



(86) Date de dépôt PCT/PCT Filing Date: 2010/03/02

(87) Date publication PCT/PCT Publication Date: 2010/09/10

(45) Date de délivrance/Issue Date: 2017/05/23

(85) Entrée phase nationale/National Entry: 2012/08/15

(86) N° demande PCT/PCT Application No.: US 2010/025943

(87) N° publication PCT/PCT Publication No.: 2010/101935

(30) Priorités/Priorities: 2009/03/02 (US61/156,758);
2009/10/07 (US61/249,493); 2010/01/19 (US61/296,270)

(51) Cl.Int./Int.Cl. *H04M 3/22* (2006.01),
H04L 12/26 (2006.01), *H04L 12/66* (2006.01),
H04M 7/12 (2006.01)

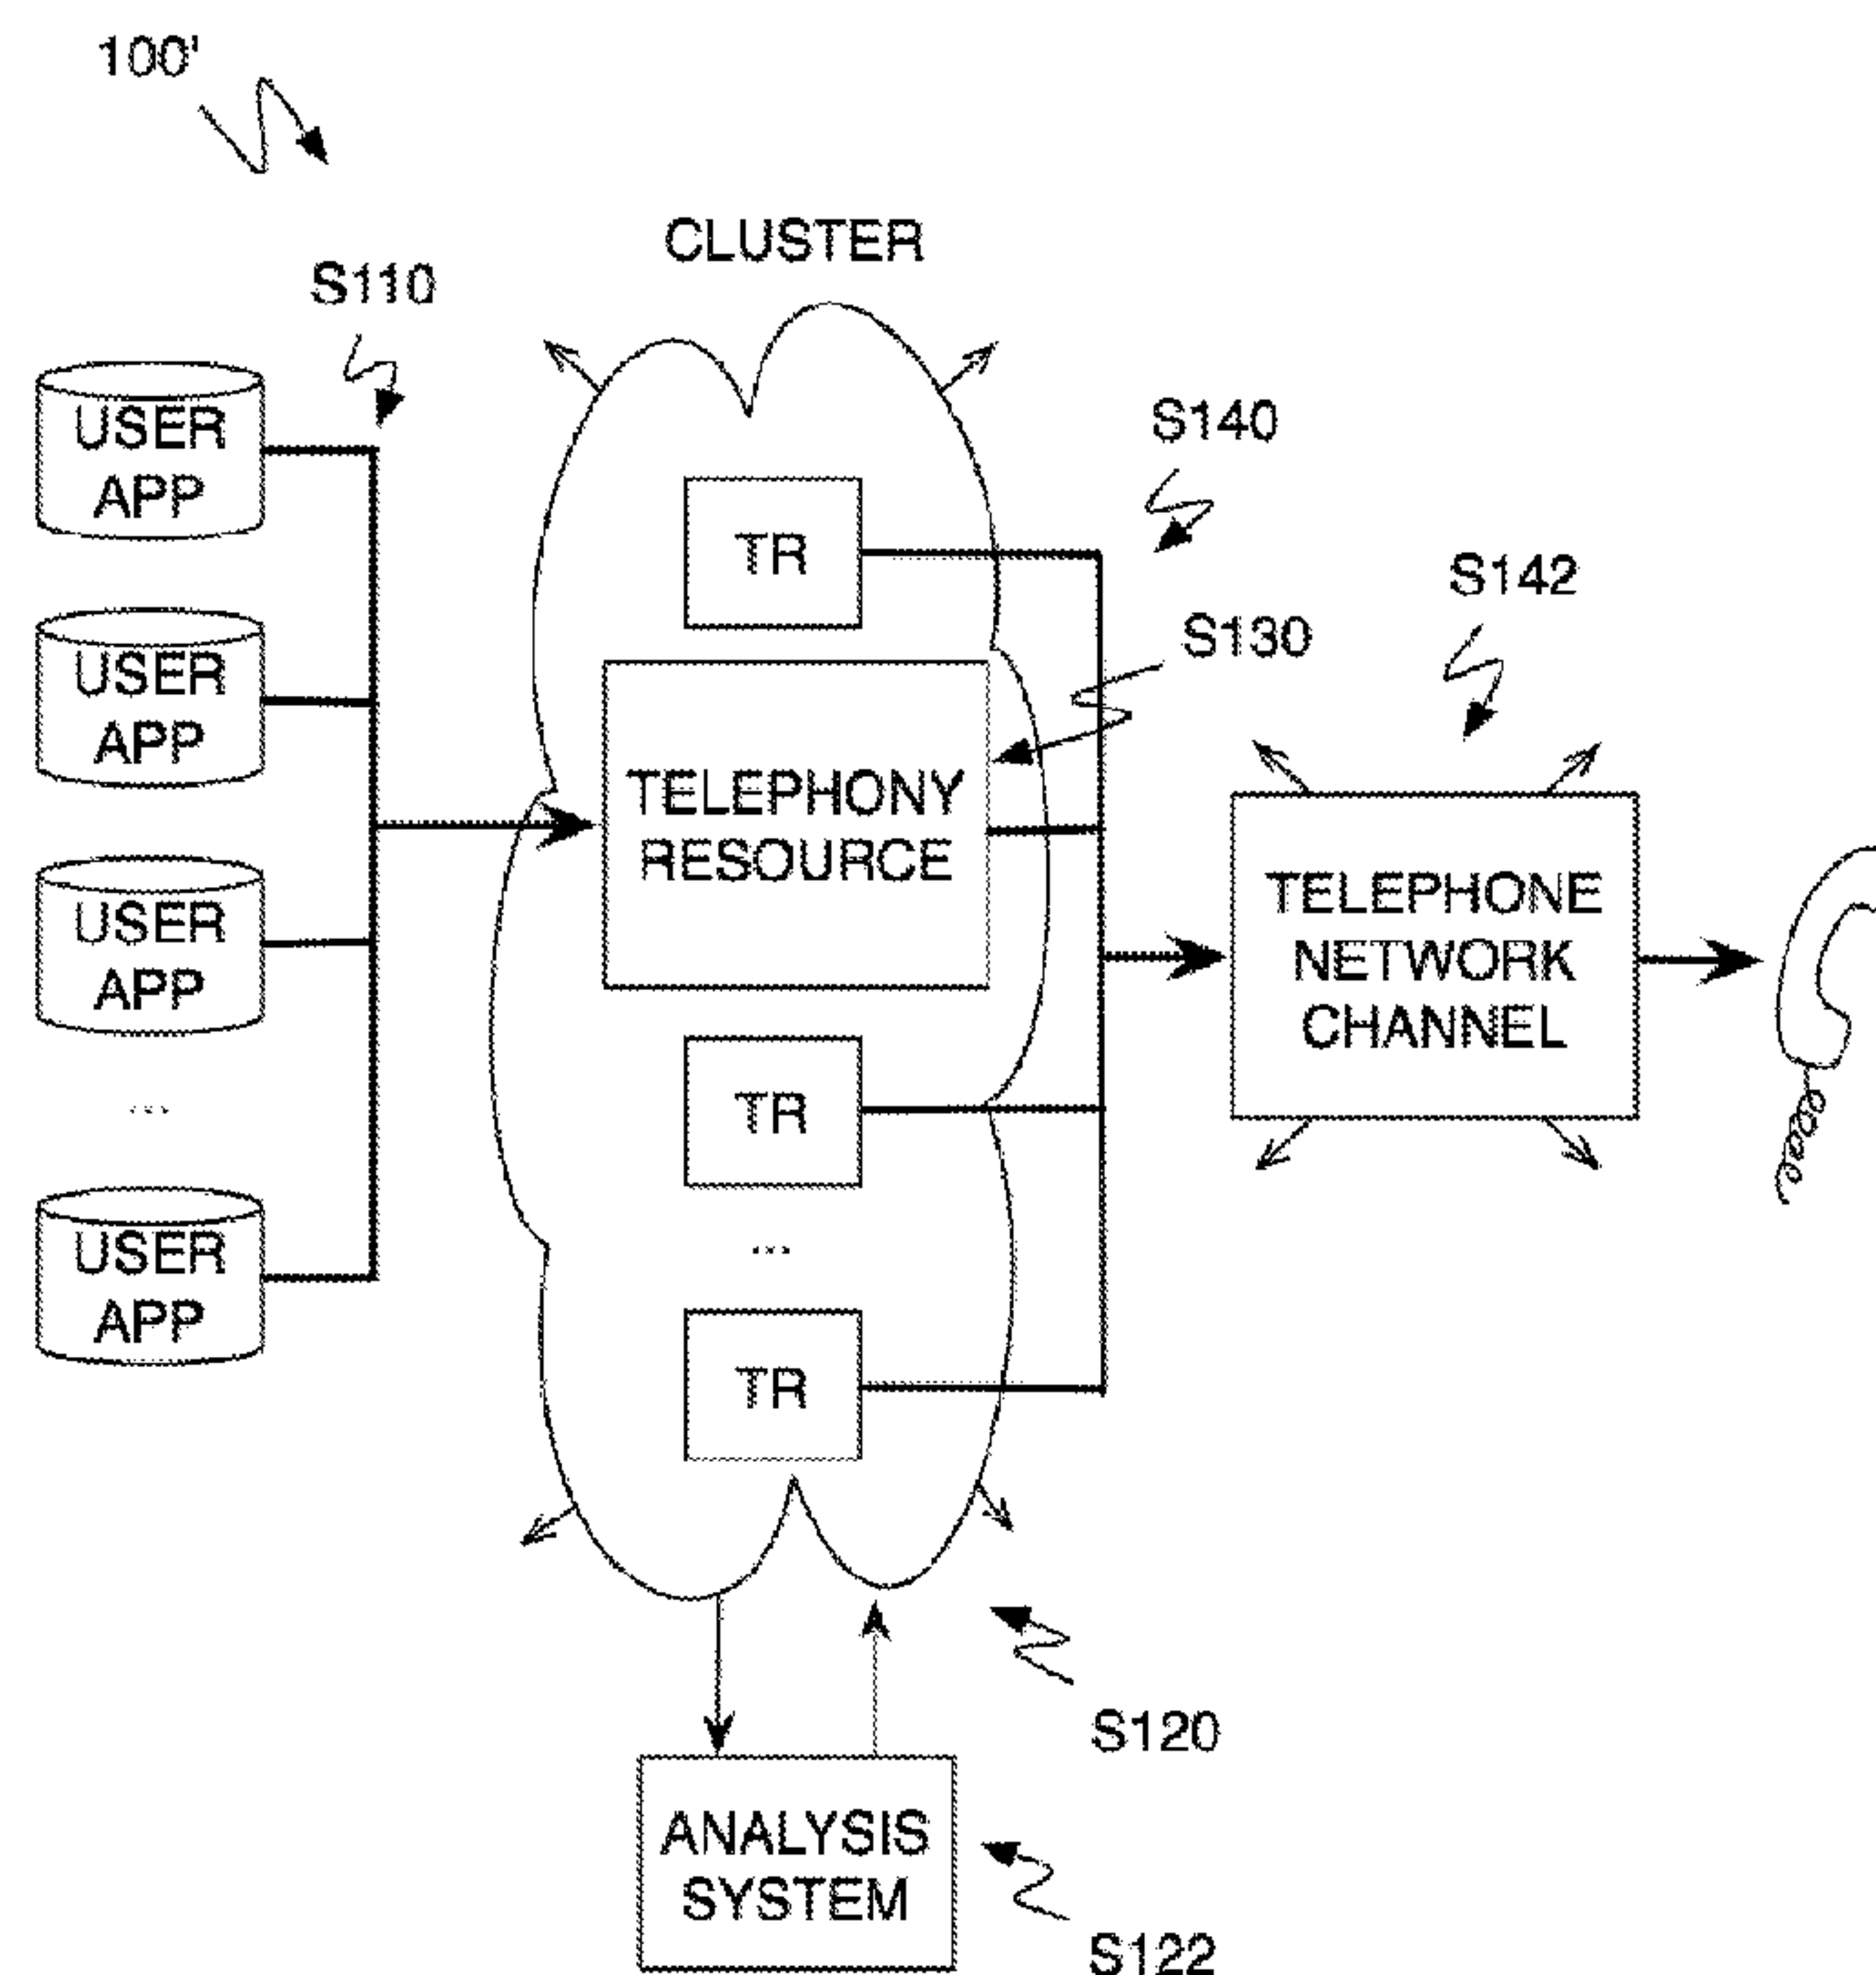
(72) Inventeurs/Inventors:
LAWSON, JEFFREY, US;
WOLTHUIS, JOHN, US;
COOKE, EVAN, US

(73) Propriétaire/Owner:
TWILIO, INC., US

(74) Agent: BERESKIN & PARR LLP/S.E.N.C.R.L.,S.R.L.

(54) Titre : PROCEDE ET SYSTEME POUR RESEAU TELEPHONIQUE PARTAGE

(54) Title: METHOD AND SYSTEM FOR A MULTITENANCY TELEPHONE NETWORK



(57) Abrégé/Abstract:

A method and system for operating a multitenant telephony system including a call queue that stores call requests received from a plurality of users; an expandable and contractible telephony resource cluster that establishes call sessions for call requests; an analysis system that calculates capacity requirements of the system; a resource allocator that manages the scaling and operation of the telephony resource cluster; and a plurality of telephony network channels that are used as telephony communication channels for call sessions.

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

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
10 September 2010 (10.09.2010)



(10) International Publication Number
WO 2010/101935 A1

(51) International Patent Classification:
H04L 12/66 (2006.01)

(21) International Application Number:
PCT/US2010/025943

(22) International Filing Date:
2 March 2010 (02.03.2010)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
61/156,758 2 March 2009 (02.03.2009) US
61/249,493 7 October 2009 (07.10.2009) US
61/296,270 19 January 2010 (19.01.2010) US

(71) Applicant (for all designated States except US):
TWILIO INC. [US/US]; Pier 38, The Embarcadero, San Francisco, CA 94107 (US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **LAWSON, Jeffrey** [US/US]; Pier 38, The Embarcadero, San Francisco, CA 94107 (US). **WOLTHUIS, John** [US/US]; Pier 38, The Embarcadero, San Francisco, CA 94107 (US). **COOKE, Evan** [US/US]; Pier 38, The Embarcadero, San Francisco, CA 94107 (US).

(74) Agent: **SCHOX, Jeffrey**; 500 3rd Street #515, San Francisco, CA 94107 (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, 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, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, RO, RS, RU, 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, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

[Continued on next page]

