



(19) **United States**

(12) **Patent Application Publication** (10) **Pub. No.: US 2007/0191033 A1**

**Marais**

(43) **Pub. Date:**

**Aug. 16, 2007**

(54) **SCALABLE WIRELESS MESSAGING SYSTEM**

(75) Inventor: **Charl Marais**, Rueschlikon (CH)

Correspondence Address:  
**Oppedahl Patent Law Firm LLC**  
**P.O. BOX 4850**  
**FRISCO, CO 80443-4850 (US)**

(73) Assignee: **SOFTWIRED AG**, Bonstetten (CH)

(21) Appl. No.: **11/276,177**

(22) Filed: **Feb. 16, 2006**

**Publication Classification**

(51) **Int. Cl.**  
**H04Q 7/20** (2006.01)

(52) **U.S. Cl.** ..... **455/466**

(57) **ABSTRACT**

A method for operating a wireless messaging system, the messaging system being configured to comprise a plurality of mobile wireless clients (1, 1', 1''), a core messaging system (7) and at least two gateway processes or gateways

(5, 5', 5'') acting as proxies on behalf of the clients (1, 1', 1''), the method comprising, in order for a client (1) to establish a connection with one of the gateways (5, 5', 5''), a two phase connect procedure with the steps of,

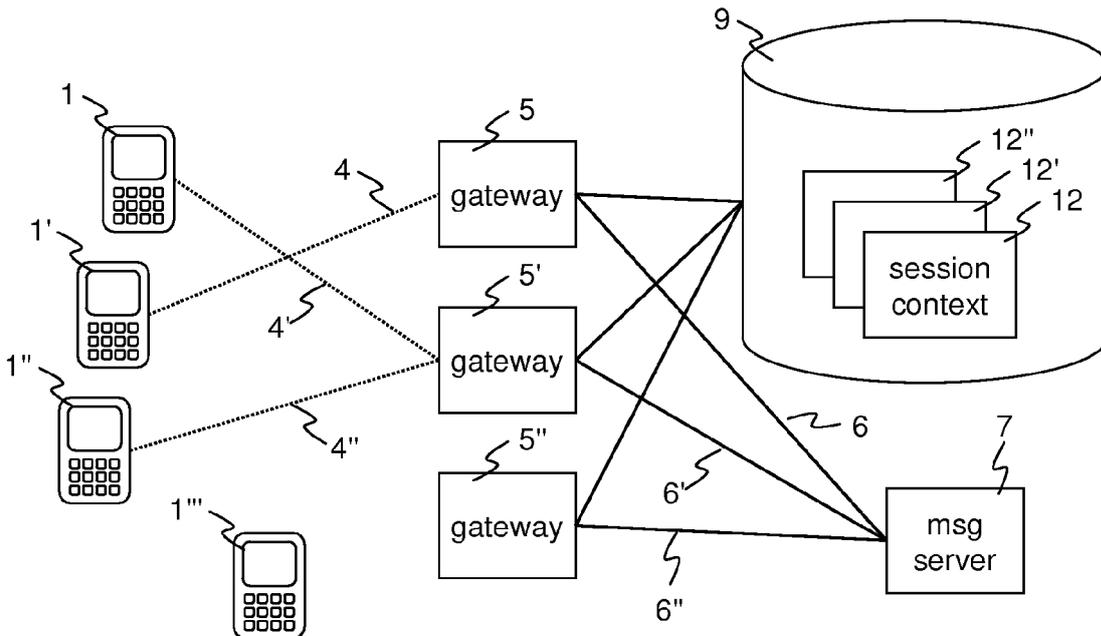
the client (1) maintaining a list of access points, each access point defining an address at one of the gateways (5, 5', 5'');

the client (1) sending, over a wireless bearer (4, 4', 4''), a phase one connection request to one of the access points;

the gateway (5, 5', 5'') that is addressed by the phase one connection request, sending, in response to the phase one connection request, over a wireless bearer (4, 4', 4''), a list of connect points to the client (1), each connect point defining an address at one of the gateways (5, 5', 5'');

the client (1) sending, over a wireless bearer (4, 4', 4''), a phase two connection request to at least one of the connect points; and

the gateway (5, 5', 5'') that is addressed by the phase two connection request responding to the phase two connection request, and the client (1) establishing a connection with the gateway (5, 5', 5'') specified by said connect point.



