



(51) International Patent Classification:
G06K 9/00 (2006.01)

(21) International Application Number:

PCT/US2018/042840

(22) International Filing Date:

19 July 2018 (19.07.2018)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

62/535,471	21 July 2017 (21.07.2017)	US
16/039,092	18 July 2018 (18.07.2018)	US

(71) Applicant: **SATORI WORLDWIDE, LLC** [US/US];
2225 East Bayshore Road, Suite 200, Palo Alto, California
94303-3220 (US).

(72) Inventors: **CHOUDHARY, Dhruv**; 555 West Middlefield
Road, Apt. M312, Mountain View, California 94043 (US).

ORSINI, Francois; 160 27th Avenue, San Francisco, California 94121 (US). **KOINOV, Anton**; 2580 Knightsbridge Lane, Santa Clara, California 95051 (US).

(74) Agent: **OVANEZIAN, Daniel E.** et al.; WOMBLE BOND DICKINSON (US) LLP, P.O. Box 7037, Atlanta, Georgia 30357-0037 (US).

(81) Designated States (*unless otherwise indicated, for every kind of national protection available*): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JO, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(54) Title: SYSTEM AND METHOD FOR COUNTING OBJECTS ON MULTIPLE ROADS USING A PAN, TILT, AND ZOOM CAMERA

1000




FIG. 10A

(57) Abstract: A method includes receiving a plurality of video frames from a camera and detecting a plurality of objects in the plurality of video frames. The method further includes determining, by a computer processing device, a plurality of trajectories, each of the plurality of trajectories corresponding to a single one of the plurality of objects, respectively. The method further includes identifying, by the computer processing device, a first thoroughfare from a first subset of the plurality of trajectories. The method further includes determining, by the computer processing device, a first number of the plurality of objects that cross a first flow cut line on the first thoroughfare.





(84) Designated States (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Published:

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

SYSTEM AND METHOD FOR COUNTING OBJECTS ON MULTIPLE ROADS USING A PAN, TILT, AND ZOOM CAMERA

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application claims the benefit of U.S. Patent Application No. 16/039,092 filed July 18, 2018, which claims priority to U.S. Provisional Patent Application No. 62/535,471 filed July 21, 2017, the entire contents of which are hereby incorporated by reference.

BACKGROUND

[0002] This specification relates to counting objects and, in particular, to systems and methods for counting objects on multiple roads using a pan, tilt, and zoom camera.

[0003] The publish–subscribe (or “PubSub”) pattern is a data communication messaging arrangement implemented by software systems where so-called publishers publish messages to topics and so-called subscribers receive the messages pertaining to particular topics to which they are subscribed. There can be one or more publishers per topic and publishers generally have no knowledge of what subscribers, if any, will receive the published messages. Because publishers may publish large volumes of messages, and subscribers may subscribe to many topics (or “channels”) the overall volume of messages directed to a particular channel and/or subscriber may be difficult to manage.

BRIEF DESCRIPTION OF THE DRAWINGS

[0004] FIG. 1A illustrates an example system that supports the PubSub communication pattern.

[0005] FIG. 1B illustrates functional layers of software on an example client device.

[0006] FIG. 2 is a diagram of an example messaging system.

[0007] FIG. 3A is a data flow diagram of an example method for writing data to a streamlet.

[0008] FIG. 3B is a data flow diagram of an example method for reading data from a streamlet.

[0009] FIG. 4A is a data flow diagram of an example method for publishing messages to a channel of a messaging system.

[0010] FIG. 4B is a data flow diagram of an example method for subscribing to a channel of a messaging system.

[0011] FIG. 4C is an example data structure for storing messages of a channel of a messaging system.

[0012] FIG. 5A is a data flow diagram of an example method for publishing and replicating messages of a messaging system.

[0013] FIG. 5B is a data flow diagram of an example method for retrieving stored messages in a messaging system.

[0014] FIGS. 5C and 5D are data flow diagrams of example methods for repairing a chain of copies of data in a messaging system.

[0015] FIG. 6 is an example data flow diagram for the application of filtering criteria in a messaging system.

[0016] FIGS. 7A - 7D are illustrations of how messages may be processed using query instructions that include a period-based parameter.

[0017] FIG. 8 is a flowchart of an operation of a messaging system.

[0018] FIG. 9 is a diagram of an example system architecture that may implement methods for counting objects on multiple roads using a pan, tilt, and zoom camera in a messaging system.

[0019] FIG. 10A is a diagram of an example system for counting objects on multiple roads using a pan, tilt, and zoom camera in a messaging system.

[0020] FIG. 10B is a diagram of an example video frame including an indication of a detected object.

[0021] FIG. 10C is a diagram of an example video frame including example trajectories and rectangle detections.

[0022] FIGS. 11A and 11B are diagrams of an example of the application of line logic to single and double lane views.

[0023] FIG. 12 is a diagram of an example threshold binary map of motion information.

[0024] FIG. 13A is a diagram of an example cascading detection system.

[0025] FIGS. 13B-D are diagrams of an example of the application of a cascading detection system.

[0026] FIG. 14 is a flowchart of an example method for counting objects on multiple roads using a pan, tilt, and zoom camera in a messaging system.

[0027] FIG. 15 is a diagram of an example system architecture that may implement counting objects on multiple roads using a pan, tilt, and zoom camera in a messaging system.

DETAILED DESCRIPTION

[0028] Elements of examples or embodiments described with respect to a given aspect of the invention can be used in various embodiments of another aspect of the invention. For example, it is contemplated that features of dependent claims depending from one independent claim can be used in apparatus, systems, and/or methods of any of the other independent claims.

[0029] The details of one or more embodiments of the subject matter described in this specification are set forth in the accompanying drawings and the description below. Other features, aspects, and advantages of the subject matter will become apparent from the description, the drawings, and the claims.

[0030] Systems and methods for counting objects on multiple roads using a pan, tilt, and zoom camera may include a messaging system. The system described herein may be implemented using a PubSub communication system as described herein that allows publishers and subscribers to publish and receive live messages. In alternative embodiments, the system and methods described herein can be implemented using other suitable (non-PubSub) real-time or near real-time network topologies.

[0031] Pan, Tilt, and Zoom (PTZ) camera video streams are being deployed by cities around the world to monitor and understand traffic patterns. One advantage of using PTZ cameras over fixed cameras is that they provide a control mechanism to customize camera angles and zoom for a variety of applications. For example, the optimal camera angle for counting vehicles may be different from the one needed to read license plates of vehicles.

[0032] Using PTZ cameras to detect and count objects may present challenges not present when using standard, fixed cameras. Because PTZ cameras have the ability to freely change angle and zoom characteristics, object detection using a machine-learning-based object detector should be robust to a wide range of angles and zooms. Additionally, slow movements of the PTZ camera should not impact the quality of the algorithms being deployed. In practice, machine learning detectors are never 100% accurate, and camera movements often produce jitter. Thus, there is a need to design systems that are aware of these peculiarities and incorporate appropriate adjustments.

[0033] Advantageously, such a system is described herein. In certain embodiments, systems and methods for counting objects on multiple roads using a pan, tilt, and zoom camera provide object detection, object tracking, line placement logic, and counting logic capabilities – each of which will be described in detail below. By utilizing the systems and methods provided herein, computer systems may operate more efficiently and accurately in a variety of contexts.

[0034] As used herein, the term “vehicle” may refer to one of, for example, a bike, a skateboard, an automotive motor vehicle, a motorcycle, a scooter, a boat, an airplane, a train, a jet, various cars, trucks, sport utility vehicles [SUVs], semi-trailers, speedboats, trains, bicycles, or other suitable conveyance. Furthermore, the term “pedestrian” as used throughout may include both humans and animals. Furthermore, the terms “thoroughfare,” “road,” and “pathway” are used throughout for convenience and may refer to any navigable way, including, but not limited to, roads, highways, freeways, thoroughfares, alleys, streets, boulevards, sidewalks, paved or unpaved trails, causeways, waterways, channels, runways, etc.

[0035] FIG. 1A illustrates an example system 100 that supports the PubSub communication pattern. Publisher clients (e.g., Publisher 1) can publish messages to named channels (e.g., “Channel 1”) by way of the system 100. A message can comprise any type of information

including one or more of the following: text, image content, sound content, multimedia content, video content, binary data, and so on. Other types of message data are possible. Subscriber clients (e.g., Subscriber 2) can subscribe to a named channel using the system 100 and start receiving messages which occur after the subscription request or from a given position (e.g., a message number or time offset). A client can be both a publisher and a subscriber.

[0036] Depending on the configuration, a PubSub system can be categorized as follows:

- One to One (1:1). In this configuration there is one publisher and one subscriber per channel. A typical use case is private messaging.
- One to Many (1:N). In this configuration there is one publisher and multiple subscribers per channel. Typical use cases are broadcasting messages (e.g., stock prices).
- Many to Many (M:N). In this configuration there are many publishers publishing to a single channel. The messages are then delivered to multiple subscribers. Typical use cases are map applications.

[0037] There is no separate operation needed to create a named channel. A channel is created implicitly when the channel is subscribed to or when a message is published to the channel. In some implementations, channel names can be qualified by a name space. A name space comprises one or more channel names. Different name spaces can have the same channel names without causing ambiguity. The name space name can be a prefix of a channel name where the name space and channel name are separated by a dot or other suitable separator. In some implementations, name spaces can be used when specifying channel authorization settings. For instance, the messaging system 100 may have `app1.foo` and `app1.system.notifications` channels where “app1” is the name of the name space. The system can allow clients to subscribe and publish to the `app1.foo` channel. However, clients can only subscribe to, but not publish to the `app1.system.notifications` channel.

[0038] FIG. 1B illustrates functional layers of software on an example client device. A client device (e.g., client 102) is a data processing apparatus such as, for example, a personal computer, a laptop computer, a tablet computer, a smart phone, a smart watch, or a server computer. Other types of client devices are possible. The application layer 104 comprises the end-user application(s) that will integrate with the PubSub system 100. The messaging layer 106 is a programmatic interface for the application layer 104 to utilize services of the system 100 such as channel subscription, message publication, message retrieval, user authentication, and user authorization. In some implementations, the messages passed to and from the messaging layer 106 are encoded as JavaScript Object Notation (JSON) objects. Other message encoding schemes are possible.

[0039] The operating system 108 layer comprises the operating system software on the client 102. In various implementations, messages can be sent and received to/from the system 100 using persistent or non-persistent connections. Persistent connections can be created using, for example, network sockets. A transport protocol such as TCP/IP layer 112 implements the Transport Control Protocol / Internet Protocol communication with the system 100 that can be used by the messaging layer 106 to send messages over connections to the system 100. Other communication protocols are possible including, for example, User Datagram Protocol (UDP). In further implementations, an optional Transport Layer Security (TLS) layer 110 can be employed to ensure the confidentiality of the messages.

[0040] FIG. 2 is a diagram of an example messaging system 100. The system 100 provides functionality for implementing PubSub communication patterns. The system comprises software components and storage that can be deployed at one or more data centers 122 in one or more geographic locations, for example. The system comprises MX nodes (e.g., MX nodes or multiplexer nodes 202, 204 and 206), Q nodes (e.g., Q nodes or queue nodes 208, 210 and 212), one or more configuration manager nodes (e.g., configuration manager 214), and optionally one or more C nodes (e.g., C nodes or cache nodes 220 and 222). Each node can execute in a virtual machine or on a physical machine (e.g., a data processing apparatus). Each MX node can serve as a termination point for one or more publisher and/or subscriber connections through the external network 216. The internal communication among MX nodes, Q nodes, C nodes, and the configuration manager can be conducted over an internal network 218, for example.

[0041] By way of illustration, MX node 204 can be the terminus of a subscriber connection from client 102. Each Q node buffers channel data for consumption by the MX nodes. An ordered sequence of messages published to a channel is a logical channel stream. For example, if three clients publish messages to a given channel, the combined messages published by the clients comprise a channel stream. Messages can be ordered in a channel stream, for example, by time of publication by the client, by time of receipt by an MX node, or by time of receipt by a Q node. Other ways for ordering messages in a channel stream are possible. In the case where more than one message would be assigned to the same position in the order, one of the messages can be chosen (e.g., randomly) to have a later sequence in the order. Each configuration manager node is responsible for managing Q node load, for example, by assigning channels to Q nodes and/or splitting channel streams into so-called streamlets. Streamlets are discussed further below. The optional C nodes provide caching and load removal from the Q nodes.

[0042] In the example messaging system 100, one or more client devices (publishers and/or subscribers) establish respective persistent connections (e.g., TCP connections) to an MX node (e.g., MX node 204). The MX node serves as a termination point for these connections. For

instance, external messages (e.g., between respective client devices and the MX node) carried by these connections can be encoded based on an external protocol (e.g., JSON). The MX node terminates the external protocol and translates the external messages to internal communication, and vice versa. The MX nodes publish and subscribe to streamlets on behalf of clients. In this way, an MX node can multiplex and merge requests of client devices subscribing for or publishing to the same channel, thus representing multiple client devices as one, instead of one by one.

[0043] In the example messaging system 100, a Q node (e.g., Q node 208) can store one or more streamlets of one or more channel streams. A streamlet is a data buffer for a portion of a channel stream. A streamlet will close to writing when its storage is full. A streamlet will close to reading and writing and be de-allocated when its time-to-live (TTL) has expired. By way of illustration, a streamlet can have a maximum size of 1 MB and a TTL of three minutes.

Different channels can have streamlets limited by different sizes and/or by different TTLs. For example, streamlets in one channel can exist for up to three minutes, while streamlets in another channel can exist for up to 10 minutes. In various implementations, a streamlet corresponds to a computing process running on a Q node. The computing process can be terminated after the streamlet's TTL has expired, thus freeing up computing resources (for the streamlet) back to the Q node, for example.