(54) Title: METHOD AND SYSTEM FOR A MULTITENANCY TELEPHONE NETWORK

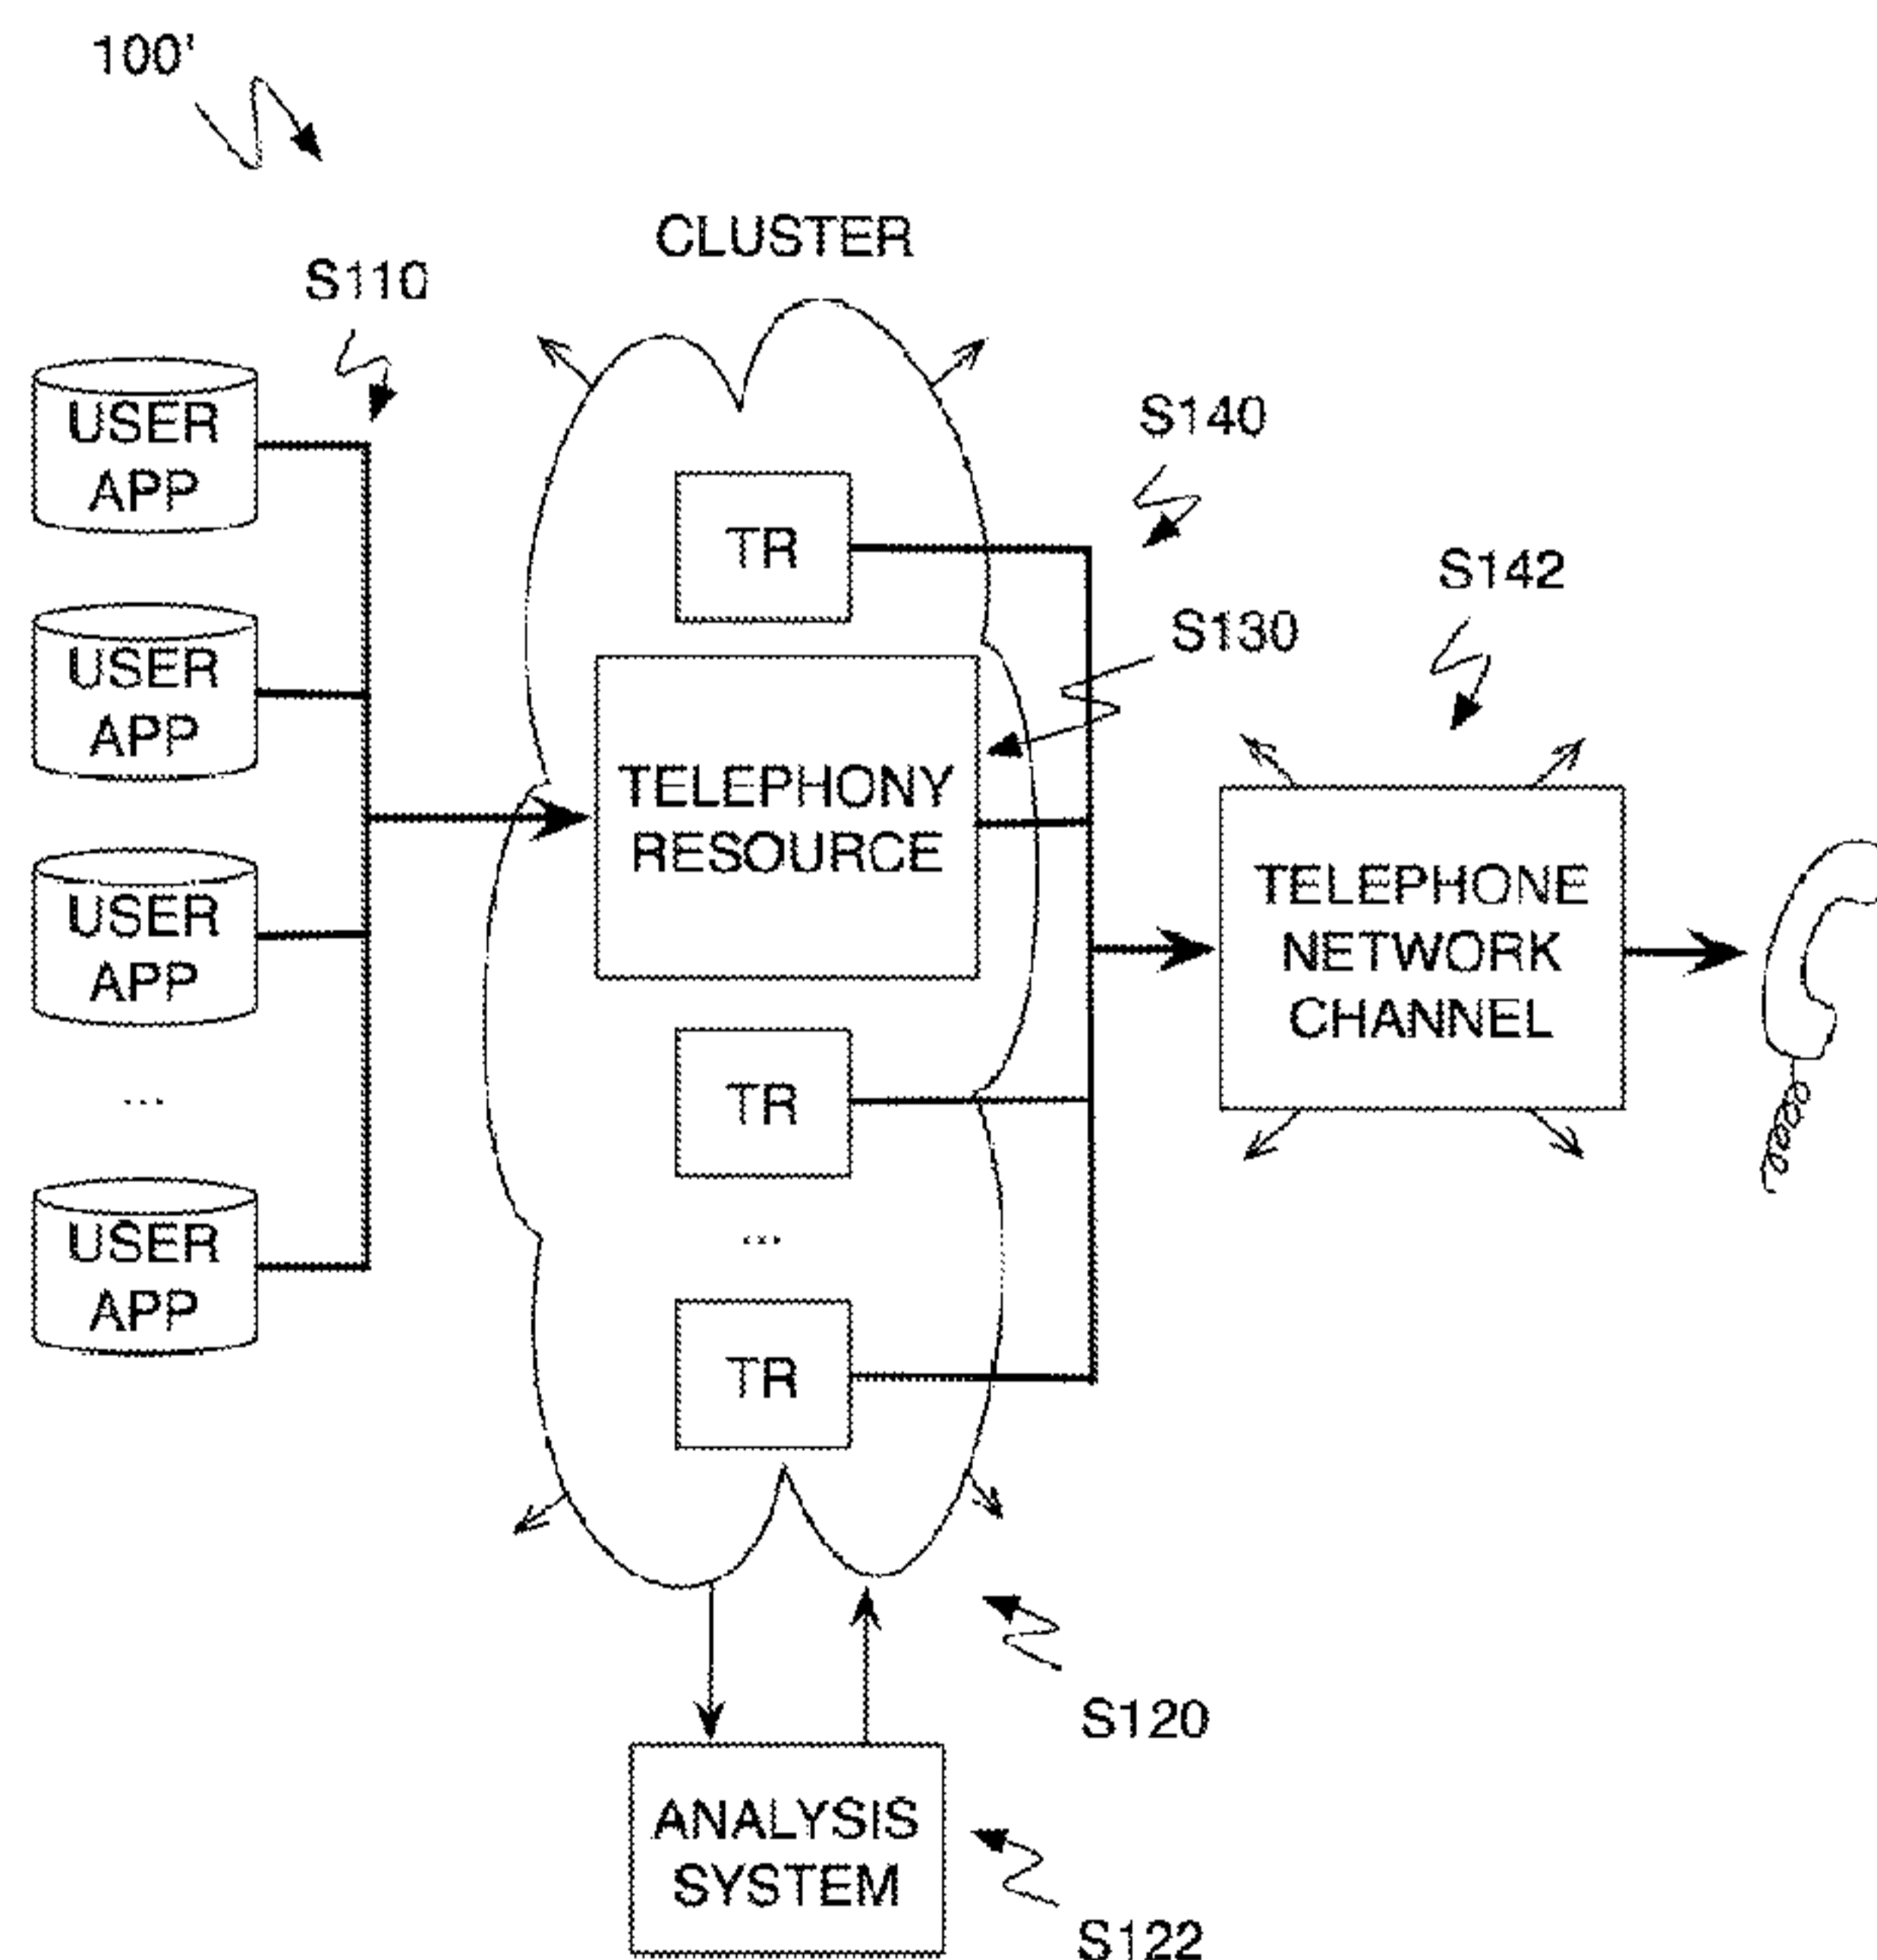


FIGURE 1

(57) Abstract: A method and system for operating a multi-tenancy telephony system including a call queue that stores call requests received from a plurality of users; an expandable and contractible telephony resource cluster that establishes call sessions for call requests; a analysis system that calculates capacity requirements of the system; a resource allocator that manages the scaling and operation of the telephony resource cluster; and a plurality of telephony network channels that are used as telephony communication channels for call sessions.

WO 2010/101935 A1 

Published:

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

METHOD AND SYSTEM FOR A MULTITENANCY TELEPHONE NETWORK

[0001]

[0002] This application is related to U.S. Patent Publication No. 2009/025,159, entitled "System and Method for Processing Telephony Sessions."

TECHNICAL FIELD

[0003] This invention relates generally to the telephony field, and more specifically to a new and useful multitenancy telephone network in the telephony field.

BACKGROUND

[0004] A telephone network has historically used a channel architecture for a telephone session or connection. This channel architecture has a foundation in the history of telephony; a physical wired connection or channel needed to be physically

connected to make a telephone call. The concept of channels is still used today. Subscribers to a telephone network are conventionally required to pay on a per channel basis. Users that wish to have a public branch exchange ("PBX"), call center, or similar telephony application typically subscribe to a service and have a fixed number of channels that are available to them and only them. As the number of channels is part of their contract, they cannot exceed that number of channels (or else the call or telephone session will fail). Since most applications only see full capacity usage on rare occasions, the user often pays for more channels than are typically used.

[0005] In contrast to the channel based architecture of the telephone network, packet based network innovations have seen a rise in recent years, such as voice over internet protocol (VOIP), internet based applications, and internet-based telephony applications. With newer technology coming to the telephony field there are unique challenges arriving for handling the hardware and software capacity demands. Dedicated hardware and software often perform tasks during a telephone call session or even act as an intermediary system for connecting a caller to an internet based application. Telephone systems generally have higher performance expectations than a website based application. While a user of a website expects a website and software to take time to load and process information, a caller experiences frustration in delays or unresponsive interactions while on the phone. Additionally, the telephony applications are still dependent on the channel based telephone system, which adds yet another barrier to scalability. The telephone network and existing telephone application software and hardware architecture limit the growing capabilities of the telephony application

field. Thus, there is a need in the telephony field to create a new and useful multitenancy telephone network. This invention provides such a new and useful system and method.

OBJECTS OF THE INVENTION

[0006] The present invention provides a system and method for providing a multitenancy telephone network for telephony applications. One objective of the present invention is to manage shared resource usage in a multi-user environment and to dynamically scale resources to satisfy capacity requirements. A related effect of this objective is that the sum total of the apparent number of resources available to each user is greater than the actual number of resources used to implement the multi-tenant telephone network. Another objective of the present invention is to efficiently use resources of a telephony platform by provisioning the processing and storage resources to satisfy capacity requirements, effectively leaving other unused resources for alternative applications, powered off for power saving, or any suitable functions. Another objective of the present invention is to make the use of a cluster of telephony resources transparent to an application of a user. This transparency is preferably preserved despite situations where operation of an application is distributed between a plurality of telephone service resources and may involve a plurality of telephone sessions on different channels. These and other objects of the invention are accomplished by the preferred embodiments of the invention, including a system for multitenancy telephone network, a method for operating a multitenant telephone

network, a method of operating a dynamic telephone network, and a method of distributing calls between telephone hardware, each described in the following sections.

BRIEF DESCRIPTION OF THE FIGURES

[0007] FIGURE 1 is a flowchart representation of a preferred embodiment for the method of operating a multitenant telephone network;

[0008] FIGURES 2-4 are schematic representations of preferred embodiments of a system for a multitenancy telephone network;

[0009] FIGURE 5 is a schematic representation of a preferred embodiment of the invention using a cluster of call transcribers;

[0010] FIGURE 6 is a flowchart of a preferred embodiment for the method for operating a dynamic telephone network;

[0011] FIGURE 7 is a flowchart of a preferred embodiment of the invention implementing a conference call; and

[0012] FIGURE 8 is a flowchart of preferred embodiment of the invention receiving an incoming call.

DESCRIPTION OF THE PREFERRED EMBODIMENTS

[0013] The following description of the preferred embodiments of the invention is not intended to limit the invention to these preferred embodiments, but rather to enable any person skilled in the art to make and use this invention.

1. System for a Multitenancy Telephone Network

[0014] As shown in FIGURES 2-4, the system 100 of the preferred embodiment includes a telephony resource cluster 110, a call queue 120, an analysis system 130, a resource allocator 140, and a plurality of telephony network channels 150. The telephony resource 110 cluster preferably includes a plurality of allocated telephony network channels 152 and/or a plurality of telephony resources 112 such as a plurality of call routers, a load balancer, and may additionally include a service application. The system functions to distribute the use of the network and system resources and dynamically adjust the system based on capacity requirements.

[0015] The telephony resource cluster 110 (or “cluster”) functions as a scalable (expandable and/or contractible) collection of resources, where at least one resource is used to create a phone call session requested by a user. The cluster 110 is preferably a collection of hardware and/or software components that can dynamically adjust to satisfy processing and/or storage requirements. The cluster 110 preferably appears as a hardware and/or software cloud to outside devices, such that management of hardware allocation and usage is handled internally by the system. In one variation shown in FIGURE 2, the telephony resource cluster 110, is preferably a plurality of telephony resources 112 which functions to provide intermediary processing tasks for a call request or call session, such as establishing a call session, converting telephony instructions into call actions, transcribing a call, or directing a call. In another variation shown in FIGURE 3, the telephony resource cluster 110 is preferably a plurality of connections to allocated telephony network channels 152, where an allocated telephony network

channel 152 is a channel of the allocated telephony network channels 152 that has been activated or designated as a channel available for a call session.

[0016] The telephony resources 112 are preferably software or hardware resources that are provisioned for a particular telephony processing task. There are preferably a plurality of telephony resources 112, and there may be a plurality of types of telephony resources that perform different dedicated tasks. A telephony resource 112 preferably includes a computer processors and/or computer storage devices. The telephony resource 112 may be a physical hardware device, virtual machine, a software program/routine, and/or any suitable combination to provide the processing and storage operations of a telephony resource 112. In some cases, a telephony resource 112 may include dedicated hardware or software. Since the telephony resources 112 share the basic functionality as either processing power or data storage, the core functionality of a telephony resource 112 may be reprovisioned such that the telephony resource 112 performs a different dedicated task. The resource allocator 140 (and more specifically the load balancer 142) preferably reprovisions telephony resources 112 to act as different parts of the resource cluster 110. For example, the cluster may include a number of text-to-speech servers and a number of call routers, but at some point in time there may be a low number of text-to-speech operations being performed and an increased number of telephony applications, and so a text-to-speech server is preferably reprovisioned as a call router. In one variation, the plurality of telephony resources 112 (i.e., the cluster 110) preferably includes a plurality of call routers 114. Additionally or alternatively, the cluster may include other hardware devices or software instances such as media

processing systems, transcription systems, text-to-speech systems, call recorders, call data storage, or any suitable hardware (physical device or virtual machine) or software. The resource allocator 140 for the cluster preferably includes a load balancer 142 that manages the distribution of processing tasks and the operation of the plurality of telephony resources 112. Additionally, the cluster may include a service application and/or a call router network that can cooperatively resolve issues that result from using a plurality of resources.

[0017] The plurality of call routers 114 functions to initiate or receive calls from telephony devices and provide telephony application related processing. Preferably, the call routers connect to an application server, which is preferably the source of the call request. The plurality of call routers 114 is preferably a dynamic number of call routers 114 that can be adjusted according to capacity requirements. As stated above, in alternative embodiments the plurality of call routers 114 may be replaced by or combined with other suitable telephony hardware or software resources such as media processing systems, transcription systems, text-to-speech systems, or other specialized hardware or software resources that are used in a telephony application. In one example, a plurality of transcription hardware or virtualized resources may be used in place of call routers for transcribing phone calls, as shown in FIGURE 5. Additionally, a call router 114 may be reprovisioned as a media processing system, transcription system, text-to-speech system, or for any suitable process, and similarly any processor may be reprovisioned to serve as a call router. The number of hardware or software resources may additionally or alternatively be allocated or deallocated so that any desired number

of resources in any suitable combination may be operated at any time. A hardware instance may be powered down, put into energy saving mode, or placed in any suitable state when deallocated. The telephony resources 112 may additionally or alternatively be operated as virtualized resources on a cloud computing platform (which may be operated by an outside party such as Elastic Compute Cloud operated by Amazon). When a telephony resources 112 such as a call router 114 is deallocated the virtualized resources may be returned to the vendor, given to other customers of the cloud computing platform, ending the virtualization of the resources, or any suitable process. A software instance may be quit or deleted when deallocated. The ratio of resources, such as the ratio of call routers to media processing systems, may be adjusted or maintained.

[0018] A call router 114 is preferably connected to a Public Switched Telephone Network (PSTN) device over the PSTN network, such that it can receive and make calls from PSTN-connected devices, such as landlines, cellular phones, satellite phones, or any other suitable PSTN-connected devices, as well as non-PSTN devices, such as Voice-Over-Internet-Protocol (VOIP) phones, SIP devices, Skype, Gtalk, or other Internet addressable voice devices. Thus the call routers 112 can preferably create connections to the telephone network of the distributed telephone controller. The call router 112 may alternatively or additionally function as or include a message router for use with telephony messaging such as SMS (Short Message Service) messages or MMS (Multi Media Messaging). The call router 112 can preferably connect to a messaging network, such that it can receive and send messages from SMS/MMS network devices, cellular

phones, computers, smartphones, or any suitable SMS/MMS network devices. The call router 112 may also send or receive text messages, multimedia messages, emails, faxes and other suitable PSTN-compatible communication messages. The call router 112 preferably communicates with the application server using an application layer protocol, more preferably using the HTTP (Hypertext Transfer Protocol), or secure HTTPS (Hypertext Transfer Protocol Secure), protocol. The application server preferably hosts a telephony application, sound file, text file, a database, and/or any suitable media, resource or file that can be used by the call router for telephone interactions. The call router 112 may additionally generate call router resources. The call router resources are preferably accessible by the application server and other devices (such as other call routers) through a call router API. The call router resource functions as an addressable representation of call router meta-data, internal call router state, or the state of a given resource used by the call router. For example a call router 114 may record a call and save the recording as a call router resource.

[0019] Additionally, the telephony resource cluster 110 of the preferred embodiment may include a service application 116 that functions as a messaging component to coordinate functionality of an application that has been distributed across various call routers 114, hardware resources, and/or software resources. The service application 116 is preferably an internal resource that is used when normal operation of an application is prevented because the operation of an application is distributed amongst various hardware and software resources of the cluster 110. The service application 116 is preferably a messaging service that offers reliable messaging where a

message is delivered to a particular destination (such as to another call router 114). The service application 116 may alternatively offer broadcasting messaging that announces a message without knowing who receives a message or if the message was received. As a first example, a hang-up service application 116 may be used to coordinate hanging up call sessions on different call routers 114. The hang-up service is preferably used to communicate to the appropriate call routers 114 to cancel outgoing calls when, for example, an application wants to dial a plurality of numbers but then hang up all unanswered calls once one of the calls is answered. As a second example, a multiple input service may gather and input commands from multiple telephone devices. So dual-tone multi-frequency (DTMF) input or voice commands may be issued by any caller and communicated to the application even if the calls are distributed over multiple call routers 114 within the cluster. This may be used in voting applications within a conference call. In this way, a telephone application does not need to actively account for processing and call handling being distributed within the cluster, and the hardware and software resources of the cluster preferably appear as a single entity to outside applications because of the internal service applications 116.

[0020] Additionally, the telephony resource cluster 110 of the preferred embodiment may include a call router network 118 that functions to allow a level of communication and synchronization between various call routers 114. The call router network 118 may additionally or alternatively be applied to other hardware or software resources. The call router network 118 is preferably used to access shared resources or as a communication channel. In one exemplary application, a voice over internet protocol

(VOIP) connection is established over the call router network 118 for mixing audio from various call routers. The VOIP connection is preferably used in implementing conference calls distributed over multiple call routers 114. As another example, the call router network 118 may additionally be used to stream audio from a call router to a realtime internet audio stream. As another example, the call router network 118 may be used to access data on another telephony resource 112 such as by using the call router API to access a call router resource. The service application 116 and the call router network 118 may additionally cooperate in synchronizing applications distributed within the cluster.

[0021] The call queue 120 of the preferred embodiment functions to manage a stack of call requests. The call queue 120 is preferably a list of outbound call requests that have not been serviced or been assigned necessary resources. The requests are preferably serviced at a rate suitable for the current capacity of the network 150 and telephony resource cluster 110. The servicing rate may alternatively be adjusted according to the capacity of the distributed telephony controller 144, the telephony resource cluster 110, and/or number of requests in the queue 120. A call request (such as one made by a telephony application) is preferably placed in the call queue 120 when capacity is exceeded or alternatively placed in the call queue 120 for every request or based on any suitable rule.

[0022] In one variation, an application preferably has associated user limits, in particular: an inter-call request rate limit (throttle) and a total limit (cap). The throttle and cap are preferably used to determine the positioning of requests in the call queue.

The limits may alternatively be assigned to an account, phone number, or any suitable entity. Telephony messages (e.e., SMS or MMS) are one variation of a call request that can additionally be placed in the call queue. Inbound and outbound telephony message can preferably be queued because inbound messages do not require immediate action unlike inbound calls. The SMS message is preferably sent after the request is serviced in the queue. SMS messages and/or MMS messages may alternatively be queued in a dedicated message queue. SMS message may have a rate limit (throttle) and total limit (caps) that varies from requests. Requests received at any rate from a user are preferably spaced in time in the call queue according to the throttle. There is preferably a latency enforced between call requests from an application. Requests of different users are preferably ordered in the queue in a staggered or alternating fashion as shown in FIGURE 6, but alternatively, users may have priority based on a service plan, first-come-first-serve policy, type of call request, and/or any suitable policy. The cap is preferably a limit on the total number of requests a user can make in a given amount of time. The user limits, handling, spacing, and/or ordering of the call queue 120 function to prevent one application from unfairly dominating the usage of the telephone network 150 or telephony resource cluster 110 at any one time. Additionally, applications may request access to telephony resources 112 as soon as possible or at some time in the future (e.g., a user schedules a call or calls for a later time). Additionally or alternatively, the user limits may be adjusted or set according to the needs of an application. An application may have particular requirements based on the nature or characteristics of

the application of the user. The user limits are preferably set according to the contract and/or pricing model that a user selects or by any suitable means.

[0023] In another variation, the call queue 120 is dedicated to requests of a single user entity. In this variation, there is preferably a plurality of individually assigned call queues 120. Call requests are preferably organized into a call queue 120 for each user. Telephony message requests alternatively have a queue for each phone number. A user requests can preferably be added to the individually assigned queue 120 at any time. Each queue is preferably serviced (i.e., dequeued) on a schedule that considers the per-user limits (such as resource limits, system-wide limits, etc.). In other words the dequeuing occurs in an alternating fashion between the plurality of call queues 120. The individually assigned call queues may additionally be for particular resources, and the dequeuing preferably occurs according to the dequeuing rate of the particular resource. The dequeuing rate is preferably related to the capacity of the resource but may alternatively be based on any suitable criteria. As with the other queuing variations, queuing may alternatively occur according to any suitable queuing methodology such as randomly, in a round-robin fashion, with fair queuing, with weighted fair queuing, based on actual resource utilization, and/or any suitable methodology. As an alternative to queuing based on account/phone number, call or message requests may be queued based on time, priority, usage history, or any suitable aspect. There may additionally be a control queue used to coordinate the dequeuing of individually assigned call queues (or message queues) 120.

[0024] As mentioned above, the call queue 120 may include an additional or alternative system for handling telephony messages (e.g., SMS or MMS messages). SMS messages preferably have additional limitations on their servicing rates and restrictions. SMS messages are preferably not only queued for sharing telephone network access with various users, but rates are also preferably implemented to prevent SMS messages from a single user from being rate limited, identified as spam. A call queue 120 for telephony messages may include at least two types of queues: a control queue and a phone number queue. The phone number queue preferably functions as a personal queue of a single user for telephone messages the user wants to send, and the control queue functions substantially similar to the multi-user queue described above for the call queue 120. The individually assigned call queue 120 may alternatively be used without the control queue, and the individually assigned call queue 120 may be based on account phone number or any suitable assignment. The control queue and phone number queue preferably functions to isolate the queuing of messages for a particular application and the messages of the plurality of messages. The content of the SMS message (the text) or MMS message (the multimedia) is preferably not stored in the call queue directly, and a reference to the SMS message content is preferably stored. This functions to reduce the load on the queue. The SMS/MMS content is preferably stored and accessed when the queued reference is serviced.