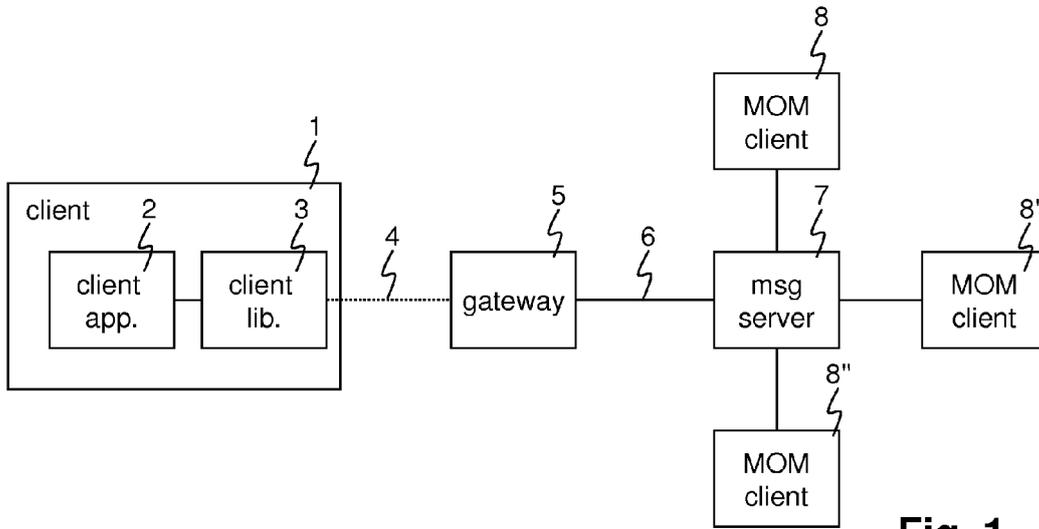


Fig. 1

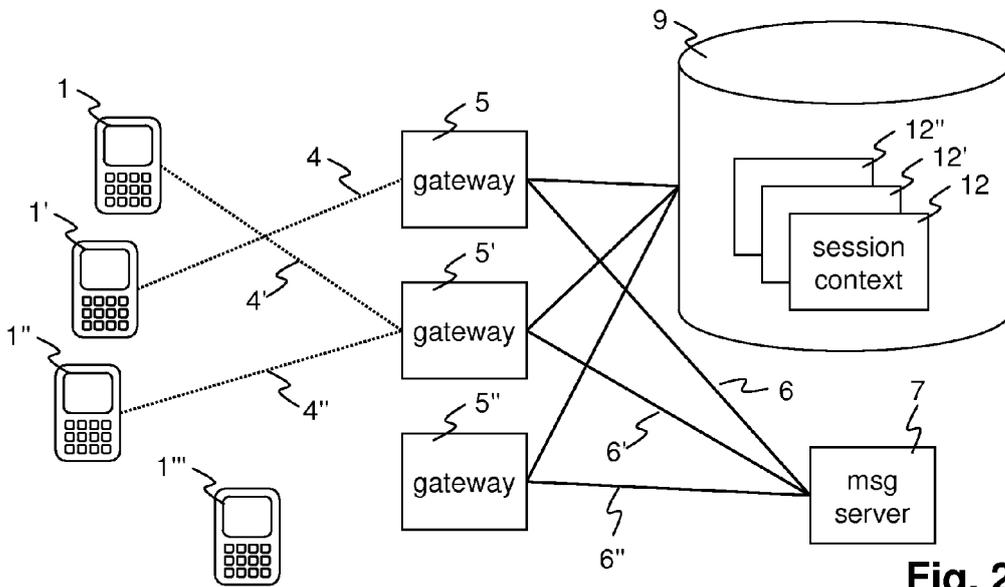


Fig. 2

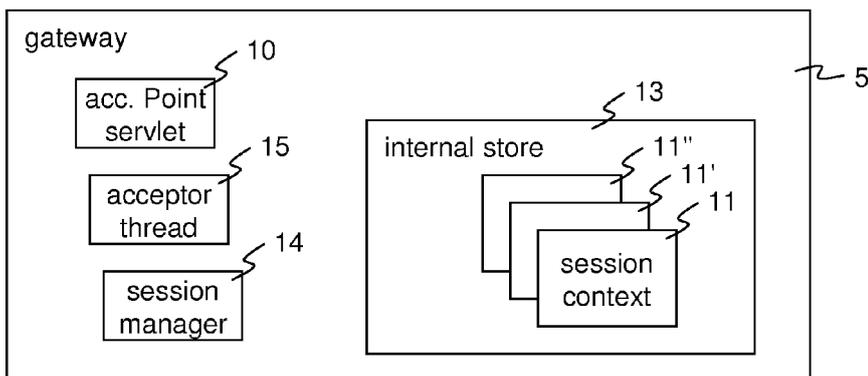


Fig. 3

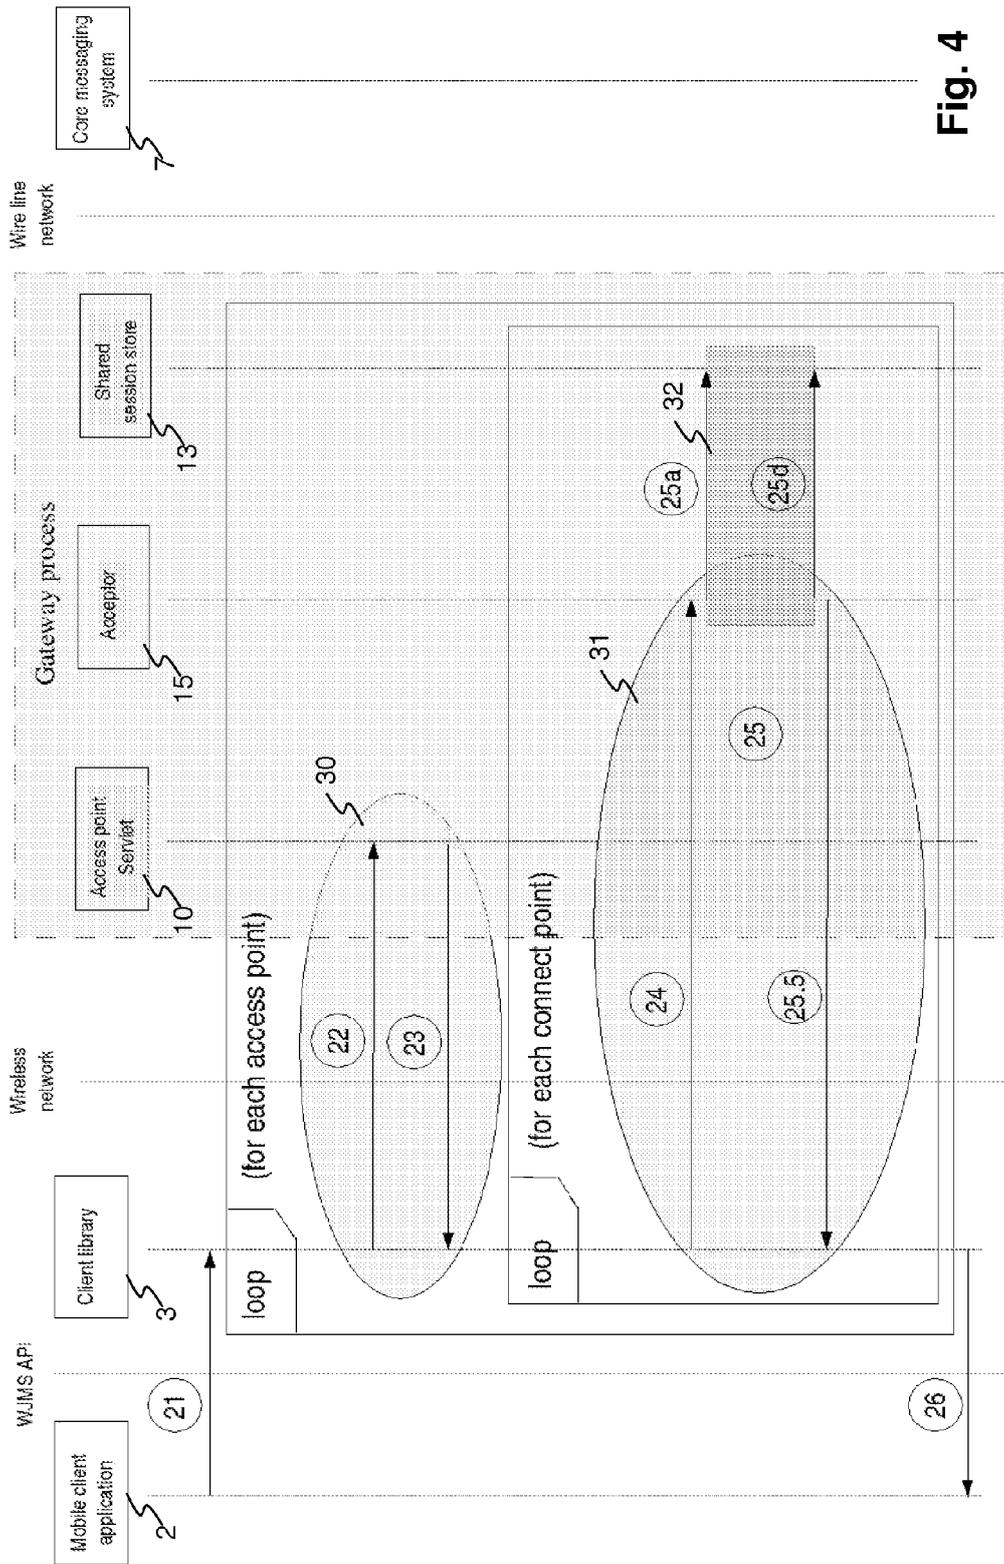


Fig. 4

## SCALABLE WIRELESS MESSAGING SYSTEM

### FIELD OF THE INVENTION

[0001] The invention relates to the field of messaging systems for wireless mobile devices and in particular to a method for operating a wireless messaging system, a corresponding computer program and a wireless messaging system for providing a gateway functionality by which wireline messaging systems can be extended to reach wireless mobile clients.

### BACKGROUND OF THE INVENTION

[0002] When working with a mobile application, that is, a program application running on a mobile device, it is desirable that the user may continue working with the application even when disconnected from the network. In such a situation, the end-to-end connection between device and central (wireline) system infrastructure is not available, due to a disruption on the unreliable wireless network component of the connection. In consequence, those applications are best built using a MOM (message oriented middleware) technology that combines client layers and a dedicated wireless gateway proxy that cooperate to shield mobile applications from wireless connection loss. Such MOM components use store and forward, session maintenance and sliding window to ensure reliable message delivery and client session integrity.

[0003] Messaging is a model in which applications are loosely coupled through the exchange of self-describing messages. This may be likened to “e-mail for applications”. Messaging allows producer applications to send messages to consumer applications, in a reliable and scalable manner. The messaging system decouples producers and consumers. It is in charge of transporting and storing messages. “Producer” and “Consumer” are just roles an application can play. An application can be both producer and consumer simultaneously, or switch roles dynamically at run-time. Two typical communication models implemented by a message system such as the commonly known JMS (Java Messaging System) are “Publish/Subscribe”, and “Point to Point messaging”.

[0004] An approach taken by Softwired Inc., with its MOM technology, iBus//Mobile and iBus//MessageServer products is to use Java and/or C# technology to extend connectivity straight through from the mobile devices to the centralized application servers. Customized applications are deployed on the mobile device, and on the server computers, using standards-based Java messaging middleware technology. A wireless data network such as GPRS, EDGE, UMTS, Mobitex, and future Wireless Broadband ‘4G’ is used to connect the mobile devices to the servers in a flexible and scalable manner.

[0005] The iBus//Mobile product defines a protocol between a mobile client layer, and a wireless proxy gateway. The mobile layer makes a JMS based API available to mobile applications, and propagates client application operations on that API across wireless networks (e.g. GSM, GPRS, UMTS) to the proxy using the protocol. Based on protocol messages received in packets across the wireless network, the gateway process is able to reflect/reproduce the mobile client operations (to the mobile API) on a proxy JMS connection between the gateway and the central messaging

infrastructure. The gateway maintains this central JMS connection on behalf of the mobile client.

