

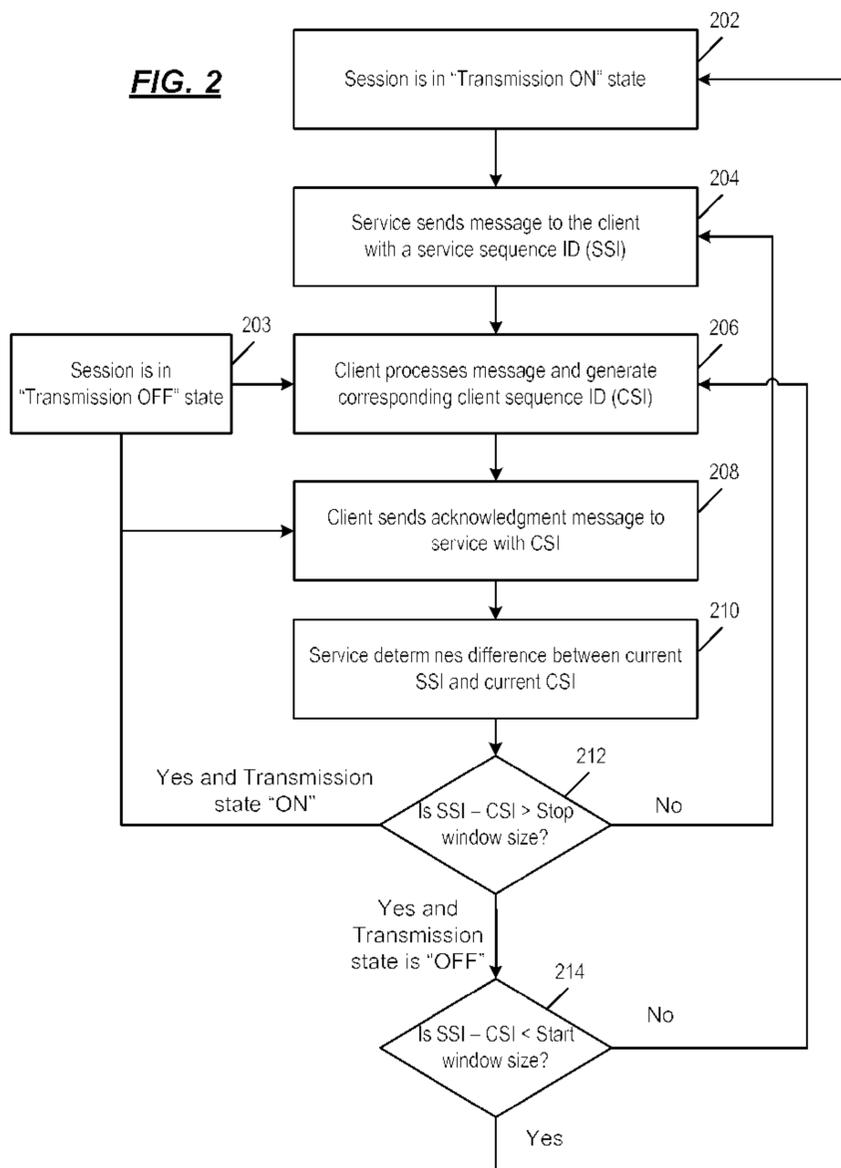


(86) **Date de dépôt PCT/PCT Filing Date:** 2014/11/06  
 (87) **Date publication PCT/PCT Publication Date:** 2015/05/14  
 (85) **Entrée phase nationale/National Entry:** 2016/05/04  
 (86) **N° demande PCT/PCT Application No.:** US 2014/064242  
 (87) **N° publication PCT/PCT Publication No.:** 2015/069837  
 (30) **Priorités/Priorities:** 2013/11/06 (US61/900,940);  
 2013/11/29 (US61/910,189)

(51) **Cl.Int./Int.Cl. H04L 12/825** (2013.01)  
 (71) **Demandeur/Applicant:**  
 CALGARY SCIENTIFIC, INC., CA  
 (72) **Inventeurs/Inventors:**  
 LEITCH, SAM ANTHONY, US;  
 ROBINSON, KEVIN GLEN, CA  
 (74) **Agent:** POWELL, MATTHEW D.

(54) **Titre : APPAREIL ET PROCEDE POUR COMMANDE DE FLUX COTE CLIENT DANS UN ENVIRONNEMENT D'ACCES A DISTANCE**

(54) **Title: APPARATUS AND METHOD FOR CLIENT-SIDE FLOW CONTROL IN A REMOTE ACCESS ENVIRONMENT**



(57) **Abrégé/Abstract:**

A client]side flow control mechanism for controlling the flow of messages communicated from a service to a client in a remote access system. The service, which may be a managed or unmanaged service, may communicate messages to the client where the

**(57) Abrégé(suite)/Abstract(continued):**

messages include a service sequence ID (SSI). As the client processes the messages from the service, the client responds with a client sequence ID (CSI). The SSI and CSI values determine the number of messages that are in transmission queues. The service compares a difference between the current SSI and CSI to a stop window value. If the difference is greater than the stop window value, transmission of messages from the service to the client is turned OFF. The client continues to send CSI acknowledgments as it processes messages. Once the difference falls below a start window value, the transmission of messages to the client is turned ON.

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property  
Organization  
International Bureau(10) International Publication Number  
**WO 2015/069837 A1**(43) International Publication Date  
14 May 2015 (14.05.2015)(51) International Patent Classification:  
*H04L 12/825* (2013.01)

(21) International Application Number:

PCT/US2014/064242

(22) International Filing Date:

6 November 2014 (06.11.2014)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

61/900,940 6 November 2013 (06.11.2013) US

61/910,189 29 November 2013 (29.11.2013) US

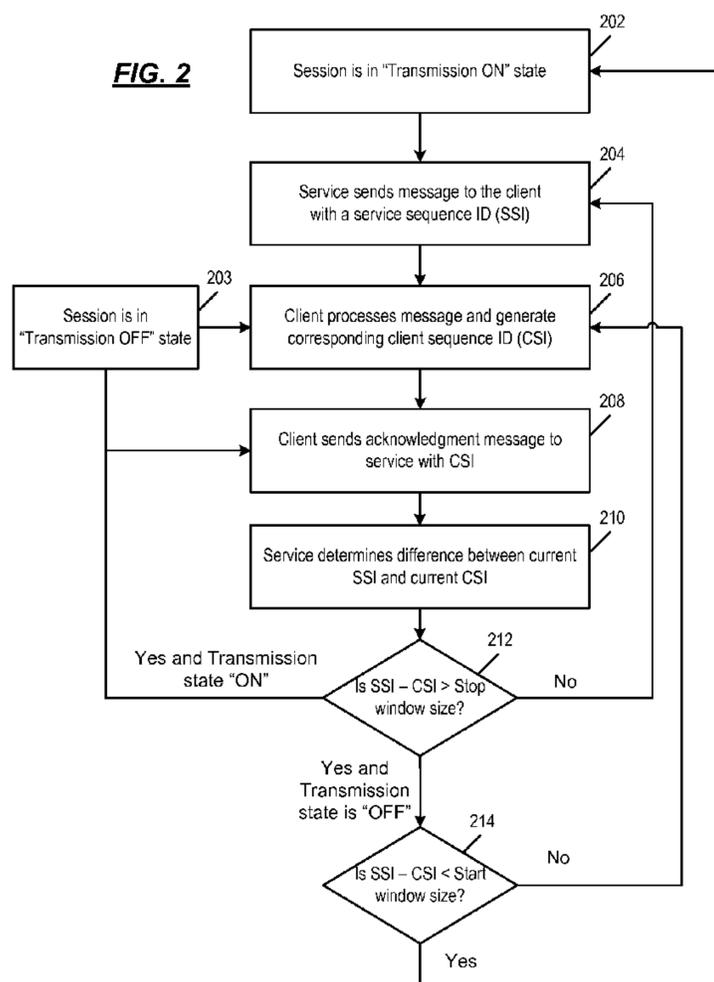
(71) Applicant: **CALGARY SCIENTIFIC, INC.** [CA/CA];  
Suite 208, 1210 - 20th Avenue SE, Calgary, Alberta T2G  
1M8 (CA).

