Office de la Propriété Intellectuelle du Canada
Un organisme d'Industrie Canada

Canadian Intellectual Property Office
An agency of Industry Canada

CA 2776973 A1 2011/04/14

(21) 2 776 973

(12) DEMANDE DE BREVET CANADIEN
CANADIAN PATENT APPLICATION

(13) A1

(86) Date de dépôt PCT/PCT Filing Date: 2010/10/05
(87) Date publication PCT/PCT Publication Date: 2011/04/14
(85) Entrée phase nationale/National Entry: 2012/04/04
(86) N° demande PCT/PCT Application No.: US 2010/051540
(87) N° publication PCT/PCT Publication No.: 2011/044174
(30) Priorité/Priority: 2009/10/05 (US61/248,593)

(51) Cl.Int./Int.Cl. *G06Q 30/00* (2012.01)

(71) Demandeur/Applicant:
CALLSPACE, INC., US

(72) Inventeurs/Inventors:
MANLEY, RICHARD J., US;
FRYAR, A. CRAIG, US;
LEE, MICHAEL, US;
SIMON, CARMICHAEL J., US

(74) Agent: GOWLING LAFLEUR HENDERSON LLP

(54) Titre : GESTION D'UNE OFFRE DE TELEPHONIE CONTEXTUALISEE
(54) Title: CONTEXTUALIZED TELEPHONY MESSAGE MANAGEMENT

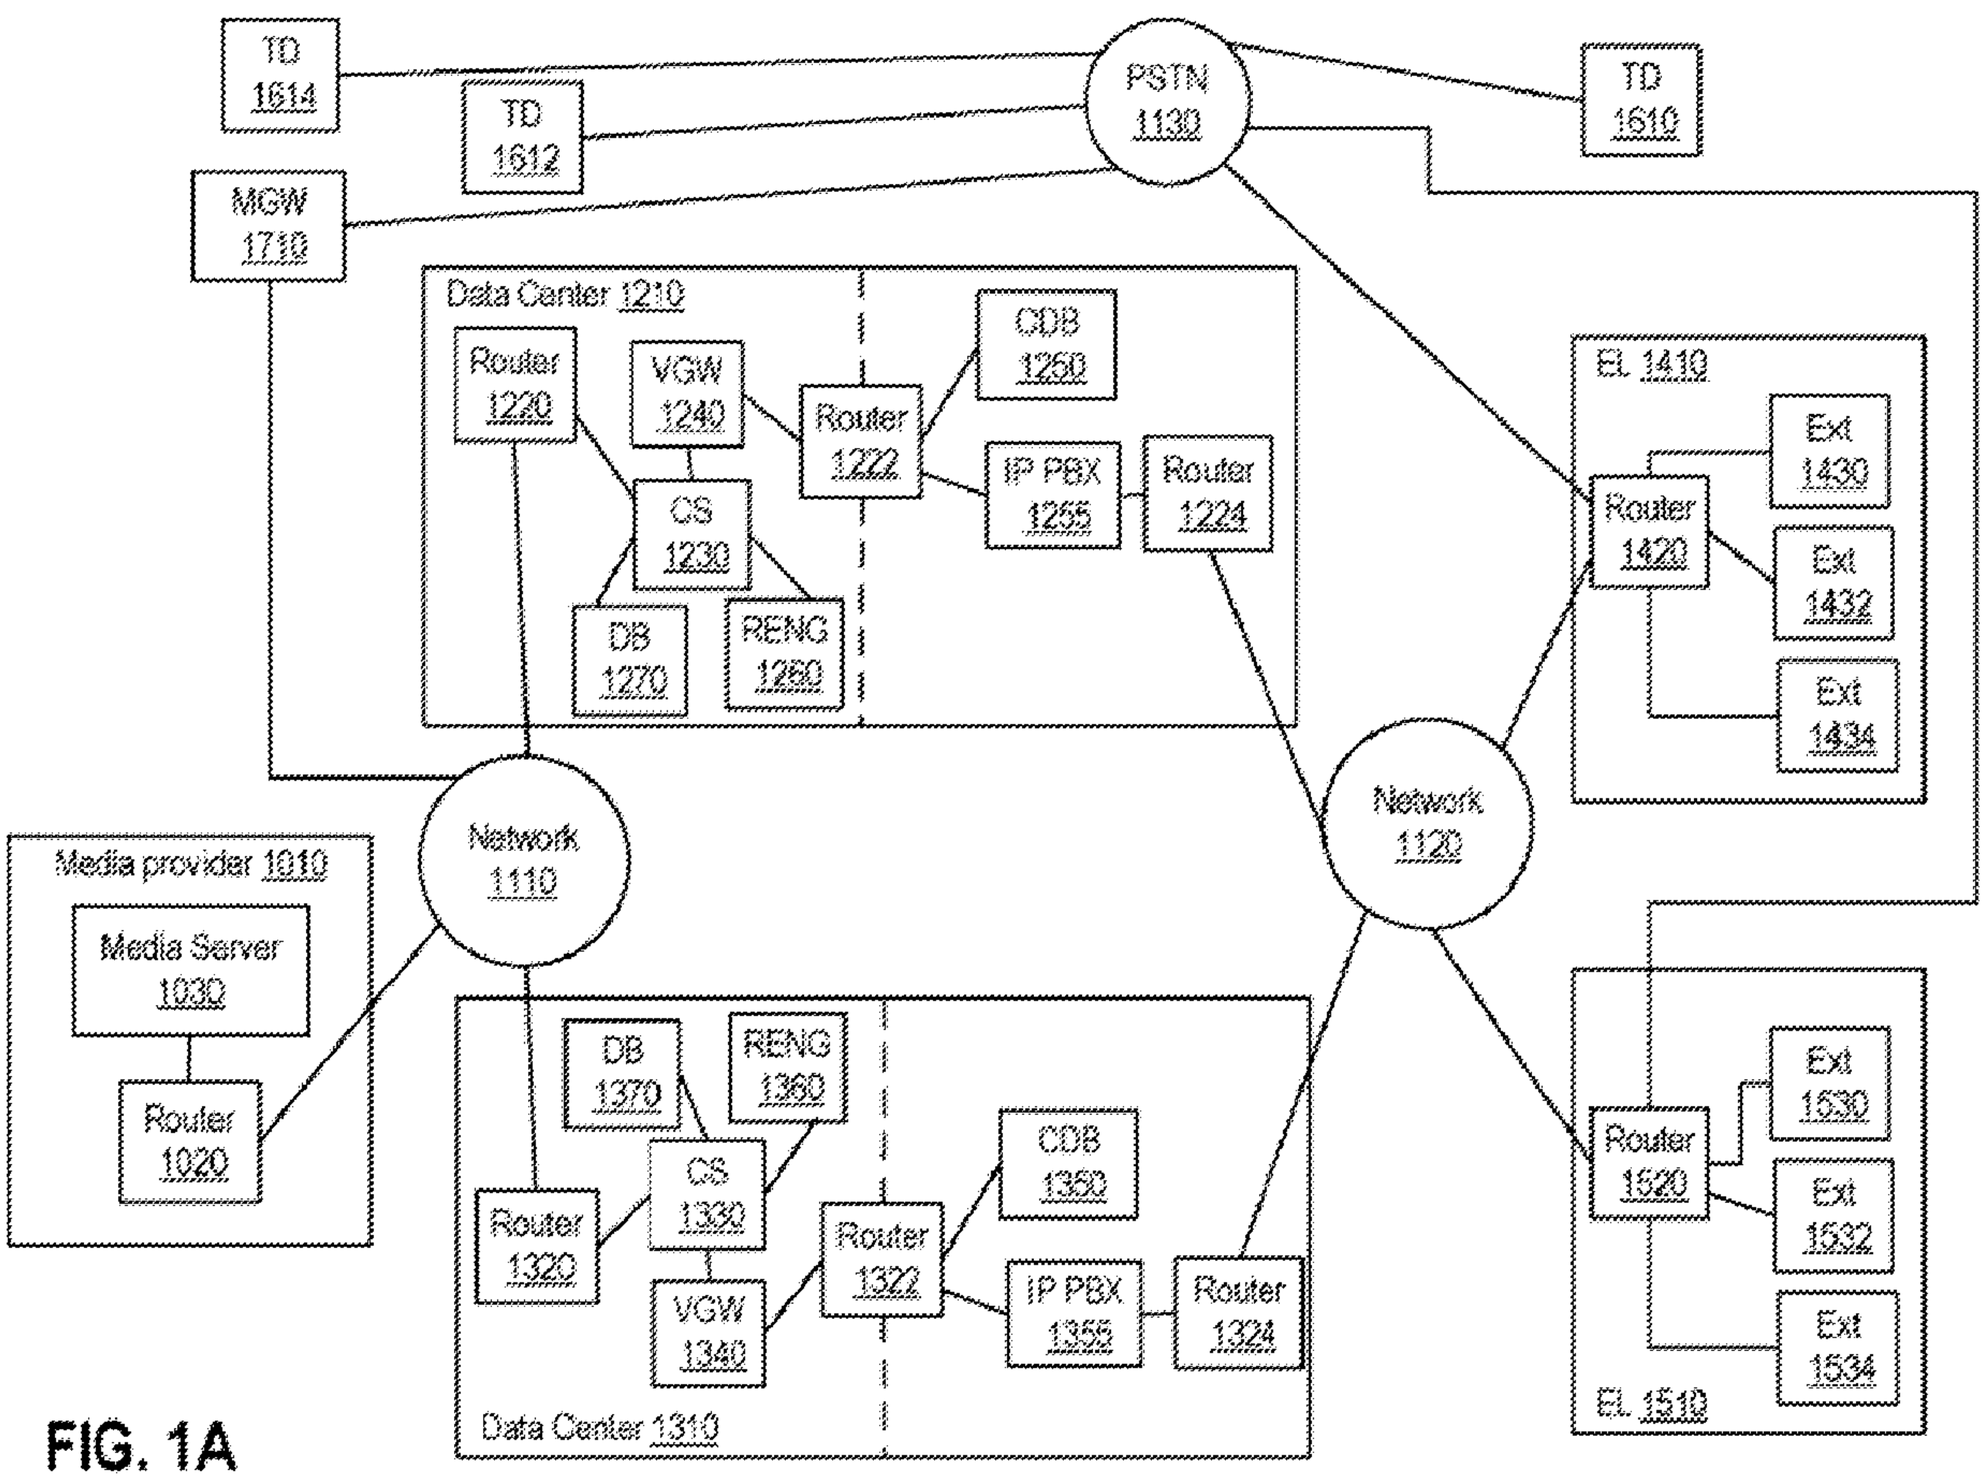

FIG. 1A

(57) Abrégé/Abstract:
In one or more embodiments, one or more methods and/or systems described can transform an inbound call into one or more call streams and/or call states that can include at least one of a contextualized or personalized message, a promotion, a coupon, an

CA 2776973 A1 2011/04/14

(21) 2 776 973

(13) A1

(57) Abrégé(suite)/Abstract(continued):
offer, a voucher, an advertisement, and an opt-in program, among others. For example, the one or more methods and/or systems described can perform: receiving identification information associated with a telephony device; determining, based on the identification information, a message for the user; and sending the message to the telephony device. For instance, the message can include a coupon or discount for a good or service. In one example, the coupon or discount can be sent via a short message service text message. In another example, the coupon or discount can include a computer-readable image that can be sent via a multimedia messaging service message.

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

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
14 April 2011 (14.04.2011)

PCT

(10) International Publication Number
WO 2011/044174 A1

(51) International Patent Classification:
G06Q 30/00 (2006.01)

(21) International Application Number:
PCT/US2010/051540

(22) International Filing Date:
5 October 2010 (05.10.2010)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
61/248,593 5 October 2009 (05.10.2009) US

(71) Applicant *(for all designated States except US)*: **CALLSPACE, INC** [US/US]; 13508 Bullick Hollow Rd., Austin, TX 78726 (US).

(72) Inventors; and
(75) Inventors/Applicants *(for US only)*: **MANLEY, Richard, J.** [US/US]; 13508 Bullick Hollow Road, Austin, TX 78726 (US). **FRYAR, A., Craig** [US/US]; 508 Kodiak Trail, Cedar Park, TX 78613 (US). **LEE, Michael** [US/US]; 614 Paden Drive, Cedar Park, TX 76813 (US). **SIMON, Carmichael, J.** [US/US]; 1818 Golden Arrow Avenue, Cedar Park, TX 78613 (US).

(74) Agent: **YUDELL, Craig, J.**; Dillon & Yudell LLP, 8911 N. Capital of Texas Hwy., Suite 2110, Austin, TX 78759 (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, LR, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

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

[Continued on next page]

(54) **Title**: CONTEXTUALIZED TELEPHONY MESSAGE MANAGEMENT

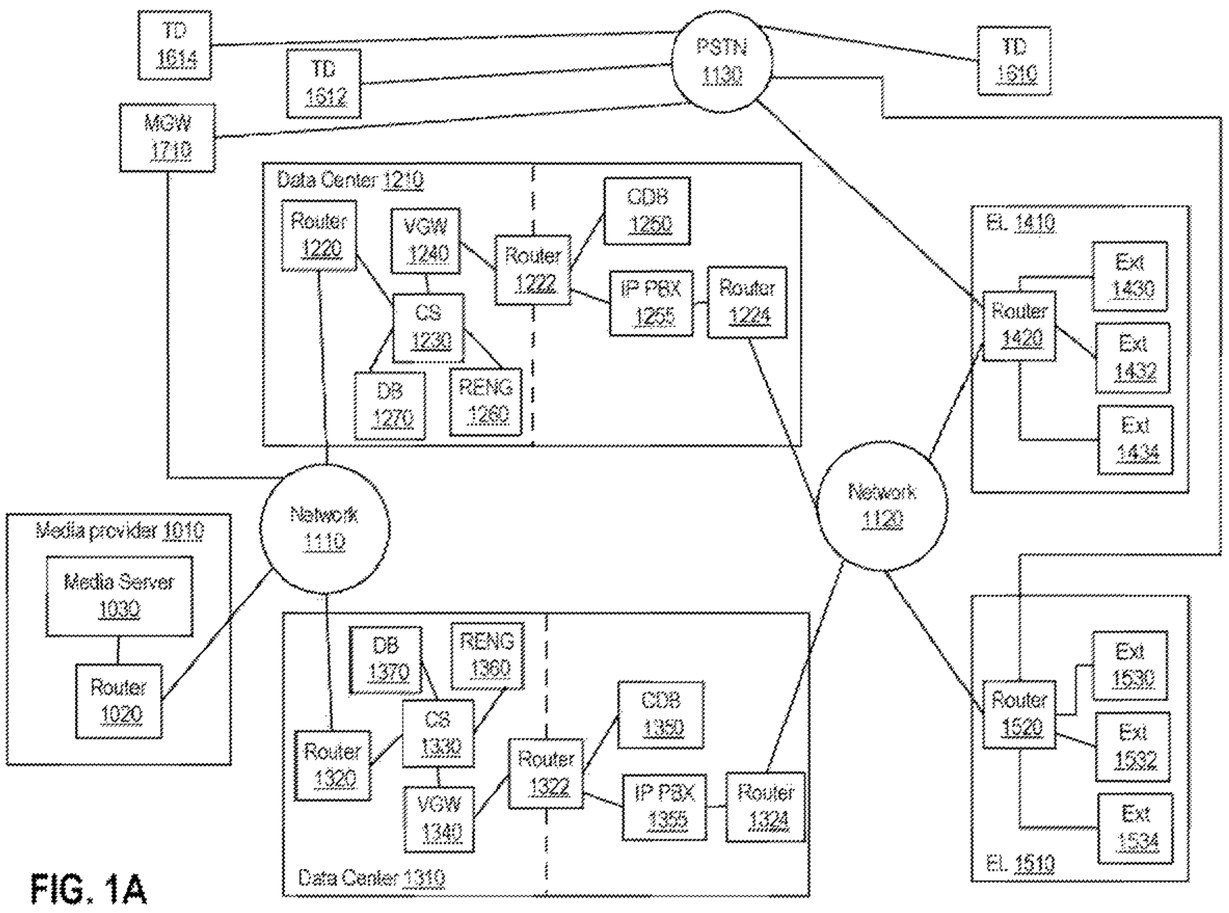

FIG. 1A

(57) **Abstract**: In one or more embodiments, one or more methods and/or systems described can transform an inbound call into one or more call streams and/or call states that can include at least one of a contextualized or personalized message, a promotion, a coupon, an offer, a voucher, an advertisement, and an opt-in program, among others. For example, the one or more methods and/or systems described can perform: receiving identification information associated with a telephony device; determining, based on the identification information, a message for the user; and sending the message to the telephony device. For instance, the message can include a coupon or discount for a good or service. In one example, the coupon or discount can be sent via a short message service text message. In another example, the coupon or discount can include a computer-readable image that can be sent via a multimedia messaging service message.

WO 2011/044174 A1

— *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))*

CONTEXTUALIZED TELEPHONY MESSAGE MANAGEMENT

Inventors:

Richard J. Manley, Craig Fryar, Michael Lee, Carmichael Simon, and William C. Archibald

[0001] This application claims benefit of U.S. Provisional Application Ser. No. 61/248,593, filed October 5, 2009, titled "Method and Apparatus For Syndicated Streaming Multimedia Messaging Network", which application is hereby incorporated by reference in its entirety as though fully and completely set forth herein.

BACKGROUND

1. Technical Field

[0002] This description relates generally to the field of managing calls within a telephony management system, and, more specifically, to dynamically providing call support services within a telephony management system.

2. Description of the Related Art

[0003] Callers place calls to entities (e.g., call centers, restaurants, government agencies, businesses, educational institutions, retailers, etc.), in order to obtain information and/or to speak to representatives within departments of an entity. In the past, callers have spent billions of minutes per day listening to silence or pre-recorded sets of repeating messages or music. Callers are subjected to navigating prompts, often outdated, presented to them by impersonal interactive voice response (IVR) telephony systems. This process is inefficient and frustrating to callers and is equally unsatisfactory to management of many of these entities. Entities have spent billions of dollars in promoting, communicating, and developing their product, service, or brand messaging in traditional marketing channels, but those entities have overlooked and/or missed various opportunities which remains untapped in the millions of calls flowing into their entities every day.

BRIEF DESCRIPTION OF THE DRAWINGS

[0004] The preferred embodiments will become apparent upon reading the following detailed description and upon reference to the accompanying drawings in which:

[0005] FIGs. 1A and 1B illustrate block diagrams of one or more network communication systems, according to one or more embodiments;

[0006] FIG. 2 illustrates a block diagram of states of a telephone call within a telephony management system, according to one or more embodiments;

[0007] FIG. 3 illustrates a method of sending a message to a telephony device;

[0008] FIG. 4 illustrates a method of transferring a call from a telephony device to an extension within the telephony management system;

[0009] FIG. 5 illustrates a method of providing a message based at least on one or more submenu choices;

[0010] FIG. 6 illustrates a method of operating a telephony management system;

[0011] FIG. 7 illustrates an exemplary block diagram of an entity location that receives messages from users' telephony devices, according to one or more embodiments;

[0012] FIGs. 8A and 8B illustrate sequence diagrams of an exemplary call flow, according to one or more embodiments;

[0013] FIG. 9 illustrates an exemplary entity-relational diagram that can be used in a rules engine, according to one or more embodiments;

[0014] FIG. 10 illustrates an exemplary block diagram of a data processing system, according to one or more embodiments;

[0015] FIG. 1P illustrates the current Call Stream for distributed IVR applications, according to one or more embodiments;

[0016] FIG. 2P illustrates the envisioned Call Stream for distributed IVR Applications, according to one or more embodiments;

[0017] FIG. 3P illustrates the envisioned components for managing, auctioning and inserting targeted messages into the slots of the CallStream of a distributed IVR application, according to one or more embodiments;

[0018] FIG. 4P illustrates an alternate view of the envisioned components for inserting targeted messages into slots of the CallStream of a distributed IVR application, according to one or more embodiments;

[0019] FIG. A1 illustrates the discrepancy between users actually reached and user that the advertiser desires to reach in traditional advertising, according to one or more embodiments;

[0020] FIG. A2 shows the relationship between the various participants in the exchange, according to one or more embodiments;

[0021] FIG. A3 is a flow chart showing the steps in an ad placement transaction, according to one or more embodiments;

[0022] FIG. A4 is a more detailed flow chart showing the operation of the exchange, according to one or more embodiments;

[0023] FIG. A5 illustrates how revenue is apportioned in the exchange, according to one or more embodiments;

[0024] FIG. A6 shows a version of the apportionment when the transaction is private within one member's network, according to one or more embodiments;

[0025] FIG. A7 shows the type of information available to generate an advertising strategy on the exchange, according to one or more embodiments;

[0026] FIG. A8 shows how the pricing structure available on the exchange is advantageous, according to one or more embodiments;

[0027] FIG. B1 illustrates an example of a voice or unified messaging system at which telephone switches and messaging systems are deployed at multiple locations, as known in the art;

[0028] FIG. B2 illustrates an example of a voice or unified messaging system at which telephone switches are deployed at multiple locations and associated with a central messaging system, as known in the art;

[0029] FIG. B3 illustrates an example of a voice or unified messaging system at which telephone switches and messaging systems are deployed at multiple locations together with an enhanced central messaging system, according to one or more embodiments;

[0030] FIG. B4 illustrates an example implementation of the enhanced central messaging system of the embodiment of FIG. B3, according to one or more embodiments;

[0031] FIG. B5 illustrates an example implementation of a messaging system deployed at a location in accordance with the embodiment of FIG. B3, according to one or more embodiments;

[0032] FIG. C1 illustrates in block diagram form a typical architecture of an electronic media system in which the system for customized electronic identification of desirable objects, according to one or more embodiments;

[0033] FIG. C2 illustrates in block diagram form one embodiment of the system for customized electronic identification of desirable objects, according to one or more embodiments;

[0034] FIGS. C3 and C4 illustrate typical network trees, according to one or more embodiments;

[0035] FIG. C5 illustrates in flow diagram form a method for automatically generating article profiles and an associated hierarchical menu system, according to one or more embodiments;

[0036] FIGS. C6-C9 illustrate examples of menu generating process, according to one or more embodiments;

[0037] FIG. C11 illustrates a hierarchical cluster tree example, according to one or more embodiments;

[0038] FIG. C12 illustrates in flow diagram form the process for determination of likelihood of interest by a specific user in a selected target object, according to one or more embodiments;

[0039] FIGs. C13A and C13B illustrate in flow diagram form the automatic clustering process, according to one or more embodiments;

[0040] While the invention may be susceptible to various modifications and alternative forms, specific embodiments thereof are shown by way of example in the drawings and will herein be described in detail. It should be understood, however, that the drawings and detailed description thereto are not intended to limit the disclosure to the particular form disclosed, but on the contrary, the disclosure is to cover all modifications, equivalents and alternatives falling within the spirit and scope of an invention as defined by appended claims.

DETAILED DESCRIPTION

[0041] In one or more embodiments, methods and/or systems described can transform inbound calls into a series call streams and/or call states that can include contextualized, personalized, relevant messages, information, promotions, coupons, offers, advertising, surveys, vouchers, opt-in programs and/or other specific messages. For example, targeting multi-site, national retail entities, embodiments described herein enable richly personalized experiences for on-hold customers and provide entity-wide brand consistency. For instance, a system described herein can provide callers with prompts during a call, to receive messages, offers, promotions, discounts, coupons and opt-in programs in a form of response messages (e.g., text messages, multimedia messages, etc.) to a caller's phone or email address. Upon receipt of a prompt during various segments of the call, the caller can respond, which can include a pressing of an indicated number or sequence of numbers, by voice acknowledgement, or other possible forms of confirmation which can include facial or gestural recognition in the case of video calls. In response to caller confirmation, a message or series of relevant messages that can include messages, offers, promotions, discounts, coupons and opt-in programs, can be transmitted to the caller's telephony device for use in future transactions that can be conducted with the entity.

[0042] Turning now to FIG. 1A, a block diagram of one or more network communications systems is illustrated, according to one or more embodiments. As shown in FIG. 1A, one or more telephony devices (TDs) 1610-1614 can be coupled to a public switched telephone network (PSTN) 1130. In one example, one or more of TDs 1610-1614 can be coupled to PSTN 1130 in a wired fashion. In another example, one or more of TDs 1610-1614 can be coupled to PSTN 1130 in a wireless fashion. For instance, telephony device (TD) 1610 can include a mobile telephony device (e.g., a cellular telephony device, a satellite telephony device, etc.), and the mobile telephony device can be coupled to PSTN 1130 via a cellular or satellite telephony network coupled to or included in PSTN 1130. In one or more embodiments, a TD can include a voice over Internet protocol (VoIP) telephony device. In various embodiments, a TD can include a personal digital assistant (PDA), a smart phone, or a hand-held computing device that allows a user of the TD to communicate with PSTN 1130 and/or other networks such as a wireless local area network (WLAN), among others.