[0006] Such a Wireless JMS (WJMS) deployment consists of Java-enabled and non Java mobile devices (PDAs, Communicators, JavaPhones, .Net Compact Framework devices etc.) with a WJMS client library installed. The devices use a wireless network to communicate with a WJMS gateway service. The gateway is a software component capable of receiving WJMS messages from a wireless network, and of relaying these messages to a server-side JMS provider. The JMS provider either is a stand-alone JMS messaging server such as Softwired’s iBus//MessageServer product, or a compliant J2EE application server such as IBM Websphere, Oracle Application Server, BEA’s WebLogic Server, SAP Web Application Server or JBoss Application Server product, which should contain a JMS layer. Such a Wireless JMS (WJMS) deployment consists of programmable mobile devices (PDAs, Communicators, JavaPhones, etc.) with the WJMS client layer installed. The devices use a wireless network to communicate with a WJMS gateway service. Provided the WJMS client layer that is compatible with the mobile platform such as JavaME, Java configurations CLDC, CDC (targeting profiles such as MIDP and IMP, .Net Compact Framework (C#)) is available, client applications can be deployed to various mobile environments (e.g. Linux, RIM Blackberry, Palm, Symbian, Windows CE etc).

[0007] Softwired’s U.S. Pat. No. 6,877,023 “Messaging system for delivering data in the form of portable message formats between message clients” and U.S. Pat No. 6,721, 779 “Messaging proxy system” disclose mechanisms for transporting the higher level WJMS messages over the wireless link while hiding the underlying communication layer. U.S. patent applications U.S. 20020120717 and 20030009511, disclose a message system for delivering data in the form of messages between message clients, the system comprising a server cluster with a group of client manager nodes and a group of independent message manager nodes. The client manager nodes have the function of managing client connections, whereas the message managers are configured to store and distribute messages.

[0008] With a large number of clients distributed over a large geographic area, and in order to ensure reliability, it is necessary to distribute the functionality of the gateway process, i.e. the interface between the wireless and the wirebound part of communications, over a number of geographically separated computers. A number of separate gateway processes may run on a single computer, or distributed over several physically separate computers.

### DESCRIPTION OF THE INVENTION

[0009] It is therefore an object of the invention to create a method for operating a wireless messaging system, a wireless messaging system and a computer program for operating a wireless messaging system of the type mentioned initially, which overcomes the disadvantages mentioned above.

[0010] It is a further object of the invention to provide a method for operating a gateway service that allows for a reliable connection procedure and connection reestablishment regardless of any individual gateway availability and connection failures.

[0011] It is a further object of the invention to provide a wireless messaging system that allows connection through multiple gateways while maintaining a consistent representation of connections and sessions, and may further be scaled to cover a large geographic area and/or large numbers of clients.

[0012] These objects are achieved by a method for operating a gateway for wireless mobile clients, a corresponding computer program and gateway according to the corresponding independent claims.

