simply arranging the sensor units 3 to have non overlapping fields of view 30.

The above has described a system for detecting the presence of people 8 in an environment, and for counting the total number of people detected during a certain window

of time. This can have a number of applications, such as marketing analysis in a retail environment; or tracking occupancy levels for safety reasons (e.g. at a sports or entertainment event); or to inform automated control of a utility such as the illumination provided by the luminaires, or heating, ventilation or air conditioning.

5 Individual embodiments of the present invention will now be described.

For convenience in the following individual sensors 3 are denoted S_n (denoting the n th sensor in the system). A measurement performed by that sensor is denoted $m_n(t)$ where t is a time denoted by an associated timestamp of that measurement generated locally at the sensor S_n based on its local clock signal 19. That is, t is a time measured based
 10 on the local clock signal 19 of that sensor and thus expressed relative to its local reference time and based on the frequency of its clock signal 19.

A flowchart for one method of measurement synchronization is shown in figure 7, first and second examples of which are described below with reference to figures 8 and 9 respectively. The description of figures 8 and 9 is interleaved with that of figure 7.

15 In the first and second examples, measurements performed by the sensors are synchronized externally at the central processing apparatus 20, on the assumption that the local reference times and/or the frequencies of different sensors may not be synchronized. As indicated above, synchronization of the sensors in this case is not achieved by adjusting the sensors (which would require some form of signalling between the external processing
 20 apparatus and the sensors), and in particular is not achieved by adjusting how they apply their respective timestamps to correct inconsistencies between their respective clock signals – rather, the inconsistencies between the clock signals are allowed to persist locally at the sensors, and accounted for centrally at the external processing apparatus instead.

Step S2 represents operations performed over an interval of time, during
 25 which each sensor S_n performs and communicates to the central processing apparatus 20 a respective set of measurements $m_n(t_1), m_n(t_2) \dots$, performed by that sensor S_n at times $t_1, t_2, \dots$ as measured locally at that sensor S_n . This is illustrated in figure 6A for sensors 3a (S_1) and 3b (S_2). The measurements pertain at least in part to the area(s) of overlap “A” between the sensor S_n and its neighbouring sensor(s) in the grid of figure 3A, as indicated by the
 30 dashed lines of figure 6A.

In the first example each measurement $m_n(t)$ is a location (scalar or vector) of a person detected anywhere in the area covered by the sensor S_n (e.g. 30a or 30b), including any area(s) A of sensor overlap with neighbouring sensor(s), denoted $x_n(t)$ below. The location of each person may for instance be with respect to the vision sensor; that is, relative

to a location of the vision sensor itself in a spatial reference frame local to that sensor. In this local reference frame, the location of the sensor may for example be (0,0) or (0,0,0). In this case, the central processing apparatus 20 converts the locations to a global spatial frame reference, based on the locations of the locations of the vision sensor recorded in the commissioning database 20 relative to a common origin (for example, relative to a floorplan); this ensures that locations originating from the same person overlap over space when expressed in the global frame of reference.

In the second example each measurement $m_n(t)$ is a light level, denoted $l_n(t)$, measured at local time t , which may for example be measured over the whole of the sensing area (e.g. 30a or 30b) including any area(s) A of sensor overlap at local time t , or over a respective sub-region of the covered area at local time t . For example, where the sensor S_n has knowledge of the area(s) of sensor overlap A between that sensor S_n and its neighbour(s), the measurement $m_n(t)$ may be a respective light level measured over area of each of those area(s) of overlap, i.e. one light level per area of sensor overlap. Alternatively, the measurement $m_n(t)$ may be a light level in a grid of light level (e.g. a 10x10 grid of light levels), for a grid of sub regions the covered area (e.g. 30a or 30b in figure 3) – in this case, the measurement $m_n(t)$ constitutes a very low resolution, monochromatic image derived from a full resolution image captured by the camera of sensor S_n at local time t .

Because the local reference times and/or clock signal frequencies are not necessarily synchronized, where two measurements $m_n(t), m_m(t')$ have respective timestamps such that $t \approx t'$ i.e. where the measurements $m_n(t), m_m(t')$ have substantially matching timestamps, that does not necessarily mean they correspond to the same physical time. That is, the timestamps from different sensors 3 may be inconsistent with one another. If when the measurements $m_n(t), m_m(t')$ are performed, there exists a timing offset – which can be expressed as a time difference $\Delta t_{n,m}$ – between the respective clock signals of S_n and S_m – arising due to those clock signals being based on different local reference times, due to them having different frequencies, or a combination of both – the fact that the measurements $m_n(t), m_m(t')$ have substantially matching timestamps actually means they were performed a time $\Delta t_{n,m}$ apart.