(72) Inventors; and

(71) Applicants : **LEITCH, Sam Anthony** [CA/US]; 21430  
SE 1st Pl, Sammamish, Washington 98074-7088 (US).  
**ROBINSON, Kevin Glen** [CA/CA]; 83 Galbraith Drive,  
SW, Calgary, Alberta T3E 4Z5 (CA).(74) Agents: **AARONSON, Lawrence A.** et al.; Meunier  
Carlin & Curfman LLC, 999 Peachtree Street, NE, Suite  
1300, Atlanta, Georgia 30309 (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, DK, DM,  
DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT,  
HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR,  
KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG,  
MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM,  
PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC,  
SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN,  
TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.(84) Designated States (unless otherwise indicated, for every  
kind of regional protection available): ARIPO (BW, GH,  
GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ,  
TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU,  
TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE,  
DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU,  
LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK,  
SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ,  
GW, KM, ML, MR, NE, SN, TD, TG).

[Continued on next page]

(54) Title: APPARATUS AND METHOD FOR CLIENT-SIDE FLOW CONTROL IN A REMOTE ACCESS ENVIRONMENT



(57) Abstract: A client]side flow control mechanism for controlling the flow of messages communicated from a service to a client in a remote access system. The service, which may be a managed or unmanaged service, may communicate messages to the client where the messages include a service sequence ID (SSI). As the client processes the messages from the service, the client responds with a client sequence ID (CSI). The SSI and CSI values determine the number of messages that are in transmission queues. The service compares a difference between the current SSI and CSI to a stop window value. If the difference is greater than the stop window value, transmission of messages from the service to the client is turned OFF. The client continues to send CSI acknowledgments as it processes messages. Once the difference falls below a start window value, the transmission of messages to the client is turned ON.

WO 2015/069837 A1

**WO 2015/069837 A1** 

---

**Published:**

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

**APPARATUS AND METHOD FOR CLIENT-SIDE FLOW CONTROL IN A REMOTE ACCESS ENVIRONMENT**

**CROSS-REFERENCE TO RELATED APPLICATION**

[0001] This application claims priority to U.S. Provisional Patent Application Nos. 61/900,940, filed November 6, 2013, entitled "METHOD FOR CLIENT-SIDE FLOW CONTROL IN A REMOTE ACCESS ENVIRONMENT," and 61/910,189, filed November 29, 2013, entitled "METHOD FOR SERVER-SERVICE SEPARATION WITH END-TO-END FLOW CONTROL IN A CLIENT-SERVER REMOTE ACCESS ENVIRONMENT," the disclosures of which are incorporated herein by reference in their entities.

**BACKGROUND**

[0002] Ubiquitous remote access to services has become commonplace as a result of the growth and availability of broadband and wireless network access. As such, users are accessing services using an ever-growing variety of client devices (e.g., mobile devices, tablet computing devices, laptop/notebook/desktop computers, etc.). A remote server may communicate messages that contain data or other information between services and client devices over a variety of networks including, 3G and 4G mobile data networks, wireless networks such as WiFi and WiMax, wired networks, etc.

[0003] A problem arises when there is enough network bandwidth to send messages from the service to a connected client, but the client cannot process the messages quickly enough. For example, with some APIs, networking is push-based. As such, during periods when the client is idle or performing background tasks, additional messages may be sent and queued

WO 2015/069837

PCT/US2014/064242

at the client. If too many are sent and the client cannot process them in a timely fashion, errors may occur.

#### SUMMARY

**[0004]** Disclosed herein are systems and methods for providing flow control in a remote access system. In accordance with aspects of the disclosure, the method may include communicating messages from a service to a client, the service being remotely accessed by the client and the messages including a service sequence ID (SSI); receiving, at the service, a client sequence ID (CSI) from the client in response to the client completing processing of a message having a specific SSI associated with the message; determining, at the service, a difference between a current SSI communicated to the client and the current CSI received from the client; and if the difference is greater than a predetermined stop window value, turning communication from the service to the client OFF.

**[0005]** In accordance with other aspects of the disclosure, a method of providing client-side flow control is described that may include communicating a message from a service to a client, the message including a first sequence number that is incremented by the service to represent a relative position of the message within a plurality of messages that are communicated from the service to the client; receiving, at the service, a response from the client indicating a completion of processing of a message having a second sequence number less than first sequence number; determining, at the service from the response, if the processing of messages at the client is lagging beyond a first predetermined value and turning communication from the service to the client OFF if the client is lagging beyond the predetermined value.

**[0006]** In accordance with yet other aspects of the disclosure, an apparatus for providing client-side flow control of communication of messages from a service to a client is described. The apparatus may include a remote access server having a server layer as a communications proxy for messages sent between the client and the service, the server layer containing an outbound client queue of messages destined for the client; and a second server executing a service layer associated with the service, the service layer having a service receive queue for messages received from the client and a service send queue of messages destined for the client. The messages destined for the client may include a service sequence ID (SSI) and the messages received from the client include a client sequence ID (CSI) indicative of the client completing processing of a message having a specific SSI associated with the message. The service may determine if a difference between a current SSI and the CSI received from the client is greater than a predetermined stop window value to turn communication from the service to the client OFF.

**[0007]** In accordance with an aspect of the disclosure, a method of providing client-side flow control in a remote access system is described. The method may include communicating messages from a service to a client, the service being remotely accessed by the client and the messages including a service sequence ID (SSI); receiving a client sequence ID (CSI) from the client in response to the client completing processing of a message having an SSI associated with the CSI; determining, at the service, a difference between a current SSI communicated to the client and the CSI received from the client; and if the difference is greater than a predetermined stop window value, turning communication from the service to the client OFF.

**[0008]** Other systems, methods, features and/or advantages will be or may become apparent to one with skill in the art upon examination of the following drawings and detailed

description. It is intended that all such additional systems, methods, features and/or advantages be included within this description and be protected by the accompanying claims.

#### BRIEF DESCRIPTION OF THE DRAWINGS

[0009] The components in the drawings are not necessarily to scale relative to each other. Like reference numerals designate corresponding parts throughout the several views.

[0010] FIG. 1 illustrates an example layered architecture in accordance with the present disclosure;

[0011] FIG. 2 illustrates a flow diagram of example operations performed within the architecture of FIG. 1;

[0012] FIG. 3 illustrates a flow diagram of example operations performed within the architecture of FIG. 1;

[0013] FIGS. 4A and 4B, illustrate example server-service models for client remote access to services in a layered architecture, such as shown in FIG. 1; and

[0014] FIG. 5 illustrates an exemplary computing device.

#### DETAILED DESCRIPTION

[0015] Unless defined otherwise, all technical and scientific terms used herein have the same meaning as commonly understood by one of ordinary skill in the art. Methods and materials similar or equivalent to those described herein can be used in the practice or testing of the present disclosure. While implementations will be described for remotely accessing services, it will become evident to those skilled in the art that the implementations are not limited thereto, but are applicable for remotely accessing any type of service, including applications and data, via a remote device.