[0013] The method for operating a wireless messaging system has the messaging system being configured to comprise a plurality of mobile wireless clients, a core messaging system and at least two gateway processes or gateways acting as proxies on behalf of the clients. In the course of normal operation of the messaging system, the gateways maintain client sessions by performing the steps of,

[0014] the gateways receiving inbound messages from the clients over connections using a wireless protocol and forwarding said inbound messages to the core messaging system over a connection using a wirebound protocol;

[0015] the gateways receiving outbound messages from the core messaging system over a connection using a wirebound protocol and forwarding said outbound messages to the clients over connections using a wireless protocol;

[0016] In order for a client to establish a connection with one of the gateways, the method comprises, a two phase connect procedure with the steps of,

[0017] the client maintaining a list of access points, each access point defining an address at one of the gateways;

[0018] the client sending, over a wireless connection, a phase one connection request to one of the access points;

[0019] the gateway that is addressed by the phase one connection request, sending, in response to the phase one connection request, over a wireless connection, a list of connect points to the client, each connect point defining an address at one of the gateways;

[0020] the client sending, over a wireless connection, a phase two connection request to at least one of the connect points; and

[0021] the gateway that is addressed by the phase two connection request responding to the phase two connection request, and the client establishing a connection with the gateway specified by said connect point.

[0022] The two phase connect allows reconfiguring of gateways, load balancing of gateways and allows multiple gateways, which means the system can both scale, be highly available and support failover functionality. The gateways may also adapt or convert messages as specified in the US patents initially cited.

[0023] In a preferred variant of the method, the list of connect points is an ordered list and is created by a load balancing unit in response to a first connection request. The order of the list is determined as to balance the load between

gateways. Two phase connect method allows the load balancing algorithm to reside centrally, within the gateways. The load balancing being central, it is e.g. easy to reconfigure the load balancing strategy without having to update client configuration or client software. Different strategy patterns and algorithms for the load balancing may be used, e.g. based on client connections, on gateway memory or on CPU load etc. The list of connect points passed to the clients is ordered so that the preferred gateway (currently having the lowest load) is first in the list, etc. Connect points to gateways that are not operational are excluded from the list of connect points.

[0024] The load balancing mechanism is preferably implemented on the gateways. Each gateway—actually, a servlet running on a container embedded with the gateway process—responding to a phase one connection request knows about the load status of the other gateways and thus can determine the order of connect points according to the current load balancing strategy. In order for each gateway to know the load status of the other gateways, the gateways periodically send or publish heartbeat messages to or between one another over the wirebound network. These messages comprise load information from the sending gateway.

[0025] In a preferred embodiment of the invention, the messaging system further comprises a common data store configured to store a session context, which session context comprises the status of a session established between the core messaging system and one of the clients, a state of the client connection associated with the session, and a session identifier. The method comprises the steps of

[0026] a first gateway maintaining a session;

[0027] the first gateway updating and maintaining, for said session, an associated session context in the common data store;

[0028] the first gateway no longer maintaining said session and no longer updating and maintaining said session context;

[0029] a second gateway retrieving, for said session, the associated session context from the common data store;

[0030] the second gateway maintaining said session.

[0031] This allows the separate gateways to maintain a consistent view of connections and sessions, and to handle incoming connection requests issued by the clients correctly.

[0032] In a further preferred embodiment of the invention, the method comprises the steps of

[0033] a gateway receiving a phase two connection request from a client, the connection request comprising a session identifier;

[0034] the gateway initiating a lock operation on the session context in the data store associated with said session identifier;

[0035] if the lock operation succeeds, the gateway locking the session context, thus preventing other gateways from reading or writing the session context;

[0036] if the lock operation does not succeed, the gateway informing the client that the connection request is rejected.

- [0037] if no session with said session identifier exists, the gateway creating in the data store, a session context with said session identifier, the gateway updating the entry to reflect an open state and the session owner to have the value of the gateway's identifier;
- [0038] if a session with said identifier already exists, the gateway reading the registered session state from the store entry;
- [0039] if the read session state is closed, the gateway updating the session entry to reflect an open state and the session owner to have the value of the gateway's identifier;
- [0040] if the read session state is open, the gateway also reading the registered session owner from the store entry, the gateway examining the lock of the registered session owner;
- [0041] if the lock of the registered owner is currently held; the gateway informing the client that the connection request is rejected.
- [0042] if the lock of the registered owner is currently not held; the gateway updating the session entry to reflect an open state and the session owner to have the value of the gateway's identifier;
- [0043] after reading and possibly updating the session entry the gateway releasing the held lock on the session context
- [0044] after unlocking the session context the gateway informing the client that the connection request is accepted or rejected.
- [0045] This is called "locking in the store", which is shared between all gateways. This on one hand serves the same purpose as the "in memory locking" explained below, but between gateways. For this purpose event notification is propagated to end points in different gateways. On the other hand, this locking prevents one gateway from 'stealing or replicating' a session which is currently active at another cluster member, since session IDs are unique within a cluster. One gateway member may want to make updates to the session store even if the session has terminated from the clients view. This shows an asynchronous view of the connection status due to wireless network latency.
- [0046] In a further preferred embodiment of the invention, the method comprises the steps of
- [0047] a gateway, after receiving a phase two connection request, determining that another gateway, which according to the session context owns the session identified by the connection request, is not operational;
- [0048] said first gateway breaking the lock and taking ownership of the session.
- [0049] This allows a robust handling of the situation in which a gateway has terminated abnormally, but a corresponding session is still registered as being owned by that terminated gateway.
- [0050] In a further preferred embodiment of the invention, the method comprises the steps of
- [0051] a gateway receiving a phase two connection request from a client, the connection request comprising a session identifier;
- [0052] if no session with said session identifier exists, the gateway creating in a local memory section a session context with said session identifier;
- [0053] the gateway initiating a lock operation on the session context in the local memory associated with said session identifier;
- [0054] if the lock operation succeeds, the gateway locking the session context, thus preventing other acceptor threads or concurrent connect requests from accessing the session context;
- [0055] if the lock operation does not succeed, the gateway informing the client that the connection request is rejected.
- [0056] This is referred to as "in memory locking". The in memory locking guarantees session data base integrity in the face of multiple mobile client connect requests originating from the same mobile device propagated across a wireless network which arrive concurrently at a single gateway. This locking prevents two concurrent session creation operations from corrupting the gateways internals. Locking is also necessary where other operations may also occur simultaneously such as session creation and session deletion concurrently contending for the same session context. This is effective and indeed necessary for efficient operation since on a wireless network propagation of packets can lead to significantly divergent event notification from the view of two end points.
- [0057] Both "locking in the store" and "in memory locking" may be implemented independently from the two phase connection procedure, that is, also a messaging system that uses e.g. a single phase or any other connection procedure.
- [0058] The computer program for operating a wireless messaging system is loadable and executable on a data processing unit in a gateway computer, the messaging system being configured to comprise a plurality of mobile wireless clients a core messaging system and at least two gateway processes or gateways for providing message forwarding from the clients to the core messaging server and vice versa. The computer program, when being executed on one of the gateway computers, performs the steps of the method according to the invention. In a preferred embodiment of the invention, the computer program product comprises a computer readable medium, having the computer program code means recorded thereon.
- [0059] Further preferred embodiments are evident from the dependent patent claims. Features of the method claims may be combined with features of the device claims and vice versa.