The timestamps are substantially matching in the sense that the time t denoted by the timestamp generated by S_n is closer to the time t' denoted by the timestamp generated by S_m than the time denoted by any other timestamp generated by S_m .

At step S4, the central processing apparatus 20 computes $\Delta t_{n,m}$ for each pair of adjacent sensors S_n , S_m by evaluating a correlation function for a plurality of different candidate (i.e. potential time differences) time differences δt (that is for $\delta t = 0,1,2$ expressed in units of whole deltas for convenience). One example of a suitable correlation function is the following:

$$\text{corr}(\delta t) = \sum_t m_n(t) \cdot m_m(t - \delta t) \quad (1)$$

This is sometimes referred to in the art as the simple cross-correlation function. In other words, the central processing apparatus 20 multiplies each measurement performed by the sensor S_n with the measurement performed by the sensor S_m whose timestamp corresponds to the timestamp of the first measurement offset by that difference value δt . Note that this is just one example of a simple correlation function, and other examples are within the scope of this disclosure, such as a normalized version of the simple correlation function of equation 1 – for example normalized to a maximum value of 1 or to have any desired value other than 1, such as the known “moravec” correlation function) – and/or a zero-mean version of the correlation function of equation (1).

Other suitable correlation functions that can be used include correlation function based on a sum of differences, for example a correlation function as follows:

$$\text{corr}(\delta t) = - \sum_j c(m_n(t), m_m(t - \delta t)) \quad (2)$$

where $c(a, b)$ is a suitable comparison function such as $c(a, b) = \|a - b\|$ (that is, the modulus of a less b) in which case the above corresponds to the known sum of absolute differences (“SAD”) cross-correlation function, or $c(a, b) = (a - b)^2$ (that is, the square of a less b), in which case the above corresponds to the known sum of squared differences (“SSD”) cross-correlation function. In general, the function $c(a, b)$ is minimized when $a = b$, such that $\text{corr}(\delta t)$ is maximized when $a = b$ (note the minus sign in equation (2)).

In general, any suitable correlation function which constitutes a measure of similarity or dependence between measurements from different sensors can be used in place of correlation functions based on equations (1) and (2).

The summation is over a suitable time window, over which correlations in the measurements from different sensors are detectable. The summation is also over measurements performed within the overlapping region A. Thus, where an event occurs and/or an object is present in the common area, both sets of measurements $m_n(t)$, $m_m(t)$ are associated with the same object and/or event, and will this exhibit detectable correlations.

That is, multiple measurements are performed over the same area over time, as the measured quantity or quantities in that same area change over time.

There exists a difference value δt_M for which the correlation function $\text{corr}(\delta t)$ is maximized. That is, for which each measurement $m_2(t - \delta t_M)$ from the second sensor S_2 timestamped with time $t - \delta t_M$ as measured locally at the second sensor has a maximum correlation with the respective first measurement $m_1(t)$ timestamped with time t as measured locally at the first sensor S_1 . The correlation is maximized because, in reality $m_2(t - \delta t_M)$ and $m_1(t)$ were performed substantially simultaneously.

The central processing apparatus estimates, for each pair of adjacent sensors S_n, S_m , the time offset $\Delta t_{n,m}$ between the respective clock signals of the sensors S_n, S_m as:

$$\Delta t_{n,m} = \delta t_M = \arg \max_{\delta t} \text{corr}(\delta t)$$

That is the time offset $\Delta t_{n,m}$ is estimated to be equal to the difference value δt_m for which the correlation function $\text{corr}(\delta t)$ is maximized. As will be readily apparent, for equation 2 the minus sign before the summation may be omitted in which case $\Delta t_{n,m} = \arg \min_{\delta t} \text{corr}(\delta t)$.

As noted, in the first example, the central processing apparatus 20 collects location information from adjacent vision sensors, from the commissioning database 26. Additionally, the unit has knowledge of the overlap of the vision sensors, from the metadata 26. The processing apparatus 20 correlates over time the reported location of person over the overlapping region for a pair of adjacent vision sensors.

