

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
11 September 2009 (11.09.2009)

PCT

(10) International Publication Number
WO 2009/109510 A1

(51) International Patent Classification:
H04L 12/58 (2006.01) H04L 12/18 (2006.01)
H04L 29/08 (2006.01)

(74) Agent: SEKAR, Anita; IBM United Kingdom Limited, Intellectual Property Law, Hursley Park, Winchester, Hampshire (GB).

(21) International Application Number:
PCT/EP2009/052272

(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, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(22) International Filing Date:
26 February 2009 (26.02.2009)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
08152198.1 3 March 2008 (03.03.2008) EP

(71) Applicant (for all designated States except US): INTERNATIONAL BUSINESS MACHINES CORPORATION [US/US]; New Orchard Road, Armonk, New York 10504 (US).

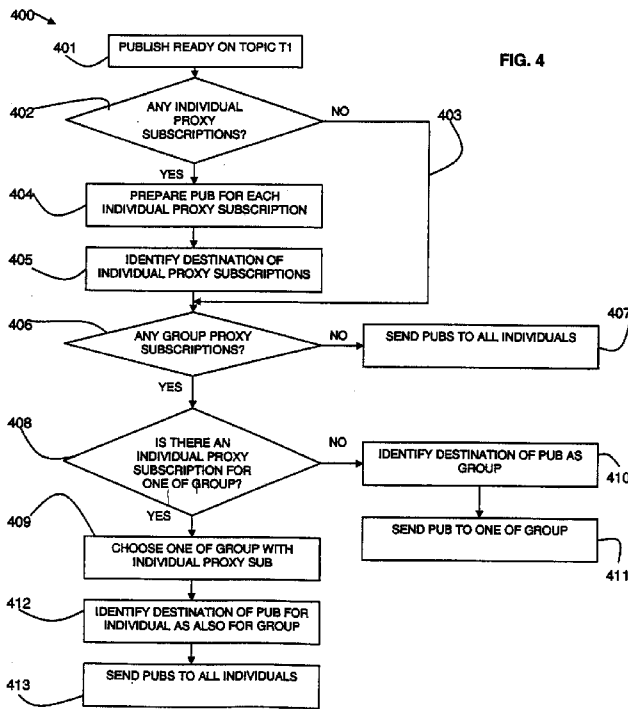
(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT,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

(72) Inventor; and

(75) Inventor/Applicant (for US only): DENNIS, Paul, Stephen [GB/GB]; 183 Bitterne Road West, Bitterne, Southampton, Hampshire SO18 1BJ (GB).

[Continued on next page]

(54) Title: METHOD AND SYSTEM FOR MESSAGE DELIVERY



(57) Abstract: A method and system for message delivery in a publish/subscribe network is provided. The method includes defining a group (202, 303, 304) of publish/subscribe engines (211-214, 311-316) as a subset of a set (201, 301, 302) of fully connectable publish/subscribe engines, wherein the group (202, 303, 304) provides a function which can be carried out by any one of the publish/subscribe engines in the group (202, 303, 304). Each of the publish/subscribe engines in the group (202, 303, 304) sends a group proxy subscription including a topic, an indication of the publish/subscribe engine sending the group proxy subscription, and an indication of the group. A publish/subscribe engine (214, 315) outside the group (202, 303, 304) sends a single publish message to one of the group (202, 303, 304) of publish/subscribe engines with a group proxy subscription for a topic. If the one of the group (202, 303, 304) of publish/subscribe engines to which the single publish message is being sent also has an individual proxy subscription, the publish message for the group (202, 303, 304) is included in the publish message for the individual proxy subscription.

WO 2009/109510 A1

Published:

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

METHOD AND SYSTEM FOR MESSAGE DELIVERY IN A PUBLISH/SUBSCRIBE NETWORK

This invention relates to the field of message delivery. In particular, the invention relates to message delivery in a publish/subscribe network in which publish/subscribe engines are
5 grouped to provide specific functionality.

In a publish/subscribe network a common technique is for each publish/subscribe engine in the network to send out proxy subscriptions for the topics that it actually has subscribers on. This proxy subscription informs the other publish/subscribe engines in the collective (a
10 group of fully connected publish/subscribe engines) that if they receive a publication for the topic then they should forward it onto the publish/subscribe engine that sent out the proxy subscription. This reduces the network traffic as only a single publication needs to be sent to the publish/subscribe engine, which is then fanned out to all the local subscribers.

Proxy subscriptions are sent between publish/subscribe engines, so each publish/subscribe
15 engine knows each other publish/subscribe engine's set of subscriptions and hence knows where to forward each publication according to its topic, avoiding the need to send all publications to all publish/subscribe engines. In conventional publish/subscribe networks, the topology does not assume any distinction between different publish/subscribe engines -
20 all are equal and each publish/subscribe engine may wish to connect to any other - and so the propagation of proxy subscription information is either hierarchical (for a hierarchical publish/subscribe engine topology) or multicast to all publish/subscribe engines.

The technique of using collectives and proxy subscriptions does not enable groups of
25 publish/subscribe engines to be defined inside a collective. A group of publish/subscribe engines may wish to serve a common goal. This goal might be delivering messages to subscribers that have a destination which can be directly accessed by any of the publish/subscribe engines in the group, for example, on a shared queue. An alternative might be if two publish/subscribe engines are desired to link between two collectives,
30 normally this would create duplicates as each one would receive the publications and attempt to forward on to the other collective.

According to a first aspect of the present invention there is provided a method for message delivery in a publish/subscribe network, comprising: defining a group of publish/subscribe engines as a subset of a set of fully connected publish/subscribe engines, wherein the group provides a function which can be carried out by any one of the publish/subscribe engines in the group; each of the publish/subscribe engines in the group sending a group proxy subscription, the group proxy subscription including a topic, an indication of the publish/subscribe engine sending the group proxy subscription, and an indication of the group.

10 According to a second aspect of the present invention there is provided a method for message delivery in a publish/subscribe network, comprising: receiving a group proxy subscription from each one of a group of publish/subscribe engines, wherein the group proxy subscription includes a topic, an indication of the publish/subscribe engine sending the group proxy subscription, and an indication of the group; receiving a publish message for the topic defined in the group proxy subscription; and forwarding the publish message to one of the publish/subscribe engines in the group.

