

(12) **UK Patent**

(19) **GB**

(11) **2567881**

(13) **B**

(45) Date of B Publication

10.02.2021

(54) Title of the Invention: **Systems and methods for processing a stream of data values**

(51) INT CL: **G06T 5/20** (2006.01) **G06T 5/00** (2006.01)

(21) Application No: **1717825.2**

(22) Date of Filing: **30.10.2017**

(43) Date of A Publication **01.05.2019**

(72) Inventor(s):
Tim Lee

(73) Proprietor(s):
**Imagination Technologies Limited
(Incorporated in the United Kingdom)
Imagination House, Home Park Estate,
KINGS LANGLEY, Hertfordshire, WD4 8LZ,
United Kingdom**

(56) Documents Cited:
GB 2510368 A1

(58) Field of Search:
As for published application 2567881 A viz:
INT CL **G06F, G06T, H04N**
Other: **WPI & EPODOC**
updated as appropriate

Additional Fields
Other: **None**

(74) Agent and/or Address for Service:
**Slingsby Partners LLP
1 Kingsway, LONDON, WC2B 6AN, United Kingdom**

GB 2567881 B

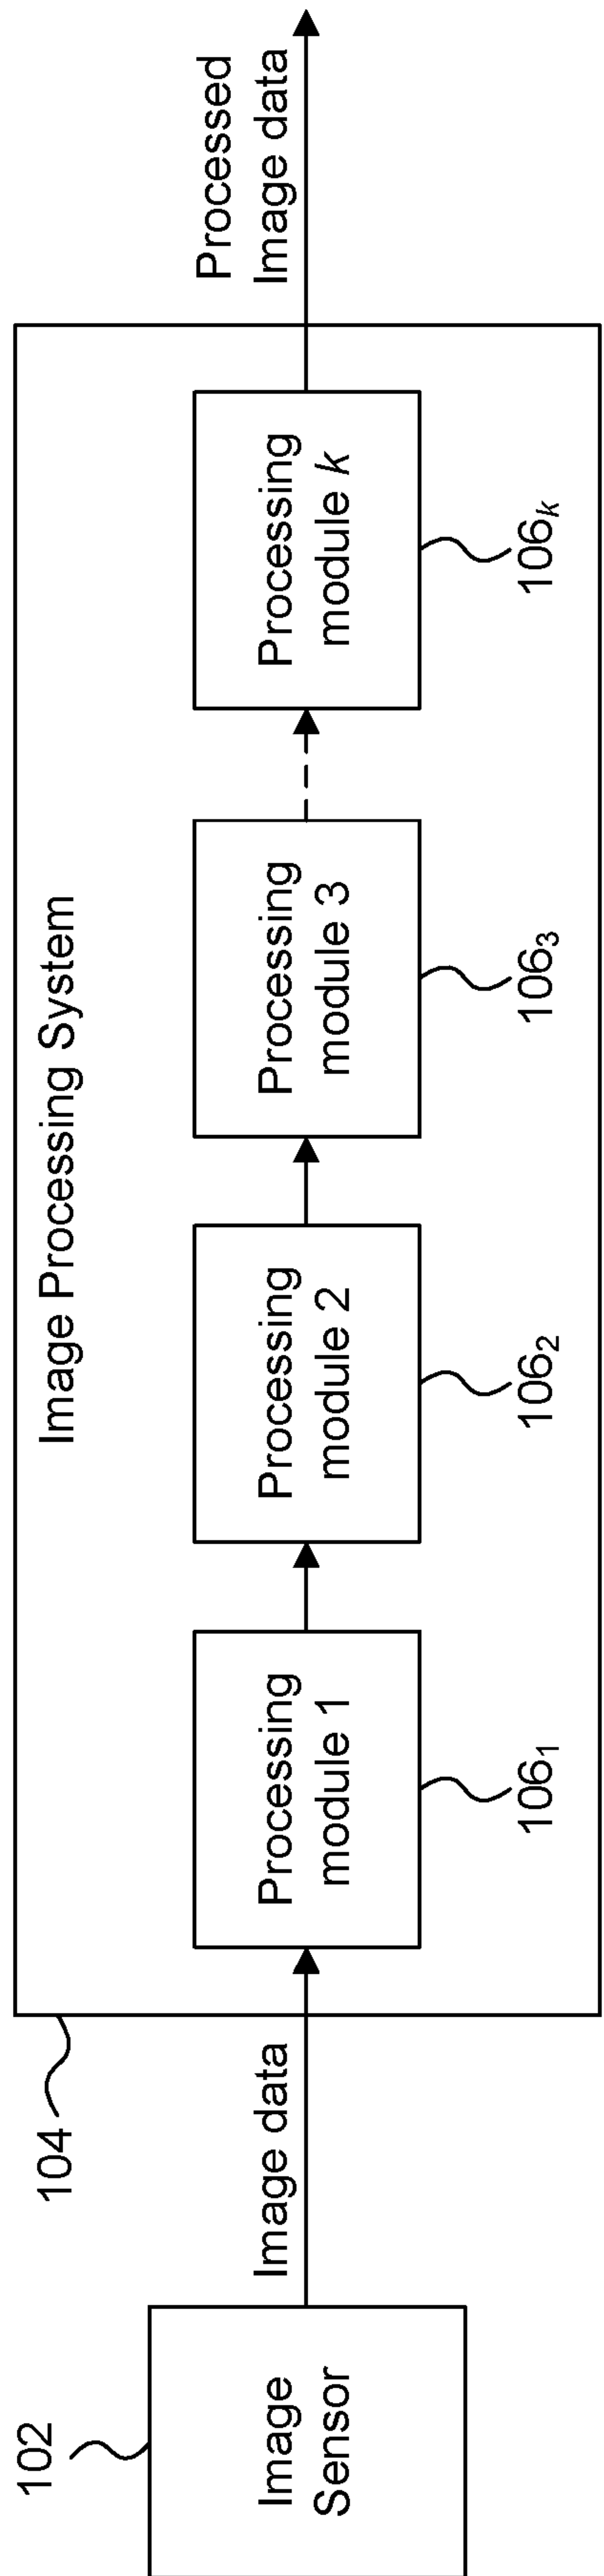


FIGURE 1a

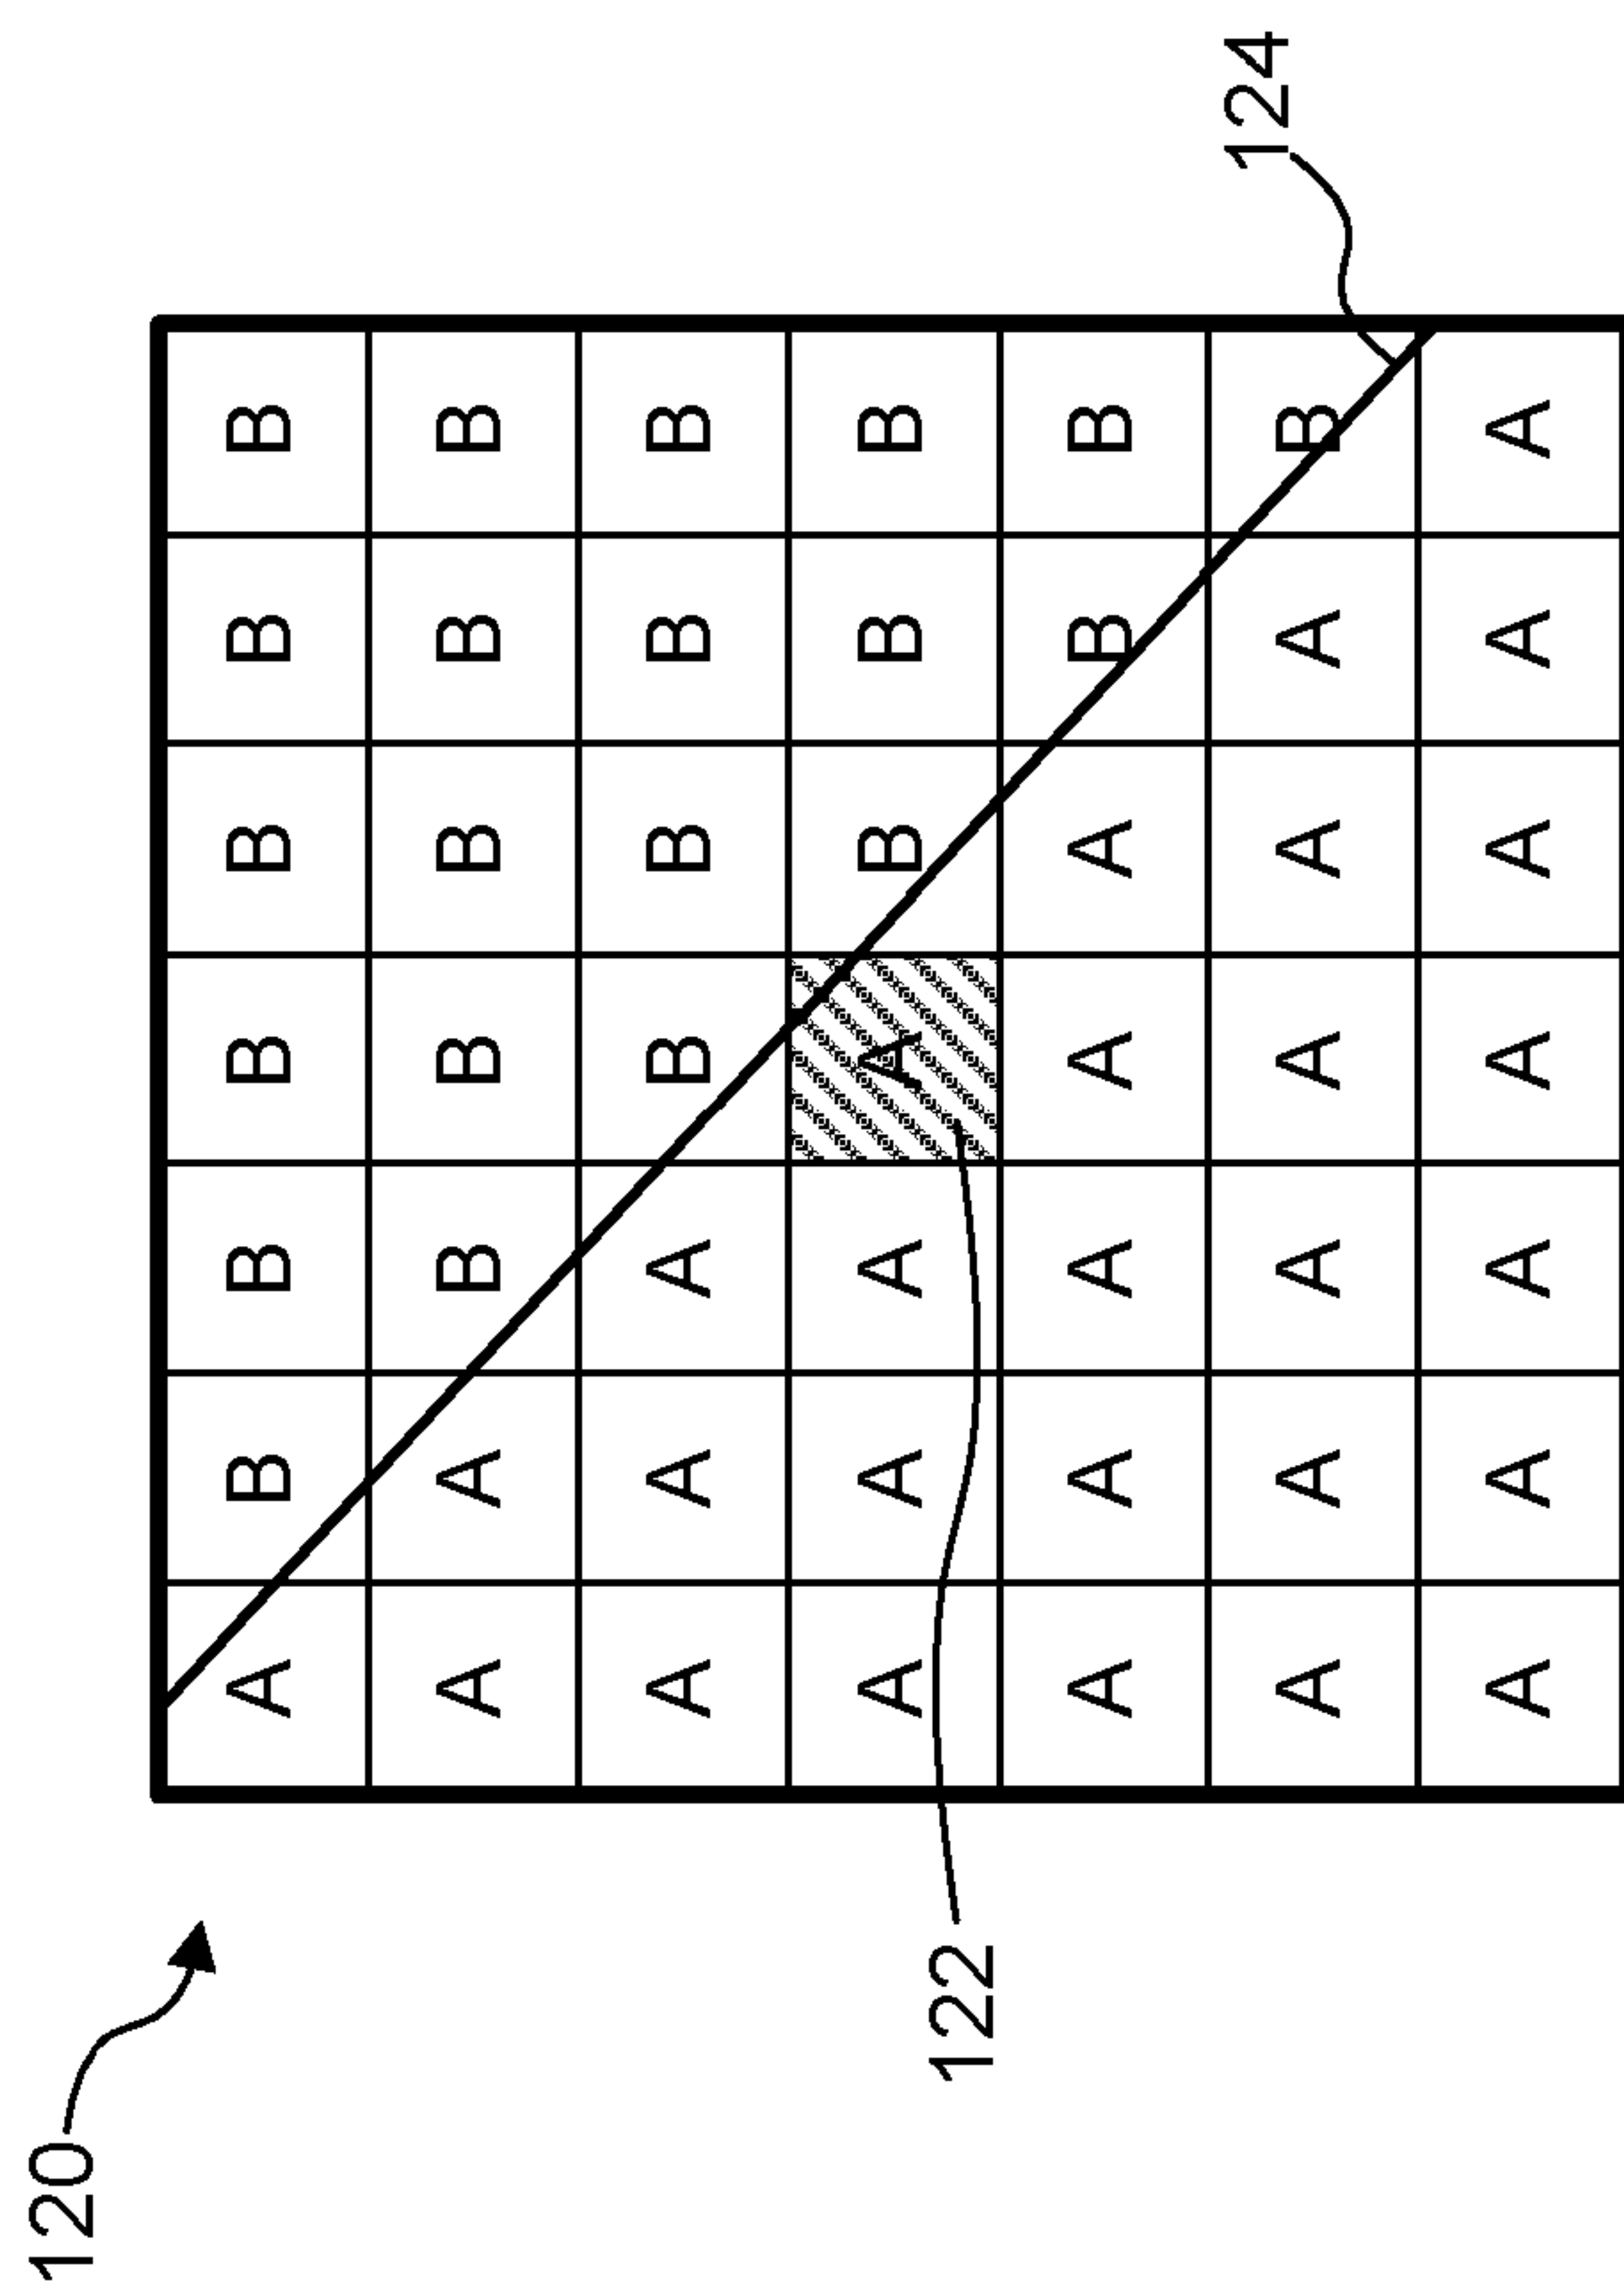


FIGURE 1b




FIGURE 2

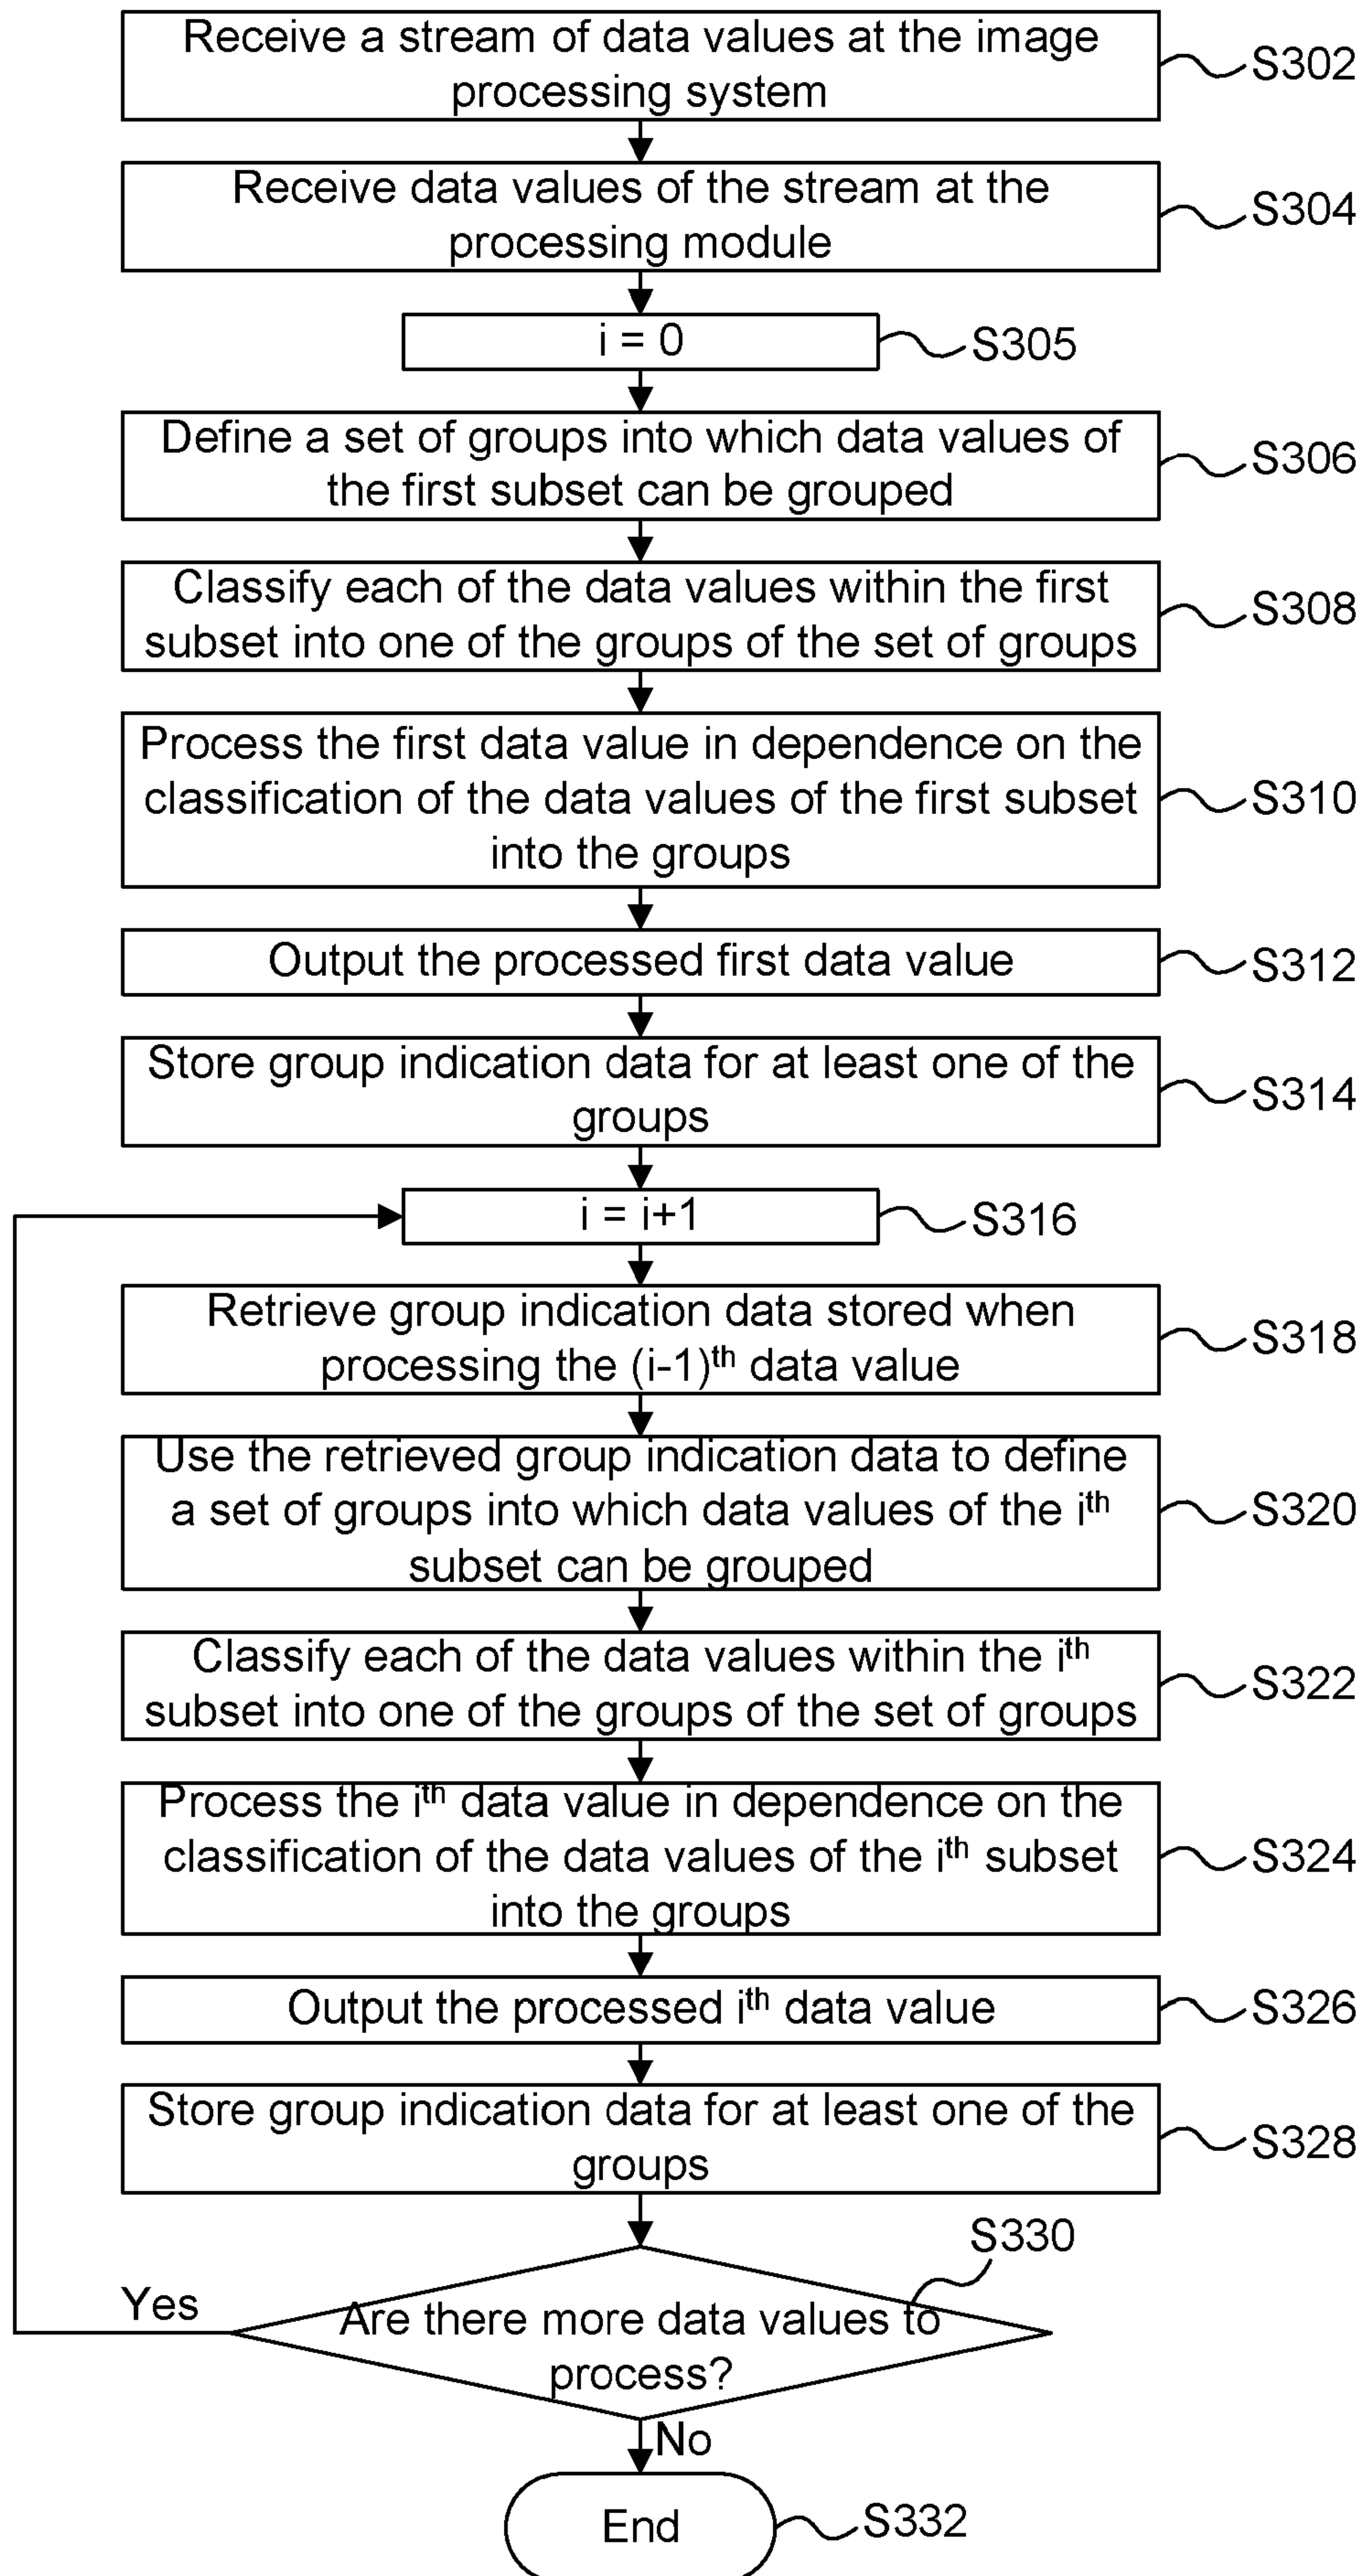


FIGURE 3




FIGURE 4




FIGURE 5




FIGURE 6




FIGURE 7

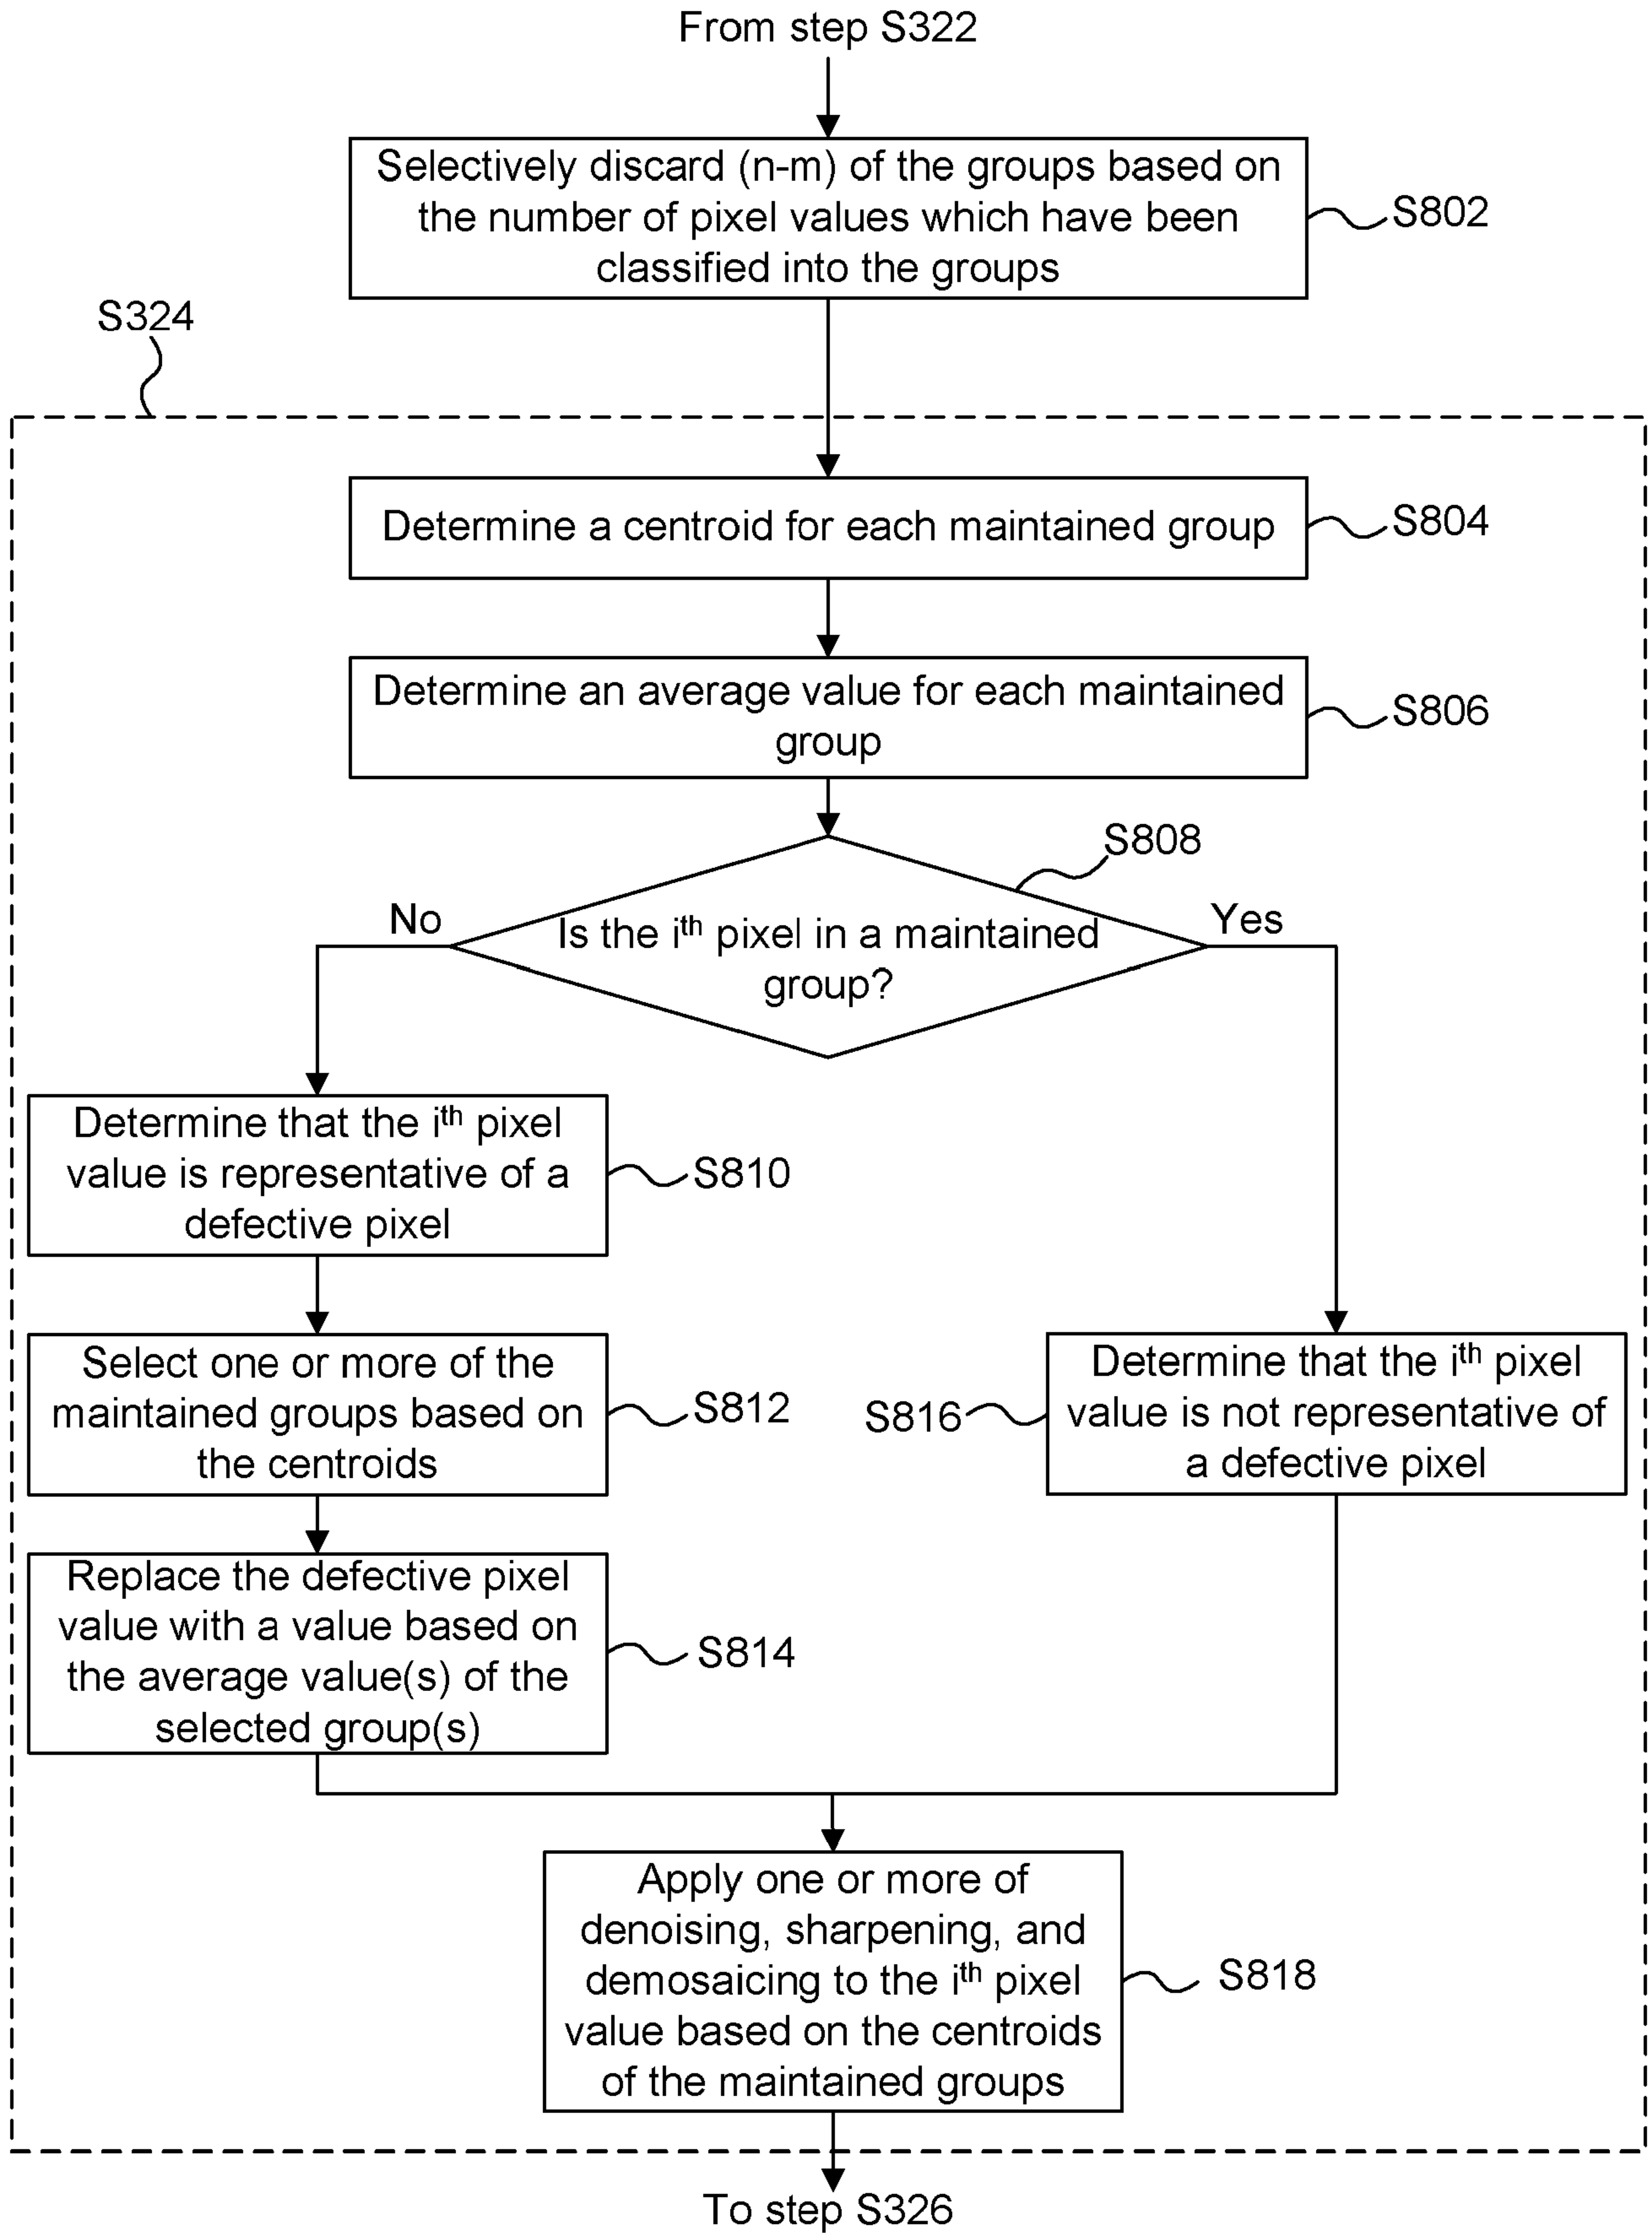




FIGURE 8




902

FIGURE 9a




904

FIGURE 9b



1002

FIGURE 10a



1004