Figure 8 shows the first example, in which a first vision sensor S_1 and a second vision sensor S_2 report respectively locations $x_1(t)$ and $x_2(t)$ to the central processing apparatus 20. The central processing apparatus 20 estimates the time shift Δt between the respective clocks of S_1 and S_2 by correlating the locations, for example based on equation 2 as:

$$\Delta t_{1,2} = \arg \max_{\delta t} - \sum_t c(x_1(t), x_2(t - \delta t))$$

where $c(x_1(t), x_2(t - \delta t)) = \|x_1(t) - x_2(t - \delta t)\|^2$, such that the sum of squared distances is minimized.

As another example, the offset can be estimated based on equation 1 as:

$$\Delta t_{1,2} = \arg \max_{\delta t} \sum_t x_1(t) \cdot x_2(t - \delta t)$$

In the example of figure 8, the time shift, $\Delta t_{1,2}$ is 3 units; that is $\delta t_M = 3$, though this is purely exemplary.

Where x_1 and x_2 are two or three dimensional location vectors, the dot in above equations (i.e. “ $\cdot$ ”) denotes the vector inner product.

5 In the second example, each vision sensor 3 communicates to the central processing apparatus 20, in addition to the location of any detected person, light levels in predefined regions within its field of view, along with its vision sensor ID and a timestamp of each measured light level. Each predefined region may for instance be with respect to the vision sensor and so the central processing apparatus 20 converts the location of these regions
10 to a global reference; this ensures that similar located regions are overlapping correctly.

The central processing apparatus 20 collects the light levels from adjacent vision sensors. Additionally, the unit has knowledge of the overlap (if any) of the vision sensors, from the metadata 21. Then, the unit correlates over time the light levels over overlapping (or nearby) regions for a pair of adjacent vision sensors. Note that the light levels
15 might change due to lighting control (e.g. as a result daylight or occupancy adaptation), so that correlations between light levels reported by different sensors are detectable due to the changing light levels.

Figure 9 shows an example of predefined regions over which vision sensor S_1 and vision sensor S_2 report light levels $l_1(t)$ and $l_2(t)$ respectively, which are the light levels
20 over the sensor overlap region A between them (see figure 3). The central processing apparatus 20 estimates the time shift Δt by correlating the locations, for example as:

$$\Delta t_{1,2} = \arg \max_{\delta t} \sum_t l_1(t) \cdot l_2(t - \delta t)$$

or as:

$$\Delta t_{1,2} = \arg \max_{\delta t} - \sum_t c(l_1(t), l_2(t - \delta t))$$

Once the respective clock signal offset $\Delta t_{n,m}$ has been computed for each pair of adjacent sensors S_n , S_m , the central processing apparatus 20 can use the estimated clock
25 signal offsets $\Delta t_{n,m}$ to account for inconsistencies in the timestamps applied to later measurements (S6). For example, the central processing apparatus is able to generate an accurate people count 66 from the later measurements for any desired time (e.g. based on people locations or other presence metrics reported by the sensors), accounting for the inconsistencies in the timestamps applied to the later measurements arising due to the clock
30 signal offsets $\Delta t_{n,m}$.

Whilst the above example uses correlations to, in effect, identify measurements $m_1(t)$ and $m_2(t - \delta t_m)$ that were performed substantially simultaneously, this is not essential. For example, if the speed at which a person is walking is known, or the distribution of the light level over time and/or space is known, correlation between non-simultaneous measurements can be detected and used to determine the clock timing offsets in an equivalent manner. One way to do so is to use such knowledge to interpolate or extrapolate the location of a person for one or both measurements thereby allowing correlation along the lines of simultaneous measurement.

In the above, multiple measurements from the first sensor S_1 are compared with multiple measurement from the second sensor S_2 at different time offsets δt . It is preferable to compare multiple measurements from both sensors in this manner as it provides more accurate results. Nevertheless, in a few circumstances comparing a single measurement from the first sensor S_1 with measurements from the second sensor S_2 – though less preferred – is sufficient.

As noted, a person or people is detected at the vision sensors based on a high resolution image captured by that vision sensor. The higher resolution used for person detection at each of the sensor units 3 may for example be at least 100x100 pixels (at least 100 pixels in each of the horizontal and vertical dimensions). Alternatively, the higher resolution used for person detection at each of the sensor units 3 may be at least 500x500 pixels (at least 500 pixels in each of the horizontal and vertical dimensions). Alternatively, the higher resolution used for person detection at each of the sensor units 3 may be at least 1000x1000 pixels (at least 1000 pixels in each of the horizontal and vertical dimensions).