[0025] A queue popper 122 (i.e., a dequeuer) is preferably a software or hardware mechanism that functions to select call requests to service from the call queue. The queue popper 122 preferably selects call requests at a preferred rate, but the queue

popper 122 may alternatively select calls requests according to capacity or available resources, or a combination thereof. There may additionally be a plurality of queue poppers 122 that function to simultaneously select call requests from the call queue 120. The number of call poppers 122 may be variable. Additional or special queue poppers 122 may be used for the additional SMS call queues. The call queue(s) 120, the queue popper(s) 122, or any suitable combination are preferably used to control the throttling (or servicing rate) of the call requests. The throttling may be performed on a per-phone number, per-account (as in a multi-tenant application), and/or according to any call/message attribute.

[0026] The analysis system 130 of the preferred embodiment functions to analyze the system to predict resource requirements. The analysis system 130 preferably monitors a plurality of aspects of the system. The analysis system 130 may monitor the current capacity such as network or hardware operation levels or trends (increasing or decreasing); usage history such as logged data to find correlations in capacity (e.g., detecting patterns); queue length and queue entry latency; analysis of applications such as historical patterns from past usage of an application; and/or any suitable aspects. Patterns in capacity needs are preferably found related to the time of day, day of the week, yearly patterns, usage patterns (such as if an increase in capacity needs by one user indicates increase in capacity needs by other users), call location, call duration of calls, and/or any suitable indicator. The analysis system 130 preferably makes distinctions between inbound and outbound capacity for telephone network channels. The analysis system preferably generates data for the resource allocator 140, a

distributed telephone controller 144, a load balancer 142, and/or additionally the call queue 120. The predictions or data from the analysis system may additionally be used for provisioning capacity of the distributed call controller, planning capacity requirements of the static capacity of the telephone network, the number of call routers, hardware or software resources within the cluster, and/or parameters of queue management. The analysis system 130 preferably compares expected and actual load, and provides data that is used to compensate for the variability in utilization of resources of the system.

[0027] The resource allocator 140 of the preferred embodiment functions to scale and manage the operation of the telephony cluster 110. The resource allocator 140 additionally preferably reprovisions telephony resources 112 of the cluster 110, allocates new telephony resources 112, deallocates telephony resources, and/or any other suitable allocation process. The resource allocator 140 may additionally control the provisioning of call queues and other devices of the system. The resource allocator 140 preferably uses data of the analysis system 130 in determining the provisioning and operation of resources. The resource allocator 140 preferably uses information from the analysis system 130 to predict required telephony resource 112 capacity. The resource allocator 140 preferably uses the predicted capacity requirements to determine how many hardware (physical or virtualized) or software resources need to be operating, and the resource allocator preferably allocates, deallocates, or reprovisions telephony resources 112 (e.g., call routers and/or other hardware or software resources) as necessary. The resource allocator 140 may additionally use startup time, operation cost, or other

parameters of hardware and software resources when determining the number and ratio of resources to have allocated at a particular time. The resource allocator 140 also preferably keeps track of the quantity of resources currently available, and makes resource availability information available to other system components, including dequeuers, load balancers etc. Such resource availability information is preferably used by other system components to adjust operation of the system components. The resource allocator 150 preferably monitors the resources and reprovisions resources in real time.

[0028] The resource allocator 140 of the preferred embodiment preferably includes a load balancer 142 that functions to distribute processing tasks amongst the call routers and other hardware. The load balancer 142 of the preferred embodiment preferably optimizes the distribution of processing tasks such that the plurality of call routers 114 is operated at or near an optimal level. The operation of the call routers 114 may be optimized for performance, energy, cost, and/or any suitable conditions. The load balancer 142 preferably directs tasks (e.g., servicing of call requests/sessions) to appropriate call routers 142 (or telephony resource 112) as the tasks are created. A task is preferably an operation of a telephony application, but may alternatively be a call request or call session. In one example, one hundred call routers 114 may provide the call router tasks for one hundred telephony applications. In a second example, one hundred call routers 114 may each handle a single call session associated with one telephony application, such as for a conference call application with one hundred participants. The resource allocator 140 preferably sends notifications as to the current

status of resources of the system (the load of resources, the number of resources, etc.) to the load balancer 142. The load balancer 142 distributes requests to currently available and running resources matching the requirements of the application being load balanced, based on data provided by the resource allocator 140.

[0029] The resource allocator 140 of the preferred embodiment may include a distributed call controller 144 that functions to controls usage and operation of the telephone network 150 by the system. The distributed call controller preferably manages the shared usage of the telephone network channels 150 by the plurality of telephony resources. The distributed call controller 144 may alternatively be a subset of multiple telephone networks if multiple network providers or carriers are used. The operation of the distributed call controller 144 preferably functions to operate an allocated number of channels for current capacity requirements of the telephone network 150. The allocated channels are preferably channels within the available static channel capacity that are in use or prepared for use. The distributed call controller preferably has less than or equal capacity as the static channel capacity at any given time. The capacity of the distributed call controller 150 can preferably be increased by allocating more resources of the telephone network to the call controller, and the capacity of the distributed call controller 144 can preferably be decreased by deallocating resources of the telephone network. As an example, a commodity hardware node may be added to the telephone network to run a telephony software stack during high capacity requirements. The distributed call controller 144 preferably uses the analysis system 130 to predict or respond to the desired capacity requirements. The telephone network 150 may

additionally be divided into inbound channels, outbound channels, and bidirectional channels that can be used for receiving calls, making calls, and both, respectively. The telephone network 150 may further include SMS or MMS inbound and outbound channels. The distributed call controller 144 preferably manages the usage of the type of channels according to predicted usage. The bidirectional channels are preferably used for flexibility in capacity requirements. As one example, if inbound call load is expected to be high, then outbound calls are preferably directed to outbound channels to leave more capacity for inbound calls. The distributed call controller 144 may additionally manage the number and usage of allocated channels according to subscription or contracts from network providers. Channels may be used allocated or deallocated to ensure that volume pricing thresholds or other network conditions are satisfied.

[0030] A telephone network with a static number of channels 150 is preferably the base infrastructure for providing users with telephone network access. Telephony sessions are preferably communicated over the telephone network and the telephony sessions preferably include telephony voice sessions and/or text/media messaging (telephony messaging). The static number of channels is preferably the total number of concurrent telephony sessions or calls that can be supported at one time. The number of channels is typically limited by the number of interconnections available to a specific carrier or network. The telephone network 150 may alternatively be composed of multiple carriers or network providers or the Public Switched Telephone Network, but the plurality of carriers or networks is preferably managed or handled as one telephone network. The static number of channels is preferably a set number for a period of time

(usually based on a contract with a telephone company), and the number is preferably large enough to provide sufficient capacity. The static number of channels preferably determines the capacity of a network and the ability of the telephone network to connect with other networks. The operation of the telephone network is preferably handled by providing applications access to a channel of the telephone network. The telephone network may have a given number of channels not being used at any given time. In one variation, the telephone network may alternatively operate unused channels in an unused-mode. The unused mode may be a full or partial hardware power down mode, a hardware sleep mode, a secondary use (such as for non-crucial uses that can preferably be interrupted with minimal adverse effects), and/or any suitable way. The unused mode would function to reduce operation cost and/or maximize the utility of unused capacity. The telephony network channels 150 are preferably Public Switched Telephone Network (PSTN) connections but may alternatively be Session Initiation Protocol (SIP) trunks or any suitable device to create a telephony network connection to a telephony device.

2. Method of Operating a Multitenant Telephone Network

[0031] As shown in FIGURE 1, the method 100' of operating a multitenant telephony network of the preferred embodiment includes the steps of multiplexing call requests of a plurality of users to a telephony resource S110, creating a first call session from the call request through the telephony resource S130, and multiplexing the call session with a plurality of additional call sessions to a telephony channel S140. The

method 100' functions to create an efficient and scalable network system for resource intensive telephone applications. The telephony resource is preferably part of a telephony resource cluster. The telephony resource cluster preferably scales to satisfy immediate capacity demands which functions to reduce operation cost and allow a wide variety of applications to use the multitenant telephony network due to the ability to handle a wide spectrum of network loads. Additionally, the method 100' functions to allow the operation of a telephony application to be distributed between a variety of multi-user, shared resources (e.g., a telephony resource) such that the specific goals of telephone applications are not limited by the multitenant telephony network. The method 100' of the preferred embodiment is preferably implemented by a system described above, but may alternatively be implemented by any suitable system.

[0032] Step S110, which includes multiplexing call requests of a plurality of users to a telephony resource, functions to share the use of a telephony resource between a plurality of users. A single telephony resource is preferably shared between a plurality of users/applications. The multiplexing preferably occurs in a form of time division multiplexing in which call requests are sent to telephony resource in an alternating fashion. The time division multiplexing is preferably based on completion of complete call sessions or processes. In other words, users take turns using the telephony resource to create a call session and run an application. For example, a first customer preferably has a call request serviced by a telephony resource and upon completion of the call session of the call request, a second user may have a call request serviced by the same telephony resource. A call request is preferably received from a user or more specifically

a telephony application residing on an external server, but the call request may alternatively be sent from any suitable source. The call request is preferably received over a packet-based communication channel, in other words a non-direct communication channel. In one variation, the call request is preferably received in a HTTP or HTTPS message but may alternatively be received through any suitable application communication protocol. Step S110 may additionally include queuing a call request of a user S112, which functions to gate or prioritize incoming call requests. The call queue is preferably used for outbound requests, while inbound calls are preferably handled immediately (or else the call session will most likely fail). Alternatively, an inbound call may be queued for full service, with a “ringing” audio played back while call is waiting in the queue to be fully serviced. A queue may, however, be used for inbound telephony messages because telephony messages such as SMS messages and MMS messages will be resent if not received on the first attempt. The call queue is preferably a list of pending call requests from a plurality of users. An additional queue may additionally or alternatively be used for telephony messages. The call requests are preferably ordered within the queue in a way that balances access to resources. Each user (e.g., account, application, or phone number) is preferably assigned an inter call request limit (a throttle) and a limit on the maximum number of call requests that can be made in a specified amount of time (a cap). Call requests are preferably selected for servicing at a specified rate or by a device (i.e., a queue popper), which may be selecting calls based on current load on the telephony resource cluster. The queue may alternatively be operated in any suitable variation such as those described above. A

queue may be assigned to each user or phone number. Queuing may alternatively occur according to any suitable queuing methodology such as randomly, in a round-robin fashion, with fair queuing, with weighted fair queuing, based on actual resource utilization, and/or any suitable methodology. A load balancer preferably distributes call requests to a telephony resource that has the least capacity. The load balancer and the call request queue preferably cooperatively distribute the load as described above.

[0033] As an additional step, method 100' preferably includes provisioning resources of the telephony resource cluster S120, which functions to scale the capacity of the telephony resource cluster to adequately multiplex a call request to a telephony resource. Step S120 may include reprovisioning an existing telephony resource of the telephony resource cluster, allocating additional resources to the telephony resource cluster, and/or deallocating resources of the telephony resource cluster, and/or re-allocating resources from one type of resource to another in realtime. The telephony resource cluster preferably includes a plurality of telephony resources performing various functions or operations as described above. For example, the telephony resource cluster may include a plurality of call routers, transcription systems, media processing systems, and text-to-speech systems. A telephony resource preferably is composed of a computer processor and/or storage resources for a first purpose. As part of S120, a resource of the telephony resource cluster a processor and/or storage device of a telephony resource is preferably reprovisioned for a new second purpose. For example a text-to-speech may be reprovisioned to function as a call router at times when more calls need to be served. Additionally, more resources may be allocated or deallocated

which may include adding new resources to the system and/or activating resources, or re-allocating resources from another customer of a shared resource environment. The resources are preferably those provided by a multitenant shared virtualized computing environment such as a cloud hosting provider (i.e., a web service that provides resizable compute capacity that allows a user to boot a machine image to create a virtual machine resource), but may alternatively be physical machines either co-located or distributed. For example, a number of resources may be operating in a powered down state. When the more capacity is required, the resources may be turned on/booted (i.e., allocated) to serve as a new resource of the telephony resource cluster. Similarly, when the telephony resource cluster has more capacity than is currently required a resource may be powered down, returned to a pool of resources for use by other companies (i.e., deallocated), or any suitable action to end current use of the resource.

[0034] Additionally, Step S120 may include analyzing resource capacity requirements S122 which functions to collect data on real-time or imminent capacity requirements. Data may be collected from the call request queue, from stored history on capacity requirements, current load of the telephony resource cluster, data from an analysis of applications, or any suitable source of predicting capacity requirements. Data from the call request queue may provide information such as number of pending call requests, the type or details of the call request, or any suitable queue related information. Stored capacity history preferably provides insight to capacity patterns such as temporal patterns throughout the day, week, or year. Current load of the telephony resource cluster preferably provides information such as the current number

of resources of the telephony resource, number of available resources of the telephony resource, the division of type of resources, the number of deallocated resources, the number of telephone network channels, etc. Application analysis data preferably is data from the telephone applications of users on expected or predicted capacity requirements. An analysis is preferably performed on the operation of the application and or gathered from a user on the expected capacity requirements of the application such as number of calls, peak time for calls, what type of calls (e.g., conference calls, SMS messages etc.). The analysis information is preferably used to control the provisioning, allocation, and deallocation of resources of Step S120. Additionally, after analyzing the capacity requirements, other components of the system such as the telephony resource cluster, telephony resources, call queue, dequeuers, resource allocator are preferably notified of relevant analysis information. Particular analysis information may be specifically sent to a component. For example, the load balancers and the dequeuers are preferably informed about available resources and adjust operation according to the capacity information.

[0035] Step S130, which includes creating a first call session from the call request through the telephony resource, functions to convert the call request into a call session using the telephony resource. Step S130 preferably additionally includes additional processing and steps specific to a particular application. In a preferred variation, a call router preferably processes telephony instructions of a call request to identify the destination phone number and then establishes a connection to the destination phone

number as part of Step S140. A transcription server may initiate recording or prepare to record a conversation of the call session.

[0036] Step S140, which includes multiplexing the call session with a plurality of additional call sessions to a telephony channel, functions to establish a telephone network connection to a telephony device. The telephony channel is preferably a PSTN (Public Switched Telephone Network) connection. This may be a physical wire or some interfacing infrastructure to connect to the PSTN. In some cases the concept of a channel is preferably subscribed to or rented from a telephone network. In one alternative, a SIP (Session Initiation Protocol) trunk may be used as an internet based gateway to a telephone network. The multiplexing preferably occurs in a form of time division multiplexing in which call sessions are connected to telephony channel in an alternating fashion. The time division multiplexing is preferably based on completion of complete call session. For example, a particular network channel may first be utilized for a call session of a first user, and upon completion of the call, a second call session may be established with the particular network channel for a second user. As part of Step S140, the telephony channels may additionally include provisioning telephony channels S142. This functions to adjust the number of available telephone network capacity of the system. By provisioning gateways to the telephone network (e.g., call routers or SIP trunks), channels or gateways to channels may be allocated or deallocated. Such scaling of telephony network channels preferably allows operation near the current telephone network capacity requirements. If such scalability was not in use then there would be a set limit on the number of channels that could be simultaneously used.

3. Method of Operating a Dynamic Telephone Network

[0037] As shown in FIGURE 6, the method 200 of providing a telephony network of the preferred embodiment includes the steps of operating a telephone network with a static number of channels S210, providing telephone network channel access to a plurality of users S220, and managing usage of channels to allow a user to access a number of channels that exceeds normal operation S230. The method functions to allow the operator of the telephone network to offer high capacity to a plurality of users, without a reduction in quality or reliability of services based on usage. This method is preferably implemented on a system substantially similar to the one described above, but any suitable system may alternatively be used. The method may additionally be used in combination with the methods herein described. The method 200 further functions to allow users to use the telephone network without a specific concern about the number of channels required for operation. The users of the telephone network are preferably operating telephony applications such as call centers, Private Branch Exchanges (PBX), phone trees, telephony phone applications, VOIP services, SMS or MMS services, and/or any suitable telephony application. The operators of the telephone network are preferably a telephone service provider such as a telephony platform provider (e.g., a internet-telephone platform provider), a telephone company (e.g., owners of a telephone network such as AT&T), and/or any suitable party. In a variation of the preferred embodiment, the method 200 may additionally include a distributed call controller, a call queue, and/or the step of assessing capacity requirements.

[0038] Step S210, which includes operating a telephone network with a static number of channels, functions to be the base infrastructure for providing users with telephone network access. The static number of channels is preferably the total number of concurrent telephony sessions or calls that can be supported at one time. The number of channels is conventionally limited by the number of interconnections available to a specific carrier or network. The telephone network may, however, be composed of multiple carriers or network providers or the Public Switched Telephone Network, but the plurality of carriers or networks is preferably managed or handled as one telephone network. The static number of channels is preferably a set number for a period of time (usually based on contract with a telephone company), and the number is preferably large enough to provide sufficient capacity. The static number of channels is preferably an indication of the capacity of a network and the ability of the telephone network to connect with other networks. The operation of the telephone network is preferably handled by providing users access to a channel of the telephone network. The telephone network may have a given number of channels not being used at any given time. In one variation, the telephone network may alternatively operate unused channels in an unused-mode. The unused mode may be a full or partial hardware power down mode, a hardware sleep mode, a secondary use (such as for non-crucial uses that can preferably be interrupted with minimal adverse effects), and/or any suitable way. The unused mode would function to reduce operation cost and/or maximize the utility of unused capacity.

[0039] As an additional alternative to the preferred embodiment, the method may include operating a distributed call controller as a subset of the telephone network S212. The distributed call controller may alternatively be a subset of multiple telephone networks if multiple network providers or carriers are used. The operation of the distributed call controller preferably functions to operate an allocated number of channels for current capacity requirements of the telephone network. The distributed call controller preferably has less than or equal capacity as the static channel capacity at any given time. The capacity of the distributed call controller can preferably be increased by allocating more resources of the telephone network to the call controller, and the capacity of the distributed call controller can preferably be decreased by deallocating resources of the telephone network. Access to the telephony network is preferably facilitated by virtualized hardware or software (such as call routers or SIP trunks). Allocation of more resources of the telephone network may additionally include a virtualization of a device to access a telephony network. For example a virtualization of a network access channel may be added to add further access capacity to the telephony network. As another example, a commodity hardware node may be added to the telephone network to run a telephony software stack during high capacity requirements.

[0040] Step S220, which includes providing telephone network channel access to a plurality of users, functions to allow a plurality of different parties to access the channels of the telephone network. The users preferably subscribe to a service of the operator of the telephone network. The users of the telephone network are preferably operating telephony applications such as call centers, Private Branch Exchanges (PBX),

phone trees, Interactive Voice Response (IVR) applications, internet-telephony applications, VOIP services, and/or any suitable telephony application. The user preferably does not subscribe to the service based on any specific number of channels. From the viewpoint of the user, the number of channels is preferably infinite or an irrelevant point for the operation of an application of the user. The user is preferably presented a per usage or time perspective (e.g., pricing and/or application usage perspective), while the telephone network is being operated on a per channel basis. The operator of the telephone network preferably converts costs associated with the operation of the telephone network (e.g. fixed capital costs of leasing from a telephone company or operation cost) into variable costs for the users. The access to the telephone network is preferably operated, leased, and/or on contract from a telephone company (such as AT&T) by a per channel basis. A lease agreement or contract may alternatively be negotiated to minimize per-channel (capacity) cost and preferably emphasize per usage or per time costs, or alternatively, any suitable leasing agreement or contract may be used. Users preferably pay by usage, a flat rate for a time period, per minute, a combination of usage and time charges, and/or any suitable pricing model.

[0041] Step S230, which includes managing usage of channels to allow a user access to a number of channels that exceeds normal operation S130, functions to provide high capacity capabilities to users while ensuring that the quality and reliability of the telephone network is not adversely affected by the usage of other users. An individual user of the plurality of users is preferably allowed to use a number of channels greater than an equal division between the plurality of users of the static

number of channels. The sum total of the maximum number of channels an individual user uses at given times may preferably be greater than the static number of channels. The given times where an individual user has access a maximum number of channels is preferably when demand on the telephone network by other users is low. Usage of the telephone network and the telephony resource cluster is preferably time based multiplex based on the completion of telephony sessions (i.e., users share the use of the resources and network). In a simplified example, a telephone network has 10 channels available and there are five users. When distributed uniformly, the users would each have 2 channels available for usage, but in one preferred embodiment all five users may have access of up to 10 channels each, assuming no other user is using the channels. During regular use of the telephony network, the user still has the ability to access the maximum number of telephone network channels but the call requests are preferably gated by user limits implemented by a call queue. In another example extending on the above example, analysis might indicate that 4 users may use 2 channels at a given point of time, thus 8 may be available to the 5th user, while keeping capacity available for the first 4 users. Managing the usage of the channels preferably includes managing the usage of resources such as by: managing a call queue, enforcing user limits, predicting and/or analyzing usage and capacity requirements, adjusting capacity based on the capacity of the distributed call controller, and/or any suitable steps of managing the resources of the telephone network. Capacity of the distributed call controller may additionally be controlled or affected by predictions and analysis and user limits may additionally be affected.