#### BRIEF DESCRIPTION OF THE DRAWINGS

[0060] The subject matter of the invention will be explained in more detail in the following text with reference to preferred exemplary embodiments, which are illustrated in the attached drawings, which schematically show:

[0061] FIG. 1 a basic system structure;

[0062] FIG. 2 a system structure with a gateway cluster;

[0063] FIG. 3 an internal structure of a gateway; and

[0064] FIG. 4 a flow diagram of a program for operating the system.

[0065] The reference symbols used in the drawings, and their meanings, are listed in summary form in the list of reference symbols. In principle, identical parts are provided with the same reference symbols in the figures.

#### DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS

[0066] FIG. 1 shows a basic system structure: Mobile clients 1, 1', 1" such as computers, cell phones, personal digital assistants, automatic field devices for measurement and control etc. each comprise a client application 2 and a client library 3. The client 1 is a computer application that does the actual work the user or owner of the client 1 is interested in. The client library 3 provides software functions with an API that allows the client application 2 to communicate over a wireless network by means of a higher level protocol such as Softwired's WJMS (Wireless Java Messaging Service) or similar. The client library 3 may establish a wireless connection 4 to a gateway 5 according to said protocol, based on a lower level protocol, in particular a bearer layer of the wireless network. The gateway 5, as well as a number of other messaging clients 8, 8', 8", are connected to a core messaging server 7 by a higher level protocol with e.g. JMS connections 6, 6', 6" running over a wirebound network.

[0067] The gateway 5 acts as a proxy on behalf of the client 1, maintaining a JMS connection to the core messaging server 7 during normal operation and also in the case when the wireless connection 4 is interrupted and there is no WJMS connection with the client 1.

[0068] Whereas henceforth reference shall be made mainly to a Java-based messaging system, it is understood that the concepts presented are applicable to other non Java-based message queuing technologies, as applied e.g. in IBM MQ Series and Microsoft Message Queuing—MSMQ.

[0069] FIG. 2 shows a system structure with a gateway cluster with a plurality of gateways 5, 5', 5" instead of a single gateway 5. Each of the clients 1, 1', 1" may connect to any one of the gateways 5, 5', 5", provided that the respective gateway 5, 5', 5" is in operational condition. Each of the gateways 5, 5', 5" may connect to a persistent data store 9 that is common to the gateways 5, 5', 5", allowing them to maintain and to share a consistent representation of the status of connections, relocating clients, and restart of a gateway 5 when necessary, using the information in the data store 9. The data store 9 implemented by a file system or a database system at one of the gateways 5, 5', 5", at the messaging server 7 or at a further, separate location, and is accessible to the gateways 5, 5', 5" over further wirebound connections over the wired network. The data store 9 stores, for the connections and sessions established by the gateways 5, 5', 5", a session store or session context 12 and provides a locking mechanism for controlling access to the session entry or locking context 12. The session entry contains a representation of (JMS) connection state resulting from mobile client application use of the client library 3 (WJMS API), and reflected by the gateway 5 on the corresponding (JMS) client connection.

[0070] The gateways 5, 5', 5" are further configured to inform one another about their operational status, or to query

one another about the status directly, or to store and retrieve such status information in the data store 9.

[0071] FIG. 3 shows an internal structure of a gateway 5. The gateway 5, depending on the context, is either a gateway computer, or only a gateway process, which may in turn comprise a plurality of process threads and sub-processes, running on a gateway computer. The different gateways 5, 5', 5" typically run on physically or geographically separated gateway computers. A gateway 5 comprises a servlet 10 for responding to phase one connection requests, acceptor threads 15 that execute a response logic for responding to phase two connection requests, and one client session manager 14 for managing mobile client sessions of the gateway service.

[0072] The client session manager 14 checks and updates session information stored for each of the sessions in a separate session store or session context 11. For example, the client session manager 14 may create, lock, look-up, unlock or delete sessions. The created sessions create producers, publish messages, remove producers, close connections etc. as defined by the JMS API. When modifying a session in this manner, the client session manager 14 updates a session context 11. During the update, access to the session context 11 is locked, as shall be explained later on. The session context 11 is stored in a local memory section 13 of the gateway 5. In a preferred embodiment of the invention, the session context 11 is mirrored in the data store 9 and represented there as a mirrored session entry 12 as well.

[0073] Further terms and concepts used are summarized in the following Glossary:

[0074] Wireless Network(s): This includes any wireless network, which carries data such as GSM, GPRS, EDGE, UMTS, Mobitex, and new generation networks such as '4G' wireless broadband.

[0075] Client: An application or object that produces and consumes messages, where messages are objects that contain the data being transferred between clients.

[0076] Sliding window: A set of unacknowledged messages between two peers. For each message in the window the peers agree, by exchanging acknowledgments, which messages have been received, and which are still pending an acknowledgment. Messages that are known to have been received by the peer are deleted from the window. At any time a peer may ask for a rebroadcast of any message in the window. Messages in the sliding window are numbered, and the peers agree as to the next sequence numbers to be sent and acknowledged. The sliding window preferably is implemented as a list of references or pointers to messages.

[0077] Access point: An HTTP URL at which a clustered gateway locator Servlet can be reached during phase 1 of connecting to a gateway cluster.

[0078] Connect point: A location at which a clustered gateway can be contacted, and a WJMS connection can be established during phase 2 of connecting. It specifies: IP Address (numeric or alphanumeric), port number, and bearer protocol that is supported.

[0079] A WJMS Connection exists between an application running on a programmable wireless device and a running instance of the gateway. End points in the gateway and the client library accept WJMS messages for transmission to the

(single) corresponding end point in the client library **3** or gateway **5**. Such transmissions occur across the bearer layer. An iBus//Mobile session defines a context within which messages (sent across a WJMS connection) are sent. The context includes:

[0080] The session identifier.

[0081] The WJMS connection.

[0082] The state of the JMS connection.

[0083] A sliding window of WJMS messages between client and gateway.