[0044] When receiving a publish request from a client device, an MX node (e.g., MX node 204) makes a request to a configuration manager (e.g., configuration manager 214) to grant access to a streamlet to write the message being published. Note, however, that if the MX node has already been granted write access to a streamlet for the channel (and the channel has not been closed to writing), the MX node can write the message to that streamlet without having to request a grant to access the streamlet. Once a message is written to a streamlet for a channel, the message can be read by MX nodes and provided to subscribers of that channel.

[0045] Similarly, when receiving a channel subscription request from a client device, an MX node makes a request to a configuration manager to grant access to a streamlet for the channel from which messages are read. If the MX node has already been granted read access to a streamlet for the channel (and the channel's TTL has not been closed to reading), the MX node can read messages from the streamlet without having to request a grant to access the streamlet. The read messages can then be forwarded to client devices that have subscribed to the channel. In various implementations, messages read from streamlets are cached by MX nodes so that MX nodes can reduce the number of times needed to read from the streamlets.

[0046] By way of illustration, an MX node can request a grant from the configuration manager that allows the MX node to store a block of data into a streamlet on a particular Q node that

stores streamlets of the particular channel. Example streamlet grant request and grant data structures are as follows:

```
StreamletGrantRequest = {
    "channel": string()
    "mode": "read" | "write"
    "position": 0
}
```

```
StreamletGrantResponse = {
    "streamlet-id": "abcdef82734987",
    "limit-size": 2000000, # 2 megabytes max
    "limit-msgs": 5000, # 5 thousand messages max
    "limit-life": 4000, # the grant is valid for 4 seconds
    "q-node": string()
    "position": 0
}
```

[0047] The StreamletGrantRequest data structure stores the name of the stream channel and a mode indicating whether the MX node intends on reading from or writing to the streamlet. The MX node sends the StreamletGrantRequest to a configuration manager node. The configuration manager node, in response, sends the MX node a StreamletGrantResponse data structure. The StreamletGrantResponse contains an identifier of the streamlet (streamlet-id), the maximum size of the streamlet (limit-size), the maximum number of messages that the streamlet can store (limit-msgs), the TTL (limit-life), and an identifier of a Q node (q-node) on which the streamlet resides. The StreamletGrantRequest and StreamletGrantResponse can also have a position field that points to a position in a streamlet (or a position in a channel) for reading from the streamlet.

[0048] A grant becomes invalid once the streamlet has closed. For example, a streamlet is closed to reading and writing once the streamlet's TTL has expired and a streamlet is closed to writing when the streamlet's storage is full. When a grant becomes invalid, the MX node can request a new grant from the configuration manager to read from or write to a streamlet. The new grant will reference a different streamlet and will refer to the same or a different Q node depending on where the new streamlet resides.

[0049] FIG. 3A is a data flow diagram of an example method for writing data to a streamlet in various embodiments. In FIG. 3A, when an MX node (e.g., MX node 202) request to write to a streamlet is granted by a configuration manager (e.g., configuration manager 214), as described

before, the MX node establishes a Transmission Control Protocol (TCP) connection with the Q node (e.g., Q node 208) identified in the grant response received from the configuration manager (302). A streamlet can be written concurrently by multiple write grants (e.g., for messages published by multiple publisher clients). Other types of connection protocols between the MX node and the Q node are possible.

[0050] The MX node then sends a prepare-publish message with an identifier of a streamlet that the MX node wants to write to the Q node (304). The streamlet identifier and Q node identifier can be provided by the configuration manager in the write grant as described earlier. The Q node hands over the message to a handler process 301 (e.g., a computing process running on the Q node) for the identified streamlet (306). The handler process can send to the MX node an acknowledgement (308). After receiving the acknowledgement, the MX node starts writing (publishing) messages (e.g., 310, 312, 314, and 318) to the handler process, which in turn stores the received data in the identified streamlet. The handler process can also send acknowledgements (316, 320) to the MX node for the received data. In some implementations, acknowledgements can be piggy-backed or cumulative. For example, the handler process can send to the MX node an acknowledgement for each predetermined amount of data received (e.g., for every 100 messages received) or for every predetermined time period (e.g., for every one millisecond). Other acknowledgement scheduling algorithms, such as Nagle's algorithm, can be used.

[0051] If the streamlet can no longer accept published data (e.g., when the streamlet is full), the handler process sends a Negative-Acknowledgement (NAK) message (330) indicating a problem, following by an EOF (end-of-file) message (332). In this way, the handler process closes the association with the MX node for the publish grant. The MX node can then request a write grant for another streamlet from a configuration manager if the MX node has additional messages to store.

[0052] FIG. 3B is a data flow diagram of an example method for reading data from a streamlet in various embodiments. In FIG. 3B, an MX node (e.g., MX node 204) sends to a configuration manager (e.g., configuration manager 214) a request for reading a particular channel starting from a particular message or time offset in the channel. The configuration manager returns to the MX node a read grant including an identifier of a streamlet containing the particular message, a position in the streamlet corresponding to the particular message, and an identifier of a Q node (e.g., Q node 208) containing the particular streamlet. The MX node then establishes a TCP connection with the Q node (352). Other types of connection protocols between the MX node and the Q node are possible.

[0053] The MX node then sends to the Q node a subscribe message (354) with the identifier of the streamlet (in the Q node) and the position in the streamlet from which the MX node wants to read (356). The Q node hands over the subscribe message to a handler process 351 for the streamlet (356). The handler process can send to the MX node an acknowledgement (358). The handler process then sends messages (360, 364, 366), starting at the position in the streamlet, to the MX node. In some implementations, the handler process can send all of the messages in the streamlet to the MX node. After sending the last message in a particular streamlet, the handler process can send a notification of the last message to the MX node. The MX node can send to the configuration manager another request for another streamlet containing a next message in the particular channel.

[0054] If the particular streamlet is closed (e.g., after its TTL has expired), the handler process can send an unsubscribe message (390), followed by an EOF message (392), to close the association with the MX node for the read grant. The MX node can close the association with the handler process when the MX node moves to another streamlet for messages in the particular channel (e.g., as instructed by the configuration manager). The MX node can also close the association with the handler process if the MX node receives an unsubscribe message from a corresponding client device.

[0055] In various implementations, a streamlet can be written into and read from at the same time instance. For example, there can be a valid read grant and a valid write grant at the same time instance. In various implementations, a streamlet can be read concurrently by multiple read grants (e.g., for channels subscribed to by multiple publisher clients). The handler process of the streamlet can order messages from concurrent write grants based on, for example, time-of-arrival, and store the messages based on the order. In this way, messages published to a channel from multiple publisher clients can be serialized and stored in a streamlet of the channel.

[0056] In the messaging system 100, one or more C nodes (e.g., C node 220) can offload data transfers from one or more Q nodes. For instance, if there are many MX nodes requesting streamlets from Q nodes for a particular channel, the streamlets can be offloaded and cached in one or more C nodes. The MX nodes (e.g., as instructed by read grants from a configuration manager) can read the streamlets from the C nodes instead.

[0057] As described above, messages for a channel in the messaging system 100 are ordered in a channel stream. A configuration manager (e.g., configuration manager 214) splits the channel stream into fixed-sized streamlets that each reside on a respective Q node. In this way, storing a channel stream can be shared among many Q nodes; each Q node stores a portion (one or more streamlets) of the channel stream. More particularly, a streamlet can be stored in, for example, registers and/or dynamic memory elements associated with a computing process on a Q node,

thus avoiding the need to access persistent, slower storage devices such as hard disks. This results in faster message access. The configuration manager can also balance load among Q nodes in the messaging system 100 by monitoring respective workloads of the Q nodes and allocating streamlets in a way that avoids overloading any one Q node.

[0058] In various implementations, a configuration manager maintains a list identifying each active streamlet, the respective Q node on which the streamlet resides, an identification of the position of the first message in the streamlet, and whether the streamlet is closed for writing. In some implementations, Q nodes notify the configuration manager and/or any MX nodes that are publishing to a streamlet that the streamlet is closed due to being full or when the streamlet's TTL has expired. When a streamlet is closed, the streamlet remains on the configuration manager's list of active streamlets until the streamlet's TTL has expired so that MX nodes can continue to retrieve messages from the streamlet.

[0059] When an MX node requests a write grant for a given channel and there is not a streamlet for the channel that can be written to, the configuration manager allocates a new streamlet on one of the Q nodes and returns the identity of the streamlet and the Q node in the StreamletGrantResponse. Otherwise, the configuration manager returns the identity of the currently open for writing streamlet and corresponding Q node in the StreamletGrantResponse. MX nodes can publish messages to the streamlet until the streamlet is full or the streamlet's TTL has expired, after which a new streamlet can be allocated by the configuration manager.

[0060] When an MX node requests a read grant for a given channel and there is not a streamlet for the channel that can be read from, the configuration manager allocates a new streamlet on one of the Q nodes and returns the identity of the streamlet and the Q node in the StreamletGrantResponse. Otherwise, the configuration manager returns the identity of the streamlet and Q node that contains the position from which the MX node wishes to read. The Q node can then begin sending messages to the MX node from the streamlet beginning at the specified position until there are no more messages in the streamlet to send. When a new message is published to a streamlet, MX nodes that have subscribed to that streamlet will receive the new message. If a streamlet's TTL has expired, the handler process 351 can send an EOF message (392) to any MX nodes that are subscribed to the streamlet.

[0061] In some implementations, the messaging system 100 can include multiple configuration managers (e.g., configuration manager 214 plus one or more other configuration managers). Multiple configuration managers can provide resiliency and prevent single point of failure. For instance, one configuration manager can replicate lists of streamlets and current grants it maintains to another "slave" configuration manager. As another example, multiple configuration

managers can coordinate operations between them using distributed consensus protocols, such as, for example, Paxos or Raft protocols.

[0062] FIG. 4A is a data flow diagram of an example method for publishing messages to a channel of a messaging system. In FIG. 4A, publishers (e.g., publisher clients 402, 404, 406) publish messages to the messaging system 100 described earlier in reference to FIG. 2. For instance, publishers 402 respectively establish connections 411 and send publish requests to the MX node 202. Publishers 404 respectively establish connections 413 and send publish requests to the MX node 206. Publishers 406 respectively establish connections 415 and send publish requests to the MX node 204. Here, the MX nodes can communicate (417) with a configuration manager (e.g., configuration manager 214) and one or more Q nodes (e.g., Q nodes 212 and 208) in the messaging system 100 via the internal network 218.

[0063] By way of illustration, each publish request (e.g., in JSON key/value pairs) from a publisher to an MX node includes a channel name and a message. The MX node (e.g., MX node 202) can assign the message in the publish request to a distinct channel in the messaging system 100 based on the channel name (e.g., “foo”) of the publish request. The MX node can confirm the assigned channel with the configuration manager 214. If the channel (specified in the subscribe request) does not yet exist in the messaging system 100, the configuration manager can create and maintain a new channel in the messaging system 100. For instance, the configuration manager can maintain a new channel by maintaining a list identifying each active streamlet of the channel’s stream, the respective Q node on which the streamlet resides, and identification of the positions of the first and last messages in the streamlet as described earlier.

[0064] For messages of a particular channel, the MX node can store the messages in one or more buffers or streamlets in the messaging system 100. For instance, the MX node 202 receives from the publishers 402 requests to publish messages M11, M12, M13, and M14 to a channel foo. The MX node 206 receives from the publishers 404 requests to publish messages M78 and M79 to the channel foo. The MX node 204 receives from the publishers 406 requests to publish messages M26, M27, M28, M29, M30, and M31 to the channel foo.

[0065] The MX nodes can identify one or more streamlets for storing messages for the channel foo. As described earlier, each MX node can request a write grant from the configuration manager 214 that allows the MX node to store the messages in a streamlet of the channel foo. For instance, the MX node 202 receives a grant from the configuration manager 214 to write messages M11, M12, M13, and M14 to a streamlet 4101 on the Q node 212. The MX node 206 receives a grant from the configuration manager 214 to write messages M78 and M79 to the streamlet 4101. Here, the streamlet 4101 is the last one (at the moment) of a sequence of streamlets of the channel stream 430 storing messages of the channel foo. The streamlet 4101

has messages (421) of the channel foo that were previously stored in the streamlet 4101, but is still open, i.e., the streamlet 4101 still has space for storing more messages and the streamlet's TTL has not expired.

[0066] The MX node 202 can arrange the messages for the channel foo based on the respective time that each message was received by the MX node 202, e.g., M11, M13, M14, M12 (422), and store the received messages as arranged in the streamlet 4101. That is, the MX node 202 receives M11 first, followed by M13, M14, and M12. Similarly, the MX node 206 can arrange the messages for the channel foo based on their respective time that each message was received by the MX node 206, e.g., M78, M79 (423), and store the received messages as arranged in the streamlet 4101. Other arrangements or ordering of the messages for the channel are possible.

[0067] The MX node 202 (or MX node 206) can store the received messages using the method for writing data to a streamlet described earlier in reference to FIG. 3A, for example. In various implementations, the MX node 202 (or MX node 206) can buffer (e.g., in a local data buffer) the received messages for the channel foo and store the received messages in a streamlet for the channel foo (e.g., streamlet 4101) when the buffered messages reach a predetermined number or size (e.g., 100 messages) or when a predetermined time (e.g., 50 milliseconds) has elapsed. For instance, the MX node 202 can store in the streamlet 100 messages at a time or in every 50 milliseconds. Other appropriate algorithms and techniques, such as Nagle's algorithm, can be used for managing the buffered messages.

[0068] In various implementations, the Q node 212 (e.g., a handler process) stores the messages of the channel foo in the streamlet 4101 in the order as arranged by the MX node 202 and MX node 206. The Q node 212 stores the messages of the channel foo in the streamlet 4101 in the order the Q node 212 receives the messages. For instance, assume that the Q node 212 receives messages M78 (from the MX node 206) first, followed by messages M11 and M13 (from the MX node 202), M79 (from the MX node 206), and M14 and M12 (from the MX node 202). The Q node 212 stores in the streamlet 4101 the messages in the order as received, e.g., M78, M11, M13, M79, M14, and M12, immediately after the messages 421 that are already stored in the streamlet 4101. In this way, messages published to the channel foo from multiple publishers (e.g., 402, 404) can be serialized in a particular order and stored in the streamlet 4101 of the channel foo. Different subscribers that subscribe to the channel foo will receive messages of the channel foo in the same particular order, as will be described in more detail in reference to FIG. 4B.