20 According to a third aspect of the present invention there is provided a computer program product stored on a computer readable storage medium for message delivery, comprising computer readable program code means for performing the steps of: defining a group of publish/subscribe engines as a subset of a set of fully connected publish/subscribe engines, wherein the group provides a function which can be carried out by any one of the publish/subscribe engines in the group; each of the publish/subscribe engines in the group sending a group proxy subscription, the group proxy subscription including a topic, an indication of the publish/subscribe engine sending the group proxy subscription, and an indication of the group.

30 According to a fourth aspect of the present invention there is provided a computer program product stored on a computer readable storage medium for message delivery, comprising computer readable program code means for performing the steps of: receiving a group proxy subscription from each one of a group of publish/subscribe engines, wherein the group proxy subscription includes a topic, an indication of the publish/subscribe engine sending the group

proxy subscription, and an indication of the group; receiving a publish message for the topic defined in the group proxy subscription; and forwarding the publish message to one of the publish/subscribe engines in the group.

5 According to a fifth aspect of the present invention there is provided a system for message delivery in a publish/subscribe network, comprising: a plurality of publish/subscribe engines defined as a group, the plurality of publish/subscribe engines being a subset of a set of fully connected publish/subscribe engines, and wherein the group provides a function which can be carried out by any one of the publish/subscribe engines in the group; each of the
10 publish/subscribe engines in the group including: means for sending a group proxy subscription, the group proxy subscription including a topic, an indication of the publish/subscribe engine sending the group proxy subscription, and an indication of the group.

15 This invention enables groups of publish/subscribe engines to be defined inside a collective, which have a specific functionality and where typically it would be desirable for a single one of these publish/subscribe engines to carry out the action. This is achieved at the same time as the publish/subscribe engines performing normal publish/subscribe work like any other publish/subscribe engine in the collective.

20 The solution proposed protects against duplicate messages being created, whilst at the same time ensuring that a member of the group will receive the publication in a timely manner so that the group functionality may be satisfied.

25 The subject matter regarded as the invention is particularly pointed out and distinctly claimed in the concluding portion of the specification. The invention, both as to organization and method of operation, together with features, and advantages thereof, may best be understood by reference to the following detailed description when read with the accompanying drawings in which:

30 Figure 1 is a block diagram of a computer system in which the present invention may be implemented;

Figure 2 is a block diagram of a first embodiment of a publish/subscribe system in accordance with the present invention;

5 Figure 3 is a block diagram of a second embodiment of a publish/subscribe system in accordance with the present invention;

Figure 4 is a flow diagram of a process carried out inside a group in accordance with an aspect of the present invention; and

10 Figure 5 is a flow diagram of a process carried out outside a group in accordance with an aspect of the present invention.

15 It will be appreciated that for simplicity and clarity of illustration, elements shown in the figures have not necessarily been drawn to scale. For example, the dimensions of some of the elements may be exaggerated relative to other elements for clarity. Further, where considered appropriate, reference numbers may be repeated among the figures to indicate corresponding or analogous features.

20 In the following detailed description, numerous specific details are set forth in order to provide a thorough understanding of the invention. However, it will be understood by those skilled in the art that the present invention may be practiced without these specific details. In other instances, well-known methods, procedures, and components have not been described in detail so as not to obscure the present invention.

25 Publish/subscribe applications are intended for situations where a single message is required by, and should be distributed to, multiple users. The advantage over other delivery methods is that the publisher is kept separated from the subscriber. This means that the publisher in a publish/subscribe application does not need to have any knowledge of either the subscriber's existence or the applications that may use the published information. Likewise, the
30 subscriber or subscriber applications do not need to know anything about the publisher application.

A publish/subscribe application has one or more publishers who publish messages from an application to a publish/subscribe engine, and a group of subscribers who subscribe to some or all of those published messages that are held on the publish/subscribe engine. The system matches the publications to the subscribers and ensures that all the messages are made
5 available and delivered to all the subscribers in a timely manner.

Geographically distributed systems or requirements for heavy workloads can be accommodated by a network of publish/subscribe engines having two-way connectivity among them.

10

Publish/subscribe engines can be linked to form a hierarchy, and publications and subscriptions can flow from any publish/subscribe engine to any other publish/subscribe engine in the hierarchy. As well as the hierarchy topology, another topology type is supported, called a collective.

15

A collective is a set of publish/subscribe engines in which the publish/subscribe engines (for example, in the form of brokers or queue managers) are totally connected. That is, every publish/subscribe engine is directly connected to every other publish/subscribe engine. Communication between publish/subscribe engines is optimized within this type of topology
20 and collectives may be connected to other stand-alone publish/subscribe engines or collectives in a large topology.

In a publish/subscribe network a technique is for each publish/subscribe engine in the network to send out proxy subscriptions for the topics that it actually has subscribers on.
25 This proxy subscription informs the other publish/subscribe engines in the collective of the publish/subscribe engine's subscriber's topic. If another publish/subscribe engine receives a publication for one of these topics then it should forward it to the publish/subscribe engine that sent out the proxy subscription. This reduces the network traffic as only a single publication needs to be sent to the publish/subscribe engine, which is then fanned out to all
30 the local subscribers.

Proxy subscriptions are sent between publish/subscribe engines, so each publish/subscribe engine knows each other publish/subscribe engine's set of subscriptions and hence knows where to forward each publication according to its topic, avoiding the need to send all publications to all publish/subscribe engines.

5

A technique of using collectives and proxy subscriptions is provided to enable groups of publish/subscribe engines to be defined inside a collective. A group of publish/subscribe engines may wish to serve a common goal or carry out a function. In one example, this goal might be delivering messages to subscribers that have a destination which can be directly
10 accessed by any of the publish/subscribe engines in the group, for example, on a shared queue. In another example, an alternative goal might be if two or more of the publish/subscribe engines are linked between two collectives. Other functions or goals can equally be provided by a group.