[0042] The method of the preferred embodiment may additionally include the step of managing a call queue of requests from the plurality of users S232. Step S232 functions to prioritize the handling of call requests from users. The call queue is preferably a program or hardware managed stack that is operated as part of a control architecture of the telephone network. The control architecture preferably manages the telephone network and usage by the plurality of users. The call queue is preferably a list of call requests awaiting service by the telephone network including telephony voice session requests and/or SMS/MMS message requests. The requests are preferably serviced at a rate suitable for the current capacity of the network and for each user. The servicing rate may alternatively be adjusted according to the capacity of the distributed call center or number of requests in the queue. A user request is preferably placed in the call queue when capacity is exceeded or alternatively placed in the call queue for every request or based on any suitable rule. A user preferably has associated user limits, in particular: a call rate limit (throttle) and a total limit (cap). The throttle and cap are preferably used to determine the positioning of requests in the call queue. Requests from a user are preferably spaced in time in the call queue according to the throttle. Requests of different users are preferably ordered in the queue in a staggered or alternating fashion as shown in FIGURE 6, but alternatively, users may have priority based on a service plan, first-come-first-serve policy and/or any suitable policy. The cap is preferably a limit on the total number of requests a user can make in a given amount of time. Subsequent requests are preferably schedule for a later time according to the cap, but requests exceeding the cap may be handled in any suitable manner. For

example, if user can make one call per second, and the user requests 100 calls, they will be scheduled equally over the next 100 seconds. Note that this cap can be described as the number of calls / time frame (1/second), or the required latency between calls in the queue (1 second). The user limits, handling, spacing, and/or ordering of the call queue function to prevent one user from unfairly dominating the usage of the telephone network at any one time. In the variation of SMS/MMS message requests, the rate of individual users is considered to prevent message filtering by a network. For the SMS/MMS variation the requests may additionally be queued in a control queue and a phone number queue. The contents of the SMS/MMS messages are preferably stored and a reference to the contents of a message is queued which functions to reduce the load on the queue. A plurality of cache servicing ports or pointers are preferably used. The servicing ports are preferably a software and/or hardware control mechanism for operating a call request from the call queue. A servicing port preferably takes a request from the call queue and connects the corresponding user application or user to a telephone network channel. The servicing port may be a direct connection, but may alternatively be a hardware or software resource such as a call router in a cluster as described above. The servicing ports are preferably less than the static number of channels to allow capacity for incoming calls, but the servicing ports may alternatively be equal to the static number of channels. In one example where there are 1000 channels of the telephone network, there may be 500 service ports. This would leave 500 channels free for incoming calls. Additionally, users may request access to telephony resources as soon as possible or at some time in the future (e.g., a user

schedules a call or calls for a later time). A queue popper is preferably a software or hardware mechanism responsible for selecting a call from the call queue to service. There may additionally be a plurality of queue poppers to select calls from a call queue. Additionally or alternatively, the user limits may be adjusted or set according to the needs of a user. A user may have particular requirements based on the nature or characteristics of the application of the user. The user limits are preferably set according to the contract and/or pricing model that a user selects or by any suitable means.

[0043] The method of the preferred embodiment may additionally include the step of predicting capacity requirements for the distributed call controller S234. Step S234 functions to assess indicators that correlate to the number of telephone network channels needed at a later point. The predicting of capacity is preferably accomplished by programmatically or mathematically (through pattern detection or any suitable algorithm) analyzing current and past information but any suitable method may alternatively be used. Patterns in capacity needs are preferably found related to the time of day, day of the week, yearly patterns, usage patterns (such as if an increase in capacity needs by one user indicates increase in capacity needs by other users), call location, call duration of calls, and/or any suitable indicator. The predictions of Step S234 may additionally be used for realtime provisioning, deprovisioning, and/or reprovisioning capacity of the distributed call controller or planning capacity requirements of the static capacity of the telephone network.

[0044] The method of the preferred embodiment may additionally include the step of reacting to capacity needs of the call queue S236. Step S236 functions to use the

call queue and other current capacity indicators to adjust the distributed call controller for the current capacity requirements or anticipated near term requirements. The call queue is preferably assessed through software or alternatively by any suitable monitoring of the call queue. The number of calls currently in the queue, the total number of users currently using the telephone network, incoming calls (that may be not be queued), the frequency of user requests, and/or any suitable characteristic of the telephone network or the call queue preferably cause a reaction to the capacity needs. The reaction is preferably for current overall capacity needs but may alternatively be for current capacity needs of an individual user or any suitable party. The reactions may include adjusting the settings of the call queue (such as call queue service rate or ordering), modifying user limits, adjusting capacity of the distributed telephone controller, and/or any suitable action. In one example, a call queue may have many calls scheduled for 100 seconds after the current time, the distributed call controller may increase capacity to accommodate the anticipated capacity requirements.

[0045] The method of the preferred embodiment may additionally include the step of analyzing capacity needs of a user and predicting the telephone network capacity needs S238. Step S238 functions to detect individual capacity needs to determine total capacity requirements of the telephone network. Capacity needs of a user are preferably acquired by analyzing a telephony application of a user. Part of the analysis preferably includes detecting periodic events that indicate capacity needs of an individual application. An example of such an event might be an application associated with a weekly TV show where callers call in around the air time of the show. The analysis may

alternatively or additionally include detecting typical call duration for an individual application. Some applications may only be used for a brief amount of time (such as when a short message is played), while other applications may require longer durations of use (such as when a user must navigate a long phone tree). Additionally, application history may be used to determine usage patterns such as by monitoring maximum, minimum, and/or average capacity requirements, frequency of requests, duration of requests, number of SMS messages sent in a particular time duration, and/or any suitable call characteristic. Usage characteristics of the individual applications of users are preferably combined with the usage characteristics of the other users to determine the total usage characteristics and capacity needs of the telephone network. Preferably, the code of the application is preferably analyzed to assess the functionality and usage patterns of the application. The application code or operation is preferably programmatically analyzed, but any suitable method may be used. Alternatively, the user and/or a second party may characterize the application and/or telephony service of the user. This characterization is preferably performed by the user while signing up, and preferably includes user expectations for the frequency of use, times of use, duration of calls, and/or any suitable characteristic of the application. The user may additionally prioritize when capacity should be highest for their application. Any suitable steps may be used to analyze an individual application.

[0046] As an additional alternative to the preferred embodiment, the method may include adjusting capacity of the distributed call controller S240. Step S240 functions to change the number of active channels of the telephone network to appropriately handle

the capacity requirements. Step S240 is preferably used in combination with Step S212, which includes operating a distributed call controller. The adjustments to the distributed call controller adjust the capacity capability that the operator offers. The capacity is preferably adjusted based on the management of the usage of channels of the telephone network. The capacity is more preferably adjusted based on the predictions and analysis of Steps of S234 and/or S236, but may alternatively be adjusted in cooperation with Step S232, Step S238, and/or for any suitable reason. When more capacity is needed, more resources, such as CPU, RAM, DISK, etc., capable of handling simultaneous channels or providing more channels, are preferably allocated to the distributed telephone controller, and conversely when less capacity is required, resources are preferably deallocated from the distributed telephone controller. The adjustment of capacity is preferably made to handle the expected or predicted capacity. The static capacity of the telephone network may alternatively or additionally be adjusted. As the telephone network capacity is typically less flexible. Adjustments to the telephone network capacity are preferably made for long-term capacity needs (e.g., on a per month basis). Any suitable adjustment to the system for more or less capacity may alternatively be used.

4. Method of Distributing Calls Between Telephony Hardware

[0047] As shown in FIGURES 7-8, the method 300 of distributing calls between telephony hardware of the preferred embodiment includes the steps of queuing a call request S310, selecting a load balancing call router S320, and connecting a call with the

selected call router S330. The method functions to balance usage of resources used in a telephony application. This method is preferably implemented on a system substantially similar to the one described above, but any suitable system may alternatively be used.

[0048] Step S310, which includes queuing a call request, functions to manage a call request until the necessary resources are available to service the call. A call request is preferably instantiated by a telephony application, a call router, a telephony device, and/or any suitable source of a call request. The call request may additionally be a SMS or MMS message request. The call request is preferably outgoing. An incoming call is preferably viewed as a more urgent call request than an outgoing request, and an incoming call may not be queued but alternatively may be passed directly to an available resource. Alternatively incoming call requests (call session initiations) may be queued, but since incoming calls have more immediacy they are preferably prioritized or the system must be have short queuing wait where a short wait is less than the time it would take for an incoming call to fail. The incoming call may alternatively be placed near the front of a queue or positioned in the queue according to separate rules appropriate for the higher priority of the call request. Similarly, a synchronous outgoing call request may be queued with high priority. A synchronous call is a call that another caller is relying on to proceed, as opposed to a new call initiated by an application in which a user will not notice a delay. Call requests are preferably ordered in the queue according to rules based on the throttle, caps, real-time urgency (priority) and/or any suitable factors.

[0049] Step S320, which includes selecting a load balancing call router, functions to identify a call router that should handle the call to preferably optimize the operation of a telephone resource cluster. The selected call router is typically the call router with the least load, but may alternatively be selected to optimize cost, energy usage, processing capability, and/or any suitable variable. Step S320 may additionally be applied to other hardware or software resources in addition to or alternatively to a call router. Call routers of a telephone resource cluster may have variable capacity and performance depending on hardware and/or software specs. The variability between the plurality call routers is preferably considered in selecting a call router. A load balancer substantially similar to the one described above is preferably the component that implements step S320, though step S320 may be performed by any suitable device. The load balancer is preferably capable of allocating and deallocating resources of the cluster, and so resources may be allocated and/or deallocated as a substep of S320. The resource allocator can preferably allocate and deallocate call routers, hardware resources, and/or software resources. The resources are preferably allocated or deallocated based on current or predicted utilization, but the resources may alternatively be allocated or deallocated as a function of other resources. For example, one media processing resource may be allocated (e.g., operating) for every five call routers. The selection of a load balancing call router preferably uses data from an analysis system. So that the step of selecting a load balancing call router may include selecting a call router that will balance load at a future time.

[0050] Step S330, which includes connecting a call with a selected call router, functions to pass control of the call to the specified resource. For an outgoing call, a call router preferably connects through a telephone network to the designated phone number. For an incoming call, the call router preferably connects to the specified telephony application; PSTN-connected device(s), such as landlines, cellular phones, satellite phones, or any other suitable PSTN-connected devices; non-PSTN devices, such as Voice-Over-Internet-Protocol (VOIP) phones, SIP devices, Skype, Gtalk, or other Internet addressable voice devices; and/or any suitable device associated with the number of the incoming call.

[0051] The method of the preferred embodiment may additionally include networking call routers that have a shared application S340. Step S340 functions to allow communication between multiple call routers. This is preferably useful in situations where functionality of an application is distributed over multiple resources (e.g., multiple call routers). The network preferably allows sharing of resources between call routers. Audio channels of call routers may additionally be mixed and shared between call routers. A VOIP channel is preferably formed over the network for bridging audio of different call routers. For example, a conference call may use the network to bridge audio of multiple call sessions from different call routers.

[0052] The method of the preferred embodiment may additionally include synchronizing applications with a service application S350. The service application functions to monitor an application distributed over a call router cluster and coordinate operation of the application. The service application may additionally be used to share

state information between the call routers. The service application preferably provides a specific functionality such as a hang up service or a multiple input service as described above. Any suitable application may be implemented by the service application such as input-gathering, multi-dialing, call splitting, call merging, and any suitable feature. Any number of service applications may be used.

[0053] As a person skilled in the art will recognize from the previous detailed description and from the figures and claims, modifications and changes can be made to the preferred embodiments of the invention without departing from the scope of this invention defined in the following claims.

CLAIMS

1. A system for operating a multitenancy telephony system comprising:
 - a call queue that stores call requests received from a plurality of users, wherein the call requests are ordered in the call queue according to a user inter-request rate limit and a user cap on the maximum number of requests in a specified time for a user;
 - a scalable telephony resource cluster that establishes call sessions for call requests;
 - an analysis system that calculates capacity requirements of the system;
 - a resource allocator that manages the scaling and operation of the telephony resource cluster; and
 - a plurality of telephony network channels that are used as telephony communication channels for call sessions.
2. The system of Claim 1, wherein the analysis system receives call queue data to calculate capacity requirements and wherein the resource allocator additionally manages allocation and operation of a plurality of call queues.
3. The system of Claim 1, wherein a call request is received from a telephony application, and wherein the analysis system analyzes the telephony application for capacity requirements, accesses past capacity requirement data, and monitors current capacity loads to calculate the capacity requirements of the system.
4. The system of Claim 1, wherein the call queue includes an additional queue for telephony message requests.

5. The system of Claim 1, wherein the telephony resource cluster includes a plurality of allocated telephony network channels that are a subset of the plurality of telephony network channels; and the resource allocator includes a distributed call controller that scales the allocated telephony network channels and connects call requests to a channel of the allocated telephony network channels.
6. The system of Claim 1, wherein the telephony resource cluster includes a plurality of telephony resources.
7. The system of Claim 6, wherein the plurality of telephony resources includes a plurality of call routers, wherein a call router makes and receives telephony communications through the telephony network channels and communicates with an application server of the user using an application layer protocol.
8. The system of Claim 7, further comprising a call router network that is a data channel between at least two call routers.
9. The system of Claim 7, further comprising a service application that coordinates the operation of a telephony application that includes at least two call sessions on at least two call routers.
10. The system of Claim 7, wherein the telephony resource cluster includes a plurality of allocated telephony network channels that are a subset of the plurality of telephony network channels; and the resource allocator includes a distributed call controller that scales the allocated telephony network channels and connects call requests to a channel of the allocated telephony network channels.

