

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
30 March 2006 (30.03.2006)

PCT

(10) International Publication Number
WO 2006/033611 A1

(51) International Patent Classification⁷: **H04L 7/00**,
G06F 1/12, H04J 3/06

(74) Agents: KYLIN, Peter et al.; Hynell Patenttjänst AB, Patron Carls väg 2, S-683 40 Hagfors/Uddeholm (SE).

(21) International Application Number:
PCT/SE2005/001307

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SM, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, YU, ZA, ZM, ZW.

(22) International Filing Date:
9 September 2005 (09.09.2005)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
0402280-2 22 September 2004 (22.09.2004) SE

(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, HU, IE, IS, IT, LT, LU, LV, MC, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

(71) Applicant (for all designated States except US):
PROSILIANT TECHNOLOGIES AKTIEBOLAG
[SE/SE]; c/o Söderberg, Vråkvägen 1, S-181 56 Lidingö (SE).

(72) Inventors; and


(75) Inventors/Applicants (for US only): **HANSSON, Kjell**
[SE/SE]; Håkonvikings väg 6, S-762 71 Rimbo (SE).
HAGSAND, Olof [SE/SE]; Älta Mossv. 31A, S-138 37 Älta (SE).

Published:

— with international search report

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(54) Title: METHOD, A COMPUTER PROGRAM PRODUCT, AND A CARRIER FOR INDICATING ONE-WAY LATENCY IN A DATA NETWORK



(57) Abstract: Disclosed herein is a method, a computer program product, and a carrier for indicating one-way latency in a data network (N) between a first node (A) and a second node (B), wherein the data network (N) lacks continuous clock synchronization, comprising: a pre-synchronisation step, a measuring step, a post-synchronisation step, an interpolation step, and generating a latency profile. The present invention also relates to a computer program product incorporating the method, a carrier comprising the computer program product, and a method for indicating server functionality based on the first aspect.

WO 2006/033611 A1

Title

Method, a computer program product, and a carrier for indicating one-way latency in a data network

5

Technical field

According to a first aspect, the present invention relates to a method for indicating one-way latency in a data network without continuous clock synchronization between a first node and a second node.

10

According to a second aspect, the present invention relates to a computer program product incorporating the first aspect.

According to a third aspect, the present invention relates to a carrier comprising the computer program product.

According to a fourth aspect, the present invention relates to a method for indicating server functionality based on the first aspect.

15

Background of invention

In the art, it is possible to achieve one-way real-time latency measurement by using synchronized clocks, such as GPS technology. It is possible to use the standard NTP protocol to achieve a level of synchronization between nodes, as described in Mills, D "Network Time Protocol (Version 3) specification, implementation and analysis", IETF RFC 1305, University of Delaware, March 1992. However, the NTP mechanism does not have a very high accuracy.

25

In US2003/0048811 A1, with the title "Methods, systems and computer program products for synchronizing clocks of nodes on a computer network", discloses an algorithm for clock synchronization between two nodes using virtual clocks, a generalization of the clock synchronization for many nodes, and using many round-trip-delays to compute an average one-trip delay. A key feature of the invention described in the US document, is that, each node manages a virtual clock for every other node it synchronizes with.

30

Summary of invention

According to the present invention and the first aspect, a
5 method for indicating one-way latency in a data network, without
continuous clock synchronization, between a first node and a second node
is disclosed. The method comprises:

- a pre-synchronisation step,
- 10 -a measuring step,
- a post-synchronisation step,
- an interpolation step, and
- generating a latency profile.

15 Based on the present invention, also a computer program
product loadable into the internal memory of a computer, comprising
software code portions for performing the method, a carrier comprising the
computer program product, and a method for indicating server
functionality are disclosed. This implies that the present invention is
20 applicable when measuring or monitoring, qualities of a server.

The present invention according to the present invention
presents a number of advantages in relation to prior art. For instance, the
accuracy of the measurements of the present invention is higher than the
accuracy of the NTP mechanism. The present invention deals with a
25 method for making one-way real-time latency measurements with a high
precision between nodes connected by a message-passed network where
clock synchronization of that precision is not available, which is in
contrast to the prior art briefly discussed above. The method may also
present individual per-packet latency values. The present invention
30 performs high precision latency measurements of packets travelling
between two nodes over a limited period of time. Since latencies may be
asymmetric, round-trip estimation may not be used, but it must be relied
on absolute and synchronous time. Also, continuous clock synchronization

does not have to be established, nor maintained. Instead, the present invention makes two synchronizations (before and after) with the single goal to interpolate the measurement results.