FIGURE 10b

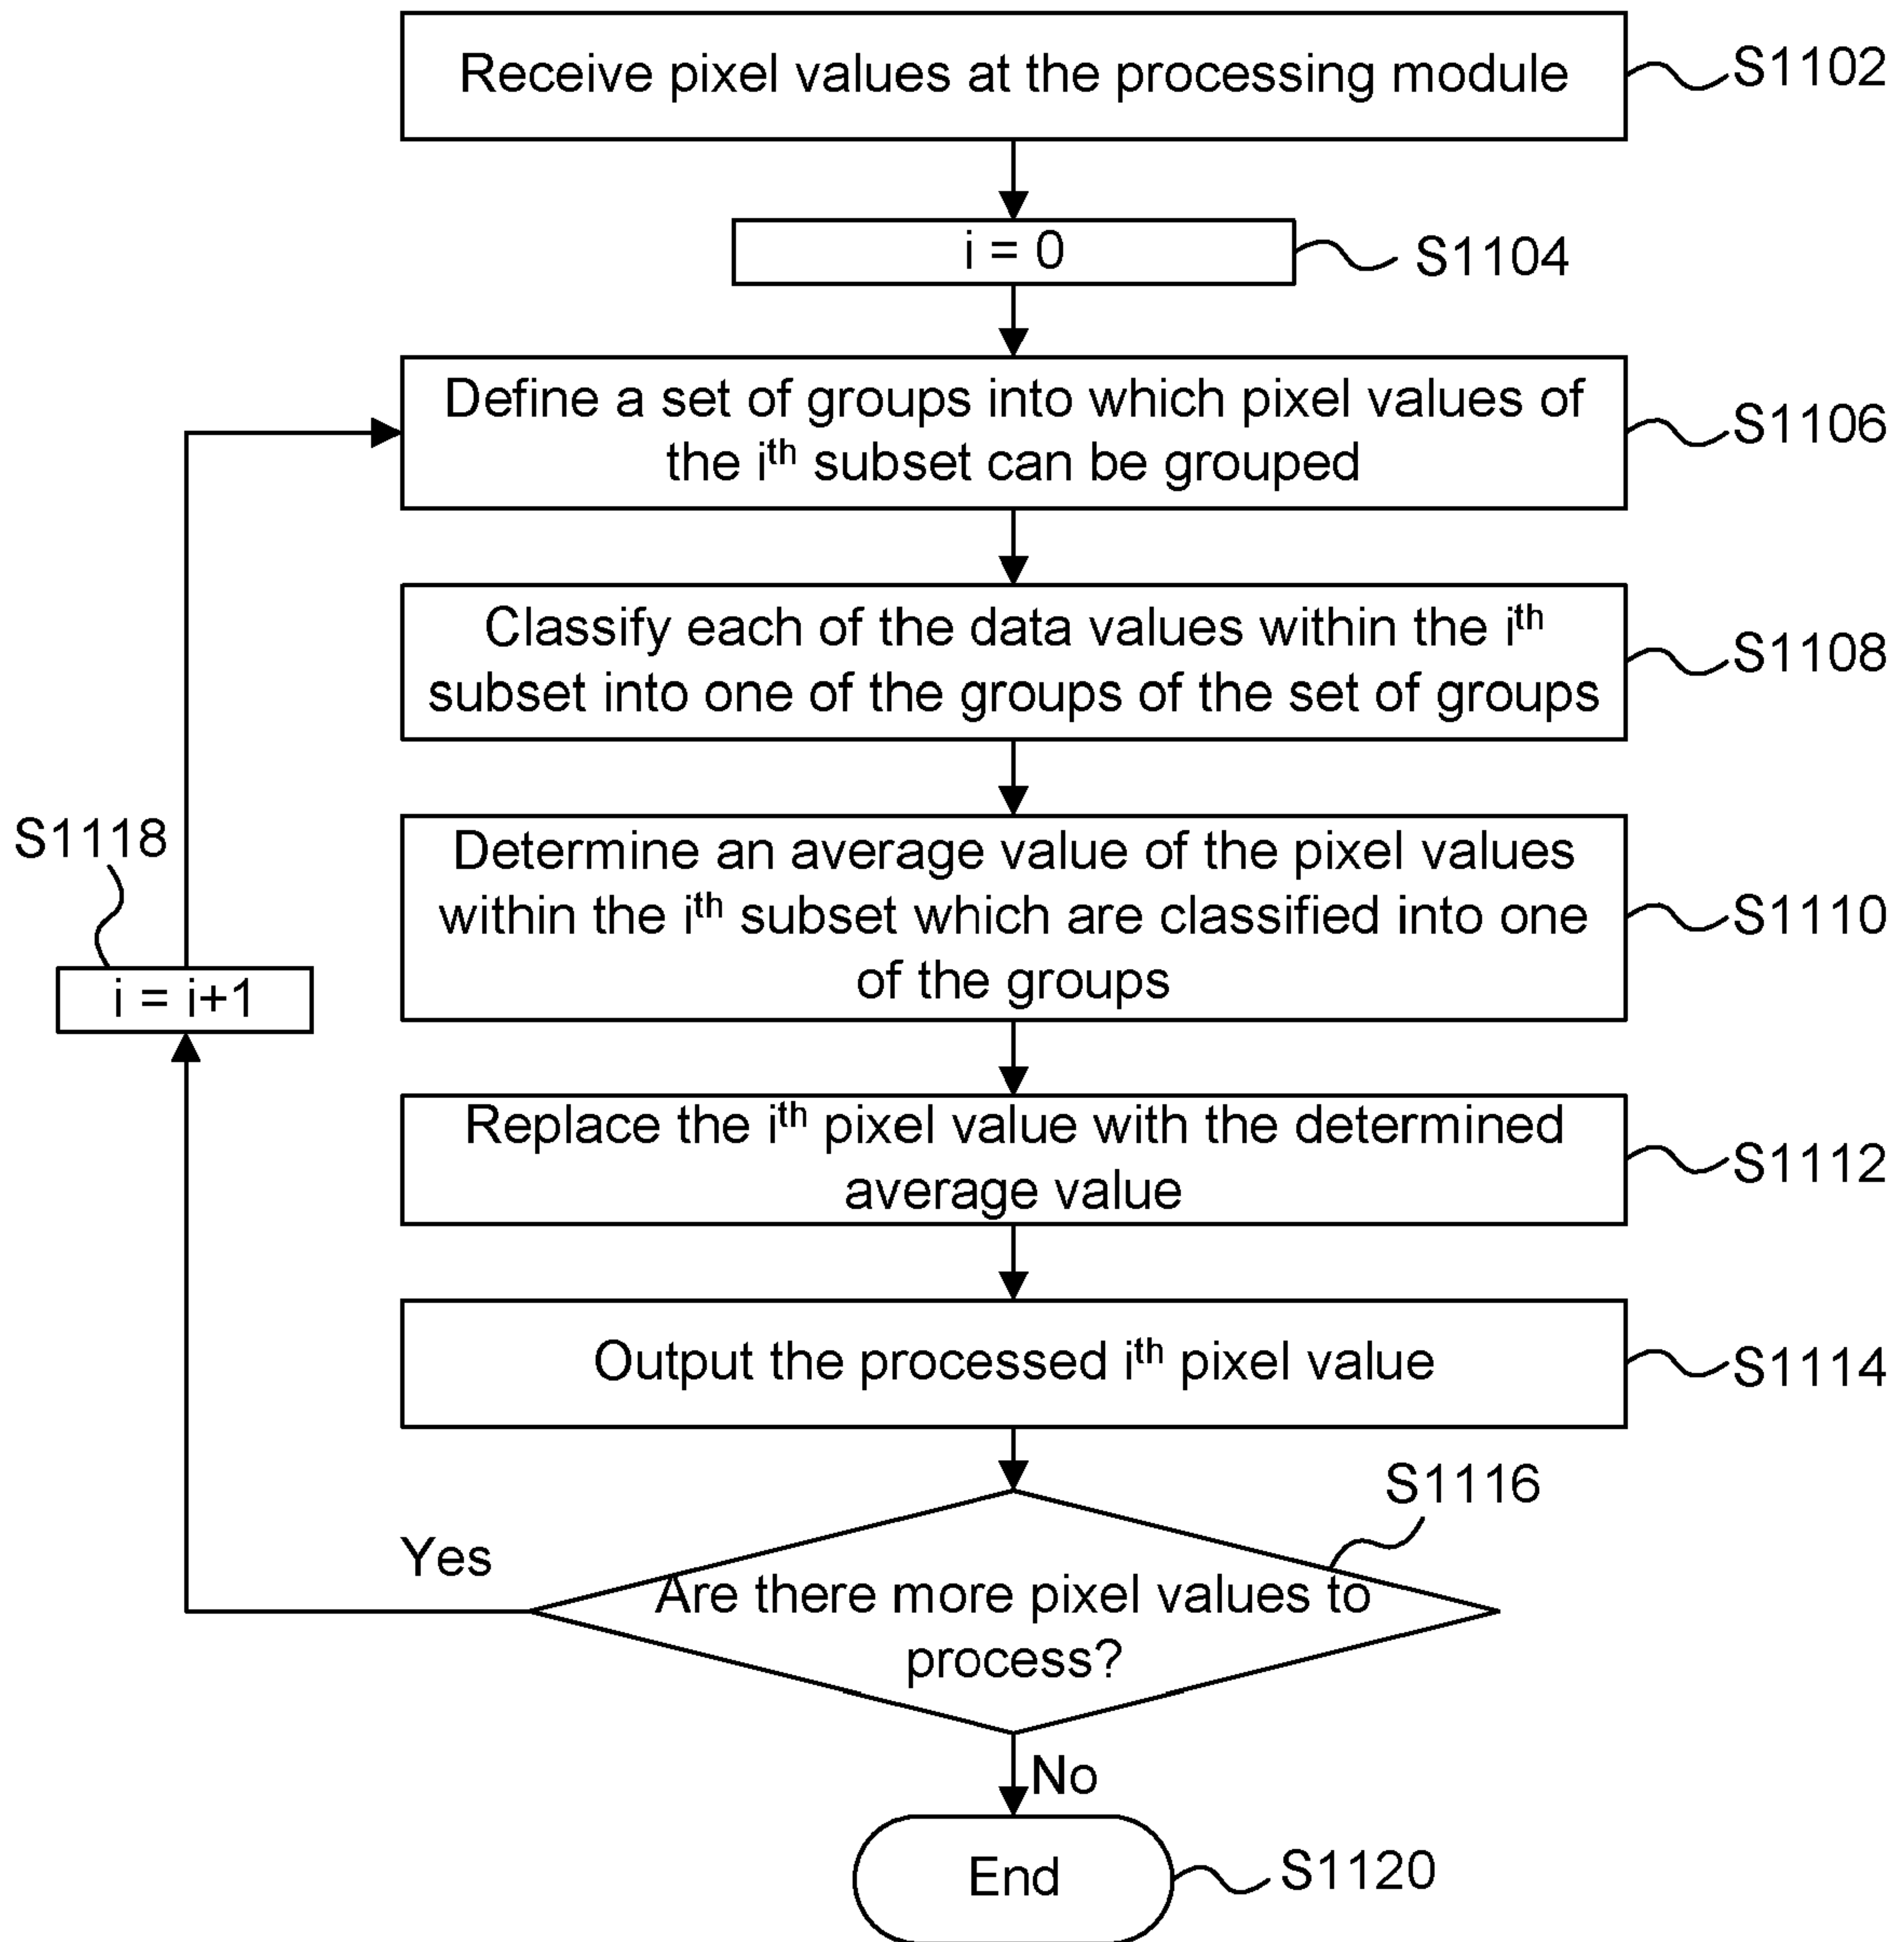


FIGURE 11

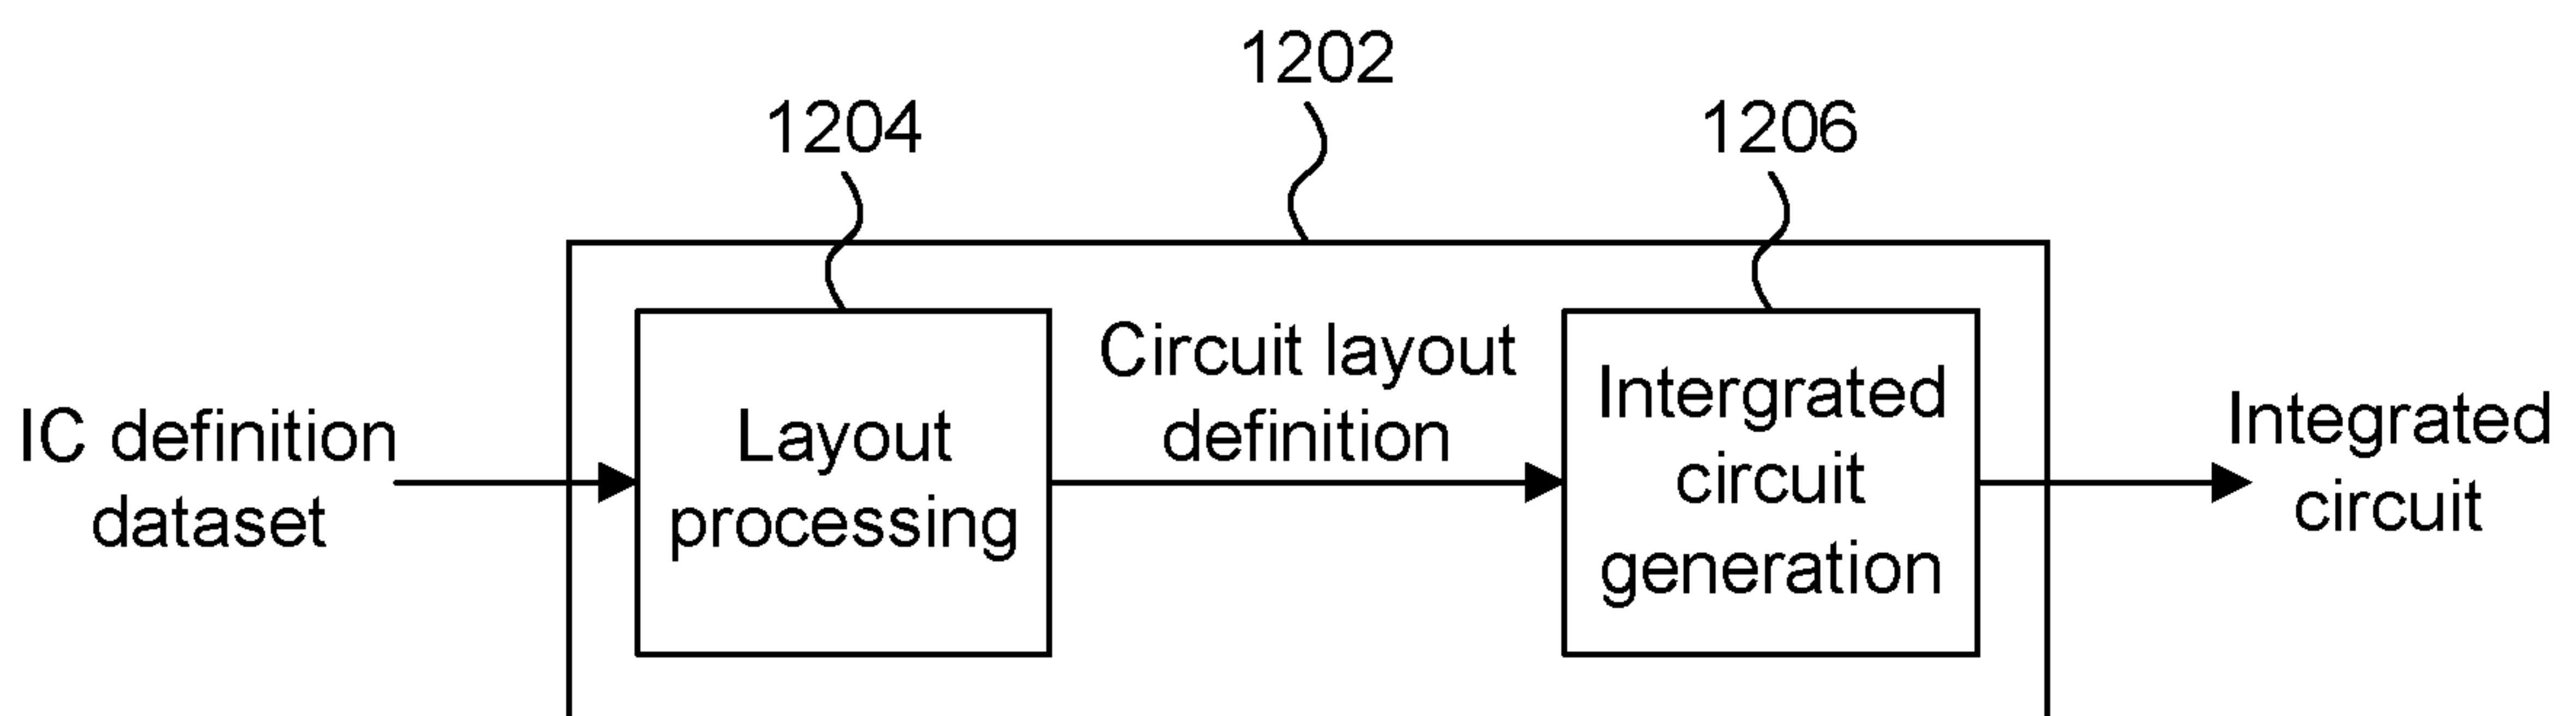


FIGURE 12



The following terms are registered trade marks and should be read as such wherever they occur in this document:

Java (Page 47)

OpenCL (Page 47)

SYSTEMS AND METHODS FOR PROCESSING A STREAM OF DATA VALUES

Background

- 5 There are many different types of data processing systems which can process a stream of data values. For example, an image processing system may be used as a camera pipeline to process pixel values originating from image sensors in a camera to provide a set of processed pixel values representing a captured image.
- 10 An image processing system, such as a camera pipeline, can be arranged to receive a stream of pixel values (e.g. in raster scan order, or any other predetermined order, such as Boustrophedon order) and perform multiple processing functions on the pixel values in order to determine processed pixel values which can then be output, e.g. for display on a display or for storage in a
- 15 memory or for transmission to another device, e.g. over a network such as the Internet. Figure 1a illustrates an image processing system 104 which is arranged to receive image data (e.g. pixel values) from an image sensor 102. As is known in the art, the image sensor 102 can include an array of sensors (e.g. charge-coupled devices (CCDs) or complementary metal–oxide–semiconductor (CMOS)
- 20 sensors) which can sense incident light at respective pixel positions to determine raw image data (e.g. unprocessed pixel values) which can be fed into the image processing system 104. In this way the image sensor 102 captures light and converts it into electrical signals (e.g. image data in the form of pixel values) to be processed in order to generate an image. The image processing system 104
- 25 comprises a number of distinct processing modules (e.g. Figure 1 illustrates k processing modules $106_1, 106_2, 106_3, \dots, 106_k$) which implement respective processing functions on the image data to determine the processed image data which represents an image and which is output from the image processing system 104. For example, the image processing system 104 may modify the pixel values
- 30 (e.g. to improve the perceptual quality of the image) using functions such as filtering and sharpening. Other functions of image processing may be to compress or encode image data into particular formats, such as the Joint Photographic Experts Group (JPEG) format or the Moving Picture Experts Group (MPEG) format.