[0069] In the example of FIG. 4A, at a time instance after the message M12 was stored in the streamlet 4101, the MX node 204 requests a grant from the configuration manager 214 to write to the channel foo. The configuration manager 214 provides the MX node 204 a grant to write

messages to the streamlet 4101, as the streamlet 4101 is still open for writing. The MX node 204 arranges the messages for the channel foo based on the respective time that each message was received by the MX node 204, e.g., M26, M27, M31, M29, M30, M28 (424), and stores the messages as arranged for the channel foo.

[0070] By way of illustration, assume that the message M26 is stored to the last available position of the streamlet 4101. As the streamlet 4101 is now full, the Q node 212 sends to the MX node 204 a NAK message, following by an EOF message, to close the association with the MX node 204 for the write grant, as described earlier in reference to FIG. 3A. The MX node 204 then requests another write grant from the configuration manager 214 for additional messages (e.g., M27, M31, and so on) for the channel foo.

[0071] The configuration manager 214 can monitor available Q nodes in the messaging system 100 for their respective workloads (e.g., how many streamlets are residing in each Q node). The configuration manager 214 can allocate a streamlet for the write request from the MX node 204 such that overloading (e.g., too many streamlets or too many read or write grants) can be avoided for any given Q node. For instance, the configuration manager 214 can identify a least loaded Q node in the messaging system 100 and allocate a new streamlet on the least loaded Q node for write requests from the MX node 204. In the example of FIG. 4A, the configuration manager 214 allocates a new streamlet 4102 on the Q node 208 and provides a write grant to the MX node 204 to write messages for the channel foo to the streamlet 4102. As shown in FIG. 4A, the Q node stores in the streamlet 4102 the messages from the MX node 204 in an order as arranged by the MX node 204: M27, M31, M29, M30, and M28 (assuming that there is no other concurrent write grant for the streamlet 4102 at the moment).

[0072] When the configuration manager 214 allocates a new streamlet (e.g., streamlet 4102) for a request for a grant from an MX node (e.g., MX node 204) to write to a channel (e.g., foo), the configuration manager 214 assigns to the streamlet its TTL, which will expire after TTLs of other streamlets that are already in the channel's stream. For instance, the configuration manager 214 can assign to each streamlet of the channel foo's channel stream a TTL of 3 minutes when allocating the streamlet. That is, each streamlet will expire 3 minutes after it is allocated (created) by the configuration manager 214. Since a new streamlet is allocated after a previous streamlet is closed (e.g., filled entirely or expired), in this way, the channel foo's channel stream comprises streamlets that each expires sequentially after its previous streamlet expires.

[0073] For example, as shown in an example channel stream 430 of the channel foo in FIG. 4A, streamlet 4098 and streamlets before 4098 have expired (as indicated by the dotted-lined gray-out boxes). Messages stored in these expired streamlets are not available for reading for

subscribers of the channel foo. Streamlets 4099, 4100, 4101, and 4102 are still active (not expired). The streamlets 4099, 4100, and 4101 are closed for writing, but still are available for reading. The streamlet 4102 is available for reading and writing, at the moment when the message M28 was stored in the streamlet 4102. At a later time, the streamlet 4099 will expire, following by the streamlets 4100, 4101, and so on.

[0074] FIG. 4B is a data flow diagram of an example method for subscribing to a channel of a messaging system. In FIG. 4B, a subscriber 480 establishes a connection 462 with an MX node 461 of the messaging system 100. Subscriber 482 establishes a connection 463 with the MX node 461. Subscriber 485 establishes a connection 467 with an MX node 468 of the messaging system 100. Here, the MX nodes 461 and 468 can respectively communicate (464) with the configuration manager 214 and one or more Q nodes in the messaging system 100 via the internal network 218.

[0075] A subscriber (e.g., subscriber 480) can subscribe to the channel foo of the messaging system 100 by establishing a connection (e.g., 462) and sending a request for subscribing to messages of the channel foo to an MX node (e.g., MX node 461). The request (e.g., in JSON key/value pairs) can include a channel name, such as, for example, “foo.” When receiving the subscribe request, the MX node 461 can send to the configuration manager 214 a request for a read grant for a streamlet in the channel foo’s channel stream.

[0076] By way of illustration, assume that at the current moment the channel foo’s channel stream 431 includes active streamlets 4102, 4103, and 4104, as shown in FIG. 4B. The streamlets 4102 and 4103 each are full. The streamlet 4104 stores messages of the channel foo, including the last message (at the current moment) stored at a position 47731. Streamlets 4101 and streamlets before 4101 are invalid, as their respective TTLs have expired. Note that the messages M78, M11, M13, M79, M14, M12, and M26 stored in the streamlet 4101, described earlier in reference to FIG. 4A, are no longer available for subscribers of the channel foo, since the streamlet 4101 is no longer valid, as its TTL has expired. As described earlier, each streamlet in the channel foo’s channel stream has a TTL of 3 minutes, thus only messages (as stored in streamlets of the channel foo) that are published to the channel foo (i.e., stored into the channel’s streamlets) no earlier than 3 minutes from the current time can be available for subscribers of the channel foo.

[0077] The MX node 461 can request a read grant for all available messages in the channel foo, for example, when the subscriber 480 is a new subscriber to the channel foo. Based on the request, the configuration manager 214 provides the MX node 461 a read grant to the streamlet 4102 (on the Q node 208) that is the earliest streamlet in the active streamlets of the channel foo (i.e., the first in the sequence of the active streamlets). The MX node 461 can retrieve messages

in the streamlet 4102 from the Q node 208, using the method for reading data from a streamlet described earlier in reference to FIG. 3B, for example. Note that the messages retrieved from the streamlet 4102 maintain the same order as stored in the streamlet 4102. However, other arrangements or ordering of the messages in the streamlet are possible. In various implementations, when providing messages stored in the streamlet 4102 to the MX node 461, the Q node 208 can buffer (e.g., in a local data buffer) the messages and send the messages to the MX node 461 when the buffer messages reach a predetermined number or size (e.g., 200 messages) or a predetermined time (e.g., 50 milliseconds) has elapsed. For instance, the Q node 208 can send the channel foo's messages (from the streamlet 4102) to the MX node 461 200 messages at a time or in every 50 milliseconds. Other appropriate algorithms and techniques, such as Nagle's algorithm, can be used for managing the buffered messages.

[0078] After receiving the last message in the streamlet 4102, the MX node 461 can send an acknowledgement to the Q node 208, and send to the configuration manager 214 another request (e.g., for a read grant) for the next streamlet in the channel stream of the channel foo. Based on the request, the configuration manager 214 provides the MX node 461 a read grant to the streamlet 4103 (on Q node 472) that logically follows the streamlet 4102 in the sequence of active streamlets of the channel foo. The MX node 461 can retrieve messages stored in the streamlet 4103, e.g., using the method for reading data from a streamlet described earlier in reference to FIG. 3B, until it retrieves the last message stored in the streamlet 4103. The MX node 461 can send to the configuration manager 214 yet another request for a read grant for messages in the next streamlet 4104 (on Q node 474). After receiving the read grant, the MX node 461 retrieves messages of the channel foo stored in the streamlet 4104, until the last message at the position 47731. Similarly, the MX node 468 can retrieve messages from the streamlets 4102, 4103, and 4104 (as shown with dotted arrows in FIG. 4B), and provide the messages to the subscriber 485.

[0079] The MX node 461 can send the retrieved messages of the channel foo to the subscriber 480 (via the connection 462) while receiving the messages from the Q nodes 208, 472, or 474. In various implementations, the MX node 461 can store the retrieved messages in a local buffer. In this way, the retrieved messages can be provided to another subscriber (e.g., subscriber 482) when the other subscriber subscribes to the channel foo and requests the channel's messages. The MX node 461 can remove messages stored in the local buffer that each has a time of publication that has exceeded a predetermined time period. For instance, the MX node 461 can remove messages (stored in the local buffer) with respective times of publication exceeding 3 minutes. In some implementations, the predetermined time period for keeping messages in the local buffer on MX node 461 can be the same as or similar to the time-to-live duration of a

streamlet in the channel foo's channel stream, since at a given moment, messages retrieved from the channel's stream do not include those in streamlets having respective times-to-live that had already expired.

[0080] The messages retrieved from the channel stream 431 and sent to the subscriber 480 (by the MX node 461) are arranged in the same order as the messages were stored in the channel stream, although other arrangements or ordering of the messages are possible. For instance, messages published to the channel foo are serialized and stored in the streamlet 4102 in a particular order (e.g., M27, M31, M29, M30, and so on), then stored subsequently in the streamlet 4103 and the streamlet 4104. The MX node retrieves messages from the channel stream 431 and provides the retrieved messages to the subscriber 480 in the same order as the messages are stored in the channel stream: M27, M31, M29, M30, and so on, followed by ordered messages in the streamlet 4103, and followed by ordered messages in the streamlet 4104.

[0081] Instead of retrieving all available messages in the channel stream 431, the MX node 461 can request a read grant for messages stored in the channel stream 431 starting from a message at particular position, e.g., position 47202. For instance, the position 47202 can correspond to an earlier time instance (e.g., 10 seconds before the current time) when the subscriber 480 was last subscribing to the channel foo (e.g., via a connection to the MX node 461 or another MX node of the messaging system 100). The MX node 461 can send to the configuration manager 214 a request for a read grant for messages starting at the position 47202. Based on the request, the configuration manager 214 provides the MX node 461 a read grant to the streamlet 4104 (on the Q node 474) and a position on the streamlet 4104 that corresponds to the channel stream position 47202. The MX node 461 can retrieve messages in the streamlet 4104 starting from the provided position, and send the retrieved messages to the subscriber 480.

[0082] As described above in reference to FIGS 4A and 4B, messages published to the channel foo are serialized and stored in the channel's streamlets in a particular order. The configuration manager 214 maintains the ordered sequence of streamlets as they are created throughout their respective times-to-live. Messages retrieved from the streamlets by an MX node (e.g., MX node 461, or MX node 468) and provided to a subscriber can be, in some implementations, in the same order as the messages are stored in the ordered sequence of streamlets. In this way, messages sent to different subscribers (e.g., subscriber 480, subscriber 482, or subscriber 485) can be in the same order (as the messages are stored in the streamlets), regardless which MX nodes the subscribers are connected to.

[0083] In various implementations, a streamlet stores messages in a set of blocks of messages. Each block stores a number of messages. For instance, a block can store two hundred kilobytes

of messages (although other sizes of blocks of messages are possible). Each block has its own time-to-live, which can be shorter than the time-to-live of the streamlet holding the block. Once a block's TTL has expired, the block can be discarded from the streamlet holding the block, as described in more detail below in reference to FIG 4C.

[0084] FIG. 4C is an example data structure for storing messages of a channel of a messaging system. As described with the channel foo in reference to FIGS. 4A and 4B, assume that at the current moment the channel foo's channel stream 432 includes active streamlets 4104 and 4105, as shown in FIG. 4C. Streamlet 4103 and streamlets before 4103 are invalid, as their respective TTLs have expired. The streamlet 4104 is already full for its capacity (e.g., as determined by a corresponding write grant) and is closed for additional message writes. The streamlet 4104 is still available for message reads. The streamlet 4105 is open and is available for message writes and reads.

[0085] By way of illustration, the streamlet 4104 (e.g., a computing process running on the Q node 474 shown in FIG. 4B) currently holds two blocks of messages. Block 494 holds messages from channel positions 47301 to 47850. Block 495 holds messages from channel positions 47851 to 48000. The streamlet 4105 (e.g., a computing process running on another Q node in the messaging system 100) currently holds two blocks of messages. Block 496 holds messages from channel positions 48001 to 48200. Block 497 holds messages starting from channel position 48201, and still accepts additional messages of the channel foo.

[0086] When the streamlet 4104 was created (e.g., by a write grant), a first block (sub-buffer) 492 was created to store messages, e.g., from channel positions 47010 to 47100. Later on, after the block 492 had reached its capacity, another block 493 was created to store messages, e.g., from channel positions 47111 to 47300. Blocks 494 and 495 were subsequently created to store additional messages. Afterwards, the streamlet 4104 was closed for additional message writes, and the streamlet 4105 was created with additional blocks for storing additional messages of the channel foo.

[0087] In this example, the respective TTL's of blocks 492 and 493 had expired. The messages stored in these two blocks (from channel positions 47010 to 47300) are no longer available for reading by subscribers of the channel foo. The streamlet 4104 can discard these two expired blocks, e.g., by de-allocating the memory space for the blocks 492 and 493. The blocks 494 or 495 could become expired and be discarded by the streamlet 4104, before the streamlet 4104 itself becomes invalid. Alternatively, streamlet 4104 itself could become invalid before the blocks 494 or 495 become expired. In this way, a streamlet can hold one or more blocks of messages, or contain no block of messages, depending on respective TTLs of the streamlet and blocks, for example.

[0088] A streamlet, or a computing process running on a Q node in the messaging system 100, can create a block for storing messages of a channel by allocating a certain size of memory space from the Q node. The streamlet can receive, from an MX node in the messaging system 100, one message at a time and store the received message in the block. Alternatively, the MX node can assemble (i.e., buffer) a group of messages and send the group of messages to the Q node. The streamlet can allocate a block of memory space (from the Q node) and store the group of messages in the block. The MX node can also perform compression on the group of messages, e.g., by removing a common header from each message or performing other suitable compression techniques.