In a preferred embodiment, the pre-synchronisation step
5 comprises sending a predetermined message from the first node to the second node. Then, at the first node, a predetermined message sent by the second node is received. The next step is to calculate an offset according to $((T_2 - T_1) + (T_3 - T_4))/2$, where T_1 is the sending time from the first node, T_2 is the receiving time at the second node, T_3 is the sending time at the
10 second node, and T_4 is the receiving time at the first node. Thereafter, a clock difference is set to the offset. Also, an absolute clock is set to T_4 .

In a preferred embodiment, the pre-synchronisation step comprises sending the predetermined message N times, and the predetermined message is received N times. However, there may be cases
15 where all N messages are not received. In such cases, there will be gaps in the correspondence of the data sent and received. This may be handled by not using the measured values associated to the missing messages. Then N , or a number less than N , round-trip-time items are generated. This is done according to $T_{i,4} - T_{i,1} - (T_{i,3} - T_{i,2})$, where i is in the interval $[1..N]$. N
20 offset items, or a number of offset items less than N , as described above is generated. The minimum round-trip-time item is retrieved and the clock difference is set to the offset related to the minimum round-trip-time item, and the absolute clock is set to $T_{i,4}$ for the i having the minimum round-trip-time item.

25 In a preferred embodiment, the method further comprises the step of measuring the overhead for making measurements at the sender.

In a preferred embodiment, the method further comprises the step of measuring the overhead for making measurements at the receiver.

In a preferred embodiment, the measuring step comprises
30 sending a predetermined message from the first node to the second node and storing the sending time for the sending. The predetermined message is received at the second node, the receiving time for the receiving is stored.

In a preferred embodiment, the measuring step comprises sending a predetermined message from the first node to the second node N times chronologically equidistantly, and the sending time for each sending is stored. The predetermined message is received at the second node, and
5 the receiving time for each receiving is stored. Hopefully, the message is received N times at the second node. Otherwise this may be handled as described above.

In a preferred embodiment, wherein the post-synchronisation step is constituted by the pre-synchronisation step as defined above.

10 In a preferred embodiment, the interpolation step comprises calculating the one-way latency in the data network between a first node and a second node according to the following relations:

15
$$\text{rate bias} = (\text{offset of the post-synchronisation step} - \text{offset of the pre-synchronisation step}) / (\text{absolute clock for post-synchronisation step} - \text{absolute clock for the pre-synchronisation step}), \text{ and}$$

20
$$\text{latency} = \text{the time at which the second node received the predetermined message} - (\text{the time at which the first node sent the predetermined message} + \text{the offset of the pre-synchronisation step} + (\text{the time at which the first node sent the predetermined message} - \text{the absolute clock of the pre-synchronisation step})).$$

25 In a preferred embodiment, the interpolation step comprises calculating the one-way latency in the data network between a first node and a second node according to the above for the messages sent between the first and second nodes.

30 In a preferred embodiment, the method further comprises an overhead term corresponding to the overhead for making measurements at at least one of the first and second nodes, and the one-way latency in the

data network between a first node and a second node is calculated according to:

latency = the time at which the second node received the predetermined message - (the time at which the first node sent the predetermined message + the offset of the pre-synchronisation step + ratebias(the time at which the first node sent the predetermined message - the absolute clock of the pre-synchronisation step)) - the overhead term.

In a preferred embodiment, in the interpolation step, the one-way latency in the data network between a first node and a second node is calculated according to the above for the N messages sent between the first and second nodes. Alternatively in case all N messages were not received, this is done for the received messages.

It also lies within the scope of the present invention that it is possible to operate in relation to more nodes than a single one. Of course, the present invention may be used to operate against a plurality of nodes.

Brief description of drawings

In Figure 1, an embodiment of two nodes, A and B, interconnected by a network N are schematically depicted.

In Figure 2, an embodiment of an architecture of a node is schematically depicted.

In Figure 3, an embodiment of a network module is schematically depicted.