[0043] Users (e.g., callers) can use TDs 1610-1614 to place calls to one or more entity locations (ELs) 1410 and 1510. As shown, PSTN 1130 can be coupled to one or more

routers 1420 and 1520 of ELs 1410 and 1510, respectively. In one example, each of routers 1420 and 1520 can be coupled to PSTN 1130 via a primary rate interface (PRI). In one or more embodiments, routers 1420 and 1520 can receive the calls from one or more of TDs 1610-1614 and can determine where and/or how to route the calls. For example, routers 1420 and 1520 can determine where and/or how to route the calls based on one or more criteria and/or configurations. As illustrated, ELs 1410 and 1510 can include one or more extensions (Exts) 1430-1434 and 1530-1534, respectively. In one example, extension (Ext) 1430 of entity location (EL) 1410 can be associated with a women's clothing department of EL 1410, Ext 1432 can be associated with a women's shoe department of EL 1410, and Ext 1434 can be associated with a children's clothing department of EL 1410. In another example, Ext 1530 of EL 1510 can be associated with a lumber department of EL 1510, Ext 1532 can be associated with a paint department of EL 1510, and Ext 1534 can be associated with a power tools department of EL 1510.

[0044] As shown, data centers 1210 and 1310 can be coupled to one or more networks 1110 and 1120. In one or more embodiments, network 1110 and/or network 1120 can include a wired network, a wireless network or a combination of wired and wireless networks. In one or more embodiments, network 1110 can include a wide are network (WAN). In one example, network 1110 can include or be coupled to a WAN that accessible by the public (e.g., the Internet). In a second example, network 1110 can form part of the Internet. In another example, network 1110 can be or include a private network.

[0045] In one or more embodiments, network 1120 can include a WAN. For example, network 1120 can be or include a private network. In one or more embodiments, network 1120 can be or include a multiprotocol label switching (MPLS) network. In one example, the MPLS network can be used to create virtual links between two nodes and can encapsulate packets of various network protocols. For instance, packets of the MPLS network can be assigned labels, and the data included in the labels can be used to forward and/or route the packets of the MPLS network. In one or more embodiments, a converged MPLS router can interface and/or communicate with the MPLS network and/or provide network routing for voice and data to each of ELs 1410 and 1510. For example, one or more of routers 1224 and 1324 can be or include a converged MPLS router.

[0046] In one or more embodiments, router 1420 can be coupled to data centers 1210 and 1310 via respective routers 1224 and 1324 and network 1120, and/or router 1520 can be

coupled to data centers 1210 and 1310 via respective routers 1224 and 1324 and network 1120. As shown, router 1224 can be coupled to an Internet protocol (IP) private branch exchange (PBX) 1255 that can be coupled to a router 1222, and a customer database (CDB) 1250 (e.g., a customer relationship management information database) can be coupled to router 1222. In one or more embodiments, an IP PBX can provide one or more of session initiation protocol (SIP) registration of one or more extensions (e.g., one or more of Exts 1430-1434 and 1530-1534), partitioning of extensions at locations, and routing calls based on one or more rules, among others.

[0047] In one or more embodiments, router 1224 can include a voice gateway (VGW) (e.g., an entity VGW). In one or more embodiments, router 1222 can include one or more of a firewall mechanism, a firewall service, and one or more firewall data structures. In one example, router 1222 can be used to separate and/or isolate one or more communications paths of one or more of router 1224, CDB 1250, and IP PBX 1255 from one or more of router 1220, VGW 1240, CS 1230, database (DB) 1270, and rules engine (RENG) 1260. In another example, router 1222 can be used to separate and/or isolate one or more communications paths of one or more of router 1220, VGW 1240, CS 1230, DB 1270, and RENG 1260 from one or more of router 1224, CDB 1250, and IP PBX 1255.

[0048] As shown, router 1220, DB 1270, RENG 1260 and VGW 1240 can be coupled to CS 1230, and VGW 1240 can be coupled to router 1222. In one or more embodiments, CS 1230 can include one or more of router 1220, DB 1270, RENG 1260 VGW 1240, and router 1222. In one or more embodiments, DB 1270 can include a relational database management system. For example, DB 1270 can include an Oracle database, a DB/2 database, a PostgreSQL database, a SQL Server database, a SQLite database, or a MySQL database, among others.

[0049] As illustrated, a media provider 1010 can be coupled to network 1110 via a router 1020. As shown, media provider 1010 can include a media server 1030 that can include various media and/or messages that can be included and/or provided in a telephone call. For example, media server 1030 can provide the various media and/or messages to a call server (e.g., CS 1230, CS 1330, etc.), so that the call server can provide the various media and/or messages to callers. In one or more embodiments, router 1020 can securely communicate with router 1220 and/or router 1320. In one example, routers 1020 and 1220 can communicate using at least one of Internet Protocol Security (IPSec), a secure socket layer (SSL), and transport layer security (TLS), among others. In a second example, routers 1020

and 1220 can communicate via an encrypted tunnel. In another example, routers 1020 and 1220 can communicate via a virtual private network (VPN). In one or more embodiments, routers 1020 and 1220 can include one or more of a firewall mechanism, a firewall service, and one or more firewall data structures.

[0050] As shown, a message gateway (MGW) 1710 can be coupled to network 1110. In one or more embodiments, MGW 1710 can provide one or messages via one or more of a short message service (SMS) and a multimedia messaging service (MMS) to one or more of TDs 1610-1614. For example, MGW 1710 can be accessible by a call server (e.g., CS 1230, CS 1330, etc.) via an application programming interface (API). For instance, the call server can use MGW 1710 and its API to fulfill sending one or more of contextualized messages, personalized messages, text messages, multimedia messages, information, promotions, coupons, offers, surveys, vouchers, and opt-in programs, among others, to one or more of TDs 1610-1614

[0051] In one or more embodiments, data center 1310 can be a secondary or backup data center to data center 1210. For example, elements 1320-1370 of data center 1310 can include one or more same and/or similar structures and/or functionalities as those described with reference to respective elements 1220-1270 of data center 1210. In one or more embodiments, DB 1370 can be synchronized with DB 1270. For example, DB 1370 can include data of DB 1270 should data center 1210 become unreachable or should DB 1270 or another element of data center 1210 fail.

[0052] Turning now to FIG. 1B, a block diagram of one or more network communications systems is illustrated, according to one or more embodiments. As shown, FIG. 1B can include elements of FIG. 1A. As illustrated, PSTN 1130 can be coupled to a router 1226 included in data center 1210 and coupled to router 1224, and PSTN 1130 can be coupled to a router 1326 included in data center 1310 and coupled to router 1324. In one or more embodiments, ELs 1410 and 1510 can be coupled to PSTN 1130 via network 1120 and routers 1224 and 1226 or routers 1324 and 1326. As illustrated, ELs 1410 and 1510 can be coupled to network 1120 via respective routers 1420 and 1520.

[0053] Turning now to FIG. 2, a block diagram of states of a telephone call within a telephony management system is illustrated, according to one or more embodiments. As shown, a call 2000 can include one or more states 2010-2080. In one or more embodiments, one state can follow another state without a specific order, and/or one or more states can be

repeated. In one example, routing state 2050 or hold state 2070 can follow conversation state 2060. In another example, routing state 2050 and/or dial plan state can be repeated. In one or more embodiments, states 2010-2080, as shown, may not depict any specific order and are enumerated for exemplary and/or identification purposes.

[0054] As illustrated, analytic information associated with one or more of states 2010-2080 can be sent to and/or collected by CS 1230, according to one or more embodiments. In one example, the analytic information can include information associated with whether or not a rule of a rules engine was invoked and/or whether or not the rule of the rules engine was matched. In a second example, the analytic information can include information associated with what media was presented to a caller and/or what time the media was presented to the caller. In a third example, the analytic information can include information associated with a received call (e.g., call time, called number, calling number, etc.). In a fourth example, the analytic information can include information associated with a call transfer (e.g., call transfer time, transfer number, calling number, etc.). In a fifth example, the analytic information can include information associated with a disconnection of a call (e.g., disconnection call time, called number, calling number, etc.). For instance, the call may have been disconnected by the caller hanging up during a call state or disconnecting after a "bye" state. In a sixth example, the analytic information can include information associated with whether or not a message was proffered to a caller and/or whether or not the message was accepted by the caller. In a seventh example, the analytic information can include information associated with providing a customized menu to a caller and/or whether or not a specific customized menu matched one or more portions of a profile of the caller.

[0055] In one or more embodiments, ringback state 2010 can include sending media to a TD that has placed a call to an entity. For example, the media sent to the TD can include audio and/or video information that can convey information associated with the entity and/or one or more goods and/or services that can be provided by the entity. For instance, call information associated with call 2000 can include one or more of an identification of the TD (caller ID) and a telephone number that the TD called (DNIS or called party) that can be used to determine and/or provide the media sent to the TD.

[0056] In one or more embodiments, a first state of a call entered by the TD upon answer of the call at the telephony management system is the "hello" state. The hello state 2020 can include providing an acknowledgement portion to the TD. For example, the

acknowledgement can be based on the call information and/or profile information associated with the call information. In one instance, the acknowledgement can thank the user of the TD for being a rewards member of the entity. In another instance, the acknowledgement can thank the user of the TD for calling the entity.

[0057] In one or more embodiments, offer state 2030 can include sending offer information to the TD. In one or more embodiments, one or more methods and/or systems described enable presentation of offer information to users of various offer types for callers to accept or decline during a course of a set of call states presented during the call. The offer types presented in the offer information include offers to receive information, discounts, coupons, reviews, directions or other types of information. Further, the offer information can be presented in various forms. For example, the offer information can include one or more of audio, a graphic or image, text, and a video that is associated with one or more of the entity, the entity location, a brand of an item and/or service for sale by the entity, a sale event, a store-wide sale event, a department sale event, hours and/or days of operation of the entity and/or an entity location. The offer information further includes instructions for the caller to accept or decline the offer for delivery of the message. In one embodiment, the offer information indicates delivery of the message to the caller's mobile telephone. In one example, the message can be sent via a SMS or a MMS.

[0058] The caller can indicate whether or not the offered message is to be sent to the TD. For example, the user can send DTMF signaling information that can indicate whether or not the message is to be sent to the TD. In another instance, the information from the user can include voice information that can indicate whether or not the message is to be sent to the TD. In one or more embodiments, the message sent to the TD can include one or more of audio information, video information, graphic information, and text information. For example, the message can include one or more of contextualized message(s), personalized message(s), information, coupon(s), promotion(s), advertisement(s) survey(s), voucher(s), directions, map(s), membership program(s), event information, brand information, product information, music, tones (e.g., sonic logos, sound marks, etc.), opt-in program(s) and other specific offer(s). In one or more embodiments, the message can be an advertisement promoting a brand for a good or service separately offered for discount, promotion, or sale. In another instance, the offer can include a discount and/or coupon for one or more goods and/or services for sale by the entity. In one or more embodiments, one or more of the offer, the discount, the coupon, and/or the advertisement can be associated with a third party in

addition to or instead of the entity. Still further, in one or more embodiments, the message can be based on one or more portions of the profile information associated with the call information.

[0059] In one or more embodiments, dial plan state 2040 can include sending a menu (e.g., an audio menu, a graphical menu, etc.) to the TD. In one example, a default menu can be sent to the TD. In another example, a menu can be built or configured using the call information and/or one or more portions of the profile information associated with the call information, and can be sent to the TD. In one or more embodiments, the default menu and/or the configured menu can include choices or options to use one or more other menus (e.g., submenus), choices or options to be transferred to a department and/or extension, and/or to receive a message. For example, the message can include one or more of contextualized message(s), personalized message(s), information, coupon(s), promotion(s), survey(s), voucher(s), directions, map(s), membership program(s), event information, brand information, product information, music, tones (e.g., sonic logos, sound marks, etc.), opt-in program(s) and other specific offers(s).

[0060] In one or more embodiments, routing state 2050 can include transferring the call to a person, department, and/or extension of the entity coupled to the telephony management system. In one example, transferring the call to a department and/or extension can include sending media (e.g., audio media, graphical media, video media, etc.) to the TD while the call is being transferred. In another example, transferring the call to a department and/or extension can include sending a message to the TD while the call is being transferred. In one or more embodiments, conversation state 2060 can include a conversation of the user of the TD and an employee or representative of the entity.

[0061] In one or more embodiments, hold state 2070 can include sending media (e.g., audio media, graphical media, video media, etc.) to the TD while the call is on hold. In one example, the audio and/or video media can be a media stream from a multicast of media streams. For instance, a determination of the media stream from the multicast of media streams can be based on the call information, a geographic location of the TD, and/or one or more portions of the profile information associated with the call information. In one or more embodiments, the profile information can include demographic information of the user of the TD. In one or more embodiments, a last state of the call, state 2080 (a "bye" state), can include sending media (e.g., audio media, graphical media, video media, etc.) to the TD near

a conclusion of the call. For example, the media can thank the user of the TD for calling the entity and/or for patronage.

[0062] Turning now to FIG. 3, a method of sending a message to a telephony device is illustrated, according to one or more embodiments. At 3010, a call can be received. For example, router 1420 can receive a call from TD 1610 via PSTN 1130. At 3020, the call can be redirected to a voice gateway. For example, router 1420 can redirect the call to VGW 1220. At 3030, call information associated with the call can be sent to a call server. For example, VGW 1220 can send the call information associated with the call to CS 1230. In one or more embodiments, the call information associated with the call can include one or more of an identification of the TD (caller ID) and a telephone number that the TD called (DNIS or called party). In one or more embodiments, VGW 1220 can communicate with CS 1230 via a socket connection over a network. In one example, VGW 1220 can communicate with CS 1230 using a hypertext transfer protocol (HTTP) via the socket. In another example, VGW 1220 can communicate with CS 1230 using a data description language, such as VoiceXML (VXML), via the socket. In one or more embodiments, VXML can provide call control. For example, VGW 1220 can receive one or more of dual-tone multi-frequency (DTMF) signaling information and voice information and provide the information corresponding to the DTMF signaling information and/or the voice information to CS 1230.

[0063] At 3040, profile information associated with the call information can be determined. In one or more embodiments, determining the profile information can include querying a database for the profile information associated with the call information. For example, CS 1230 can perform a structured query language (SQL) query of DB 1270 using the call information to determine the profile information. For instance, the profile information associated with multiple users can be indexed by one or more of an identification of a TD and a telephone number that the TD called. In one or more embodiments, DB 1270 can be periodically updated with caller information from CDB 1250.

[0064] In one or more embodiments, the profile information associated with the call information can include information associated with the user of the TD. For example, the profile information can include information associated with the user such as one or more of a time zone, a country, a state, a city, a zip code, a telephony device type, a telephony device operating system, a telephony device phone number, a distance from an entity location, at least one movie rented, at least one television show watched, a number of children, at least

one investment, at least one product purchase, a marital status, an age, at least one interest, at least one hobby, a gender, a birthday, an income, at least one past transaction, a content rating, a membership status, and at least one past department of an entity called, among others. At 3050, messaging and routing for the call can be determined based on one or more of the determined profile information and the telephone number that the TD called. In one or more embodiments, CS 1230 can utilize rules engine 1260 to determine messaging and routing for the call. For example, one or more portions of the profile information can match and/or satisfy one or more rules of rules engine 1260 to determine messaging and routing for the call.

[0065] At 3060, a message can be determined. In one example, the message can be associated with one or more of the entity, the entity location, a brand of an item and/or service for sale by the entity, a sale event, a store-wide sale event, a department sale event, hours and/or days of operation of the entity and/or an entity location. In a second example, the message can include an advertisement. In another example, the message can include a discount and/or coupon for one or more goods and/or services for sale by the entity. In one or more embodiments, the message can be based on one or more portions of the profile information associated with the call information. In one or more embodiments, one or more of a message, a coupon, an offer, an advertisement, and a discount can be associated with a third party in addition to or instead of the entity.

[0066] At 3070, a proffer of a message can be sent to the TD. In one or more embodiments, the proffer of the message can convey a request of the user to provide input to accept or decline the message. In one example, the proffer of the message can convey a request of the user to input a number on a keypad of the TD to indicate that the message is to be sent to the TD. For instance, the request of the user to input a number on a keypad of the TD can convey: "Press 'one' to receive the message, or press 'two' to continue." In another example, the proffer of the message can convey a request of the user to speak a word into a sound input device of the TD to indicate that the message is to be sent to the TD. For instance, the request of the user to speak a word can convey: "Say 'yes' to receive the message, or say 'no' to continue."

[0067] At 3080, the user input can be received. For example, CS 1230 can receive, via VGW 1240, information from the user indicating whether or not the message is to be sent to the TD. In one instance, the information from the user can include DTMF signaling information that

can indicate whether or not the user is to be sent to the TD. In another instance, the information from the user can include voice information that can indicate whether or not the message is to be sent to the TD. At 3090, it can be determined whether or not to send the message to the TD. If the message is not to be sent to the TD, the method can proceed to 3120. In one or more embodiments, an amount of time transpiring without receiving user input can indicate that the message is not to be sent to the TD.

[0068] If the message is to be sent to the TD, the message can be sent to the TD at 3100. In one or more embodiments, the message can be sent to the TD can include one or more of audio information, video information, graphic information, and text information. In one example, the message can be sent via a SMS or a MMS. For instance, CS 1230 can send a SMS message or a MMS message that includes the message via MGW 1710 to TD 1610. In another example, the message can be sent via a data network. In one instance, the message can be sent via a data network such as public-accessible wide area network such as the Internet. In another instance, the message can be sent via a data network such as a cellular telephone network that supports data communications other than telephone conversations. In one example, the message can be sent via an email message to an email address of the user of the TD. In another example, the message can be sent to an application that is configured to execute on the TD.

[0069] In one or more embodiments, the message can include one or more of textual information, image information, video information, and audio information. In one example, a text message can include the message. In another example, one or more graphics and/or images can include the message. For instance, the one or more graphics and/or images can include a bar code (e.g., a computer-readable one or two-dimensional bar code) that can be used as a coupon or discount for one or more goods and/or services. In one or more embodiments, the message can be used to identify one or more of the TD and the user of the TD. For example, the bar code that can be used as a coupon or discount for one or more goods and/or services can include information that can be used to identify one or more of the TD and the user of the TD.

[0070] At 3120, a menu for the call can be built. For example, CS 1230 can retrieve multiple pieces of media from storage and/or DB 1270 and combine the multiple pieces of media to build the menu for the call. In one instance, the multiple pieces of media can be ordered based on one or more portions of the profile information. For example, a piece of media

associated with a department where the user spends more time and/or more money can be prioritized to be included in the menu before other pieces of media. For instance, a piece of media may include "Press 'one' for the ladies' shoes department" and may be prioritized to be included in the menu before another piece of media such as "Press 'two' for the men's department", based on information of the caller's profile. In this fashion, the menu options (e.g., digits "one", "two", "three", etc.) can be dynamically mapped and/or planned contextually.

[0071] In one or more embodiments, building the media for the various call states can include cross fading multiple pieces of media. For instance, the multiple pieces of media can include multiple audio files, and building various media sequences and/or the menu can include combining the multiple audio files such that the various media sequences and/or the menu appears to be a continuous audio stream to a human ear.

[0072] At 3130, the menu can be provided to the TD. At 3140, user input can be received. For example, CS 1230 can receive information from the user indicating a department and/or extension. In one instance, the information from the user can include DTMF signaling information that can indicate the department and/or extension. In another instance, the information from the user can include voice information that can indicate the department and/or extension.

[0073] At 3150, a selection, based on the user input, can be determined. For example, the menu can include two or more departments and/or extensions, and the selection of a department and/or extension can be determined from the user input. If the selection is for a first department and/or extension, then the call can be transferred to the first department and/or extension at 3160. If the selection is for a second department and/or extension, then the call can be transferred to the second department and/or extension at 3170. If the selection is for a third department and/or extension, then the call can be transferred to the third department and/or extension at 3180.

[0074] Turning now to FIG. 4, a method of transferring a call from a telephony device to an extension within the telephony management system is illustrated, according to one or more embodiments. At 4010, audio signals and/or data can be streamed and/or transmitted to the TD. For example, CS 1230 can stream and/or transmit audio that is specific to the user of the TD or that is sourced from one or more audio streams. In one instance, the stream of audio can be based on one or more portions of the profile information associated with the call

information. For example, the stream of audio can include information that is specific to and/or customized for the caller based on one or more portions of the profile information associated with the call information. In another instance, the stream of audio can be one of a multicast of audio streams. In one or more embodiments, CS 1230 can be a source of the multicast of audio streams.

[0075] In one or more embodiments, an audio stream of the multicast of audio streams can be determined based on one or more portions of the profile information associated with the call information. For example, a portion of the profile information can include an age. In one instance, a first audio stream of the multicast of audio streams can be determined if the age is within thirteen to sixteen years of age. In another instance, a second audio stream of the multicast of audio streams can be determined if the age is within forty-three to fifty-five years of age. In one or more embodiments, an audio stream of the multicast of audio streams can be determined based the extension that the call is being transferred. For example, a first audio stream of the multicast of audio streams can be determined if the call is being transferred to Ext 1430, and a second audio stream of the multicast of audio streams can be determined if the call is being transferred to Ext 1432.

[0076] In one or more embodiments, a message can be included in the audio stream. At 4020, it can be determined whether or not to send a proffer of a message. If it is determined not to send the proffer of a message, then the method can proceed to 4080. If it is determined to send the proffer of a message, a message can be determined at 4030. In one or more embodiments, the message can be determined in a same or similar fashion as determined at 3060 of FIG. 3. At 4040, the proffer of the message can be sent to the TD. In one or more embodiments, the proffer of the message can be sent to the TD in a same or similar fashion as 3070 of FIG. 3.

[0077] At 4050, user input can be received. In one or more embodiments, the user input can be received in a same or similar fashion as user input can be received with reference to method element 3080 of FIG. 3. At 4060, it can be determined whether or not to send the message to the TD. In one or more embodiments, determining whether or not to send the message to the TD can be performed in a same or similar fashion as 3090 of FIG. 3. If the message is not to be sent to the TD, the method can proceed to 4080. If the message is to be sent to the TD, the message can be sent to the TD at 4070. In one or more embodiments, the

message can be sent to the TD in a same or similar fashion as the message can be sent to the TD with reference to method element 3100 of FIG. 3.

[0078] At 4080, call transfer information can be sent to a VGW. For example, CS 1230 can send the call transfer information to VGW 1240. For instance, the call transfer information can include extension information and an address. In one or more embodiments, the address can include an Internet protocol (IP) address. For example, the address can be an IP address of a router that manages communications for an entity location. For instance, the IP address can be an IP address of router 1420 that manages communications for EL 1410, and the call transfer information can be of form similar to "1430@IPaddress", where IPaddress is an IP address of router 1420 and "1430" can be used by router 1420 to route the call transfer to Ext 1430. In one or more embodiments, the IP address of the router that manages communications for the entity location can be a publicly routable IP address. In one or more embodiments, the IP address of the router that manages communications for the entity location can be a private IP address. For example, the private IP address may be associated with one or more IP addresses described in Request for Comments (RFC) 1918 which is available from the Internet Engineering Task Force (IETF). For instance, network 1120 can be a private network that uses private IP addresses. At 4090, the call transfer information can be sent to an IP PBX to transfer the call. For example, VGW 1240 can send the call transfer information to IP PBX 1255 via a SIP.

[0079] In one or more embodiments, CS 1230 can be or include a SIP agent that can continue to stream the audio to the TD and maintain a SIP state while other one or more portions of methods and/or processes are occurring. At 4100, a call transferred message can be received. For example, CS 1230 can receive a SIP message indicating that the call has been transferred to Ext 1430. At 4110, the audio that is streamed to the TD can be ceased or stopped.