An image processing system may be implemented in a pipeline arrangement, whereby multiple processing modules of the image processing system perform respective functions on pixel values representing an image. The processing modules 106_1 to 106_k operate on pixel values in turn, in a sequence, and may operate independently of each other. There are many different processing functions which may be performed by respective ones of the processing modules 106. Some examples of the different processing functions which may be applied by a processing module are: adjusting gains applied to the pixel values, adjusting the white balance of the pixel values, detecting defective pixels within the image, correcting defective pixels, filtering (e.g. denoising), sharpening and demosaicing, to give just a few examples. A person skilled in the art would be aware of many different functions that may be performed by different processing modules within the image processing system 104. The exact processing modules to implement, their configuration and the order in which they are implemented in the image processing system 104 (e.g. for use in a camera pipeline) are design choices which are made when the image processing system 104 is designed.

Some of the processing modules operate on individual pixel values, without needing to consider other pixel values. For example, the adjustment of gains and the adjustment of the white balance can be performed for a pixel value without considering other pixel values. In contrast, some of the other processing modules operate on a kernel (i.e. a block or subset) of pixel values. For example, 2D filtering operations such as denoising, defective pixel detection and correction operations, sharpening operations and demosaicing operations operate on a kernel of pixel values in order to determine how to process a particular pixel value. For example, Figure 1b shows a kernel 120 of pixel values of an image. The kernel 120 includes forty nine pixel values, arranged as a 7x7 block of pixel values, centred on a pixel value 122, which is a current pixel value being processed. Since, in a camera pipeline the pixel values are received row-by-row (e.g. in raster scan order), in order for a processing module 106 to have access to a block of pixel values including pixel values from multiple rows, a line store module may be implemented in the image processing system 104 prior to a processing module 106. Each row of pixel values may for example include

hundreds or thousands of pixel values. A line store module can store multiple rows of pixel values, and can provide a contiguous block of pixel values (e.g. the block 120) to the processing module 106 at a time. A line store module is a large block to implement in hardware because it typically needs to be able to store
5 thousands of pixel values at a given time.

Many of the processing functions which may be implemented by the processing modules 106 on blocks of pixel values may involve performing some sort of edge analysis of the image represented by the pixel values within a block. Some
10 examples of edge detection algorithms (such as Sobel edge detection and Canny edge detection) use differential methods. It is noted that the detection of edges within a kernel is super-linearly hard, e.g. the Canny algorithm has a complexity $\mathcal{O}(n^2 \log n)$ for a kernel containing n pixels. For example, a denoising filtering function may apply spatial smoothing to the pixel values within a kernel.

15 Denoising is a known process in image processing. For example, the pixel value 122 may be filtered based on other pixel values within the kernel 120, e.g. the pixel value 122 may be smoothed so that it is based on a weighted sum of a plurality of the pixel values within the kernel 120 (e.g. where the weights depend upon the distance between the pixel value 122 and the pixel value being
20 summed). In a simple example, the pixel value 122 may be based on a mean of a plurality of the pixel values within the kernel 120. Spatial smoothing helps to reduce the appearance of random noise in the image. However, it can be perceptually detrimental to an image to smooth pixel values over a distinct edge in an image. This can have the effect of blurring the edge in the image. So a
25 processing module 106 implementing a denoising filter will attempt to determine the presence of an edge within the kernel 120. There are many different algorithms for analysing a kernel of pixel values to determine an edge. The example shown in Figure 1b shows an edge 124 between two materials (A and B) in the scene being imaged, where the edge 124 passes through the kernel 120.
30 The pixels in the kernel 120 have been labelled as either "A" or "B" to represent which material is visible in the image at the respective pixel position. The denoising applied to pixel value 122 may be implemented as a weighted sum of (or an average of) only the pixel values labelled A (and not including the pixel values labelled B), where the weights may depend upon the distance to the pixel

122. In this way the processing module 106 can perform a bilateral filtering process to implement the denoising.

Similarly, a processing module 106 which implements sharpening will identify the edge 124 in the kernel 120 and will apply sharpening to pixel values near the edge to increase the differences in pixel values on opposing sides of the edge 124, thereby sharpening the appearance of the edge 124. Sharpening is a known process in image processing.

Defective pixel detection is a known process in image processing. A processing module 106 which implements defective pixel detection on the pixel value 122 aims to detect whether the pixel value is representative of a defective pixel by identifying whether it is significantly and unexpectedly different to nearby pixel values. In order to do this the processing module 106 will identify the edge 124, so that it can compare the pixel value 122 to other nearby pixel values which are expected to be similar because they relate to the same material in the scene (material A).

Defective pixel correction is a known process in image processing. Defective pixel correction may be implemented in the same, or a separate, processing module to one which implements defective pixel detection. If the pixel value 122 has been identified as being representative of a defective pixel, then a processing module 106 which implements defective pixel correction on the pixel value 122 aims to determine a suitable value to assign to the pixel value to replace the defective pixel value. For example, an average (e.g. the mean or the median) of the pixel values within the kernel 120 which represent the same material (material A) could be used to replace the defective pixel value 122. In order to do this the processing module 106 will identify the edge 124 so that it knows which pixel values within the kernel 120 are representative of the same material (material A) as is supposed to be represented by pixel value 122.

Demosaicing is a known process in image processing. A processing module 106 which implements demosaicing aims to reconstruct a full colour image from incomplete colour samples which are output from the image sensor 102.

Interpolation may be used for demosaicing, but similar to denoising, the interpolation is only performed over pixel values which are determined to be representing the same material. So in order to perform demosaicing for the pixel value 122, a processing module 106 will identify the edge 124 so that it knows which pixel values within the kernel 120 are representative of the same material (material A) and can therefore be used in the interpolation of the demosaicing process.

Typically, edge detection methods for detecting edges within a kernel can only detect a single edge within the kernel.

Summary

This Summary is provided to introduce a selection of concepts in a simplified form that are further described below in the Detailed Description. This Summary is not intended to identify key features or essential features of the claimed subject matter, nor is it intended to be used to limit the scope of the claimed subject matter.

There are a number of problems associated with processing a stream of pixel values in a camera pipeline using a sequence of processing modules, such as described in the background section above and shown in Figure 1a.

For example, since each of the processing modules 106 in the image processing system 104 operates independently, the different processing modules 106 may be designed separately. It is possible that inconsistent edge determinations may be made in the different processing modules. Inconsistent edge determinations across the different processing modules can result in artefacts appearing in the image.

30

As mentioned above, standard methods for detecting edges within a kernel can only detect a single edge within the kernel. There may be situations in which a kernel includes more than one edge, and standard edge detection algorithms might not be able to detect all of the edges within the kernel. An alternative

12 06 18

approach is to use a clustering approach in which data values within a kernel are grouped (or “clustered”) into one or more groups. Clustering algorithms are able to group the data values within a kernel into more than two groups, and can cope with corners within the kernel. However, clustering is a problem which is known to be “NP hard”, i.e. non-deterministic polynomial-time hard. As such it is difficult to implement a clustering algorithm multiple times in multiple processing modules for each pixel value of an image in a camera pipeline in real-time as images are captured by the image sensor 102. In fact, it is so difficult that a clustering algorithm has not been used in this way in the prior art. As an example, current state of the art cameras tend to be configured to capture images with many millions of pixels.

Furthermore, different processing modules might be implementing functions with at least partially contradictory aims. For example, a denoising function will tend to reduce differences between pixel values, whereas a sharpening function will tend to increase differences between pixel values. If each block is implemented separately it can be difficult to achieve a suitable balance between the contradictory processes, which may tend to fight against each other. Therefore, tuning of the image processing system 104 can be difficult because changes to one processing module 106 may have unforeseen consequences on the operation of another processing module 106.

According to examples described herein, there is provided a data processing system configured to process a stream of data values, the data processing system comprising:

a store configured to store group indication data which indicates one or more groups into which data values can be grouped; and

a processing module configured to process, in each of a plurality of iterations, a respective particular data value of the stream by operating on a respective particular subset of data values of the stream, by, in each of the plurality of iterations:

retrieving, from the store, group indication data for at least one group defined for a previous subset of data values;

using the retrieved group indication data to define, for the particular subset of data values in the current iteration, a set of groups into which data values within the particular subset can be grouped;

classifying each of the data values within the particular subset into one of the groups of the set of groups;

processing the particular data value using one or more of the data values of the particular subset in dependence on the classification of the data values of the particular subset into the groups; and

storing, in the store, group indication data for at least one group of the set of groups defined for the particular subset of data values in the current iteration, for use in a subsequent iteration;

wherein the processing module is configured to output the processed data values.

Each of the data values within the particular subset may be classified into one of the groups of the set of groups based on the value of that data value. The group indication data for a group may indicate a range of data values which are to be classified into the group. For example, the group indication data for a group may comprise an indicative value within the range of data values which are to be classified into the group. The extent of the range of data values around the indicative value may be defined by a calibration process. The calibration process may be a static calibration process, a dynamic calibration process or a real-time calibration process. Said storing group indication data for at least one group may comprise storing, as the indicative value for a group, an average value of the data values which are classified as being part of that group.

The group indication data for a group may comprise indications of the upper and lower bounds of the range of data values which are to be classified into the group.

The stream of data values may represent a one-dimensional array of data values. For example, each of the subsets of data values within the stream of data values may represent a contiguous block of data values within the one-dimensional array. For example, the data processing system may be an audio processing system, and wherein the data values may be audio samples of an audio signal. As another

example, the data processing system may be a signal processing system, and the data values may be signal samples of a transmitted signal.

In other examples, the stream of data values represents a two-dimensional array of data values. For example, each of the subsets of data values within the stream of data values may represent a contiguous block of data values within the two-dimensional array. For example, the data processing system may be an image processing system, wherein the data values may be pixel values.

Said processing the particular data value may comprise performing, in a consolidated operation, multiple processing functions which depend upon the classification of data values of the particular subset into the groups. For example, the multiple processing functions may comprise two or more of: (i) defective pixel detection, (ii) defective pixel correction, (iii) denoising, (iv) sharpening, and (v) demosaicing.

The group indication data may be stored in, and retrieved from, the store for up to m groups, and the defined set of groups into which data values within the particular subset can be grouped may comprise up to n groups, where $n \geq m$.

The processing module may be configured to, in each of the plurality of iterations: after the classification of the data values of the particular subset into the groups, and prior to said processing the particular data value, selectively discard one or more of the groups based on the number of data values which are classified into each of the groups, such that up to m groups are maintained. The defined set of groups into which data values within the particular subset can be grouped may comprise: the groups for which group indication data is retrieved from the store, and groups defined by gaps around the groups for which group indication data is retrieved from the store. In some examples, $m \leq n \leq 2m + 1$. In a particular example, $m = 3$ and $n = 7$.

30

Said processing the particular pixel value may comprise: if the particular pixel value is within a maintained group, determining that the particular pixel value is not representative of a defective pixel; and if the particular pixel value is not within a

maintained group, determining that the particular pixel value is representative of a defective pixel.

The processing module may be configured to, in each of the plurality of iterations:

5 determine, for each of at least one group of the set of groups, a centroid representative of a position of the group; and determine, for each of the at least one group of the set of groups, an average value of the data values within the group. For example, where the data values are pixel values, said processing the particular pixel value may comprise applying one or more of denoising, sharpening
10 and demosaicing to the particular pixel value based on the at least one determined centroid of the at least one group. Processing the particular pixel value may comprise, if the particular pixel value is determined to be representative of a defective pixel: selecting one or more of the maintained groups based on the determined centroids of the groups, and replacing the particular pixel value with a
15 value based on the determined one or more average values for the one or more selected groups.

There is provided a method of processing a stream of data values in a data processing system, the method comprising:

20 processing, in each of a plurality of iterations, a respective particular data value of the stream by operating on a respective particular subset of data values of the stream, by, in each of the plurality of iterations:

retrieving group indication data which indicates at least one group defined for a previous subset of data values;

25 using the retrieved group indication data to define, for the particular subset of data values in the current iteration, a set of groups into which data values within the particular subset can be grouped;

classifying each of the data values within the particular subset into one of the groups of the set of groups;

30 processing the particular data value using one or more of the data values of the particular subset in dependence on the classification of the data values of the particular subset into the groups; and

storing group indication data for at least one group of the set of groups defined for the particular subset of data values in the current iteration, for use in a subsequent iteration;

wherein the method further comprises outputting the processed data values.

In some examples, there is provided an image processing system configured to process a stream of pixel values, the image processing system comprising:

a processing module configured to implement processing of a particular pixel value of the stream by operating on a particular subset of pixel values of the stream, by:

classifying each of the pixel values within the particular subset into a group of a set of groups;

processing the particular pixel value using one or more of the pixel values of the particular subset in dependence on the classification of the pixel values of the particular subset into the groups, wherein said processing the particular pixel value comprises performing, in a consolidated operation, multiple processing functions which depend upon the classification of pixel values of the particular subset into the groups; and outputting the processed particular pixel value.

In some examples, there is provided an image processing system configured to process pixel values, the image processing system comprising a processing module configured to:

receive a plurality of pixel values; and

implement processing of a particular pixel value by operating on a particular subset of the received pixel values, by:

defining a set of one or more groups into which pixel values within the particular subset can be grouped;

classifying each of the pixel values within the particular subset into one of the groups of the set of one or more groups based on the value of that pixel value;

processing the particular pixel value using one or more of the pixel values of the particular subset in dependence on the classification of the pixel values of the particular subset into the one or more groups; and outputting the processed particular pixel value.

5

In some examples, there is provided an image processing system configured to process pixel values, the image processing system comprising a processing module configured to:

10 receive a plurality of pixel values, each of the received pixel values having a first number of bits; and

implement processing of a particular pixel value by operating on a particular subset of the received pixel values, by:

classifying each of the pixel values within the particular subset into a group of a set of one or more groups;

determining an average value in respect of the pixel values within the particular subset which are classified into one of the one or more groups, wherein the determined average value has a second number of bits, wherein said second number is greater than said first number;

replacing the particular pixel value based on the determined average value; and

outputting the processed particular pixel value.

10

The data processing systems and image processing systems described herein may be embodied in hardware on an integrated circuit. There may be provided a method of manufacturing, at an integrated circuit manufacturing system, a data processing system and/or an image processing system as described herein.

15

There may be provided an integrated circuit definition dataset that, when processed in an integrated circuit manufacturing system, configures the system to manufacture a data processing system or an image processing system as described herein. There may be provided a non-transitory computer readable storage medium having stored thereon a computer readable description of an integrated circuit that, when processed, causes a layout processing system to generate a circuit layout description used in an integrated circuit manufacturing system to manufacture a data processing system or an image processing system as described herein.

20

25

There may be provided an integrated circuit manufacturing system comprising:

a non-transitory computer readable storage medium having stored thereon a computer readable integrated circuit description that describes a data processing system or an image processing system as described herein;

30

a layout processing system configured to process the integrated circuit description so as to generate a circuit layout description of an integrated circuit embodying the data processing system or the image processing system as described herein; and

an integrated circuit generation system configured to manufacture the data processing system or the image processing system according to the circuit layout description.

5 Although methods described herein are best suited to being implemented in dedicated hardware, the methods could be implemented in software running on general purpose hardware. Therefore, there may be provided computer program code for performing any of the methods described herein. There may be provided non-transitory computer readable storage medium having stored thereon computer
10 readable instructions that, when executed at a computer system, cause the computer system to perform any of the methods described herein.

The above features may be combined as appropriate, as would be apparent to a skilled person, and may be combined with any of the aspects of the examples
15 described herein.

Brief Description of the Drawings

Examples will now be described in detail with reference to the accompanying
20 drawings in which:

Figure 1a shows an example image processing system of the prior art;
Figure 1b shows a 2D array of pixel values representing an image;
Figure 2 shows an image processing system according to examples described
25 herein;
Figure 3 is a flow chart for a method of processing a stream of data values;
Figure 4 shows a 2D array of data values, highlighting a first subset of the data values which will be used for processing a first data value in the 2D array;
Figure 5 shows a 2D array of data values, highlighting a second subset of the data
30 values which will be used for processing a second data value in the 2D array;
Figure 6 illustrates ranges of data values spanned by different groups of data values;
Figure 7 shows a calibration curve used to determine the extent of a range of a group as a function of the average value of the group;

Figure 8 is a flow chart showing some of the steps involved in processing a pixel value in examples described herein;

Figure 9a shows a representation of a first example image which is received at a processing module;

5 Figure 9b shows a representation of the first example image when it is output from the processing module;

Figure 10a shows a representation of a second example image which is received at the processing module;

10 Figure 10b shows a representation of the second example image when it is output from the processing module;

Figure 11 is a flow chart for a method of implementing dynamic range enhancement; and

Figure 12 shows an integrated circuit manufacturing system for generating an integrated circuit embodying a data processing system.

15

The accompanying drawings illustrate various examples. The skilled person will appreciate that the illustrated element boundaries (e.g., boxes, groups of boxes, or other shapes) in the drawings represent one example of the boundaries. It may be that in some examples, one element may be designed as multiple elements or that
20 multiple elements may be designed as one element. Common reference numerals are used throughout the figures, where appropriate, to indicate similar features.

Detailed Description

25 The following description is presented by way of example to enable a person skilled in the art to make and use the invention. The present invention is not limited to the embodiments described herein and various modifications to the disclosed embodiments will be apparent to those skilled in the art.

30 Embodiments will now be described by way of example only. According to embodiments described herein, a stream of data values (e.g. pixel values) is processed in a data processing system (e.g. an image processing system such as a camera pipeline). As an example, processed pixel values may represent an image.