In Figure 4, an embodiment of a requestor node pre-synchronization flowchart is schematically depicted.

In Figure 5, an embodiment of a responder node pre-synchronization flowchart is schematically depicted.

In Figure 6, an embodiment of a flowchart of requesting node in the measurement phase is schematically depicted.

In Figure 7, an embodiment of a flowchart of the responding node in the measurement phase is schematically depicted.

In Figure 8, an embodiment of a flowchart of the interpolation method is schematically depicted.

5

Description of preferred embodiments

In a first embodiment, a system with two nodes A and B interconnected by a communication network N is depicted in Figure 1. The nodes

10 communicate by sending messages (packets) over the network N. A measurement is performed from a node A to a node B, where A is called a requesting node, and B is called a responding node. Each node may work both as a requesting node and a responding node. A node can also perform measurements with more than one other node. For example, A can
15 perform a measurement with a third node C (not disclosed in Figure 1) at the same time. The network N may be an inter-network running the IP protocol. This enables any nodes with an IP-interface and an IP protocol stack to communicate with each other over N.

In figure 2, an embodiment of a node is shown. The computer
20 node is equipped with a network interface card that can communicate using IP. Such a node has a CPU, memory buses, disks, etc, that enables it to operate as a computer. The node runs an operating system, in which the system software can be implemented. This embodiment is implemented as a software module running in an operating system of such a node.

25 In Figure 3, an embodiment of a network module is shown. The software module implementing the method described in this document needs to have access to a network module. The network module shown in Figure 3 typically consists of a network interface card, a device driver, an IP stack and a socket API. The network interface card enables the node to
30 physically connect to an access network. The device driver contains software enabling the IP stack to access the network services on the network interface card. The IP stack contains full implementation of the communication protocols that enables the node to communicate over the

internet. This may be the set of protocols referred to as TCP/IP. The socket API is a functional interface that the system module can access in order to send and receive packets to and from the network.

5 In an embodiment, a system module implementing the invention may be implemented as a user application in an operating system. It requires a socket API to access the network in order to send and receive packets over the network.

The nodes communicate with messages over the network. There are two kinds of messages:

- 10
- Synchronization messages
 - Measurement messages

Both types of messages may be encapsulated over the IP protocol using the UDP/IP transport protocol or some other non-reliable datagram service. In an embodiment, both types of messages are encoded
15 with the RTP protocol.

A synchronization message is either a request (syncreq) or response (syncresp). The request message is sent by the requesting node and received by a responding node. A response is sent by a responding node when it receives a syncreq message. The syncresp message is
20 received by the requesting node.

The syncreq message contains the following fields: a sequence number and a time-stamp T1.

The syncresp message contains the following fields: a sequence number and three timestamps: T1, T2, and T3.

25 The semantics of the message fields are as follows:

- Sequence number – The requesting node sets the sequence number incrementally (0, 1, 2, etc). The responder copies the sequence number from a syncreq to a syncresp message. The sequence number is used to detect packet loss, reordering or duplications on
30 the network.
- Timestamp T1. The time when the syncreq message was sent by the requesting node.

- Timestamp T2. The time when the syncreq message was received by the responding node.
- Timestamp T3. The time the syncresp message was sent by the responding node.

5

The measurement messages are sent from the requesting node to the responding node only. The measurement message contains a sequence field and a timestamp field T1.

The semantic of the message fields are as follows:

10

- The sequence number. The requesting node sets the sequence number incrementally (0, 1, 2, etc).
- Timestamp T1. The time when the measurement message was sent by the requesting node.

Now referring to the inventive method, both nodes have high accuracy
15 clocks that are not synchronized with each other. High accuracy means that they are linear with respect to each other over a limited time period on the order of minutes, and that they have high resolution, at least to the level of 1 microsecond. That is, the clocks have different rates, but the rate difference is constant over time.

20 The method is divided into five steps:

- P1 - Synchronization1
- P2 - Measurement
- P3 - Synchronization2
- 25 • P4 - Interpolation and
- Generating a latency profile.

In Table 1 below an embodiment of constants used to parameterise the method are given. The values given to the constants are merely an
30 example; the method can be used also for other values.