[0089] As described above, a streamlet (a data buffer) residing on a Q node stores messages of a channel in the messaging system 100. To prevent failure of the Q node (a single point failure) that can cause messages being lost, the messaging system 100 can replicate messages on multiple Q nodes, as described in more detail below.

[0090] FIG. 5A is a data flow diagram of an example method 500 for publishing and replicating messages of the messaging system 100. As described earlier in reference to FIG. 4A, the MX node 204 receives messages (of the channel foo) from the publishers 406. The configuration manager 214 can instruct the MX Node 204 (e.g., with a write grant) to store the messages in the streamlet 4102 on the Q node 208. In FIG. 5A, instead of storing the messages on a single node (e.g., Q node 208), the configuration manager 214 allocates multiple Q nodes to store multiple copies of the streamlet 4102 on these Q nodes.

[0091] By way of illustration, the configuration manager 214 allocates Q nodes 208, 502, 504, and 506 in the messaging system 100 to store copies of the streamlet 4102. The configuration manager 214 instructs the MX node 204 to transmit the messages for the channel foo (e.g., messages M27, M31, M29, M30, and M28) to the Q node 208 (512). A computing process running on the Q node 208 stores the messages in the first copy (copy #1) of the streamlet 4102. Instead of sending an acknowledgement message to the MX node 204 after storing the messages, the Q node 208 forwards the messages to the Q node 502 (514). A computing process running on the Q node 502 stores the messages in another copy (copy #2) of the streamlet 4102. Meanwhile, the Q node 502 forwards the messages to the Q node 504 (516). A computing process running on the Q node 504 stores the messages in yet another copy (copy #3) of the streamlet 4102. The Q node 504 also forwards the message to the Q node 506 (518). A computing process running on the Q node 506 stores the messages in yet another copy (copy #4) of the streamlet 4102. The Q node 506 can send an acknowledgement message to the MX node 204, indicating that all the messages (M27, M31, M29, M30, and M28) have been stored successfully in streamlet copies #1, #2, #3 and #4.

[0092] In some implementations, after successfully storing the last copy (copy #4), the Q node 506 can send an acknowledgement to its upstream Q node (504), which in turns sends an acknowledgement to its upstream Q node (502), and so on, until the acknowledgement is sent to the Q node 208 storing the first copy (copy #1). The Q node 208 can send an acknowledgement message to the MX node 204, indicating that all messages have been stored successfully in the streamlet 4102 (i.e., in the copies #1, #2, #3 and #4).

[0093] In this way, four copies of the streamlet 4102 (and each message in the streamlet) are stored in four different Q nodes. Other numbers (e.g., two, three, five, or other suitable number) of copies of a streamlet are also possible. In the present illustration, the four copies form a chain of copies including a head copy in the copy #1 and a tail copy in the copy #4. When a new message is published to the streamlet 4102, the message is first stored in the head copy (copy #1) on the Q node 208. The message is then forwarded downstream to the next adjacent copy, the copy #2 on the Q node 502 for storage, then to the copy #3 on the Q node 504 for storage, until the message is stored in the tail copy the copy #4 on the Q node 506.

[0094] In addition to storing and forwarding by messages, the computing processes running on Q nodes that store copies of a streamlet can also store and forward messages by blocks of messages, as described earlier in reference to FIG. 4C. For instance, the computing process storing the copy #1 of the streamlet 4102 on Q node 208 can allocate memory and store a block of, for example, 200 kilobytes of messages (although other sizes of blocks of messages are possible), and forward the block of messages to the next adjacent copy (copy #2) of the chain for storage, and so on, until the block messages is stored in the tail copy (copy #4) on the Q node 506.

[0095] Messages of the streamlet 4102 can be retrieved and delivered to a subscriber of the channel foo from one of the copies of the streamlet 4102. FIG. 5B is a data flow diagram of an example method 550 for retrieving stored messages in the messaging system 100. For instance, the subscriber 480 can send a request for subscribing to messages of the channel to the MX node 461, as described earlier in reference to FIG. 4B. The configuration manager 214 can provide to the MX node 461 a read grant for one of the copies of the streamlet 4102. The MX node 461 can retrieve messages of the streamlet 4102 from one of the Q nodes storing a copy of the streamlet 4102, and provide the retrieved messages to the subscriber 480. For instance, the MX node 461 can retrieve messages from the copy #4 (the tail copy) stored on the Q node 506 (522).

[0096] As for another example, the MX node 461 can retrieve messages from the copy #2 stored on the Q node 502 (524). In this way, the multiple copies of a streamlet (e.g., copies #1, #2, #3, and #4 of the streamlet 4102) provide replication and redundancy against failure if only one copy of the streamlet were stored in the messaging system 100. In various implementations,

the configuration manager 214 can balance workloads among the Q nodes storing copies of the streamlet 4102 by directing the MX node 461 (e.g., with a read grant) to a particular Q node that has, for example, less current read and write grants as compared to other Q nodes storing copies of the streamlet 4102.

[0097] A Q node storing a particular copy in a chain of copies of a streamlet may fail, e.g., a computing process on the Q node storing the particular copy may freeze. Other failure modes of a Q node are possible. An MX node can detect a failed node (e.g., from non-responsiveness of the failed node) and report the failed node to a configuration manager in the messaging system 100 (e.g., configuration manager 214). A peer Q node can also detect a failed Q node and report the failed node to the configuration manager. For instance, an upstream Q node may detect a failed downstream Q node when the downstream Q node is non-responsive, e.g., fails to acknowledge a message storage request from the upstream Q node as described earlier. It is noted that failure of a Q node storing a copy of a particular streamlet of a particular channel stream does not have to be for publish or subscribe operations of the particular streamlet or of the particular channel stream. Failure stemming from operations on another streamlet or another channel stream can also alert a configuration manager about failure of a Q node in the messaging system 100.

[0098] When a Q node storing a particular copy in a chain of copies of a streamlet fails, a configuration manager in the messaging system 100 can repair the chain by removing the failed node, or by inserting a new node for a new copy into the chain, for example. FIGS. 5C and 5D are data flow diagrams of example methods for repairing a chain of copies of a streamlet in the messaging system 100. In FIG. 5C, for instance, after detecting that the Q node 504 fails, the configuration manager 214 can repair the chain of copies by redirecting messages intended to be stored in the copy #3 of the streamlet 4102 on the Q node 502 to the copy #4 of the streamlet 4102 on the Q node 506. In this example, a message (or a block of messages) is first sent from the MX node 204 to the Q node 208 for storage in the copy #1 of the streamlet 4102 (572). The message then is forwarded to the Q node 502 for storage in the copy #2 of the streamlet 4102 (574). The message is then forwarded to the Q node 506 for storage in the copy #4 of the streamlet 4102 (576). The Q node 506 can send an acknowledgement message to the configuration manager 214 indicating that the message has been stored successfully.

[0099] Here, a failed node can also be the node storing the head copy or the tail copy of the chain of copies. For instance, if the Q node 208 fails, the configuration manager 214 can instruct the MX node 204 first to send the message to the Q node 502 for storage in the copy #2 of the streamlet 4102. The message is then forwarded to the next adjacent copy in the chain for storage, until the message is stored in the tail copy.

[00100] If the Q node 506 fails, the configuration manager 214 can repair the chain of copies of the streamlet 4102 such that the copy #3 on the Q node 504 becomes the tail copy of the chain. A message is first stored in the copy #1 on the Q node 208, then subsequently stored in the copy #2 on the Q node 502, and the copy #3 on the Q node 504. The Q node 504 then can send an acknowledgement message to the configuration manager 214 indicating that the message has been stored successfully.

[00101] In FIG. 5D, the configuration manager 214 replaces the failed node Q node 504 by allocating a new Q node 508 to store a copy #5 of the chain of copies of the streamlet 4102. In this example, the configuration manager 214 instructs the MX node 204 to send a message (from the publishers 406) to the Q node 208 for storage in the copy #1 of the streamlet 4102 (582). The message is then forwarded to the Q node 502 for storage in the copy #2 of the streamlet 4102 (584). The message is then forwarded to the Q node 508 for storage in the copy #5 of the streamlet 4012 (586). The message is then forwarded to the Q node 506 for storage in the copy #4 of the streamlet 4102 (588). The Q node 506 can send an acknowledgement message to the configuration manager 214 indicating that the message has been stored successfully.

[00102] FIG. 6 is a data flow diagram 600 illustrating the application of selective filtering, searching, transforming, querying, aggregating and transforming of messages in real time to manage the delivery of messages into and through each channel and on to individual subscribers. Users operating applications on client devices, such as, for example, smartphones, tablets, and other internet-connected devices, act as subscribers (e.g., subscriber 480 in FIG. 4B, subscriber 602 in FIG. 6). The applications may be, for example, consumers of the messages to provide real-time information about news, transportation, sports, weather, or other subjects that rely on published messages attributed to one or more subjects and/or channels. Message publishers 604 can be any internet-connected service that provides, for example, status data, transactional data or other information that is made available to the subscribers 602 on a subscription basis.

[00103] In some versions, for example, the relationship between publishers and channels is 1:1, that is there is one and only one publisher that provides messages into that particular channel. In other instances, the relationship may be many-to-one (more than one publisher provides messages into a channel), one-to-many (a publisher's messages are sent to more than one channel), or many-to-many (more than one publisher provides messages to more than one channel). Typically, when a subscriber subscribes to a channel, they receive all messages and all message data published to the channel as soon as it is published. The result, however, is that many subscribers can receive more data (or data that requires further processing) than is useful. The additional filtering or application of functions against the data places undue processing

requirements on the subscriber application and can delay presentation of the data in its preferred format.

[00104] A filter 606 can be created by providing suitable query instructions at, for example, the time the subscriber 602 subscribes to the channel 608. The filter 606 that is specified can be applied to all messages published to the channel 608 (e.g., one message at a time), and can be evaluated before the subscriber 602 receives the messages (e.g., see Block 2 in FIG. 6). By allowing subscribers 602 to create query instructions a priori, that is upon subscribing to the channel 608 and before data is received into the channel 608, the burden of filtering and processing messages moves closer to the data source, and can be managed at the channel level. As a result, the messages are pre-filtered and/or pre-processed before they are forwarded to the subscriber 602. Again, the query instructions need not be based on any a priori knowledge of the form or substance of the incoming messages. The query instructions can be used to pre-process data for applications such as, for example, real-time monitoring services (for transportation, healthcare, news, sports, weather, etc.) and dashboards (e.g., industrial monitoring applications, financial markets, etc.) to filter data, summarize data and/or detect anomalies. One or more filters 606 can be applied to each channel 608.

[0105] The query instructions can implement real-time searches and queries, aggregate or summarize data, or transform data for use by a subscriber application. In some embodiments, including those implementing JSON formatted messages, the messages can be generated, parsed and interpreted using the query instructions, and the lack of a pre-defined schema (unlike conventional RDBMS / SQL-based applications) means that the query instructions can adapt to changing business needs without the need for schema or application layer changes. This allows the query instructions to be applied selectively at the message level within a channel, thus filtering and/or aggregating messages within the channel. In some instances, the queries may be applied at the publisher level – meaning channels that receive messages from more than one publisher may apply certain filters against messages from specific publishers. The query instructions may be applied on a going-forward basis, that is on only newly arriving messages, and/or in some cases, the query instructions may be applied to historical messages already residing in the channel queue.

[0106] The query instructions can be applied at either or both of the ingress and egress side of the PubSub service. On the egress side, the query instructions act as a per-connection filter against the message channels, and allows each subscriber to manage their own set of unique filters. On the ingress side, the query instructions operate as a centralized, system-wide filter that is applied to all published messages.

[0107] For purposes of illustration and not limitation, examples of query instructions that may be applied during message ingress include:

- A message may be distributed to multiple channels or to a different channel (e.g., based on geo-location in the message, or based on a hash function of some value in the message).
- A message may be dropped due to spam filtering or DoS rules (e.g., limiting the number of messages a publisher can send in a given time period).
- An alert message may be sent to an admin channel on some event arriving at any channel (e.g., `cpu_temp > threshold`).

[0108] For purposes of illustration and not limitation, examples of query instructions that may be applied during message egress include:

- Channels that contain events from various sensors where the user is only interested in a subset of the data sources.
- Simple aggregations, where a system reports real time events, such as cpu usage, sensor temperatures, etc., and we would like to receive some form of aggregation over a short time period, irrespective of the number of devices reporting or the reporting frequency, e.g., `average(cpu_load)`, `max(temperature)`, `count(number_of_users)`, `count(number_of_messages) group by country`.
- Transforms, where a system reports real time events and metadata is added to them from mostly static external tables, e.g., adding a city name based on IP address, converting an advertisement ID to a marketing campaign ID or to a marketing partner ID.
- Adding default values to event streams where such values do not exist on certain devices.
- Advanced aggregations, where a system reports real time events, and combines some mostly static external tables data into the aggregation in real time, e.g., grouping advertisement clicks by partners and counting number of events.
- Counting number of user events, grouping by a/b test cell allocation.

[0109] In some embodiments, the query instructions may be used to define an index or other suitable temporary data structure, which may then be applied against the messages as they are received into the channel to allow for the reuse of the data element(s) as searchable elements. In such cases, a query frequency may be maintained to describe the number of times (general, or in a given period) that a particular data element is referred to or how that element is used. If the frequency that the data element is used in a query exceeds some threshold, the index may be stored for subsequent use on incoming messages, whereas in other instances in which the index

is used only once (or infrequently) it may be discarded. In some instances, the query instruction may be applied to messages having arrived at the channel prior to the creation of the index. Thus, the messages are not indexed according to the data elements described in the query instructions but processed using the query instructions regardless, whereas messages arriving after the creation of the index may be filtered and processed using the index. For queries or other subscriptions that span the time at which the index may have been created, the results of applying the query instructions to the messages as they are received and processed with the index may be combined with results of applying the query instructions to non-indexed messages received prior to receipt of the query instructions.