1 / 8

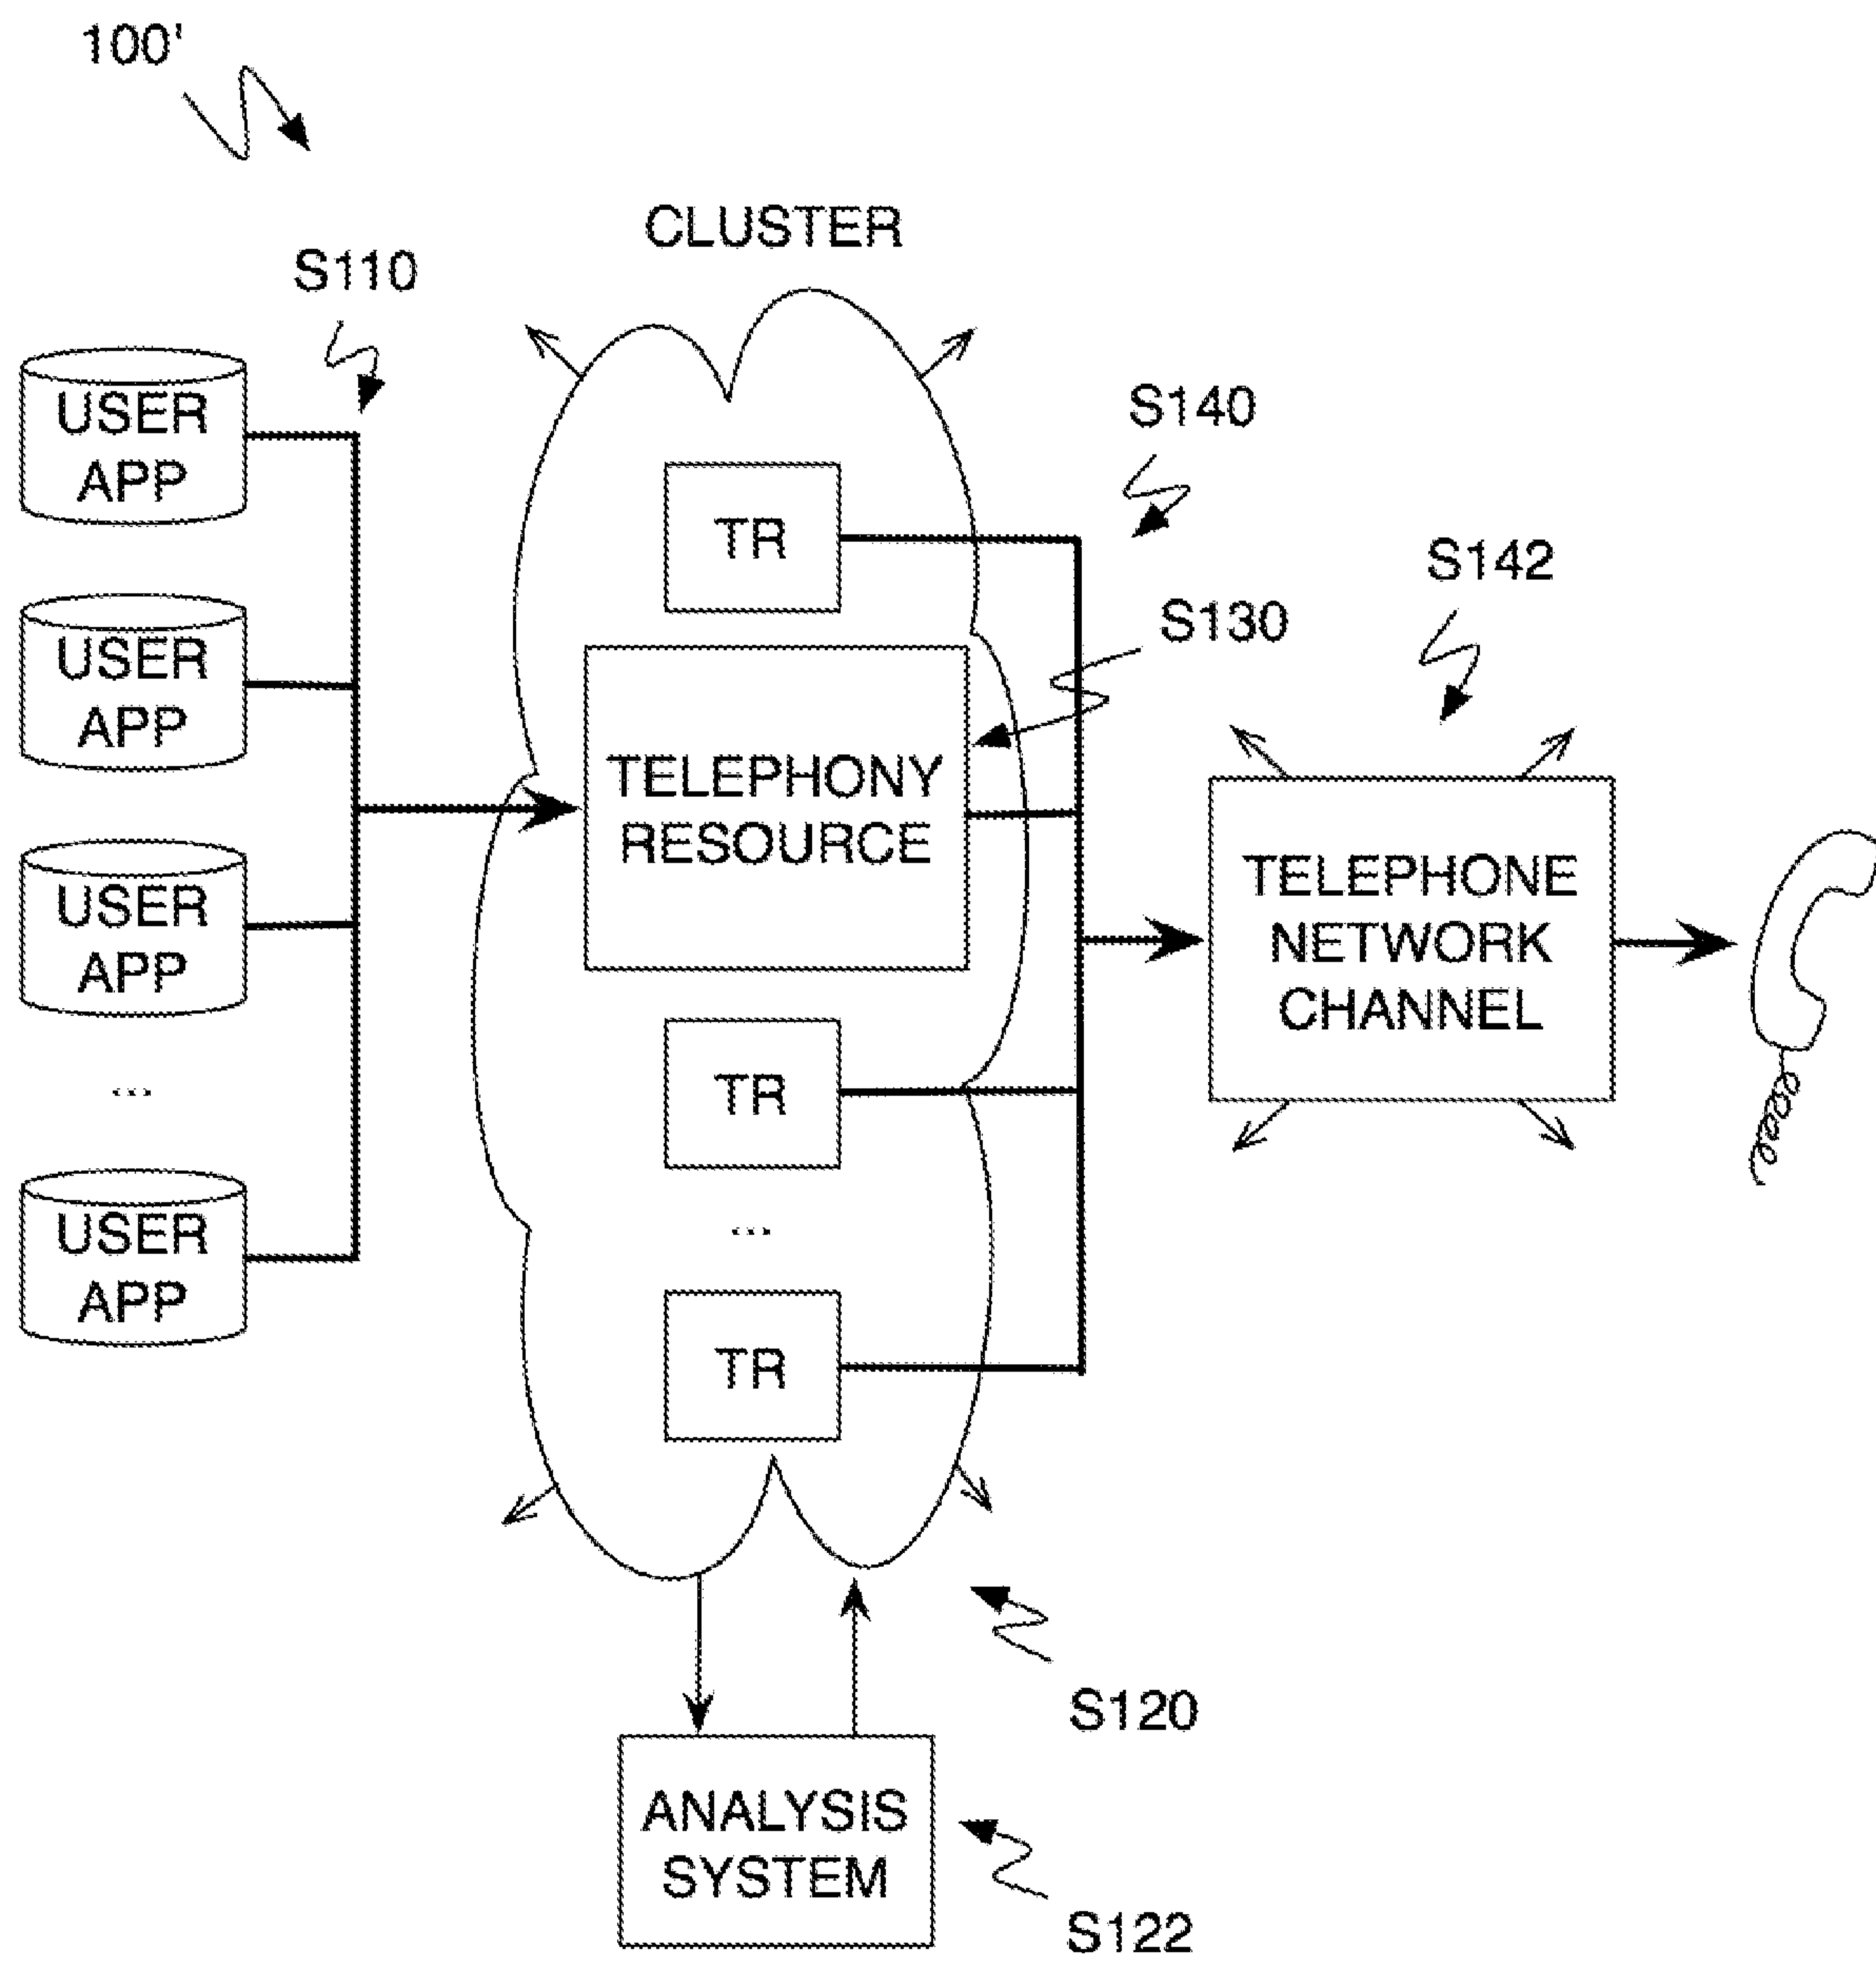


FIGURE 1

2 / 8

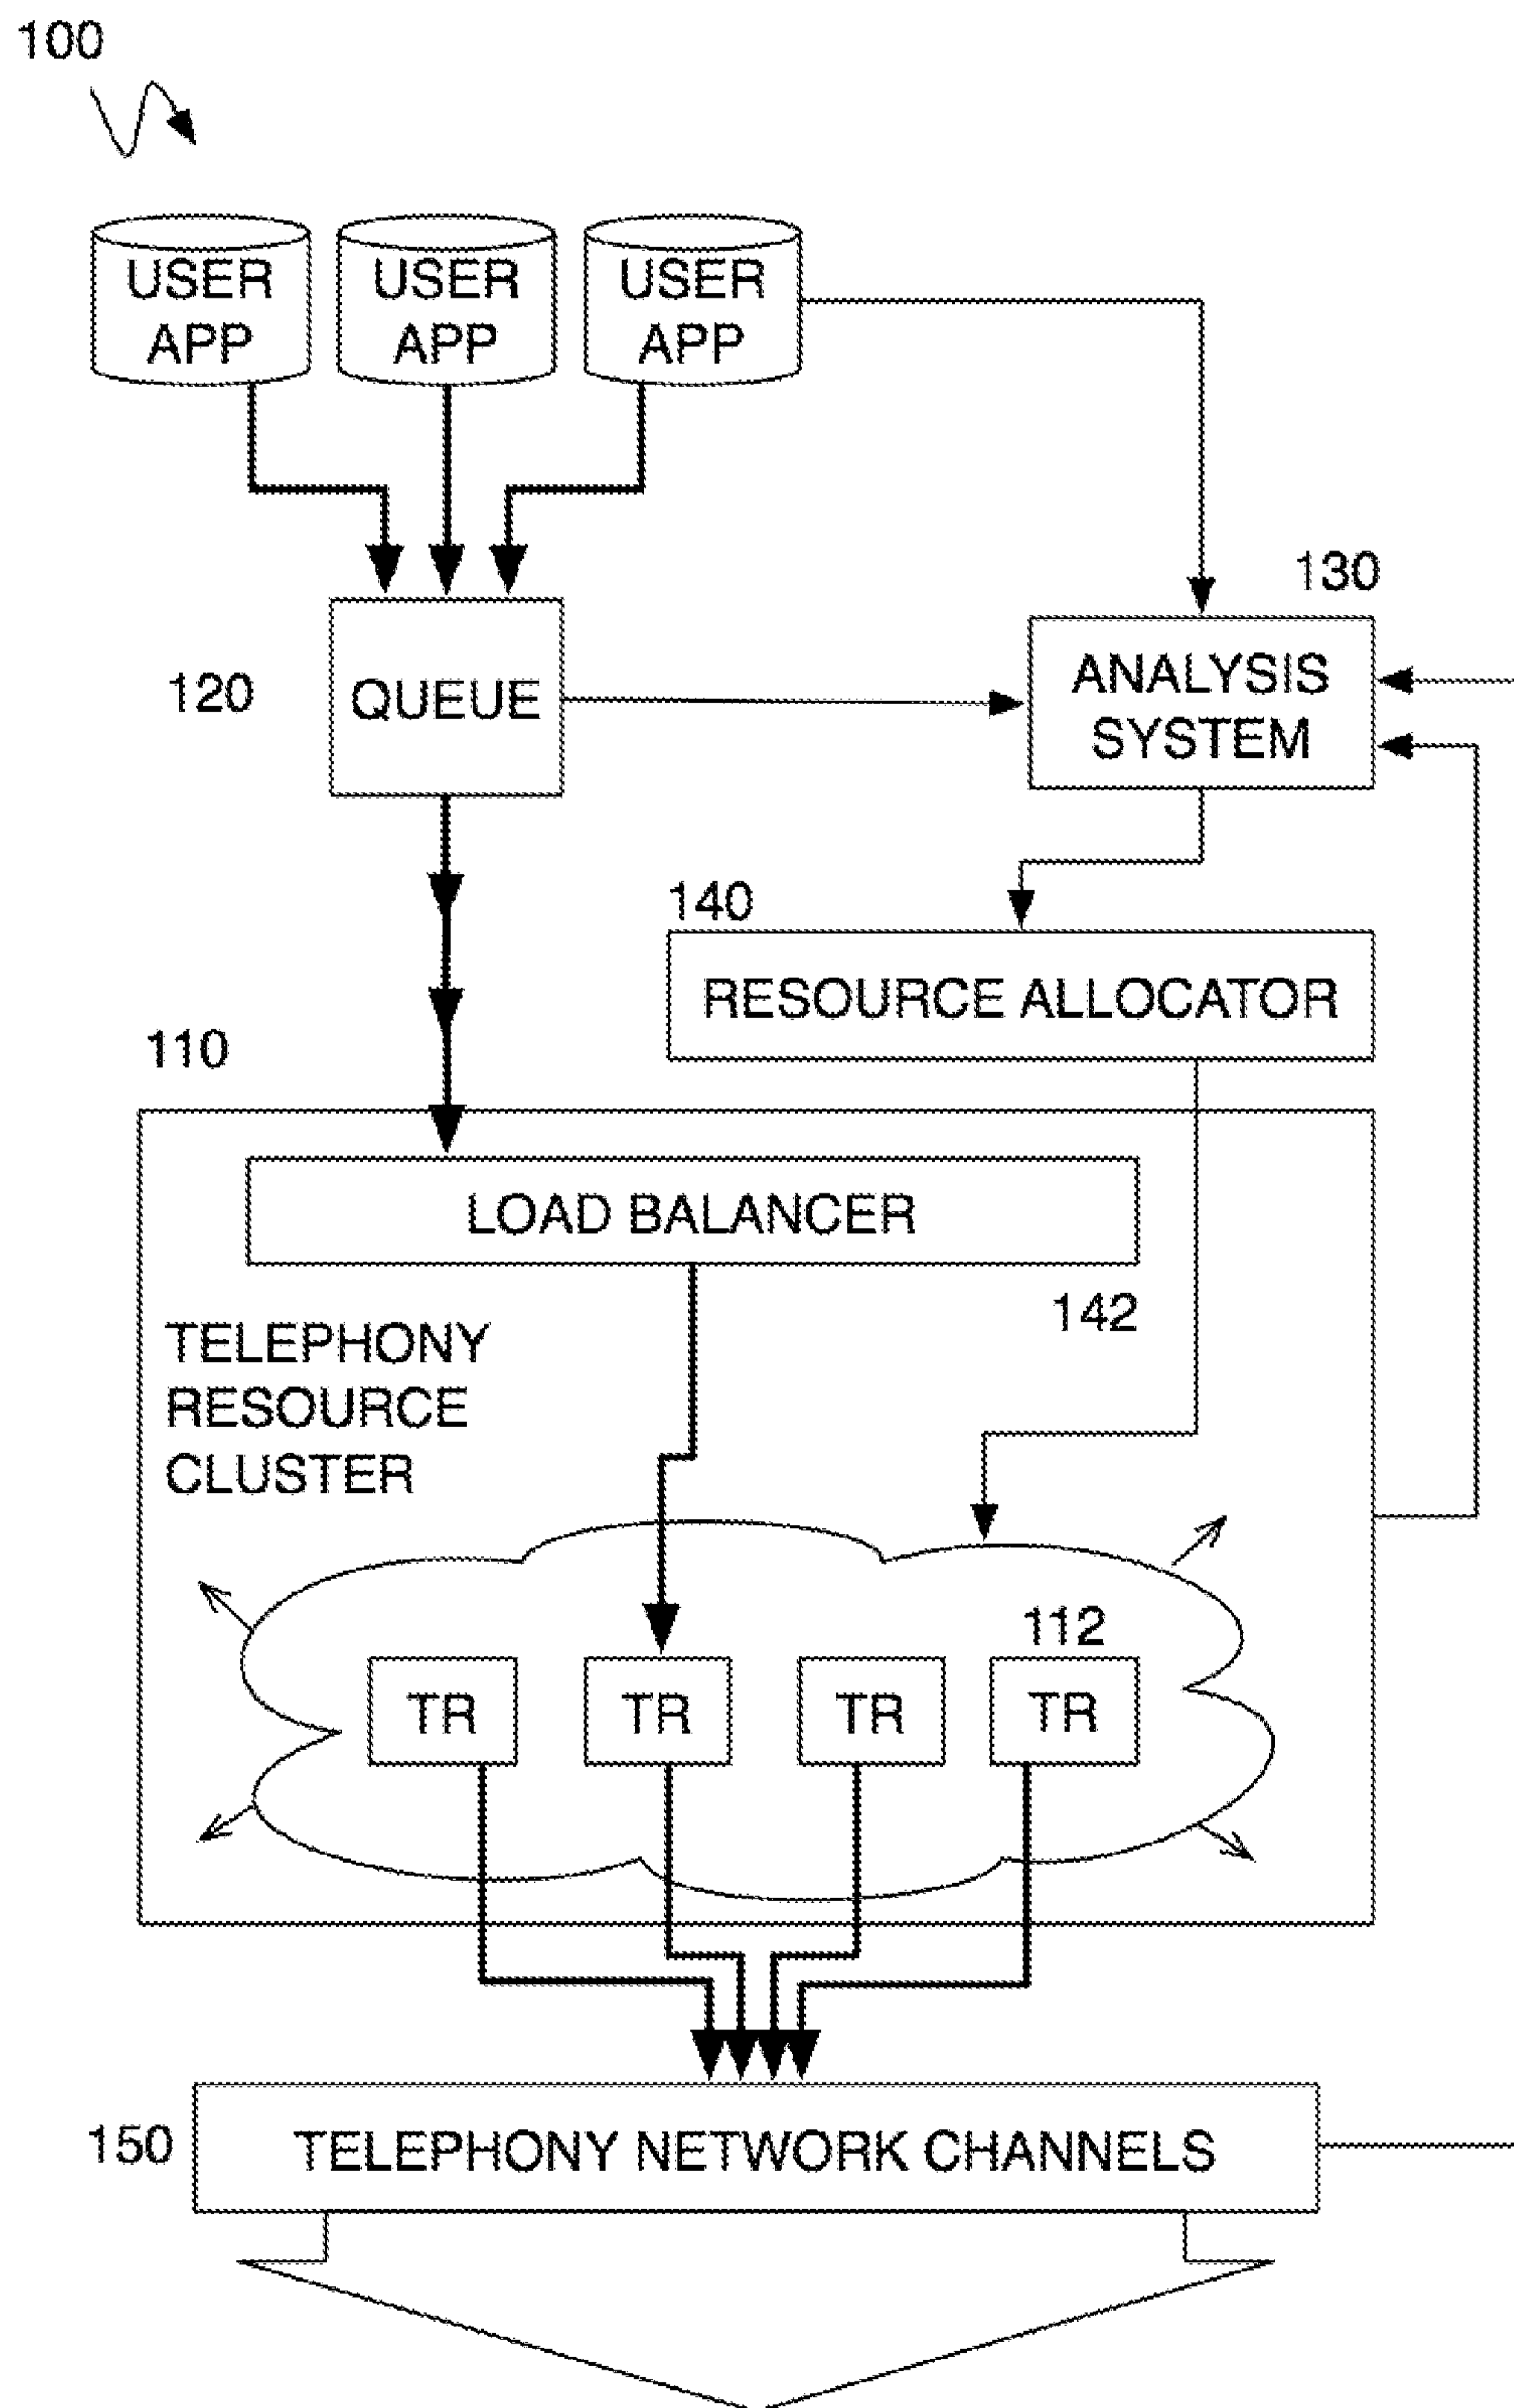