[0016] With the above overview as an introduction, reference is now made to FIG. 1, which illustrates a layered architecture in accordance with the present disclosure. The layered architecture implements a client flow control mechanism which allows a client 102 to control the flow of information from a service executing on an application server 106 through a remote access server 104. As illustrated in FIG. 1, there are three layers in the system: a client layer 112, a server layer 114 and a service layer 116. In accordance with the present disclosure, the layered architecture includes messaging queues that account for the particularities of client ability to process messages containing data or other information from the application server 106. The client layer 112 resides in the client 102 and is used to connect to the server layer 114. Typically, the client layer 112 includes a client application, e.g., a web browser, dedicated application, etc., used to provide a user interface at the client 102. The client 102 may be wireless handheld devices such as, for example, an IPHONE, an ANDROID-based device, a tablet device or a desktop/notebook personal computer that are connected by a communication network 125 to the remote access server 104.

[0017] The remote access server 104 may include a server remote access program that is used to connect the client 102 to the application server 106, which may be, e.g., a medical application. An example of the server remote access program is PUREWEB, available from Calgary Scientific, Inc. of Calgary, Alberta, Canada. The server remote access program may optionally provide for connection marshalling and application process management across the architecture of FIG. 1. The remote access server 104 may be connected to the application server 106 over a communication connection 126. If the application server 106 is executing on a computing device other than the remote access server 104, the communication connection 126 may be a communication network. For example, the communication connection 126 may be a TCP/IP communications network, a VPN connection, a dedicated connection, etc. If the

application server 106 is executing on the same node or computer as the remote access server, the communication connection 126 may be a TCP/IP socket connection, shared memory space, shared file location, etc.

**[0018]** The server layer 114 serves to proxy communications from the client 102 to the service layer 116, as described below. The client layer 112 contains a client receive queue 120 of messages received from the remote access server 104, which are sent on behalf of the application server 106 to which the client 102 is logically connected. The client layer 112 also contains a client send queue 122 of messages it sends to a receive queue 123 in server layer 114 of the remote access server 104. The messages may be ultimately destined to the application server 106, for example.

**[0019]** An instance of the service layer 116 (e.g., associated with the service) connects to the server layer 114, which proxies communications from the application server 106 to the client 102 logically associated with that service. The service layer 116 contains a service receive queue 128 of messages it receives from the server 104 on behalf of the client 102. Each message may be labeled with the unique identifier of the client 102. The service layer 116 also contains a service send queue 130 of messages it sends to a client queue 124 in the server layer 114 of the remote access server 104. Each message in the service send queue 130 may also be labeled with the unique identifier of the client 102. Additionally details regarding types of services is provided with reference to FIGS. 4A and 4B.

**[0020]** The client flow control mechanism of the present disclosure may be included as part of a service Software Development Kit (SDK) implemented in the service layer 116 that causes the service to stop sending messages when the network connection between server 104 and client 102 is saturated as determined by message tracking system. As will be described below with reference to FIG. 2, a sequence counter may be added to track the current service

WO 2015/069837

PCT/US2014/064242

sequence ID (SSI) sent to a client and current client sequence ID (CSI) received from the client for each connected session. The client layer 112 listens for messages from the service execution on the application server 106 that have the service sequence ID (SSI). Whenever the client receives and processes a message having an SSI, it sends the client sequence ID (CSI) as an acknowledgement to the service. Transmission to the client 102 may be turned OFF when difference between the service sequence ID (SSI) and the client sequence ID (CSI) is larger than a predetermined OffWindowSize value. In other words, the combination of the SSI and CSI indicate the number of messages in queues 120, 124 and 130 is greater than the OffWindowSize value. Transmission to the client 102 may be turned ON when the difference between the service sequence ID (SSI) the client sequence ID (CSI) is below an OnWindowSize value. As such, the service will be able to determine how far behind the client 102 is in its message processing. The OnWindowSize value allows the client to "catch up" before the service sends more data.

Service sequence ID: SSI  
Client sequence ID: CSI

## Flow control message

	CSI	SSI
Client transmit on (initially)	0	0
Service sends SequenceID		
Client receives/processes	0	1

## Acknowledgement

Client sends update SequenceID		
Service receives/processes	1	1

If SSI-CSI > stop window size  
Client transmit OFF

If SSI-CSI < Start window size  
Client transmit ON

Stop window size  $\geq$  Start window size

[0021] FIG. 2 is an operational flow diagram 200 of providing client-side flow control. At 202, a session begins having its transmission state "ON." As such, communication between the client 102 and the service execution on the application server 106 is established and messages are flowing. At 204, the service sends a message to the client with the service sequence ID (SSI). The SSI is incrementing identifier sent by the application server 106. As described above, each time the service at the application server 106 sends a packet of messages to the client 102, it sends a sequence ID to client 102. At 206, the inbound message having the SSI is processed by the client and the client sequence ID (CSI) is generated. As noted above, the CSI may be a reflection of the SSI contained in the message processed by the client 102. At 208, the client sends acknowledgment message to the service with the CSI. In accordance with aspects of the present disclosure, the CSI send by the client 102 is sent after client 102 processes the message from the application server 106. For example, if the SSI is "1," then the client 102 may reflect "1" as the CSI in the acknowledgement message to the application server 106.

[0022] At 210, the service determines the difference between the current SSI and the current CSI. The application server 106 maintains the current SSI and CSI values. A difference between these values may be used by the application server 106 to determine if the client 102 is overwhelmed or otherwise lagging behind the application server 106. At 212, if the difference between the SSI the CSI is larger than a stop window size (e.g., the OffWindowSize value) and the current state transmission state is "ON," then the session is set to an "OFF" transmission state at 203. This change of transmission state will stop the application server 106 from sending additional messages to the client 102. The application server 106 will continue to track the SSI and CSI values, because the client 102 may send messages updating the CSI as it continues to process messages. While the transmission state is OFF, messages are queued in

WO 2015/069837

PCT/US2014/064242

the service send queue 130 in the service layer 116. Processing then continues at 206 where the client 102 continues process queued messages in the client receive queue 120.

**[0023]** If, at 212, the difference between the SSI and the CSI is not greater than the stop window size, then processing returns at 204 where the service sends a next message to the client with an incremented SSI. The application server 106 will send the next message in the service send queue 130 to the client 102. Processing continues, as described above.

**[0024]** If at 212, if the current transmission state is "OFF," then at 214 it is determined if the difference between the SSI and the CSI is less than a start window size (e.g., the OnWindowSize value). If not, then processing returns to 206, where the client 102 processes the next message in the client receive queue 120. If at 214 the difference between the SSI and the CSI is less than the start window size, then the client is ready to process messages from service, and the transmission state is set to "ON" at 202.

**[0025]** In the above operational flow, example stop window (OffWindowSize) value is 5. It is noted that this value may be adjusted based on weights assigned to the communication session between client 102 and application server 106. An example start window (OnWindowSize) value is 2. In accordance with some implementations, the stop window value and the start window value are different in order to prevent a situation where transmission rapidly toggles back and forth between ON and OFF.

**[0026]** In accordance with yet other implementations, the start and stop window values can be dynamically adjusted based on, e.g., historical performance of the client 102, a type of network being used (e.g., Ethernet vs. cellular data), etc. The difference between the start window value and stop window value may be made bigger as transmissions between the application server 106 and client 102 are successful and made smaller where there is trouble in the transmissions.

**[0027]** Thus, as described above, the application server 106 makes all flow control decisions based on information received from the client 102. Further, the process 200 prevents delays in informing in the application server 106 that there is a delay at the client 102. This allows the message flow to adapt to quality/performance of the client 102. For example, even if the client 102 slows down and cannot process messages quickly enough, the flow control mechanism of FIG. 2 will stop the transmission of messages, even though the network is fully capable of sending messages.

**[0028]** In accordance with the above, the operational flow of FIG. 2 may also be used to account for transmission characteristics of the communication connection 126. For example, the communication connection 126 may not provide the high bandwidth of an internal system bus of a computing device. Therefore, the operational flow 200 address latency that may occur during communication between the application server 106 and the remote access server 104.