15

Referring to Figure 1, an exemplary system for implementing a publish/subscribe engine includes a data processing system 100 suitable for storing and/or executing program code including at least one processor 101 coupled directly or indirectly to memory elements through a bus system 103. The memory elements can include local memory employed during actual execution of the program code, bulk storage, and cache memories which
20 provide temporary storage of at least some program code in order to reduce the number of times code must be retrieved from bulk storage during execution.

20

The memory elements may include system memory 102 in the form of read only memory (ROM) 104 and random access memory (RAM) 105. A basic input/output system (BIOS)
25 106 may be stored in ROM 104. System software 107 may be stored in RAM 105 including operating system software 108. Software applications 110 may also be stored in RAM 105.

25

The system 100 may also include a primary storage means 111 such as a magnetic hard disk drive and secondary storage means 112 such as a magnetic disc drive and an optical disc
30 drive. The drives and their associated computer-readable media provide non-volatile storage of computer-executable instructions, data structures, program modules and other data for the system 100. Software applications may be stored on the primary and secondary storage

30

means 111, 112 as well as the system memory 102.

The computing system 100 may operate in a networked environment using logical connections to one or more remote computers via a network adapter 116.

5

Input/output devices 113 can be coupled to the system either directly or through intervening I/O controllers. A user may enter commands and information into the system 100 through input devices such as a keyboard, pointing device, or other input devices (for example, microphone, joy stick, game pad, satellite dish, scanner, or the like). Output devices may include speakers, printers, etc. A display device 114 is also connected to system bus 103 via an interface, such as video adapter 115.

10

Two example embodiments are used to illustrate the described system showing different functions of groups. It should be understood that the described method and system are not limited to the example functions of the groups, but may be applied to groups carrying out any group function.

15

A first embodiment of the described system is provided with reference to Figure 2. Figure 2 shows an example publish/subscribe network 200 with a collective 201 of four publish/subscribe engines (B1, B2, B3, B4) 211, 212, 213, 214. For example, the publish/subscribe engines may be queue managers. In this example, a group within the collective 201 is defined in the form of a queue sharing group (QSG1) 202 with a shared queue 203. The queue sharing group 202 includes three of the publish/subscribe engines (B1, B2, B3) 211, 212, 213.

20

25

In this first embodiment of the described system, the case is considered in which the destination of the publications for a particular subscriber is the shared queue 203.

There are two forms of proxy subscription that are used in the described systems.

30

- The first one is the standard proxy subscription that is sent from a publish/subscribe engine that details the topic that the proxy is for and the publish/subscribe engine that

is interested in it. This is referred to as an “individual proxy subscription” in order to distinguish it from the second type of proxy subscription.

- The second type of proxy is similar, but is qualified with a group name as well. This second type is referred to as a “group proxy subscription”.

5

For the sake of this description, these proxies are represented as “T1@B1” for an individual proxy subscription (referring to a specified topic at a specified publish/subscribe engine) and “T1@B1@GRP1” for a group proxy subscription (referring to a specified topic at a specified publish/subscribe engine at a specified group).

10

The description of the processing that takes place can be divided into two parts: one for inside (associated with) a group (in this example the queue sharing group 202); and one for outside (not associated with) a group. These two parts can be understood independently, those outside the group are not concerned with what the group is doing; they just have to send the group the correct number and type of publications. Figure 2 is used to explain how the proxy subscriptions are used by those inside the group and those outside.

15

The publish/subscribe engine B4 214 just needs to know that the publish/subscribe engines B1-B3, 211-213 are in a “group”, what that actual group does is not important to the publish/subscribe engine B4 214. In this example, the publish/subscribe engine B4 214 does not know that the publish/subscribe engines B1-B3, 211-213 of the group are part of a queue sharing group.

20

Taking B2 212 to start with, there is a subscriber called S2 222 that is connected to B2 212 and has subscribed to topic T1 and the destination is the shared queue 203. As S2 222 is using a shared queue 203, B2 212 will send the full subscription (not a proxy) to the other publish/subscribe engines in the queue sharing group 202 (B1 211 and B3 213). This means that B1 211 and B3 213 also have a full subscription for S2 222.

25

30

On B1 211 there is a subscriber S1 221 that has subscribed to topic T1 with a private destination, for example, a local queue 231. B1 211 will ensure that the other

publish/subscribe engines B2 212, B3 213 in the queue sharing group 202 have a proxy subscription, so they know to send B1 211 any publications on T1.

5 Similarly, B3 213 also has a private subscriber S3 223 on topic T1 with a destination of a private destination in the form of local queue 233. B3 213 sends out a proxy subscription for this topic.

10 The subscriptions at each of the publish/subscribe engines B1 - B4 211-214 are shown in callouts 241-244. The callout 241 for publish/subscribe engine B1 211 has the following subscriptions:

“S1 on T1” – the full subscription for its local subscriber S1 221;

“S2 on T1” – the full subscription of S2 222 on the shared queue 203;

“T1@B3” – the proxy subscription from B3 213 for S3 223 with the local destination queue 233 on B3 213.

15

The callout 242 for publish/subscribe engine B2 212 has the following subscriptions:

“S2 on T1” – for the full subscription of S2 222 on the shared queue 203;

“T1@B1” – the proxy subscription from B1 211 for S1 221 with the local destination queue 231 on B3 211; and

20

“T1@B3” – the proxy subscription for from B3 213 for S3 223 with the local destination queue 233 on B3 213.

The callout 243 for publish/subscribe engine B3 213 has the following subscriptions:

“S3 on T1” – the full subscription for its local subscriber S3 223;

25

“S2 on T1” – the full subscription of S2 222 on the shared queue 203;

“T1@B1” – the proxy subscription from B1 211 for S1 221 with the local destination queue 231 on B1 211.

The callout 244 for publish/subscribe engine B4 214 has the following subscriptions:

30

“T1@B1” – the individual proxy subscription from B1 211 for S1 221 with the local destination queue 231 on B1 211;

“T1@B3” – the individual proxy subscription from B3 213 for S3 223 with the local destination queue 233 on B3 213;

“T1@B1@QSG1” – the group proxy subscription from B1 211 for the shared subscribers with the destination of the shared queue 203;

5 “T1@B2@QSG1” – the group proxy subscription from B2 212 for the shared subscribers with the destination of the shared queue 203;

“T1@B3@QSG1” – the group proxy subscription from B3 213 for the shared subscribers with the destination of the shared queue 203.