[0080] Turning now to FIG. 5, a method of providing a message based at least on one or more submenu choices is illustrated, according to one or more embodiments. At 5010, an initial menu can be built with one or more submenu choices and can include one or more department and/or extension choices. In one or more embodiments, the initial menu can be built in a similar fashion to the menu built in method element 3120 of FIG. 3. In one or more embodiments, the submenu choices of the initial menu can include choices that are user selectable to navigate to another menu associated with a choice. In one example, a submenu choice can be for a hardware department that can lead to a submenu that can include two or

more submenu items or options of plumbing hardware, fasteners (e.g., screws, bolts, nuts, nails, glue, etc.), and electrical hardware (e.g., wire, electrical outlets, circuit breakers, etc.), among others. In another example, a submenu choice can be for a children's department that can lead to a submenu that can include two or more submenu items or options of children's shoes, boys' clothing, girls' clothing, toddlers' clothing, and pre-teens' clothing, among others.

[0081] At 5020, the initial menu with submenu choices can be provided to the TD. For example, CS 1230 can provide the initial menu with submenu choices to TD 1610. At 5030, user input can be received. For example, CS 1230 can receive user input from TD 1610, and the user input can include a choice from the initial menu. In one instance, the user input can include DTMF signaling information that can indicate a choice from the initial menu. In another instance, the user input can include voice information that can indicate a choice from the initial menu.

[0082] At 5040, it can be determined whether or not the user input indicates a transfer. For example, a menu choice from the initial menu can be associated with a department and/or extension. For instance, the menu choice can be associated with Ext 1430. If the user input indicates a menu choice that is associated with a department and/or extension, the call can be transferred to the department and/or extension at 5150. If the user input does not indicate a menu choice that is associated with a department and/or extension, the method can proceed to 5050, where the submenu choice can be stored.

[0083] At 5060, it can be determined whether or not to send a proffer of a message. In one or more embodiments, CS 1230 can determine whether or not to send the proffer of a message to the user of the TD based on a number of submenu choices that have been selected by the user, an ordered sequence of two or more submenu choices that have been selected by the user, an ability of the TD to receive a message, a distance the TD is from an entity location, an amount of time transpiring during the call, and/or one or more portions of the profile information associated with the call information.

[0084] In one or more embodiments, determining whether or not to send the proffer of a message to the user of the TD can include determining whether or not the TD includes a capability to receive a message. In one example, CS 1230 can search a database using the call information associated with the call to determine whether or not the TD includes the capability to receive the message. In one instance, the database can include telephone numbers (e.g., caller IDs) and corresponding one or more capabilities associated with the

telephone numbers, and CS 1230 can search the database using the telephone number (e.g., caller ID) of TD 1610 to determine whether or not TD 1610 can receive the message. In another instance, the database can include telephone numbers (e.g., caller IDs) and corresponding one or more communication methods and/or systems that can be used to send the message to the user of the TD. For example, CS 1230 can search the database using the telephone number (e.g., caller ID) of TD 1610 to determine whether or not an email address is associated with the telephone number of TD 1610, where the email address can be used to send the message to the user of TD 1610.

[0085] In one or more embodiments, determining whether or not to send the proffer of a message to the user of the TD can include determining a distance the TD is from an entity location. In one example, CS 1230 can search a database using the call information associated with the call to determine distance the TD is from an entity location. For instance, the database can include telephone numbers (e.g., caller IDs) and corresponding geographic locations associated with the telephone numbers, and CS 1230 can search the database using the telephone number (e.g., caller ID) of TD 1610 to determine the geographic location of TD 1610. Using the geographic location (e.g., latitude, longitude, and/or altitude) of TD 1610 and a geographic location of an entity location, CS 1230 can determine a distance TD 1610 is from the entity location.

[0086] In a second example, CS 1230 can query a system (e.g., a cellular telephone system, a satellite telephony system, a VoIP telephony system, etc.) using the call information associated with the call to determine a distance the TD is from an entity location. In one instance, the system can provide a geographic location of the TD or a geographic location of a cellular telecommunications antenna communicating with the TD. In another instance, CS 1230 can provide a geographic location of an entity location and the telephone number (e.g., caller ID) of TD 1610 to the system, and the system can respond by providing a distance TD 1610 is from the entity location.

[0087] In another example, CS 1230 can query TD 1610 for a geographic location of TD 1610, and TD 1610 can respond by providing its geographic location that can be usable by CS 1230 to determine a distance TD 1610 is from an entity location. In one instance, TD 1610 can include a global positioning (GPS) device that can determine a geographic location of TD 1610 that can be provided to CS 1230. In a second instance, a system (e.g., a cellular telephone system) can provide a geographic location of TD 1610 or a geographic location of

a cellular telecommunications antenna communicating with TD 1610 to TD 1610. In another instance, the user of TD 1610 can provide a geographic location of TD 1610 to TD 1610. In one example, the user can provide one or more of a zip code, an address, a town or city name, and an intersection of streets, among others, to TD 1610.

[0088] If it is determined not to send the proffer of a message, then the method can proceed to 5130. If it is determined to send the proffer of a message, a message can be determined at 5070, where a message can be determined based at least on one or more submenu choices. In one example, the message can be determined based on a number of submenu choices that have been selected by the user. In a second example the message can be determined based on an ordered sequence of two or more submenu choices that have been selected by the user. In one or more embodiments, the message can also be determined based on one or more portions of the profile information associated with the call information and/or a distance from an entity location. In one example, a first message can include a first discount (e.g., a first coupon) for one or more goods and/or services to a first TD (e.g., TD 1610) based on a first distance the first TD is from the entity location (e.g., EL 1410), and a second message can include a second discount (e.g., a second coupon) for the one or more goods and/or services to a second TD (e.g., TD 1612) based on a second distance the second TD is from the entity location. For instance, the second distance can be greater than the first distance, and the second discount can be greater than the first discount. In a second example, the message can include one or more of a map and directions to an entity location.

[0089] At 5080, the proffer of the message to the TD can be sent to the TD. In one or more embodiments, the proffer of the message can be sent to the TD in a same or similar fashion as 3070 of FIG. 3. At 5090, user input can be received. In one or more embodiments, the user input can be received in a same or similar fashion as user input can be received with reference to method element 3080 of FIG. 3. At 5100, it can be determined whether or not to send the message to the TD. In one or more embodiments, determining whether or not to send the message to the TD can be performed in a same or similar fashion as 3090 of FIG. 3. If the message to the TD is not to be sent to the TD, the method can proceed to 5130. If the message to the TD is to be sent to the TD, the message can be sent to the TD at 5120. In one or more embodiments, the message can be sent to the TD in a same or similar fashion as the message can be sent to the TD with reference to method element 3100 of FIG. 3.

[0090] At 5130, a submenu can be built. In one or more embodiments, the submenu can be built in a similar fashion to the menu built in method element 5010. At 5140, the submenu can be provided to the TD. In one or more embodiments, the submenu can be provided to the TD in a similar fashion to providing the initial menu to the TD in method element 5020. In one or more embodiments, the method can proceed to 5030. In one or more embodiments, a submenu of an initial menu or another submenu may include two or more department and/or extension options without further submenu options. In one or more embodiments, a submenu of an initial menu or another submenu may include one or more submenu options and/or one or more department and/or extension options that have already been presented to the user of the TD during the call.

[0091] Turning now to FIG. 6, a method of operating a telephony management system is illustrated, according to one or more embodiments. At 6010, the call server (e.g., CS 1230, CS 1330, etc.) can receive information. For example, the received information can include analytic information that can be stored or can be request information. In one instance, the analytic information can include information of and/or resulting from one or more of states 2010-2080. For example, the analytic information can include a result from a condition evaluation. In another instance, the request information can include a request for information from a rules engine (e.g., RENG 1260, RENG 1360, etc.). In one or more embodiments, the received information can be from a device different from the call server or from a method or process utilized by the call server.

[0092] At 6020, content of the received information can be determined. If the received information includes analytic information, then the analytic information can be stored at 6030. In one example, the analytic information can be stored in a log file (e.g., a log file of the call server). In another example, the analytic information can be stored in a database (e.g., DB 1270, DB 1370, etc.). In one or more embodiments, the method can proceed to 6010.

[0093] If the received information includes request information, then the method can proceed to 6040 where a database (e.g., DB 1270, DB 1370, etc.) can be queried. In one or more embodiments, querying the database can include sending information to the database. For example, one or more of an identification of a TD (caller ID) and a telephone number that the TD called (DNIS or called party) can be sent to the database, where the database can use the

one or more of the identification of the TD and the telephone number that the TD called to retrieve associated information. At 6050, a result can be received from the database.

[0094] At 6060, a rules engine (e.g., RENG 1260, RENG 1360, etc.) can be queried. In one or more embodiments, querying the rules engine can include sending information to the rules engine. In one example, the rules engine can be sent information to determine whether or not to proffer a message to a user of a TD. In a second example, the rules engine can be sent information to determine a message to be sent to a user of a TD. In a third example, the rules engine can be sent information associated with a call sequence or call state. In another example, the rules engine can be sent information associated with the result from the database query.

[0095] At 6070, a result from the rules engine can be received. In one or more embodiments, the result from the rules engine can include a default result (e.g., no rule matched the query information) or a result based on the query information matching one or more rules. At 6080, the result from the rules engine can be sent to the device, method, or process that sent the request for the result. In one or more embodiments, the method can proceed to 6010.

[0096] Turning now to FIG. 7, an exemplary block diagram of an entity location that receives messages from users' telephony devices is illustrated, according to one or more embodiments. In one or more embodiments, EL 1410 can include one or more point of sale terminals (POSes) 7010-7016 coupled to router 1420, and each of POSes 7010-7016 can be or include a computing device. In one or more embodiments, a point of sale terminal (POS) (e.g., a POS of POSes 7010-1016) can receive message information, or information associated with the message information, sent from a TD (e.g., a TD of TDs 1610-1616). For example, the message information, or information associated with the message information, can include a coupon that can be used to provide a discount on compensation (e.g., cost) for one or more goods and/or services messaged for sale by an entity. For instance, POS 7010 can receive a coupon via message information 7250, or information associated with the message information 7250, and a user of TD 1610 can receive ten percent (10%) off a purchase of ladies shoes, a brand of ladies shoes, clothing, a brand of clothing, one or more goods and/or services messaged for sale by the entity, etc. In one or more embodiments, message information, or information associated with the message information, sent from a TD to a POS can include an identifier and/or an index that is usable to identify one or more of a message, a coupon, the TD, and a user of the TD, among others.

[0097] In one or more embodiments, EL 1410 can include an optical input device (OID) 7110 coupled to POS 7010, where OID 7110 can sense and/or discern one or more images from a display 7230 of TD 1610 and/or from a paper medium. For example, OID 7110 can transform the one or more images from display 7230 of TD 1610 and/or from the paper medium and provide data associated with the one or more images to POS 7010, where the data is usable by POS 7010 and conveys message information to POS 7010. In one or more embodiments, OID 7110 and/or POS 7010 can use one or more of pattern recognition, artificial intelligence, and computer vision systems and/or methods to determine information conveyed by the one or more images from display 7230 of TD 1610 and/or from the paper medium. In one or more embodiments, determining the information conveyed by the one or more images can include transforming optical signals from the one or more images into data that can be used by a computing device and/or stored in a memory medium included in or coupled to the computing device.

[0098] In one example, OID 7110 can sense and/or discern a computer-readable image (CRI) 7280 displayed by display 7230 of TD 1610 and provide information that is included and/or encoded in CRI 7280 to POS 7010. For instance, CRI 7280 can include a barcode, and OID 7110 can sense and/or discern the barcode displayed by display 7230 and provide information that is included and/or encoded in the barcode to POS 7010. In one or more embodiments, the information included and/or encoded in CRI 7280 can include information associated with message information 7250. In one or more embodiments, message information 7250 can be received by TD 1610 when CS 1230 sends a message to TD 1610.

[0099] In one or more embodiments, EL 1410 can include a radio frequency device (RFD) 7112 coupled to POS 7012. For example, RFD 7112 can include one or more of a radio frequency receiver and a radio frequency transmitter configured to communicate with a TD, such as TD 1612. As shown, TD 1612 can include a RFD 7360 that can include one or more of a radio frequency receiver and a radio frequency transmitter configured to communicate with another RFD, such as RFD 7112. In one or more embodiments, RFD 7360 can transmit message information 7350 to RFD 7112, and RFD 7112 can provide message information 7350 to POS 7012. In one example, a user of TD 1612 can direct TD 1612 to transmit message information 7350 to RFD 7112. In another example, RFD 7112 can transmit a query to TD 1612, via RFD 7360, and TD 1612 can respond to the query by transmitting message information 7350 to RFD 7112, via RFD 7360. In one or more embodiments,

message information 7350 can be received by TD 1612 when CS 1230 sends a message to TD 1612.

[00100] In one or more embodiments, RFDs 7112 and 7360 can communicate using one or more ISM (industrial, scientific and medical) bands. For instance, an ISM band can include a frequency range of 6.765-6.795Mhz, 433.05-434.79Mhz, 902-928Mhz, 2.4-2.5Ghz, 5.725-5.875Ghz, or 24.0-24.25Ghz, among others. In one or more embodiments, RFDs 7112 and 7360 can communicate using one or more of IEEE 802.11, IEEE 802.15, IEEE 802.15.4, ZigBee, 6LoWPAN, spread spectrum of frequencies, frequency modulation of a carrier wave, and amplitude modulation of a carrier wave, among others. In one or more embodiments, TD 1612 can be identified by information used in communicating with RFD 7112 and/or POS 7012. In one example, TD 1612 can be identified by a media access control (MAC) address used in communicating with RFD 7112 and/or POS 7012. In another example, TD 1612 can be identified by an IP address or a portion of an IP address used in communicating with RFD 7112 and/or POS 7012. For instance, IPv6 (Internet protocol version 6) can be used in communicating with RFD 7112 and/or POS 7012, and TD 1612 can be identified by IP address or a portion of an IP address used by TD 1612 can be used to identify TD 1612. In one or more embodiments, one or more of the MAC address, the IP address, and the portion of the IP address can be included in analytic information that can be sent CS 1230 for storage and/or analysis.

[00101] In one or more embodiments, EL 1410 can include an infrared device (IRD) 7114 coupled to POS 7014. For example, IRD 7114 can include one or more of an infrared receiver and an infrared transmitter configured to communicate with another IRD, such as an IRD 7470. As shown, TD 1614 can include IRD 7470 which can include one or more of an infrared receiver and an infrared transmitter configured to communicate with another IRD, such as IRD 7114. In one or more embodiments, IRD 7470 can transmit message information 7450 to IRD 7114, and IRD 7114 can provide message information 7450 to POS 7014. In one example, a user of TD 1614 can direct TD 1614 to transmit message information 7450 to IRD 7114. In another example, IRD 7114 can transmit a query to TD 1614, via IRD 7470, and TD 1614 can respond to the query by transmitting message information 7450 to IRD 7114, via IRD 7470. In one or more embodiments, message information 7450 can be received by TD 1614 when CS 1230 sends a message to TD 1612.

[00102] In one or more embodiments, an entity can use an entity associate 7610 (e.g., an employee of the entity, a contractor of the entity, a volunteer of the entity, etc.) to read message text 7580 from a display 7530 of a TD 1616, and entity associate 7610 can provide message text 7580 to POS 7016. For example, entity associate 7610 can provide message text 7580 to POS 7016 via a keyboard or keypad of or coupled to POS 7016. In one or more embodiments, message text 7580 can include information associated with message information 7550. For example, information associated with message information 7550 can be included and/or encoded in message text 7580. In one or more embodiments, message information 7550 can be received by TD 1616 when CS 1230 sends a message to TD 1616.

[00103] As shown, each of TDs 1610-1616 can include respective processors 7220-7520 coupled to respective memories 7240-7540 and coupled to respective displays 7230-7250. In one or more embodiments, memories 7240-7540 can include respective message information 7250-7550 and respective program instructions 7290-7590 that can be executed on a respective processor to implement one or more methods and/or systems described herein. As illustrated, TD 1612 can include RFD 7360 coupled to processor 7320, and TD 1614 can include IRD 7470 coupled to processor 7420.

[00104] Turning now to FIGs. 8A and 8B, sequence diagrams of an exemplary call flow is illustrated, according to one or more embodiments. As shown, a caller can dial into an entity VGW (e.g., router 1420), and the entity VGW can forward the call to an entity IP PBX (e.g., IP PBX 1255). The entity IP PBX can route the call to a TMS (telephony management system) VGW that can invoke an application on a call server (e.g., CS 1230) which implements one or more methods described herein. In one or more embodiments, the TMS VGW can provide call information to the call server via VXML (Voice eXtensible Markup Language). For example, the call information can include an identification of the TD (ANI or caller ID) and a telephone number that the TD called (DNIS or called party).

[00105] The application executing on the call server can query a customer database (e.g., CDB 1250) or a database (e.g., DB 1270) that includes caller information from the customer database to determine at least one entry that includes caller profile information associated with the identification of the TD. Using the caller profile information and one or more of analytic information and business rules, the application executing on the call server can determine and/or produce media that includes a proffer of a message to provide to the caller. In one or more embodiments, the media can include the multiple messages that are cross-

faded together. For example, cross-fading the multiple pieces of media can include combining the multiple audio files such that the media appears to be a continuous audio stream to a human ear. For instance, the media can include information such as "Thank you for calling Corkies Department Store at Barton River Mall. Please press 'one' to receive a coupon to receive ten percent off any pair of ladies shoes, or press 'two' to continue."

[00106] The caller can provide user input to the TD that indicates a request for or an acceptance of the proffer for the message. For example, the user can press "one" on a keypad of the TD, and the TMS VGW can receive DTMF signaling information indicating the "one" on the keypad of the TD. The TMS VGW can provide VXML data to the call server that indicates the "one" on the keypad of the TD, and the call server can provide the message to the TD. In one example, the call server can send a coupon for ten percent off any pair of ladies shoes to the TD via MGW 1710. In another example, the call server can send a coupon for ten percent off any pair of ladies shoes to an email address associated with the caller via a data network. In one or more embodiments, the coupon can include text and/or graphics usable to receive ten percent off any pair of ladies shoes and/or to identify one or more of the TD and the user of the TD.

[00107] Using the caller profile information and one or more of analytic information and business rules, the application executing on the call server can determine and/or produce media (e.g., a menu) to the caller. In one or more embodiments, the media can include the multiple pieces of media that are cross-faded together. For instance, the media can include information such as "Please make a selection from the following menu items. Press 'one' for our ladies department; press 'two' for our teens' department; press 'three' for our children's department; press 'four' for our men's department; or press 'zero' for operator assistance." The caller can provide user input to the TD that indicates the ladies department. For example, the user can press "one" on a keypad of the TD, and the TMS VGW can receive DTMF signaling information indicating the "one" on the keypad of the TD. The TMS VGW can provide VXML data to the call server that indicates the "one" on the keypad of the TD, and the call server can forward (e.g., transfer) the call to an entity telephone (e.g., Ext 1432) via the TMS VGW and the entity VGW.

[00108] Turning now to FIG. 9, an exemplary entity-relational diagram that can be used in a rules engine is illustrated, according to one or more embodiments. As shown, the entity-relational diagram (ERD) can include tables 9010-9150. In one or more embodiments, the

rules engine can be implemented using a database (e.g., DB 1270, DB 1370, etc.), and tables 9010-9150 can store information associated with implementing the rules engine. For example, the database can be or include a relational database system (RDBMS), where one or more tables of the ERD can be linked to one or more other tables of the ERD via one or more primary keys and/or foreign keys. For instance, Department table 9020 can be linked to Store table 9010 via a foreign key (e.g., Department.StoreID). In another instance, StoreMenu table 9040 can be linked to Store table 9010 via a first foreign key (e.g., StoreMenu.StoreID) and can be linked to Department table 9020 via a second foreign key (e.g., StoreMenu.DepartmentID).

[00109] Turning now to FIG. 10, an exemplary block diagram of a data processing system is illustrated, according to one or more embodiments. As shown, a data processing system (DPS) 10010 can include a processor 10020, a memory 10030 coupled to processor 10020, and one or more network interfaces 10060-10066 coupled to processor 10020. As illustrated, memory 10030 can include one or more applications (APPs) 10040 and 10042, an operating system (OS) 10045, and/or a configuration 10050. In one or more embodiments, one or more of APPs 10040 and 10042 and OS 10045 can store program instructions that are executable by processor 10020 to implement on or more processes and/or methods described herein. In one or more embodiments, configuration 10050 can store configuration information usable by one or more of APPs 10040 and 10042 and OS 10045 to implement on or more processes and/or methods described herein.

[00110] In one or more embodiments, one or more of network interfaces 10060-10066 can be configured to be coupled to one or more of a router, a network, a telephony system, a telephony network, a telephony device, and another DPS, among others. In one or more embodiments, one or more of the devices and/or systems described herein can include same and/or similar structures and/or functionalities as those described with reference to DPS 10010. For example, one or more of: routers 1020, 1220-1226, 1420 1320-1326, and 1520; VGWs 1240 and 1340; CSs 1230 and 1330; media server 1030; DBs 1270 and 1370; RENGs 1260 and 1360; CDBs 1250 and 1350; IP PBXs 1255 and 1355; and POSes 7010-7016 can include same and/or similar structures and/or functionalities as those described with reference to DPS 10010.

[00111] The exemplary system reduced to practice is initially intended to allow for syndicated streaming-multimedia content to be delivered to telephony end-points in multi-

location branded outlets. For example, an earlier prototypical / predecessor system delivered corporate-wide consistent branded audio to on-hold callers at hundreds of locations at a major world-wide hotel brand. This content can be managed at a corporate-wide level, can be co-created and managed with co-marketing partners, can be modified at particular times, for particular locations or even for particular callers. The content can consist of brand messaging, survey questions, event promotions, program promotions, informational or emergency bulletins, and 3rd party advertising for those partners that wish access to those people who call into that major hotel brand.

[00112] Similarly, we could envision a very large multi-location retailer (such as a Wal-mart) centrally manage the entire call experience at over thousands of locations, including the introductory message ("The Hello"), every choice point ("Press 1 for Electronics, Press 2 for Women's Apparel, …"), in-transit routing messages ("You are being transferred to the Electronics department"), and on-hold with context ("While you are on hold for Electronics, would you like to hear about the new Apple iPod nano?").

[00113] The system enables the corporation (say, a retailer such as a Wal-mart) to reserve "slots" in the CallStream for corporate use, allocate some for regional use, allocate more for store use, and let 3rd party advertisers bid for the remaining slots.

[00114] Fig. 1P shows the current CallStream, with either 'dead air' identical messaging from a "tape loop" for every caller. The 'tape loops' are created by the store manager, are inconsistent, and often old.

[00115] Fig. 2P shows an envisioned CallStream with syndicated messages that are new and consistent across a brand's locations, and can be customized to the caller at each branch point and each transit point.

[00116] Figures 3P and 4P illustrate the system, as reduced to practice, consists of: Content Management System. Allows for the creation, management, targeting, placement and bids for streaming multimedia content.

[00117] Master IVR Control. Allows for the centralized management of all in-store or multi-location telephony end-points, their related slots, and their related IVR programming.

[00118] The Streaming Multi-media Object Selection Service, which selects which piece of media (say, audio) will be played for any particular caller, at any particular point of interaction with the brand's store by phone.

[00119] The Streaming Multi-media Ad Marketplace, which allows for 3rd party advertisers to bid for 'slots' in the CallStream which match a particular media buy profile.

[00120] The Streaming Multi-media Unified Messaging Fabric, which is a telephony framework, incorporating the Master IVR and the telephony end-points which allow centralized control of a heterogeneous, very distributed, and previously unmanaged telephony network.

[00121] The sub-systems to be discussed in more detail in this disclosure are: (a) the Unified/Syndicated Messaging System for IVR, (b) the Streaming Multi-media Selection Service, and (c) the Streaming Multi-media Ad Marketplace.

[00122] It is the company's plan, and the inventor's contention that the system should not be limited to this particular implementation, but that the system is easily extended to support the following:

[00123] Multiple Types of Content. The system will easily scale to include other streaming multi-media content types, such as, but not limited to video and dynamic time-based 2D and 3D models (motion models).

[00124] Multiple Intentions of Content. This disclosure focuses heavily on advertising (for bidding) and matching of information (targeting/selection scenarios. Other purposed content, such as transmission of audio icons, watermarks, digital data, coupons, blended infomercials, and promotional material would all fall under the purview of an identically architected system.

[00125] Multiple Types of Interaction and Feedback. This disclosure does not focus heavily upon voice or keyed feedback; this is not intended to limit the scope of invention with respect to creating context for selection and delivery of multi-media objects. One extension

envisioned is the delivery of an SMS (Mobile phone Text) coupon upon the press of a key requesting such delivery during the presentation of a streaming multimedia object.

[00126] Multiple Types of End-points for Delivery of Content. The current disclosure focuses on multi-location telephony end-points, but the rest of the invention will work equally well with VOIP end-points, IP Radio end-points, IP TV end-points, and advertiser supported calling card end points.