In the second example, the light levels reported over for each portion of the other area may effectively constitute a lower resolution image, as also noted above. The lower resolution of the second example used by each of the sensor units 3 to externally report the images to the processing apparatus 20 may be no more than 10x10 pixels (no more than ten pixels in each of the horizontal and vertical dimensions) – with each pixel in the lower resolution image being a light level measured over a respective part of the area covered by the sensor. In embodiments, the lower resolution used by each of the sensor units 3 to externally report the images to the processing apparatus 20 may be no more than 25x25 pixels (no more than twenty-five pixels in each of the horizontal and vertical dimensions). Alternatively, the lower resolution used by each of the sensor units 3 to externally report the images to the processing apparatus 20 may be no more than 50x50 pixels (no more than fifty pixels in each of the horizontal and vertical dimensions). In embodiments the lower

resolution used for reporting may be reduced by at least ten times in each dimension compared to the higher resolution used for detection (ten times fewer pixels in each of the horizontal and vertical directions). Alternatively or in addition, the lower resolution used for reporting may be reduced by at least fifty times in each dimension compared to the higher resolution used for detection (fifty times fewer pixels in each of the horizontal and vertical directions). Alternatively the lower resolution used for reporting may be reduced by at least one hundred times in each dimension compared to the higher resolution used for detection (one hundred times fewer pixels in each of the horizontal and vertical directions).

The sensor unit 3 may automatically and unconditionally report the lower resolution image to the external processing apparatus 20 each time an image captured, or may report it periodically. Alternatively the sensor unit 3 may automatically report the lower resolution image to the external processing apparatus 20 only in response to an event, e.g. whenever a local automated check performed by the local image processing code 12a determines the image does not conform to an empirical or analytical expectation, or in response to the local image processing code 12a detecting a debug sequence signalled in the illumination from one or more of the luminaires 4.

Figure 10 shows a flowchart for another synchronization method, which is performed in the third example described below. The third example is illustrated in figure 11, the description of which is interleaved with that of figure 10.

At step S52 a synchronization code is embedded in the visible light outputted by at least one of the luminaires 4 at a first time, and used by the sensors 3 as a reference point for synchronization. That is, the synchronization code defines a reference time that is global across the system.

As shown in figure 11, the illumination of one of the luminaires 4 falls within the fields of view of two adjacent vision sensors 3a, 3b. Hence the synchronization code when embedded in its illumination, is detected by both sensors 3a, 3b (S54).

The synchronization code is a modulation sequence, shown as 17 and can be embedded in the illumination by modulating any characteristic of the light (e.g. using anyone or more of amplitude, frequency and/or phase). For example, the luminaire 4 may dim its illumination according to the sequence 17 and each vision sensor 3a, 3b reports its time with respect to the sensor perceived dimming sequence, as explained below.

Information from each vision sensor 3a, 3b is synchronized at the central processing apparatus 20 during a synchronization phase. This phase may be performed repeatedly at certain intervals chosen so as not to disturb users (e.g. during intervals of

global un-occupancy of the environment). During this synchronization phase, the luminaire is dimmed with the sequence 17 that allows for timing retrieval (e.g. based on gold codes). Each vision sensor 3a, 3b knows (e.g. receive from the central processing apparatus 20) the sequence 17 and thus can determine time locally with respect to the start of the sequence 17.

5 Alternatively, the dimming sequence can be transmitted using visible light communication, using on so-called “visible light communication” techniques and thus a continuous (and more accurate) synchronization can be achieved, as the synchronization code when embedded using visible light communication is not perceptible to a human eye. In order to render the modulation substantially imperceptible it is generally advantageous to use
10 a modulation technique that does not introduce frequency components below 100 Hz.

 Thereafter (S56), at each of a plurality of second times, each vision sensor communicates to the central processing apparatus 20: a people counting metric of the kind described above (e.g. a people count, location information etc.) and a timestamp denoting that second time relative to the start of the dimming sequence (i.e. the first time), along with
15 its vision sensor ID. That is, as an elapsed time from the synchronization code. The timestamp is thus generated relative to the global reference time, such that the timestamps outputted all the vision sensors 3 in the system. That is, the timestamps are globally synchronized across the sensor system.

 The people counting metric has lower information content than the image(s)
20 from which it is generated, in this case one or more images captured by the camera of the vision sensor. The person(s) is not identifiable in the people counting metric even if they are identifiable in that image(s).

 The central processing apparatus 20 arranges the received information in chronological time with respect to the start of the dimming sequences for each vision sensor
25 3.