**[0029]** Accordance with the other implementations, the operational flow 200 may provide for interoperability among services and/or clients that do not recognize or send an SSI and/or CSI. For example, if the application server 106 does not send SSI, the client 102, may be adapted such that it does not send a CSI. Also, if the service sends an SSI, but client 102, does not respond with CSI, then the service will assume that the client is not enabled.

**[0030]** In some implementations, two or more client devices may collaboratively interact in a collaborative session with the service using the remote access server 104. The remote access server 104 may utilize a state model to present a synchronized view of display data associated with the application server 106. Flow control may be implemented on a per-client basis. Where transmission is OFF to a particular client, that client may skip messages to re-synchronize its view with other clients in the collaborative session when transmission returns to an ON state.

**[0031]** FIG. 3 illustrates an operational flow diagram 300 of providing reliable messaging associated between a service and a logically connected client utilizing the SSI and CSI. For example, in some implementations, the CSI responses from the client may be used so that the application server 106 knows what messages have been successfully received by the client 102. At 302, the service sends a message to the client with the SSI. Next, at 302, the client sends a message to the service with the CSI value. If the application server 106 receives this message, then at 306, the application server 106 can safely remove the message associated with the associated SSI value from the service send queue 130. However, in instances where the CSI is not received, the service may resend the message with the corresponding SSI (308). Similarly, the process 302-308 may be used at the client. For example, every time the application server 106 sends a message to the client 102, the received SSI value may be used to safely remove the messages associated with the previous SSI value from the client send queue 122. Thus, the SSI and CSI values may be used to provide reliable communications between the application server 106 and the client 102.

**[0032]** Reference is now made to FIGS. 4A and 4B, which illustrate examples of alternative models for client remote access to services in a layered architecture, such as that shown in FIG. 1 in which the client-side flow mechanism described above may be implemented. Fig. 4A illustrates a “managed service” 115 in which the services are deployed on the same system node or computing device as the remote access and application server to which client devices communicate. In this case, the service is managed in that the server controls the application/process life cycle by starting and stopping the service as client connect and disconnect. Fig. 4B illustrates an “unmanaged service” 117 in which services are deployed on servers at a different system node from the integrated remote access and application server the application/process life cycle is no longer managed by the remote access server.

**[0033]** In both FIGS. 4A and 4B, the server remote access program may provide for connection marshalling and application process management. Further both the managed service 115 and unmanaged service 117 may implement the client-side flow control of FIG. 2 and/or reliable messaging of FIG. 3. For example, in the environment of FIG. 4B is implemented, if the communication connection 126 experiences conditions that cause communication problems between the unmanaged service 117 and the remote access server 104, the tracking of the SSI and CSI values by the flow control mechanism of FIG. 2 may be used by the unmanaged service 117 to turn off communications until the conditions are more favorable.

**[0034]** As shown in FIG. 4A, the client 102 has the client layer 112 that may communicate to a remote access and application server 103 that includes the server layer 114 and the service layer 116. The client 102 may be connected by a communication network 125 to the remote access and application server 103. The remote access and application server 103 may include a server remote access program that executes in the server layer 114. The server remote access program is used to connect the client 102 to a managed service 115 (e.g., an application) executing in the service layer 116. Within the remote access and application server 103, the server remote access program in the server layer 114 may be connected to the service in the service layer 116 using a TCP socket connection and by, e.g., a system bus of the remote access and application server 103. Thus, bandwidth between the server remote access program and the service is extremely high. An example of the client 102 and the remote access and application server 103 is shown in FIG. 5.

**[0035]** Referring now to FIG. 4B, there is illustrated an example of an unmanaged service deployment. In such environments, a remote access server 104 includes the server layer 114 in which the server remote access program executes. An application server 106

includes the service layer 116 in which the service or application executes. In the environment of FIG. 4B, the service (shown as unmanaged service 117) is provided on a computing device other than a computing device on which the server remote access program executes and communicates over the communication connection 126 to the remote access server 104. Thus, the service is on a node separate from the server. An example of the remote access server 104 and the application server 106 is shown in FIG. 5.

**[0036]** In FIG. 4B, the client 102 connects to the remote access server 104 over communication connection 125. The application server 106 may connect to the remote access server 104 at a predetermined Internet Protocol (IP) address and/or socket, or using a Uniform Resource Locator (URL) associated with the remote access server 104 to register the service or application with the server remote access program executing on the remote access server 104. An external entity (the end user, or another process or application) launches the unmanaged service 117 outside of the remote access server 104. The unmanaged service 117, on startup, connects to the remote access server 104 using a server-service socket connection and establishes the session as a queued and unmanaged application to which a client may connect.

**[0037]** FIG. 5 shows an exemplary computing environment in which example embodiments and aspects may be implemented. The computing system environment is only one example of a suitable computing environment and is not intended to suggest any limitation as to the scope of use or functionality.

**[0038]** Numerous other general purpose or special purpose computing system environments or configurations may be used. Examples of well-known computing systems, environments, and/or configurations that may be suitable for use include, but are not limited to, personal computers, servers, handheld or laptop devices, multiprocessor systems, microprocessor-based systems, network personal computers (PCs), minicomputers, mainframe

computers, embedded systems, distributed computing environments that include any of the above systems or devices, and the like.

**[0039]** Computer-executable instructions, such as program modules, being executed by a computer may be used. Generally, program modules include routines, programs, objects, components, data structures, etc. that perform particular tasks or implement particular abstract data types. Distributed computing environments may be used where tasks are performed by remote processing devices that are linked through a communications network or other data transmission medium. In a distributed computing environment, program modules and other data may be located in both local and remote computer storage media including memory storage devices.

**[0040]** With reference to FIG. 5, an exemplary system for implementing aspects described herein includes a computing device, such as computing device 500. In its most basic configuration, computing device 500 typically includes at least one processing unit 502 and memory 504. Depending on the exact configuration and type of computing device, memory 504 may be volatile (such as random access memory (RAM)), non-volatile (such as read-only memory (ROM), flash memory, etc.), or some combination of the two. This most basic configuration is illustrated in FIG. 5 by dashed line 506.

**[0041]** Computing device 500 may have additional features/functionality. For example, computing device 500 may include additional storage (removable and/or non-removable) including, but not limited to, magnetic or optical disks or tape. Such additional storage is illustrated in FIG. 5 by removable storage 508 and non-removable storage 510.

**[0042]** Computing device 500 typically includes a variety of tangible computer readable media. Computer readable media can be any available tangible media that can be

accessed by device 500 and includes both volatile and non-volatile media, removable and non-removable media.

**[0043]** Tangible computer storage media include volatile and non-volatile, and removable and non-removable media implemented in any method or technology for storage of information such as computer readable instructions, data structures, program modules or other data. Memory 504, removable storage 508, and non-removable storage 510 are all examples of computer storage media. Tangible computer storage media include, but are not limited to, RAM, ROM, electrically erasable program read-only memory (EEPROM), flash memory or other memory technology, CD-ROM, digital versatile disks (DVD) or other optical storage, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices, or any other medium which can be used to store the desired information and which can be accessed by computing device 500. Any such computer storage media may be part of computing device 500.

**[0044]** Computing device 500 may contain communications connection(s) 512 that allow the device to communicate with other devices. Computing device 500 may also have input device(s) 514 such as a keyboard, mouse, pen, voice input device, touch input device, etc. Output device(s) 516 such as a display, speakers, printer, etc. may also be included. All these devices are well known in the art and need not be discussed at length here.