[00127] Further, the inventor's would like to make note that the particular methods of media selection or targeting are illustrative, and should not be considered limiting.

[00128] Syndicated Messaging Architecture and System for IVR Applications

[00129] Messaging systems are well known in the art. One of the simplest forms of messaging system is a voice messaging system. Computer Telephony Integration (CTI) is a term, which refers to the integration of computer architectures with telephony systems. A voice messaging system is an example of a CTI system, and a further example is a unified messaging system.

[00130] Voice or unified messaging systems have conventionally been deployed in one of three implementations.

[00131] In a first voice messaging implementation a voice messaging (or unified messaging) system is co-located with a subscriber's telephone switch, i.e. a subscriber's private branch exchange (PBX). Within a multi-site organization or enterprise this will tend to result in each physical location (which has its own telephone switch) having a separate voice messaging (or unified messaging) system. For an enterprise with many distributed locations, such an arrangement requires the deployment of many (possibly small) systems resulting in considerable IT administrative costs.

[00132] FIG. B1 shows two distributed locations each with their own telephone switch, specifically a PBX 102 and 106 respectively, and respective dedicated unified or voice messaging systems 104 and 108. This system provides, at least, the following three typical voice-messaging functions: call answering, automated attendant and subscriber access.

[00133] 1. The call answering feature is initiated when calls are forwarded by the telephone system as a result of the called party not answering (e.g. due to a busy or no-answer condition).

[00134] 2. The automated attendant feature offers callers a menu for automatically routing a call to the desired answering point, without the need for operator intervention.

[00135] 3. The subscriber access feature allows subscribers, or mailbox owners, of the voice messaging system to call into the system and retrieve their messages over the phone.

[00136] In a second voice messaging implementation, the voice messaging system utilizes a deployment of a central messaging system serving multiple remote locations. This form of deployment requires long-distance telephone connections to be provided between the remote private branch exchanges (PBXs), serving remote locations, and the centralized messaging system.

[00137] FIG. B2 shows two locations each having a respective PBX 202 and 204 being served by a single centralized unified or voice-messaging system 206 associated with a single centralized PBX 208. The long distance telephone connections generally need to be provided regardless of their use (i.e. it is necessary to purchase enough capacity to handle busy periods). In addition some voice messaging features are lost in the centralized arrangement of FIG. B2. The automated attendant application provides callers into an office or building with the ability to connect to a user. This capability is not provided in the centralized arrangement of FIG. B2: the system providing the automated attendant function handles all users for all systems, and so does not provide service equivalent to a human attendant at the remote site. Additionally, different offices may be located in different countries with different language requirements. A simple example is that a first office may be in Canada and need both English and French languages, whereas another office served from the same centralized facility may have entirely, non-overlapping, language requirements, for example located in Mexico and requiring Spanish. A common and unique numbering plan across all mailboxes across all sites is also needed.

[00138] In a third implementation the second implementation is modified such that the long-distance telephone connectivity is provided by an organization's (existing) data (IP)

Wide Area network (WAN), possibly using Voice over Internet Protocol (VoIP). This deployment puts a strict requirement on the bandwidth and more critically the quality of service (e.g. the network latency) characteristics of the corporation's WAN. This most generally is not, and will not, be met for the majority of corporations now or in the near future.

[00139] It is an aim of the present invention to provide an improved messaging system. A further aim is to provide such a system offering increased deployment flexibility.

[00140] In a first aspect the present invention provides a messaging system comprising at least one distributed front-end messaging system and a centralized data store associated with said at least one distributed front-end messaging system, in which system the centralized data store includes means for storing data associated with users of the at least one distributed front-end messaging system, the at least one distributed front-end messaging system further including a respective at least one cache means for storing at least a portion of the centralized data associated with users of said at least one distributed front-end messaging system such that at least one messaging function can be provided to users of said at least one distributed front-end messaging system in dependence on the data stored in said cache means. The messaging system is preferably a voice messaging system or unified messaging system.

[00141] There may be provided a plurality of distributed front-end messaging systems each associated with a respective plurality of users and each including a cache means, wherein the centralized data store is adapted to store data associated with all users of said front-end messaging systems.

[00142] There may be further provided a centralized front-end messaging system associated with said centralized data store. The centralized front-end messaging system may be associated with a plurality of users, data associated with said users being stored in the centralized data store.

[00143] The centralized front-end messaging system may provide at least one messaging function for users of said at least one distributed front-end messaging system. The centralized front-end messaging system may be adapted to identify the front-end messaging system of a user. The centralized messaging system may be adapted to identify the front-end messaging system of a user in dependence on a called number, a calling number, or a unique user identifier.

[00144] The centralized messaging system may provide access to all stored data associated with said at least one distributed front-end messaging system associated with the user.

[00145] The centralized data store may store configuration data and message data associated with all users.

[00146] The at least one voice messaging function may include call answering. Said at least one messaging function may be a subscriber access function.

[00147] Each front-end messaging system may be associated with a respective Voice Mail Domain. Each front-end messaging system may be associated with a telecommunications switch.

[00148] In a further aspect the present invention provides a method of configuring a messaging system comprising: storing, at a centralized location, data associated with all users of the messaging system; storing, at least one distributed location, at least part of said data associated with users at the at least one distributed location, and providing at least one messaging function to users at the at least one distributed front-end messaging system in dependence on the data stored at the distributed location.

[00149] The step of storing the data at said at least one distributed location may comprise the step of caching the data at the centralized location. The method may further comprise step of providing at least one messaging function to users at the at least one distributed front-end messaging system in dependence on data stored at the centralized locations. The method may further comprise the step of accessing the centralized location directly.

[00150] In further aspects, the present invention also provides a messaging system comprising a plurality of remote servers and a centralized data store associated with said plurality of remote servers, in which system the centralized data store includes means for storing messages and data associated with all users of the plurality of remote servers, and the plurality of remote servers are each associated with a respective cache means for storing at least a portion of the data associated with users of said remote server such that at least one voice messaging function can be provided to users of said voice server independent of the centralized message store. The messaging system is preferably a unified messaging or voice messaging system.

[00151] The at least one voice messaging function may include call answering.

[00152] The plurality of voice servers may be associated with a plurality of Voice Mail Domains. The number of voice servers may correspond to the number of Voice Mail Domains. Each voice server may be associated with a telecommunications switch. The telecommunications switch may be a public branch exchange.

[00153] The centralized data store may further be provided with a collection of front-end servers or a single server, said interface being connected to a dedicated telecommunications

switch, such a private branch exchange. Such switch may provide subscriber access to the centralized data store.

[00154] The centralized data store may provide a common message store for all voice mail domains and a common directory store for all voice mail domains.

[00155] There may be provided for subscriber access to the centralized data store via said interface.

[00156] In accordance with embodiments of the invention the messaging function is architecturally split into a number of separate components, which, critically, can be implemented in at least two separate computer systems. These systems can be deployed:

[00157]

[00158] i. In separate physical locations connected using a data network. The message storage can be centralized within a single (or few) data storage facilities, providing security, simplicity and management costs benefits.

[00159] ii. These systems provide the beneficial external user characteristics normally derived from multiple distributed systems (e.g. a local automated attendant is provided, specific language support etc).

[00160] iii. In addition these systems may be engineered to provide reliable voice/unified messaging support without strict requirements on data network availability and quality of service (QoS).

[00161] Such an architecture allows much more flexibility of deployment for voice messaging or unified messaging systems.

[00162] Multi-media Ad Marketplace

[00163] In one embodiment, the invention is an Internet Advertising process for matching advertisers with via media'slots' via their respective brokers. The matching process includes providing an exchange (on which the brokers represent the advertisers and owners of the media space) and holding a real time auction on the exchange when a media space is accessed by a user. Advertisers specify their bids and targets to the exchange, via the broker, in advance. The exchange determines if the ad available and the user are appropriately matched with the bids offered and targets specified.

[00164] In particular versions, advertiser targeting strategy and correlating bid price is based on information about the user and/or media slot owner, which is considered in behalf of the broker by the exchange at the time the slot is accessed. The information includes one or more of following categories of targeting: contextual targeting, behavioral targeting,

demographic targeting, technology targeting, time and space targeting, those categories of targeting include but are not limited to: advertiser frequency caps, media site owner caps, user time zone, specific time zone, country, state, city, zip or postal-code DMAsite content ratings, site audio effects, site format, other site specific data, user phone type, user audio system; or, ad type, age, gender, income, channel, category, keyword, mobile, desktop app.

[00165] The auction is run for every single impression. The impression is the most granular unit of on-the-fly advertising. Preferably, the auction is held within a very short perceptive interval such as (for example) 100 msec, and if possible much more quickly.

[00166] In other versions of the invention, the exchange supports brokers accepting advertiser payments and paying media slot owners based on differing valuation basis. The valuation bases supported by the exchange include: CPM (cost per thousand impressions), CPC (cost per click-through), CPL (cost per lead), or CPA (cost per action or cost per acquisition).

[00167] Targeted Selection of Streaming Multimedia Objects

[00168] This sub-system of the invention relates to customized electronic identification of desirable objects, such as news articles, or advertisements from an inventory, in an electronic media environment, and in particular to a system that is given or automatically constructs both a "target profile" for each target object in the electronic media based, for example, on the frequency with which each word appears in an article relative to its overall frequency of use in all articles, as well as a "target profile interest summary" for each user, which target profile interest summary describes the user's interest level in various types of tar get objects. The system then evaluates the target profiles against the users' target profile interest summaries to generate a user-customized rank ordered listing of target objects most likely to be of interest to each user so that the user can select from among these potentially relevant target objects, which were automatically selected by this system from the plethora of target objects that are profiled on the electronic media. Users' target profile interest summaries can be used to efficiently organize the distribution of information in a large scale system consisting of many users interconnected by means of a communication network.

[00169] Targeted Selection of Streaming Multimedia Objects: The Problem

[00170] It is a problem in the field of electronic media to enable a user to access information of relevance and interest to the user without requiring the user to expend an excessive amount of time and energy searching for the information. Alternatively, it is a problem to deliver relevant advertising or offers to consumers on a stand-alone basis, or interleaved with other multi-media content.

[00171] Electronic media, such as on-line information sources, provide a vast amount of information to users, typically in the form of "articles," each of which comprises a publication item or document that relates to a specific topic. The difficulty with electronic media is that the amount of information available to the user is overwhelming and the article repository systems that are connected on-line are not organized in a manner that sufficiently

simplifies access to only the articles of interest to the user. Presently, a user either fails to access relevant articles because they are not easily identified or expends a significant amount of time and energy to conduct an exhaustive search of all articles to identify those most likely to be of interest to the user. Furthermore, even if the user conducts an exhaustive search, present information searching techniques do not necessarily accurately extract only the most relevant articles, but also present articles of marginal relevance due to the functional limitations of the information searching techniques. There is also no existing system which automatically estimates the inherent quality of an article or other target object to distinguish among a number of articles or target objects identified as of possible interest to a user.

[00172] Therefore, in the field of information retrieval, there is a long-standing need for a system which enables users to navigate through the plethora of information. With commercialization of communication networks, such as the Internet, the growth of available information has increased. Customization of the information delivery process to the user's unique tastes and interests is the ultimate solution to this problem. However, the techniques which have been proposed to date either only address the user's interests on a superficial level or provide greater depth and intelligence at the cost of unwanted demands on the user's time and energy. While many researchers have agreed that traditional methods have been lacking in this regard, no one to date has successfully addressed these problems in a holistic manner and provided a system that can fully learn and reflect the user's tastes and interests. This is particularly true in a practical commercial context, such as on-line services available on the Internet. There is a need for an information retrieval system that is largely or entirely passive, unobtrusive, undemanding o f the user, and yet both precise and comprehensive in its ability to learn and truly represent the user's tastes and interests. Present information retrieval systems require the user to specify the desired information retrieval behavior through cumbersome interfaces.

[00173] Users may receive information on a computer network either by actively retrieving the information or by passively receiving information that is sent to them. Just as users of information retrieval systems face the problem of too much information, so do users who are targeted with electronic junk mail by individuals and organizations. An ideal system would protect the user from unsolicited advertising, both by automatically extracting only the most relevant messages received by electronic mail, and by preserving the confidentiality of the user's preferences, which should not be freely available to others on the network.

[00174] Researchers in the field of published article information retrieval have devoted considerable effort to finding efficient and accurate methods of allowing users to select

articles of interest from a large set of articles. The most widely used methods of information retrieval are based on keyword matching: the user specifies a set of keywords which the user thinks are exclusively found in the desired articles and the information retrieval computer retrieves all articles which contain those keywords. Such methods are fast, but are notoriously unreliable, as users may not think of the right keywords, or the keywords may be used in unwanted articles in an irrelevant or unexpected context. As a result, the information retrieval computers retrieve many articles, which are unwanted by the user. The logical combination of keywords and the use of wild-card search parameters help improve the accuracy of keyword searching but do not completely solve the problem of inaccurate search results.

[00175] Starting in the 1960's, an alternate approach to information retrieval was developed: users were presented with an article and asked if it contained the information they wanted, or to quantify how close the information contained in the article was to what they wanted. Each article was described by a profile, which comprised either a list of the words in the article or, in more advanced systems, a table of word frequencies in the article. Since a measure of similarity between articles is the distance between their profiles, the measured similarity of article profiles can be used in article retrieval. For example, a user searching for information on a subject can write a short description of the desired information. The information retrieval computer generates an article profile for the request and then retrieves articles with profiles similar to the profile generated for the request. These requests can then be refined using "relevance feedback", where the user actively or passively rates the articles retrieved as to how close the information contained therein is to what is desired. The information retrieval computer then uses this relevance feedback information to refine the request profile and the process is repeated until the user either finds enough articles or tires of the search.

[00176] A number of researchers have looked at methods for selecting articles of most interest to users. An article titled "Social Information filtering: algorithms for automating 'word of mouth'" was published at the CHi-95 Proceedings by Patti Maes et al and describes the Ringo information retrieval system which recommends musical selections. The Ringo system requires active feedback from the users—users must manually specify how much they like or dislike each musical selection. The Ringo system maintains a complete list of users ratings of music selections and makes recommendations by finding which selections were liked by multiple people. However, the Ringo system does not take advantage of any available descriptions of the music, such as structured descriptions in a data base, or free text, such as that contained in music reviews. An article titled "Evolving agents for personalized

information filtering", published at the Proc. 9th IEEE Conf on AI for Applications by Sheth and Maes, described the use of agents for information filtering which use genetic algorithms to learn to categorize Usenet news articles. In this system, users must define news categories and the users actively indicate their opinion of the selected articles. Their system uses a list of keywords to represent sets of articles and the records of users' interests are updated using genetic algorithms.

[00177] A number of other research groups have looked at the automatic generation and labeling of clusters of articles for the purpose of browsing through the articles. A group at Xerox Parc published a paper titled "Scatter/gather: a cluster-based approach to browsing large article collections" at the 15 Ann. Int'l S1GIR '92, ACM 318-329 (Cutting et al. 1992). This group developed a method they call "scatter/gather" for performing information retrieval searches. In this method, a collection of articles is "scattered" into a small number of clusters, the user then chooses one or more of these clusters based on short summaries of the cluster. The selected clusters are then "gathered" into a subcollection, and then the process is repeated. Each iteration of this process is expected to produce a small, more focused collection. The cluster "summaries" are generated by picking those words which appear most frequently in the cluster and the titles of those articles closest to the center of the cluster. However, no feedback from users is collected or stored, so no performance improvement occurs over time.

[00178] Apple's Advanced Technology Group has developed an interface based on the concept of a "pile of articles". This interface is described in an article titled "A 'pile' metaphor for supporting casual organization of information in Human factors in computer systems" published in CHI '92 Conf. Proc. 627-634 by Mander, R. G. Salomon and Y. Wong. 1992. Another article titled "Content awareness in a file system interface: implementing the 'pile' metaphor for organizing information" was published in 16 Ann. Int'l S1GIR '93, ACM 260-269 by Rose E. D. et al. The Apple interface uses word frequencies to automatically file articles by picking the pile most similar to the article being filed. This system functions to cluster articles into sub-piles, determine key words for indexing by picking the words with the largest TF/IDF (where TF is term (word) frequency and IDF is the inverse document frequency) and label piles by using the determined key words.

[00179] Numerous patents address information retrieval methods, but none develop records of a user's interest based on passive monitoring of which articles the user accesses. None of the systems described in these patents pre sent computer architectures to allow fast retrieval of articles distributed across many computers. None of the systems described in these patents

address issues of using such article retrieval and matching methods for purposes of commerce or of matching users with common interests or developing records of users' interests. U.S. Pat. No. 5,321,833 issued to Chang et al. teaches a method in which users choose terms to use in an information retrieval query, and specify the relative weightings of the different terms. The Chang system then calculates multiple levels of weighting criteria. U.S. Pat. No. 5,301,109 issued to Landauer et al. teaches a method for retrieving articles in a multiplicity of languages by constructing "latent vectors" (SVD or PCA vectors) which represent correlations between the different words. U.S. Pat. No. 5,331,554 issued to Graham et al. discloses a method for retrieving segments of a manual by comparing a query with nodes in a decision tree. U.S. Pat. No. 5,331,556 addresses techniques for deriving morphological part-of-speech information and thus to make use of the similarities of different forms of the same word (e.g. "article" and "articles").

[00180] Therefore, there presently is no information retrieval and delivery system operable in an electronic media environment that enables a user to access information of relevance and interest to the user without requiring the user to expend an excessive amount of time and energy. As an extension, there is further very little in the world of targeted advertising and promotion delivery that automates the delivery of relevant advertising to a multimedia consumer.

[00181] Targeted Selection of Streaming Multimedia Objects: The Solution

[00182] The above-described problems are solved and a technical advance achieved in the field by the system for customized electronic identification of desirable objects in an electronic media environment, which system enables a user to access target objects of relevance and interest to the user without requiring the user to expend an excessive amount of time and energy. Profiles of the target objects are stored on electronic media and are accessible via a data communication network. In many applications, the target objects are informational in nature, and so may themselves be stored on electronic media and be accessible via a data communication network.

[00183] Relevant definitions of terms for the purpose of this description include: (a.) an object available for access by the user, which may be either physical or electronic in nature, is termed a "target object", (b.) a digitally represented profile indicating that target object's attributes is termed a "target profile", (c.) the user looking for the target object is termed a "user", (d.) a profile holding that user's attributes, including age/zip code/etc. is termed a "user profile", (e.) a summary of digital profiles of target objects that a user likes and/or dislikes, is termed the "target profile interest summary" of that user, (f) a profile consisting of

a collection of attributes, such that a user likes target objects whose profiles are similar to this collection of attributes, is termed a “search profile” or in some contexts a “query” or “query profile,” (g.) a specific embodiment of the target profile interest summary which comprises a set of search profiles is termed the “search profile set” of a user, (h.) a collection of target objects with similar profiles, is termed a “cluster,” (i.) an aggregate profile formed by averaging the attributes of all tar get objects in a cluster, termed a “cluster profile,”(j.) a real number determined by calculating the statistical variance of the profiles of all target objects in a cluster, is termed a “cluster variance,” (k.) a real number determined by calculating the maximum distance between the profiles of any two target objects in a cluster, is termed a “cluster diameter.”

[00184] The system for electronic identification of desirable objects of the present invention automatically constructs both a target profile for each target object in the electronic media based, for example, on the frequency with which each word appears in an article relative to its overall frequency of use in all articles, as well as a “target profile interest summary” for each user, which target profile interest summary describes the user’s interest level in various types of target objects. The system then evaluates the target profiles against the users’ target profile interest summaries to generate a user-customized rank ordered listing of tar get objects most likely to be of interest to each user so that the user can select from among these potentially relevant target objects, which were automatically selected by this system from the plethora of target objects available on the electronic media.

[00185] Because people have multiple interests, a target profile interest summary for a single user must represent multiple areas of interest, for example, by consisting of a set of individual search profiles, each of which identifies one of the user’s areas of interest. Each user is presented with those target objects whose profiles most closely match the user’s interests as described by the user’s target profile interest summary. Users’ target profile interest summaries are automatically updated on a continuing basis to reflect each user’s changing interests. In addition, target objects can be grouped into clusters based on their similarity to each other, for example, based on similarity of their topics in the case where the target objects are published articles, and menus automatically generated for each cluster of target objects to allow users to navigate throughout the clusters and manually locate target objects of interest. For reasons of confidentiality and privacy, a particular user may not wish to make public all of the interests recorded in the user’s target profile interest summary, particularly when these interests are determined by the user’s purchasing patterns. The user may desire that all or part of the target profile interest summary be kept confidential, such as

information relating to the user's political, religious, financial or purchasing behavior; indeed, confidentiality with respect to purchasing behavior is the user's legal right in many states. It is therefore necessary that data in a user's target profile interest summary be protected from unwanted disclosure except with the user's agreement. At the same time, the user's target profile interest summaries must be accessible to the relevant servers that perform the matching of target objects to the users, if the benefit of this matching is desired by both providers and consumers of the target objects. The disclosed system provides a solution to the privacy problem by using a proxy server which acts as an intermediary between the information provider and the user. The proxy server dissociates the user's true identity from the pseudonym by the use of cryptographic techniques. The proxy server also permits users to control access to their target profile interest summaries and/or user profiles, including provision of this information to marketers and advertisers if they so desire, possibly in exchange for cash or other considerations. Marketers may purchase these profiles in order to target advertisements to particular users, or they may purchase partial user profiles, which do not include enough information to identify the individual users in question, in order to carry out standard kinds of demographic analysis and market research on the resulting database of partial user profiles.

[00186] In the preferred embodiment of the invention, the system for customized electronic identification of desirable objects uses a fundamental methodology for accurately and efficiently matching users and target objects by automatically calculating, using and updating profile information that describes both the users' interests and the target objects' characteristics. The target objects may be published articles, purchasable items, or even other people, and their properties are stored, and/or represented and/or denoted on the electronic media as (digital) data. Examples of target objects can include, but are not limited to: a newspaper story of potential interest, a movie to watch, an item to buy, e-mail to receive, or another person to correspond with. In all these cases, the information delivery process in the preferred embodiment is based on determining the similarity between a profile for the target object and the profiles of target objects for which the user (or a similar user) has provided positive feedback in the past. The individual data that describe a target object and constitute the target object's profile are herein termed "attributes" of the target object. Attributes may include, but are not limited to, the following: (1) long pieces of text (a newspaper story, a movie review, a product description or an advertisement), (2) short pieces of text (name of a movie's director, name of town from which an advertisement was placed, name of the language in which an article was written), (3) numeric measurements (price of a product,

rating given to a movie, reading level of a book), (4) associations with other types of objects (list of actors in a movie, list of persons who have read a document). Any of these attributes, but especially the numeric ones, may correlate with the quality of the target object, such as measures of its popularity (how often it is accessed) or of user satisfaction (number of complaints received).

[00187] The preferred embodiment of the system for customized electronic identification of desirable objects operates in an electronic media environment for accessing these target objects, which may be news, electronic mail, other published documents, or product descriptions. The system in its broadest construction comprises three conceptual modules, which may be separate entities distributed across many implementing systems, or combined into a lesser subset of physical entities. The specific embodiment of this system disclosed herein illustrates the use of a first module which automatically constructs a "target profile" for each target object in the electronic media based on various descriptive attributes of the target object. A second module uses interest feedback from users to construct a "target profile interest summary" for each user, for example in the form of a "search profile set" consisting of a plurality of search profiles, each of which corresponds to a single topic of high interest for the user. The system further includes a profile processing module which estimates each user's interest in various target objects by reference to the users' target profile interest summaries, for example by comparing the target profiles of these target objects against the search profiles in users' search profile sets, and generates for each user a customized rank-ordered listing of target objects most likely to be of interest to that user. Each user's target profile interest summary is automatically updated on a continuing basis to reflect the user's changing interests.

[00188] Target objects may be of various sorts, and it is sometimes advantageous to use a single system that delivers and/or clusters target objects of several distinct sorts at once, in a unified framework. For example, users who exhibit a strong interest in certain novels may also show an interest in certain movies, presumably of a similar nature. A system in which some target objects are novels and other target objects are movies can discover such a correlation and exploit it in order to group particular novels with particular movies, e.g., for clustering purposes, or to recommend the movies to a user who has demonstrated interest in the novels. Similarly, if users who exhibit an interest in certain World Wide Web sites also exhibit an interest in certain products, the system can match the products with the sites and thereby recommend to the marketers of those products that they place advertisements at those sites, e.g., in the form of hypertext links to their own sites.