It has been appreciated that where a processing module processes a particular data value from the stream by operating on a respective subset of data values surrounding the particular data value (e.g. a kernel of data values centred on the particular data value that is currently being processed), the subset of data values which are operated on for a particular data value in the stream will at least partially overlap with the subset of data values which are operated on for another (e.g. the next) data value in the stream. Therefore, some information relating to the subset of data values for processing one data value in the stream is relevant to, and can be used for, the subset of data values for processing the next data value in the stream. In other words, information from processing a current data value in the stream (e.g. group indication data which indicates one or more groups into which data values can be grouped) can be used to simplify the task of processing the next data value in the stream (e.g. for defining a set of groups into which data values can be grouped). In this way, the processing involved in performing clustering on the data values within a subset of data values can be shared between consecutive data values in the stream. This reduces the latency and power consumption of the image processing system, in contrast to prior art systems in which each data value in the stream is processed separately.

In particular, the Bayesian principle of using prior knowledge can simplify the clustering problem. Groups which are determined for clustering of data values within a previous subset (the preceding subset) are likely to be similar to the groups which are determined for clustering of data values within a current subset because the two subsets partially overlap with other (e.g. the two subsets may substantially overlap with each other). Therefore, we can use the previous groups defined for the previous data value as a starting point for defining the groups to use for the current data value. When the current data value has been processed, group indication data for the groups that have been used for the current data value can be stored, so that in the next iteration (i.e. when processing the next data value of the stream) the group indication data can be retrieved and used to define the groups into which data values within the next subset can be grouped. This greatly simplifies the clustering problem so that it is deterministically solvable in real-time (i.e. it is no longer an NP hard problem) even on devices which have tight

restrictions on processing power and size (e.g. mobile devices such as handheld cameras, smart phones and tablets, etc.).

Therefore, in some examples described herein, an edge detection algorithm is not
5 used to identify groups of data values within a kernel, and instead a clustering
algorithm is used to identify the groups of data values within the kernel. Edge
detection usually results in a binary split of the data values within a kernel into two
groups: one group on either side of the detected edge. For example, as described
above, Figure 1b shows the pixel values being classified as either A or B based on
10 their positions relative to the edge 124. In this way, the edge 124 is thinned (or
thresholded) down to a single pixel width. However, in reality, edges in images
are often soft, and may have finite widths such that some pixels can lie within an
edge rather than on one side or the other, such that a set of “fuzzy”, or half-toned,
pixel values may exist along the edge. Edge detectors will split these edge pixel
15 values into one of the two groups defined by the edge. Making the wrong choice
or ascribing an intermediate value (lying on or near the edge 124) to either group
A or B (as shown in Figure 1b) may lead to mis-classification of pixels, which in
turn may lead to poor performance and the introduction of errors along or near the
edge 124 in the processed image data. However, in contrast, a clustering
20 algorithm may generate a separate group for the intermediate pixel values which
are between the values of pixels within groups A and B. In particular, a clustering
algorithm can easily group the pixel values within a kernel into more than two
groups. Furthermore, in a clustering algorithm, the grouping of the pixel values
within a kernel does not depend upon the spatial position of the pixel values within
25 the kernel. In this sense the clustering algorithm is free to classify a pixel value
that is within a kernel to a group without considering the spatial position of the
pixel value within the kernel. Instead, a set of groups are defined into which pixel
values within a particular subset can be grouped, and then each of the pixel values
within the particular subset are classified into one of the groups based on the
30 value of that pixel value. For example, the groups may be defined by defining a
range of pixel values for each of the groups, wherein each of the pixel values
within a particular subset is grouped into one of the groups by comparing the pixel
value to one or more of the defined pixel value ranges for the groups.

The clustering approach used in examples described herein can cope with: (i) any angle of edge (i.e. the angle of an edge is not considered), (ii) corners (e.g. if a quadrant within the kernel is different to the rest of the kernel), (iii) more than one edge within the kernel, and (iv) soft edges (e.g. edges with a finite width which may span over one or more pixels, such that some pixels may be located within an edge). This is in contrast to standard edge detectors which cannot cope with any of the four situations mentioned above.

Furthermore, since data is shared between consecutive data values in the stream, it is possible to use a smaller kernel size and still retain good processing results compared to prior art systems in which each data value in the stream is processed separately. For example, a 5x5 kernel could be used rather than a 7x7 kernel, which approximately halves the number of data values in the kernel. A reduction in the size of the kernels tends to reduce the silicon area of the image processing system and reduce the processing power consumed by the image processing system.

Furthermore, in examples described herein, multiple processing functions (which would typically be performed by separate processing modules) which depend upon the classification of pixel values of a subset into groups, are performed in a consolidated operation (e.g. by a single processing module). This means that the clustering process only needs to be performed once for the multiple processing functions. The multiple processing functions may for example comprise two or more of: (i) defective pixel detection, (ii) defective pixel correction, (iii) denoising, (iv) sharpening, and (v) demosaicing. Consolidating multiple functions can reduce the number of processing modules and the number of line store modules in the image processing system, which will reduce the silicon area and power consumption of the image processing system. Furthermore, by applying the multiple processing functions in a consolidated operation (e.g. with a single processing module) it is simpler to tune the operation of the functions compared to when the different functions are implemented by separate blocks. The tuning is particularly simplified for functions which may tend to fight against each other, such as denoising and sharpening. The tuning is also simplified for functions which replicate some functionality, such as defective pixel detection and defective

pixel correction. Furthermore, by performing, in a consolidated operation, multiple processing functions which depend upon the classification of data values of the particular subset into the groups, the classification of the data values into the groups is consistent for the multiple processing functions, thereby avoiding
5 artefacts which may occur due to inconsistent clustering decisions for different functions.

Furthermore, in some examples, a pixel value can be replaced with an average of pixel values within a group of similar pixel values. This averaging process can be
10 used to enhance (i.e. increase) the dynamic range of the pixel values. For example, an average value can be determined with more bits than the number of bits in each of the individual input pixel values.

The processing modules in the image processing system may be implemented in
15 hardware, e.g. in fixed function circuitry. Dedicated fixed function hardware is typically more efficient at processing data values (in terms of processing speed and power consumption) compared to implementing the processing modules by running software on general purpose hardware. Implementing the processing modules in hardware is not as flexible as implementing the processing modules in
20 software in the sense that the operation is difficult (or impossible) to change after the hardware has been manufactured, whereas software can easily be changed by loading different software code into a memory storing the software to be executed. However, in a data processing system such as a camera pipeline or other system which processes a stream of data values in a fixed manner in real-time, the
25 advantages of implementing the processing modules in hardware (in terms of processing speed and power consumption) tend to outweigh the disadvantage of the inflexibility of a hardware implementation. By implementing multiple processing functions in a single consolidated operation, the number of processing modules is reduced in the image processing system, thereby significantly reducing
30 the size (in terms of silicon area) of the hardware implementing the data processing system. For example, four or five distinct processing modules in a prior art system could be consolidated down to just one processing module in examples described herein.

Examples are described below with reference to an image processing system, but concepts described herein could be applied to other types of data processing systems (e.g. audio processing systems and signal processing systems) which process streams of data values. Figure 2 illustrates an image processing system 204 which is arranged to receive image data (e.g. pixel values) from an image sensor 202. The image data received from the image sensor 202 may be referred to as “raw” image data or “unprocessed” image data. As described above, an image sensor, such as the image sensor 202, can include an array of sensors (e.g. charge-coupled devices (CCDs) or complementary metal–oxide–
10 semiconductor (CMOS) sensors) which can sense incident light at respective pixel positions to determine raw image data (e.g. unprocessed pixel values) which can be fed into the image processing system 204. The image processing system 204 could be described as a camera pipeline. The system shown in Figure 2 may be implemented in a device, such as a camera. In some examples, the image sensor
15 202 may capture a sequence of images, which can be used to represent frames of a video sequence. The image data may be in one of many different possible formats, for example the pixel values may be monochrome pixel values, wherein each pixel value relates to the same, single colour channel. In other examples, the pixel values may comprise Red, Green and Blue (RGB) values to represent
20 the intensity of the respective three colour channels at the respective pixel positions. In some examples, the raw image data may relate to multiple colour channels, but the raw image data for each particular pixel position may include only a single colour value (e.g. a red, green or blue colour value), arranged in a pattern such as a Bayer pattern, and in these examples a demosaicing process
25 may be performed such that a full set of colour values (e.g. three colour values, such as red, green and blue colour values) are determined for each pixel position. A demosaicing process is described below. Other colour formats may be used in other examples. The image sensor 202 may provide the pixel values to the image processing system 204 in a raster scan order, such that pixel values for lines of
30 pixels (e.g. rows of pixels) are received and processed in the image processing system 204 as they are received from the image sensor 202 in real-time.

The image processing system 204 may comprise a number of processing modules 206, 210 and 214 for applying respective processing functions to the image data.

The processing modules may, for example, be implemented in dedicated hardware modules (e.g. in fixed function circuitry) such that they can be optimised for performing specific processing functions on the pixel values in an efficient manner (e.g. “efficient” in terms of at least one of latency, power consumption and silicon area). The image processing system 204 also comprises a store 212 which is coupled to the processing module 210. As described in more detail below, the store 212 is configured to store group indication data which indicates one or more groups into which pixel values can be grouped. The processing module(s) 206 and 214 are represented with dashed boxes because the number of these modules may vary in different implementations and that number may be zero in some implementations. The processing module(s) 206 and 214 may perform any suitable functions on pixel values, e.g. adjusting gains of the pixel values, adjusting the white balance of the pixel values, etc. If a processing module requires access to pixel values from more than one line of the image (e.g. a 2D filtering module) then a line store module (or “line store bank”) can be implemented in the image processing system 204 prior to the processing module. For example, processing module 210 requires access to multiple rows of pixel values, e.g. the processing module 210 may be configured to implement a 2D filtering process, defective pixel detection, defective pixel correction, sharpening and/or demosaicing. As described above, in these types of processes, a particular pixel value can be processed based on the pixel values within a kernel of pixel values including the particular pixel value (e.g. centred on the particular pixel value). The line store module 208 is implemented to store pixel values for multiple lines of pixels and can provide pixel values from multiple lines to the processing module 210. Processed image data is output from the image processing system 204, and may be used in any suitable manner, e.g. output to a display, stored in memory, transmitted to another device, etc.

Operation of the image processing system 204 is described with reference to the flow chart shown in Figure 3. The image sensor 202 captures image data, e.g. pixel intensity values for different pixels within an image. Techniques for capturing these pixel intensity values at the image sensors are known in the art and the details of this process are beyond the scope of this description. In step S302 the image processing system 204 receives a stream of data values from the image

sensor 202. The pixel values are fed from the image sensor 202 to the image processing system 204 in a particular order, thereby forming a stream of pixel values, e.g. raster scan order in which rows of pixel values are received in turn at the image processing system 204. It would be possible for the pixel values to be passed to the image processing system 204 in a different order, e.g. Boustrophedon order, and it is possible for multiple lines of pixel values to be passed to the image processing system 204 at once, in parallel. However, since the image sensor 202 and the image processing system 204 may be designed by different parties, it is convenient for there to be a presumed, or standard, order in which pixel values are passed from the image sensor 202 to the image processing system 204, and raster scan order is the usual order used. Therefore, in the examples described herein we refer to receiving rows of pixel values in a raster scan order, but it should be apparent that other orders could be used, and that columns of pixel values could be received rather than rows, so where we refer to “rows” of pixel values it is to be understood that we are generally referring to “lines” of pixels values, which could for example be rows or columns.

As described above, the image processing system 204 may, or may not, perform some processing on the pixel values in the processing module(s) 206 before the pixel values are stored in the line store module 208. As described above, the line store module 208 allows pixel values from multiple rows to be provided to the processing module 210 together.

Figure 4 shows an example in which the stream of pixel values represents a two-dimensional array of pixel values 402 forming an image, wherein the array includes six rows of pixel values, with each row including seven pixel values. Obviously, in more complex examples, there may be many more pixel values within an image, e.g. millions of pixel values arranged such that there are hundreds or thousands of rows and hundreds or thousands of columns of pixels.

The processing module 210 processes a particular pixel value 404 by operating on the pixel values within the 5x5 kernel 406, which surrounds and includes pixel value 404 (e.g. the kernel 406 is centred on pixel value 404). In step S304, the processing module 210 receives the twenty five pixel values in the kernel 406 from

the line store module 208. The kernel 406 defines a subset of pixel values which represents a contiguous block of pixel values within the two-dimensional array of pixel values representing the image.

- 5 In step S305 a subset index i is set such that $i = 0$, for the processing of a first subset (e.g. the kernel 406).

In step S306, the processing module 210 defines a set of one or more groups into which pixel values of a first subset (i.e. pixel values within the kernel 406) can be
10 grouped. This can be done by analysing the pixel values of the subset 406 and performing a clustering process to group the pixel values into groups based on their pixel values. This may involve an edge detection process to identify edges between different materials (i.e. sharp changes in appearance) in the image. Figure 4 shows two edges (408 and 410) in the scene, such that three distinct
15 areas are formed within the kernel 406. Therefore, three groups are defined (A, B and C).

In step S308 the processing module 210 classifies each of the pixel values of the subset 406 into one of the groups (A, B or C) of the set of groups. In Figure 4
20 each pixel value is denoted with either "A", "B" or "C" to represent which group the pixel value has been classified into.

In step S310 the processing module 210 processes the pixel value 404 in dependence on the classification of the pixel values of the subset 406 into the
25 groups (A, B or C). The processing performed by the processing module 210 may, for example, involve one or more of defective pixel detection, defective pixel correction, denoising, sharpening and demosaicing. It is noted that the denoising (or filtering) referred to herein is performed in the spatial domain, i.e. it is spatial denoising (or spatial filtering). If the processing module 210 performs more than
30 one processing function, these processing functions may be performed in a consolidated manner, such that contradictory aims of the different processing functions can be controlled in a consistent manner, and such that tuning of the image processing system is simplified.

In step S312 the processed pixel value 404 is outputted from the processing module 210 for use in the image processing system 204. The processed pixel value may be used by other processing modules 214 within the image processing system 204, and/or outputted from the image processing system 204.

5

In step S314 the processing module 210 stores group indication data in the store 212, which can be used when processing the next pixel value. The group indication data provides an indication of the three groups (A, B and C) which were used when processing the pixel value 404. For example, step S314 may involve
 10 storing, as an indicative value for a group, an average value μ (e.g. the mean value or the median value) of the pixel values which were classified as being part of that group in step S308. Therefore, as an example, three indicative values μ_A , μ_B and μ_C may be stored to indicate the three groups A, B and C. As described in more detail below, the group indication data which is stored for a group may
 15 include a centroid position for the group, which may be used as part of a defective pixel correction process to choose a suitable group from which an average value can be used to replace a defective pixel value.

Figure 5 shows a two dimensional array of pixel values 502 forming an image,
 20 which is the same image as that described above with reference to Figure 4. Figure 5 illustrates the situation when the processing module 210 is processing the next pixel value 504 after the pixel value 404 shown in Figure 4 (where “next” refers to a raster scan order in this example). In order to process the pixel value 504, the processing module 210 operates on a kernel 506 which defines a subset
 25 of pixel values in the image. Since the image is the same in Figures 4 and 5, the edges 408 and 410 are in the same place in both Figures. It can be appreciated that the subset of pixel values 506 partially overlaps with the subset of pixel values 406 which were used to process the pixel value 404. In step S316, the index i is incremented, such that on the first iteration step S316 involves setting the subset
 30 index to 1.

In step S318 the processing module 210 retrieves the group indication data stored for the previous subset (i.e. the $(i - 1)^{\text{th}}$ subset) from the store 212. The group indication data indicates, for a group, the range of values which fall into that group.

The group indication data may take different forms in different examples, but as described above the group indication data for a group may represent an indicative value for the group (e.g. determined as an average (e.g. mean or median) value of the pixel values which were included in that group in the previous iteration). As an example, three indicative values (μ_A , μ_B and μ_C) may be retrieved in step S318, which have values represented in Figure 6. Figure 6 illustrates a scale 600 of possible pixel values, ranging from 0 to the maximum possible pixel value ("Max") which is, for example, set by the number of bits in a pixel value. For example, if a pixel value has eight bits, then the maximum integer pixel value which can be represented is 255, so the full range of possible pixel values would span from 0 to 255. It is clear that, in other examples, other numbers of bits could be used for the pixel values such that the value of Max can be different. Furthermore, in some examples, pixel values may be in a format in which they are allowed to be negative such that zero would not be at the bottom of the range. The group indication data for a group which is retrieved in step S318 may include an indication of a centroid position for the group.

In the example described above, a pixel value comprises a single colour component. For example, in a monochrome image, each pixel value relates to the same, single colour channel. For simplicity and clarity, an example is described herein in relation to the pixel values of a monochrome image. However, as explained in more detail below, in another example, an image may include pixel values which collectively relate to multiple colour channels, but it may be the case that each particular pixel value relates to a single one of those colour channels. For example, an image sensor may generate pixel values in accordance with a predetermined pattern (e.g. a Bayer pattern, such that some pixel values represent Red colour values, some pixel values represent Blue colour values, and some pixel values represent Green colour values). A demosaicing process can be applied to generate full colour data (e.g. three colour components for each pixel). In a further example, each pixel value may comprise a plurality of colour components (e.g. after a demosaicing process has been performed). As an example, a pixel value could comprise three components (e.g. red, green and blue components), each of which may for example have eight bits. The concepts used in the methods described herein, which are primarily described with reference to a

monochrome image comprising pixel values relating to a single colour channel, could be used for processing chromatic images in which pixel values comprise data relating to a plurality of colour channels (either before or after demosaicing). It is noted that, in examples described herein, where a pixel value has multiple components the classification of the pixel into one of the groups is determined once per pixel (not once per component), i.e. each component of the pixel value is classified into the same group. This ensures that consistent clustering decisions are made across all the colour channels.

10 In step S320, the processing module 210 uses the retrieved group indication data (e.g. the indicative values μ_A , μ_B and μ_C) to define a set of groups into which data values within the i^{th} subset 506 can be grouped. For example, a calibration process can be used to define an extent of a range of data values around an indicative value μ for a group. For example, Figure 7 is a graph showing a calibration curve 702. The calibration curve shows a typical standard deviation that is experienced due to noise in the image sensor 202, as a function of the pixel value. As the pixel value increases the standard deviation tends to increase. The calibration curve 702 can be pre-set (e.g. during design time) based on a likely noise experienced by a generic image sensor, such that the calibration curve is not specifically tailored for the particular image sensor 202. As an alternative, the calibration process may be performed on the sensor at manufacture time (e.g. when it has been manufactured) and the calibration curve 702 may be stored at the manufacture time (this may be referred to as a “static” calibration process). Alternatively, the calibration curve 702 could be set by testing the specific image sensor 202 to detect a standard deviation for different pixel values. In other words, the sensor noise curve (or “calibration curve”) may be generated from an analysis of the noise in a real-time image. This could be done by using the image sensor 202 to determine pixel values for an object with a uniform colour and brightness, and determining differences in the pixel values. By altering the colour and brightness of the object, typical standard deviations could be detected across the range of pixel values, thereby determining the calibration curve 702. For example, the calibration of the sensor could be performed based on earlier, recently processed images (this may be referred to as a “dynamic” calibration process). Some examples may use a “real-time” calibration process in which the

calibration is determined for an iteration and results of the calibration are stored with the group indication data, such that they can be used on the next iteration.