 As another example, a combinations of the third and second embodiments may be used. For example, in the event of an occupancy change, a luminaire may be turned on following a pre-defined dimming sequence, which is both detectable at the vision sensors 3 in accordance with the third embodiment and which leads to correlations in their outputs
30 detectable at the central processing apparatus 20 in accordance with the second embodiment, with both types of synchronization being performed to maximize accuracy.

 The central processing apparatus 20 uses the presence information and synchronized time stamps to provide an estimated people count over the total area covered by the sensors 3, for any desired time.

Each of the second times can be conveyed, to the external processing apparatus 20, relative to the first time (of the synchronization code) in a number of ways, for example as a difference value obtained by subtracting the first time from the second time; or by supplying the first time and the second time to the external apparatus 20 separately, as in
5 that the external processing apparatus 20 is able to compute the second time relative to the first time. In the case of the latter, these do not need to be outputted to the people counting apparatus separately.

Figure 6 shows a schematic block diagram of the sensing system, in which the vision sensors are integrated in the luminaires in the manner of figure 2B, to provide an
10 accurate people count 64 based on the synchronized data (synchronized according to one or more of the above embodiments).

It will be appreciated that the above embodiments have been described by way of example only.

For example, whilst in the above sensor units comprise sensor devices in the
15 form of visible or infrared cameras, other types of sensor device (ultrasound, radar etc.) are also viable.

Other variations to the disclosed embodiments can be understood and effected by those skilled in the art in practicing the claimed invention, from a study of the drawings, the disclosure, and the appended claims. In the claims, the word "comprising" does not
20 exclude other elements or steps, and the indefinite article "a" or "an" does not exclude a plurality. A single processor or other unit may fulfil the functions of several items recited in the claims. The mere fact that certain measures are recited in mutually different dependent claims does not indicate that a combination of these measures cannot be used to advantage. A computer program may be stored/distributed on a suitable medium, such as an optical storage
25 medium or a solid-state medium supplied together with or as part of other hardware, but may also be distributed in other forms, such as via the Internet or other wired or wireless telecommunication systems. Any reference signs in the claims should not be construed as limiting the scope.

CLAIMS:

1. A person detection sensor unit (3) comprising:
 - a communications interface (7);
 - a sensor device (6) configured to capture over time sensor data from an area covered by the sensor device;
 - 5 a processor (11) configured to implement the following steps:
 - detecting in a first portion of the sensor data captured at a first time a predetermined synchronization code, and measuring the first time based on a clock signal available at the sensor unit;
 - detecting in a second portion of the sensor data captured at a second time, at
 - 10 least one person present in the area, and measuring the second time based on the clock signal;
 - based on said detection at the sensor unit of the at least one person, generating from the second portion of the sensor data presence data for the second time pertaining to the detected at least one person; and
 - outputting via the communications interface the presence data for the second
 - 15 time and associated timing data, which conveys the second time as measured at the sensor unit relative to the first time as measured at the sensor unit.
2. A sensor unit according to claim 1, wherein the presence data for the second time comprises:
 - 20 a presence count indicating a number of people detected by the sensor unit in the covered area at the second time; and/or
 - a presence score indicating a likelihood that there is a person or people in the covered area at the second time; and/or
 - a number of person location identifiers, each identifying a location of a person
 - 25 detected in the covered area at the second time.
3. A sensor unit according to claim 2, wherein each person location identifier is a two or three dimensional location vector.

4. A sensor unit according to claim 1, 2 or 3, wherein sensor device is a photo sensor device configured to sense visible and/or non-visible radiation.

5. A sensor unit according to claim 4, wherein the photo sensor device is an image capture device and the sensor data is image data, the first portion of the image data being one or more first images captured by the image capture device of the area and the second portion being one or more second image captured by the image capture device of the area.

6. A sensor unit according to claim 4 or 5, wherein the processor is configured to detect the synchronization code embedded in the radiation as amplitude and/or phase modulations.

7. A sensor unit according to any preceding claim, wherein the processor is configured to determine a difference between the first time as measured at the sensor unit and the second time as measured at the sensor unit, wherein the associated timing data comprise the determined difference.

8. A person detection system according to any of claims 1 to 6, wherein the associated timing data comprises a first timestamp of the first time and a second stamp of the second time, and thereby conveys the first time relative to the second time.