FIGURE 2

3 / 8

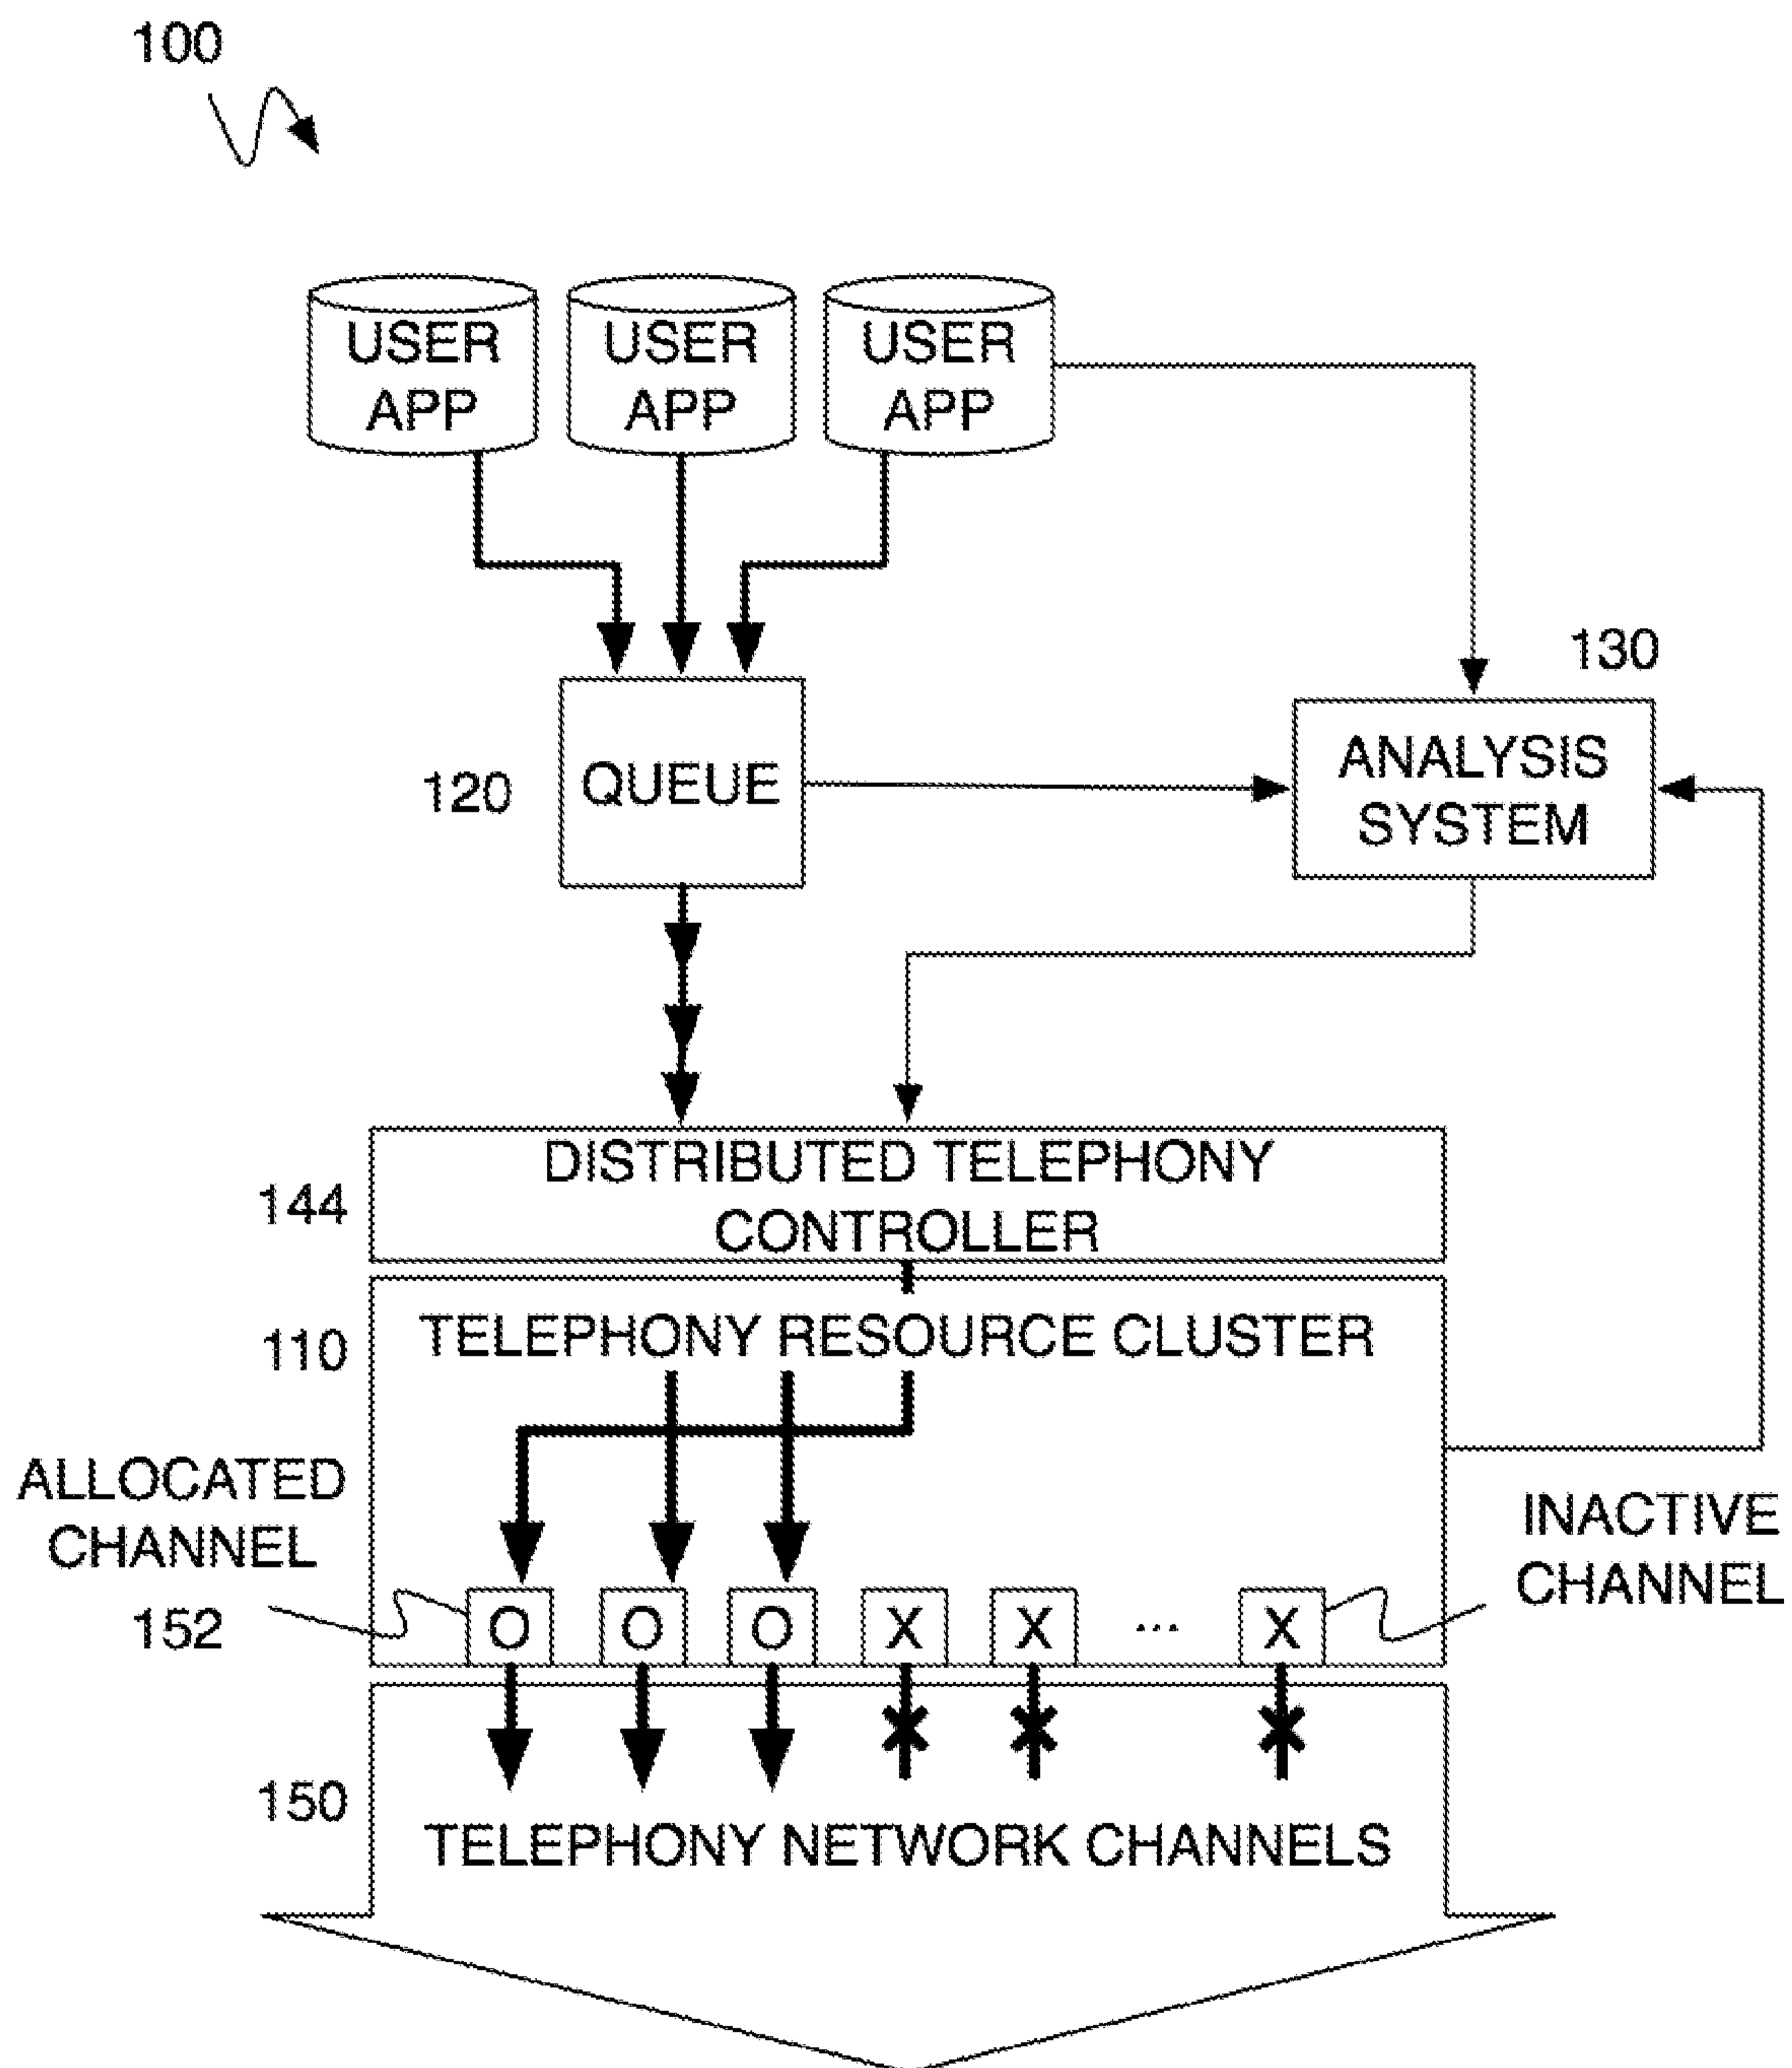


FIGURE 3

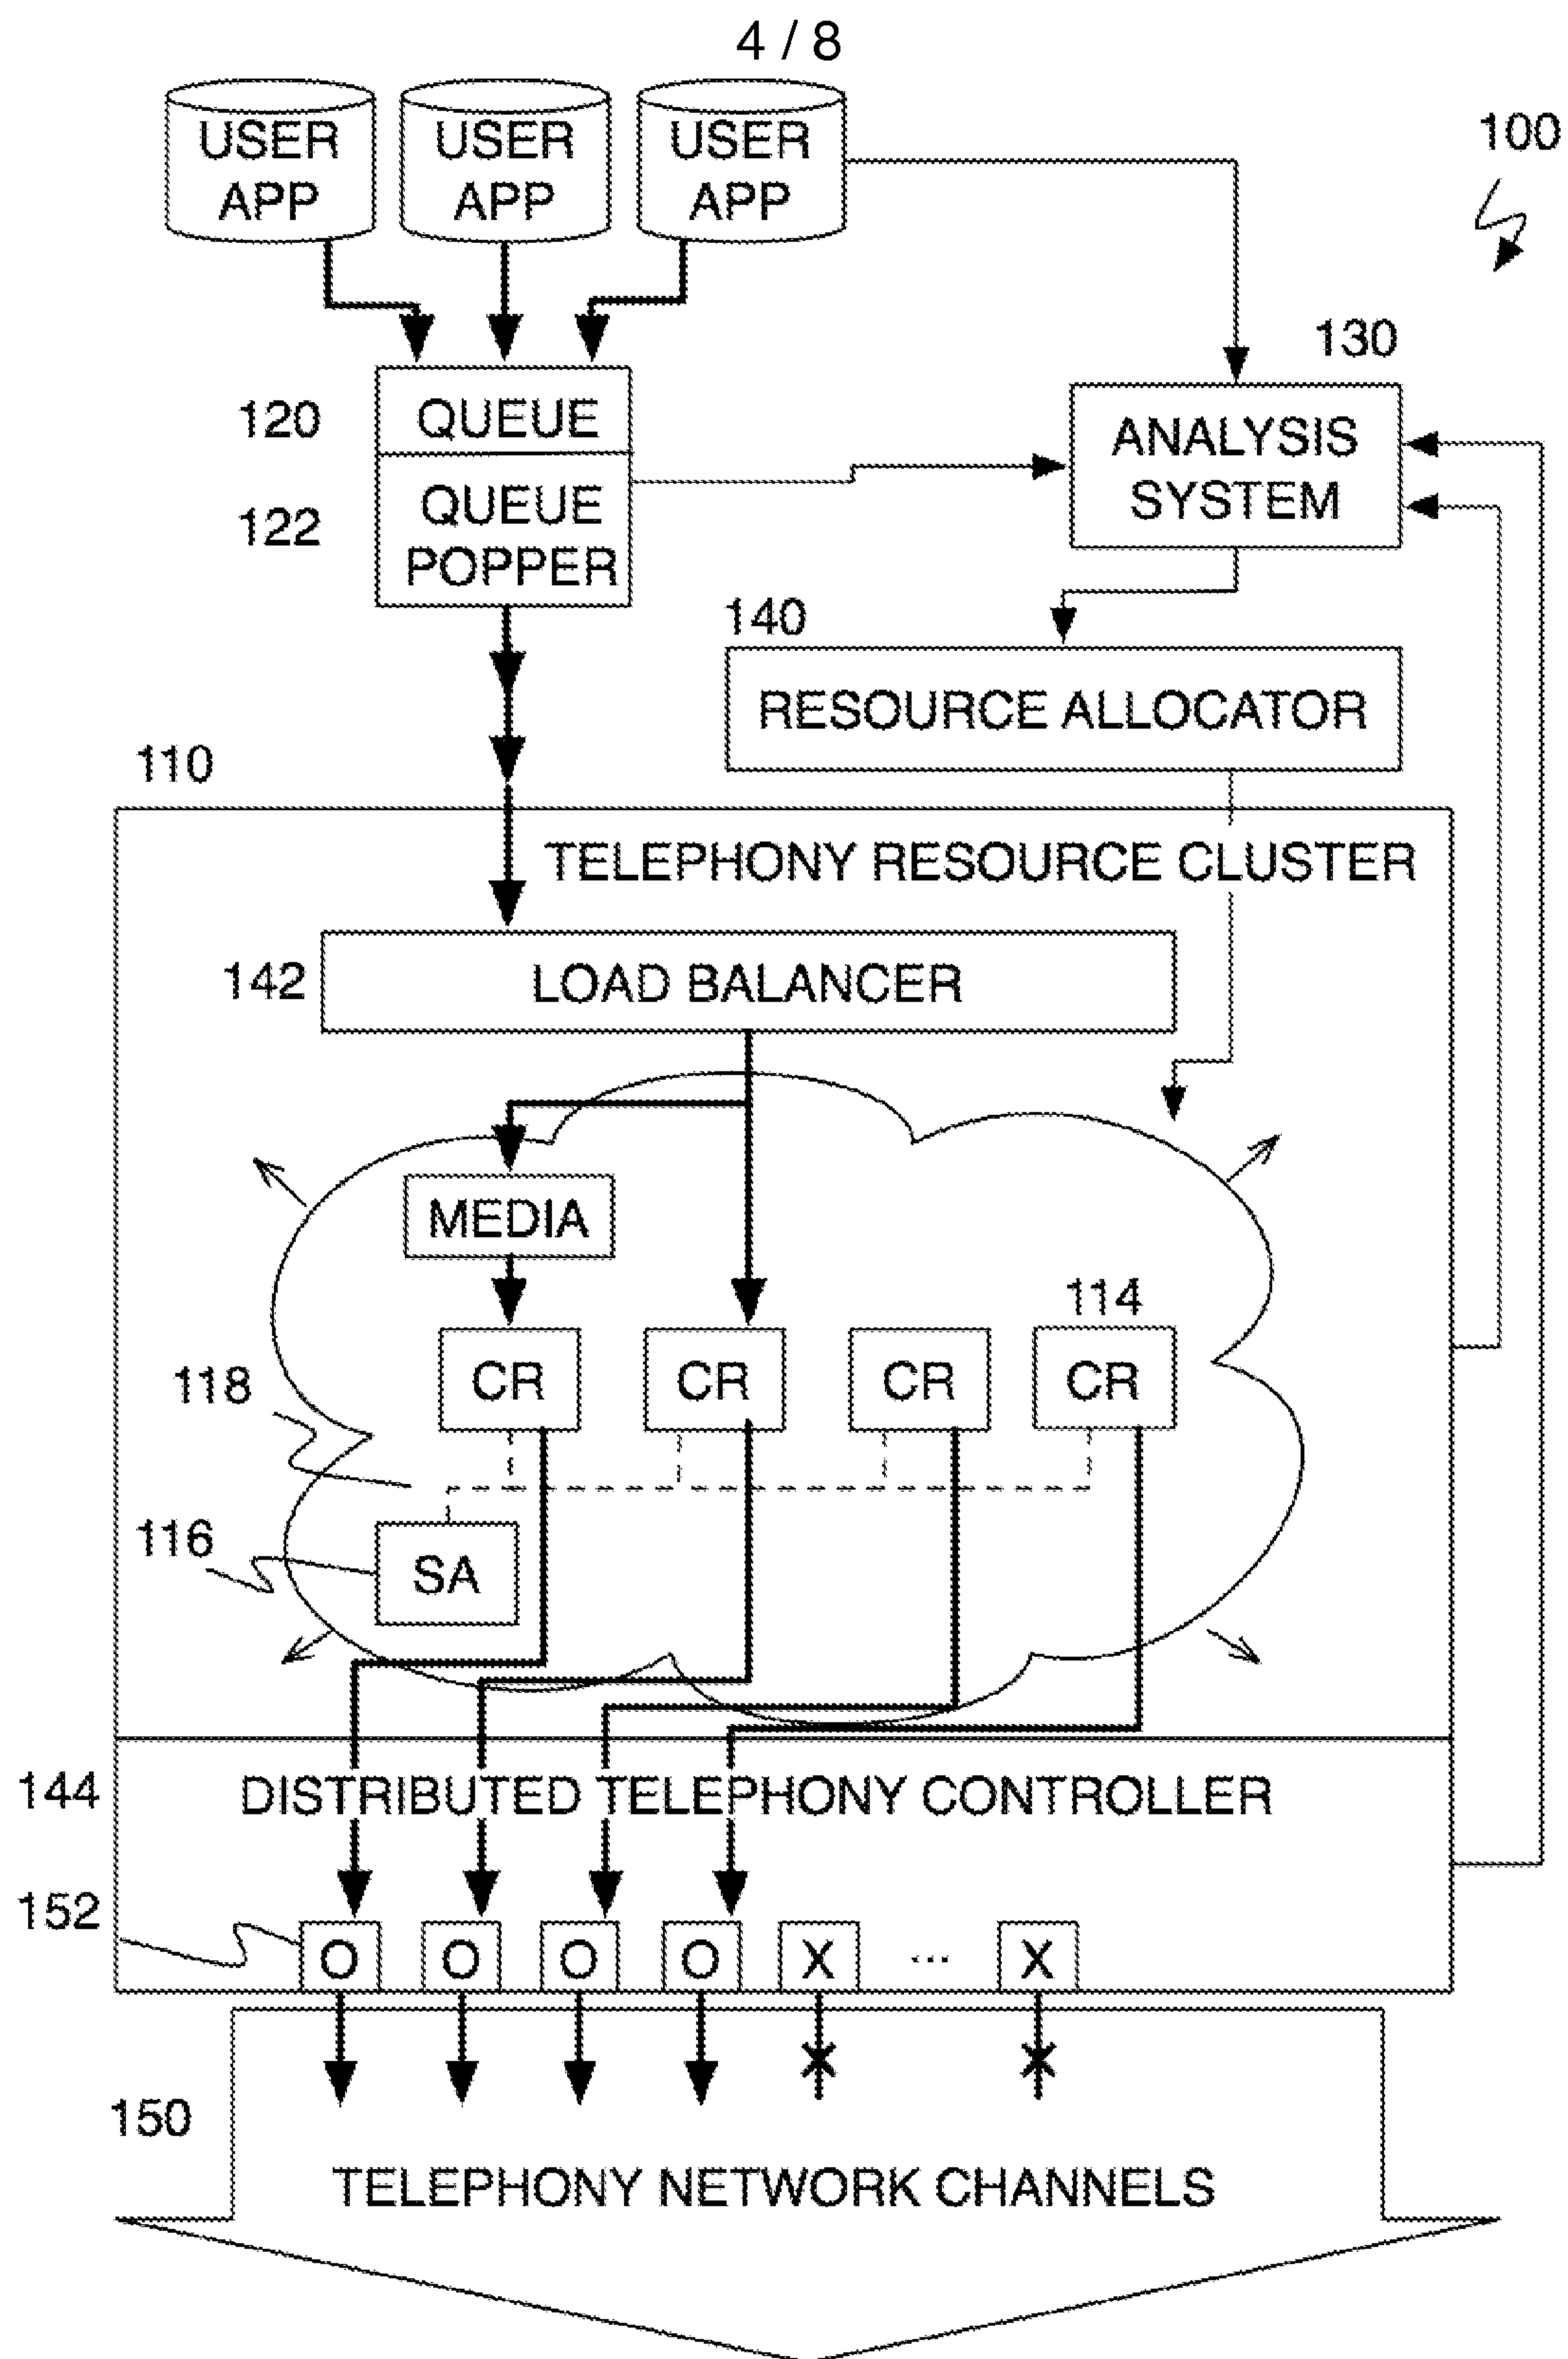