[0084] The session ID is created by the client application **2**. The client application **2** ensures that the session ID is unique. An iBus//Mobile session persists through connection loss and reestablishment, caused by the unreliable nature of wireless networks.

[0085] Given the state of the session and the set of unacknowledged messages represented by the sliding window, the client library can reestablish connection with a gateway after a loss of connection or after a changeover from another gateway.

[0086] The following sections give an overview of the operation of the mobile gateway clustering system and method, henceforth called "iBus Mobile gateway". Clustering functionality enables separate iBus Mobile gateway processes to cooperatively provide the gateway service between JMS wire-line networks and wireless networked devices. The clustered release scales up the existing single gateway version of iBus Mobile such that multiple gateways cooperate to service mobile client messaging needs.

What Functionality Implements Gateway Clustering?

[0087] 1. All clustered gateways are aware of all other gateways in the same cluster. Each gateway reads the cluster membership from its configuration file.

[0088] 2. An HTTP Servlet (the cluster connect servlet) running in each Gateway can answer client requests for the IPAddresses of cluster members.

[0089] 3. Mobile clients connect to the cluster in a two phase way. First, requesting the list of IPAddresses where cluster members can be found, and then sequentially trying each member in the list for the proper WJMS connection.

[0090] 4. This documentation distinguishes between access point URLs and connect point URLs. Access points are the locations at which cluster connect servlets for the cluster can be contacted. Connect points are returned in the lists of connect URLs that clients receive from cluster connect servlets.

[0091] 5. All gateways share the same session store so any gateway can respond to any client request for a WJMS connection. Note that the session DB store must be located on an NFS or similar file system location or a shared database. All gateways will read and write to that location.

[0092] 6. Connection establishment always follows the two phases. First get the cluster membership list and then try each member in turn. Regardless of four connection scenarios;

[0093] a) Initial client connection.

[0094] b) Normal client reconnect (after orderly client shutdown).

[0095] c) Session re-establishment after connection loss.

[0096] d) Fail-over scenario where one (or more) gateway(s) fail.

[0097] 7. The clients will try members in the order of the connect points list received during phase one. So the algorithm of where to direct client load is centrally controlled (within the cluster).

[0098] 8. Gateway Clustering supports load balancing. Various load-balancing strategies are configurable such as load balance according to client connections, cpu load on the gateway processor, memory available to the gateway process, etc.

[0099] 9. Sharing the session store location among all cluster members reduces client session re-establishment to an operation similar to re-establishment in the stand-alone gateway case; the cluster member to which, the client reconnects simply reloads the session from store.

[0100] 10. A client session must be active at one gateway only. Cluster members share information about which session is active on, which Gateway, through the store. Each client session has one owner gateway, while the session is active at the owner other gateways are locked out of servicing connection requests for the session.

[0101] 11. Gateway Failover—gateways monitor each other's presence (via the store), such that if one gateway fails then it's store will no longer be active. Other members are then no longer locked out of handling the sessions owned by the failed gateway, client sessions will fail-over to other gateways.

Components of a Gateway Cluster

1. The Gateways

[0102] By definition an iBus//Mobile cluster includes multiple gateway processes. Any gateway that is a member of the cluster is able to accept a WJMS connection request from any mobile device. Cluster membership is configured into the configuration files of each individual gateway.

1.A. The Gateway Locator

[0103] A gateway component called the 'clusteredGatewayLocator' must be enabled in the gateway config. The class that is usually entered as the clustered gateway locator is: `ch.softwired.gateway.clustering.SimpleGatewayLocator`.

[0104] The SimpleGatewayLocator reads the cluster membership from the configuration file of the local gateway process. The locator will not change, or adjust the cluster membership once it is instantiated, so reconfiguring the cluster requires restarting cluster members. For every member the SimpleGatewayLocator needs the 'ipaddress', 'port' and the supported network protocol (bearer) at which that gateway is accepting WJMS connections, to be in the local gateway config file. A SimpleGatewayLocator can be disabled in the configuration file by setting its 'enabled' property to 'false', this has the effect of preventing the local gateway cluster connection Servlet from answering phase

one client connect requests. Note that if the locator is disabled but other gateways are configured to include the local gateway then the gateway is still a cluster member, it addresses phase two connection requests and not phase one requests.

#### 1.B. The Cluster Store

[0105] All gateways in a cluster share one session database, so their stores must all be physically located at the same place, and that file system must be accessible to every cluster member, gateway stores have the following functionality that is specific to clustering:

[0106] 1. The store maintains an 'active store instance' lock in the store location. Each store periodically updates its lock.

[0107] 2. All cluster members monitor the store locks of all other Gateways to ascertain that other gateways have not failed.

[0108] 3. Any session that is connected to the cluster is owned by one gateway (at any point in time).

[0109] 4. Ownership is registered in the session entry in the data store.

[0110] 5. Any gateway that wants to read or write the data for a session must first acquire the lock for that session.

[0111] 6. Usually the session lock is held by the gateway to which the session is currently connected (the owner).

[0112] 7. If a client disconnects and then reconnects to the cluster, the gateway to which it reconnects will have to first become the session owner.

[0113] 8. If the session lock is still held by the previous owner (which is possible if the disconnect was not 'orderly' and previous owner has not yet timed out the session, or if the previous owner has failed and its active store lock has not yet timed out) then the reconnect gateway will reject the connection attempt, from the cluster point of view the session is still running.

#### 2. The JMS Domain

[0114] All gateway processes in one cluster must connect to the same JMS messaging domain. Note that iBus//Mobile clustering addresses gateway scaling only, not JMS scaling. Additional mobile client connections can be serviced by deploying, and configuring additional gateways. The requirement that all cluster members connect to the same JMS domain implies that the JMS service must be capable of scaling to service the messaging load of all clients of all gateways.

#### 3. Cluster Connection Servlets (Access Points)

[0115] Usually each gateway process includes a Servlet which responds to phase 1 connect requests. Phase 1 requests are HTTP requests made by the clustering enabled QOS to each Servlet, in sequence, until one access point returns a list of cluster connect points. Each time a client requests a WJMS connect point in the cluster the Servlet obtains a list of cluster members from the local clustered gateway locator.

[0116] As each cluster member deploys one (or possibly more) Servlets, clients can request connect points from a number of access point URLs. The precise http URL at

which a gateway's cluster connection servlet is deployed can vary from one installation to another, provided the client applications know the URLs for the access points of the cluster. Obviously the cluster connect points that the client receives in phase 1 must support the bearer protocol that the client wants to use (TCP, HTTP, UDP etc). Hence the client request must include the bearer in the query string of the http servlet request.