<i>Constant name</i>	<i>Description</i>	<i>Example values</i>
SNR	Number of syncreq messages sent.	50
NM	Number of measurement messages sent.	10000
DT	Delay between sending of measurement messages.	20ms

Table 1.

In Table 2 below, variables used in this method are explained.

Variable name	Description
NSREQ	Number of syncreq messages sent.
NSRESP	Number of syncresp messages received.
T1	Time when message was sent by requesting node.
T2	Time when message was received by responding node.
T3	Time when message was sent by responding node.
T4	Time when message was received by requesting node.
RTT	Round-trip-time
RTTMIN	The smallest RTT value during a synchronization phase.
CABS0	Wall clock of a synchronization message in the P1 phase
CDIFF0	The difference/offset between the two clocks at a synchronization message in the P1 phase
CABS1	Wall clock of a synchronization message in the P3 phase
CDIFF2	The difference/offset between the two clocks at a synchronization message in the P3 phase
SEQ	Sequence number set by requesting node.
A[]	Vector containing T1 for all measurement messages.

B[]	Vector containing T2 for all measurement messages.
L[]	Vector containing the resulting one-way latencies, or the latency profile.
Ks	Overhead of sending a message
Kr	Overhead of receiving a message
RATEBIAS	Difference in rate between the two clocks
VALID[]	Vector of boolean values determining the validity of the entries in A[], B[] and L[]

Table 2.

The output of the present invention is a latency profile, which is the vector containing the resulting one-way latencies, or L[].

In Figure 4, an embodiment of a requestor node pre-synchronization flowchart is schematically depicted. The node sends a syncreq to the responding node. It sets the sequence number and the T1 timestamp in the syncreq message. Then it waits for a reply to come back from the responding node, or for a timeout to occur. If a syncreq message was received, a timestamp T4 is registered when the syncresp message was received. Together with the three timestamps T1, T2 and T3, the module tries to find the message with the smallest round-trip-time. This message is used to find the two values CABS0 and CDIFF0 and is used in the interpolation method P4. The method uses two variables NSREQ and NSRESP to record the number of sent syncreq messages and received syncresp messages, respectively. These variables are used as a terminating condition. If the module sends 2SNR syncreq messages without having received SNR syncresp messages, this is an error. As soon as the module has received SNR syncresp messages, it continues to the next phase, P2A. SNR is a predefined constant, typically 50 messages. The method may also use the variables RTT and RTT_MIN. RTT_MIN is preset to a large value, and is used to find the syncreq/syncresp pair with the smallest round-trip-time. This measurement is then used to compute the CABS and CDIFF values. In other words, we claim that the best measurement is the one with the smallest RTT. Many other methods use the mean value. Note that

the method described in Figure 4 may be implemented somewhat differently. For example, the sending and receiving of messages can be made concurrently, not sequentially as is shown in the figure. In that case, two processes are created, one that is sending syncreq messages regularly, and one that is waiting for syncresp messages. In that case, a timeout need not be made. Instead, a delay between the sending of syncreq messages need to be introduced.

In Figure 5, an embodiment of a responder node pre-synchronization flowchart is schematically depicted. The node waits for a syncreq from the requesting node. When such a message is received, it creates a syncresp message, copies the sequence number and T1 from the syncreq message, records T2 and T3, and sends the syncresp message back to the requesting node. If the received message is not a syncreq message, it is assumed that it is a measurement message which is handled in P2B. The size of the vectors is equal to the number of measurement messages sent.

The measurement phase consists of the requesting node periodically sending measurement messages to the responding node. The responding node records the timestamps of the time of sending and the time of receiving the messages in two vectors A[] and B[], respectively. The size of the vectors is equal to the number of measurement messages sent, NM. The two vectors are later used in P4.

In Figure 6, an embodiment of a flowchart of requesting node in the measurement phase is schematically depicted. The requesting node sends NM messages (for example 10000) with interval DT between each packet (for example 20ms). Each syncreq message will contain SEQ, the sequence number; and T1, the time the message was sent. The overhead of sending a message Ks is computed initially. This is the difference in time from when the timestamp was taken and when the message was actually sent. Ks may be set to 0 if the node lacks the capability to compute this time.