9. A sensor system comprising:
a plurality of person detection sensors units;
a transmitting unit (4) configured to emit at a first time a synchronization code detectable by the sensor units and embodying the synchronization code;
wherein each of the sensor units is configured according to any preceding claim;

10. A sensor system according to claim 9, further comprising:
a people counting apparatus;
wherein each of the sensor units is configured to output respective presence data for the second time and associated timing data which conveys, to the people counting apparatus the second time relative to the first time; and

wherein the people counting apparatus is configured to use the respective presence data to estimate a people count for a total area covered by the sensor units

11. A sensor system according to claim 9 or 10, wherein the transmitting unit is a luminaire (4) configured to emit illumination at the first time in which the synchronization code is embedded, the sensor device of each sensor unit being a photo sensor device.

12. A sensor system according to claim 11, wherein the synchronization code is embedded using visible light communication, whereby it is imperceptible to a human eye.

13. A people detection method implemented by a person detection sensor unit (3) of a sensor system, the method comprising:

capturing over time sensor data from an area covered by the sensor unit;

detecting in a first portion of the sensor data captured at a first time a

predetermined synchronization code;

measuring the first time based on a clock signal available at the sensor unit;

detecting in a second portion of the sensor data captured at a second time at least one person present in the area;

measuring the second time based on the clock signal;

based on said detection at the sensor unit of the at least one person, generating from the second portion of the sensor data presence data for the second time pertaining to the detected at least one person; and

outputting, to a processing apparatus external to the sensor unit, the presence data for the second time and associated timing data, which conveys to the external processing apparatus the second time as measured at the sensor unit relative to the first time as measured at the sensor unit.

14. A method according to claim 13, wherein the presence data for the second time comprises:

a presence count indicating a number of people detected by the sensor unit in the covered area at the second time; and/or

a presence score indicating a likelihood that there is a person or people in the covered area at the second time; and/or

a number of person location identifiers, each identifying a location of a person detected in the covered area at the second time.

15. A computer program code comprising executable code stored on a computer
5 readable storage medium and configured when executed to implement the method of claim
13 or 14.

1/11




FIG. 1




FIG. 2




FIG. 2A




FIG. 2B




FIG. 3

FIG. 3A




FIG. 4





FIG. 4A






Fig 5




Fig 6





Fig 6A

7/11

Fig 7






Fig 8

9/11

Fig 9

10/11

Fig 10




Fig 11

INTERNATIONAL SEARCH REPORT

International application No

PCT/EP2016/080097

A. CLASSIFICATION OF SUBJECT MATTER
INV. H05B37/02
ADD.

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
H05B

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

EPO-Internal, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	WO 96/25021 A1 (MYTECH CORP [US]) 15 August 1996 (1996-08-15) page 1, line 9 - page 16, line 33; figures 2,7	1-15
A	WO 2014/187717 A1 (KONINKL PHILIPS NV [NL]) 27 November 2014 (2014-11-27) the whole document	1,9,13,15
A	WO 2015/052613 A1 (NEODELIS S R L [IT]; TORINO POLITECNICO [IT]) 16 April 2015 (2015-04-16) page 7, line 4 - line 20	1,13
A	EP 2 709 428 A2 (SENSITY SYSTEMS INC [US]) 19 March 2014 (2014-03-19) paragraph [0051] - paragraph [0054]	1,13



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

21 February 2017

Date of mailing of the international search report

28/02/2017

Name and mailing address of the ISA/

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040,
Fax: (+31-70) 340-3016

Authorized officer

Henderson, Richard

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/EP2016/080097

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
WO 9625021	A1	15-08-1996	CA 2239094 A1 15-08-1996
		DE 69615006 D1 11-10-2001	
		DE 69615006 T2 06-06-2002	
		DK 0809922 T3 27-12-2001	
		EP 0809922 A1 03-12-1997	
		JP H11501118 A 26-01-1999	
		US 5640143 A 17-06-1997	
		WO 9625021 A1 15-08-1996	
WO 2014187717	A1	27-11-2014	CN 105230129 A 06-01-2016
		EP 3000286 A1 30-03-2016	
		JP 2016519414 A 30-06-2016	
		US 2016095190 A1 31-03-2016	
		WO 2014187717 A1 27-11-2014	
WO 2015052613	A1	16-04-2015	CN 105830535 A 03-08-2016
		EP 3056065 A1 17-08-2016	
		US 2016242265 A1 18-08-2016	
		WO 2015052613 A1 16-04-2015	
EP 2709428	A2	19-03-2014	CN 103687200 A 26-03-2014
		EP 2709428 A2 19-03-2014	
		JP 2014064274 A 10-04-2014	
		KR 20140034712 A 20-03-2014	
		KR 20150089983 A 05-08-2015	
		US 2014084795 A1 27-03-2014	
		US 2016366753 A1 15-12-2016	