[00189] The ability to measure the similarity of profiles describing target objects and a user's interests can be applied in two basic ways: filtering and browsing. Filtering is useful when large numbers of target objects are described in the electronic media s pace. These target objects can for example be articles that are received or potentially received by a user, who only has time to read a small fraction of them. For example, one might potentially receive all items on the AP news wire service, all items posted to a number of news groups, all advertisements in a set of newspapers, or all unsolicited electronic mail, but few people have the time or inclination to read so many articles. A filtering system in the system for customized electronic identification of desirable objects automatically selects a set of articles that the user is likely to wish to read. The accuracy of this filtering system improves over time by noting which articles the user reads and by generating a measurement of the depth to which the user reads each article. This information is then used to update the user's target profile interest summary. Browsing provides an alternate method of selecting a small subset of a large number of target objects, such as articles. Articles are organized so that users can actively navigate among groups of articles by moving from one group to a larger, more general group, to a smaller, more specific group, or to a closely related group. Each individual article forms a one-member group of its own, so that the user can navigate to and from individual article s as well as larger groups. The methods used by the system for customized electronic identification of desirable objects allow articles to be grouped into clusters and the clusters to be grouped and merged into larger and larger clusters. These hierarchies of clusters then form the basis for menuing and navigational systems to allow the rapid searching of large numbers of articles. This same clustering technique is applicable to any type of target objects that can be profiled on the electronic media.

[00190] There are a number of variations on the theme of developing and using profiles for article retrieval, with the basic implementation of an on-line news clipping service representing the preferred embodiment of the invention. Variations of this basic system are disclosed and comprise a system to filter electronic mail, an extension for retrieval of target objects such as purchasable items which may have more complex descriptions, a system to automatically build and alter menuing systems for browsing and searching through large numbers of target objects, and a system to construct virtual communities of people with common interests. These intelligent filters and browsers are necessary to provide a truly passive, intelligent system interface. A user interface that permits intuitive browsing and filtering represents for the first time an intelligent system for determining the affinities between users and target objects. The detailed, comprehensive target profiles and user-

specific target profile interest summaries enable the system to provide responsive routing of specific queries for user information access. The information maps so produced and the application of users' target profile interest summaries to predict the information consumption patterns of a user allows for pre-caching of data at locations on the data communication network and at times that minimize the traffic flow in the communication network to thereby efficiently provide the desired information to the user and/or conserve valuable storage space by only storing those target objects (or segments thereof) which are relevant to the user's interests.

[00191] Multi-media Ad Marketplace

[00192] The invention incorporates a real-time, auction-based exchange for online display advertising in which a member of the exchange (usually an advertising network, broker, or agency) buys on the exchange for its advertisers and sells on the exchange for its media slot owners. Generally, the novel exchange represents a highly organized and sophisticated primary or secondary market. Like any advanced marketplace, the exchange and its operators do not participate directly in the marketplace; they create and protect the marketplace. The exchange itself is initially to only benefit from transactions on a flat fee basis, so it doesn't benefit from the value of transactions (or by click fraud as some large Internet players do currently). Its goal is to provide brokers a fair opportunity to exchange advertising. The exchange allows the market (AKA: the members and the exchange participants) to dictate pricing according to perceived supply and demand. The exchange only regulates and provides detailed information to both buyers and sellers to make sure that the all of the buyers are aware of current and historically relevant information, which results in buyers and sellers having the ability to assign value based on all critically relevant information.

[00193] The exchange also functions as the clearinghouse between members, so the members need not worry about credit risk or collection. The exchange is operated on an impression basis (every time an ad appears, or is presented). This allows participants to change and determine value for the unique attributes of every single ad request. This uniquely granular approach provides each member of the exchange and their clients the necessary control to uphold this sophisticated, novel exchange. As shown in FIG. A2, the exchange's clients are members, operating as brokers. The members have advertisers, which are almost always product owners or their agents, who are buying advertising space in accordance with a strategy. And members also have media slot owners, which are almost always retailers or their agents, who are selling advertising space, or spots, within their larger streaming multimedia content framework, such as an IVR. The outlet'smedia streams will most likely

have ads placed in various places throughout the outlet. Preferably, the designated area will show various ad opportunities that are all the same length. The outlet will place a snippet of code (defined herein as an ad spot) in his call stream code for each location on the site he'd like it to appear. That snippet of code only need be unique for each format (ads of certain length), but it may be unique for every location.

[00194] An exchange member's client, an outlet, may have a single call stream or 1000 (for example) within a particular IVR application. He may place a particular spot code (code snippet) for the spot format at the front of every open slot. Thus he would have one ad spot in (say) 1000 different locations. However, the system can track those 1000 different locations (each referred to as an media outlet slots) whenever the "Call Stream's" content is available. This level of granularity in media slot outlet inventory in conjunction with running an auction for every ad impression provides members with unprecedented levels of control. Simply put, an outlet has spots available for advertising, which advertisers compete to obtain. In the novel exchange, the competition takes place in a real-time auction in response to a caller accessing the exchange via an outlet's spot on aoutlet's voice application. The member's advertisers specify in advance how much they are willing to pay based on a user's profile, or the user's past behavior, or the page content, or other factors. Advertisers can bid on a CPM, CPC, CPA, or CPL basis. When a user accessesa streamingmultimedia application, an auction is held among all the advertisers--preferably in less than 100 milliseconds, for every single ad impression. The highest bid for a particular spot wins, and that ad is shown in the spot.

[00195] In order to keep the market fair, balanced, and open, the market provides its members flexibility. Members may possibly allow all, some, or even none of their traffic requests to participate on the exchange or instead just use the exchange technology to match their advertisers with their own media slot owners. Usually, the members specify that the auction favor a match between their own outlets and advertisers, unless the profit of using traffic from the Exchange exceeds a certain threshold. This guarantees members the best possible liquidity--fewer ad campaigns go unfilled, less inventory goes unsold, and members make the highest profit, always.

[00196] The Exchange serves its members (i.e. brokers with seats on the exchange), who in turn serve their advertisers and media slot owners. Advertisers and media slot owners typically do not deal directly with the Exchange, they work through a member. The member's advertisers and media slot owners need not even be aware that their advertising network or broker is trading on the Exchange.

[00197] As shown in FIG. A3, the Exchange is the central marketplace operating at the server level. The member is the advertising network or advertising broker that holds a seat on the exchange and acts on behalf of the broker's clients, the advertiser or outlet, in executing transactions. There are many seats on the exchange, and each member has its own advertisers and outlets. End users are the peopleaccessing the audio applications that the advertisers wish to reach by placing ads on the outlet's sites.

[00198] On the advertiser side of FIG. A2:

[00199] The advertiser is the entity paying a member to run ads. Advertisers are usually product owners or their agents. A member can have any number of advertisers.

[00200] A campaign is a general advertising effort of an advertiser, and may be focused on a specific product line. An advertiser can have any number of campaigns.

[00201] A buy is the budgeting and scheduling level of the campaign structure. The buy controls how much money is spent and when it is spent. A campaign can have any number of buys.

[00202] A strategy is the targeting and bidding level of the campaign structure. The strategy level is also where the creatives (the ads themselves) are assigned and grouped. A buy can have any number of strategies, and any number of creatives.

[00203] On the media slot owner or outlet side of FIG. A2:

[00204] The media slot owner is the entity that owns the streaming multimedia application that accepts advertisements.

[00205] An outlet can have any number of audio applications.

[00206] The ad spot is the snippet of configuration code on an audio application defining where an ad can appear. It is typically created by the outlet (using the exchange interface) and then is placed in at least one location on one page on one of the outlet sites. The ad code specifies the format (length), acceptable content, and other parameters. A spot is always a subset of the outlet, and there can be any number of spots on an outlet.

[00207] The slot is the specific location in an audio application where an ad or piece of streaming can appear (begin to play). It is typically defined by the ad code created by the outlet (within the exchange interface) and placed at that location. The ad code specifies the format (length), acceptable content, and other parameters. The length (size)of any candidate media spot is always a subset of the slot, and there can be any number of objects in an ad slot.

[00208] A possible transaction on the Exchange is illustrated in FIG. A3. There are three levels of participation in the exchange. The exchange server level is the exchange itself. Brokers are the members of the exchange, and only members communicate with the

exchange directly. The inventors have a working exchange, (herein called by "CallSpaceXN") which is used herein to illustrate the novel exchange. The clients are the advertisers and outlets who communicate with the brokers. Outlets, using their member's Exchange interface, place their ad spot inventory in the exchange. Advertisers, using their member's Exchange interface, specify in advance the targeting they want and how much they are willing to pay. In the figure, an end user/site visitor accesses an audio application owned by an outlet in Member Bravo's network. The media slot on the outlet contains ad computer code, which makes annetworkrequest to the CallSpaceXN exchange when the visitor's accesses a CallStream. This triggers a single-pass auction among all of the interested advertisers within the CallSpaceXN exchange from both the outlet member's advertisers (the private network) and all other qualified members. The auction preferably takes less than 100 milliseconds. In every transaction, the exchange determines the bid within the private network (in which the member supplies both the advertiser and outlet) that benefits the publishing member most according to price and the bidder from all of the other members of the exchange (in which the member only supplies either the advertiser or the outlet) that benefits the outlet member most according to price. Then before declaring only one bidder the winner, the preferences of the outlet member are examined, and the transaction that benefits that member most, according to his setting (which could be price or something else), is declared winner.

[00209] Transaction Flow

[00210] In FIG. A4, the flow of a single transaction is illustrated:

[00211] A. A visitor accesses a specific CallStreamwithin a specific audio application, which Bravo outlet represents at (point 1 in FIG. A4).

[00212] B. A request is made to Bravo outlet to deliver content. (point 2)

[00213] C. Bravo outlet requests CallSpaceXN to return anCallSpaceXN ad if the minimum CPM is met or exceeded. (point 3).

[00214] D. The ad director server group immediately references all CallSpaceXN data with additional correlated data from RAM about the viewer and/or the page the viewer has just accessed. That RAM and storage may be in a group of machines dedicate to providing the data, Information Provider Cache Group (IPCs). The ad director may also filter out which bidders it would like to hear from in the auction (point 4). The ad director then opens the auction for bidding and sends a multicast message, including all of the auction information for this auction (point 3) which includes user information and site information, to the selected bidder server groups (point 5).

[00215] E. The bidder server group responds to the ad director server group with all bids from advertisers bidding with qualified criteria (point 6). Bids may come from Member bravo, alpha up to member N, basically any interested member of the exchange.

[00216] F. The ad director group declares a winner (point 7) according to the publishing members preset values. Then, depending on the advertiser's creative's settings, CallSpaceXN sends a request to either the CallSpaceXN provided ad serving option (which may include) a content distribution network (CDN) (point 8) or a third-party ad server to serve the ad, and then logs the winning advertiser's information, the outlet's information (include site, spot, context, and time) for the given impression, and all known user information. This transaction occurs for every single impression in the network.

[00217] FIG. A5 outlines an executed transaction. However prior to the exchange running any transactions the participating members must have completed the following:

[00218] 1. Seat Holder/Member Bravo signed up as a member in the novel exchange. As a part of sign-up they made specifications, including the definition of the highest bidder. The highest bidder may be the bidder which results in the highest payment to the member or, according to the outlet's settings; it may be the bidder that results in the highest payment to Bravo Outlet 1.

[00219] 2. The bravo advertiser 1 made his bid for all of his targeting strategies in advance. He placed his bid using the user interface provided by member Bravo, which interface may be provided by CallSpaceXN. The UI accesses and updates the CallSpaceXN database. The CallSpaceXN database updates the bidder server group upon change.

[00220] 3. The Bravo outlet adds CallSpaceXN ad code to his website such that web requests are made when an end-user/site visitor lands on a page, aCallSpaceXN request is made. The outlet also specifies his default category and the minimum CPM he's willing to accept for any given impression on that spot.

[00221] In FIG. A5, the flow of a single transaction is illustrated. The user visited Seat Bravo's outlet, which triggered the auction. The auction was run and advertiser bidding the amount that benefited Seat Bravo the most was Alpha's advertiser. The advertiser paid $1.00 CPM for the impression. The two members took the margin percentages that they predetermined and set (in each case the members used the system default setting of an 18% margin on the publishing members account and 17% margin on the advertising members), which means the outlet received the remaining $0.65. The exchange preferably tracks all useful details of every auction, provides complete accounting to its members, and handles collections and payments from the members to the Exchange. Members could have specified

the following exchange participation settings which could alter the supply and/or demand for each impression and adjust the mechanics of auction:

[00222] Keep this media slot owner private. This selection is made by a member acting as the outlet member. This requires the exchange to hold a private auction for the member's own advertisers on this particular site in order to keep the transactions on this site only private within the Bravo network.

[00223] Keep all of mymedia slot owners private. None of the other members' advertisers will be allowed to bid on an impression from any of the member's media slot outlets.

[00224] Keep all of my media slot owners and all of my advertisers private. This requires the exchange to exclude all other members from bidding or accepting bids from the member. This means this member has created an exclusive economy (see FIG. A6-private network).

[00225] Keep this advertiser private. This selection is made by a member acting as the advertising member. This requires the exchange only permit this specific advertiser to bid on impressions the member's media slot outlets make available. The advertiser will never bid on an impression from a media slot outlet provided from another member.

[00226] Keep all of my advertisers private. None of the member's advertisers will ever bid on an impression from any media slot outlet provided from other members.

[00227] Other arrangements can be supported as well.

[00228] A primary benefit to members of the Exchange is liquidity: fewer campaigns go unfilled, and less inventory goes unsold. At a macro level, since all parties (advertisers, outlets, and brokers) involved are more efficient than current models commonly allow, The market prospers and produces more value, and ultimately end-users are more satisfied because they see more relevance and value in the ads they are shown.

[00229] While staying neutral, the Exchange supports methods to generate revenue for itself as well. A member preferably must buy his seat on the Exchange. This can be a one-time, upfront payment that will cover the cost of the additional hardware and interface customization needed to serve the member. The cost of the seat can vary, based on the member's projected volume and quality of its traffic. The Exchange can also charge all members a small transaction fee for every auction it runs. The inventors have found a transaction fee varying from under $0.0005 CPM to $0.15 CPM based on the volume of transactions from the member is sustainable. If one member provides the buyer and another member provides the seller, the fee can be split between the two members. If a single member provides both the buyer and the seller, that member pays twice the fee. The Exchange

preferably should not participate on a percentage or revenue share basis in the transaction, because it advantageously should remain a disinterested, neutral marketplace for its members.

[00230] This novel exchange also creates novel sub exchange for information and user data. In step D mentioned insectionTransaction Flow, information providers can provide data to the exchange that can be used by advertisers on demand.

[00231] The novel exchange can partner with any company that has user data available. It can be the recent search history or other audio application behavioral data. It can also be information about the user, such as demographic information.

[00232] In advance, CallSpaceXN creates an agreement with information providers for a specific revenue share. For instance, CallSpaceXN agrees to pay the information provider 20% of the winning bids for all auctions that use the personal data. Then, an Advertiser A specifies in advance of any given auction that he'd like to target his ad exclusively to a specific demographic or behavioral group, say 34 year-old females. When the auction is run, in the CallSpaceXN system, the ad director requests all information available for this site and user for this impression from the IPC group. The auction multi-cast then sent to the bidders contains all available user data. In this case, Advertiser A will only bid on auctions in which it is known the user is both 34 and female.

[00233] In this sub market, there is economic pressure on the information providers to move prices downward because if the same information is available from two or more information providers the system will only use and pay for the data that is least expensive. For instance, using our example, if two information providers had made available to the exchange that a particular user was a 34 year-old female and one provider required a 25% revenue share and the other required a 15% revenue share then the novel exchange would only use the data from the provider with the 15% revenue share.

[00234] In this sub market there may be instances when two or more information providers are required to fulfill a given bid on a given impression. For instance, there may be one information provider that provides that a given user is female, but another information provider may provide that the same user is 34 years old. In that case, all of the available information is made available in the auction for bidding, but each information providers whose information is used in the winning bid will get their revenue percentage divided by the number of information providers. For example, in an auction where the winning bid required the uses of data, and one provider contracted at a revenue share of 15% supplies the age of 34 and another information provider contracted at a rate of 25% supplies the gender. The first

information provider would receive 15%/2 and the other information provider would receive 25%/2, essentially 12.5%.

[00235] The exchange also improves the collection process for advertisers and outlets. Currently, when Internet publishers, for example, sell ad space through brokers, the media slot owner generally doesn't get paid until the advertiser pays the broker. Dealing only through members and instituting a regular billing/disbursement cycle, Bills are sent and invoices paid by the exchange. Since the exchange has leverage on the advertiser side, the collection process works smoothly and predictably.

[00236] An important element in the Exchange based market place is the advertiser campaign. The Exchange uniquely takes advantage of the capabilities the CallSpaceXN can provide to achieve targeted advertising. The inventors have offered the following capabilities on the Exchange for advertisers:

[00237] Create one or more advertising campaigns. It is at the campaign level that an advertiser can exclude specific outlets.

[00238] Create at least one buy for each campaign. A buy is where the advertiser can allocate budget dollars and define start and end dates.

[00239] Create one or more "strategies" for each buy. A strategy is where the advertiser defines the matching criteria he wants.

[00240] To each strategy they assign one or more ads, referred to as a "creative" or "creatives". They can have any number of strategies for a buy.

[00241] Turn on a campaign and "go live," and then watch their results in real-time, fine-tuning their strategies as needed.

[00242] Targeting is a particularly useful tool that takes advantage of the information (much like an Internet "cookie") passed to an audio application about a user when the user accesses that audio application. Examples of the kind of targeting capabilities that can be used are shown in FIG. A7. Advertisers can specify in advance just what sort of advertising opportunity they want. The matching strategies they can use include the following targeting categories:

[00243] Contextual-advertisers target based on the context of the content of the audio application where the ad may be played. There are at least four levels of contextual relevance, including specific audio outlet, channels, a variety of categories, and by any number of keywords.

[00244] Behavioral-advertisers target based on the audio outlet's visitor's recent behavior or activity. For example, advertisers can select viewers based on the viewer's recent search engine queries and enter words or phrases to match.

[00245] Demographic-advertisers target based on the demographic information of the user, such as age, gender, income, and geography.

[00246] Technology-advertisers target based on the technology use and preferences of the viewer. For example, advertisers can select viewers based on their phones or audio players (VOIP, PSTN, Cell Phone, Phone OS)

[00247] Time and space-advertisers target based on the geographic location and time of day of the viewer or advertiser. Frequency-capping also prevents any one viewer from seeing any one ad too many times in a day.

[00248] The Exchange is developed with modular targeting capabilities, so individual targeting methods can constantly be added under the categories of targeting. The individual targeting methods include but are not limited to:

[00249] Advertiser frequency caps — is control performed by advertiser or member which limits the number of times that any single user can view any certain buy within a certain period. For example a 2 per 24 cap means that the advertiser will only show the ad 2 times per user per 24 hours;

[00250] Media slot outlet or publisher caps — is a control performed by the outlet member which limits the number of times a user can receive any ads from a specified spot. For example a 1 per 24 cap means that a spot will only show 1 ad per 24 hours;

[00251] User time zone — advertisers control if a buy runs according to the end-users time zone. For example, advertiser targeting business-to-business products 9:00 am to 4:00 pm will show the ad to a user in London at 9:00 am to 4:00 pm GMT and likewise a user in Los Angeles at 9:00 am to 4:00 pm PST;

[00252] Specific time zone — advertisers control if a buy runs according to the advertisers specified time zone. For example, an advertiser requesting end-users to respond to a call center that only has operators standing by in New York City from 9:00 am to 4:00 pm. The ad will show to users in New York City from 9:00 am to 4:00 pm EST and likewise a user in Los Angeles at 6:00 am to 1:00 pm PST;

[00253] Country, State, city, postal code or DMA — advertisers control which area's end-users are eligible to view the ad based on the current IP address of the user;

[00254] Advertiser content ratings — outlet can exclude ads based on ratings such as Language, Nudity, Violence, Animation, Alcohol, Audio, Dating/Romance, Expandable,

Gambling, Guns, Network, Political, Sex/Diet Drugs, Tobacco, and/or Video which can be determined by the advertiser, member, or in some cases by the system;

[00255] Site content ratings — advertisers can exclude their ads from appearing on audio applications that contain certain content based criteria such as Language, Nudity, Violence, Animation, Alcohol, Audio, Dating/Romance, Expandable, Gambling, Guns, Network, Political, Sex/Diet Drugs, Tobacco, and/or Video;[Format— Any number of audio or video formats.

[00256] Rich media—3D immersive environments and on-line games are ripe for on-demand product placement

[00257] Other Audio Application and Outlet Specific Data;

[00258] User Audio Device--advertisers can target ads according to the audio device through which the end-users access the audio application and the ads;

[00259] Age--advertisers can target the age of the specific user viewing the ad.

[00260] Gender--advertisers can target the gender of the specific user viewing the ad.

[00261] Income--advertisers can target the income of the specific user viewing the ad.

[00262] Channel--the advertisers can specify that ads will only play through audio outlets that fall within a certain channel. Eventually audio applications will be divided into over two dozen channels. For example, an advertiser may specify that he will only play his ad within audio applications that are related to sports and recreation.

[00263] Categories--the advertisers can specify that ads will only play in audio applications that fall within a certain subset of the channels. All websites are divided into approximately hundreds of categories. For example, an advertiser may specify that he will only play his ad in audio applications that are related to sports & recreation: football.

[00264] Keyword--advertisers can specify that ads will only play in audio applications that are related to a certain keyword or group of keywords. For example, the advertiser can specify that the ads will only play in audio application CallStreams containing the keyword of “mortgage” and “refinance”.

[00265] As described above, the targeting allows the advertiser to buy a single impression, which is matched to the person or the specific content or circumstance of that impression: the context of the page, the past behavior of the end user, the profile of the end user, the geographic location of the end user, the time of day, the number of times the end user has already heard the ad, and so on and on. The advertiser can specify the impression he wants, and at the same time specify (bid) what it is worth to him. When it gets that impression, he knows what he got, and he knows that he did not overpay.

[00266] Another novel capability the Exchange system can offer is a guarantee that an advertiser will not overpay for under-performing inventory. The Exchange can contain an application that tracks the performance of every single ad spot--every location for every ad on every CallStreamin every audio application in the Exchange. Just before an auction, this application can check the history of the spot. If the spot performs about as well as other spots in its category, the advertiser's bid is made. However, if the spot under-performs its peers, the advertiser's bid is reduced accordingly, to reflect the relative value of that spot.

[00267] The Exchange further offers flexibility in revenue generation for all parties as well as risk reduction. A broker can make an additional profit by letting his advertisers pay on one basis, but paying his outlets on another. This is called "arbitrage."

[00268] Many advertisers pay on a CPM (cost per thousand) basis, where they simply pay a certain amount for every thousand impressions. Outlets (Analogous to publishers in web media) are almost always paid on a CPM basis. But some advertisers prefer to shift the risk from themselves to the outlet or the Exchange and pay only on a CPC (cost per click) basis, where the advertiser pays only when an ad is clicked on; a CPL (cost per lead) basis, where an advertiser pays only when a lead is gotten; or on a CPA (cost per acquisition) basis where an advertiser pays only when a sale is generated. Since outlets are almost always paid on a CPM by the member, the member assumes the risk of number of sales, leads, or clicks, not being sufficient to make a profit.

[00269] A shrewd broker may be willing to sell an advertiser this "insurance policy" thus distributing his risk to the member. The member may accept payment on a $25.00 CPA basis, for example, and show some number of ads and pay the outlets on a CPM basis. If the number of ads he shows costs less than $25.00 for every action that happens, the broker keeps the difference between the $25.00 he gets from the advertisers and whatever he had to pay the media slot outlet. The challenge, for the broker, is to know where to run those ads and what to pay for them. If done so successfully, the member can benefit with higher than usual margins.

[00270] Advertisers often want to avoid spending their entire daily budget too quickly. Metering is the process of trying to distribute spending evenly throughout the life of the buy. Metering is part of the Exchange.A first level insures that an advertiser's budget is evenly spent. This first level of metering limits the number of times an advertiser's ad is shown in order to spread his budget evenly over time. It its most basic form, metering simply turns the bidding off and on to maintain the Exchanges default spending versus evenness ratio.A

higher, advanced level of metering gives the advertiser more control over the flow of impressions by letting him set two variables.