[0110] For purposes of illustration and not limitation, one use case for such a filtering application is a mapping application that subscribes to public transportation data feeds, such as the locations of all buses across a city. The published messages may include, for example, geographic data describing the location, status, bus agency, ID number, route number, and route name of the buses. Absent pre-defined query instructions, the client application would receive individual messages for all buses. However, query instructions may be provided that filter out, for example, inactive routes and buses and aggregate, for example, a count of buses by agency. The subscriber application receives the filtered bus data in real time and can create reports, charts and other user-defined presentations of the data. When new data is published to the channel, the reports can be updated in real time based on a period parameter (described in more detail below).

[0111] The query instructions can be provided (e.g., at the time the subscriber subscribes to the channel) in any suitable format or syntax. For example, the following illustrates the structure of several fields of a sample subscription request Protocol Data Unit (PDU) with the PDU keys specific to adding a filter to a subscription request:

```
{
  "action": "subscribe",
  "body": {
    "channel": "ChannelName"
    "filter": "QueryInstructions"
    "period": [1-60, OPTIONAL]
  }
}
```

In the above subscription request PDU, the “channel” field can be a value (e.g., string or other appropriate value or designation) for the name of the channel to which the subscriber wants to subscribe. The “filter” field can provide the query instructions or other suitable filter commands, statements, or syntax that define the type of key/values in the channel message to return to the subscriber. The “period” parameter specifies the time period in, for example, seconds, to retain messages before returning them to the subscriber (e.g., an integer value from 1 to 60, with a default of, for example, 1). The “period” parameter will be discussed in more detail below. It is noted that a subscription request PDU can include any other suitable fields, parameters, or values.

[0112] One example of a query instruction is a “select” filter, which selects the most recent (or “top”) value for all (e.g., “select.*”) or selected (e.g., “select.name”) data elements. In the example below, the Filter column shows the filter value sent in the query instructions as part of a subscription as the filter field. The Message Data column lists the input of the channel message data and the message data sent to the client as output. In this example, the value for the “extra” key does not appear in the output, as the “select” filter can return only the first level of results and does not return any nested key values.

Filter	Message Data
SELECT *	Input <pre>{ "name": "art", "eye": "blue"}, { "name": "art", "age": 11}, { "age": 12, "height": 190}</pre> Output <pre>{ "name": "art", "age": 12, "eye": "blue", "height": 190}</pre>

SELECT top.*	Input <pre>{ "top": { "age": 12, "eyes": "blue" } }, { "top": { "name": "joy", "height": 168 }, "extra": 1 }, { "top": { "name": "art" } }</pre> Output <pre>{ "name": "art", "age": 12, "eye": "blue", "height": 168 }</pre>
--------------	---

[0113] For aggregative functions, all messages can be combined that satisfy the query instructions included in the GROUP BY clause. The aggregated values can then be published as a single message to the subscriber(s) at the end of the aggregation period. The number of messages that are aggregated depends on, for example, the number of messages received in the channel in the period value for the filter. For instance, if the period parameter is set to 1, and 100 messages are received in one second, all 100 messages are aggregated into a single message for transmission to the subscriber(s). As an example, a query instruction as shown below includes a filter to aggregate position data for an object, grouping it by obj_id, with a period of 1:

SELECT * WHERE (<expression with aggregate function>) GROUP BY obj_id

In this example, all messages published in the previous second with the same obj_id are grouped and sent as a batch to the subscriber(s).

[0114] In some embodiments, a MERGE(*) function can be used to change how aggregated message data is merged. The MERGE(*) function can return a recursive union of incoming messages over a period of time. The merge function may be used, for example, to track location data for an object, and the subscriber is interested in the most recent values for all key/value pairs contained in a set of aggregated messages. The following statement shows an exemplary syntax for the MERGE(*) function:

SELECT [expr] [name,]MERGE(*)[*] [AS name] [FROM expr] [WHERE expr]
[HAVING expr] GROUP BY name

[0115] The following examples illustrate how the MERGE(*) function may be applied within query instructions to various types of channel messages. In the following examples, the Filter column shows the filter value included in the query instructions as part of a subscription request as the FILTER field. The Message Data column lists the Input channel message data and the resulting message data sent to the subscriber as Output. The filter returns the most recent values of the keys identified in the input messages, with the string MERGE identified as the column

name in the output message data. The first example below shows the MERGE(*) function in a filter with a wildcard, for the message data is returned using the keys from the input as column names in the output.

Filter	Message Data
SELECT MERGE(*)	Input <pre>{ "name": "art", "age": 10 }, { "name": "art", "age": 11, "items": [0] }</pre> Output <pre>{ "MERGE": { "name": "art", "age": 11, "items": [0] } }</pre>

The next example illustrates the use of the MERGE(*) function in a filter using a wildcard and the “AS” statement with a value of MERGE. The output data includes MERGE as the column name.

Filter	Message Data
SELECT MERGE(*).*	Input <pre>{ "name": "art", "age": 12, "items": [0], "skills": { "work": ["robots"] } }, { "name": "art", "age": 13, "items": ["car"],</pre>

	<pre>"skills": { "home": ["cooking"] }</pre> Output <pre>{ "name": "art", "age": 13, "items": ["car"], "skills": { "work": ["robots"], "home": ["cooking"] } }</pre>
SELECT MERGE(top.*) AS merge	Input <pre>{"top": {}, "garbage": 0}, {"top": {"name": "art", "eyes": "blue"}}, {"top": {"name": "joy", "height": 170}}</pre> Output <pre>{"merge": {"name": "joy", "eyes": "blue", "height": 170}}</pre>

[0116] Generally, for aggregative functions and for filters that only include a SELECT(expr) statement, only the latest value for any JSON key in the message data from the last message received can be stored and returned. Therefore, if the most recent message received that satisfies the filter statement is missing a key value identified in a previously processed message, that value is not included in the aggregate, which could result in data loss. However, filters that also include the MERGE(*) function can retain the most recent value for all keys that appear in messages to an unlimited JSON object depth. Accordingly, the most recent version of all key values can be retained in the aggregate.

[0117] The MERGE(*) function can be used to ensure that associated values for all keys that appear in any message during the aggregation period also appear in the final aggregated message. For example, a channel may track the physical location of an object in three dimensions: x, y, and z. During an aggregation period of one second, two messages are published to the channel, one having only two parameters: OBJ{x:1, y:2, z:3} and OBJ{x:2, y:3}. In the second message, the z value did not change and was not included in the second message. Without the MERGE(*) function, the output result would be OBJ{x:2, y:3}. Because the z value was not present in the last message in the aggregation period, the z value was not included in the final aggregate. However, with the MERGE(*) function, the result is OBJ{x:2, y:3, z:3}.

[0118] The following table shows one set of rules that may be used to aggregate data in messages, depending on the type of data. For arrays, elements need not be merged, but instead JSON values can be overwritten for the array in the aggregate with the last array value received.

Type of JSON Data	Data to Aggregate {msg1}, {msg2}	Without MERGE(*)	With MERGE(*)
Additional key/value	{a: 1, b: 2} , {c:3}	{c: 3}	{a:1, b:2, c:3}
Different value datatype	{a:2}, {a: "2"}	{a:"2"}	{a:"2"}
Missing key/value	{a:2}, {}	{a:2}	{a:2}
null value	{a:2}, {a:null}	{a:null}	{a:null}
Different key value	{a:{b:1}}, {a:{c:2}}	{a:{c:2}}	{a:{b:1, c:2}}
Arrays	{a:[1, 2]}, {a:[3, 4]}	{a:[3, 4]}	{a:[3, 4]}

[0119] The query instructions can be comprised of one or more suitable filter commands, statements, functions, or syntax. For purposes of illustration and not limitation, in addition to the SELECT and MERGE functions, the query instructions can include filter statements or functions, such as, for example, ABS(expr), AVG(expr), COALESCE(a[, b...]), CONCAT(a[, b...]), COUNT(expr), COUNT_DISTINCT(expr), IFNULL(expr1, expr2), JSON(expr), MIN(expr[, expr1,...]), MAX(expr[, expr1,...]), SUBSTR(expr, expr1[, expr2]), SUM(expr), MD5(expr), SHA1(expr), FIRST_VALUE(expr) OVER (ORDER BY expr1), and/or LAST_VALUE(expr) OVER (ORDER BY expr1), where “expr” can be any suitable expression

that is capable of being processed by a filter statement or function, such as, for example, a SQL or SQL-like expression. Other suitable filter commands, statements, functions, or syntax are possible for the query instructions.

[0120] According to the present invention, non-filtered queries can translate to an immediate copy of the message to the subscriber, without any JSON or other like processing. Queries that include a SELECT filter command (without aggregation) can translate into an immediate filter. In instances in which the messages are formatted using JSON, each message may be individually parsed and any WHERE clause may be executed directly on the individual message as it arrives, without the need for creating indices or other temporary data structures. If the messages pass the WHERE clause filter, the SELECT clause results in a filtered message that can be converted back to its original format or structure (e.g., JSON) and sent to the subscriber.

[0121] Aggregative functions, such as, for example, COUNT(), SUM(), AVG(), and the like, can translate into an immediate aggregator. In instances in which the messages are formatted using JSON, each message may be individually parsed and any WHERE clause may be executed directly on the individual message as it arrives, without the need for creating indices or other temporary data structures. If a WHERE clause is evaluated, messages passing such criteria are aggregated (e.g., aggregates in the SELECT clause are executed, thereby accumulating COUNT, SUM, AVG, and so forth) using the previous accumulated value and the value from the individual message. Once per aggregation period (e.g., every 1 second), the aggregates are computed (e.g., $AVG = SUM/COUNT$), and the SELECT clause outputs the aggregated message, which can be converted to its original format or structure (e.g., JSON) and sent to the subscriber.

[0122] More complex aggregative functions, such as, for example, GROUP BY, JOIN, HAVING, and the like, can be translated into a hash table aggregator. Unlike SELECT or other like functions that can use a constant memory, linearly expanding memory requirements can be dependent upon the results of the GROUP BY clause. At most, grouping by a unique value (e.g., SSN, etc.) can result in a group for each individual message, but in most cases grouping by a common data element (e.g., user_id or other repeating value) can result in far fewer groups. In practice, each message is parsed (from its JSON format, for example). The WHERE clause can be executed directly on the individual message as it arrives, without creating indices or other temporary structures. If the WHERE clause is satisfied, the GROUP BY expressions can be computed directly and used to build a hash key for the group. The aggregative functions in the SELECT clause can be executed, accumulating COUNT, SUM, AVG, or other functions using the previous accumulated value specific for the hash key (group) and the value from the individual message. Once per aggregation period (e.g., every 1 second), the aggregates are

computed (e.g., $AVG = SUM/COUNT$) for each hash key (group), and the SELECT clause can output the aggregated message for each hash key to be converted back to its original format or structure (e.g., JSON) and sent to the subscriber (e.g., one message per hash key (group)).

[0123] In embodiments in which the aggregation period is limited (e.g., 1 second - 60 seconds) and the network card or other hardware/throughput speeds may be limited (e.g., 10/gbps), the overall maximal memory consumption can be calculated as time * speed (e.g., 1GB per second, or 60GB per minute). Hence, the upper bound is independent of the number of subscribers. In certain implementations, each message only need be parsed once (e.g., if multiple filters are set by multiple clients) and only if needed based on the query instructions, as an empty filter does not require parsing the message.

[0124] Referring to FIG. 7A, subscriptions can include a “period” parameter, generally defined in, for example, seconds and in some embodiments can range from 1 to 60 seconds, although other time increments and time ranges are possible. The period parameter(s) can be purely sequential (e.g., ordinal) and/or time-based (e.g., temporal) and included in the self-described data and therefore available for querying, aggregation, and the like. For example, FIG. 7A illustrates the filter process according to the present invention for the first three seconds with a period of 1 second. In the present example, the subscription starts at $t=0$. The filter created from the query instructions is applied against all messages received during each 1-second period (e.g., one message at a time). The results for each period are then batched and forwarded to the subscriber. Depending on the query instructions used, the messages can be aggregated using the aggregation functions discussed previously before the message data is sent to the subscriber.

[0125] In some cases, the process defaults to sending only new, incoming messages that meet the query instructions on to the subscriber. However, a subscriber can subscribe with history and use a filter, such that the first message or messages sent to the subscriber can be the historical messages with the filter applied. Using the period of `max_age` and/or a “next” parameter provides additional functionality that allows for retrieval and filtering of historical messages.

[0126] More particularly, a `max_age` parameter included with the query instructions can facilitate the retrieval of historical messages that meet this parameter. FIG. 7B illustrates an example of a `max_age` parameter of 2 seconds (with a period of 1 second) that is provided with the query instructions. The filter created from the query instructions is applied to the historical messages from the channel that arrived from $t-2$ through $t=0$ ($t=0$ being the time the subscription starts), and to the messages that arrived in the first period (from $t=0$ to $t+1$). These messages can be sent in a single batch to the subscriber (as Group 1). The filter is applied to each message in each subsequent period (e.g., from $t+1$ to $t+2$ as Group 2) to batch all messages that meet the query instructions within that period. Each batch is then forwarded on to the subscriber.

[0127] When a subscriber subscribes with a “next” parameter to a channel with a filter, the filter can be applied to all messages from the next value up to the current message stream position for the channel, and the results can be sent to the subscriber in, for example, a single batch. For example, as illustrated in FIG. 7C, a next parameter is included with the query instructions (with a period of 1 second). The next parameter instructs the process to apply the filter created from the query instructions to each message from the “next position” up through the current stream position (e.g., up to $t=0$) and to the messages that arrived in the first period (from $t=0$ to $t+1$). These messages can be sent in a single batch to the subscriber (as Group 1). The filter is applied to each message in each subsequent period (e.g., from $t+1$ to $t+2$ as Group 2) to batch all messages that meet the query instructions within that period. Each batch is then forwarded on the subscriber.