10 Looking at the subscriptions on B4 214, there is an individual proxy subscription from B1 for T1 and also from B3 for the same topic. In addition, there are a number of group proxy subscriptions of the form TOPIC@QMGR@GROUP which are due to the shared subscriber S2 222. Each publish/subscribe engine B1 - B3 211-213 in the queue sharing group 202 will send out a group proxy subscription for the shared subscribers independently of their private
15 subscribers, and also independently of the other publish/subscribe engines in the queue sharing group 202. This mean that because there are three publish/subscribe engines B1 - B3 211-213 in the queue sharing group 203 there are three group proxy subscriptions at B4 214, one from each of B1-B3. The publish/subscribe engines B1 - B3 211-213 in the queue sharing group 203 do not send these proxy subscriptions to the other publish/subscribe
20 engines in the queue sharing group 203, they already have full subscriptions.

Firstly, the process outside the group is considered. A publisher P4 224 on B4 214 is published for topic T1. There will be two publications that are sent out from B4 214 to satisfy the proxy subscriptions shown in the callout 244.

25 Following the normal rules for proxy subscriptions, where the publish/subscribe engine B4 214 on which the initial publication is occurring will forward the publication to the other publish/subscribe engines that have provided a proxy subscription. In this case, a publish message will be sent to both B1 211 and B3 213 to satisfy the individual proxy subscriptions
30 of “T1@B1” and “T1@B3”.

When these publications are sent to the publish/subscribe engines B1 211, B3 213 to satisfy the individual proxy subscriptions, the destination of the publication will be identified. So for B1 211 the message will identify that it is to satisfy the subscription on B1 211 and similarly for B3.

5

In addition, one of these messages will also be identified as being destined for the queue sharing group QSG1 202. It is assumed that it is the message that is sent to B1 211 that is to perform this action. This means that the message sent to B3 213 will be identified as having a destination of B3 213 only, whereas the message sent to B1 211 will be identified as being

10 destined for both B1 211 and QSG1 202. Only a single message needs to be sent to satisfy the whole of the group QSG1 202, so although there are three group proxy subscriptions from the queue sharing group QSG1 202, only one message is required, which can actually be combined with satisfying the private workload of one of the members of the group.

15 Secondly, the process inside the group is considered. B1 211 will receive a publication from B4 214 that is identified as being for B1 211 and QSG1 202. B1 211 will therefore satisfy its private subscribers with this publication and also the shared subscribers (that is the functionality of the group). B3 213 will receive a publication destined for just B3 213 and will therefore only satisfy its private subscribers.

20

As has already been stated, the publish/subscribe engines outside a group do not need to know what the group is going to do with the publications. The publish/subscribe engines outside the group follow the following rules:

- They send a publication for each of the normal proxy subscriptions and also mark

25 one of these as being for the group;

- They send a publication for each of the normal proxy subscriptions and a separate publication for the group as a separate message, to one of the publish/subscribe engines in the group; or
- If there are no normal proxy subscriptions for the group, they send a unique message

30 to one of the group.

This means that it is possible to create other groups that carry out a different function, without needing to change the publish/subscribe engines outside the group.

5 A second embodiment of the described system is provided with reference to Figure 3 in which a group with functionality is a group that is used as a gateway between two collectives. Normally, it would only be possible to have a single link, due to the fact that duplicate messages would be created if there was more than one link. Using the described method it is possible to create a group that provides higher availability between collectives, whilst not risking duplicates being created. Figure 3 is used to explain how this type of
10 group works.

Referring to Figure 3, a publish/subscribe environment 300 is shown in which there are two collectives, COL1 301 and COL2 302.

15 The first collective COL1 301 has three publish/subscribe engines B1 311, B2 312, B5 315. Two of the publish/subscribe engines B1 311, B2 312 of the first collective COL1 301 are in a gateway group GW1 303.

20 The second collective COL2 302 has three publish/subscribe engines B3 313, B4 314, B6 316. Two of the publish/subscribe engines B3 313, B4 314 of the second collective COL2 302 are in a gateway group GW2 304.

In COL2 302 there is a subscriber S1 326 on B6 306. It is assumed that a publisher P4 325 is on B5 315 in COL1 301. It is described how, using proxy subscriptions, the publication is
25 delivered to S1 326.

Before looking at the publication, the proxy subscriptions that have been created are considered. The proxy subscriptions at each of the publish/subscribe engines B1-B6 311-316 are shown in callouts 342-346.

30 As B3 313 and B4 314 are in the same collective as B6 316, they will have both received a proxy subscription "T1@B6" from B6 316 for topic T1. This is shown in the callouts 343,

344 for B3 313 and B4 314. As both B3 313 and B4 314 belong to a “Gateway Group” GW2 304, they will both send group proxy subscriptions to the publish/subscribe engine gateways in the other collective. In this example, there are two gateway publish/subscribe engines B1 311, B2 312 in the other collective 301 which also form a “Gateway Group” 303.

5 The design would equally work with a single publish/subscribe engine on the other side.

B1 311 and B2 312 both receive group proxy subscriptions from B3 313 and B4 314.

Because B3 313 and B4 314 are in a group they send group proxy subscriptions, hence B1 311 and B2 312 have proxy subscriptions of “T1@B3@GW2” and “T1@B4@GW2” shown
10 in their callouts 341, 342.

B1 311 and B2 312 will then send out group proxy subscriptions to the other publish/subscribe engines in COL1 301, but not to those in the same group as themselves.

Therefore B5 315 will receive group proxy subscriptions from B1 311 and B2 312 resulting
15 in the group proxy subscriptions of “T1@B1@GW1” and “T1@B2@GW1” in the callout 345 for B5 315.