[00271] Syndicated Messaging Architecture and System for IVR Applications: Preferred Embodiments

[00272] The present invention is described herein by way of reference to a particular example, and in particular to the example of a unified/voice messaging system. However the invention is not limited in its applicability to the specifics of the examples and embodiments described herein. The primary purpose of this component of the larger system is to provide best of class unified messaging to world class organizations while inserting the ability to manage ad insertion to callers.

[00273] Referring to FIG. B3 there is shown a unified/voice messaging system deployed in a geographically distributed manner and implemented in accordance with a preferred embodiment of the present invention.

[00274] At a first distributed or remote location there is provided a PBX 302 , connected to an associated voice or unified messaging front-end 306 . At a second distributed or remote location there is provided a PBX 304 , connected to an associated voice or unified messaging front-end 308 . Each of the voice or unified messaging front-ends 306 or 308 may be considered to be remote or distributed voice or unified messaging systems. Each of the remote voice or unified messaging systems 306 or 308 is connected to a data network 310 , through which there is provided a connection to a common, centralized voice or unified messaging system back-end 312 . The voice or unified messaging system back-end 312 may be considered to be a centralized voice or unified messaging system, or a centralized data facility. The centralized voice or unified messaging system 312 is, in the described embodiment, further provided with a 'clustered' front end, which provides, as will be described further herein below, front-end functionality equivalent to the functionality provided by the front-end voice or unified messaging system 306 and 308 . The 'clustered' front end is connected to a PBX 314 .

[00275] The voice or unified messaging front-end 306 effectively defines a first voice mail domain VMD# 1 . All users associated with the PBX 302 belong to VMD# 1 . As is discussed further herein below, certain of the voice or unified messaging functionality for VMD# 1 is also provided centrally by the voice or unified messaging back-end 312 . As such, the voice or unified messaging back-end 312 forms part of VMD# 1 .

[00276] The voice or unified messaging front-end 308 effectively defines a second voice mail domain VMD# 2 . All users associated with the PBX 304 belong to VMD# 2 . As is discussed further herein below, certain of the voice or unified messaging functionality for VMD# 2 is also provided centrally by the voice or unified messaging back-end 312 . As such, the voice or unified messaging back-end 312 forms part of VMD# 2 .

[00277] In the preferred embodiment of the present invention there is further defined a third voice mail domain VMD# 3 . All users associated with PBX 314 belong to VMD# 3 . All of the voice or unified messaging functionality for VMD# 3 is provided centrally by the voice or unified messaging back-end 312 and clustered front-end 316 .

[00278] Thus, in the embodiment of FIG. B3, users associated with remote PBX 302 are part of VMD# 1 , users associated with remote PBX 304 are part of VMD# 2 , and users associated with PBX 314 are part of VMD# 3 . Further remote locations may exist defining further voice mail domains, and individual voice mail domains may comprise multiple voice or unified messaging systems.

[00279] The functionality of the architecture of the system of FIG. B3 in accordance with the preferred embodiment of the invention is now further described.

[00280] The front-end voice/unified messaging functionality of the voice or unified messaging front-ends 306 and 308 provide certain voice/unified messaging functions local to the respective PBX 302 and 304 for the respective voice mail domains independent of access to the centralized voice/unified messaging system back-end 312 . The clustered front-end 316 provides certain voice/unified messaging functions for users of VMD# 3 . The back-end 312 provides certain voice/unified messaging functionality for all voice mail domains.

[00281] It should be noted that although the embodiment specifically discloses the provision of PBXs at remote locations, more generally the remote locations can be considered to be provided with a switching capability, or telecommunications switch. The invention is not limited to PBX implementations.

[00282] The front-end or remote system functionality may include, for example, call answering and automated attendant functionality. These front-end or remote systems may also provide subscriber access capability. However the implementation of subscriber access may depend on the data networking characteristics, i.e. quality of service (QoS) of the connection from the front-end to the back-end.

[00283] The implementation of the centralized data facility or voice/unified messaging system back-end system or central system of FIG. B3 in accordance with an example embodiment of the invention is shown in FIG. B4.

[00284] Referring to FIG. B4, in a preferred embodiment the voice or unified messaging system back-end 312 comprises a message store 402 , a directory 404 , and a plurality of servers 406 .

[00285] In FIG. B4 there is shown four front-end servers 406 a to 406 d. The number of servers 406 is implementation dependent, and depends upon the overall capacity of the system needed to be supported. Additional servers may be provided to allow for redundancy to improve system reliability. The servers provide access to the voice or unified messaging system functionality for VMD# 3 , accessing the system via PBX 314 .

[00286] For providing centralized voice or unified messaging functionality, the centralized voice or unified messaging system is configured for all voice mail domains. As shown in FIG. B4 the centralized voice or unified messaging system is provided with a message store 402 and a directory 404 . The message store 402 stores messages, such as voice mails and e-mails, associated with system users for all voice mail domains. The directory 404 stores descriptive attributes associated with those system users, and system configuration data. As shown in FIG. B4, the directory 404 stores VMD objects defining system configuration data for each voice mail domain. The directory 404 additionally stores subscriber objects, defining the voice mail domain for each of the system subscribers.

[00287] As shown in FIG. B4, there is provided a central group of servers 406 in the centralized system. The servers provide access to the voice or unified messaging system functionality.

[00288] The implementation of the distributed voice/unified messaging system front-end or remote system of FIG. B3 in accordance with an example embodiment of the invention is shown in FIG. B5.

[00289] Each PBX, such as PBX 302 , is associated with a remote voice server 504 , which may also be considered to be a remote front-end voice server. As can be seen, the remote voice server 504 is connected to the data network 310 . The remote voice server 504 is associated with a cache message store 506 and a cache directory store 508 , illustrated in FIG. B5. In accordance with a preferred embodiment of the present invention, and as discussed further herein below, information associated with the users connected to the PBX 302 , and stored in the corresponding message store 402 and directory store 404 of FIG. B4, are copied to the caches 506 and 508 of FIG. B5, such that a local copy of such information is available. As will be further discussed herein below, remote voice or unified messaging functionality is provided by the front-end voice server 504 and caches 506 and 508 , generally designated by reference numeral 502 and considered to be the voice or unified messaging front-end. It

should be further noted that each PBX may be associated with more than one front-end voice server 504 .

[00290] The invention provides, in a first embodiment, mechanisms to allow call answering and automated attendant functionality to operate when supported by low-guarantee network connectivity from the messaging system back-end. This is achieved by providing the necessary functionality at the front end. As such the provision of the functionality is not dependent upon any permanently available back-end or central connectivity. This is achieved by providing the necessary data, logic and control parameters to achieve such functionality in the caches. The back-end or central systems provide the primary message storage and directory storage and configuration data. All data, messages, user properties and system configuration information is primarily stored at the back-end or central system.

[00291] The invention provides in a second embodiment a mechanism for reliable subscriber access through direct connectivity to a large 'clustered' back-end system, in the centralized data storage facility, as provided for in FIG. B4.

[00292] In the following description, the invention is considered in two parts, although it should be understood that much of the functionality is shared. The invention is also described by way of reference to examples of specific feature implementations, but is not limited to such.

[00293] In a first embodiment, there is described in accordance with a preferred implementation of the invention the implementation of a call answering function and an automated attendant function supported remotely from the distributed centralized data facility, via a low QOS data network, as illustrated in FIGS. B4 and B5. This relates, more generally, to a description of those functions, which can be provided independent of access to the back-end.

[00294] Call answering is a relatively simple process with only limited requirement for access to user and system data. The call answering process requires access to data of the following classes:

[00295] I. System properties, including system configuration data such as language availability; and

[00296] II. User properties, such as user configuration data such as 'find-me' rules and user status (e.g. extended absence), as well as associated audio greetings.

[00297] The front-end systems associated with each PBX need to maintain a cached copy of this data to allow a high degree of quality and reliability in the call answering process regardless of the connectivity to the back-end.

[00298] Preferably a cache management process operates in the background (at a pre-defined time interval, e.g. 5 minutes) and creates a local (and persisted) cache of both system and user properties from the back-end primary stores. As shown in FIG. B4, the directory includes voice mail domain (VMD) objects and subscriber objects. For VMD objects the system configuration parameters for each voice mail domain are defined. For subscriber objects the VMD identity for each subscriber is defined. This information is cached at the remote system. The cache management is preferably provided by the front-end or remote system. A time-stamp is preferably stored for user greetings in the message store. Before an audio greeting is played, the front-end system may check the time of last-update on the back-end primary version of the greeting. If the front-end cached version of a greeting is correct, then it is used. Otherwise a copy may be fetched from the back-end. In general, call answering results in recording an audio message. This can be submitted for delivery in the background with little real-time network capability.

[00299] Automated attendant, when supported by a cache of user and system properties as shown in FIG. B5, can also work in the local front-end system.

[00300] It is also possible to 'dynamically cache' information to reduce the effects of network latency and 'smooth-off' bursts of high network bandwidth requirements. As an example, the n+1th message can be pre-fetched when the nth message is played. This serves, to a point, to provide improved operation on poor networks but cannot cover all real-world network (e.g. total WAN failure) conditions.

[00301] Subscribers of VMD# 3 access voice or unified messaging functionality in a conventional manner, all functionality being provided by the primary data stored in the centralized voice or unified messaging system back-end 312 .

[00302] As discussed above, voice messaging or unified messaging functionality in VMD# 1 and VMD# 2 is preferably provided by the localized cached information. Where functionality requires direct centralized access, this can be provided by access to the centralized data facility via the data network. Such access may be required for certain subscriber access functions.

[00303] In a second embodiment, there is described in accordance with a preferred embodiment of the invention the implementation of a subscriber access function for all voice mail domains supported at a shared centralized front-end system, as illustrated in FIG. B4, whilst exhibiting the same user characteristics as if calling a local system in the arrangement of FIG. B1. The advantage of such an arrangement is that subscriber access functions may be provided if the link for the remote systems to the centralized data facility is not available.

[00304] In the preferred embodiment, all subscribers of all voice mail domains may call into PBX 314, which may have a toll-free number associated therewith, for subscriber access functions. The system of FIG. B3 may be an enterprise system of a single organization, and the toll-free number may be an enterprise wide number for all voice mail domains within the enterprise.

[00305] Existing voice mail systems contain considerable system configuration data, which defines the operation of the system. This information, in the Avaya Unified Messenger and MMA systems, is stored within a directory object called a voice mail domain object, in the directory of FIG. B4. In the preferred embodiments of the present invention, as discussed above, the front-end systems maintain a cache of this information, while the back-end systems manage the primary copy of the configuration data. In existing systems, and as mentioned hereinabove, multiple front-end systems can be part of a single voice mail domain.

[00306] However, critically, in order to provide centralized subscriber access, the front-end systems 316 co-located with the centralized data stores 312 need to be able to be part of multiple voice mail domains. A call coming into the clustered front-end 316 may be associated with any VMD, and there is therefore a need to handle the call appropriately for any VMD. The single (large) centralized system of FIG. B4 must be able to operate as part of all voice mail domains for which subscriber mailboxes are hosted. All voice mail domains are defined by directory objects stored in the directory systems within the back-end systems and as a consequence this configuration data is available to both local and centralized front-end systems. This access to the configuration data of all voice mail domains is provided by the network interface of FIG. B4, and stored locally in the caches of FIG. B5.

[00307] Subscribers calling into the large centralized front-end system (via the toll-free PBX access) must be identified. This identification can be based, for example, on various options:

[00308]

[00309] a) Identification of the called number. This works if the subscribers of each voice mail domain are given a different number to call (possibly an 800 'toll-free' number). In such case their locally valid mailbox number is enough to identify the subscribers, as their voice mail domain can be known. The centralized voice servers may therefore include a look-up table matching dialed numbers to voice mail domains.

[00310] b) By logging into the system using an enterprise/organization wide numeric address, and not their usual mailbox/extension number (which is only valid in conjunction

with their voice mail domain). This address allows the user to be identified, from which a VMD can be identified.

[00311] c) By identifying and recognizing the caller's number, for example the subscribers cell phone, home phone or work phone numbers may be stored in a look-up table.

[00312] The clustered or centralized voice servers 406 of FIG. B4 can be adapted to provide identification of the calling party as described above to interpret the call for the appropriate voice mail domain.

[00313] Once the voice mail domain (VMD) and the subscriber are identified the centralized front-end system of FIG. B4 needs to provide an interaction, which is identical to that experienced by a subscriber calling into a local front-end system of FIG. B1. The front-end system 306 thus effectively shifts its modus operandi based on the voice mail domain information, i.e. the system operates as if it were the voice mail domain identified. All system parameter information is available and used to provide the required interface.

[00314] Once logged on, subscribers are presented with the correct addressing options. Addressing a message to a user by their mailbox/extension number must work correctly. A user within location 1 (e.g. voice mail domain VMD# 1) addressing a message to "4003 ", needs to have this resolved to the 4003 mailbox (at location 1). A user from location 2 (e.g. voice mailbox domain VDM# 2) may address to the same number and must have this resolved to a different mailbox. The provision of the clustered front-end 316 of FIG. B3 and FIG. B4 provides for such features to work correctly.

[00315] Certain voice/unified messaging scenarios result in the system launching outbound calls. To allow centralized systems to perform this correctly, telephone numbers must be stored in their canonical form. Both types of front-end systems (distributed and centralized) are preferably configured with rules to ensure calls are correctly dialed.

[00316] In summary, therefore, in embodiments the present invention provides for the provisions of distributed messaging system functionality and centralized messaging system functionality, in which the distributed functionality is provided by cached stores from the centralized system, and in which users of the distributed systems may access certain functionality through direct access to the centralized system.

[00317] Targeted Selection of Streaming Multimedia Objects

[00318] Measuring Similarity

[00319] This section describes a general procedure for automatically measuring the similarity between two target objects, or, more precisely, between target profiles that are automatically generated for each of the two target objects. This similarity determination

process is applicable to target objects in a wide variety of contexts. Target objects being compared can be, as an example but not limited to: textual documents, human beings, movies, mutual funds or multi-media streaming advertising. It is assumed that the target profiles, which describe the target objects, are stored at one or more locations in a data communication network on data storage media associated with a computer system. The computed similarity measurements serve as input to additional processes, which function to enable human users to locate desired target objects using a large computer system. These additional processes estimate a human user's interest in various target objects, or else cluster a plurality of target objects in to logically coherent groups. The methods used by these additional processes might in principle be implemented on either a single computer or on a computer network. Jointly or separately, they form the underpinning for various sorts of database systems and information retrieval systems.

[00320] Target Objects and Attributes

[00321] In classical Information Retrieval (IR) technology, the user is a literate human and the target objects in question are textual documents stored on data storage devices interconnected to the user via a computer network. That is, the target objects consist entirely of text, and so are digitally stored on the data storage devices within the computer network. However, there are other target object domains that present related retrieval problems that are not capable of being solved by present information retrieval technology, which are applicable to targeting of articles and advertisements to readers of an on-line newspaper:

[00322] (a.) the user is a film buff and the target objects are movies available on videotape.

[00323] (b.) the user is a consumer and the target objects are used cars being sold.

[00324] (c.) the user is a consumer and the target objects are products being sold through promotional deals.

[00325] (d.) the user is an investor and the target objects are publicly traded stocks, mutual funds and/or real estate properties.

[00326] (e.) the user is a student and the target objects are classes being offered.

[00327] (f.) the user is an activist and the target objects are Congressional bills of potential concern.

[00328] (g.) the user is a net-surfer and the target objects are links to pages, servers, or newsgroups available on the World Wide Web which are linked from pages and articles on-line newspaper.

[00329] (h.) the user is a philanthropist and the target objects are charities.

[00330] (i.) the user is ill and the target objects are ads for medical specialists.

[00331] (j.) the user is an employee and the target objects are classifieds for potential employers.

[00332] (k.) the user is an employer and the target objects are classifieds for potential employees.

[00333] (j.) the user is a lonely heart and the target objects are classifies for potential conversation partners.

[00334] (m) the user is in search of an expert and the target objects are users, with known retrieval habits, of an document retrieval system.

[00335] (n) the user is in need of insurance and the target objects are classifieds for insurance policy offers.

[00336] In all these cases, the user wishes to locate some small subset of the target objects—such as the target objects that the user most desires to rent, buy, investigate, meet, read, give mammograms to, insure, and so forth. The task is to help the user identify the most interesting target objects, where the user's interest in a target object is defined to be a numerical measurement of the user's relative desire to locate that object rather than others.

[00337] The generality of this problem motivates a general approach to solving the information retrieval problems noted above. It is assumed that many target objects are known to the system for customized electronic identification of desirable objects, and that specifically, the system stores (or has the ability to reconstruct) several pieces of information about each target object. These pieces of information are termed "attributes": collectively, they are said to form a profile of the target object, or a "target profile." For example, where the system for customized electronic identification of desirable objects is activated to identify selection of interest, a particular category of on-line products for review or purchase by the user, it can be appreciated that there are certain unique sets of attributes which are pertinent to the particular product category of choice. For the application as part of a movie critic column (where the system identifies movie titles and reviews which are most interesting to the users), the system is likely to be concerned with values of attributes such as these:

[00338] (a.) title of movie,

[00339] (b.) name of director,

[00340] (c.) Motion Picture Association of America (MPAA) child-appropriateness rating (0=G, 1=PG, . . .),

[00341] (d.) date of release,

[00342] (e.) number of stars granted by a particular critic,

[00343] (f.) number of stars granted by a second critic,

[00344] (g.) number of stars granted by a third critic,

[00345] For example, a customized financial news column may be presented to the user in the form of articles which are of interest to the user. In this case, however, an accordingly those stocks which are most interesting to the user may be presented as well.

[00346] (h.). full text of review by the third critic,

[00347] (i.). list of customers who have previously rented this movie,

[00348] (j.) list of actors.

[00349] Each movie has a different set of values for these attributes. This example conveniently illustrates three kinds of attributes. Attributes c-g are numeric attributes, of the sort that might be found in a database record. It is evident that they can be used to help the user identify target objects (movies) of interest. For example, the user might previously have rented many Parental Guidance (PG) films, and many films made in the 1970's. This generalization is useful: new films with values for one or both attributes that are numerically similar to these (such as MPAA rating of 1, release date of 1975) are judged similar to the films the user already likes, and therefore of probable interest. Attributes a-b and h are textual attributes. They too are important for helping the user locate desired films. For example, perhaps the user has shown a past interest in films whose review text (attribute h) contains words like "chase,""explosion,""explosions,""hero,""gripping," and "superb." This generalization is again useful in identifying new films of interest. Attribute i is an associative attribute. It records associations between the target objects in this domain, namely movies, and ancillary target objects of an entirely different sort, namely humans. A good indication that the user wants to rent a particular movie is that the user has previously rented other movies with similar attribute values, and this holds for attribute I just as it does for attributes a-h. For example, if the user has often liked movies that customer C_{17} and customer C_{190} have rented, then the user may like other such movies, which have similar values for attribute i. Attribute j is another example of an associative attribute, recording associations between target objects and actors. Notice that any of these attributes can be made subject to authentication when the profile is constructed, through the use of digital signatures; for example, the target object could be accompanied by a digitally signed note from the MPAA, which note names the target object and specifies its authentic value for attribute c.

[00350] These three kinds of attributes are common: numeric, textual, and associative. In the classical information retrieval problem, where the target objects are documents (or more generally, coherent document sections extracted by a text segmentation method), the system might only consider a single, textual attribute when measuring similarity: the full text of the

target object. However, a more sophisticated system would consider a longer target profile, including numeric and associative attributes:

[00351] (a.) full text of document (textual),

[00352] (b.) title (textual),

[00353] (c.) author (textual),

[00354] (d.) language in which document is written (textual),

[00355] (e.) date of creation (numeric),

[00356] (f.) date of last update (numeric),

[00357] (g.) length in words (numeric),

[00358] (h.) reading level (numeric),

[00359] (i.) quality of document as rated by a third_party editorial agency (numeric),

[00360] (j.) list of other readers who have retrieved this document (associative).

[00361] As another domain example, consider a domain where the user is an advertiser and the target objects are potential customers. The system might store the following attributes for each target object (potential customer):

[00362] (a.) first two digits of zip code (textual),

[00363] (b.) first three digits of zip code (textual),

[00364] (c.) entire five-digit zip code (textual),

[00365] (d.) distance of residence from advertiser's nearest physical storefront (numeric),

[00366] (e.) annual family income (numeric),

[00367] (f.) number of children (numeric),

[00368] (g.) list of previous items purchased by this potential customer (associative),

[00369] (h.) list of filenames stored on this potential customer's client computer (associative),

[00370] (i.) list of movies rented by this potential customer (associative),

[00371] (j.) list of investments in this potential customer's investment portfolio (associative),

[00372] (k.) list of documents retrieved by this potential customer (associative),

[00373] (l.) written response to Rorschach inkblot test (textual),

[00374] (m.) multiple-choice responses by this customer to 20 self-image questions (20 textual attributes).

[00375] As always, the notion is that similar consumers buy similar products. It should be noted that diverse sorts of information are being used here to characterize consumers, from their consumption patterns to their literary taste s and psychological peculiarities, and that

this fact illustrates both the flexibility and power of the system for customized electronic identification of desirable objects of the present invention. Diverse sorts of information can be used as attributes in other domains as well (as when physical, economic, psychological and interest-related questions are used to profile the applicants to a dating service, which is indeed a possible domain for the present system), and the advertiser domain is simply an example.

[00376] As a final domain example, consider a domain where the user is an stock market investor and the target objects are publicly traded corporations. A great many attributes might be used to characterize each corporation, including but not limited to the following:

[00377] (a.) type of business (textual),

[00378] (b.) corporate mission statement (textual),

[00379] (c.) number of employees during each of the last 10 years (ten separate numeric attributes),

[00380] (d.) percentage growth in number of employees during each of the last 10 years,

[00381] (e.) dividend payment issued in each of the last 40 quarters, as a percentage of current share price,

[00382] (f.) percentage appreciation of stock value during each of the last 40 quarters, list of shareholders (associative),

[00383] (g.) composite text of recent articles about the corporation in the financial press (textual).

[00384] It is worth noting some additional attributes that are of interest in some domains. In the case of documents and certain other domains, it is useful to know the source of each target object (for example, refereed journal article vs. UPI newswire article vs. Usenet newsgroup posting vs. question-answer pair from a question-and-answer list vs. tabloid newspaper article vs. . . .); the source may be represented as a single-term textual attribute. Important associative attributes for a hypertext document are the list of documents that it links to, and the list of documents that link to it. Documents with similar citations are similar with respect to the former attribute, and documents that are cited in the same places are similar with respect to the latter. A convention may optionally be adopted that any document also links to itself especially in systems where users can choose whether or not to retrieve a target object, a target object's popularity (or circulation) can be usefully measured as a numeric attribute specifying the number of users who have retrieved that object. Related measurable numeric attributes that also indicate a kind of popularity include the number of replies to a target object, in the domain where target objects are messages posted to an

electronic community such as an computer bulletin board or newsgroup, and the number of links leading to a target object, in the domain where target objects are interlinked hypertext documents on the World Wide Web or a similar system. A target object may also receive explicit numeric evaluations (another kind of numeric attribute) from various groups, such as the Motion Picture Association of America (MPAA), as above, which rates movies' appropriateness for children, or the American Medical Association, which might rate the accuracy and novelty of medical research papers, or a random survey sample of users (chosen from all users or a selected set of experts), who could be asked to rate nearly anything. Certain other types of evaluation, which also yield numeric attributes, may be carried out mechanically. For example, the difficulty of reading a text can be assessed by standard procedures that count word and sentence lengths, while the vulgarity of a text could be defined as (say) the number of vulgar words it contains, and the expertise of a text could be crudely assessed by counting the number of similar texts its author had previously retrieved and read using the invention, perhaps confining this count to texts that have high approval ratings from critics. Finally, it is possible to synthesize certain textual attributes mechanically, for example to reconstruct the script of a movie by applying speech recognition techniques to its soundtrack or by applying optical character recognition techniques to its closed-caption subtitles.

[00385] Decomposing Complex Attributes

[00386] Although textual and associative attributes are large and complex pieces of data, for information retrieval purposes they can be decomposed into smaller, simpler numeric attributes. This means that any set of attributes can be replaced by a (usually larger) set of numeric attributes, and hence that any profile can be represented as a vector of numbers denoting the values of these numeric attributes. In particular, a textual attribute, such as the full text of a movie review, can be replaced by a collection of numeric attributes that represent scores to denote the presence and significance of the words "aardvark,""aback,""abacus," and so on through "zymurgy" in that text. The score of a word in a text may be defined in numerous ways. The simplest definition is that the score is the rate of the word in the text, which is computed by computing the number of times the word occurs in the text, an d dividing this number by the total number of words in the text. This sort of score is often called the "term frequency" (TF) of the word. The definition of term frequency may optionally be modified to weight different portions of the text unequally: for example, any occurrence of a word in the text's title might be counted as a 3-fold or more generally k-fold occurrence (as if the title had been repeated k times within the text), in order to reflect a

heuristic assumption that the words in the title are particularly important indicators of the text's content or topic.