[0117] An example of a URL that addresses a cluster connect Servlet is:

[0118] `http://192.168.168.10:8080/ibus/cluster/Simple-GatewayLocator/Servlet?protocol=tcp`

#### 4. Client Application QOS

[0119] For a mobile client to do the two phase cluster connection a CLUSTER\_CONNECT protocol object is needed on the iBus protocol stack. The CLUSTER\_CONNECT object is responsible for connecting to a cluster access point (Servlet via HTTP) to obtain the gateway cluster connect points. Once a connect point list is obtained CLUSTER\_CONNECT will pass each connect point (ipaddress and port) down the stack to the CONNECT object until one of these phase 2 WJMS connections succeeds, and the client connection to the cluster is completed. Note that the application can pass the URLs of multiple access point Servlets to CLUSTER\_CONNECT, in case any individual servlet is down.

[0120] Phase 1 of connecting to a cluster must indicate the connection protocol (TCP, UDP, HTTP) that the WJMS connections (to be established in phase 2) will use. The client application passes the list of access points to the CLUSTER\_CONNECT object as a property on the QOS parameters: CLUSTER\_SERVLET\_URL, the presence of this property indicates to the iBus//Mobile client library that the client is connecting to a gateway cluster. The format of the access points list specified by CLUSTER\_CONNECT\_URL is:

[0121] `http://hostname/servletpath?protocol=(tcp|http|udp) { |http:// hostname/servletpath?protocol=(tcp|http|udp) . . . }`

[0122] In the following, a step by step description of a connection setup is given: FIG. 4 shows a flow diagram of a program for operating the system accordingly. The process steps are labeled with circled numbers 21-26 and are explained in more detail in the following. Phase one steps are encircled by an oval 30, phase two steps are encircled by a further oval 31. Shaded box 32 denotes mutual exclusion across acceptors and clustered gateways.

1. Client Application Initiates a Connect Request to iBus//Mobile Client Library (Label 21).

[0123] The request specifies:

[0124] 1. The session ID of the iBus//Mobile client sessions.

[0125] 2. The bearer type that the client wants to use.

[0126] 3. A list of access points at which clustered gateway locator Servlets should be running.

[0127] 2. Client Library Requests a Set of Connect Points (Label 22).

[0128] (Phase 1 requests)

[0129] This phase one request is made across HTTP.

[0130] The client library will try each access point in order until a reply is received.

[0131] 3. Clustered Gateway Locator Servlet Responds (Label 23).

[0132] (Phase 1 reply.)

[0133] On receiving a phase 1 request the contacted locator determines which cluster members the client is best directed to.

[0134] The factors influencing the 'best member' consideration:

[0135] 1. The load that individual members are currently reporting.

[0136] 2. The bearer that the client application has specified.

[0137] The locator builds the list of connect points and replies.

[0138] 4. Client Library Requests WJMS Connection from Connect Points (Label 24).

[0139] (Phase 2 requests.)

[0140] On receiving the phase 1 reply the client library tries each received cluster member (at the provided connect point) in turn. This phase two request is made using the bearer protocol specified by the connect point. The client library will try each connect point in order until a connection is established. Each connect point is tried more than once before requests are issued to the next.

5. Clustered Gateway Processes a Connect Request (Label 25).

[0141] On receiving a connection request the gateway locks the session context in the session store (label 25a). If this lock operation fails then so does the connect request.

[0142] The session context is locked regardless of session existence, i.e. no distinction between initial connect requests and reconnect requests. While the session (or session context) is locked no other gateway, (or acceptor thread in the same gateway) can read or write the session state in the store.

[0143] The accepting gateway checks the session state, a closed state indicates that the session is not active anywhere in the cluster. The accepting gateway is free to create (or open) the session.

[0144] An open session state indicates that the session is active on the cluster.

[0145] The accepting gateway reads the session owner, i.e. an identifier specifying a gateway process, that is a cluster member, from the session store.

[0146] The accepting gateway reads the run status of the session owner.

[0147] If the registered owner is not running then the accepting gateway is free to create (or open) the session.

[0148] If the registered owner is running then the connect request fails. The session is already (still) running.

[0149] After confirming that the session is not already active in the cluster the accepting gateway registers itself as the owner of the open session. After processing the connect request the accepting gateway unlocks the session in the store (label 25d).

[0150] The accepting gateway replies to the connect request, either connection accepted, or rejected. (Phase 2 reply, label 25.5)

[0151] In each gateway, multiple acceptor threads handle incoming requests for WJMS connections, i.e. phase two connection requests. This allows multiple clients to connect to each gateway simultaneously, which is important during a standalone gateway (or whole cluster) restart. So client connection establishment is multithreaded both within a gateway process, and between cluster members. An acceptor thread calls a client session manager for handling the session-related parts of the method.

[0152] As multiple member gateways may receive connection requests from the same client device simultaneously they will all contend for the same session information in the store.

[0153] There follows a brief summary of the concepts involved in the locking mechanism:

[0154] Session context: the context within which a WJMS message (and the corresponding JMS message) is sent. This context is altered by client operations versus the WJMS API, for instance create a JMS session on a JMS connection, or create a MessageProducer on an existing Destination within an existing JMS session, now use the MessageProducer to send a message.

[0155] Session identifier: A system wide unique string of characters that identifies the client application and the session context.

[0156] Session entry: Exists in the session store. Is the location in the store at which the session context (of the corresponding session identifier) is stored. If a session does not exist, i.e. the client app. with the corresponding session id has never yet connected then the entry does not exist either.

[0157] Session lock context: To avoid two threads (or Gateways) simultaneously updating the session context all operations on the store are single threaded wrt. the session id. The session lock context is the locking object on which the lock mechanism operates. The session lock context can exist regardless of session existence. The lock operation (on the session lock context) must be executed first regardless of session existence.

6. Client Library Returns to the Application (Label 26).

[0158] On receiving a (connection accepted) reply to the connect request the client library passes a positive return value "Connection established" to the client application.

[0159] On receiving a (connection rejected) reply the client library continues processing connect points (from the current access point) and other access points (if all current

connect points failed). On exhausting all connect points received from all access points the client library passes a negative return value "Could not establish a connection" to the client application.

[0160] It may happen that a Gateway to which the session is connecting discovers that another Gateway, which claims to own the session is not running, i.e. has not updated the lock for a certain time. This may have been triggered by the client reconnecting elsewhere because the client no longer receives heartbeats from a former connection with the other Gateway. In such a situation, the gateway is allowed to break the locking of the session context. This is commonly called stealing a session.