The calibration curve 702 can be used to define ranges of pixel values around the indicative values μ_A , μ_B and μ_C , such that three groups 602_A, 602_B and 602_C can be defined. The extents of the ranges of the groups 602_A, 602_B and 602_C are shown with dashed lines in Figure 6. The extents of the ranges may be determined in different ways in different examples. To give some examples, a range around an indicative value μ may be given by $\pm\sigma$, $\pm 2\sigma$ or $\pm\sigma^2$, where σ is the standard deviation as indicated by the calibration curve 702 for the indicative value μ (such that σ^2 is the variance). It can be seen that the size of the value range 602_A is smaller than the size of the value range 602_B, which itself is smaller than the size of the value range 602_C. This is because, according to the calibration curve 702, the range of a group extends further from the indicative value μ for larger indicative values μ . There may be a respective calibration curve for each colour component (e.g. red, green and blue components), or there may be a single calibration curve which is used for all of the colour components. In the example described above the extent of a range depends only on the expected noise (e.g. as indicated by the calibration curve). However, in some examples the variation may be dependent upon factors (e.g. surface shading or texture) other than the noise. Furthermore, in some examples, a calibration curve is not used to define the ranges. For example, the extents for each range could be of equal size, which may be fixed or variable.

In the example given above, the indicative value μ for a group is the average value of the data values in that group, such that the group range extends either side of the indicative value μ . In other examples, the indicative value could represent another value within the range of values of a group, e.g. the indicative value for a group could be the minimum or maximum of the values within a group such that the group range extends on only one side of the indicative value (e.g. only above the indicative value if the indicative value is a minimum value of a group, and only below the indicative value if the indicative value is a maximum value of a group). Furthermore, although in the example described above, the group indication data that is stored for a group in the store 212 is just the indicative value μ , (e.g. where

the extents of the group around the indicative value are inferred from the calibration curve 702), in other examples it would be possible to store, as the group indication data for a group, the upper and lower bounds of the range of pixel values which are to be classified into the group. This would mean that more data
5 (e.g. twice as much data in a naïve system) would be stored in the store 212 to indicate the groups, but there would be no need to use a calibration curve. In both of these examples, the group indication data for a group indicates (perhaps indirectly) a range of data values which are to be classified into the group.

10 Three groups have been defined which correspond with the groups 602_A, 602_B, 602_C that were used in the previous iteration. Further groups are defined by the gaps between these groups, such that up to seven groups are determined in this example, as indicated by the ranges 604₁ to 604₇. Therefore, in this example, the processing module 210 has defined seven groups 604 which collectively span the
15 entire range of possible data values. The groups 604 are contiguous so they do not overlap with each other. This means that for every possible pixel value (from zero to Max in the example shown in Figure 6) one, and only one, group 604 is defined. In some cases, fewer than seven groups may be defined, e.g. if group 602_A started at zero.

20

It is noted that in this 1D example, where the data values are single-component values, the groups 602_A, 602_B and 602_C are simple to define, and the gaps between those groups create further well-defined groups. In other examples in which the data values have multiple components (e.g. Red, Green and Blue
25 components), groups can be defined as regions (e.g. cuboids, or other simple shapes) surrounding the indicative values (μ) for the groups. The regions are chosen so that they do not overlap with each other. In these multi-component examples, the group/region definitions can be updated in much the same way as the groups in the 1D example described in detail herein. However, the gaps
30 between regions are not so well defined. One option for handling the gaps in a multi-component example is to not define groups for the gaps. This would mean that some possible pixel values would not fall into one of the groups. Another option for handling the gaps in a multi-component example is to define a single “gap” group, whereby any pixel value which does not fall into one of the regions

(e.g. cuboids) surrounding an indicative value for a group is categorised into the “gap group”. In this way “good space” can be tracked by the regions (e.g. cuboids) and the rest of the value space can be flagged as “bad” (or “noise”, “broken”, “needs fixing”).

5

In general, group indication data for up to m groups may be retrieved from the store 212 in step S318, wherein the defined set of groups into which pixel values can be grouped comprises up to n groups, where $n \geq m$, with the extra defined groups being defined by the gaps between and around the retrieved groups. In general $m \leq n \leq 2m + 1$, and as an example n may be equal to $2m + 1$. In the 1D example given above in which there is a gap either side of each of the retrieved groups, $m = 3$ and $n = 7$.

In step S322 the processing module 210 classifies each of the pixel values within the i^{th} subset 506 into one of the groups 604 defined in step S320. Each of the pixel values within the subset 506 is classified into one of the groups 604 based on the value of that pixel value. In other words, if the pixel value falls into the value range for a particular group then the pixel value is classified into that particular group. In Figure 5 each pixel value is denoted with either “A”, “B” or “C” to represent which group the pixel value has been classified into. Group A corresponds to group 604₂, group B corresponds to group 604₄ and group C corresponds to group 604₆.

Steps S318 to S322 are much simpler to implement than steps S306 and S308. Steps S318 to S322 turn the clustering problem into one which is not NP hard, and is solvable in real-time even on devices such as handheld cameras, smartphone and tablets which have tight restrictions on size and power consumption. This is achieved by applying the Bayesian approach such that knowledge of the groups used in the previous iteration is used as a starting point (i.e. as “priors”) for defining the groups to use for the current iteration. Steps S318 and S320 define the groups 604 without needing to analyse the pixel values in the subset 506, and so these steps can be performed very quickly on hardware which covers a small silicon area, with little power consumption. Furthermore, once the groups (604₁ to 604₇) have been defined, it is simple to classify each of the pixel values into one of

30

the groups by comparing the pixel value to the ranges of the values covered by the groups. The classification of a pixel value from the subset 506 into one of the groups does not depend upon the position of the pixel value within the subset 506. Therefore the classification process does not need to analyse the positions of the pixel values, and can simply compare the pixel value to the ranges of values defined for each of the groups. Since the classification of a pixel value does not depend upon the position of the pixel value and does not look for specific edges, it can be a simpler process than an edge detection process which attempts to identify the positions of the edges 408 and 410 within the image in order to determine regions within the kernel 506 which are associated with the different materials in the scene, and can then classify each pixel value based on its position and the positions of the determined regions. One other advantage to the grouping method of steps S318 to S322 is that the orientation of edges in the image does not affect the effectiveness of the grouping process. This is in contrast to some known edge detection algorithms such as a Sobel filter which only looks for edges at certain angles. For example, a Canny filter typically uses a vertical Sobel filter and a horizontal Sobel filter, and then attempts to combine them into a general filter; whereas by contrast, the clustering approach taken in the examples described herein can cope with all angles and curved edges and corners, etc.

As described in more detail below, after the pixel values in the subset 506 have been classified into the seven groups 604_1 to 604_7 , and prior to processing the pixel value 504 in step S324 some of the groups may be discarded based on the number of data values which are classified into each of the groups 604. With reference to the example described above in which group indication data is retrieved in step S318 for up to m groups and in which up to n groups are defined in step S320, where $n > m$, then some of the groups may be discarded, such that m groups are maintained. For example, if m groups are retrieved and n groups are defined then $(n - m)$ groups may be discarded. For example, after step S322, the population of the seven groups 604_1 to 604_7 can be analysed to determine which three of the groups 604 are the most highly populated, and those three groups will be maintained (e.g. groups 604_2 , 604_4 and 604_6 may be maintained), whilst the other four groups will be discarded (e.g. groups 604_1 , 604_3 , 604_5 and 604_7 may be discarded).

In some examples, different numbers of groups are discarded, and in some other examples no groups are discarded at this stage.

5 It is noted that in the example described above, the group which includes pixel 504 itself may have been discarded. However, in some examples (e.g. examples in which defective pixel correction is not being implemented), the group which includes pixel 504 is considered to be an interesting group (when processing pixel 504 itself), worthy of maintaining, even if it is not one of the most highly populated
10 groups. So in these examples, the group including pixel 504 is maintained irrespective of the population of that group, and then one or more other groups may also be maintained based on their respective populations.

15 Furthermore, in some examples, it is possible to merge two or more groups together to form a merged group. For example, rather than discarding a group, the group may be merged with another nearby group (“nearby” in terms of the values represented by the group). For example, two small groups with very similar values may sometimes be better represented as one group.

20

In step S324 the processing module 210 processes the pixel value 504 in dependence on the classification of the pixel values of the subset 506 into the groups. The processing performed by the processing module 210 may for example involve one or more of: (i) defective pixel detection and/or defective pixel
25 correction), (ii) denoising, (iii) sharpening and (iv) demosaicing. For example, the processing performed by the processing module 210 may for example involve two or more of: (i) defective pixel processing (e.g. defective pixel detection and/or defective pixel correction), (ii) denoising, (iii) sharpening and (iv) demosaicing. In some examples, the processing module performs a plurality of processing
30 functions including denoising and sharpening. As described in more detail below with reference to the example flow chart shown in Figure 8, the processing module 210 may perform more than one processing function in a consolidated operation. This can help to control contradictory aims of the different processing functions in a consistent manner which can be easily tuned.

In step S326 the processed pixel value 504 is outputted from the processing module 210 for use in the image processing system 204. The processed pixel value may be used by other processing modules 214 within the image processing system 204, and/or outputted from the image processing system 204.

In step S328 the processing module 210 stores group indication data in the store 212, which can be used when processing the next pixel value. The group indication data provides an indication of the maintained groups (e.g. groups 604₂, 604₄ and 604₆) which were used when processing the pixel value 504 in step S324. For example, step S328 may involve storing, as an indicative value for a group, an average value μ (e.g. the mean value, the median value, a weighted mean value or a mean of a trimmed set of the pixel values in the group (e.g. the mean of the interquartile of pixel values in a group) to give just some examples of average values which could be used) of the pixel values which were classified as being part of that group in step S322. Therefore, as an example, three indicative values μ_A , μ_B and μ_C may be stored to indicate the three groups 604₂, 604₄ and 604₆). The group indication data which is stored for a group may include an indicative position (e.g. a centroid position) for the group, which may be used as part of a defective pixel correction process to choose a suitable group from which an average value can be used to replace a defective pixel value. Other examples of indicative positions include a weighted centroid position and an average trimmed centroid for the group. Furthermore, group indication data which is stored for a group may include an indication of the spread of the data values which are classified as being part of that group, which may be used to define a range for the group. Useful indications of the spread of the data values include a standard deviation, a range, and an interquartile range to give just some examples. Furthermore, group indication data which is stored for a group may include an indication of relevance of the group (e.g. the number of members of the group, i.e. the number of pixel values which have been included in that group).

In some examples, the pixel values may be grouped into fewer than three distinct groups (e.g. if all of the pixel values within a subset have the same pixel value then only one group is populated). In these situations, group indication data for

fewer than three (i.e. fewer than m) groups is stored in step S328. This is why in some cases, group indication data for fewer than three (i.e. fewer than m) groups may be retrieved in step S318 on some iterations.

5 In step S330 the processing module 210 determines whether there are any more pixel values to process, and if there are, then the method passes back to step S316 in which the index i is incremented and the method repeats by performing steps S318 to S328 for the next pixel value. If it is determined in step S330 that there are no more pixel values to process then the method ends in step S332.

10

In the method shown in Figure 3 each processed pixel value is outputted as it has been determined (e.g. in steps S312 and S326). It would be possible to output groups of processed pixel values in batches when a whole batch of pixel values has been processed. For example, when a row of pixel values have been

15

determined then all of the processed pixel values in that row could be outputted together. Furthermore, the method could wait until all of the pixel values within an image have been processed before any of them are outputted.

Figure 8 is a flow chart of an example showing steps that are performed by the processing module 210 to perform multiple processing functions in a consolidated operation in step S324. The multiple processing functions comprise two or more of: defective pixel detection, defective pixel correction, spatial denoising, sharpening and demosaicing.

20

25 Following the classification step S322, the method passes to step S802 in which, as described above, one or more of the groups that were defined in step S320 are selectively discarded, such that m groups are maintained. As described in the example given above, the groups with the lowest numbers of pixel values are discarded. Therefore, the least populated groups are discarded. Following step
30 S802 m groups are maintained and the method passes to step S324, which comprises the steps S804 to S818 as described in more detail below.

In step S804 the processing module 210 determines a centroid for each of the maintained groups. The centroid for a group is representative of a position of the

group. This position is a spatial position of a group within the image. A centroid may be determined for a group by finding an average (e.g. mean) position of the pixel values which have been classified into the group. Alternatively, a centroid may be determined for a group by finding a weighted average (e.g. a weighted mean) position for the pixel values which have been classified into the group, e.g. in which higher weights are applied to pixel values which are closer to the centre of the group compared to the weights applied to pixel values which are further from the centre of the group (e.g. where the distance may be measured as a Euclidean distance or a Manhattan distance).

10

In step S806 the processing module 210 determines an average value for each of the maintained groups. The average value for a group may be a mean, or a median, of the pixel values which have been classified into that group. The average value for a group provides a representative value of the pixel values within the group.

15

In step S808 the processing module 210 determines whether the current pixel value being processed (e.g. pixel value 504) is within one of the maintained groups. The current pixel value is the i^{th} pixel value, and may be referred to as the particular pixel value.

20

In the example shown in Figure 8, defective pixel correction is one of the processes being performed. If the current pixel value is not within one of the maintained groups then the method passes to step S810 in which it is determined that the current pixel value is representative of a defective pixel. For the i^{th} pixel value 504 to not be within one of the maintained groups, it must be dissimilar to most of the other pixel values in the subset 506, which is interpreted as meaning that the pixel value 504 is a defective pixel. In other words, if the pixel value 504 did not represent a defective pixel it is unlikely that the pixel value would not be in one of the maintained groups. In this way, the defective pixel detection function has been performed.

25
30

In step S812 the processing module 210 selects one or more of the maintained groups, based on the centroids of the maintained groups determined in step S804.

For example, the processing module 210 may select the closest of the maintained groups to the pixel value 504, based on the centroids of the maintained groups determined in step S804. For example, with reference to Figure 5, if pixel value 504 was determined to be representative of a defective pixel (i.e. it was not classified as being part of group A) then the centroids of the groups A, B and C would be considered to determine which was closest to the position of the pixel value 504 in the image. In this example, the centroid for group A would be the closest centroid to the position of the pixel value 504. In this example, the centroids for groups B and C would be further than the centroid for group A from the position of the pixel value 504. Therefore, in this example, group A would be selected in step S812 for the pixel value 504 (if the pixel value was determined to be representative of a defective pixel).

In step S814, the defective pixel is replaced with a value based on the determined average value(s) for the selected group(s). If a single group was selected in step S812 then the processing module 210 replaces the i^{th} pixel value 504 with the average value of the selected group which was determined in step S806. In the example shown in Figure 5, the pixel value 504 could be replaced with an average value (e.g. a mean or a median value) of the pixel values in group A. The average of the pixel values which are included in the closest maintained group is a suitable value to assign to the pixel value 504 if it is determined to be defective. If there are multiple (e.g. two) substantially equally close maintained groups, and the pixel value 504 is not within either of them, then it is probable that the pixel value 504 is on or near an edge between the closest of the maintained groups. In this scenario, multiple groups are selected in step S812, and in step S814 a suitable value to replace the pixel value 504 with is a weighted average of the averages of the closest maintained groups. The weights of the weighted average may be determined based on the distances between the position of the pixel value 504 and the centroid positions of the (e.g. two) closest maintained groups.

30

In this way, the defective pixel correction function has been performed in steps S812 and S814. Following step S814 the method passes to step S818 which is described below.

If, in step S808, it is determined that the current pixel value 504 is within one of the maintained groups then the method passes to step S816 in which it is determined that the current pixel value 504 is not representative of a defective pixel. For the i^{th} pixel value 504 to be within one of the maintained groups, it must be similar to
5 at least some of the other pixel values in the subset 506, and this is interpreted as meaning that the pixel value 504 is not a defective pixel. The method then passes to step S818.

In step S818 the processing module 210 applies one or more of denoising,
10 sharpening and demosaicing to the i^{th} pixel value based on the determined centroids of the maintained groups. The centroids of the groups can be used to indicate whether the i^{th} pixel is close to an edge (e.g. edge 408 or 410). For example, the i^{th} pixel may be located at a position (0,0). If the centroid of another group (i.e. a group that does not include the i^{th} pixel) is close to the position (0,0)
15 then there is an edge near to the i^{th} pixel. What constitutes “close to” in this context is an implementation decision, and may be different in different examples.

If the i^{th} pixel value is not close to an edge then spatial denoising can be applied based on the pixel values within the same group as the pixel value 504 to reduce
20 the appearance of random noise in the image, and in this case little or no sharpening may be applied to the pixel value 504. The denoising may involve replacing the pixel value 504 with an average of the pixel values within the same group (e.g. group A). The average may, for example, be a mean, a weighted mean and a mean of a trimmed set of pixel values (e.g. a mean of the interquartile
25 range). The pixel values that are received at the processing module 210 may have b bits, where to give some examples, b may be 8, 10 or 12. These pixel values may originate from a sensor, and it may be the case that one or more of the least significant bits (LSBs), e.g. the two least significant bits, of the pixel values are essentially representing noise. This noise may be a combination of one or
30 more of: noise resulting from analogue to digital conversion (ADC noise), thermal noise, quantisation noise, power plane noise and random Gaussian noise. In accordance with the Central Limit theorem, all of these types of noise get reduced in a deterministic fashion by a denoising process which takes the average of a group of pixel values. The clustering approach described herein enables a pixel

value to be replaced with an average of similar pixel values from its local area, i.e. an average of the group of which the pixel value is a part. This enables the precision of the pixel values to be increased, i.e. the noise in the pixel values can be reduced. In particular, the signal to noise ratio (SNR) of the pixel values can be better (i.e. higher) than the SNR of the pixel values output from the sensor. This is achieved without any temporal averaging, such that it can be achieved in real-time on pixel values of an image, e.g. in a camera pipeline. For example as described above, one or more of the LSBs of the incoming pixel values received at the processing module 210 may be dominated by noise (e.g. two of the LSBs of the input pixel values may be dominated by noise); whereas the pixel values output from the processing module may have fewer LSBs which are dominated by noise (e.g. none or one of the LSBs of the output pixel values may be dominated by noise).

If the i^{th} pixel value is close to an edge then sharpening may be more strongly applied to enhance the sharpness of the edge, e.g. such that the difference between the two groups either side of the edge is accentuated. This can be done by biasing the i^{th} pixel value away from the average value of the pixels in the group on the other side of the edge. In this case the spatial denoising does not filter over the edge so as to avoid blurring the edge, and instead the spatial denoising just filters over pixel values which have been classified into the same group as the i^{th} pixel value.

Some demosaicing techniques are known in the art which use groups of similar colour values. Having formed the groups of pixels using the novel methods described herein, those groups can be used in existing algorithms that use groups, such as demosaicing. In order to prevent, or at least reduce, the introduction of chroma artefacts, demosaicing is applied based on the pixel values which have been classified into the same group as the i^{th} pixel value (and not based on the pixel values which have been classified into other groups). As described above, the classification of the pixel values into the groups is consistent for different colour channels. An example of demosaicing is described with reference to a Bayer image in which 50% of the pixel values are for the Green channel, 25% of the pixel values are for the Red channel, and 25% of the pixel values are for the

Blue channel. The outputs of the demosaicing process are Red, Green and Blue values for each of the pixels. The Green plane is processed first to determine values for the green component of the pixel values of the current group (i.e. the group including the i^{th} pixel) which did not initially have a Green pixel value. This can be done by interpolating, predicting or guessing the Green component values. Therefore, at the end of this first stage, we have a Green component value for each pixel in the current group being processed. Intensity differences between pixels of the current group (i.e. the group including the i^{th} pixel) in the Green colour channel can then be used to interpolate in the Red and Blue colour channels in order to determine the Red and Blue component values for pixels of the current group which did not initially have Red or Blue pixel values respectively. The determination of the pixel values of the Red and Blue planes in this second pass can be performed with a bi-linear filter to reconstruct the missing 75% of data, based on weightings determined by the Green plane. Failure to accurately detect/locate an edge (i.e. failure to correctly assign the pixel values to the groups) may generate erroneous pixels/colours because pixel values from other groups would then distort the determined pixel values for the current group.