[00387] However, for lengthy textual attributes, such as the text of an entire document, the score of a word is typically defined to be not merely its term frequency, but its term frequency multiplied by the negated logarithm of the word's "global frequency," as measured with respect to the textual attribute in question. The global frequency of a word, which effectively measures the word's un-informativeness, is a fraction between 0 and 1, defined to be the fraction of all target objects for which the textual attribute in question contains this word. This adjusted score is often known in the art as TF/IDF ("term frequency times inverse document frequency"). When global frequency of a word is taken into account in this way, the common, uninformative words have scores comparatively close to zero, no matter how often or rarely they appear in the text. Thus, their rate has little influence on the object's target profile. Alternative methods of calculating word scores include latent semantic indexing or probabilistic models.

[00388] Instead of breaking the text into its component words, one could alternatively break the text into overlapping word bigrams (sequences of 2 adjacent words), or more generally, word n-grams. These word n-grams may be scored in the same way as individual words. Another possibility is to use character n-grams. For example, this sentence contains a sequence of overlapping character 5-grams which starts "for e", "or ex", "r exa", "exam", "examp", etc. The sentence may be characterized, imprecisely but usefully, by the score of each possible character 5-gram ("aaaaa", "aaaab", . . . "zzzzz") in the sentence. Conceptually speaking, in the character 5-gram case, the textual attribute would be decomposed into at least 26^5 =11,881,376 numeric attributes. Of course, for a given target object, most of these numeric attributes have values of 0, since most 5-grams do not appear in the target object attributes. These zero values need not be stored anywhere. For purposes of digital storage, the value of a textual attribute could be characterized by storing the set of character 5-grams that actually do appear in the text, together with the nonzero score of each one. Any 5-gram that is no t included in the set can be assumed to have a score of zero. The decomposition of textual attributes is not limited to attributes whose values are expected to be long texts. A simple, one-term textual attribute can be replaced by a collection of numeric attributes in exactly the same way. Consider again the case where the target objects are movies. The "name of director" attribute, which is textual, can be replaced by numeric attributes giving the scores for "Federico-Fellini,""Woody-Allen,""Terence-Davies," and so forth, in that attribute. For these one-term textual attributes, the score of a word is usually defined to be its rate in the

text, without any consideration of global frequency. Note that under these conditions, one of the scores is 1, while the other scores are 0 and need not be stored. For example, if Davies did direct the film, then it is "Terence-Davies" whose score is 1, since "Terence-Davies" constitutes 100% of the words in the textual value of the "name of director" attribute. It might seem that nothing has been gained over simply regarding the textual attribute as having the string value "Terence-Davies." However, the trick of decomposing every non-numeric attribute into a collection of numeric attributes proves useful for the clustering and decision tree methods described later, which require the attribute values of different objects to be averaged and/or ordinally ranked. Only numeric attributes can be averaged or ranked in this way.

[00389] Just as a textual attribute may be decomposed into a number of component terms (letter or word n-grams), an associative attribute may be decomposed into a number of component associations. For instance, in a domain where the target objects are movies, a typical associative attribute used in profiling a movie would be a list of customers who have rented that movie. This list can be replaced by a collection of numeric attributes, which give the "association scores" between the movie and each of the customers known to the system. For example, the 165th such numeric attribute would be the association score between the movie and customer #165, where the association score is defined to be 1 if customer #165 has previously rented the movie, and 0 otherwise. In a subtler refinement, this association s core could be defined to be the degree of interest, possibly zero, that customer #165 exhibited in the movie, as determined by relevance feedback (as described below). As another example, in a domain where target objects are companies, an associative attribute indicating the major shareholders of the company would be decomposed into a collection of association scores, each of which would indicate the percentage of the company (possibly zero) owned by some particular individual or corporate body. Just as with the term scores used in decomposing lengthy textual attributes, each association score may optionally be adjusted by a multiplicative factor: for example, the association score between a movie and customer #165 might be multiplied by the negated logarithm of the "global frequency" of customer #165, i.e., the fraction of all movies that have been rented by customer #165. Just as with the term scores used in decomposing textual attributes, most association scores found when decomposing a particular value of an associative attribute are zero, and a similar economy of storage may be gained in exactly the same manner by storing a list of only those ancillary objects with which the target object has a nonzero association score, together with their respective association scores.

[00390] Similarity Measures

[00391] What does it mean for two target objects to be similar? More precisely, how should one measure the degree of similarity? Many approaches are possible and any reasonable metric that can be computed over the set of target object profiles can be used, where target objects are considered to be similar if the distance between their profiles is small according to this metric. Thus, the following preferred embodiment of a target object similarity measurement system has m any variations.

[00392] First, define the distance between two values of a given attribute according to whether the attribute is a numeric, associative, or textual attribute. If the attribute is numeric, then the distance between two values of the attribute is the absolute value of the difference between the two values. (Other definitions are also possible: for example, the distance between prices p1 and p2 might be defined by $|(p1-p2)|/(\max(p1,p2)+1)$, to recognize that when it comes to customer interest, $50 00 and $5020 are very similar, whereas $3 and $23 are not.) If the attribute is associative, then its value V may be decomposed as described above into a collection of real numbers, representing the association scores between the target object in question and various ancillary objects. V may therefore be regarded as a vector with components V_1, V_2, V_3, etc., representing the association scores between the object and ancillary objects 1, 2, 3, etc., respectively. The distance between two vector values V and U of an associative attribute is then computed using the angle distance measure, $\arccos(Vu^t/.\mathrm{sqrt}((Vv^t)(UU^t))$. (Note that the three inner products in this expression have the form $XY^t=X_1+Y_1+X_2Y_2+X_3Y_3+\ldots$, and that for efficient computation, terms of the form X_iY_i may be omitted from this sum if either of the scores X_i and Y_i is zero.) Finally, if the attribute is textual, then its value V may be decomposed as described above into a collection of real numbers, representing the scores of various word n-grams or character n-grams in the text. Then the value V may again be regarded as a vector, and the distance between two values is again defined via the angle distance measure. Other similarity metrics between two vectors, such as the dice measure, may be used instead. It happens that the obvious alternative metric, Euclidean distance, does not work well: even similar texts tend not to overlap substantially in the content words they use, so that texts encountered in practice are all substantially orthogonal to each other, assuming that TF/IIDF scores are used to reduce the influence of non-content words. The scores of two words in a textual attribute vector may be correlated; for example, “Kennedy” and “JFK” tend to appear in the same documents.

[00393] Thus it may be advisable to alter the text somewhat before computing the scores of terms in the text, by using a synonym dictionary that groups together similar words. The

effect of this optional pre-alteration is that two texts using related words are measured to be as similar as if they had actually used the same words. One technique is to augment the set of words actually found in the article with a set of synonyms or other words, which tend to co-occur with the words in the article, so that "Kennedy" could be added to every article that mentions "JFK." Alternatively, words found in the article may be wholly replaced by synonyms, so that "JFK" might be replaced by "Kennedy" or by "John F. Kennedy" wherever it appears. In either case, the result is that documents about Kennedy and documents about JFK are adjudged similar. The synonym dictionary may be sensitive to the topic of the document as a whole; for example, it may recognize that "crane" is likely to have a different synonym in a document that mentions birds than in a document that mentions construction. A related technique is to replace each word by its morphological stem, so that "staple", "stapler", and "staples" are all replaced by "staple." Common function words ("a", "and", "the" . . .) can influence the calculated similarity of texts without regard to their topics, and so are typically removed from the text before the scores of terms in the text are computed. A more general approach to recognizing synonyms is to use a revised measure of the distance between textual attribute vectors V and U, namely arccos (AV(AU) ^{t} /sqrt (AV(AV) ^{t} AU(AU) ^{t}), where the matrix A is the dimensionality-reducing linear transformation (or an approximation thereto) determined by collecting the vector values of the textual attribute, for all target objects known to the system, and applying singular value decomposition to the resulting collection. The same approach can be applied to the vector values of associative attributes. The above definitions allow us to determine how close together two target objects are with respect to a single attribute, whether numeric, associative, or textual. The distance between two target objects X and Y with respect to their entire multi-attribute profiles P_X and P_Y is then denoted d(X,Y) or d(P_X, P_Y) and defined as:

[00394] (((distance with respect to attribute a)(weight of attribute a)) ^{k} +((distance with respect to attribute b)(weight of attribute b)) ^{k} +((distance with respect to attribute c)(weight of attribute c)) ^{k} + . . .) ^{k}

[00395] where k is a fixed positive real number, typically 2, and the weights are non-negative real numbers indicating the relative importance of the various attributes. For example, if the target objects are consumer goods, and the weight of the "color" attribute is comparatively very small, then price is not a consideration in determining similarity: a user who likes a brown massage cushion is predicted to show equal interest in the same cushion manufactured in blue, and vice-versa. On the other hand, if the weight of the "color" attribute is comparatively very high, then users are predicted to show interest primarily in products

whose colors they have liked in the past: a brown massage cushion and a blue massage cushion are not at all the same kind of target object, however similar in other attributes, and a good experience with one does not by itself inspire much interest in the other. Target objects may be of various sorts, and it is sometimes advantageous to use a single system that is able to compare tar get objects of distinct sorts. For example, in a system where some target objects are novels while other target objects are movies, it is desirable to judge a novel and a movie similar if their profiles show that similar users like them (an associative attribute). However, it is important to note that certain attributes specified in the movie's target profile are undefined in the novel's target profile, and vice versa: a novel has no "cast list" associative attribute and a movie has no "reading level" numeric attribute. In general, a system in which target objects fall into distinct sorts may sometimes have to measure the similarity of two target objects for which somewhat different sets of attributes are defined. This requires an extension to the distance metric d(*,*) defined above. In certain applications, it is sufficient when carrying out such a comparison simply to disregard attributes that are not defined for both target objects: this allows a cluster of novels to be matched with the most similar cluster of movies, for example, by considering only those attributes that novels and movies have in common.

[00396] However, while this method allows comparisons between (say) novels and movies, it does not define a proper metric over the combined space of novel s and movies and therefore does not allow clustering to be applied to the set of all target objects. When necessary for clustering or other purposes, a metric that allows comparison of any two target objects (whether of the same or different sorts) can b e defined as follows. If a is an attribute, then let Max(a) be an upper bound on the distance between two values of attribute a; notice that if attribute a is an associative or textual attribute, this distance is an angle determined by arccos, so that Max(a) may be chosen to be 180 degrees, while if attribute a is a numeric attribute, a sufficiently large number must be selected by the system designers. The distance between two values of attribute a is given as before in the case where both values are defined; the distance between two undefined values is taken to be zero; finally, the distance between a defined value and an undefined value is always taken to be Max(a)/2. This allows us to determine how close together two target objects are with respect to an attribute a, even if attribute a does not have a defined value for both target objects. The distance d(*,*) between two target objects with respect to their entire multi-attribute profiles is then given in terms of these individual attribute distances exactly as before. It is assumed that one attribute in such a system specifies the sort of target object ("movie", "novel", etc.), and that this attribute may

be highly weighted if target objects of different sorts are considered to be very different despite any attributes they may have in common.

[00397] Examples

[00398] Advertiser on the Exchange

[00399] An advertiser has a one day campaign with a $2,400 dollar budget. He turns on advanced metering with granularity set at one hour and a max spending cap of five times the time unit. The system will try to spend $100 per hour throughout the day. However, if in the second hour of the day the advertiser is only able to spend $40 due to market conditions, the remaining $60 will be added to the next hour's spending. If that hour falls short in spending again, the remainder will again be added to the next hour until the advertiser's accumulation reaches $500 in one hour. The spending cap will not allow the advertiser to spend more than 5 times the budget per time unit.

[00400] Another level of metering has been developed which controls the number of impressions shown by adjusting the bid: it adjusts the bid downward if too many impressions are being shown, or upward if too few are being shown. The result is that the advertiser spends a fixed amount of money over a fixed time, but may show many more impressions than with the basic or advanced metering. This highest level of metering should be used only with highly targeted campaigns; however, since lowering the bid can result are dramatically lower results.

[00401] Unique capabilities that benefit outletsare also been developed as part of the Exchange. Preferably, all outlets in the Exchange work through their member-broker, using that member's branded-version of a web-based user interface (or another interface which uses the Exchange's secure API). Through this web-based interface, the outlet is able to define each of his audio applications and every one of his ad spots, watch their performance in real-time, and manage his inventory, minimum pricing, and other factors for the highest revenues possible.

[00402] On the outlet side of the Exchange:

[00403] Submit any number of websites to the Exchange, describing the sites in detail, telling the system about their content and what they will and will not accept in advertising subject matter. Each new website will need to be approved by the member's administrator, however, before it can "go live" in the Exchange.

[00404] Describe each advertising spot, each place when an ad can appear, in terms of its format, acceptable content, minimum CPM payments, and other factors.

[00405] Set ad configuration code from the Exchange interface into his audio application at the spot where he wants the ad to appear. That puts the spot up for auction the next time a viewer lands on that page.

[00406] Manage which advertisers or which individual ads appear on his site--or do not appear.

[00407] Outlets get two important advantages in working through a member of the Exchange.

[00408] First, an auction guarantees that the outlet always gets the highest price any advertiser is willing to pay at that moment for that opportunity. In a traditional fixed-rate system there is often an advertiser who would have paid more for the ad. Or in the case where no advertiser was willing to pay the fixed rate, some advertiser may have at least have been willing to pay something a little less.

[00409] Second, an auction sells off much more of the outlet's inventory, which makes him more money overall. There is no such thing as bad inventory--it is only a matter of fair pricing.

[00410] Another novel feature of the Exchange which benefits the outlet is value-based pricing, as illustrated in FIG. A8. Under any fixed-price advertising model, whether it is on the web or in print, television or radio, there are frequently advertisers who would pay more than published rates if they could know more about the opportunity; that is, if they could be assured of targeting their exact audience. Conversely, even when the published rates are higher than most advertisers are willing to pay to reach the audience, there are likely some advertisers who would pay some price to reach that audience.

[00411] As shown in FIG.A8, in the auction-based exchange, every ad opportunity is available for auction to all advertisers in the Exchange. Consequently, with advanced targeting and value pricing selected by the advertisers, and with the inherent knowledge base of the characteristics of the ad opportunity, the exchange will sell every ad for the highest price any advertiser is willing to pay. That is value pricing for the highest price.

[00412] Moreover aoutlet sells more of his inventory than he would under any fixed rate scheme. With fixed-rates, the outlet can usually sell off a lot of his premium inventory. But how does he price the remainder? Usually he does not want the buyers of his premium inventory to see low pricing on his remainder, so he offloads it to another seller, or it goes unsold completely. An advertiser who was willing to pay $1.00 CPM to show an ad to a viewer the first three times may not want to show that same ad to that same viewer another three times for $1.00--but he might for $0.50, or $0.25. Or a large ad campaign running on

the outlet's site simply ran out of budget--how does the outlet sell off the next few thousand impressions before another campaign gets underway? The auction-based model gets the outlet the highest price any advertiser is willing to pay--whatever that price may be--on all of his inventory. This effect is clearly illustrated by comparing the fixed price model on the left with the value based model on the right.

[00413] Another novel feature is that outlet working through an Exchange member have almost no risk. Here is how:

[00414] The outlet sets a minimum price he is willing to accept for a specific spot. For example, if the advertiser knows he can get $0.25 CPM from some non-member network, he specifies that as the minimum.

[00415] Someone on behalf of the outlet then enters pass through ad code from another network, broker or other source as the backup ad source if the minimum price is not met in an auction. When a viewer lands on that page, the Exchange runs the auction for each impression on each spot. If an advertiser meets or beats the outlet's minimum, the Exchange shows that advertiser's ad. If no advertiser offers a bid that results in aoutlet payment equal to or above the outlet's minimum, The Exchange shows the pass through ad indicated by the outlet.

[00416] In the first instance the outlet made more than he would have from the other source; in the second instance he lost nothing for trying.

[00417] The forgoing discloses how to create a novel Internet Advertising exchange system, and various novel features that are of particular utility. The implementation of this system may be accomplished in a variety of ways that those skilled in the art will appreciate.

[00418] Therefore the implementation details are not considered part of the novelty or unique to the operation of the system as disclosed.

[00419] Utilizing the Similarity Measurement for Matching Buyers and Sellers

[00420] A simple application of the similarity measurement is a system to match buyers with sellers in small-volume markets, such as used cars and other used goods, artwork, or employment. Sellers submit profiles of the goods (target objects) they want to sell, and buyers submit profiles of the goods (target objects) they want to buy. Participants may submit or withdraw these profiles at any time. The system for customized electronic identification of desirable objects computes the similarities between seller-submitted profiles and buyer-submitted profiles, and when two profiles match closely (i.e., the similarity is above a threshold), the corresponding seller and buyer are notified of each other's identities. To

prevent users from being flooded with responses, it may be desirable to limit the number of notifications each user receives to a fixed number, such as ten per day.

[00421] Utilizing the Similarity Measurement for Filtering on Relevance

[00422] A filtering system is a device that can search through many target objects and estimate a given user's interest in each target object, so as to identify those that are of greatest interest to the user. The filtering system uses relevance feed back to refine its knowledge of the user's interests: whenever the filtering system identifies a target object as potentially interesting to a user, the user (if an on-line user) provides feedback as to whether or not that target object really is of interest. Such feedback is stored long-term in summarized form, as part of a database of user feedback information, and may be provided either actively or passively. In active feedback, the user explicitly indicates his or her interest, for instance, on a scale of −2 (active distaste) through 0 (no special interest) to 10 (great interest). In passive feedback, the system infers the user's interest from the user's behavior. For example, if target objects are textual documents, the system might monitor which documents the user chooses to read, or not to read, and how much time the user spends reading them. A typical formula for assessing interest in a document via passive feedback, in this domain, on a scale of 0 to 10, might be:

[00423] +2 if the second page is viewed,

[00424] +2 if all pages are viewed,

[00425] +2 if more than 30 seconds was spent viewing the document,

[00426] +2 if more than one minute was spent viewing the document,

[00427] +2 if the minutes spent viewing the document are greater than half the number of pages.

[00428] If the target objects are electronic mail messages, interest points might also be added in the case of a particularly lengthy or particularly prompt reply. If the target objects are purchasable goods, interest points might be added for target objects that the user actually purchases, with further points in the case of a large-quantity or high-price purchase. In any domain, further points might be added for target objects that the user accesses early in a session, on the grounds that users access the object s that most interest them first. Other potential sources of passive feedback include an electronic measurement of the extent to which the user's pupils dilate while the user views the target object or a description of the target object. It is possible to combine active and passive feedback. One option is to take a weighted average of the two ratings. Another option is to use passive feedback by default, but to allow the user to examine and actively modify the passive feedback score. In the scenario

above, for instance, an uninteresting article may sometimes remain on the display device for a long period while the user is engaged in unrelated business; the passive feedback score is then inappropriately high, and the user may wish to correct it before continuing. In the preferred embodiment of the invention, a visual indicator, such as a sliding bar or indicator needle on the user's screen, can be used to continuously display the passive feedback score estimated by the system for the target object being viewed, unless the user has manually adjusted the indicator by a mouse operation or other means in order to reflect a different score for this target object, after which the indicator displays the active feedback score selected by the user, and this active feedback score is used by the system instead of the passive feedback score. In a variation, the user cannot see or adjust the indicator until just after the user has finished viewing the target object. Regardless how a user's feedback is computed, it is stored long-term as part of that user's target profile interest summary.

[00429] Utilizing the Similarity Measurement for Topical Interest Determination

[00430] Relevance feedback only determines the user's interest in certain target objects: namely, the target objects that the user has actually had the opportunity to evaluate (whether actively or passively). For target objects that the user has not yet seen, the filtering system must estimate the user's interest. This estimation task is the heart of the filtering problem, and the reason that the similarity measurement is important. More concretely, the preferred embodiment of the filtering system is a news clipping service that periodically presents the user with news articles of potential interest. The user provides active and/or passive feedback to the system relating to these presented articles. However, the system does not have feedback information from the user for articles that have never been presented to the user, such as new articles that have just been added to the database, or old articles that the system chose not to present to the user. Similarly, in the dating service domain where target objects are prospective romantic partners, the system has only received feedback on old flames, not on prospective new loves.

[00431] As shown in Fig. C12, the evaluation of the likelihood of interest in a particular target object for a specific user can automatically be computed. The interest that a given target object X holds for a user U is assumed to be a sum of two quantities: q(J, X), the intrinsic "quality" of X, plus f(U, X), the "topical interest" that users like U have in target objects like X. For any target object X, the intrinsic quality measure q(U, X) is easily estimated at steps 1201 - 1203 directly from numeric attributes of the target object X. The computation process begins at step 1201, where certain designated numeric attributes of target object X are specifically selected, which attributes by their very nature should be

positively or negatively correlated with users' interest. Such attributes, termed "quality attributes," have the normative property that the higher (or in some cases lower) their value, the more interesting a user is expected to find them. Quality attributes of target object X may include, but are not limited to, target object X's popularity among users in general, the rating a particular reviewer has given target object X, the age (time since authorship—also known as out-datedness) of target object X, the number of vulgar words used in target object X, the price of target object X, and the amount of money that the company selling target object X has donated to the user's favorite charity. At step 1202 , each of the selected attributes is multiplied by a positive or negative weight indicative of the strength of user U's preference for those target objects that have high values for this attribute, which weight must be retrieved from a data file storing quality attribute weights for the selected user. At step 1203 , a weighted sum of the identified weighted selected attributes is computed to determine the intrinsic quality measure q(U, X). At step 1204 , the summarized weighted relevance feedback data is retrieved, wherein some relevance feedback points are weighted more heavily than others and the stored relevance data can be summarized to some degree, for example by the use of search profile sets. The more difficult part of determining user U's interest in target object X is to find or compute at step 1205 the value of f(U, X), which denotes the topical interest that users like U generally have in target objects like X. The method of determining a user's interest relies on the following heuristic: when X and Y are similar target objects (have similar attributes), and U and V are similar users (have similar attributes), then topical interest f(U, X) is predicted to have a similar value to the value of topical interest f(V, Y). This heuristic leads to an effective method because estimated values of the topical interest function f(*, *) are actually know n for certain arguments to that function: specifically, if user V has provided a relevance-feedback rating of r(V, Y) for target object Y, then insofar as that rating represents user V's true interest in target object Y, we have r(V, Y)=q(V, Y)+f(V, Y) and can estimate f(V, Y) as r(V, Y)−q(V, Y). Thus, the problem of estimating topical interest at all points becomes a problem of interpolating among these estimates of topical interest at selected points, such as the feedback estimate of f(V, Y) a s r(V, Y)−q(V, Y). This interpolation can be accomplished with any standard smoothing technique, using as input the known point estimates of the value of the topical interest function f(*, *), and determining as output a function that approximates the entire topical interest function f(*, *)

[00432] Not all point estimates of the topical interest function f(*, *) should be given equal weight as inputs to the smoothing algorithm. Since passive relevance feedback is less reliable

than active relevance feedback, point estimates made from passive relevance feedback should be weighted less heavily than point estimates made from active relevance feedback, or even not used at all. In most domains, a user's interests may change over time and, therefore, estimates of topic al interest that derive from more recent feedback should also be weighted more heavily. A user's interests may vary according to mood, so estimates of topical interest that derive from the current session should be weighted more heavily for the duration of the current session, and past estimates of topical interest made at approximately the current time of day or on the current weekday should be weighted more heavily. Finally, in domains where users are trying to locate target objects of long-term interest (investments, romantic partners, pen pals, employers, employees, suppliers, service providers) from the possibly meager information provided by the target profiles, the users are usually not in a position to provide reliable immediate feedback on a target object, but can provide reliable feedback at a later date. An estimate of topical interest f(V, Y) should be weighted more heavily if user V has had more experience with target object Y. Indeed, a useful strategy is for the system to track long-term feedback for such target objects. For example, if target profile Y was created in 1990 to describe a particular investment that was available in 1990, and that was purchased in 1990 by user V, then the system solicits relevance feedback from user V in the years 1990, 1991, 1992, 1993, 1994, 1995, etc., and treats these as successively stronger indications of user V's true interest in target profile Y, and thus as indications of user V's likely interest in new investments whose current profiles resemble the original 1990 investment profile Y. In particular, if in 1994 and 1995 user V is well-disposed toward his or her 1990 purchase of the investment described by target profile Y, then in those years and later, the system tends to recommend additional investments when they have profiles like target profile Y, on the grounds that they too will turn out to be satisfactory in 4 to 5 years. It makes these recommendations both to user V and to users whose investment portfolios and other attributes are similar to user V's. The relevance feedback provided by user V in this case may be either active (feedback=satisfaction ratings provided by the investor V) or passive (feedback=difference between average annual return of the investment and average annual return of the Dow Jones index portfolio since purchase of the investment, for example).