[0161] As a result of this connection mechanism, a client can be active on only one gateway at any point in time; otherwise the session is "still running" elsewhere in the cluster. All client connect operations, reconnects and failovers are handled according to the same procedure. Failure of one gateway will cause the affected clients to switch over to another gateway automatically.

[0162] While the invention has been described in present preferred embodiments of the invention, it is distinctly understood that the invention is not limited thereto, but may be otherwise variously embodied and practiced within the scope of the claims.

#### LIST OF DESIGNATIONS

- [0163] 1 client
- [0164] 2 client application
- [0165] 3 client library
- [0166] 4 wireless connection
- [0167] 5 gateway
- [0168] 6 wirebound connection
- [0169] 7 messaging server
- [0170] 8 messaging client
- [0171] 9 data store
- [0172] 10 servlet
- [0173] 11 session context
- [0174] 12 mirrored context in the data store
- [0175] 13 local memory section
- [0176] 14 client session manager
- [0177] 15 acceptor threads

1. A method for operating a wireless messaging system, the messaging system being configured to comprise a plurality of mobile wireless clients, a core messaging system and at least two gateway processes or gateways, the gateways acting as proxies on behalf of the clients and in the course of normal operation of the messaging system maintaining sessions by performing the steps of,

the gateways receiving inbound messages from the clients over a wireless bearer using a wireless protocol and forwarding said inbound messages to the core messaging system over a connection using a wirebound protocol;

the gateways receiving outbound messages from the core messaging system over a connection using a wirebound protocol and forwarding said outbound messages to the clients over connections using a wireless protocol; the method comprising, in order for a client to establish a connection with one of the gateways, a two phase connect procedure with the steps of,

the client maintaining a list of access points, each access point defining an address at one of the gateways;

the client sending, over a wireless connection, a phase one connection request to one of the access points;

the gateway that is addressed by the phase one connection request, sending, in response to the phase one connection request, over a wireless bearer, a list of connect points to the client, each connect point defining an address at one of the gateways;

the client sending, over a wireless bearer, a phase two connection request to at least one of the connect points; and

the gateway that is addressed by the phase two connection request responding to the phase two connection request, and the client establishing a connection with the gateway specified by said connect point.

2. The method of claim 1, further comprising the step of

the client, if a phase one connection request is unsuccessful, sending a further phase one connection request to a further access point selected from the list of access points.

3. The method of claim 1, further comprising the step of

the list of connect points being an ordered list, and the client, if the preceding phase two connection request is unsuccessful, sending a further phase two connection request to the next connect point in the current list of connect points.

4. The method of claim 3, further comprising the step of

a load balancing unit creating the list of connect points in response to a first connection request, the order of the list being determined as to balance the load between gateways.

5. The method of claim 4, further comprising the step of

the load balancing unit determining a load parameter characterizing, for each of the available gateways, a current load, and the load balancing unit ordering the list of connect points in ascending order according to the load parameter of the gateways associated with the connect points.

6. The method of claim 1, wherein the access points are defined as a HTTP request.

7. The method of claim 1, wherein the connect points are defined by a protocol, an address, and a port to be used.

8. The method of claim 1, the messaging system further comprising a common data store configured to store a session context, which session context comprises the status of a session established between the core messaging system and one of the clients, a state of the client connection associated with the session, and a session identifier, the method comprising the steps of

a first gateway maintaining a session;  
 the first gateway updating and maintaining, for said session, an associated session context in the common data store;  
 the first gateway no longer maintaining said session;  
 a second gateway retrieving, for said session, the associated session context from the common data store;  
 the second gateway maintaining said session.

**9.** The method of claim 8, comprising the steps of

a gateway receiving a phase two connection request from a client, the connection request comprising a session identifier;

the gateway initiating a lock operation on the session context in the data store associated with said session identifier;

if the lock operation succeeds, the gateway locking the session context, thus preventing other gateways from accessing the session context;

if the lock operation does not succeed, the gateway informing the client that the connection request is rejected;

if no session with said session identifier exists, the gateway creating in the data store, an entry for the session context with said session identifier.

**10.** The method of claim 9, comprising the steps of

a gateway, after receiving a phase two connection request, determining that another gateway, which according to the session context owns the session identified by the connection request, is not operational; said first gateway breaking the lock and taking ownership of the session.

**11.** The method of claim 1, comprising the steps of

a gateway receiving a phase two connection request from a client, the connection request comprising a session identifier;

the gateway initiating a lock operation on the session context in the local memory associated with said session identifier;

if the lock operation succeeds, the gateway locking the session context, thus preventing other acceptor threads triggered by concurrent connect requests for the same session identifier from accessing the session context;

if the lock operation does not succeed, the gateway informing the client that the connection request is rejected;

if no session with said session identifier exists, the gateway creating in a local memory section a session context with said session identifier.

**12.** A wireless messaging system, the messaging system being configured to comprise a plurality of mobile wireless

clients, a core messaging system and at least two gateway processes or gateways, the gateways acting as proxies on behalf of the clients and comprising a client session manager for, in the course of normal operation of the messaging system, maintaining client sessions by

receiving inbound messages from the clients over wireless bearers using a wireless protocol and forwarding said inbound messages to the core messaging system over a connection using a wirebound protocol;

receiving outbound messages from the core messaging system over a connection using a wirebound protocol and forwarding said outbound messages to the clients over a wireless bearer using a wireless protocol; the wireless messaging system further comprising, in order for a client to establish a connection with one of the gateways, a distributed logic for implementing a two phase connect procedure, said distributed logic comprising:

in the client, a client library providing an application program interface for, among others, accepting, access points, each access point defining an address at one of the gateways;

in the client, the client library comprising a phase one routine for sending, over a wireless bearer, a phase one connection request to one of the access points;

in the gateway that is addressed by the phase one connection request, a servlet for sending, in response to the phase one connection request, over a wireless bearer, a list of connect points to the client, each connect point defining an address at one of the gateways;

the client library also comprising a phase two routine for sending, over a wireless bearer, a phase two connection request to at least one of the connect points; and

the gateway that is addressed by the phase two connection request comprising a response logic for responding to the phase two connection request, and for establishing a connection between the client and the gateway.

**13.** The system of claim 12, further comprising a load balancing unit for creating the list of connect points in response to a first connection request, the order of the list being determined as to balance the load between gateways

**14.** A computer program for operating a wireless messaging system, which computer program is loadable and executable on a data processing unit in a gateway computer, the messaging system being configured to comprise a plurality of mobile wireless clients, a core messaging system and at least two gateway processes or gateways for providing message forwarding from the clients to the messaging server and vice versa, and which computer program, when being executed on one the gateway computers, performs the steps of the method of claim 1.

\* \* \* \* \*