**[0045]** It should be understood that the various techniques described herein may be implemented in connection with hardware or software or, where appropriate, with a combination of both. Thus, the methods and apparatus of the presently disclosed subject matter, or certain aspects or portions thereof, may take the form of program code (i.e., instructions) embodied in tangible media, such as floppy diskettes, CD-ROMs, hard drives, or any other machine-readable storage medium wherein, when the program code is loaded into

WO 2015/069837

PCT/US2014/064242

and executed by a machine, such as a computer, the machine becomes an apparatus for practicing the presently disclosed subject matter. In the case of program code execution on programmable computers, the computing device generally includes a processor, a storage medium readable by the processor (including volatile and non-volatile memory and/or storage elements), at least one input device, and at least one output device. One or more programs may implement or utilize the processes described in connection with the presently disclosed subject matter, e.g., through the use of an application programming interface (API), reusable controls, or the like. Such programs may be implemented in a high level procedural or object-oriented programming language to communicate with a computer system. However, the program(s) can be implemented in assembly or machine language, if desired. In any case, the language may be a compiled or interpreted language and it may be combined with hardware implementations.

[0046] Although the subject matter has been described in language specific to structural features and/or methodological acts, it is to be understood that the subject matter defined in the appended claims is not necessarily limited to the specific features or acts described above. Rather, the specific features and acts described above are disclosed as example forms of implementing the claims.

WO 2015/069837

PCT/US2014/064242

**WHAT IS CLAIMED:**

1. A method of providing client-side flow control in a remote access system, comprising:
  - communicating messages from a service to a client, the service being remotely accessed by the client and the messages including a service sequence ID (SSI);
  - receiving, at the service, a client sequence ID (CSI) from the client in response to the client completing processing of a message having a specific SSI associated with the message;
  - determining, at the service, a difference between a current SSI communicated to the client and the current CSI received from the client; and
  - if the difference is greater than a predetermined stop window value, turning communication from the service to the client OFF.
  
2. The method of claim 1, further comprising:
  - subsequently determining, at the service, a difference between the current SSI and a subsequent CSI received from the client; and
  - if the difference is less than a predetermined start window value, turning communication from the service to the client ON.
  
3. The method of claim 2, wherein the predetermined start window value is less than the predetermined stop window value.
  
4. The method of any of claims 2-3, further comprising dynamically adjusting the start window value or the stop window value in accordance with one of network conditions and a processing capacity of the client.

WO 2015/069837

PCT/US2014/064242

5. The method of any of claims 1-4, further comprising maintaining the current SSI and CSI values at the service to perform the determining.

6. The method of any of claims 1-5, further comprising:

providing a client layer at the client;

providing a service layer associated with the service, the service layer having a service receive queue for messages received from the client and a service send queue of messages destined for the client; and

providing a server layer as a communications proxy for messages sent between the client and the service, the server layer containing an outbound client queue of messages destined for the client.

7. The method of claim 6, further comprising queue messages in the service send queue when the communication from the service to the client is OFF.

8. The method of any of claims 6-7, further comprising performing the determining in a service layer SDK provided in the service layer.

9. The method of any of claims 6-8, further comprising:

detecting a failure of communication between the client and the service by determining that a message have a particular SSI was not acknowledged by the client; and

upon detecting the failure, resending the message having the particular SSI.

WO 2015/069837

PCT/US2014/064242

10. The method of any of claims 6-9, further comprising resending the message from the service send queue.

11. The method of any of claims 1-10, further comprising:  
detecting a failure of communication between the client and the service using the SSI;  
and  
upon detecting the failure, resending a message having a particular SSI determined to have been unacknowledged by the client.

12. The method of any of claims 1-11, further comprising:  
providing the service as an unmanaged service executing on an application server;  
providing a server remote access program executing on a remote access server that communicates with the service over a first communication connection and with the client over a second communication connection;  
receiving a connection from the client at the remote access server over the second communication connection; and  
communicating messages from the service to the client in accordance with the determining.

13. The method of claim 12, further comprising providing the application server and the remote access server as different nodes.

14. A method of providing client-side flow control, comprising:

WO 2015/069837

PCT/US2014/064242

communicating a message from a service to a client, the message including a tracking mechanism that is incremented by the service to represent a relative position of the message within a plurality of messages that are communicated from the service to the client;

receiving, at the service, a response from the client indicating a completion of processing of a message according to the tracking mechanism;

determining, at the service from the response, if the processing of messages at the client is lagging beyond a first predetermined value and

turning communication from the service to the client OFF if the client is lagging beyond the predetermined value.

15. The method of claim 14, further comprising:

receiving a subsequent response at the service; and

determining, at the service from the subsequent response, if the processing of messages at the client is no longer lagging beyond a second predetermined value; and

turning communication from the service to the client ON.

16. The method of claim 15, wherein the second predetermined value is less than the first predetermined value.

17. The method of any of claims 15-16, further comprising dynamically adjusting the first predetermined value or the second predetermined value in accordance with one of network conditions and a processing capacity of the client.

WO 2015/069837

PCT/US2014/064242

18. An apparatus for providing client-side flow control of communication of messages from a service to a client, comprising:

a remote access server having a server layer as a communications proxy for messages sent between the client and the service, the server layer containing an outbound client queue of messages destined for the client;

a second server executing a service layer associated with the service, the service layer having a service receive queue for messages received from the client and a service send queue of messages destined for the client,

wherein the messages destined for the client include a service sequence ID (SSI) and the messages received from the client include a client sequence ID (CSI) indicative of the client completing processing of a message having a specific SSI associated with the message, and

wherein the service determines if a difference between a current SSI and the CSI received from the client is greater than a predetermined stop window value to turn communication from the service to the client OFF.

19. The apparatus of claim 18, wherein the service further determines if a difference between the current SSI and the CSI is less than a predetermined start window value to turn communication from the service to the client ON.

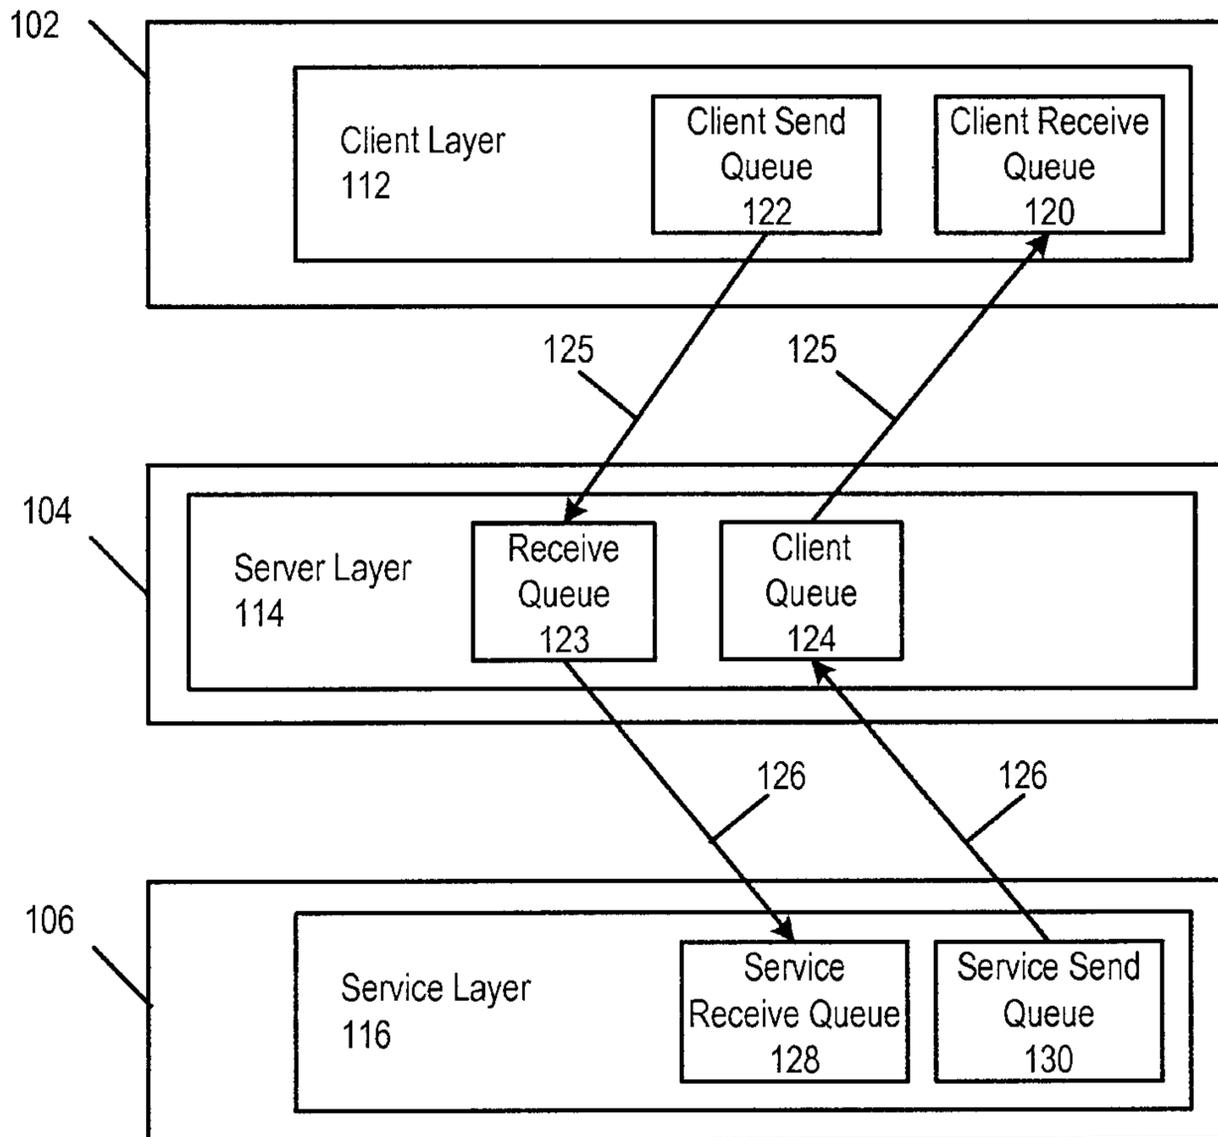
20. The apparatus of any of claims 18-19, wherein the predetermined start window value is less than the predetermined stop window value.