[0128] When a subscriber subscribes with a next parameter, chooses to receive historical messages on a channel, and includes a filter in the subscription, the subscriber can be updated to the current message stream position in multiple batches. FIG. 7D illustrates an example of a max_age parameter of 2 seconds (with a period of 1 second) and a next parameter that can be combined into one set of query instructions. The filter created from the query instructions is applied to the historical messages from the channel that arrived from the end of the history to the “next” value of the subscription (i.e., from 2 seconds before the next value up to the next value), to the messages from the next value to the current stream position (e.g., up to $t=0$), and to the messages that arrived in the first period (from $t=0$ to $t+1$). These messages can be sent in a single batch to the subscriber (as Group 1). The filter is applied to each message in each subsequent period (e.g., from $t+1$ to $t+2$ as Group 2) to batch all messages that meet the query instructions within that period. Each batch is then forwarded on the subscriber. Consequently, historical messages can be combined with messages that start at a particular period indicator and batched for transmission to the subscriber.

[0129] The query instructions can define how one or more filters can be applied to the incoming messages in any suitable manner. For example, the resulting filter(s) can be applied to any or all messages arriving in each period, to any or all messages arriving across multiple periods, to any or all messages arriving in select periods, or to any or all messages arriving on a continuous or substantially continuous basis (e.g., without the use of a period parameter such that messages are not retained before returning them to the subscriber). Such filtered messages can be batched in any suitable manner or sent individually (e.g., one message at a time) to subscribers. In particular, the filtered messages can be sent to the subscriber in any suitable format or syntax. For example, the following illustrates the structure of several fields of a sample channel PDU that contains the message results from a filter request:

```

{
    "action": "channel/data",
    "body": {
        "channel": ChannelName
        "next": ChannelStreamPosition
        "messages": [ChannelData]+      // Can be one or more messages
    }
}

```

In the above channel PDU, the “channel” field can be a value (e.g., string or other appropriate value or designation) of the channel name to which the subscriber has subscribed. The “next” field can provide the channel stream position of the batch of messages returned in the channel PDU. The “messages” field provides the channel data of the messages resulting from application of the specified filter. One or more messages can be returned in the “messages” field in such a channel PDU. It is noted that a channel PDU can include any other suitable fields, parameters, values, or data.

[0130] FIG. 8 is a flowchart of an example method for applying query instructions to published messages for publishers and subscribers of a messaging system. The method can be implemented using, for example, an MX node (e.g., MX node 204, MX node 461) and a Q node (e.g., Q node 212, Q node 208) of the messaging system 100, for example. The method begins by receiving query instructions from one or more subscribers (block 802). The one or more subscribers are subscribed to a channel of a plurality of channels. The query instructions may be cached and implemented in real time, or, in some instances, stored at the respective message nodes (e.g., MX node 204). Messages are received from one or more publishers (Block 804). Each message is associated with a particular one of the plurality of channels. The query instructions are then applied to the messages for the channel as the messages are received (Block 806). The messages resulting from the application of the query instructions are sent to the corresponding subscribers (Block 808). The messages received by the subscribers from the channel are thereby limited to those that satisfy the query instructions. In certain embodiments, as will be further appreciated with respect to FIGS. 9-15, the present embodiments may include system architecture for counting objects on multiple roads using a pan, tilt, and zoom camera.

[0131] In accordance with the present techniques, FIG. 9 is a diagram of an example system architecture 900 that may be utilized to implement, for example, the present counting objects on multiple roads using a pan, tilt, and zoom camera techniques. In accordance with at least some of the present embodiments, the system architecture 900 may include one or more global

network services, and, in some embodiments, may correspond to the system 100 discussed above with respect to FIGS. 1A and 1B. For example, as depicted, the system architecture 900 may include a server electronic device 901 (or numerous server electronic devices 901) and a number of client electronic devices 916, 918, 920, 922, and 924 that may be in communication with the server electronic device 901. The server electronic device 901 may support the PubSub communication pattern, as described earlier in reference to FIGS. 1A through 5D.

[0132] In some embodiments, the server electronic device 901 may be referred to as a PubSub system or a PubSub messaging system. As illustrated, the server electronic device 901 may include, for example, a channel 902, a channel 904, a channel 906, a channel 908, a channel 910, a channel 912, and a channel 914, although any suitable number of channels can be supported by the server electronic device 901. The messages published to channels 902, 904, 906, 908, 910, 912, and 914 (e.g., channel streams) may be divided into streamlets which may be stored within Q nodes or one more databases of the server electronic device 901, as generally described, for example, earlier in reference to FIGS. 1A through 5D. C nodes of the messaging system may be used to offload data transfers from one or more Q nodes (e.g., to cache some of the streamlets stored in the Q nodes). As will be further appreciated with respect FIG. 10, the server electronic device 901 may include detector, tracker, counting, and line logic, which may be utilized to count objects on multiple roads using a pan, tilt, and zoom camera.

[0133] In certain embodiments, the client electronic devices 916, 918, 920, 922, and 924 may establish respective persistent connections (e.g., TCP communications or other similar communications channels) to one or more MX nodes. The one or more MX nodes may serve as termination points for these connections, as described earlier in reference to FIGS. 1A through 5D. As further illustrated, each of the client electronic devices 916, 918, 920, 922, and 924 may include one or more respective application components 926, 928, 930, 932, and 934, which may, for example, allow users to subscribe to and publish to the channels 902, 904, 906, 908, 910, 912, and 914 of the server electronic device 901. In some embodiments, the client electronic devices 916, 918, 920, 922, and 924 may include one or more monitoring or other data collection electronic devices (e.g., a pan, tilt, and zoom camera) that may collect, for example, real-time or near real-time video streams (e.g., video frames) to be transmitted to the server electronic device 901.

[0134] FIG. 10A is a diagram of an example system 1000 for counting objects on multiple roads using a pan, tilt, and zoom camera in a messaging system, in accordance with the present techniques. The system 1000 may include detector logic 1002, tracker logic 1004, counter logic 1006, and line logic 1008. In another embodiment, system 1000 may include ensemble logic 1010.

[0135] In one embodiment, detector logic 1002 may include one or more detectors (e.g., D1 1003a, D2 1003b, and D3 1003c). While three detectors are depicted, system 1000 may include any number of detectors. In one embodiment, detectors may be intelligent microservices (e.g., “bots”) in a PubSub system, as described herein. As described herein, intelligent microservices may be server-side applications that operate on a data stream (e.g., a live data stream) of a channel in a PubSub system. Publishers may create and add intelligent microservices to a marketplace. Subscribers may search for and access intelligent microservices associated with particular live data streams and/or channels via the marketplace. Each component and subcomponent of system 1000 may be implemented via intelligent microservices.

[0136] In one embodiment, detector logic 1002 takes a video frame (e.g., video frame 1012a from video frames 1012a-n) as input (1005), divides the frame 1012a into regions of interest (RoIs) and performs detection on each RoI. Detector logic 1002 may return indications of detected objects. In one embodiment, indications of detected objects may be depicted as bounding boxes over the detected objects in the video frame, as shown with respect to 1001 of FIG. 10B.

[0137] In one embodiment, each detector D1 1003a – D3 1003c may perform detection operations on a different RoI from the same video frame. In another embodiment, each detector D1 1003a – D3 1003c may perform detection operations on a different RoI from the different video frames. In yet another embodiment, each detector D1 1003a – D3 1003c may perform detection operations on the same RoI from a video frame. In another embodiment, detector logic 1002 takes a video frame (e.g., video frame 1012a from video frames 1012a-n) as input (1005), and does not divide the frame 1012a into RoIs, and instead performs detection on the complete video frame 1012a. In one embodiment, detectors D1 1003a – D3 1003c may be a variety of different types of detectors. For example, detectors based on machine learning may be used in conjunction with, or in place of, blob-type detectors. Additional details describing detectors are provided with respect to FIGS. 12 and 13A-D, below.

[0138] In one embodiment, multiple detectors (e.g., D1 1003a – D3 1003c) work in parallel on each incoming video frame 1012a-n to produce a set of detections. The results from each detector may then be merged to produce a single set of detections. Detectors may vary in terms of accuracy and speed, and thus a naive pooling of predictions can produce false positives. Advantageously, the present invention provides for the use of a cascading system, in which detectors sequentially mask the image RoIs where objects are detected. Such a cascading system may be implemented by ensemble logic 1010, and is described further with respect to FIGS. 13A-D.

[0139] In other embodiments, multiple detectors (e.g., D1 1003a – D3 1003c) do not work in parallel on each incoming video frame 1012a-n to produce detections. The results from each detector need not be merged to produce a single set of detections in this case. Thus, ensemble logic 1010 (and the cascading system described with respect to FIGS. 13A-D) may be optional.

[0140] Tracker logic 1004 receives the object detections from the current video frames (e.g., video frames that are being received in real, or substantially-real, time) and matches them to current active trajectories. In one embodiment, a trajectory is a series of locations of an object tracked across consecutive video frames. An active trajectory may be a trajectory that is currently being tracked. The tracker logic 1004 may run in real-time and is able to track hundreds, or thousands, of objects per video frame. In one embodiment, tracker logic 1004 may execute a variety of tracking algorithms. One such example algorithm is provided:

- For current trajectories, tracker logic 1004 predicts the next position using a sequential discounted filter of order 3 (although filters of order more than 3 will also work equally well, they may add computational time). If, for a given trajectory, the displacement vector is $s(t-1)$ between the locations in the last two frames, then:

$$s(t) = (\alpha * s(t-1) + \alpha/2 * s(t-2) + \alpha/4 * s(t-3)) / (7\alpha/4)$$

- For a known delta change in camera angle, tracker logic 1004 adjusts the predicted positions. For a displacement ‘z’ in camera angle, adjusted position may be calculated as:

$$s(t) = s(t) - z$$

- For camera zoom, tracker logic 1004 may adjust the predicted positions using a suitable perspective transform.
- Tracker logic 1004 may match the detections to the predicted trajectory positions using a nearest neighbor algorithm.
- For detections for which no active trajectory is found, tracker logic may instantiate a new trajectory with the rectangle detection as the first point in the trajectory.
- Active trajectories for which no detection was found may be marked dormant if no detection is found for some number (e.g., 5) of consecutive frames.
- Tracker logic 1004 may add the matched detections to their respective trajectories. If rectangles are used to indicate the detected objects, the size of the rectangle added may be changed to be the average of the rectangle size of the last frame and the rectangle size detected in the current frame. This averaging may suppress random jumps in rectangle sizes.

FIG. 10C shows example trajectories (e.g., 1009) and rectangle detections (e.g., 1007).

[0141] In one embodiment, tracker logic 1004 may determine a directionality for each of the trajectories. Trajectory directionalities may represent the direction of travel of the trajectories. Tracking logic 1004 may identify roads and pathways based on the trajectories and their corresponding directionalities, by identifying sets of trajectories (e.g., with corresponding directionalities) that are similar. In one embodiment, a similarity threshold may be used to identify distinct roads and/or pathways. For example, a road may be identified if a threshold number of trajectories having approximately the same directionality (e.g., within a directionality threshold) are determined to be within a specified angle threshold of each other. In one embodiment, trajectories that are within an angle threshold, but have opposite directionalities, may represent two sides of a road, with opposing traffic on each side. In such a scenario, two roads may be identified, one for each direction of traffic, and a separate flow cut line may be determined for each road.

[0142] Referring back to FIG. 10A, in one embodiment, counter logic 1006 counts the number of objects (trajectories) that cross determined flow cut lines. As described herein, a flow cut line is an imaginary line that cuts the flow of traffic in such a way that all trajectories on a road are near perpendicular to the line. In one embodiment, a separate flow cut line may be determined for each detected road or path.

[0143] In one embodiment, line logic 1008 may identify the directions of major traffic flow (e.g., identified visually as roads or paths), and calculate a flow cut line that is perpendicular to this direction of flow and cuts through all trajectories. The line logic 1008 may run periodically to find the flow cut lines, and may be succeeded by longer periods of object counting. In one embodiment, an algorithm used by line logic 1008 may provide a variety of information corresponding to the flow cut line, including, but not limited to: (1) the centroid of the flow cut line; (2) the slope of the flow cut line; and (3) the length of the flow cut line.

[0144] An example algorithm used by the line logic 1008 is provided below:

- Gather trajectories for any number of video frames (e.g., 500 frames). During the gathering phase, the algorithm can track objects and record the first and the last point of the trajectory. The vector from the first point to the last point may provide information about the flow direction, which may help to significantly reduce the dimensionality of the problem. The algorithm may also filter out any trajectories with less than some number of locations (e.g., 10 locations).
- Apply K-means clustering to this two-dimensional vector data after normalizing the vectors to unit magnitude. In one embodiment, the number of clusters may be decided based on the compactness of cluster assignment.

- The problem may be constrained by starting with any number of clusters (e.g., 4 clusters). If any two cluster centroids have an angle less than 90 degrees, the data may be re-clustered with three clusters. This process may be repeated until all cluster centroids have angles more than 90 degrees between them.
- The cluster centroids may provide the direction of maximum flow in the video. The slope of the flow cut line for each centroid can be calculated as the slope perpendicular to the centroid slope.
- Cluster assignments for all of the trajectories vectors may then be obtained. Line logic 1008 may create a set of all of the trajectory points for a given cluster, and generate the convex hull.
- The centroid of the convex hull provides the centroid of the flow cut line.
- The length of the flow cut line is given by the distance between the two points where the flow cut line intersects the convex hull. The line logic 1008 can add appropriate “padding” on both sides of the flow cut line to sufficiently extend the line beyond the convex hull boundaries.

An example of the application of the line logic 1008 to single 1101 and double 1103 lane views is illustrated in FIGS. 11A and 11B, respectively. FIG. 11A shows a single lane view 1101, including a convex hull 1102 and flow cut line 1104. FIG. 11B shows a double lane view 1103, including a first convex hull 1106 and a first flow cut line 1108 of a first lane, and a second convex hull 1110 and a second flow cut line 1112 of a second lane.