When a publication happens at B5 315 on topic T1, the publish/subscribe engine B5 315
20 look at the proxy subscriptions 345 (“T1@B1@GW1” and “T1@B2@GW1”) that it has and as they are group proxy subscriptions, B5 315 will send a single message to either B1 312 or B2 322, indicating that it is for group GW1 303.

It is assumed it is sent to B2 312. B2 312 will then see that the publication is for GW1 303
25 and because of the group proxy subscriptions that it has (“T1@B3@GW2” and “T1@B4@GW2”), and the type of group, B2 312 will send the message to either B3 313 or B4 314.

It is assumed it is sent to B3 313. B3 313 will see that the message is destined for GW2 304,
30 but because it is coming from the partner gateway GW1 303 it will send the publication out according to the proxy subscriptions that it has (“T1@B6”). In this case, just to B6 316, but not back to GW1 303.

In this way it is possible to have a highly available link between two collectives 301, 302 without risking duplicate messages, or being reliant on any particular gateway publish/subscribe engine. The gateway between the collectives will still function if a single publish/subscribe engine in one (or both) of the gateway groups fails. The design is not
5 limited to just two publish/subscribe engines in each gateway group; more publish/subscribe engines could be added to the groups to provide an even higher level of availability.

Referring to Figure 4 a flow diagram 400 shows the described process from the perspective of a publication outside a group.

10

A publication is ready 401 on topic T1 at a publish/subscribe engine. It is determined 402 if there are any individual proxy subscriptions on the topic. If not, the process loops 403 to later in the process, prior to step 406. If there are individual proxy subscriptions on the topic, a publication is prepared 404 for each individual proxy subscription on the topic. The
15 destination publish/subscribe engine is identified 405 in the publication for each of the individual proxy subscriptions.

It is then determined 406 if there are any group proxy subscriptions. If not, the prepared publications for the individual proxy subscriptions are sent 407.

20

If there are group proxy subscriptions, it is determined if there is an individual proxy subscription for one of the publish/subscribe engines in the group 408. If so, one of the publish/subscribe engines with an individual proxy subscription is chosen 409. If not, a group destination is identified 410 in the publication and the publication is sent 411 to the
25 chosen publish/subscribe engine.

If there is an individual proxy subscription for a publish/subscribe engine in the group, identify 412 the previously prepared publication for this publish/subscribe engine as also being for the group. Send 413 the publication to all individual proxy publish/subscribe
30 engines.

As an alternative, two publications may be sent to the same publish/subscribe engine, one as a publication destined to a distinct publish/subscribe engine and a separate publish for the group as separate message.

5 Referring to Figure 5, a flow diagram 500 shows the described process from the perspective of a received publication inside a group.

A publication is received 501 at a publish/subscribe engine. The destination identified in the publication is checked 502. It is determined 503 if there is an individual destination. If not,
10 the process loops 504 to skip the next step. If there is an individual destination, the publication is sent 505 to the private subscriber on the publish/subscribe engine.

It is then determined 506 if there is a group destination. If not, the process loops 507 to skip the next step. If there is a group destination, then the publish/subscribe engine carries out
15 508 the functionality of the group. The process then ends 509.

Any publish/subscribe engine outside the group that is sending a publication to a member of the group does not need to know what the group is going to do with the publication. This enables new groups to be created that have specific functionality without having to change
20 all of the other publish/subscribe engines in the collective.

The described process groups a subset of the publish/subscribe engines together to provide some specific functionality (a common goal) while still performing conventional publish/subscribe operations. This builds on the recognition that there are occasions where
25 different publish/subscribe engines should have different roles within the network. The described process implements different types of propagation of proxy subscription within a single network (i.e. some "group proxy subscriptions" are used as well as standard proxy subscriptions). Different processing is performed inside the group compared with outside the group, to make best use of shared queues within a group of publish/subscribe engines.

30

The invention can take the form of an entirely hardware embodiment, an entirely software embodiment or an embodiment containing both hardware and software elements. In an

embodiment, the invention is implemented in software, which includes but is not limited to firmware, resident software, microcode, etc.

5 The invention can take the form of a computer program product accessible from a computer-usable or computer-readable medium providing program code for use by or in connection with a computer or any instruction execution system. For the purposes of this description, a computer usable or computer readable medium can be any apparatus that can contain, store, communicate, propagate, or transport the program for use by or in connection with the instruction execution system, apparatus or device.

10

The medium can be an electronic, magnetic, optical, electromagnetic, infrared, or semiconductor system (or apparatus or device) or a propagation medium. Examples of a computer-readable medium include a semiconductor or solid state memory, magnetic tape, a removable computer diskette, a random access memory (RAM), a read only memory (ROM), a rigid magnetic disk and an optical disk. Current examples of optical disks include compact disk read only memory (CD-ROM), compact disk read/write (CD-R/W), and DVD.

15

Improvements and modifications can be made to the foregoing without departing from the scope of the present invention.

CLAIMS

1. A method for message delivery in a publish/subscribe network, comprising:
defining a group (202, 303, 304) of publish/subscribe engines (211-214, 311-316) as
5 a subset of a set (201, 301, 302) of fully connectable publish/subscribe engines, wherein the
group (202, 303, 304) provides a function which can be carried out by any one of the
publish/subscribe engines (211-214, 311-316) associated with the group (202, 303, 304);
each of the publish/subscribe engines (211-214, 311-316) associated with the group
(202, 303, 304) sending a group proxy subscription, the group proxy subscription comprising
10 a topic, an indication of the publish/subscribe engine sending the group proxy subscription,
and an indication of the group (202, 303, 304).
2. A method as claimed in claim 1, comprising a publish/subscribe engine (214, 315)
not associated with the group (202, 303, 304) sending a publish message to one of the group
15 (202, 303, 304) of publish/subscribe engines (211-214, 311-316) with a group proxy
subscription for a topic.
3. A method as claimed in claim 1 or claim 2, wherein, if the one of the group (202,
303, 304) of publish/subscribe engines (211-214, 311-316) to which the publish message is
20 being sent also has an individual proxy subscription, including the publish message for the
group (202, 303, 304) in the publish message for the individual proxy subscription.
4. A method as claimed in any one of the preceding claims, wherein a publish message
identifies a destination of one of: an individual publish/subscribe engine, a group (202, 303,
25 304) of publish/subscribe engines (211-214, 311-316), or both an individual
publish/subscribe engine and a group (202, 303, 304) of publish/subscribe engines (211-214,
311-316).
5. A method as claimed in claim 4, wherein a publish/subscribe engine receiving a
30 publish message with a destination identified as a group (202, 303, 304) of publish/subscribe
engines (211-214, 311-316), carries out the function of the group (202, 303, 304).