**WO 2015/069837**

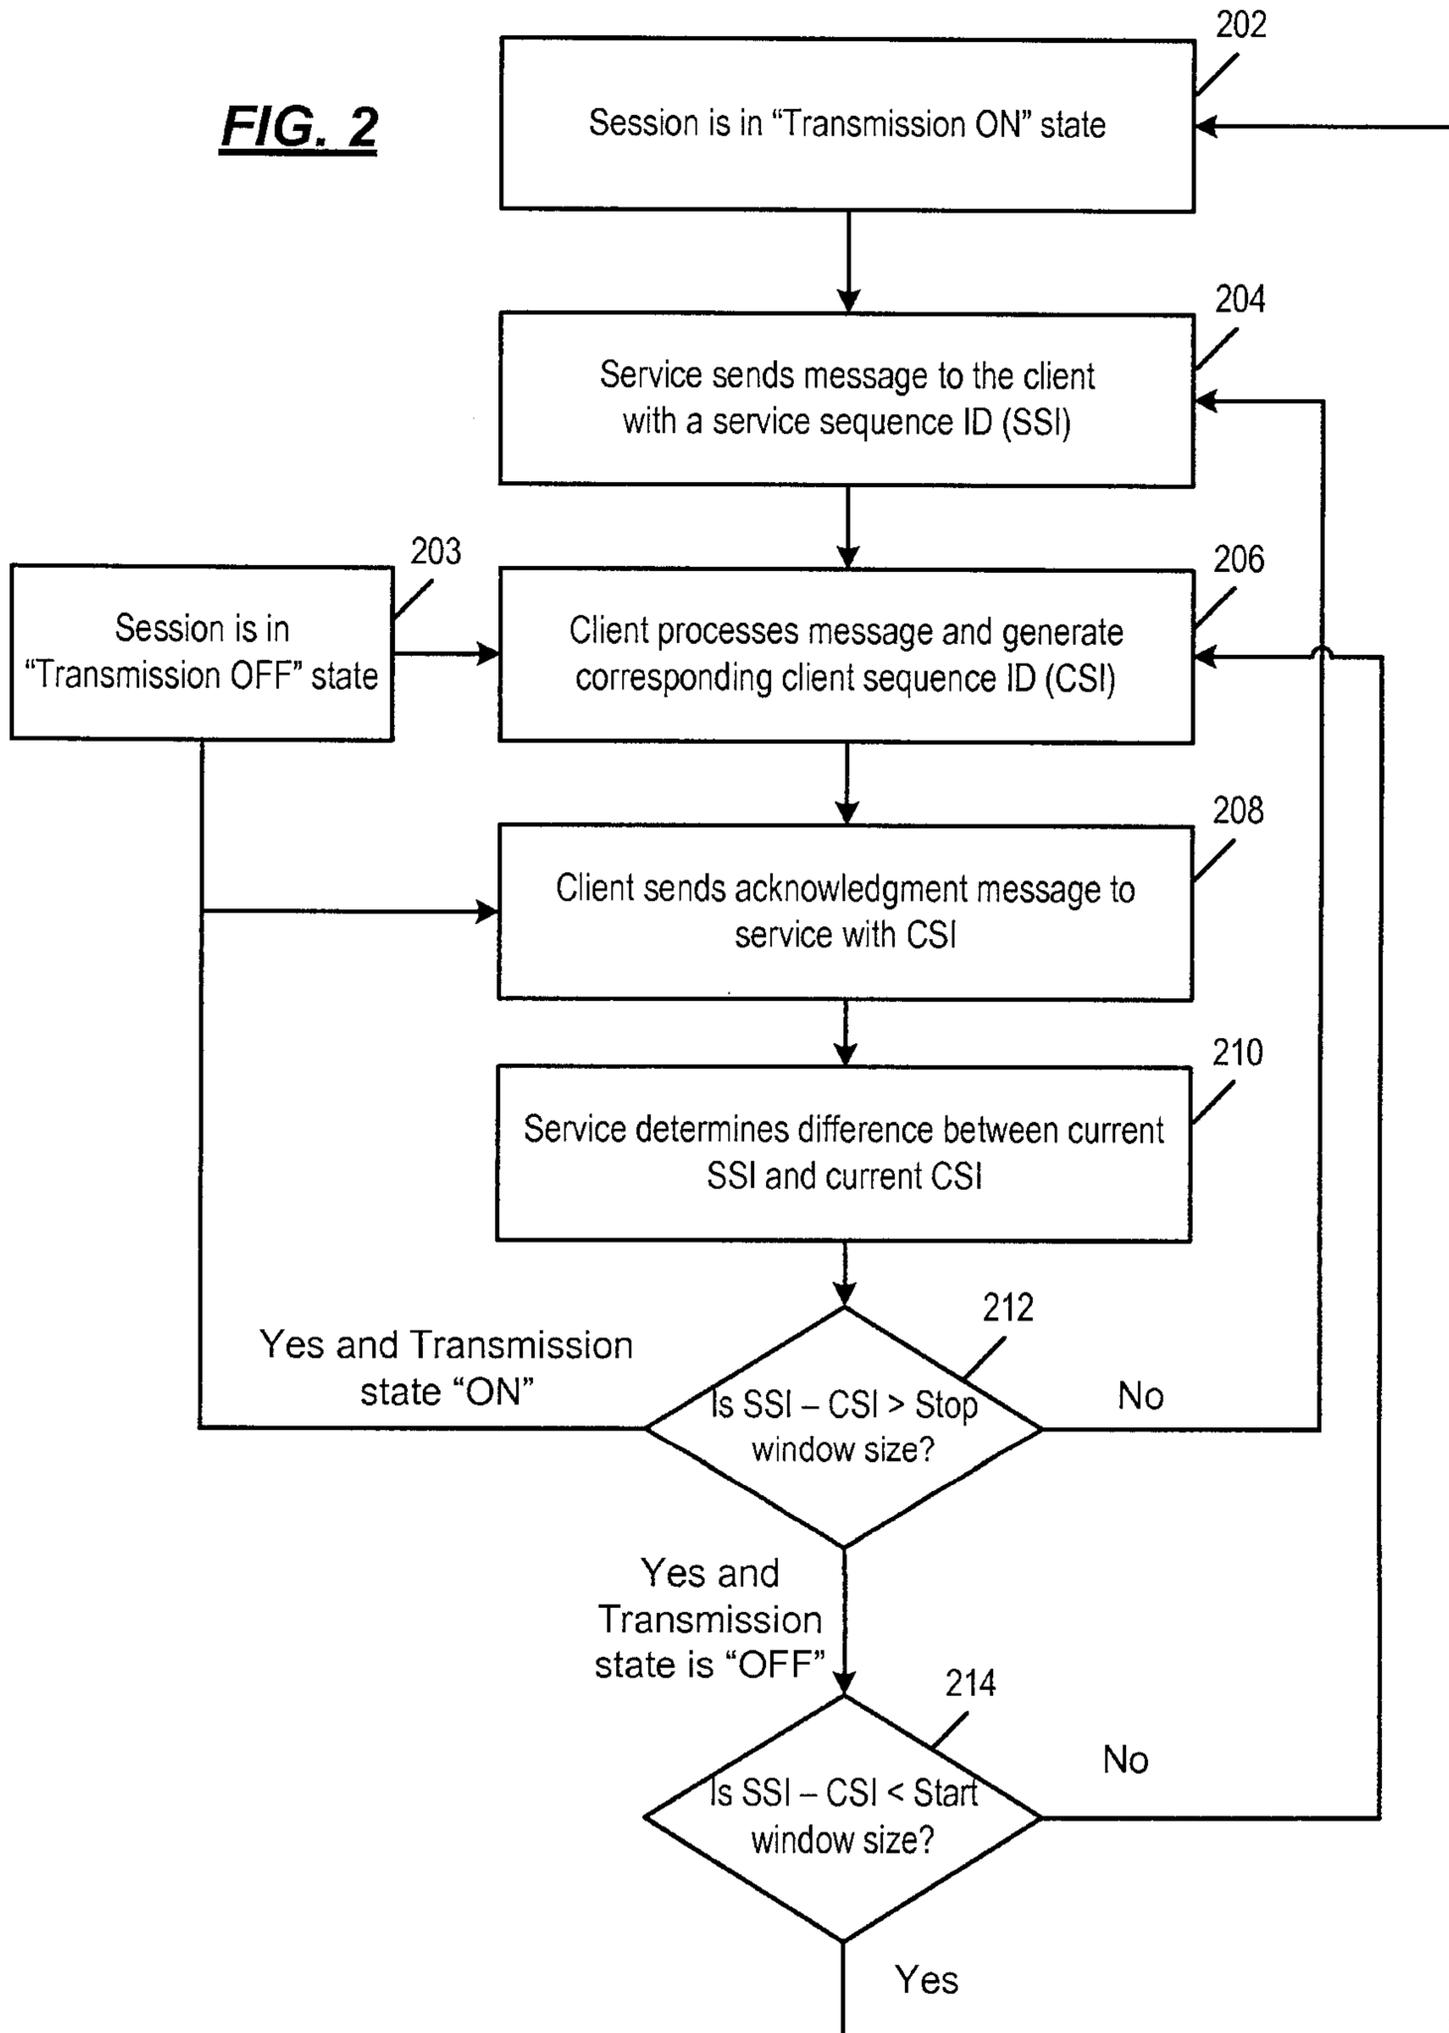
**PCT/US2014/064242**

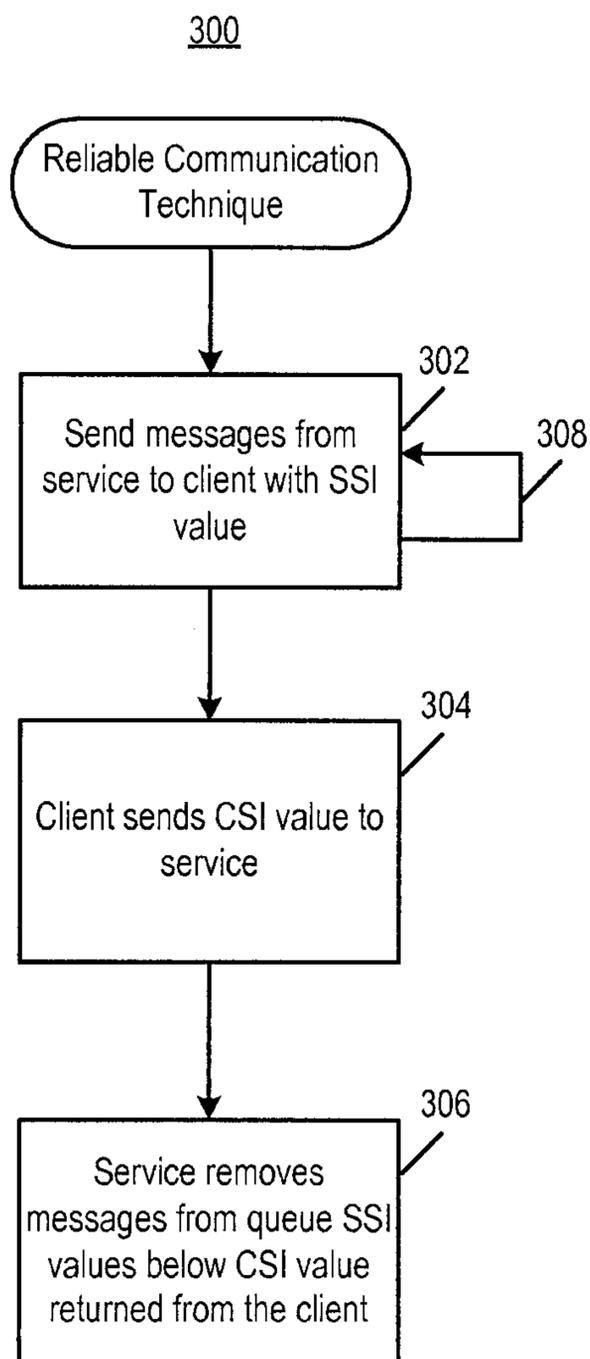
21. The apparatus of any of claims 19-20, wherein the start window value or the stop window value is dynamically adjusted in accordance with one of network conditions and a processing capacity of the client.

22. The apparatus of claim 18, wherein messages are queued at in service send queue when the communication from the service to the client is OFF.