In Figure 7, an embodiment of a flowchart of the responding node is shown. The responding node stores the sending timestamp T1 in a

vector A, and the receiving timestamp T2 in the vector B. The sequence number is used as an index in the vector. The overhead of sending a message Kr is computed initially. This is the difference in time from when the timestamp was taken and when the message was actually sent. Kr may
 5 be set to 0 if the node lacks the capability to compute this time.

The second synchronisation phase is in this embodiment similar to phase P1 described above. The differences are as follows:

1. The two processes are called P3A and P3B instead of P1A and P1B, respectively.
- 10 2. The resulting variables are named CABS1 and CDIFF1 instead of CABS0 and CDIFF0, respectively.
3. After successful completion of the processes, both flowchart goes to P4 instead of to P2A and P2B.

In the interpolation phase, the measurements collected in phase P2 in
 15 the vectors A[] and B[] and the synchronization values CABS0, CDIFF0, CABS1 and CDIFF1 in phases P1 and P3 are used to interpolate a sequence of one-way latency values. The method itself can be performed on the requesting node, the responding node, or some other node, and can be performed at any time after the other three phases. For example, this
 20 phase can be made as a post processing stage in a server. However, the data must be transferred to the place where the method is implemented. The end result of the method is a vector L[], i.e. the latency profile, with size NM containing the true one-way latency values of the measurement between the requesting and responding node.

25 In Figure 8, an embodiment of a flowchart of the interpolation method is schematically depicted. First the difference in rate RATEBIAS is computed as follows:

$$\text{RATEBIAS} = (\text{CDIFF1} - \text{CDIFF0}) / (\text{CABS1} - \text{CABS0}),$$

30

The method iteratively computes the values of the one-way latency vector L[] from values collected or computed, as follows:

$$L[i] = B[i] - (A[i] + \text{CDIFF0} + \text{RATEBIAS} * (A[i] - \text{CABS0})) - K_s - K_r$$

Claims

1. Method for indicating one-way latency in a data network (N), without continuous clock synchronization, between a first node (A) and a second node (B), comprising:

- 5 -a pre-synchronisation step,
 -a measuring step,
 -a post-synchronisation step,
 -an interpolation step, and
 -generating a latency profile.

10

2. Method according to claim 1, wherein the pre-synchronisation step comprises:

- sending a predetermined message from the first node (A) to the second node (B),
5 -receiving at the first node (A) a predetermined message sent by the second node (B),
 -calculating an offset according to
 $((T_2 - T_1) + (T_3 - T_4))/2$,

where T_1 is the sending time from the first node (A), T_2 is the receiving time

- 0 at the second node (B), T_3 is the sending time at the second node (B), and
 T_4 is the receiving time at the first node (A), and
 -setting a clock difference to the offset, and
 -setting an absolute clock to T_4 .

3. Method according to claim 2, wherein the pre-synchronisation step comprises

- sending the predetermined message N times,
 -receiving the predetermined message,
 -generating N round-trip-time items according to
 $T_{i,4} - T_{i,1} - (T_{i,3} - T_{i,2})$, where i is in the interval [1..N],
 -generating N offset items,
 -retrieving the minimum round-trip-time item,

- setting the clock difference to the offset related to the minimum round-trip-time item, and
- setting the absolute clock to $T_{i,4}$ for the i having the minimum round-trip-time item.

5

4. Method according to claim 1, further comprising:

- measuring the overhead for making measurements at the sender.

5. Method according to claim 1, further comprising:

- 10 -measuring the overhead for making measurements at the receiver.

6. Method according to claim 1, wherein the measuring step comprises:

- 15 -sending a predetermined message from the first node (A) to the second node (B),
- storing the sending time for the sending,
 - receiving the predetermined message at the second node (B), and
 - storing the receiving time for the receiving.

20 7. Method according to claim 6, wherein the measuring step comprises -sending a predetermined message from the first node (A) to the second node (B) N times chronologically equidistantly,

- storing the sending time for each sending,
- receiving the predetermined message at the second node (B), and
- 25 -storing the receiving time for each receiving.

8. Method according to claim 1, wherein the post-synchronisation step is constituted by the pre-synchronisation step according to any one of claims 2 to 5.