FIGURE 4

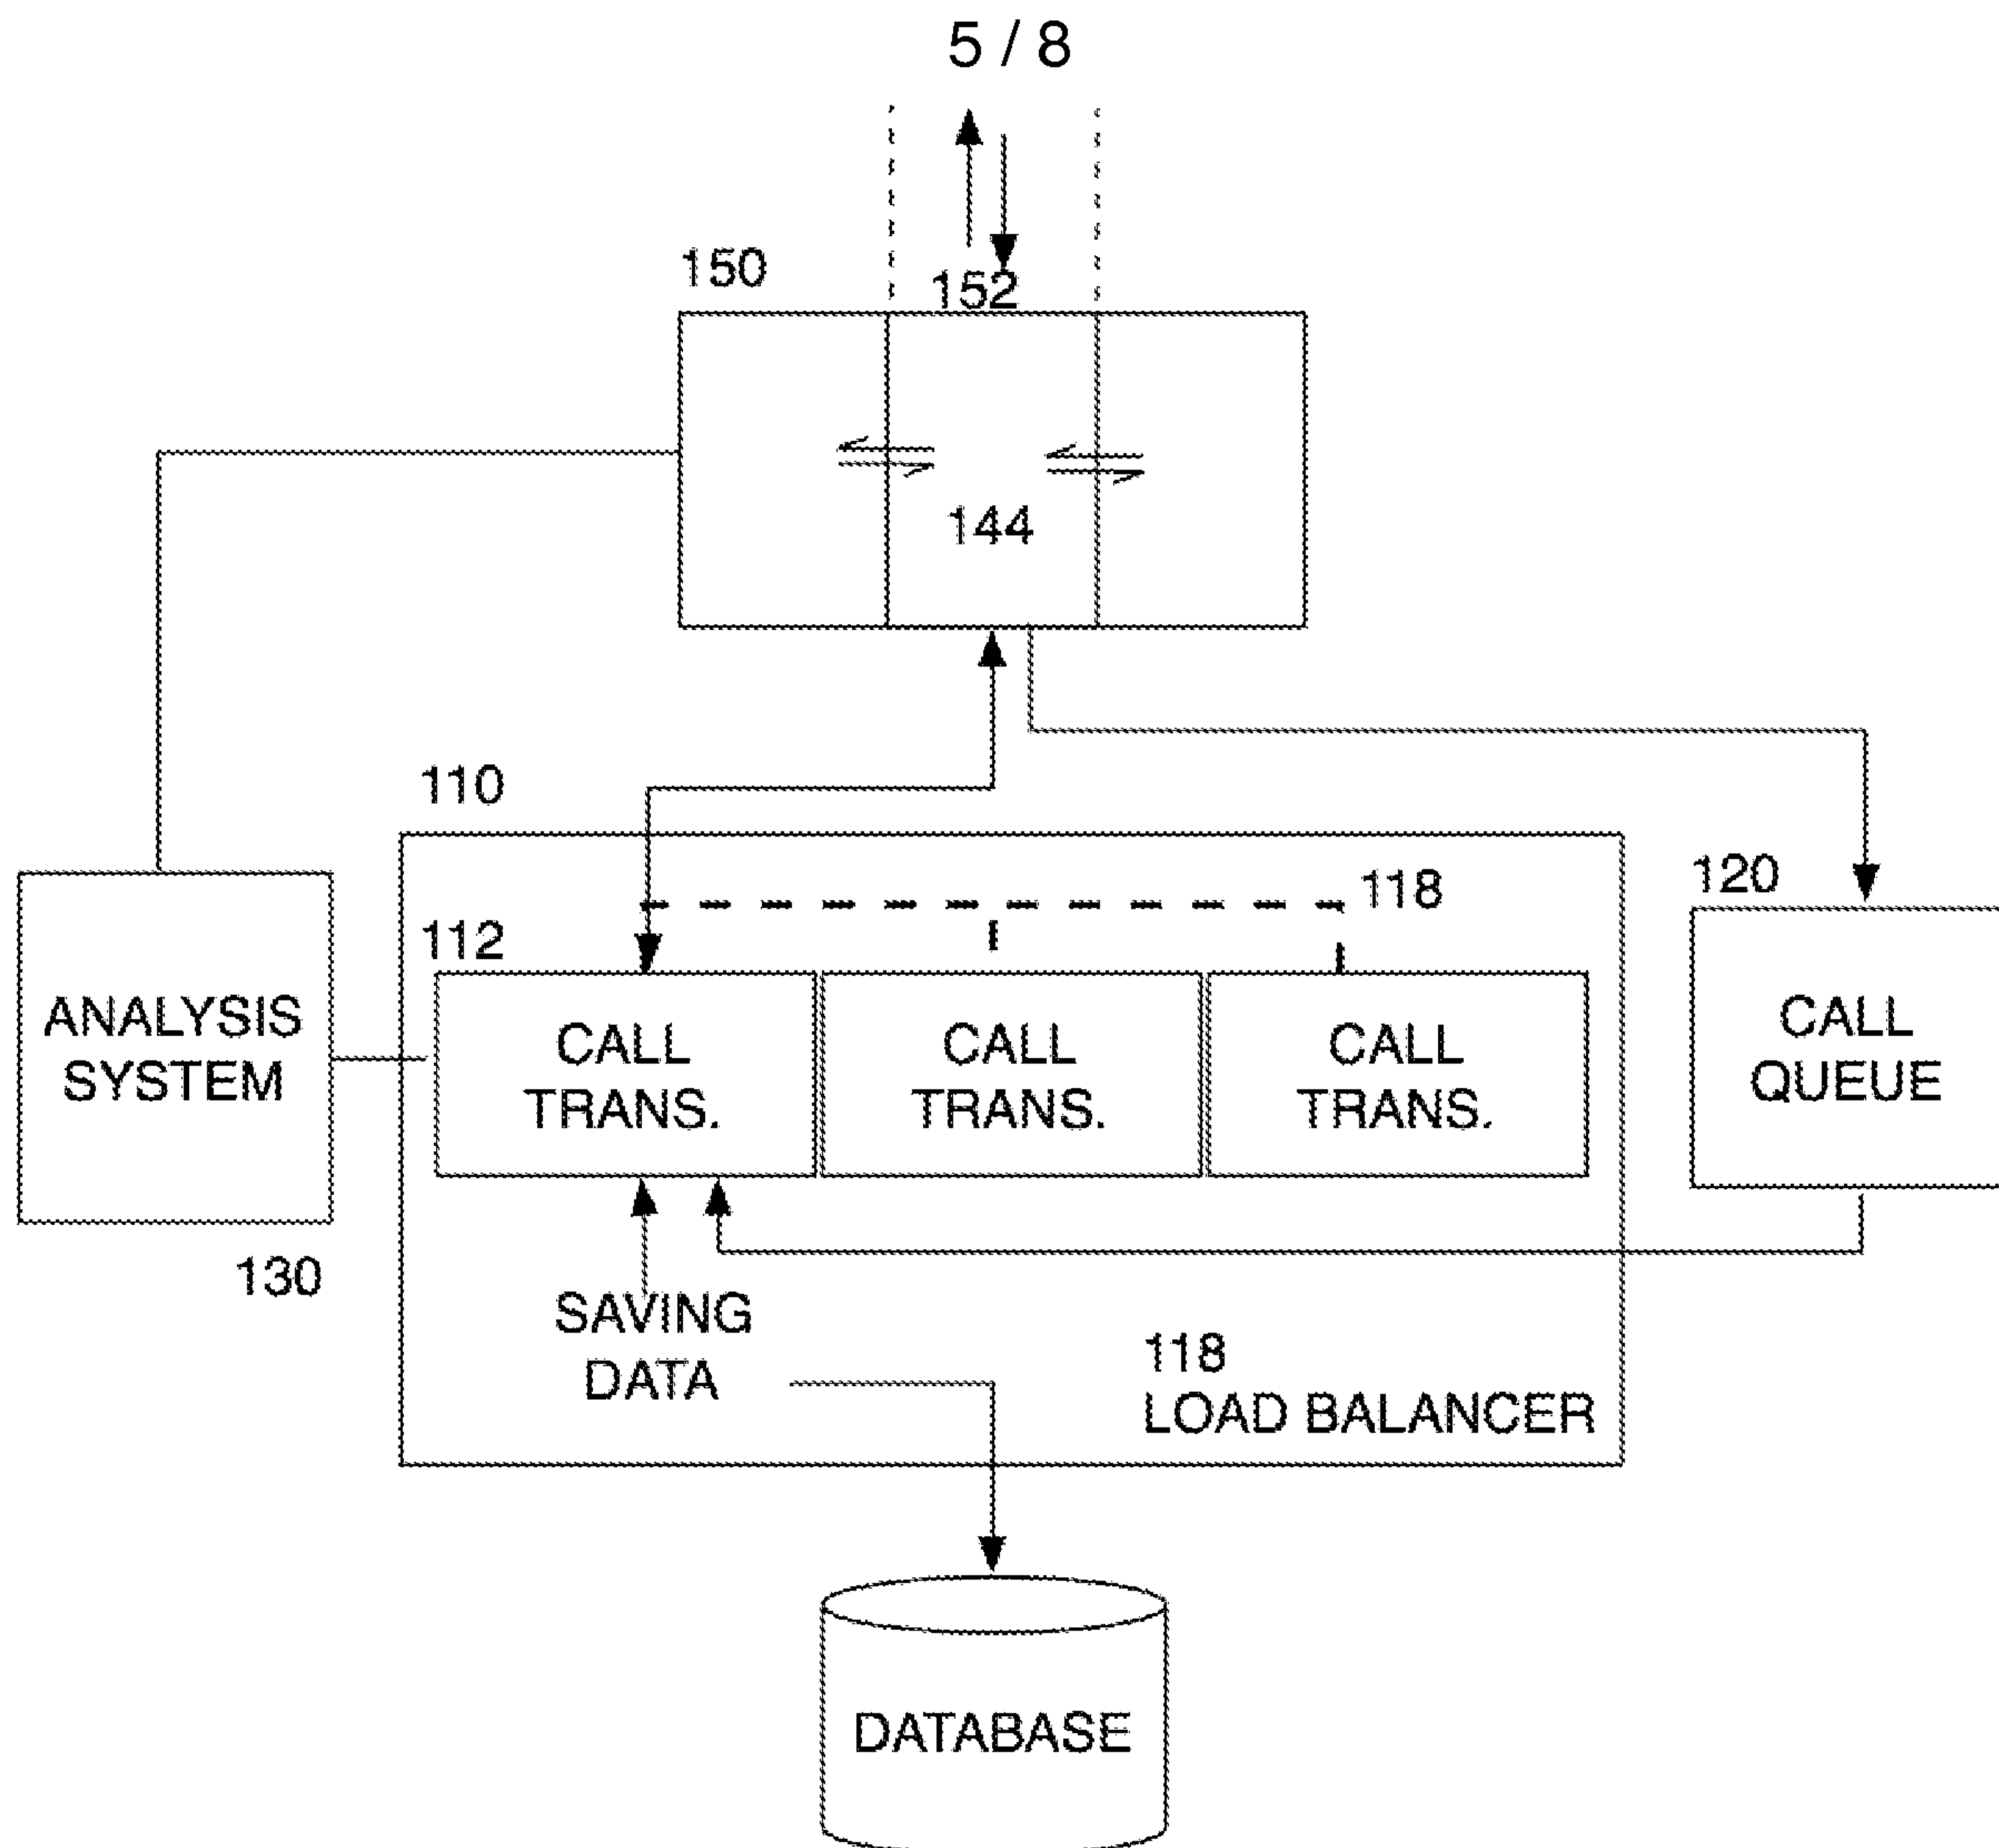


FIGURE 5

6 / 8

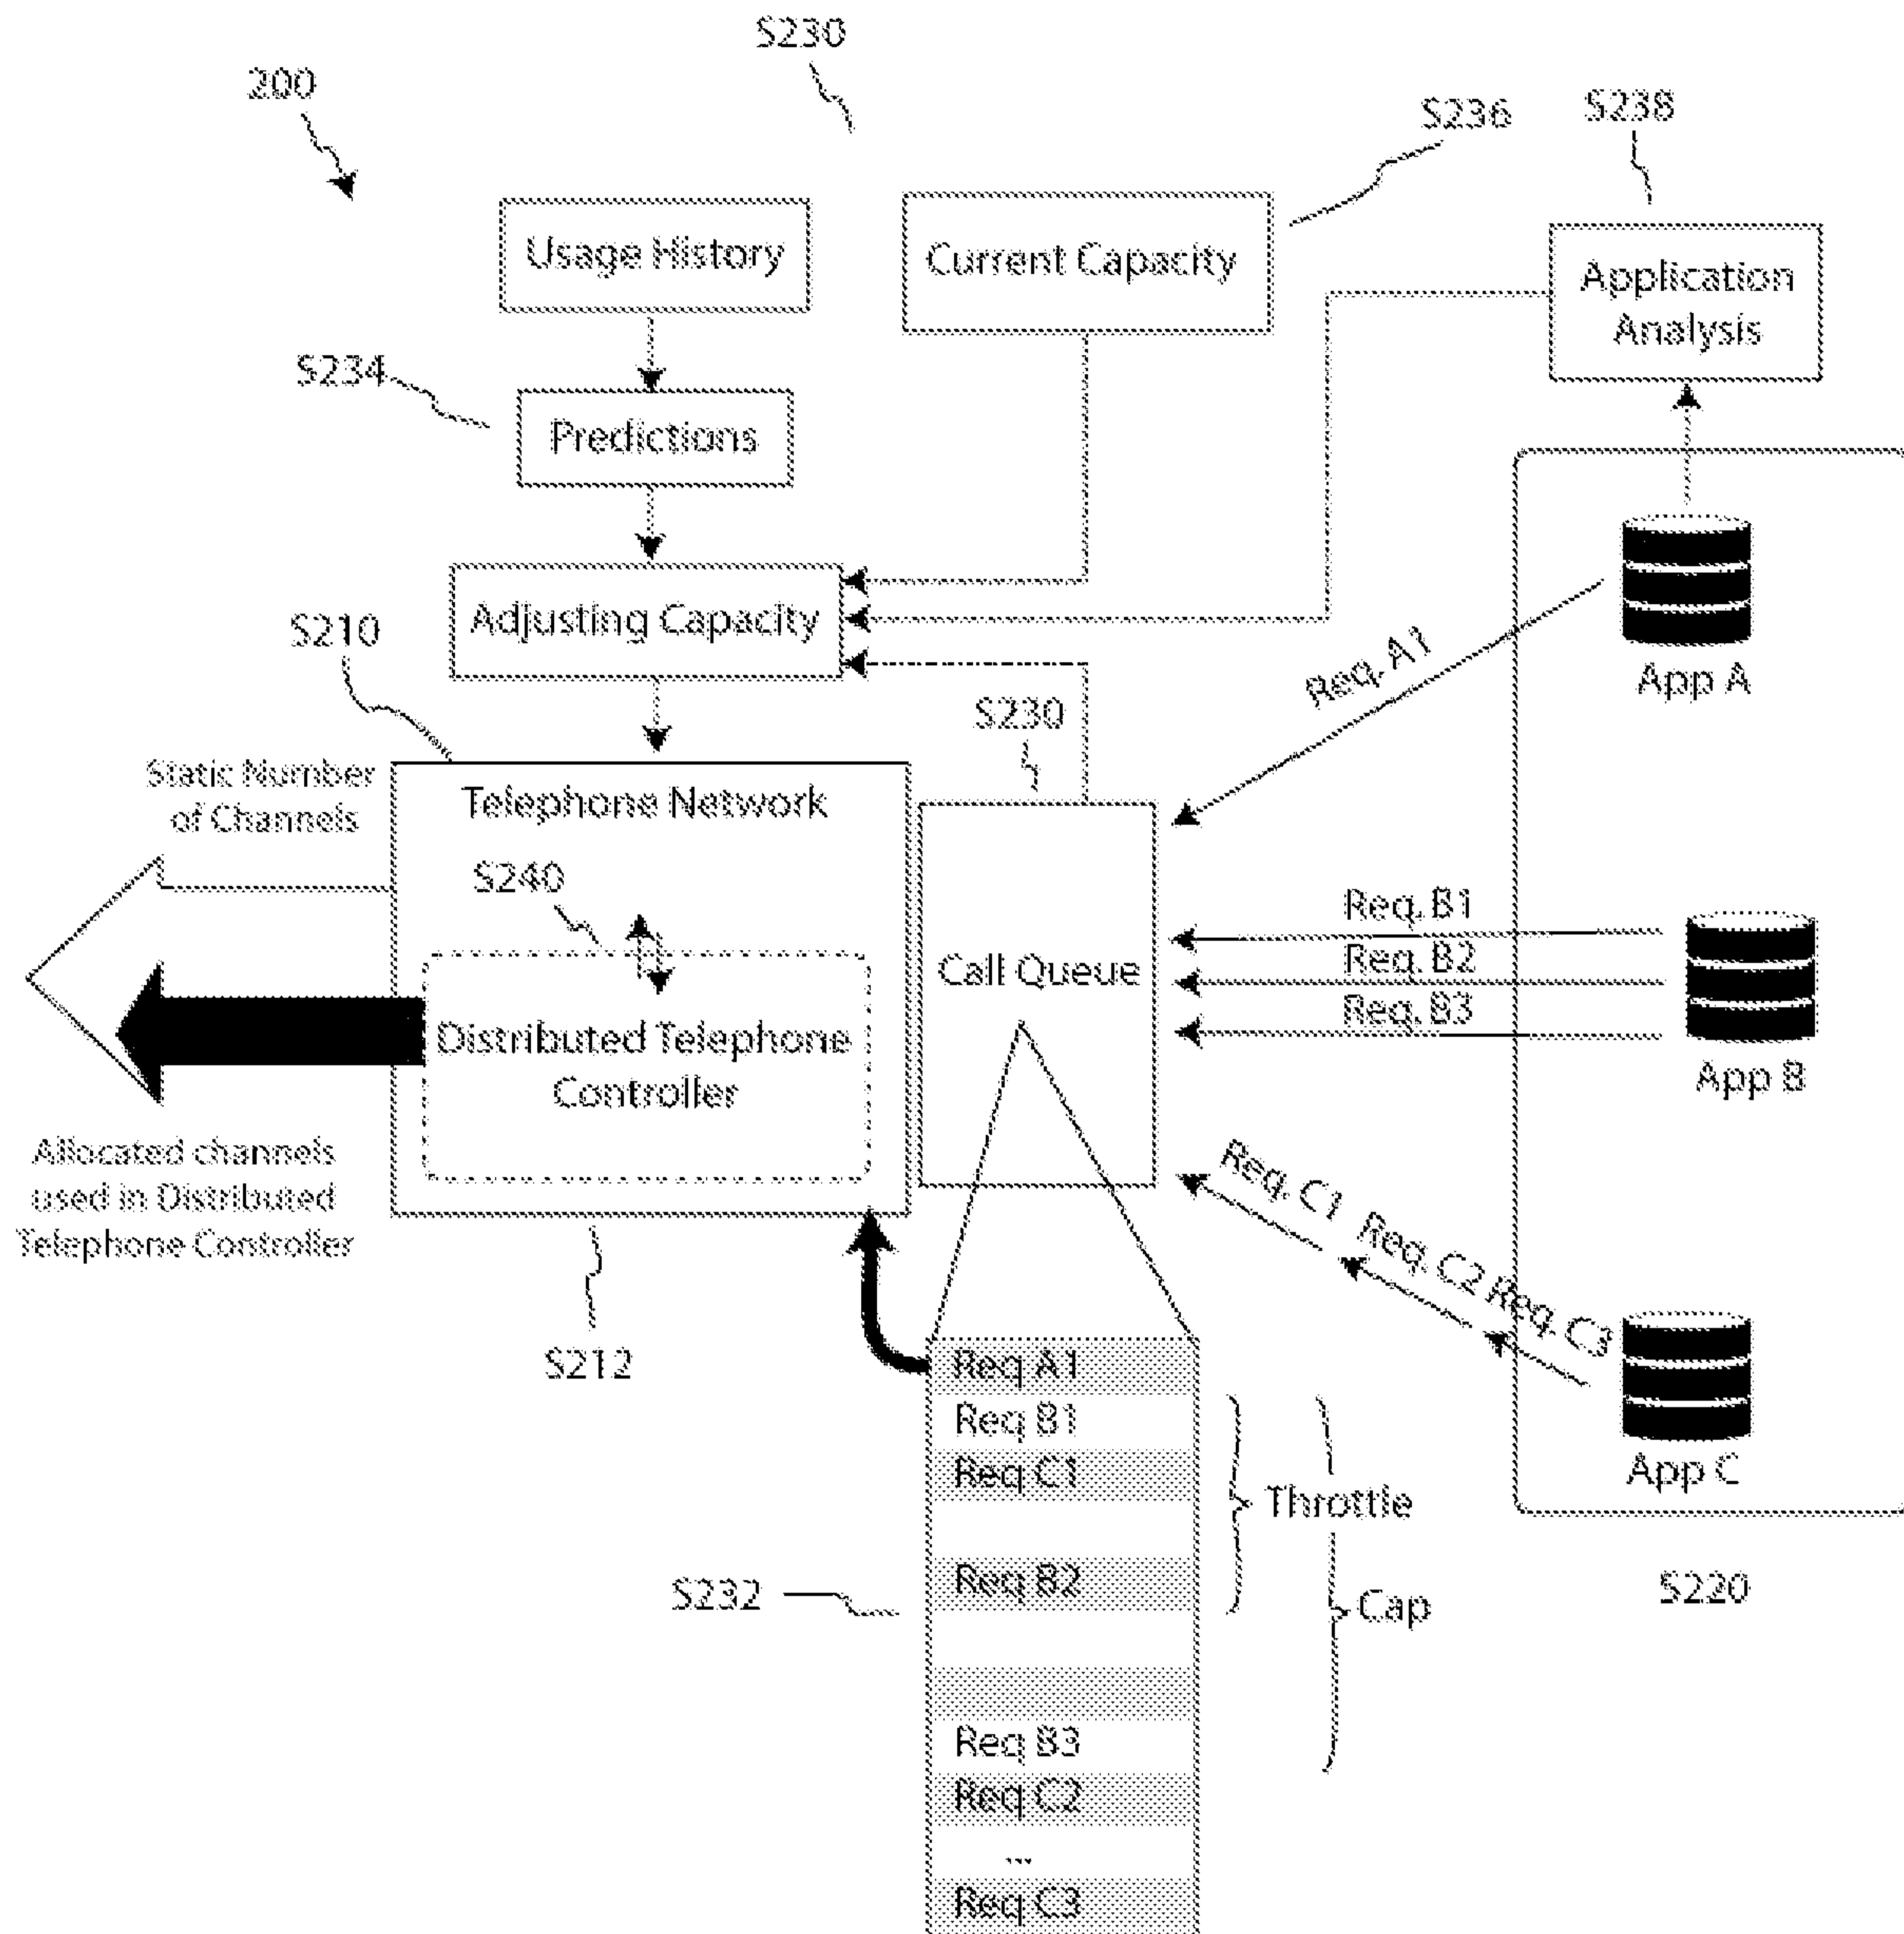