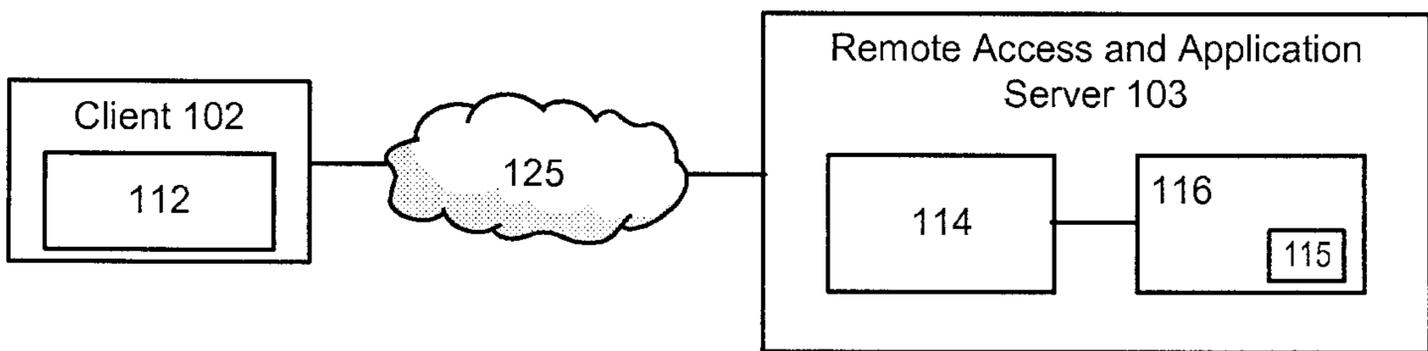


**FIG. 1**

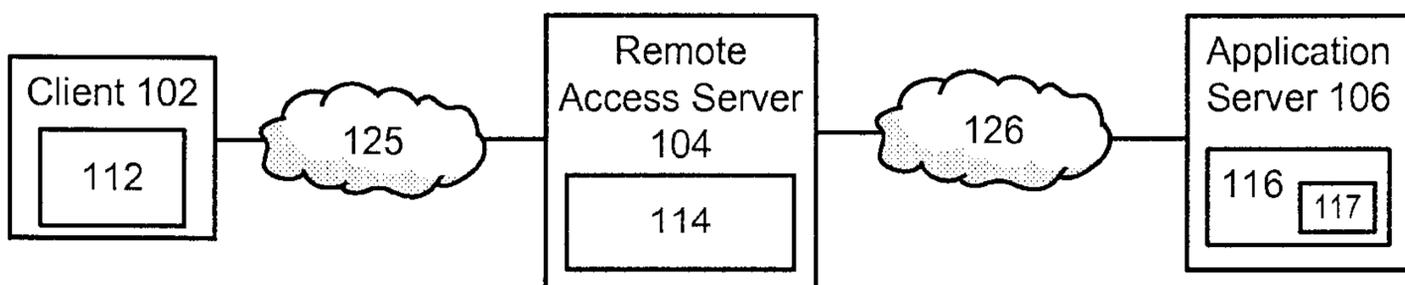




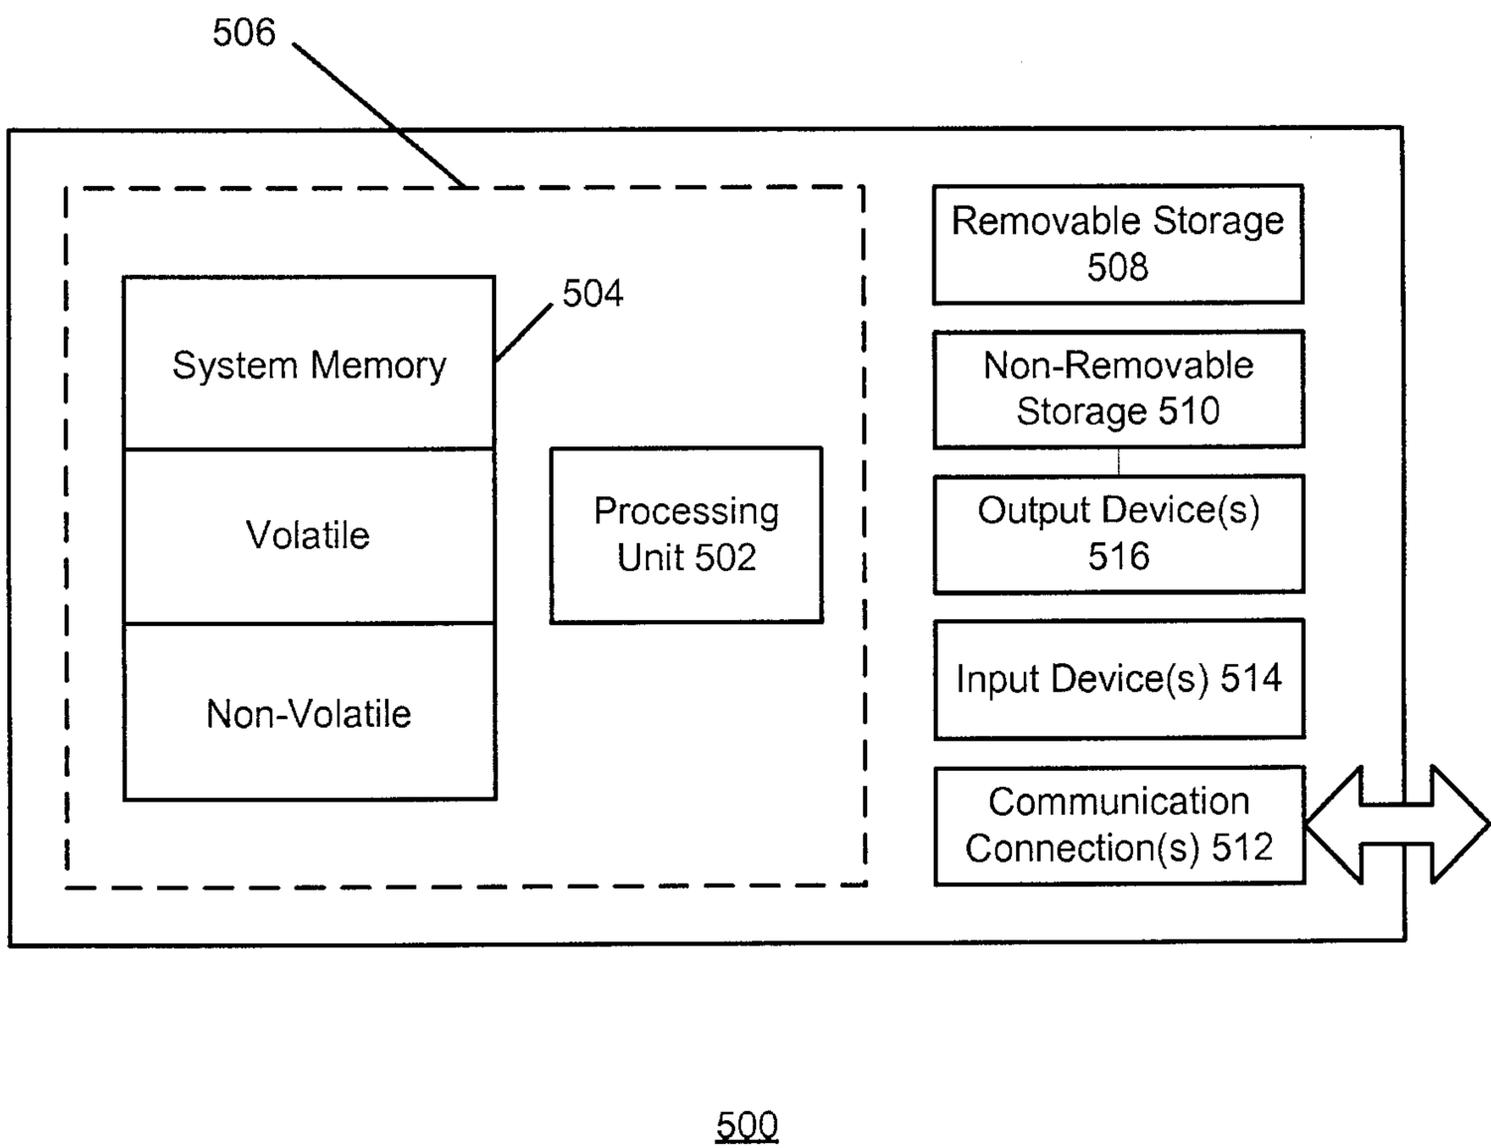
**FIG. 3**



**FIG. 4A**



**FIG. 4B**



**FIG. 5**

**FIG. 2**