6. A method as claimed in any one of the preceding claims, wherein the group (202, 303, 304) comprises a queue sharing group (202) with a shared queue (203) within a collective (201), and the function comprises sending the publish message from the shared queue (203) to a subscriber.

5

7. A method as claimed in any one of claims 1 to 5, wherein the group (202, 303, 304) comprises a gateway group (303, 304) within a first collective (301) for connection to a second collective (302), and the function comprises linking the first and second collectives (301, 302).

10

8. A method as claimed in any one of the preceding claims, wherein a proxy subscription is sent from a publish/subscribe engine to other publish/subscribe engines to inform the other publish/subscribe engines that if they receive a publish message for the topic associated with the proxy subscription, they should forward it to the publish/subscribe engine sending the proxy subscription.

15

9. A method as claimed in any one of the preceding claims, wherein a group proxy subscription is sent from each one of a group (202, 303, 304) of publish/subscribe engines (211-214, 311-316) to publish/subscribe engines (214, 315) not associated with the group (202, 303, 304), to inform the publish/subscribe engines not associated with the group (202, 303, 304) that if they receive a publish message for the topic associated with the proxy subscription, they should forward it to one of the publish/subscribe engines (211-214, 311-316) associated with group (202, 303, 304).

20

10. A method for message delivery in a publish/subscribe network, comprising:
receiving a group proxy subscription from each one of a group (202, 303, 304) of publish/subscribe engines (211-214, 311-316), wherein the group proxy subscription comprises a topic, an indication of the publish/subscribe engine sending the group proxy subscription, and an indication of the group (202, 303, 304);

25

receiving a publish message for the topic associated with the group proxy subscription; and

30

forwarding the publish message to one of the publish/subscribe engines (211-214, 311-316) associated with group (202, 303, 304).

- 5 11. A computer program comprising program code means adapted to perform all the steps of any of claims 1 to 10 when said program is run on a computer.
12. An apparatus for message delivery in a publish/subscribe network, comprising:
means for defining a group (202, 303, 304) of publish/subscribe engines (211-214, 311-316) as a subset of a set (201, 301, 302) of fully connectable publish/subscribe engines,
10 wherein the group (202, 303, 304) is operable to provide a function which can be carried out by any one of the publish/subscribe engines (211-214, 311-316) associated with the group (202, 303, 304); and
means for controlling each of the publish/subscribe engines (211-214, 311-316) associated with the group (202, 303, 304) to send a group proxy subscription, the group
15 proxy subscription comprising a topic, an indication of the publish/subscribe engine sending the group proxy subscription, and an indication of the group (202, 303, 304).
13. A system for message delivery in a publish/subscribe network, comprising:
a plurality of publish/subscribe engines (211-214, 311-316) definable as a group
20 (202, 303, 304), the plurality of publish/subscribe engines (211-214, 311-316) being a subset of a set (201, 301, 302) of fully connectable publish/subscribe engines, and wherein the group (202, 303, 304) is operable to provide a function which can be carried out by any one of the publish/subscribe engines (211-214, 311-316) associated with group (202, 303, 304);
each of the publish/subscribe engines (211-214, 311-316) associated with group
25 (202, 303, 304) comprising:
means for sending a group proxy subscription, the group proxy subscription comprising a topic, an indication of the publish/subscribe engine sending the group proxy subscription, and an indication of the group (202, 303, 304).
- 30 14. A system as claimed in claim 13, comprising:
a publish/subscribe engine (214, 315) not associated with the group (202, 303, 304) comprising:

means for receiving group proxy subscription for a topic from each publish/subscribe engine in a group (202, 303, 304); and

means for sending a publish message to one of the group (202, 303, 304) of publish/subscribe engines (211-214, 311-316) with a group proxy subscription for a topic.

5

15. A system as claimed in claim 13, wherein the publish/subscribe engine (214, 315) not associated with the group (202, 303, 304) also comprises:

means for determining if the one of the group (202, 303, 304) of publish/subscribe engines (211-214, 311-316) to which the publish message is being sent also has an

10 individual proxy subscription; and

means for including the publish message for the group (202, 303, 304) in the publish message for the individual proxy subscription.

16. A system as claimed in any one of claims 13 to 15, wherein a publish message is operable to identify a destination of one of: an individual publish/subscribe engine, a group (202, 303, 304) of publish/subscribe engines (211-214, 311-316), or both an individual publish/subscribe engine and a group (202, 303, 304) of publish/subscribe engines (211-214, 311-316).

20 17. A system as claimed in claim 16, wherein a publish/subscribe engine in the group (202, 303, 304) comprises:

means for determining a destination identified in a publish message as a group (202, 303, 304) of publish/subscribe engines (211-214, 311-316); and

means for carrying out the function of the group (202, 303, 304).

25

18. A system as claimed in any one of claims 13 to 17, wherein the group (202, 303, 304) comprises a queue sharing group (202) with a shared queue within a collective, and the function comprises sending the publish message from the shared queue to a subscriber.

30 19. A system as claimed in any one of claims 13 to 17, wherein the group (202, 303, 304) comprises a gateway group (303, 304) within a first collective for connection to a second collective, and the function comprises linking the first and second collectives.