[0145] In one embodiment, line logic 1008 may determine an optimal camera zoom ratio for system 100. For example, if the flow cut line is very close to the corners of an image, or is clipped by the edges of the image, it is possible that some trajectories may be unable to intersect, which may decrease counting accuracy. To alleviate such a problem, the camera can be zoomed out and the analysis may be repeated. Line logic 1008 may apply a random search algorithm to choose a zoom ratio that helps identify a flow cut line that is not clipped by the edges of the image and is located a threshold distance away from the edges. For purposes of illustration and not limitation, the present system may use a threshold of, for example, 10% of the width and height of the image to determine the proximity of a line to the edges.

[0146] In one embodiment, processing logic of system 1000 may determine an optimal camera angle for a PTZ camera to ensure maximum accuracy and minimum number of false positives.

This may be done by monitoring a coverage metric, which may be defined as follows:

Coverage = (Motion pixels masked by detector) / (Pixels outside of the detection mask).

This coverage metric may be maximized via a random search algorithm, in which many different camera angles are tried within the permissible scope of the PTZ camera and the angle with the

best coverage is selected. In one embodiment, line logic 1008 determines the optimal camera angle. In other embodiments, other components or subcomponents of system 1000 make the determination.

[0147] FIG. 12 is a diagram of an example threshold binary map 1200 of motion information. In one embodiment, at least one of detectors D1 1003a – D3 1003c of FIG. 10A is blob detector. In one embodiment, this type of detector does not use a machine learning algorithm to detect objects. Instead, a blob detector may use motion information from video frames. Motion information may be obtained by subtracting successive video frames and applying suitable thresholding to identify moving “blobs” of objects. Such a threshold may be calculated so as to remove any points with low intensity. For example, any points with magnitude less than, for example, 10% of the maximum magnitude may be removed, although other appropriate values for the threshold percentage can be used. Such a detector may not know what the blobs represent, but it can draw rectangles around these blobs and use them as detections. FIG. 12 shows an example binary map of motion, including several blobs (e.g., blob 1202).

[0148] In one embodiment, the blob detector tracks objects moving between video frames and may not differentiate between various types of objects, and hence may detect a number of false positives. Advantageously, a variety of techniques may be used to increase the accuracy of such a blob detector. For example, the blob detector logic may:

- Reject blobs with an area smaller than ‘x’ % (e.g., a size threshold) of the total image (e.g., video frame) area. For example, in a 640x360 image with $x=0.02$, the processing logic may reject blobs smaller than an area of 460. The present algorithm can work for any other values of x.
- Reject blobs that have a density ‘d’ that is below a density threshold. In one embodiment, density is defined as the percentage of nonzero pixels inside a convex hull that envelops a blob. In one example, blobs with a density below $d=0.4$ can be rejected. Any suitable value of density ‘d’ may be used.
- Reject blobs based on an aspect ratio threshold. Aspect ratio is defined by the ratio of width to height. In one embodiment, processing logic may reject blobs that are very thin or very wide. In one embodiment, an acceptable range for the aspect ratio is about 0.75 to about 1.5. In other embodiments, other ranges may be used.
- Reject blobs based on orientation. The angle of the ellipse that encapsulates the blob is a measure of the direction in which the blob is moving. If the angle of the ellipse is significantly different from the angle of the flow of direction of traffic, the blob may be rejected. For example, if the angle of an ellipse is more than 45 degrees (the orientation

threshold) away from the flow of traffic, processing logic may reject the blob. In other embodiments, other orientation thresholds may be used.

[0149] FIG 13A is a diagram of an example cascading detection system 1300. The present disclosure provides for multiple detectors that work on incoming frames in parallel to produce a set of rectangle detections. After the multiple detectors process individual video frames, the results from each detector may be merged to produce a single set of detections. As described above, detectors vary in terms of accuracy and speed, and thus a naive pooling of predictions can produce an unacceptable number of false positives. Advantageously, the present disclosure describes a cascading system, in which detectors sequentially mask image RoIs in which objects are detected.

[0150] For purposes of illustration and not limitation, for the 3-detector system D1 1302, D2 1304, and D3 1306, the cascade ensemble first uses detector D1 1302 to generate rectangle detections. Detector D2 1304 is then evaluated only in RoIs in which D1 1302 did not detect anything. Similarly, D3 1306 is evaluated for RoIs in which D1 1302 and D2 1304 did not detect anything. Thus, each detector feeds a detection mask into the next detector. Advantageously, such a configuration significantly reduces the computational complexity of each of the detectors, because each does not evaluate all of the possible RoIs in each of the video frames. It is noted that any suitable number of detectors can be used in the cascaded ensemble, although only three are depicted.

[0151] FIGS. 13B-D are diagrams of an example of the application of a cascading detection system. A special case of the cascade ensemble of FIG. 13A may occur when D3 1306 is a blob detector, as described with respect to FIG. 12. In such a case, D1 1302 and D2 1304 may produce a detection mask, and D3 1306 is a blob detector that identifies any motion information outside of this detection mask. FIGS. 13B-D illustrate such a configuration, where the FIG. 13B is the raw binary map of motion (e.g., similar to that in FIG. 12). FIG. 13C shows the detected RoIs (e.g., 1308) produced by D1 1302 and D2 1304. FIG. 13D shows the binary map masked by the detected RoIs. In this example, D1 1302 and D2 1304 missed detecting the car 1310, which can now be tracked by the blob detector (e.g., the larger blob in FIG. 13D).

[0152] FIG. 14 is a flowchart of an example method for counting objects on multiple roads using a pan, tilt, and zoom camera in a messaging system. In certain embodiments, the method 1400 may be performed by processing logic that may include hardware such as one or more computer processing devices, software (e.g., instructions running/executing on a computer processing device), firmware (e.g., microcode), or a combination thereof, such as the server electronic device 901 discussed above with respect to FIG. 9.

[0153] The method 1400 may begin with processing logic receiving a plurality of video frames from a camera. In one embodiment, the camera may be a pan tilt, and zoom camera. In another embodiment, the camera may be a standard camera that does not have pan, tilt, and/or zoom capabilities. In one embodiment, the PTZ camera is operatively connected to a publisher device in a PubSub system, as described herein. At block 1404, processing logic may detect a plurality of objects in the plurality of video frames. In one embodiment, the objects are all objects that are moving in the plurality of video frames. To detect the plurality of objects, processing logic may perform blob detection on the plurality of video frames. Advantageously, to increase the accuracy of the blob detection (e.g., to end up with only a subset of the detected objects that are vehicles and/or pedestrians, for example), processing logic may remove a subset of the plurality of objects based on at least one of: a blob area threshold, a blob density threshold, a blob aspect ratio threshold, or a blob orientation threshold as described herein with respect to FIG. 12. This technique may be used to ensure that a remaining subset of the plurality of objects includes only vehicles and/or pedestrians.

[0154] In another embodiment, processing logic may divide each of the plurality of video frames into regions of interest and perform object detection by a first detector on each of the regions of interest. In one embodiment, the first detector is a blob detector. In other embodiments, the first detector may be a machine-learning based detector. In yet other embodiments, any other suitable detector-type may be used as the first detector. Processing logic may continue by determining that no objects were detected in a subset of the regions of interest and performing object detection by a second detector on only the subset of the regions of interest. In one embodiment, the second detector is a blob detector. In other embodiments, the second detector may be a machine-learning based detector. In yet other embodiments, any other suitable detector-type may be used as the second detector. As described herein, the first detector and/or the second detector may be intelligent microservices of a PubSub system. Additional details regarding cascading detectors are provided with respect to FIGS. 13A-D.

[0155] At block 1406, processing logic (e.g., a computer processing device) determines a plurality of trajectories, each of the plurality of trajectories corresponding to a single one of the plurality of objects, respectively. In one embodiment, a trajectory is a series of locations of an object tracked across consecutive video frames. At block 1408, processing logic (e.g., the computer processing device) identifies a first thoroughfare from a first subset of the plurality of trajectories. In one embodiment, the first subset of the plurality of trajectories includes only objects that have the same (e.g., within a directionality threshold) directionalities. At block 1410, processing logic (e.g., the computer processing device) determines a first number of the plurality of objects that cross a first flow cut line on the first thoroughfare. In one embodiment,

processing logic may calculate the flow cut line, as described herein. For example, processing logic may calculate a centroid, a slope, and a length of the flow cut line. The flow cut line may be perpendicular to the first subset of trajectories in one embodiment.

[0156] Processing logic may optimize the position of a flow cut line by modifying a camera zoom of the PTZ camera. For example, processing logic may determine that the flow cut line is within a threshold of an edge of the plurality of video frames and modify a camera zoom setting of the camera based on the determination. Processing logic may also optimize an angle of the PTZ camera to maximize detection. In one embodiment, processing logic may determine a coverage metric for the camera and modify a camera angle of the camera based on the coverage metric.

[0157] Processing logic may repeat the above operations to determine any number of roads/pathways, and count any number of objects on them. For example, in one embodiment, processing logic may identify a second thoroughfare from a second plurality of trajectories, wherein the second plurality of trajectories has a second corresponding directionality that is different than the first corresponding directionality of the first subset (e.g., two sides of a freeway, with traffic moving in opposite directions on each side). Processing logic may then determine a second number of the plurality of objects that cross a second flow cut line on the second thoroughfare. Processing logic may continue to track objects in accordance with the above techniques by receiving a new video frame from the camera, detecting a new object in the new video frame, and matching the new object to the first subset of the plurality of trajectories.

[0158] FIG. 15 is a block diagram of an example computing device 1500 that may perform one or more of the operations described herein, in accordance with the present embodiments. The computing device 1500 may be connected to other computing devices in a LAN, an intranet, an extranet, and/or the Internet. The computing device 1500 may operate in the capacity of a server machine in client-server network environment or in the capacity of a client in a peer-to-peer network environment. The computing device 1500 may be provided by a personal computer (PC), a set-top box (STB), a server, a network router, switch or bridge, or any machine capable of executing a set of instructions (sequential or otherwise) that specify actions to be taken by that machine. Further, while only a single computing device 1500 is illustrated, the term “computing device” shall also be taken to include any collection of computing devices that individually or jointly execute a set (or multiple sets) of instructions to perform the methods discussed herein.

[0159] The example computing device 1500 may include a computer processing device (e.g., a general purpose processor, ASIC, etc.) 1502, a main memory 1504, a static memory 1506 (e.g., flash memory and a data storage device 1508), which may communicate with each other via a bus 1510. The computer processing device 1502 may be provided by one or more general-

purpose processing devices such as a microprocessor, central processing unit, or the like. In an illustrative example, computer processing device 1502 may comprise a complex instruction set computing (CISC) microprocessor, reduced instruction set computing (RISC) microprocessor, very long instruction word (VLIW) microprocessor, or a processor implementing other instruction sets or processors implementing a combination of instruction sets. The computer processing device 1502 may also comprise one or more special-purpose processing devices such as an application specific integrated circuit (ASIC), a field programmable gate array (FPGA), a digital signal processor (DSP), network processor, or the like. The computer processing device 1502 may be configured to execute the operations described herein, in accordance with one or more aspects of the present disclosure, for performing the operations and blocks discussed herein.

[0160] The computing device 1500 may further include a network interface device 1512, which may communicate with a network 1514. The data storage device 1508 may include a machine-readable storage medium 1516 on which may be stored one or more sets of instructions, *e.g.*, instructions for carrying out the operations described herein, in accordance with one or more aspects of the present disclosure. Instructions implementing module 1518 may also reside, completely or at least partially, within main memory 1504 and/or within computer processing device 1502 during execution thereof by the computing device 1500, main memory 1504 and computer processing device 1502 also constituting computer-readable media. The instructions may further be transmitted or received over the network 1514 via the network interface device 1512.

[0161] While machine-readable storage medium 1516 is shown in an illustrative example to be a single medium, the term “computer-readable storage medium” should be taken to include a single medium or multiple media (*e.g.*, a centralized or distributed database and/or associated caches and servers) that store the one or more sets of instructions. The term “computer-readable storage medium” shall also be taken to include any medium that is capable of storing, encoding or carrying a set of instructions for execution by the machine and that cause the machine to perform the methods described herein. The term “computer-readable storage medium” shall accordingly be taken to include, but not be limited to, solid-state memories, optical media and magnetic media.

[0162] Embodiments of the subject matter and the operations described in this specification can be implemented in digital electronic circuitry, or in computer software, firmware, or hardware, including the structures disclosed in this specification and their structural equivalents, or in combinations of one or more of them. Embodiments of the subject matter described in this specification can be implemented as one or more computer programs, *i.e.*, one or more modules

of computer program instructions, encoded on computer storage medium for execution by, or to control the operation of, data processing apparatus. Alternatively or in addition, the program instructions can be encoded on an artificially-generated propagated signal, e.g., a machine-generated electrical, optical, or electromagnetic signal that is generated to encode information for transmission to suitable receiver apparatus for execution by a data processing apparatus. A computer storage medium can be, or be included in, a computer-readable storage device, a computer-readable storage substrate, a random or serial access memory array or device, or a combination of one or more of them. Moreover, while a computer storage medium is not a propagated signal, a computer storage medium can be a source or destination of computer program instructions encoded in an artificially-generated propagated signal. The computer storage medium can also be, or be included in, one or more separate physical components or media (e.g., multiple CDs, disks, or other storage devices).

[0163] The operations described in this specification can be implemented as operations performed by a data processing apparatus on data stored on one or more computer-readable storage devices or received from other sources.

[0164] The term “computer processing device” encompasses all kinds of apparatus, devices, and machines for processing data, including by way of example a programmable processor, a computer, a system on a chip, or multiple ones, or combinations, of the foregoing. Although referred to as a computer processing device, use of the term also encompasses embodiments that include one or more computer processing devices. The computer processing device can include special purpose logic circuitry, e.g., an FPGA (field programmable gate array) or an ASIC (application-specific integrated circuit). The computer processing device can also include, in addition to hardware, code that creates an execution environment for the computer program in question, e.g., code that constitutes processor firmware, a protocol stack, a database management system, an operating system, a cross-platform runtime environment, a virtual machine, or a combination of one or more of them. The computer processing device and execution environment can realize various different computing model infrastructures, such as web services, distributed computing and grid computing infrastructures.