Following step S818 the method passes to step S326 described above.

By applying multiple processing functions within the same processing module, consistent decisions are made about how a pixel value is to be processed and the processing of pixel values is deterministic and simple to tune. For example, the denoising and sharpening functions are performed in a consolidated operation so that the functions don't apply contradictory operations to a particular pixel value. Similarly chroma errors are less likely to be produced from the demosaicing process, leading to a lower acceptable level of spatial denoising, and hence a higher level of image definition.

In the example shown in Figure 8 defective pixel correction is one of the processes being performed. In some other examples, defective pixel fixing is not one of the processes being performed, and as such steps S808, S810, S812, S814 and S816 are not performed.

Figure 8 shows an example in which multiple processing functions are applied within the same processing module 210, as incorporated into the method shown in Figure 3 in which the pixel values are classified into groups according to a clustering approach as described above. However, in alternative examples, the concept of applying multiple processing functions within the same processing module in a consolidated manner is not reliant on the manner in which the pixel values are classified into groups. Therefore, in some examples, an edge detection process could be used to classify the pixel values within a subset into one or more groups, and then the multiple processing functions could be applied in a consolidated manner, wherein the multiple processing functions depend upon the classification of pixel values of the subset into the groups. Therefore the method may comprise detecting an edge within a current subset (e.g. using known edge detection techniques), wherein each of the pixel values within the current subset is classified into one of the groups of the set of groups based on the relative position of the pixel value compared to the detected edge.

Figures 9a to 10b illustrate examples in which the processing module 210 is used to apply defective pixel detection, defective pixel correction, denoising and sharpening to two test images. These examples do not implement demosaicing, but demosaicing could be included in other examples.

Figure 9a shows a representation of a first example image 902 which is received at the processing module 210. The image 902 has been artificially generated as a test image so that it has a large number of defective pixels (approximately one in every 300 pixels is defective in the image 902, whereas a real, poor quality sensor might typically have one in every 2000 pixels being defective), so the test image 902 is worse in terms of defective pixels than a standard image captured by a real image sensor. The image 902 has random noise applied to it, which can be seen in the 'speckled' nature of the otherwise flat regions. Furthermore, the image 902 has been downsampled so the edges between different regions have been blurred.

Figure 9b shows an image 904 which is a representation of the first example image 902 when it is output from the processing module 210. It can be seen that

almost all of the defective pixels have been fixed. In particular, all single and double defective pixels have been fixed and most triple and quad defective pixel groups have been fixed. It is also apparent that the noise in the flat regions of the image has been reduced. The edges between the different coloured regions have been sharpened.

Figure 10a shows a representation of a second example image 1002 which is received at the processing module 210. The image 1002 has also been artificially generated as a test image so that it has a large number of defective pixels (again, approximately one in every 300 pixels is defective in the image 1002). The image 1002 has random noise applied to it, which can be seen in the 'speckled' nature of the otherwise flat regions. Furthermore, the image 1002 has been downsampled so the edges between different regions have been blurred. It is further noted that the edges in image 1002 are circular so they are at different orientations at different positions in the image.

Figure 10b shows an image 1004 which is a representation of the second example image 1002 when it is output from the processing module 210. It can be seen that almost all of the defective pixels have been fixed. Again, all single and double defective pixels have been fixed and most triple and quad defective pixel groups have been fixed. It is also apparent that the noise in the flat regions of the image has been reduced. The edges between the different coloured regions have been sharpened, at all edge orientations.

The examples shown in Figures 9a to 10b illustrate how effective the processing can be. Furthermore, this is achieved with a single processing module to implement defective pixel detection, defective pixel correction, denoising and sharpening, rather than with separate processing modules for each separate function. Furthermore, the kernel used in the processing module is a 5x5 kernel in the example used to process the images in Figure 9a to 10b, which is smaller than may have been used previously in one processing module implementing one of the processing functions. Implementing fewer and smaller processing modules leads to smaller silicon area and lower power consumption (as functionality is not replicated in numerous processing modules), and the different processing

functions do not make contradictory decisions, so the resulting images are better and tuning is simpler. This is achieved applying the Bayesian principle to the grouping process as described above, such that knowledge of the grouping is carried over from one iteration to the next iteration.

5

In the examples shown in Figures 4 and 5 the subsets of pixel values 406 and 506 are 5x5 blocks of pixel values. In other examples, the subsets may have different sizes.

10 Furthermore, in the examples described above, the subsets of pixel values are square, e.g. kernels 406 and 506 are square. However, in other examples, the subsets of pixel values may be non-square blocks of pixel values. For example, it may be beneficial to use subsets which represent blocks of pixel values which are wider than they are tall. Just to give an example, subsets could be 3x5 blocks of
15 pixel values (i.e. three rows and five columns). Since the line store module 208 stores rows of data values, (e.g. as they arrive in raster scan order), a reduction in the number of rows means that the line store module 208 can have a reduced size, and a delay incurred in waiting for the line store module 208 to receive enough pixel values for the processing module 210 to process pixel values can be
20 reduced. Whilst it may be advantageous to reduce the number of rows in the subsets for the reasons given above, the same is not true for reducing the number of columns in the subsets (when the data values are received in rows, e.g. in raster scan order or Boustrophedon order), so it can be beneficial to have subsets which have more columns than rows (i.e. they are wider than they are tall). This
25 allows the subsets to retain a large overall size, whilst reducing the number of rows in the subsets. The size and shape of the subsets may be designed in dependence on the processing functions performed by the processing module 210.

30 In the examples described above, the subset of pixel values for a particular pixel value includes and surrounds the particular pixel value, e.g. the subset is centred on the particular pixel value. Pixel values which are on, or near, the edge of the image may be handled in different ways. For example, the size and/or shape of the kernel (or "subset") for a pixel on or near the image edge could be modified to

account for the edge of the image. Alternatively, the image could be extended (e.g. by reflection of existing pixel values across the image edge) such that the size and shape of the kernel (or “subset”) for a pixel on or near the image edge does not need to be modified. The treatment of edge pixels can vary in different
 5 implementations.

In the methods described above with reference to Figure 3, after the data values are received at the processing module, there is a “start-up” procedure (including steps S306 to S314) for the first subset (which can be considered to be for index
 10 $i = 0$), and then there is a “recurrent” procedure (including steps S318 to S330) which are performed for subsequent subsets (e.g. for $i \geq 1$). As described above, the recurrent procedure is much simpler to implement than the start-up procedure. In some examples, the start-up procedure is not performed, such that the recurrent procedure can be performed for each of the subsets. This is simpler to implement,
 15 and means that the same procedure is performed for all of the pixel values, so the hardware can be simplified. The recurrent procedure could start with the first pixel in the image, e.g. using a single group with the average value (μ) set to 0.5. This may cause artifacts in the first one or more pixels to be processed, but it would be simple to implement. Alternatively, the recurrent procedure could start one or
 20 more pixels (e.g. two pixels) before (e.g. to the left of) the first pixel to be processed in the image, e.g. by making use of the reflected region of the image extending beyond the edges of the image. Starting with a single group with $\mu = 0.5$, would mean that the first iteration has a single group into which all of the pixel values are classified, and the group indication data (e.g. average of the pixel
 25 values) for the group may be carried forward to the second iteration. In the second iteration, multiple groups, e.g. three groups, may be defined (e.g. a first group surrounding the average value carried forward from first iteration, a group for values below the first group, and a group for values above the first group). Pixel values can be classified into the three groups, and the group indication data
 30 (e.g. averages of the pixel values) for the three groups can be carried forward to the third iteration. The third iteration may be for the first “real” (i.e. non-reflected) pixel to be processed, e.g. the top left pixel of the image. The recurrent procedure defined above (e.g. steps S318 to S332) can then be performed for the first pixel in the image based on the group indication data (e.g. the three averages of the

pixel values) for the groups, and then the process can proceed as described above.

5 So the start-up procedure (i.e. steps S306 to S314) is not necessarily performed in all examples. In examples in which initial iterations are performed for reflected pixels which do not lie in the image, these initial iterations can be performed during the horizontal blanking interval, such that there is little or no perceptible latency caused by performing the initial iterations. In some examples, more than two initial iterations (or just one initial iteration) may be performed before iterations are
10 performed for real pixel values in the image. By making use of image reflection, the group indication data for the groups will become well trained by the time data is processed for real pixel values of the image. These initial iterations for reflected pixels, might only be performed prior to performing an iteration for the first pixel of the image (e.g. the top left pixel of the image, when the image is processed in
15 raster scan order). For subsequent rows of pixels, the iteration for the left-most pixel of a row can use the group indication data stored for the left-most pixel of the line above. This means that the group indication data for the left-most pixel of a line may be stored in the store 212 at least until it is used for processing the left-most pixel of the line below.

20 The examples given above are described in terms of an image processing system 204 which processes pixel values. The same principles can apply to other types of data processing system, e.g. where data values are not pixel values.

25 For example, the stream of data values may represent a one-dimensional array of data values. In this case, each of the subsets of data values within the stream of data values represents a contiguous block of data values within the one-dimensional array.

30 For example, a 1D stream of data values could be used if the data processing system is an audio processing system and the data values are audio samples of an audio signal. For example, the audio signal could be received and sampled at a microphone and the audio samples can be passed to the audio processing system. The audio processing system may perform some processing functions,

some of which may implement a clustering process according to the Bayesian approach as described herein, whereby group indication data is stored in one iteration for use in defining a set of groups in a subsequent iteration when processing the next data value in the stream. Processed audio samples of the audio signal can be outputted, e.g. via a speaker.

In another example, a 1D stream of data values could be used if the data processing system is a signal processing system and the data values are signal samples of a transmitted signal. For example, the signal could be received over a wired or wireless channel, at a device in which the signal processing system is implemented. The device can sample the received signal and then the signal processing system may perform some processing functions, some of which may implement a clustering process according to the Bayesian approach as described herein whereby group indication data is stored in one iteration for use in defining a set of groups in a subsequent iteration when processing the next data value in the stream. The processed signal samples can be used to extract the data from the received signal.

In the examples described above, a pixel value may be replaced by an average of a plurality of pixel values. For example, a pixel value is replaced by an average value of one or more of the maintained groups if the pixel value is determined to be representative of a defective pixel, i.e. if the pixel value is not in one of the maintained groups. Furthermore, if a pixel is within one of the maintained groups (e.g. if it is determined not to be representative of a defective pixel) then denoising may be applied in which the pixel value is replaced by an average value of the group in which it is located. When pixel values are replaced by averages of multiple pixel values, dynamic range enhancement can be applied whereby the number of bits used to represent the pixel value can be increased. The average value could be, for example, the mean, a weighted mean or a mean of a trimmed set of the pixel values in the group (e.g. the mean of the interquartile of pixel values in a group) to give just some examples of average values which could be used. By determining an average of pixel values which have been classified into the same group, the number of bits of accuracy of the average value can be greater than the number of bits of each individual pixel value. This is because the

averaging process reduces the noise. In other words, the averaging process increases the signal to noise ratio (SNR). This means that the SNR of the pixel values output from the processing module 210 may be greater than the SNR of the pixel values captured by an image sensor (e.g. the image sensor 202) and

5 provided to the image processing system 204. Therefore, the average value may be determined with a greater number of bits than the pixel value which it is to replace. Increasing the number of bits in this manner increases the dynamic range of the pixel values. The clustering approach allows the average value to be determined from a group of pixel values which are similar to the pixel value to be

10 replaced (similar in terms of location, i.e. they are all within the same subset of pixel values, and similar in terms of pixel values because they have been grouped into the same group). Therefore, the averaging reduces the random noise in a deterministic way, thereby improving the signal to noise ratio. For example, random errors which may be reduced include errors resulting from analogue to

15 digital conversion (ADC noise), thermal noise, quantisation noise, power plane noise and random Gaussian noise, to give some examples. Averaging over the whole kernel would blur or soften edges within the kernel. However, by averaging over the pixel values within a single group of similar pixel values, the edges between groups within the kernel are not damaged. The outputted pixel values

20 (i.e. the averaged pixel values) have more bits per pixel than the incoming pixel values, and the outputted bits have better SNR than the incoming pixel values.

More generally, a dynamic range enhancement method could be applied without necessarily being part of the example method shown in Figure 8. For example,

25 Figure 11 illustrates a flow chart for a method of implementing a dynamic range enhancement which may be performed by the processing module 210 using the clustering approach described above. In step S1102 pixel values are received at the processing module 210. Each of the received pixel values has a first number of bits, e.g. each received pixel value may have b_{input} bits where to give just some

30 examples, b_{input} may be 8, 10 or 12, etc. In step S1104 an index i is set to zero. In step S1106 a set of groups are defined into which pixel values of the i^{th} subset can be grouped. The i^{th} subset may for example include the pixel values within a kernel (e.g. the kernel 406 or the kernel 506), which may be centred on the i^{th} pixel value. The groups may be defined as described in the examples above, e.g. in

step S306 or in step S320. In step S1108 each of the data values within the i^{th} subset is classified into one of the defined groups. As described above, the classification of pixel values into groups may be based on a comparison of the pixel value with value ranges defined for each of the groups, or may be based on
5 detecting one or more edges within the kernel and classifying the pixel values based on the positions of the pixel values within the kernel relative to the detected edge(s).

In step S1110 an average value of the pixel values within the i^{th} subset which are
10 classified into one of the groups is determined. For example, step S1110 may involve determining an average value of the pixel values which have been classified into the same group as the i^{th} pixel value (i.e. the pixel value at the centre of the i^{th} subset). As described above, the average value could be, for example, the mean, a weighted mean or a mean of a trimmed set of the pixel
15 values in the group (e.g. the mean of the interquartile of pixel values in a group) to give just some examples of average values which could be used. The determined average value has a second number of bits, $b_{average}$, which is greater than the first number of bits, b_{input} , i.e. $b_{average} > b_{input}$. For example, $b_{average}$ could be one, two, three or four greater than b_{input} , just to give some examples. As
20 explained above, more bits can be used for the average value (compared to the number of bits of the input pixel values) because averaging the pixel values over a particular group of similar pixel values increases the signal to noise ratio of the values. By using more bits for the average value, the average value has a greater precision than the input pixel values. Since the averaging is performed over
25 similar pixel values (i.e. pixel values classified into the same group), the averaging does not blur the edges in the image.

In step S1112 the i^{th} pixel value is replaced based on the determined average value. For example, the i^{th} pixel value may be replaced with the determined
30 average value.

In step S1114 the processed i^{th} pixel value is outputted, e.g. from the processing module 210 for use in the image processing system 204. The processed pixel

value may be used by other processing modules 214 within the image processing system 204, and/or outputted from the image processing system 204.

5 In step S1116 the processing module 210 determines whether there are any more pixel values to process, and if there are, then the method passes to step S1118 in which the index i is incremented. The method then passes back to step S1106 and the method repeats by performing steps S1106 to S1116 for the next pixel value. If it is determined in step S1116 that there are no more pixel values to process then the method ends in step S1120.

10

In the method shown in Figure 11 each processed pixel value is outputted as it has been determined (in step S1114). It would be possible to output groups of processed pixel values in batches when a whole batch of pixel values has been processed. For example, when a row of pixel values have been determined then
15 all of the processed pixel values in that row could be outputted together.

Furthermore, the method could wait until all of the pixel values within an image have been processed before any of them are outputted.

20 According to the method shown in Figure 11, the number of bits of an outputted pixel value is greater than the number of bits of an input pixel value. Therefore, the dynamic range of the pixel values can be increased. For example, each of the outputted pixel values may have one or two more bits than the corresponding input pixel values. Based on a typical group having 8 or 16 pixel values within it, according to the central limit theorem, 1.5 to 2 more bits of accuracy can be
25 achieved by averaging the pixel values. The averaging process described above allows the amount of data in an output pixel value can be greater than the amount of data in the corresponding input pixel value because information is shared over multiple pixel values.

30 It is worth noting that this is not the same as super-resolution which generates a higher density of pixels than the source image, because with the dynamic range enhancement described herein, the same pixel density is maintained but the number of useful bits per pixel for each pixel is increased.

It is further worth noting that this is not the same as temporal denoising, which operates in the temporal domain, rather than the spatial domain (as in the averaging process described above), though the accuracy, effect and impact of temporal denoising will be improved whenever the spatial noise reduction is enhanced.

An image processing system is described herein and shown in Figure 2 as comprising a number of functional blocks. This is schematic only and is not intended to define a strict division between different logic elements of such entities. Each functional block may be provided in any suitable manner. It is to be understood that intermediate values described herein as being formed by an image processing system need not be physically generated by the image processing system at any point and may merely represent logical values which conveniently describe the processing performed by the image processing system between its input and output.

The processing systems described herein (i.e. the image processing systems and data processing systems) may be embodied in hardware on an integrated circuit. The processing systems described herein may be configured to perform any of the methods described herein. Generally, any of the functions, methods, techniques or components described above can be implemented in software, firmware, hardware (e.g., fixed logic circuitry), or any combination thereof. The terms "module," "functionality," "component", "element", "unit", "block" and "logic" may be used herein to generally represent software, firmware, hardware, or any combination thereof. In the case of a software implementation, the module, functionality, component, element, unit, block or logic represents program code that performs the specified tasks when executed on a processor. The algorithms and methods described herein could be performed by one or more processors executing code that causes the processor(s) to perform the algorithms/methods. Examples of a computer-readable storage medium include a random-access memory (RAM), read-only memory (ROM), an optical disc, flash memory, hard disk memory, and other memory devices that may use magnetic, optical, and other techniques to store instructions or other data and that can be accessed by a machine.

The terms computer program code and computer readable instructions as used herein refer to any kind of executable code for processors, including code expressed in a machine language, an interpreted language or a scripting language. Executable code includes binary code, machine code, bytecode, code
5 defining an integrated circuit (such as a hardware description language or netlist), and code expressed in a programming language code such as C, Java (RTM) or OpenCL (RTM). Executable code may be, for example, any kind of software, firmware, script, module or library which, when suitably executed, processed, interpreted, compiled, executed at a virtual machine or other software
10 environment, cause a processor of the computer system at which the executable code is supported to perform the tasks specified by the code.

A processor, computer, or computer system may be any kind of device, machine or dedicated circuit, or collection or portion thereof, with processing capability such
15 that it can execute instructions. A processor may be any kind of general purpose or dedicated processor, such as a CPU, GPU, System-on-chip, state machine, media processor, an application-specific integrated circuit (ASIC), a programmable logic array, a field-programmable gate array (FPGA), or the like. A computer or computer system may comprise one or more processors.

It is also intended to encompass software which defines a configuration of hardware as described herein, such as HDL (hardware description language) software, as is used for designing integrated circuits, or for configuring
20 programmable chips, to carry out desired functions. That is, there may be provided a computer readable storage medium having encoded thereon computer readable program code in the form of an integrated circuit definition dataset that when processed (i.e. run) in an integrated circuit manufacturing system configures
25 the system to manufacture an image processing system or data processing system configured to perform any of the methods described herein, or to manufacture an image processing system or data processing system comprising
30 any apparatus described herein. An integrated circuit definition dataset may be, for example, an integrated circuit description.