FIG. 1




FIG. 2




FIG. 3




FIG. 4

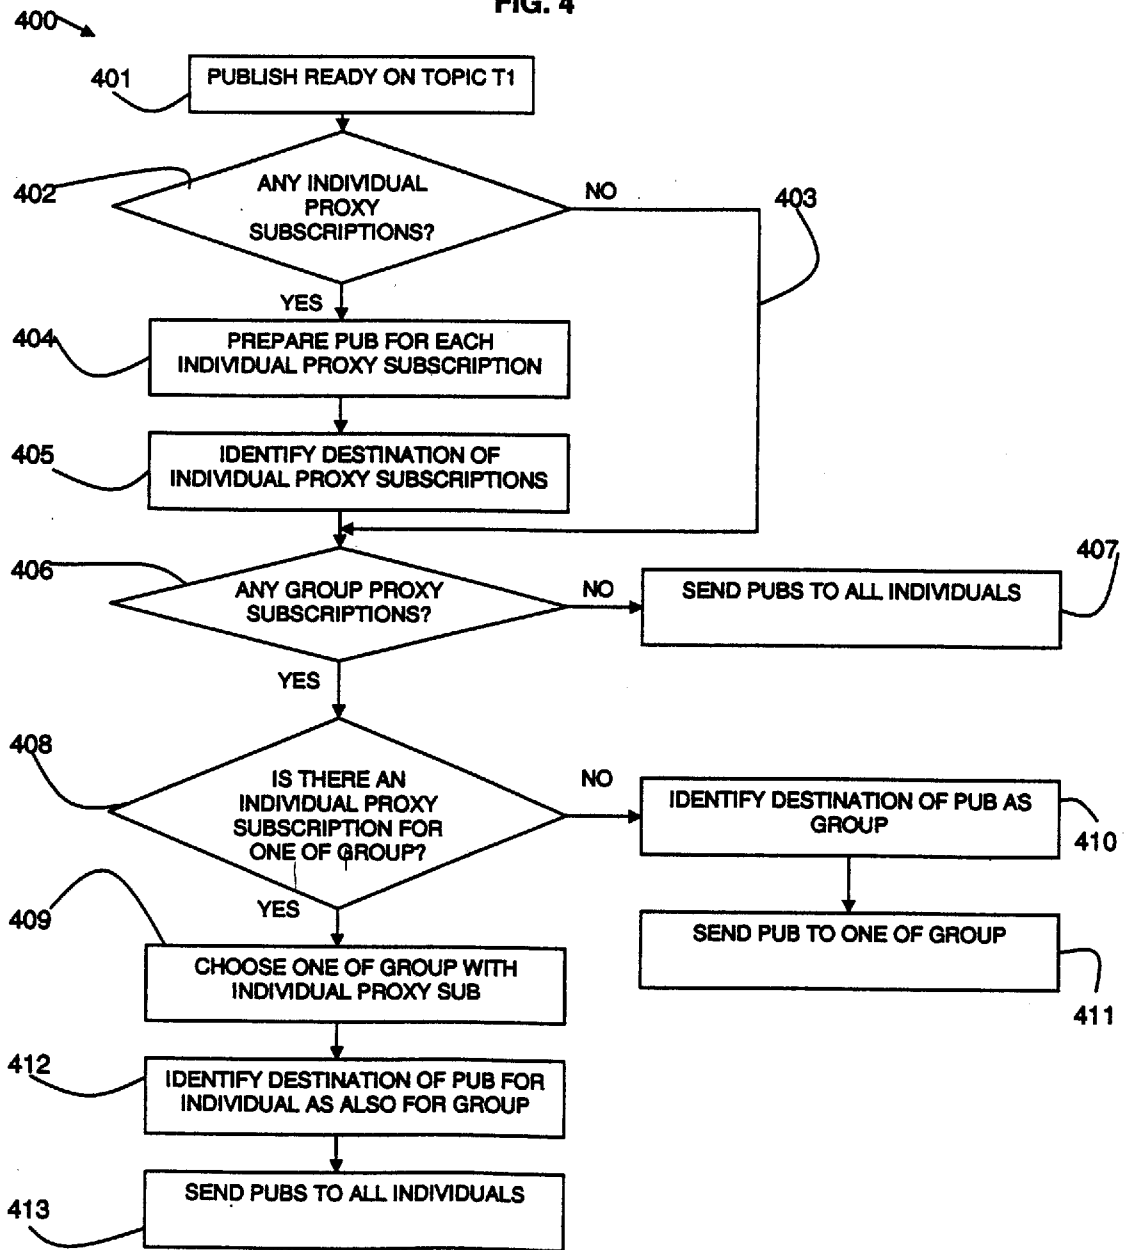
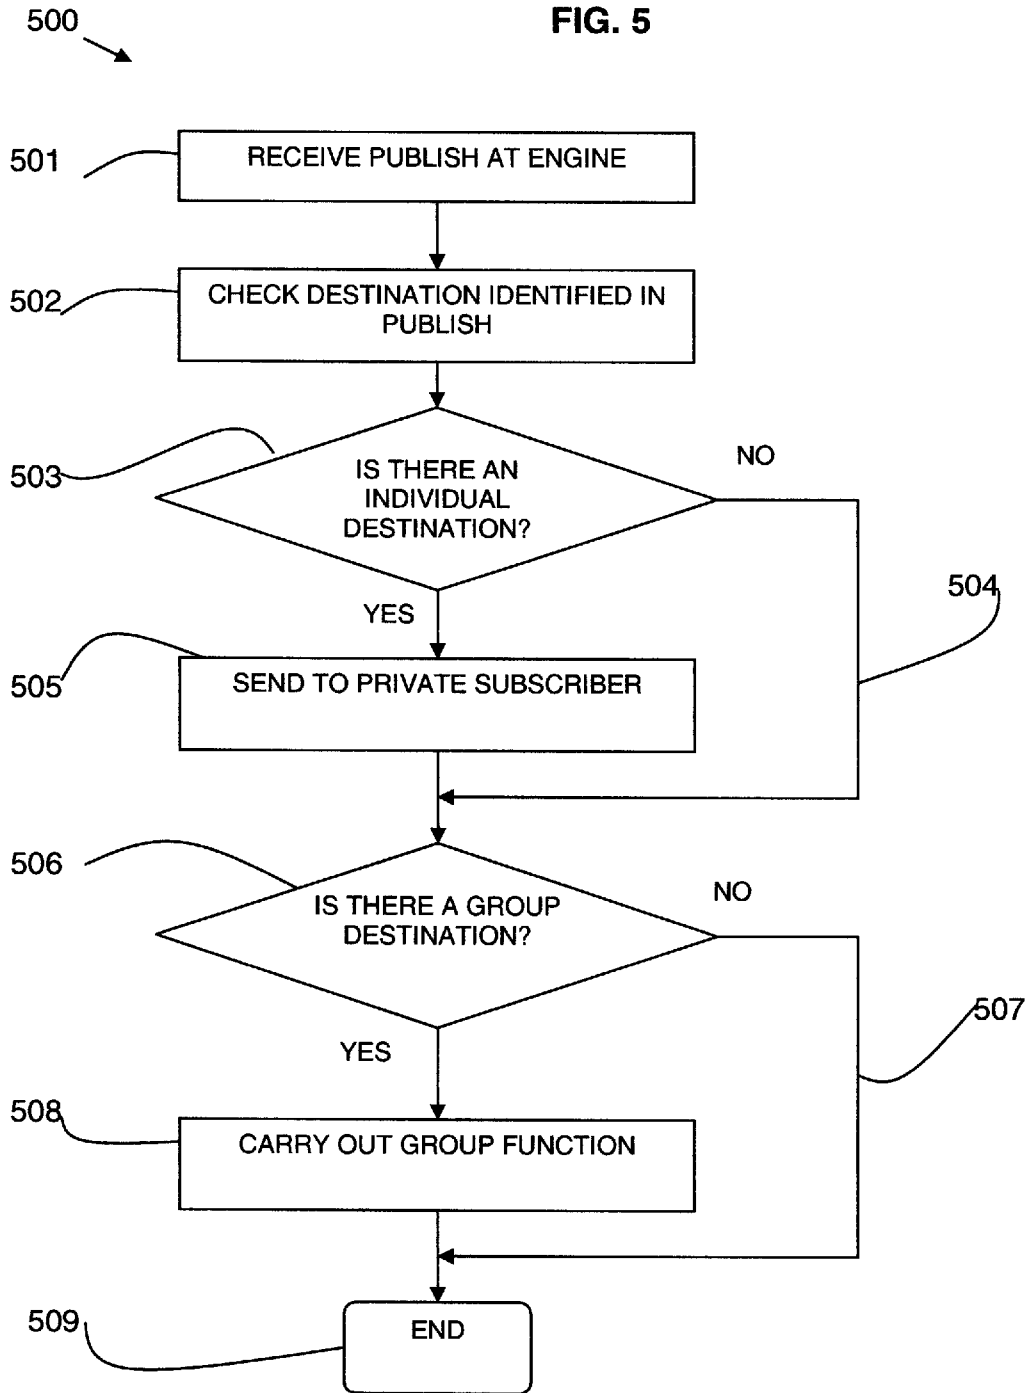