FIGURE 6

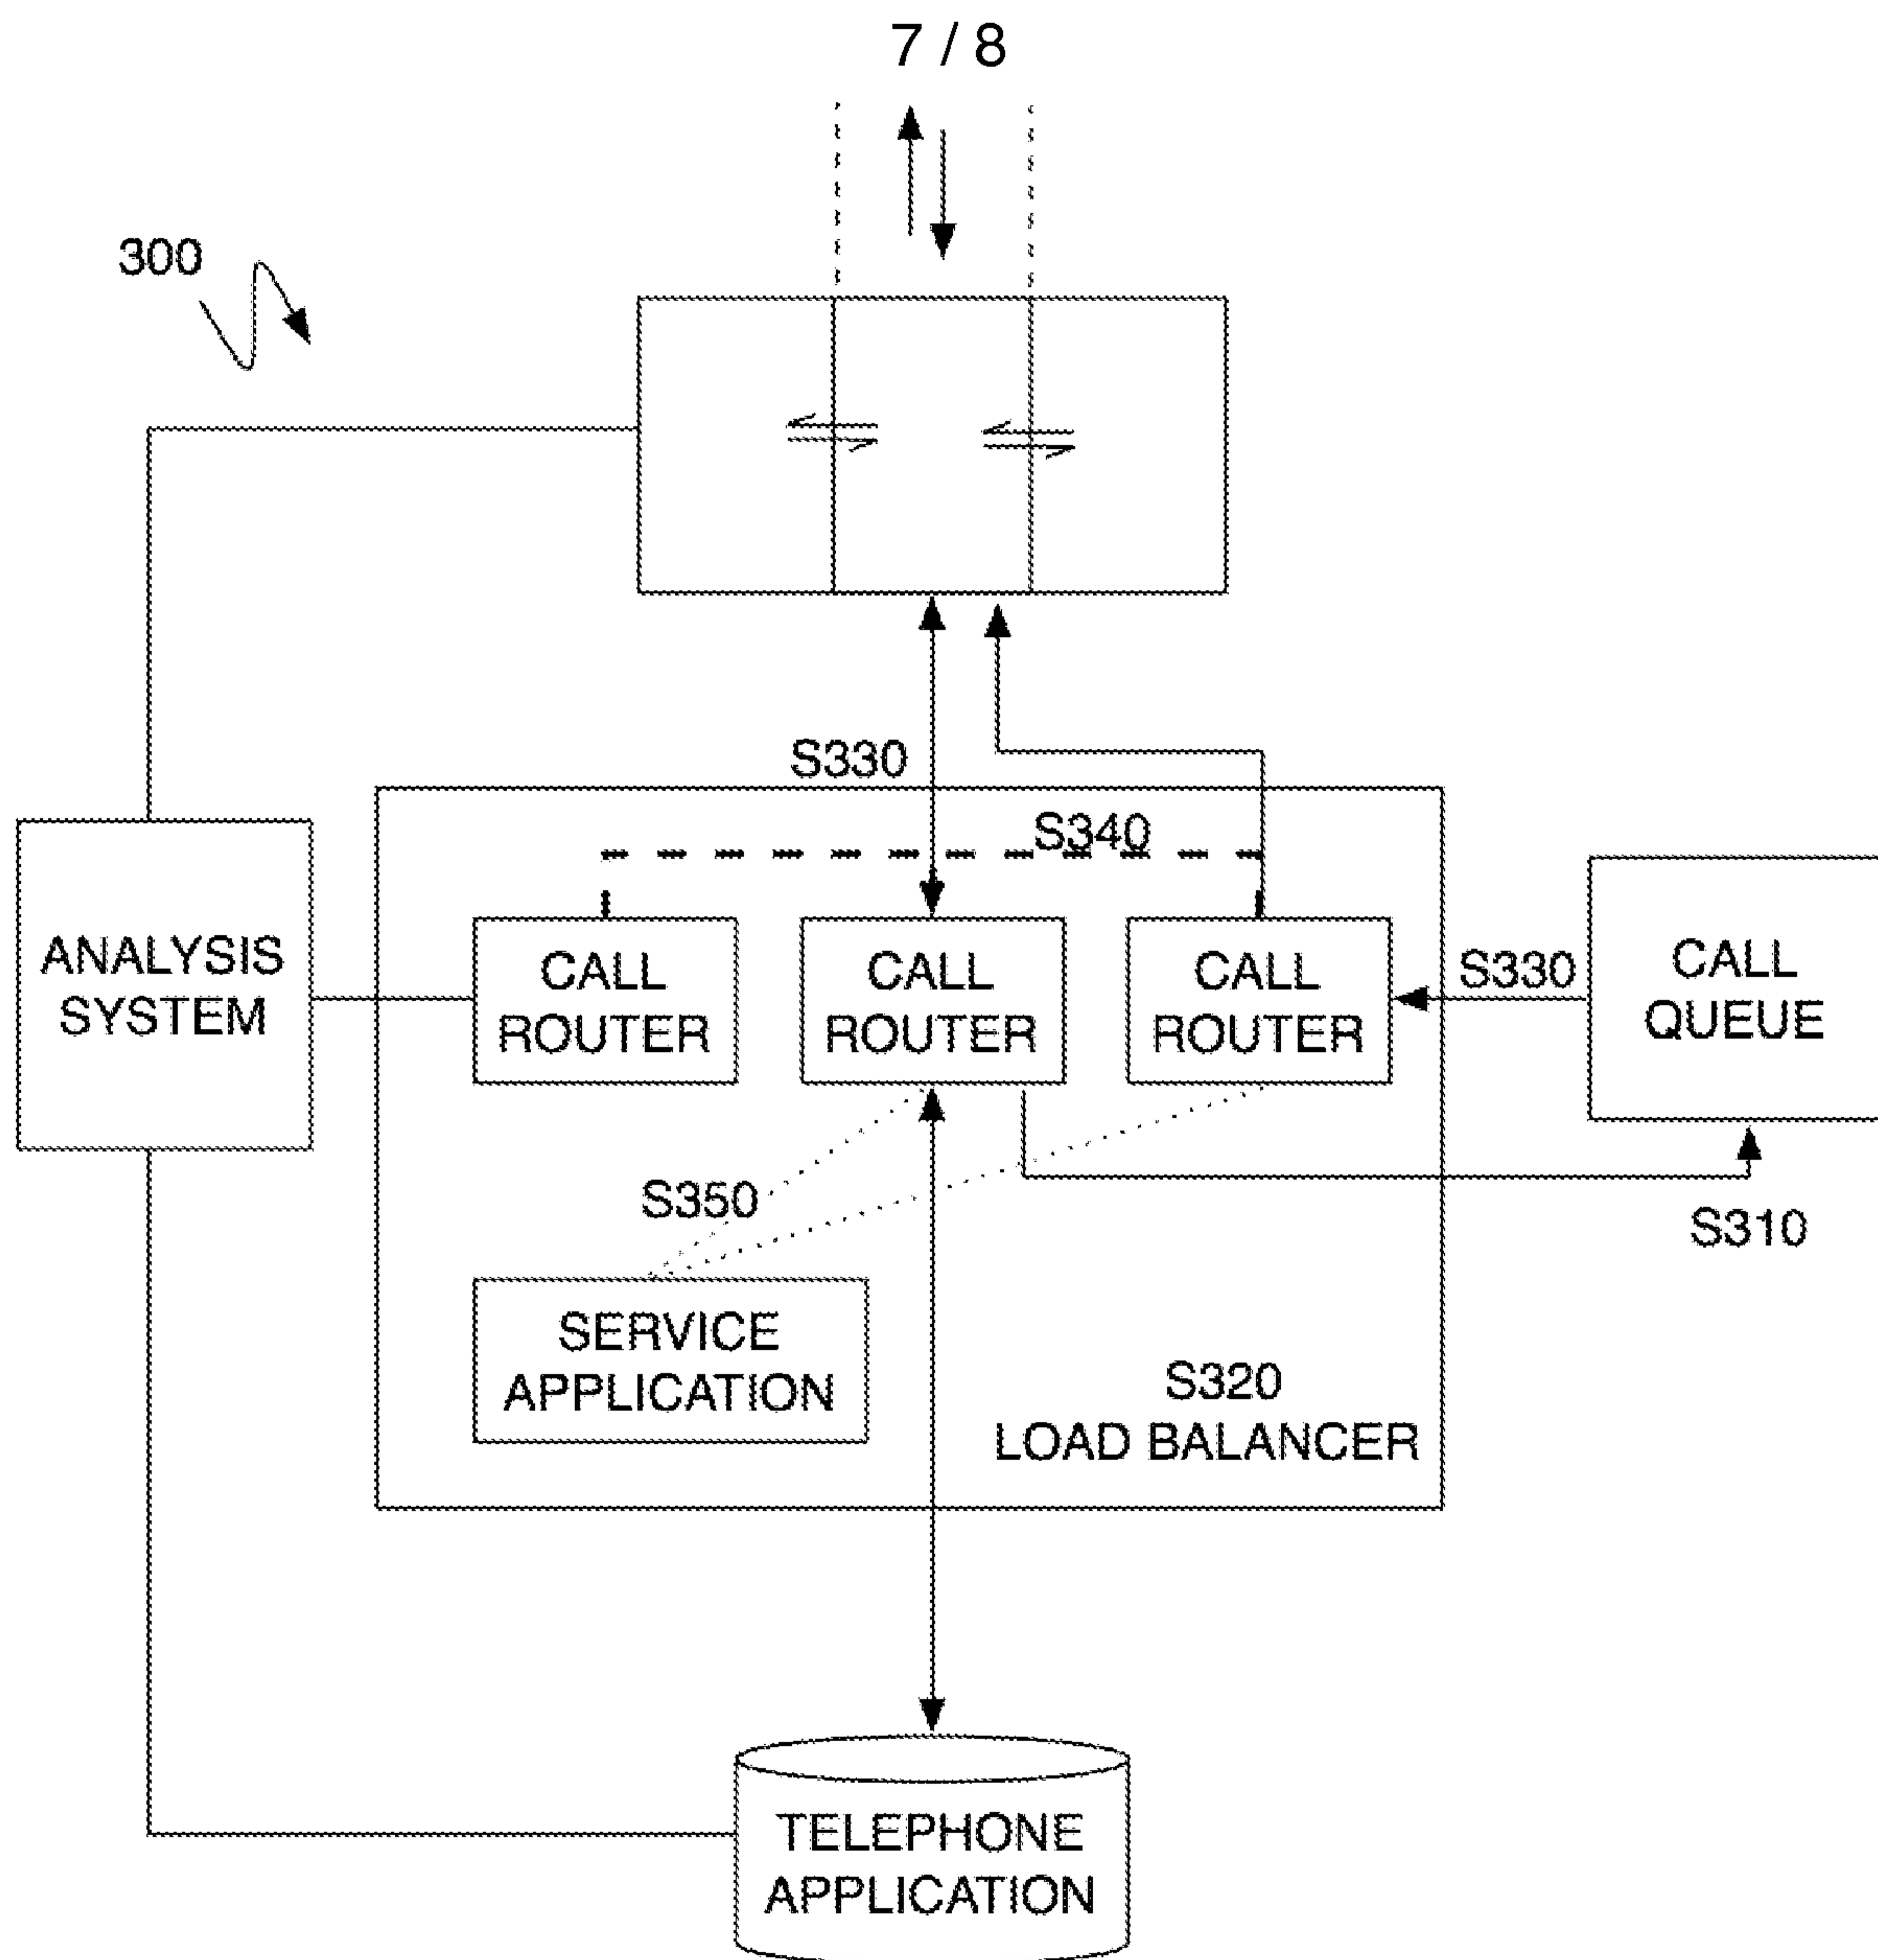


FIGURE 7

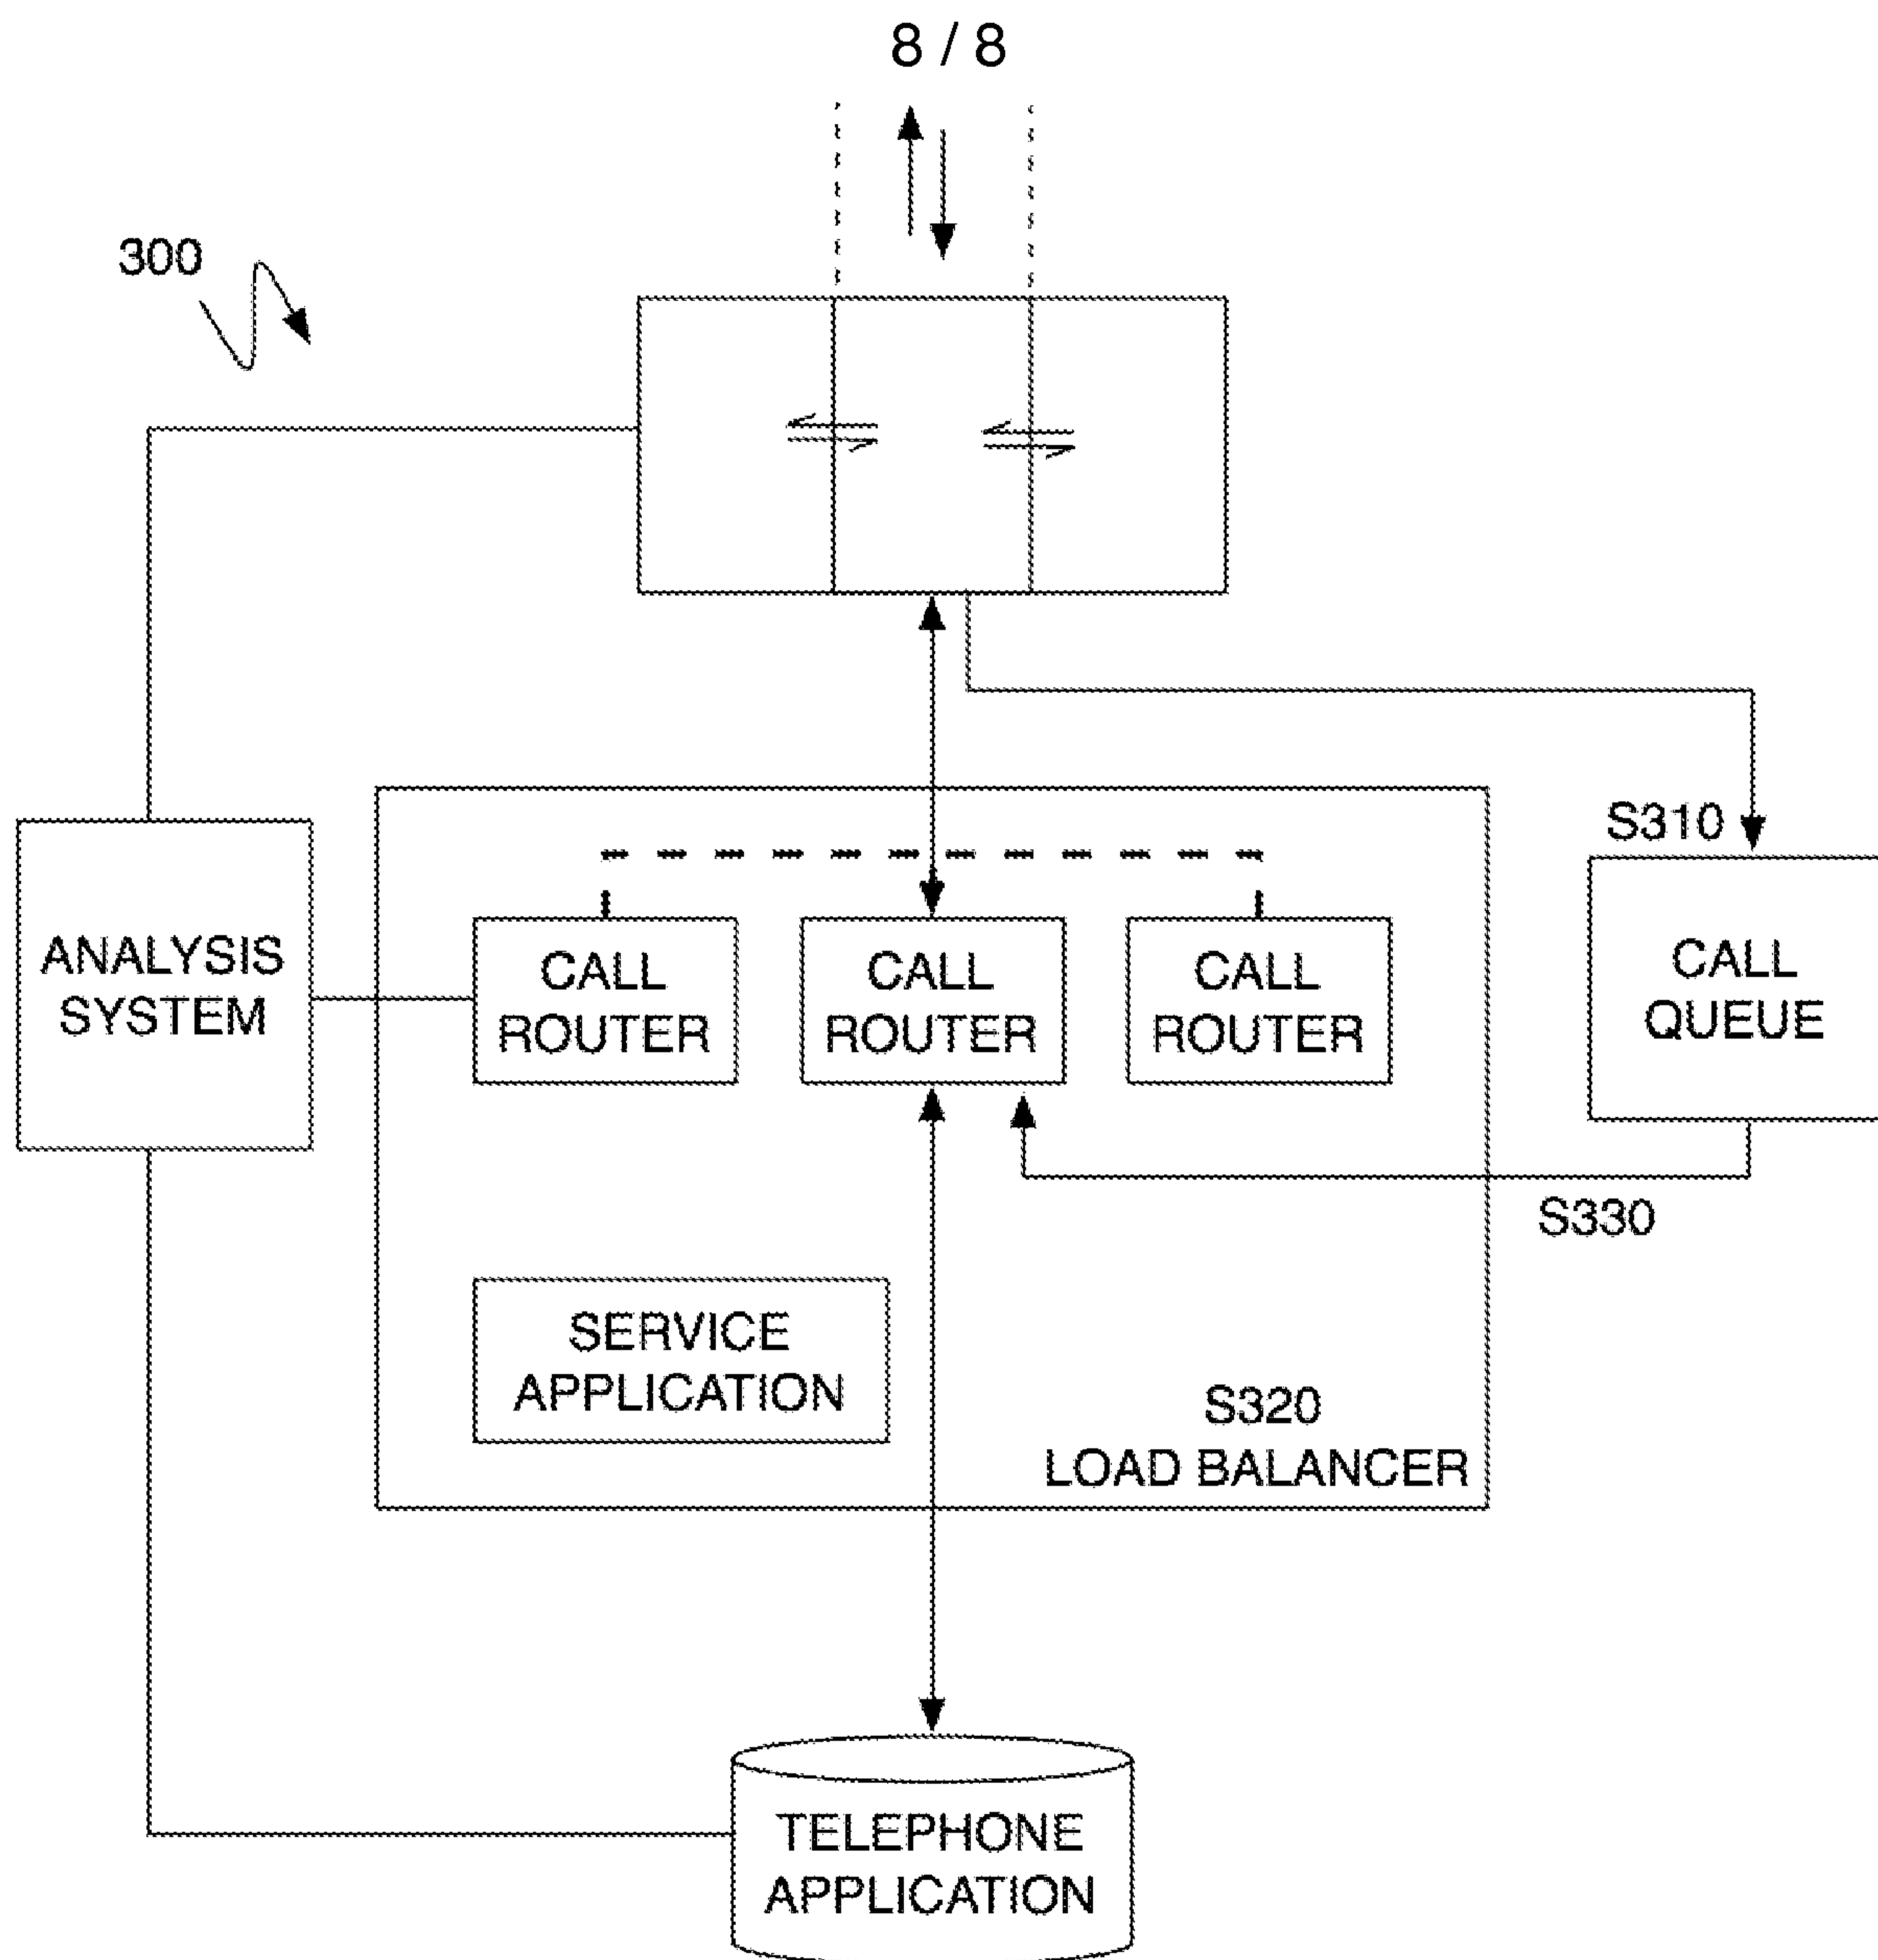


FIGURE 8