Therefore, there may be provided a method of manufacturing, at an integrated circuit manufacturing system, an image processing system or data processing system as described herein. Furthermore, there may be provided an integrated circuit definition dataset that, when processed in an integrated circuit manufacturing system, causes the method of manufacturing an image processing system or data processing system to be performed.

An integrated circuit definition dataset may be in the form of computer code, for example as a netlist, code for configuring a programmable chip, as a hardware description language defining an integrated circuit at any level, including as register transfer level (RTL) code, as high-level circuit representations such as Verilog or VHDL, and as low-level circuit representations such as OASIS (RTM) and GDSII. Higher level representations which logically define an integrated circuit (such as RTL) may be processed at a computer system configured for generating a manufacturing definition of an integrated circuit in the context of a software environment comprising definitions of circuit elements and rules for combining those elements in order to generate the manufacturing definition of an integrated circuit so defined by the representation. As is typically the case with software executing at a computer system so as to define a machine, one or more intermediate user steps (e.g. providing commands, variables etc.) may be required in order for a computer system configured for generating a manufacturing definition of an integrated circuit to execute code defining an integrated circuit so as to generate the manufacturing definition of that integrated circuit.

An example of processing an integrated circuit definition dataset at an integrated circuit manufacturing system so as to configure the system to manufacture an image processing system will now be described with respect to Figure 12.

Figure 12 shows an example of an integrated circuit (IC) manufacturing system 1202 which is configured to manufacture an image processing system as described in any of the examples herein. In particular, the IC manufacturing system 1202 comprises a layout processing system 1204 and an integrated circuit generation system 1206. The IC manufacturing system 1202 is configured to receive an IC definition dataset (e.g. defining an image processing system as

described in any of the examples herein), process the IC definition dataset, and generate an IC according to the IC definition dataset (e.g. which embodies an image processing system as described in any of the examples herein). The processing of the IC definition dataset configures the IC manufacturing system 1202 to manufacture an integrated circuit embodying an image processing system as described in any of the examples herein.

The layout processing system 1204 is configured to receive and process the IC definition dataset to determine a circuit layout. Methods of determining a circuit layout from an IC definition dataset are known in the art, and for example may involve synthesising RTL code to determine a gate level representation of a circuit to be generated, e.g. in terms of logical components (e.g. NAND, NOR, AND, OR, MUX and FLIP-FLOP components). A circuit layout can be determined from the gate level representation of the circuit by determining positional information for the logical components. This may be done automatically or with user involvement in order to optimise the circuit layout. When the layout processing system 1204 has determined the circuit layout it may output a circuit layout definition to the IC generation system 1206. A circuit layout definition may be, for example, a circuit layout description.

The IC generation system 1206 generates an IC according to the circuit layout definition, as is known in the art. For example, the IC generation system 1206 may implement a semiconductor device fabrication process to generate the IC, which may involve a multiple-step sequence of photo lithographic and chemical processing steps during which electronic circuits are gradually created on a wafer made of semiconducting material. The circuit layout definition may be in the form of a mask which can be used in a lithographic process for generating an IC according to the circuit definition. Alternatively, the circuit layout definition provided to the IC generation system 1206 may be in the form of computer-readable code which the IC generation system 1206 can use to form a suitable mask for use in generating an IC.

The different processes performed by the IC manufacturing system 1202 may be implemented all in one location, e.g. by one party. Alternatively, the IC

manufacturing system 1202 may be a distributed system such that some of the processes may be performed at different locations, and may be performed by different parties. For example, some of the stages of: (i) synthesising RTL code representing the IC definition dataset to form a gate level representation of a circuit to be generated, (ii) generating a circuit layout based on the gate level representation, (iii) forming a mask in accordance with the circuit layout, and (iv) fabricating an integrated circuit using the mask, may be performed in different locations and/or by different parties.

10 In other examples, processing of the integrated circuit definition dataset at an integrated circuit manufacturing system may configure the system to manufacture an image processing system without the IC definition dataset being processed so as to determine a circuit layout. For instance, an integrated circuit definition dataset may define the configuration of a reconfigurable processor, such as an
15 FPGA, and the processing of that dataset may configure an IC manufacturing system to generate a reconfigurable processor having that defined configuration (e.g. by loading configuration data to the FPGA).

In some embodiments, an integrated circuit manufacturing definition dataset, when
20 processed in an integrated circuit manufacturing system, may cause an integrated circuit manufacturing system to generate a device as described herein. For example, the configuration of an integrated circuit manufacturing system in the manner described above with respect to Figure 12 by an integrated circuit manufacturing definition dataset may cause a device as described herein to be
25 manufactured.

In some examples, an integrated circuit definition dataset could include software which runs on hardware defined at the dataset or in combination with hardware defined at the dataset. In the example shown in Figure 12, the IC generation
30 system may further be configured by an integrated circuit definition dataset to, on manufacturing an integrated circuit, load firmware onto that integrated circuit in accordance with program code defined at the integrated circuit definition dataset or otherwise provide program code with the integrated circuit for use with the integrated circuit.

The implementation of concepts set forth in this application in devices, apparatus, modules, and/or systems (as well as in methods implemented herein) may give rise to performance improvements when compared with known implementations.

5 The performance improvements may include one or more of increased computational performance, reduced latency, increased throughput, and/or reduced power consumption. During manufacture of such devices, apparatus, modules, and systems (e.g. in integrated circuits) performance improvements can be traded-off against the physical implementation, thereby improving the method
10 of manufacture. For example, a performance improvement may be traded against layout area, thereby matching the performance of a known implementation but using less silicon. This may be done, for example, by reusing functional blocks in a serialised fashion or sharing functional blocks between elements of the devices, apparatus, modules and/or systems. Conversely, concepts set forth in this
15 application that give rise to improvements in the physical implementation of the devices, apparatus, modules, and systems (such as reduced silicon area) may be traded for improved performance. This may be done, for example, by manufacturing multiple instances of a module within a predefined area budget.

20 The applicant hereby discloses in isolation each individual feature described herein and any combination of two or more such features, to the extent that such features or combinations are capable of being carried out based on the present specification as a whole in the light of the common general knowledge of a person skilled in the art, irrespective of whether such features or combinations of features
25 solve any problems disclosed herein. In view of the foregoing description it will be evident to a person skilled in the art that various modifications may be made within the scope of the invention.

Claims

1. A data processing system configured to process a stream of data values, the data processing system comprising:

5 a store configured to store group indication data which indicates one or more groups into which data values can be grouped; and

a processing module configured to process, in each of a plurality of iterations, a respective particular data value of the stream by operating on a respective particular subset of data values of the stream, by, in each of the

10 plurality of iterations:

retrieving, from the store, group indication data for at least one group defined for a previous subset of data values;

15 using the retrieved group indication data to define, for the particular subset of data values in the current iteration, a set of groups into which data values within the particular subset can be grouped;

classifying each of the data values within the particular subset into one of the groups of the set of groups;

20 processing the particular data value using one or more of the data values of the particular subset in dependence on the classification of the data values of the particular subset into the groups; and

storing, in the store, group indication data for at least one group of the set of groups defined for the particular subset of data values in the current iteration, for use in a subsequent iteration;

25 wherein the processing module is configured to output the processed data values.

2. The data processing system of claim 1 wherein each of the data values within the particular subset is classified into one of the groups of the set of groups based on the value of that data value.

30

3. The data processing system of claim 1 or 2 wherein the group indication data for a group indicates a range of data values which are to be classified into the group.

4. The data processing system of claim 3 wherein the group indication data for a group comprises an indicative value within the range of data values which are to be classified into the group.

5 5. The data processing system of claim 4 wherein the extent of the range of data values around the indicative value is defined by a calibration process.

6. The data processing system of claim 5 wherein the calibration process is a static calibration process, a dynamic calibration process or a real-time calibration
10 process.

7. The data processing system of any of claims 4 to 6 wherein said storing group indication data for at least one group comprises storing, as the indicative value for a group, an average value of the data values which are classified as
15 being part of that group.

8. The data processing system of claim 3 wherein the group indication data for a group comprises indications of the upper and lower bounds of the range of data values which are to be classified into the group.
20

9. The data processing system of any preceding claim wherein said storing group indication data for at least one group comprises storing at least one of: (i) an indication of the spread of the data values which are classified as being part of that group, (ii) an indicative position of the data values which are classified as being
25 part of that group, and (iii) an indication of the number of members of the group.

10. The data processing system of any preceding claim wherein a particular subset of data values represents a block of data values which includes and surrounds the corresponding particular data value.
30

11. The data processing system of any preceding claim wherein the stream of data values represents a one-dimensional array of data values.

12. The data processing system of claim 11 wherein each of the subsets of data values within the stream of data values represents a contiguous block of data values within the one-dimensional array.

5 13. The data processing system of claim 11 or 12 wherein the data processing system is an audio processing system, and wherein the data values are audio samples of an audio signal.

10 14. The data processing system of claim 11 or 12 wherein the data processing system is a signal processing system, and wherein the data values are signal samples of a transmitted signal.

15 15. The data processing system of any of claims 1 to 10 wherein the stream of data values represents a two-dimensional array of data values.

16. The data processing system of claim 15 wherein each of the subsets of data values within the stream of data values represents a contiguous block of data values within the two-dimensional array.

20 17. The data processing system of claim 15 or 16 wherein the data processing system is an image processing system, and wherein the data values are pixel values.

25 18. The data processing system of claim 17 wherein the data processing system is configured to receive lines of pixel values, and wherein the data processing system comprises a line store module configured to store lines of pixel values such that the subsets of pixel values can span over more than one line.

30 19. The data processing system of claim 18 wherein the lines of pixel values are rows of pixel values, wherein the data processing system is configured to receive the stream of pixel values according to a raster scan order, and wherein the subsets of pixel values represent blocks of pixel values within the two dimensional array which are wider than they are tall.

20. The data processing system of any preceding claim wherein said processing the particular data value comprises performing, in a consolidated operation, multiple processing functions which depend upon the classification of data values of the particular subset into the groups.

5

21. The data processing system of claim 20 when dependent upon claim 17 wherein the multiple processing functions comprise two or more of: (i) defective pixel detection, (ii) defective pixel correction, (iii) denoising, (iv) sharpening, and (v) demosaicing.

10

22. The data processing system of any preceding claim wherein said group indication data is stored in, and retrieved from, the store for up to m groups, and wherein the defined set of groups into which data values within the particular subset can be grouped comprises up to n groups, where $n \geq m$, and wherein the processing module is configured to, in each of the plurality of iterations:

15

after the classification of the data values of the particular subset into the groups, and prior to said processing the particular data value, selectively discard one or more of the groups based on the number of data values which are classified into each of the groups, such that up to m groups are maintained.

20

23. The data processing system of claim 22 wherein the defined set of groups into which data values within the particular subset can be grouped comprises:

the groups for which group indication data is retrieved from the store, and groups defined by gaps around the groups for which group indication data

25

is retrieved from the store.

24. The data processing system of claim 22 or 23 wherein $m \leq n \leq 2m + 1$.

25. The data processing system of claim 24 wherein $m = 3$ and $n = 7$.

30

26. The data processing system of any of claims 22 to 25 when dependent upon claim 17, wherein said processing the particular pixel value comprises:

if the particular pixel value is within a maintained group, determining that the particular pixel value is not representative of a defective pixel; and

if the particular pixel value is not within a maintained group, determining that the particular pixel value is representative of a defective pixel.

27. The data processing system of any preceding claim wherein the processing
5 module is configured to, in each of the plurality of iterations:

determine, for each of at least one group of the set of groups, a centroid representative of a position of the group; and

determine, for each of the at least one group of the set of groups, an average value of the data values within the group.

10

28. The data processing system of claim 27 wherein the data values are pixel values, and wherein said processing the particular pixel value comprises applying one or more of denoising, sharpening and demosaicing to the particular pixel value based on the at least one determined centroid of the at least one group.

15

29. The data processing system of claim 27 or 28 when dependent upon claim 26 wherein said processing the particular pixel value comprises, if the particular pixel value is determined to be representative of a defective pixel:

20 selecting one or more of the maintained groups based on the determined centroids of the groups, and

replacing the particular pixel value with a value based on the determined one or more average values for the one or more selected groups.

30. A data processing system configured to perform the method of any of
25 claims 32 to 60.

31. The data processing system of any preceding claim wherein the data processing system is embodied in hardware on an integrated circuit.

30 32. A method of processing a stream of data values in a data processing system, the method comprising:

processing, in each of a plurality of iterations, a respective particular data value of the stream by operating on a respective particular subset of data values of the stream, by, in each of the plurality of iterations:

retrieving group indication data which indicates at least one group defined for a previous subset of data values;

using the retrieved group indication data to define, for the particular subset of data values in the current iteration, a set of groups into which data values within the particular subset can be grouped;

classifying each of the data values within the particular subset into one of the groups of the set of groups;

processing the particular data value using one or more of the data values of the particular subset in dependence on the classification of the data values of the particular subset into the groups; and

storing group indication data for at least one group of the set of groups defined for the particular subset of data values in the current iteration, for use in a subsequent iteration;

wherein the method further comprises outputting the processed data values.

33. The method of claim 32 wherein each of the data values within the particular subset is classified into one of the groups of the set of groups based on the value of that data value.

34. The method of claim 32 or 33 wherein the group indication data for a group indicates a range of data values which are to be classified into the group.

35. The method of claim 34 wherein the group indication data for a group comprises an indicative value within the range of data values which are to be classified into the group.

36. The method of claim 35 wherein the extent of the range of data values around the indicative value is defined by a calibration process.

37. The method of claim 36 wherein the calibration process is a static calibration process, a dynamic calibration process or a real-time calibration process.

38. The method of any of claims 35 to 37 wherein said storing group indication data for at least one group comprises storing, as the indicative value for a group, an average value of the data values which are classified as being part of that group.

5

39. The method of claim 34 wherein the group indication data for a group comprises indications of the upper and lower bounds of the range of data values which are to be classified into the group.

10

40. The method of any of claims 32 to 39 wherein said storing group indication data for at least one group comprises storing at least one of: (i) an indication of the spread of the data values which are classified as being part of that group, (ii) an indicative position of the data values which are classified as being part of that group, and (iii) an indication of the number of members of the group.

15

41. The method of any of claims 32 to 40 wherein a particular subset of data values represents a block of data values which includes and surrounds the corresponding particular data value.

20

42. The method of any of claims 32 to 41 wherein the stream of data values represents a one-dimensional array of data values.

25

43. The method of claim 42 wherein each of the subsets of data values within the stream of data values represents a contiguous block of data values within the one-dimensional array.

30

44. The method of claim 42 or 43 wherein the data processing system is an audio processing system, and wherein the data values are audio samples of an audio signal.

45. The method of claim 42 or 43 wherein the data processing system is a signal processing system, and wherein the data values are signal samples of a transmitted signal.

46. The method of any of claims 32 to 41 wherein the stream of data values represents a two-dimensional array of data values.

47. The method of claim 46 wherein each of the subsets of data values within the stream of data values represents a contiguous block of data values within the two-dimensional array.

48. The method of claim 46 or 47 wherein the data processing system is an image processing system, and wherein the data values are pixel values.

10

49. The method of claim 48 wherein the data processing system comprises a line store module and wherein the method comprises:

receiving lines of pixel values; and

storing received lines of pixel values in the line store module, such that the

15

subsets of pixel values can span over more than one line.

50. The method of claim 49 wherein the lines of pixel values are rows of pixel values, wherein said receiving lines of pixel values comprises receiving the stream of pixel values according to a raster scan order, and wherein the subsets of pixel values represent blocks of pixel values within the two dimensional array which are wider than they are tall.

20

51. The method of any of claims 32 to 50 wherein said processing the particular data value comprises performing, in a consolidated operation, multiple processing functions which depend upon the classification of data values of the particular subset into the groups.

25

52. The method of claim 51 when dependent upon claim 48 wherein the multiple processing functions comprise two or more of: (i) defective pixel detection, (ii) defective pixel correction, (iii) denoising, (iv) sharpening, and (v) demosaicing.

30

53. The method of any of claims 32 to 52,

wherein in each of the plurality of iterations, the group indication data is retrieved, and stored, for up to m groups,

wherein the retrieved groups are used to define a set of up to n groups into which data values within the particular subset can be grouped, where $n \geq m$, and

5 wherein the method further comprises, in each of the plurality of iterations:

after the classification of the data values of the particular subset into the groups, and prior to said processing the particular data value, selectively discarding one or more of the groups based on the number of data values which are classified into each of the groups, such that up to m groups are maintained.

54. The method of claim 53 wherein the defined set of groups into which data values within the particular subset can be grouped comprises:

the groups for which group indication data is retrieved, and
 15 groups defined by gaps around the groups for which group indication data is retrieved.

55. The method of claim 53 or 54 wherein $m \leq n \leq 2m + 1$.

20 56. The method of claim 55 wherein $m = 3$ and $n = 7$.

57. The method of any of claims 53 to 56 when dependent upon claim 48, wherein said processing the particular pixel value comprises:

if the particular pixel value is within a maintained group, determining that
 25 the particular pixel value is not representative of a defective pixel; and

if the particular pixel value is not within a maintained group, determining that the particular pixel value is representative of a defective pixel.

58. The method of any of claims 32 to 57 wherein the method comprises, in
 30 each of the plurality of iterations:

determining, for each of at least one group of the set of groups, a centroid representative of a position of the group; and

determining, for each of the at least one group of the set of groups, an average value of the data values within the group.

59. The method of claim 58 wherein the data values are pixel values, and wherein said processing the particular pixel value comprises applying one or more of denoising, sharpening and demosaicing to the particular pixel value based on the at least one determined centroid of the at least one group.

60. The method of claim 58 or 59 when dependent upon claim 57 wherein said processing the particular pixel value comprises, if the particular pixel value is determined to be representative of a defective pixel:

10 selecting one or more of the maintained groups based on the determined centroids of the groups, and

replacing the particular pixel value with a value based on the determined one or more average values for the one or more selected groups.

15 61. A method of manufacturing, using an integrated circuit manufacturing system, a data processing system as claimed in any of claims 1 to 31.

62. Computer readable code configured to cause the method of any of claims 32 to 61 to be performed when the code is run.

20 63. A computer readable storage medium having encoded thereon the computer readable code of claim 62.

25 64. An integrated circuit definition dataset that, when processed in an integrated circuit manufacturing system, configures the integrated circuit manufacturing system to manufacture a data processing system as claimed in any of claims 1 to 31.

30 65. A non-transitory computer readable storage medium having stored thereon a computer readable description of an integrated circuit that, when processed in an integrated circuit manufacturing system, causes the integrated circuit manufacturing system to manufacture a data processing system as claimed in any of claims 1 to 31.

66. An integrated circuit manufacturing system configured to manufacture a data processing system as claimed in any of claims 1 to 31.

67. An integrated circuit manufacturing system comprising:

5 a non-transitory computer readable storage medium having stored thereon a computer readable description of an integrated circuit that describes a data processing system as claimed in any of claims 1 to 31;

a layout processing system configured to process the integrated circuit description so as to generate a circuit layout description of an integrated circuit
10 embodying the data processing system as claimed in any of claims 1 to 31; and
an integrated circuit generation system configured to manufacture the data processing system according to the circuit layout description.