FIG. 5



INTERNATIONAL SEARCH REPORT

International application No
PCT/EP2009/052272

A. CLASSIFICATION OF SUBJECT MATTER
INV. H04L12/58 H04L29/08 H04L12/18

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)
H04L

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 practical, search terms used)

EPO-Internal

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X Y	<p>US 6 336 119 B1 (BANAVAR GURUDUTH S [US] ET AL) 1 January 2002 (2002-01-01) abstract</p> <p>column 1, lines 30-34 column 1, line 64 - column 2, line 7 column 2, line 30 - column 3, line 59 column 4, line 45 - column 6, line 23 column 7, lines 6-56 figures 1,2,4A,4B,5,6A,6B</p> <p style="text-align: center;">----- -/--</p>	<p>1-3,8-15</p> <p>4-7, 16-19</p>

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 document 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

12 May 2009

Date of mailing of the international search report

28/05/2009

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

Di Felice, M

INTERNATIONAL SEARCH REPORT

International application No

PCT/EP2009/052272

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	<p>WO 2004/072800 A (PROGRESS SOFTWARE CORP [US]; CULLEN WILLIAM [US]) 26 August 2004 (2004-08-26) abstract paragraph [0002] paragraph [0007] paragraphs [0033] - [0040] paragraph [0070] paragraph [0080]</p>	4,5,16, 17
Y	<p>US 2005/015619 A1 (LEE WING [US]) 20 January 2005 (2005-01-20) abstract paragraph [0004] paragraph [0006] paragraph [0016]</p>	6,18
Y	<p>G.FOX ET AL.: "A Scalable Durable Grid Event Service" TECHNICAL REPORT, COMMUNITY GRIDS LAB, [Online] 2001, pages 1-11, XP002527189 Indiana University, USA Retrieved from the Internet: URL:http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.15.2070&rep=rep1&type=pdf [retrieved on 2009-05-11] abstract page 1, lines 23-35 page 1, line 51 - page 3, line 15 figure 1</p>	7,19
L	<p>AHMET FATIH MUSTACOGLU ET AL: "A novel event-based consistency model for supporting collaborative cyberinfrastructure based scientific research" COLLABORATIVE TECHNOLOGIES AND SYSTEMS, 2007. CTS 2007. INTERNATIONAL SYMPOSIUM ON, IEEE, PISCATAWAY, NJ, USA, 25 May 2007 (2007-05-25), pages 271-277, XP031319342 ISBN: 978-0-9785699-1-4 page 277, citation [24]</p>	
A	<p>R.SZAROWSKI: "Hybrid Publish-Subscribe: A Compromise Approach for Large-Scale" CAISE '03 - THE 15TH CONFERENCE ON ADVANCED INFORMATION SYSTEMS ENGINEERING, SHORT PAPER PROCEEDINGS INFORMATION SYSTEMS FOR A CONNECTED SOCIETY, 16 June 2003 (2003-06-16), - 20 June 2003 (2003-06-20) pages 241-244, XP002527190 Klagenfurt/Velden, Austria, ISBN 86-435-0549-8 the whole document</p>	1-19

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/EP2009/052272

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 6336119	B1	01-01-2002	NONE
WO 2004072800	A	26-08-2004	GB 2417160 A 15-02-2006
US 2005015619	A1	20-01-2005	NONE