[0165] A computer program (also known as a program, software, software application, script, or code) can be written in any form of programming language, including compiled or interpreted languages, declarative, procedural, or functional languages, and it can be deployed in any form, including as a stand-alone program or as a module, component, subroutine, object, or other unit suitable for use in a computing environment. A computer program may, but need not, correspond to a file in a file system. A program can be stored in a portion of a file that holds other programs or data (e.g., one or more scripts stored in a markup language resource), in a

single file dedicated to the program in question, or in multiple coordinated files (e.g., files that store one or more modules, sub-programs, or portions of code). A computer program can be deployed to be executed on one computer or on multiple computers that are located at one site or distributed across multiple sites and interconnected by a communication network.

[0166] The processes and logic flows described in this specification can be performed by one or more programmable processors executing one or more computer programs to perform actions by operating on input data and generating output. The processes and logic flows can also be performed by, and apparatus can also be implemented as, special purpose logic circuitry, e.g., an FPGA (field programmable gate array) or an ASIC (application-specific integrated circuit).

[0167] Processing devices suitable for the execution of a computer program include, by way of example, both general and special purpose microprocessors, and any one or more processors of any kind of digital computer. Generally, a processing device will receive instructions and data from a read-only memory or a random access memory or both. The essential elements of a computer are a processor for performing actions in accordance with instructions and one or more memory devices for storing instructions and data. Generally, a computer will also include, or be operatively coupled to receive data from or transfer data to, or both, one or more mass storage devices for storing data, e.g., magnetic disks, magneto-optical disks, optical disks, or solid state drives. However, a computer need not have such devices.

[0168] Moreover, a computer can be embedded in another device, e.g., a smart phone, a mobile audio or video player, a game console, a Global Positioning System (GPS) receiver, or a portable storage device (e.g., a universal serial bus (USB) flash drive), to name just a few. Devices suitable for storing computer program instructions and data include all forms of non-volatile memory, media and memory devices, including, by way of example, semiconductor memory devices, e.g., EPROM, EEPROM, and flash memory devices; magnetic disks, e.g., internal hard disks or removable disks; magneto-optical disks; and CD-ROM and DVD-ROM disks. The processing device and the memory can be supplemented by, or incorporated in, special purpose logic circuitry.

[0169] To provide for interaction with a user, embodiments of the subject matter described in this specification can be implemented on a computer having a display device, e.g., a CRT (cathode ray tube) or LCD (liquid crystal display) monitor, for displaying information to the user and a keyboard and a pointing device, e.g., a mouse, a trackball, a touchpad, or a stylus, by which the user can provide input to the computer. Other kinds of devices can be used to provide for interaction with a user as well; for example, feedback provided to the user can be any form of sensory feedback, e.g., visual feedback, auditory feedback, or tactile feedback; and input from the user can be received in any form, including acoustic, speech, or tactile input. In addition, a

computer can interact with a user by sending resources to and receiving resources from a device that is used by the user; for example, by sending web pages to a web browser on a user's client device in response to requests received from the web browser.

[0170] Embodiments of the subject matter described in this specification can be implemented in a computing system that includes a back-end component, e.g., as a data server, or that includes a middleware component, e.g., an application server, or that includes a front-end component, e.g., a client computer having a graphical user interface or a Web browser through which a user can interact with an implementation of the subject matter described in this specification, or any combination of one or more such back-end, middleware, or front-end components. The components of the system can be interconnected by any form or medium of digital data communication, e.g., a communication network. Examples of communication networks include a local area network ("LAN") and a wide area network ("WAN"), an inter-network (e.g., the Internet), and peer-to-peer networks (e.g., ad hoc peer-to-peer networks).

[0171] The computing system can include clients and servers. A client and server are generally remote from each other and typically interact through a communication network. The relationship of client and server arises by virtue of computer programs running on the respective computers and having a client-server relationship to each other. In some embodiments, a server transmits data (e.g., an HTML page) to a client device (e.g., for purposes of displaying data to and receiving user input from a user interacting with the client device). Data generated at the client device (e.g., a result of the user interaction) can be received from the client device at the server.

[0172] A system of one or more computers can be configured to perform particular operations or actions by virtue of having software, firmware, hardware, or combinations of them installed on the system that in operation causes or cause the system to perform the actions. One or more computer programs can be configured to perform particular operations or actions by virtue of including instructions that, when executed by data processing apparatus, cause the apparatus to perform the actions.

[0173] While this specification contains many specific implementation details, these should not be construed as limitations on the scope of any inventions or of what may be claimed, but rather as descriptions of features specific to particular embodiments of particular inventions. Certain features that are described in this specification in the context of separate embodiments can also be implemented in combination in a single embodiment. Conversely, various features that are described in the context of a single embodiment can also be implemented in multiple embodiments separately or in any suitable subcombination. Moreover, although features may be described above as acting in certain combinations and even initially claimed as such, one or

more features from a claimed combination can in some cases be excised from the combination, and the claimed combination may be directed to a subcombination or variation of a subcombination.

[0174] Similarly, while operations are depicted in the drawings in a particular order, this should not be understood as requiring that such operations be performed in the particular order shown or in sequential order, or that all illustrated operations be performed, to achieve desirable results. In certain circumstances, multitasking and parallel processing may be advantageous. Moreover, the separation of various system components in the embodiments described above should not be understood as requiring such separation in all embodiments, and it should be understood that the described program components and systems can generally be integrated together in a single software product or packaged into multiple software products.

[0175] Thus, particular embodiments of the subject matter have been described. Other embodiments are within the scope of the following claims. In some cases, the actions recited in the claims can be performed in a different order and still achieve desirable results. In addition, the processes depicted in the accompanying figures do not necessarily require the particular order shown, or sequential order, to achieve desirable results. In certain implementations, multitasking and parallel processing may be advantageous.

CLAIMS

What is claimed is:

1. A method, comprising:
receiving a plurality of video frames from a camera;
detecting a plurality of objects in the plurality of video frames;
determining, by a computer processing device, a plurality of trajectories, each of the plurality of trajectories corresponding to a single one of the plurality of objects, respectively;
identifying, by the computer processing device, a first thoroughfare from a first subset of the plurality of trajectories; and
determining, by the computer processing device, a first number of the plurality of objects that cross a first flow cut line on the first thoroughfare.
2. The method of claim 1, wherein the first subset of the plurality of trajectories has a first corresponding directionality, the method further comprising:
identifying a second thoroughfare from a second plurality of trajectories, wherein the second plurality of trajectories has a second corresponding directionality that is different than the first corresponding directionality; and
determining a second number of the plurality of objects that cross a second flow cut line on the second thoroughfare.
3. The method of claim 1, wherein detecting the plurality of objects comprises performing blob detection on the plurality of video frames.
4. The method of claim 3, further comprising removing a subset of the plurality of objects based on at least one of: a blob area threshold, a blob density threshold, a blob aspect ratio threshold, or a blob orientation threshold, wherein a remaining subset of the plurality of objects comprises vehicles or pedestrians.
5. The method of claim 1, wherein detecting the plurality of objects comprises:
dividing each of the plurality of video frames into regions of interest; and
performing object detection by a first detector on each of the regions of interest.
6. The method of claim 5, further comprising:
determining that no objects were detected in a subset of the regions of interest; and

performing object detection by a second detector on only the subset of the regions of interest.

7. The method of claim 6, wherein at least one of the first detector and the second detector comprises intelligent microservices of a publish-subscribe system.
8. The method of claim 1, wherein the camera is a pan, tilt, and zoom camera, further comprising:
 - determining a coverage metric for the pan, tilt, and zoom camera; and
 - modifying a camera angle of the pan, tilt, and zoom camera based on the coverage metric.
9. The method of claim 1, further comprising calculating a centroid, a slope, and a length of the flow cut line, wherein the first flow cut line is perpendicular to the first subset of trajectories.
10. The method of claim 1, further comprising:
 - determining that the first flow cut line is within a threshold of an edge of the plurality of video frames; and
 - modifying a camera zoom of the camera based on the determining.
11. The method of claim 1, further comprising:
 - receiving a new video frame from the camera;
 - detecting a new object in the new video frame; and
 - matching the new object to the first subset of the plurality of trajectories.
12. A publish-subscribe system, comprising:
 - a computer processing device to:
 - receive a plurality of video frames from a camera;
 - detect a plurality of objects in the plurality of video frames;
 - determine a plurality of trajectories, each of the plurality of trajectories corresponding to a single one of the plurality of objects, respectively;
 - identify a first thoroughfare from a first subset of the plurality of trajectories; and
 - determine a first number of the plurality of objects that cross a first flow cut line on the first thoroughfare.

13. The system of claim 12, wherein to detect the plurality of objects the processing device is further to perform blob detection on the plurality of video frames.
14. The system of claim 13, the processing device further to remove a subset of the plurality of objects based on at least one of: a blob area threshold, a blob density threshold, a blob aspect ratio threshold, or a blob orientation threshold, wherein a remaining subset of the plurality of objects comprises vehicles or pedestrians.
15. The system of claim 12, wherein to detect the plurality of objects the processing device is further to:
 - divide each of the plurality of video frames into regions of interest; and
 - perform object detection by a first detector on each of the regions of interest.
16. The system of claim 15, the processing device further to:
 - determine that no objects were detected in a subset of the regions of interest; and
 - perform object detection by a second detector on only the subset of the regions of interest.
17. The system of claim 12, wherein the camera is a pan, tilt, and zoom camera, the processing device further to:
 - determine a coverage metric for the pan, tilt, and zoom camera; and
 - modify a camera angle of the pan, tilt, and zoom camera based on the coverage metric.
18. The system of claim 12, the processing device further to:
 - receive a new video frame from the camera;
 - detect a new object in the new video frame; and
 - match the new object to the first subset of the plurality of trajectories.
19. A non-transitory computer-readable storage medium including instructions that, when executed by a computer processing device, cause the computer processing device to:
 - receive a plurality of video frames from a camera;
 - detect a plurality of objects in the plurality of video frames;
 - determine, by the computer processing device, a plurality of trajectories, each of the plurality of trajectories corresponding to a single one of the plurality of objects, respectively;

identify, by the computer processing device, a first thoroughfare from a first subset of the plurality of trajectories; and

determine, by the computer processing device, a first number of the plurality of objects that cross a first flow cut line on the first thoroughfare.

20. The non-transitory computer-readable storage medium of claim 19, wherein to detect the plurality of objects the processing device is further to:

divide each of the plurality of video frames into regions of interest;
perform object detection by a first detector on each of the regions of interest;
determine that no objects were detected in a subset of the regions of interest; and
perform object detection by a second detector on only the subset of the regions of interest.

1/24




FIG. 1A




FIG. 1B

2/24




FIG. 2


300

FIG. 3A

4/24


350

FIG. 3B

5/24




FIG. 4A

450




FIG. 4B

490




FIG. 4C

8/24

500




FIG. 5A

9/24


550

FIG. 5B

10/24

570




FIG. 5C

11/24


580

FIG. 5D




FIG. 6




FIG. 7A




FIG. 7B




FIG. 7C





FIG. 7D

**FIG. 8**




FIG. 9

1000




FIG. 10A

18/24




FIG. 10B




FIG. 10C

1101



FIG. 11A

1103





FIG. 11B

1200



1202

FIG. 12

1300




FIG. 13A




FIG. 13B




FIG. 13C

1308

1310






FIG. 13D

1400**FIG. 14**

24/24

**FIG. 15**

INTERNATIONAL SEARCH REPORT

International application No

PCT/US2018/042840

A. CLASSIFICATION OF SUBJECT MATTER

INV. G06K9/00

ADD.

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

G06K G06T

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

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

EPO-Internal, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	ERHAN BAS ET AL: "Automatic Vehicle Counting from Video for Traffic Flow Analysis", INTELLIGENT VEHICLES SYMPOSIUM, 2007 IEEE, IEEE, PI, 1 June 2007 (2007-06-01), pages 392-397, XP031126976, ISBN: 978-1-4244-1067-5	1-6,9, 11-16, 18-20
Y	abstract	8,17
A	figures II, IV section section II ----- -/--	7,10



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents :

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

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

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

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

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

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

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

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

"&" document member of the same patent family

Date of the actual completion of the international search

9 October 2018

Date of mailing of the international search report

15/10/2018

Name and mailing address of the ISA/

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

Authorized officer

Martinez, Francis

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2018/042840

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	SHUBHAM JAIN ET AL: "Panoptes : servicing multiple applications simultaneously using steerable cameras", INFORMATION PROCESSING IN SENSOR NETWORKS, 18 April 2017 (2017-04-18), pages 119-130, XP055508292, 2 Penn Plaza, Suite 701 New York NY 10121-0701 USA DOI: 10.1145/3055031.3055085 ISBN: 978-1-4503-4890-4 abstract figure 1	8,17
A	----- ROGERIO S FERIS: "Chapter 3: Case Study: IBM Smart Surveillance System", 15 December 2010 (2010-12-15), INTELLIGENT VIDEO SURVEILLANCE: SYSTEMS AND TECHNO,, PAGE(S) 47 - 76, XP008147108, ISBN: 978-1-4398-1328-7 figure 3.1 section 3.3.2	7,12
A	----- ARUN HAMPAPUR ET AL: "Searching surveillance video", 2007 IEEE CONFERENCE ON ADVANCED VIDEO AND SIGNAL BASED SURVEILLANCE, AVSS 2007. 5-7 SEPT. 2007 - LONDON, UK, 1 September 2007 (2007-09-01), pages 75-80, XP055321785, DOI: 10.1109/AVSS.2007.4425289 ISBN: 978-1-4244-1696-7 section 3.3 -----	7,12