30

9. Method according to claim 1, wherein the interpolation step comprises -calculating the one-way latency in the data network between a first node (A) and a second node (B) according to:

rate bias = (offset of the post-synchronisation step – offset of the pre-synchronisation step)/(absolute clock for post-synchronisation step – absolute clock for the pre-synchronisation), and

5

latency = the time at which the second node (B) received the predetermined message – (the time at which the first node (A) sent the predetermined message + the offset of the pre-synchronisation step + (the time at which the first node (A) sent the predetermined message - the absolute clock of the pre-synchronisation step)).

10

10. Method according to claim 1, wherein the interpolation step comprises -calculating the one-way latency in the data network (N) between a first node (A) and a second node (B) according to claim 9 for the N messages sent between the first and second nodes (A, B).

15

11. Method according to claim 1, further comprising:

-an overhead term corresponding to the overhead for making measurements at at least one of the first and second nodes (A, B),

20

-calculating the one-way latency in the data network between a first node (A) and a second node (B) according to:

latency = the time at which the second node (B) received the predetermined message – (the time at which the first node (A) sent the predetermined message + the offset of the pre-synchronisation step + ratebias(the time at which the first node (A) sent the predetermined message - the absolute clock of the pre-synchronisation step)) - the overhead term.

25

12. Method according to claim 11, wherein the interpolation step comprise -calculating the one-way latency in the data network between a first node (A) and a second node (B) according to claim 9 for the N messages sent between the first and second nodes (A, B).

30

13. A computer program product loadable into the internal memory of a computer, comprising software code portions for performing the method(s) of any of the previous claims, when run on a computer.

- 5 14. A carrier comprising the computer program product according to claim 13.

15. Method for indicating server functionality according to any one of claims 1 to 12.

1/4




Fig. 1




Fig. 2

2/4




Fig. 3




Fig. 4

3/4




Fig. 5




Fig. 6

4/4




Fig. 7




Fig. 8

INTERNATIONAL SEARCH REPORT

International application No.

PCT/SE 2005/001307

A. CLASSIFICATION OF SUBJECT MATTER

IPC: see extra sheet

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)

IPC: H04L, H04J, G06F

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

SE,DK,FI,NO classes as above

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

EPO-INTERNAL, WPI DATA, PAJ

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 20030048811 A1 (EDWARD ADAMS ROBIE, JR, ET AL), 13 March 2003 (13.03.2003), [0006]-[0014] --	1-15
A	EP 1215559 A2 (CHRYSLIS-ITS INC), 19 June 2002 (19.06.2002), [0011]-[0013] --	1-15
A	US 20030048754 A1 (LEON BRUCKMAN), 13 March 2003 (13.03.2003), [0012]-[0047] --	1-15
A	US 20010000071 A1 (RICHARD ALLEN NICHOLS), 29 March 2001 (29.03.2001), [0010]-[0013] --	1-15

☒ 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

1 December 2005

Date of mailing of the international search report

06-12-2005

Name and mailing address of the ISA/

Swedish Patent Office

Box 5055, S-102 42 STOCKHOLM

Facsimile No. +46 8 666 02 86

Authorized officer

Anders Edlund/MN

Telephone No. +46 8 782 25 00

INTERNATIONAL SEARCH REPORT

International application No.
PCT/SE 2005/001307

C (Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 6545979 B1 (ANDRE POULIN), 8 April 2003 (08.04.2003), column 1, line 55 - column 3, line 13 -- -----	1-15

INTERNATIONAL SEARCH REPORT

International application No.
PCT/SE2005/001307

Continuation of second sheet

H04L 7/00 (2006.01)

G06F 1/12 (2006.01)

H04J 3/06 (2006.01)

INTERNATIONAL SEARCH REPORT

Information on patent family members

29/10/2005

International application No.

PCT/SE 2005/001307

US	20030048811	A1	13/03/2003	CA	2360009 A	11/03/2003
EP	1215559	A2	19/06/2002	US	20030123491 A	03/07/2003
US	20030048754	A1	13/03/2003	NONE		
US	20010000071	A1	29/03/2001	US	6363073 B	26/03/2002
				AU	2327499 A	09/08/1999
				CN	1289489 A,T	28/03/2001
				EP	1050125 A	08/11/2000
				JP	2002502146 T	22/01/2002
				US	6157646 A	05/12/2000
				WO	9938285 A	29/07/1999
US	6545979	B1	08/04/2003	NONE		