[00433] To effectively apply the smoothing technique, it is necessary to have a definition of the similarity distance between (U, X) and (V, Y), for any users U and V and any target objects X and Y. We have already seen how to define the distance d(X, Y) between two target objects X and Y, given their attributes. We may regard a pair such as (U, X) as an extended object that bears all the attributes of target X and all the attributes of user U; then

the distance between (U, X) and (V, Y) may be computed in exactly the same way. This approach requires user U, user V, and all other users to have some attributes of their own stored in the system: for example, age (numeric), social security number (textual), and list of documents previously retrieved (associative). It is these attributes that determine the notion of "similar users." Thus it is desirable to generate profiles of users (termed "user profiles") as well as profiles of target objects (termed "target profiles"). Some attributes employed for profiling users may be related to the attributes employed for profiling target objects: for example, using associative attributes, it is possible to characterize target objects such as X by the interest that various users have shown in them, and simultaneously to characterize users such as U by the interest that they have shown in various target objects. In addition, user profiles may make use of any attributes that are useful in characterizing humans, such as those suggested in the example domain above where target objects are potential consumers. Notice that user U's interest can be estimated even if user U is a new user or an off-line user who has never provided any feedback, because the relevance feedback of users whose attributes are similar to U's attributes is taken into account.

[00434] For some uses of filtering systems, when estimating topical interest, it is appropriate to make an additional "presumption of no topical interest" (or "bias toward zero"). To understand the usefulness of such a presumption, suppose the system needs to determine whether target object X is topically interesting to the user U, but that users like user U have never provided feedback on target objects even remotely like target object X.

[00435] The presumption of no topical interest says that if this is so, it is because users like user U are simply not interested in such target objects and therefore do not seek them out and interact with them. On this presumption, the system should estimate topical interest f(U, X) to be low. Formally, this example has the characteristic that (U, X) is far away from all the points (V, Y) where feedback is available. In such a case, topical interest f(U, X) is presumed to be close to zero, even if the value of the topical interest function f(*, *) is high at all the faraway surrounding points at which its value is known. When a smoothing technique is used, such a presumption of no topical interest can be introduced, if appropriate, by manipulating t he input to the smoothing technique. In addition to using observed values of the topical interest function f(*, *) as input, the trick is to also introduce fake observations of the form topical interest f(V, Y)=0 for a lattice of points (V, Y) distributed throughout the multidimensional space. These fake observations should be given relatively low weight as inputs to the smoothing algorithm. The more strongly they are weighted, the stronger the presumption of no interest.

[00436] The following provides another simple example of an estimation technique that has a presumption of no interest. Let g be a decreasing function from non-negative real numbers to non-negative real numbers, such as $g(x)=e^{x}$ or $g(x)=\min(1, x^{-k})$ where $k>1$. Estimate topical interest f(U, X) with the following g-weighted average:

[00437] $f(U,X) = \Sigma((r(V,Y) - q(V,Y) * g(\text{distance } \varphi(U,X) \wedge (V,Y))/ \Sigma g(\text{distance } \varphi (U,V) \wedge (V,Y))$

[00438] Here the summations are over all pairs (V, Y) such that user V has provided feedback r(V, Y) on target object Y, i.e., all pairs (V, Y) such that relevance feedback r(V, Y) is defined. Note that both with this technique and with conventional smoothing techniques, the estimate of the topical interest f(U, X) is not necessarily equal to r(U, X)−q(U, X), even when r(U, X) is defined.

[00439] Utilizing the Similarity Measurement: Adjusting Weights and Residue Feedback

[00440] The method described above requires the filtering system to measure distances between (user, target object) pairs, such as the distance between (U, X) and (V, Y). Given the means described earlier for measuring the distance between two multi-attribute profiles, the method must therefore associate a weight with each attribute used in the profile of (user, target object) pairs, that is, with each attribute used to profile either users or target objects. These weights specify the relative importance of the attributes in establishing similarity or difference, and therefore, in determining how topical interest is generalized from one (user, target object) pair to another. Additional weights determine which attributes of a target object contribute to the quality function q, and by how much.

[00441] It is possible and often desirable for a filtering system to store a different set of weights for each user. For example, a user who thinks of two-star films as having materially different topic and style from four-star films wants to assign a high weight to “number of stars” for purposes of the similarity distance measure d(*, *); this means that interest in a two-star film does not necessarily signal interest in an otherwise similar four-star film, or vice-versa. If the user also agrees with the critics, and actually prefers four-star films, the user also wants to assign “number of stars” a high positive weight in the determination of the quality function q. In the same way, a user who dislikes vulgarity wants to assign the “vulgarity score” attribute a high negative weight in the determination of the quality function q, although the “vulgarity score” attribute does not necessarily have a high weight in determining the topical similarity of two films.

[00442] Attribute weights (of both sorts) may be set or adjusted by the system administrator or the individual user, on either a temporary or a permanent basis. However, it is often

desirable for the filtering system to learn attribute weights automatically, based on relevance feedback. The optimal attribute weights for a user U are those that allow the most accurate prediction of user U's interests. That is, with the distance measure and quality function defined by these attribute weights, user U's interest in target object X, q(U, X)+f(U, X), can be accurately estimated by the techniques above. The effectiveness of a particular set of attribute weights for user U can therefore be gauged by seeing how well it predicts user U's known interests.

[00443] Formally, suppose that user U has previously provided feedback on target objects X_1, X_2, X_3, . . . X_n, and that the feedback ratings are r(U, X_1), r(U, X_2), r(U, X_3), . . . r(U, X_n). Values of feedback ratings r(*,*) for other users and other target objects may also be known. The system may use the following procedure to gauge the effectiveness of the set of attribute weights it currently stores for user U: (i) For each 1<=I<=n, use the estimation techniques to estimate q(U, X_1)+f(U, X_i) from all known values of feedback ratings r. Call this estimate a_i. (ii) Repeat step (i), but this time make the estimate for each 1<=i<=n without using the feedback ratings r(U, X_j) as input, for any j such that the distance d(X_i, X_j) is smaller than a fixed threshold. That is, estimate each q(U, X_i)+f(U, X_i) from other values of feedback rating r only; in particular, do not use r(U, X_i) itself. Call this estimate b_i. The difference $a_i - b_i$ is herein termed the "residue feedback r_{res}(U, X_i) of user U on target object X_i." (iii) Compute user U's error measure, $(a_1 - b_1)^2 + (a_2 - b_2)^2 + (a_3 - b_3)^2 + \ldots + (a_n - b_n)^2$.

[00444] A gradient-descent or other numerical optimization method may be used to adjust user U's attribute weights so that this error measure reaches a (local) minimum. This approach tends to work best if the smoothing technique used in estimation is such that the value of f(V, Y) is strongly affected by the point estimate r(V, Y)−q(V, Y) when the latter value is provided as input. Otherwise, the presence or absence of the single input feedback rating r(U, X_i), in steps (i)-(ii) may not make a_i and b_i very different from each other. A slight variation of this learning technique adjusts a single global set of at tribute weights for all users, by adjusting the weights so as to minimize not a particular user's error measure but rather the total error measure of all users. These global weights are used as a default initial setting for a new user who has not yet provided any feedback. Gradient descent can then be employed to adjust this user's individual weights over time.

[00445] Even when the attribute weights are chosen to minimize the error measure for user U, the error measure is generally still positive, meaning that residue feedback from user U has not been reduced to 0 on all target objects. It is useful to note that high residue feedback

from a user U on a target object X indicates that user U liked target object X unexpectedly well given its profile, that is, better than the smoothing model could predict from user U's opinions on target objects with similar profiles. Similarly, low residue feedback indicates that user U liked target object X less than was expected. By definition, this unexplained preference or dis-preference cannot be the result of topical similarity, and therefore must be regarded as an indication of the intrinsic quality of target object X. It follows that a useful quality attribute for a target object X is the average amount of residue feedback $r_{res}(V, X)$ from users on that target object, averaged over all users V who have provided relevance feedback on the target object. In a variation of this idea, residue feedback is never averaged indiscriminately over all users to form a new attribute, but instead is smoothed to consider users' similarity to each other. Recall that the quality measure q(U, X) depends on the user U as well as the target object X, so that a given target object X may be perceived by different users to have different quality. In this variation, as before, q(U, X) is calculated as a weighted sum of various quality attributes that are dependent only on X, but then an additional term is added, namely an estimate of $r_{res}(U, X)$ found by applying a smoothing algorithm to known values of $r_{res}(V, X)$. Here V ranges over all users who have provided relevance feedback on target object X, and the smoothing algorithm is sensitive to the distances d(U, V) from each such user V to user U.

[00446] Utilizing the Similarity Measurement: Clustering

[00447] A method for defining the distance between any pair of target object's was disclosed above. Given this distance measure, it is simple to apply a standard clustering algorithm, such as k-means, to group the target objects into a number of clusters, in such a way that similar target objects tend to be grouped in the same cluster. It is clear that the resulting clusters can be used to improve the efficiency of matching buyers and sellers in the application described in section "Matching Buyers and Sellers" above: it is not necessary to compare every buy profile to every sell profile, but only to compare buy profiles and sell profiles that are similar enough to appear in the same cluster. As explained below, the results of the clustering procedure can also be used to make filtering more efficient, and in the service of querying and browsing tasks.

[00448] The k-means clustering method is familiar to those skilled in the art. Briefly put, it finds a grouping of points (target profiles, in this case, whose numeric coordinates are given by numeric decomposition of their attributes as described above) to minimize the distance between points in the clusters and the centers of the clusters in which they are located. This is done by alternating between assigning each point to the cluster which has the nearest center

and then, once the points have been assigned, computing the (new) center of each cluster by averaging the coordinates of the points (target profiles) located in this cluster. Other clustering methods can be used, such as "soft" or "fuzzy" k-means clustering, in which objects are allowed to belong to more than one cluster. This can be cast as a clustering problem similar to the k-means problem, but now the criterion being optimized is a little different:

$$\sum_i \sum_c i_{iC} d(x_i, \bar{x}_c)$$

[00449] where C ranges over cluster numbers, i ranges over target objects, x_i is the numeric vector corresponding to the profile of target object number i, $\bar{c}$ is the mean of all the numeric vectors corresponding to target profiles of target objects in cluster number C, termed the "cluster profile" of cluster C, d(*, *) is the metric used to measure distance between two target profiles, and i_{iC} is a value between 0 and 1 that indicates how much target object number i is associated with cluster number C, where i is an indicator matrix with the property that for each i, SUM SUB C I SUB iC=1. For k-means clustering, $i_{iC}$is either 0 or 1.

[00450] Any of these basic types of clustering might be used by the system:

[00451] 1) Association-based clustering, in which profiles contain only associative attributes, and thus distance is defined entirely by associations. This kind of clustering generally (a) clusters target objects based on the similarity of the users who like them or (b) clusters users based on the similarity of the target objects they like. In this approach, the system does not need any information about target objects or users, except for their history of interaction with each other.

[00452] 2) Content-based clustering, in which profiles contain only non-associative attributes. This kind of clustering (a) clusters target objects based on the similarity of their non-associative attributes (such as word frequencies) or (b) clusters users base d on the similarity of their non-associative attributes (such as demographics and psychographics). In this approach, the system does not need to record any information about users' historical patterns of information access, but it does need information about the intrinsic properties of users and/or target objects. !

[00453] 3) Uniform hybrid method, in which profiles may contain both associative and non-associative attributes. This method combines 1a and 2a, or 1b and 2b. The distance d(P_X

, P_Y) between two profiles P_X and P_Y may be computed by the general similarity-measurement methods described earlier.

[00454] 4) Sequential hybrid method. First apply the k-means procedure to do 1a, so that articles are labeled by cluster based on which user read them, then use supervised clustering (maximum likelihood discriminant methods) using the word frequencies to do the process of method 2a described above. This tries to use knowledge of who read what to do a better job of clustering based on word frequencies. One could similarly combine the methods 1b and 2b described above.

[00455] Hierarchical clustering of target objects is often useful. Hierarchical clustering produces a tree which divides the target objects first into two large clusters of roughly similar objects; each of these clusters is in turn divided into two or more smaller clusters, which in turn are each divided into yet smaller clusters until the collection of target objects has been entirely divided into "clusters" consisting of a single object each. The node d denotes a particular target object d, or equivalently, a single-member cluster consisting of this target object. Target object d is a member of the cluster (a, b, d), which is a subset of the cluster (a, b, c, d, e, f), which in turn is a subset of all target objects. A tree would be produced from a set of target objects. In a cluster , each letter represents a target object, and axes x 1 and x 2 represent two of the many numeric attributes on which the target objects differ. Such a cluster tree may be created by hand, using human judgment to form clusters and sub-clusters of similar objects, or may be created automatically in either of two standard ways: top-down or bottom-up. In top-down hierarchical clustering, the set of all target objects would be divided into the clusters (a, b, c, d, e, f) and (g, h, i, j, k). The clustering algorithm would then be reapplied to the target objects in each cluster, so that the cluster (g, h, i, j, k) is sub-partitioned into the clusters (g, k) and (h, i, j), and so on to arrive at a tree. In bottom-up hierarchical clustering, the set of all target objects in would be grouped into numerous small clusters, namely (a, b), d, (c, f), e, (g,k), (h, i), and j. These clusters would then themselves be grouped into the larger clusters (a, b, d), (c, e, f), (g, k), and (h, i, j), according to their cluster profiles. These larger clusters would themselves be grouped into (a, b, c, d, e, f) and (g, k, h, i, j), and so on until all target objects had been grouped together, resulting in a tree. Note that for bottom-up clustering to work, it must be possible to apply the clustering algorithm to a set of existing clusters. This requires a notion of the distance between two clusters. The method disclosed above for measuring the distance between target objects can be applied directly, provided that clusters are profiled in the same way as target objects. It is only necessary to adopt the convention that a cluster's profile is the average of the target profiles of all the

target objects in the cluster; that is, to determine the cluster's value for a given attribute, take the mean value of that attribute across all the target objects in the cluster. For the mean value to be well-defined, all attributes must be numeric, so it is necessary as usual to replace each textual or associative attribute with its decomposition into numeric attributes (scores), as described earlier. For example, the target profile of a single Woody Allen film would assign "Woody-Allen" a score of 1 in the "name-of-director" field, while giving "Federico-Fellini" and "Terence-Davies" scores of 0. A cluster that consisted of 20 films directed by Allen and 5 directed by Fellini would be profiled with scores of 0.8, 0.2, and 0 respectively, because, for example, 0.8 is the average of 20 ones and 5 zeros.

[00456] Utilizing the Similarity Measurement: Searching for Target Objects

[00457] Given a target object with target profile P, or alternatively given a search profile P, a hierarchical cluster tree of target objects makes it possible for the system to search efficiently for target objects with target profiles similar to P. It is only necessarily to navigate through the tree, automatically, in search of such target profiles. The system for customized electronic identification of desirable objects begins by considering the largest, top-level clusters, and selects the cluster whose profile is most similar to target profile P. In the event of a near-tie, multiple clusters may be selected. Next, the system considers all sub-clusters of the selected clusters, and this time selects the sub-cluster or sub-clusters whose profiles are closest to target profile P. This refinement process is iterated until the clusters selected on a given step are sufficiently small, and these are the desired clusters of target objects with profiles most similar to target profile P. Any hierarchical cluster tree therefore serves as a decision tree for identifying target objects. In pseudo-code form, this process is as follows:

[00458] 1. Initialize list of identified target objects to the empty list at step X A 00

[00459] 2. Initialize the current tree T to be the hierarchical cluster tree of all objects at step X A 01 and at step X A 02 scan the current cluster tree for target objects similar to P, using the process detailed in X B . At step X A 03 , the list of target objects is returned.

[00460] 3. At step X B 00 , the variable I is set to 1 and for each child sub-tree Ti of the root of tree T, is retrieved.

[00461] 4. At step X B 02 , calculate d(P, pi), the similarity distance between P and pi,

[00462] 5. At step X B 03 , if d(P, pi)<t, a threshold, branch to one of two options

[00463] 6. If tree Ti contains only one target object at step X B 04 , add that target object to list of identified target objects at step X B 05 and advance to step X B 07 .

[00464] 7. If tree Ti contains multiple target objects at step X B 04 , scan the ith child sub-tree for target objects similar to P by invoking the steps of the process of X B recursively and

then recurse to step 3 (step X A 01) with T bound for the duration of the recursion to tree Ti, in order to search in tree Ti for target objects with profiles similar to P.

[00465] In step 5 of this pseudo-code, smaller thresholds are typically used at lower levels of the tree, for example by making the threshold an affine function or other function of the cluster variance or cluster diameter of the cluster pi. If the cluster tree is distributed across a plurality of servers, as described in the section of this description titled "Network Context of the Browsing System", this process may be executed in distributed fashion as follows: steps 3-7 are executed by the server that stores the root node of hierarchical cluster tree T, and the recursion in step 7 to a sub-cluster tree T_i involves the transmission of a search request to the server that stores the root node of tree T_i, which server carries out the recursive step upon receipt of this request. Steps 1-2 are carried out by the processor that initiates the search, and the server that executes step 6 must send a message identifying the target object to this initiating processor, which adds it to the list.

[00466] Assuming that low-level clusters have been already been formed through clustering, there are alternative search methods for identifying the low-level cluster whose profile is most similar to a given target profile P. A standard back-propagation neural net is one such method: it should be trained to take the attributes of a target object as input, and produce as output a unique pattern that can be used to identify the appropriate low-level cluster. For maximum accuracy, low-level clusters that are similar to each other (close together in the cluster tree) should be given similar identifying patterns. Another approach is a standard decision tree that considers the attributes of target profile P one at a time until it can identify the appropriate cluster. If profiles are large, this may be more rapid than considering all attributes. A hybrid approach to searching uses distance measurements as described above to navigate through the top few levels of the hierarchical cluster tree, until it reaches an cluster of intermediate size whose profile is similar to target profile P, and then continues by using a decision tree specialized to search for low-level sub-clusters of that intermediate cluster.

[00467] One use of these searching techniques is to search for target objects that match a search profile from a user's search profile set. This form of searching is used repeatedly in the news clipping service, active navigation, and Virtual Community Service applications, described below. Another use is to add a new target object quickly to the cluster tree. An existing cluster that is similar to the new target object can be located rapidly, and the new target object can be added to this cluster. If the object is beyond a certain threshold distance from the cluster center, then it is advisable to start a new cluster. Several variants of this

incremental clustering scheme can be used, and can be built using variants of subroutines available in advanced statistical packages. Note that various methods can be used to locate t he new target objects that must be added to the cluster tree, depending on the architecture used. In one method, a "web-crawler" program running on a central computer periodically scans all servers in search of new target objects, calculates the target profiles of these objects, and adds them to the hierarchical cluster tree by the above method. In another, whenever a new target object is added to any of the servers, a software "agent" at that server calculates the target profile and adds it to the hierarchical cluster tree by the above method.

[00468] Utilizing the Similarity Measurement: Rapid Profiling

[00469] In some domains, complete profiles of target objects are not always easy to construct automatically. When target objects are multi-media games e.g., an attribute such as genre (a single textual term such as "action", "suspense/thriller", "word games", etc.) may be a matter of judgment and opinion. More significantly, if each title has an associated attribute that records the positive or negative relevance feedback to that title from various human users (consumers), then all the association scores of any newly introduced titles are initially zero, so that it is initially unclear what other titles are similar to the new title with respect to the users who like them. Indeed, if this associative attribute is highly weighted, the initial lack of relevance feedback information may be difficult to remedy, due to a vicious circle in which users of moderate-to-high interest are needed to provide relevance feedback but relevance feedback is needed to identify users of moderate-to-high interest.

[00470] Fortunately, however, it is often possible in principle to determine certain attributes of a new target object by extraordinary methods, including but not limited to methods that consult a human. For example, the system can in principle determine the genre of a title by consulting one or more randomly chosen individuals from a set of known human experts, while to determine the numeric association score between a new title and a particular user, it can in principle show the title to the that user and obtain relevance feedback. Since such requests inconvenience people, however, it is important not to determine all difficult attributes this way, but only the ones that are most important in classifying the article. "Rapid profiling" is a method for selecting those numeric attributes that are most important to determine. (Recall that all attributes can be decomposed into numeric attributes, such as association scores or term scores.) First, a set of existing target objects that already have complete or largely complete profiles are clustered using a k-means algorithm. Next, each of the resulting clusters is assigned a unique identifying number, and each clustered target object is labeled with the identifying number of its cluster. Standard methods then allow

construction of a single decision tree that can determine any target object's cluster number, with substantial accuracy, by considering the attributes of the target object, one at a time. Only attributes that can if necessary be determined for any new target object are used in the construction of this decision tree. To profile a new target object, the decision tree is traversed downward from its root as far as is desired. The root of the decision tree considers some attribute of the target object. If the value of this attribute is not yet known, it is determined by a method appropriate to that attribute; for example, if the attribute is the association score of the target object with user #4589, then relevance feedback (to be used as the value of this attribute) is solicited from user #4589, perhaps by the ruse of adding the possibly uninteresting target object to a set of objects that the system recommends to the user's attention, in order to find out what the user thinks of it. Once the root attribute is determined, the rapid profiling method descends the decision tree by one level, choosing one of the decision sub-trees of the root in accordance with the determined value of the root attribute. The root of this chosen sub-tree considers another attribute of the target object, whose value is likewise determined by an appropriate method. The process c an be repeated to determine as many attributes as desired, by whatever methods are available, although it is ordinarily stopped after a small number of attributes, to avoid the burden of determining too many attributes.

[00471] It should be noted that the rapid profiling method can be used to identify important attributes in any sort of profile, and not just profiles of target objects. In particular, recall that the disclosed method for determining topical interest through similarity requires users as well as target objects to have profiles. New users, like new target objects, may be profiled or partially profiled through the rapid profiling process. For example, when user profiles include an associative attribute that records the user's relevance feedback on all target objects in the system, the rapid profiling procedure can rapidly form a rough characterization of a new user's interests by soliciting the user's feedback on a small number of significant target objects, and perhaps also by determining a small n umber of other key attributes of the new user, by on-line queries, telephone surveys, or other means. Once the new user has been partially profiled in this way, the methods disclosed above predict that the new user's interests resemble the known interests of other users with similar profiles. In a variation, each user's user profile is subdivided into a set of long-term attributes, such as demographic characteristics, and a set of short-term attributes that help to identify the user's temporary desires and emotional state, such as the user's textual or multiple-choice answers to questions whose answers reflect the user's mood. A subset of the user's long-term attributes are

determined when the user first registers with the system, through the use of a rapid profiling tree of long-term attributes. In addition, each time the user logs on to the system, a subset of the user's short-term attributes are additionally determined, through the use of a separate rapid profiling tree that asks about short-term attributes.

[00472] Utilizing the Similarity Measurement: Market Research

[00473] A technique similar to rapid profiling is of interest in market research (or voter research). Suppose that the target objects are consumers. A particular attribute in each target profile indicates whether the consumer described by that target profile h as purchased product X. A decision tree can be built that attempts to determine what value a consumer has for this attribute, by consideration of the other attributes in the consumer's profile. This decision tree may be traversed to determine whether additional users are likely to purchase product X. More generally, the top few levels of the decision tree provide information, valuable to advertisers who are planning mass-market or direct mail campaigns, about the most significant characteristics of consumers of product X.

[00474] Similar information can alternatively be extracted from a collection of consumer profiles without recourse to a decision tree, by considering attributes one at a time, and identifying those attributes on which pro duct X's consumers differ significantly from its non-consumers. These techniques serve to characterize consumers of a particular product; they can be equally well applied to voter research or other survey research, where the objective is to characterize those individuals from a given set of surveyed individuals who favor a particular candidate, hold a particular opinion, belong to a particular demographic group, or have some other set of distinguishing attributes. Researchers may wish to purchase batches of analyzed or unanalyzed user profiles from which personal identifying information has been removed. As with any statistical database, statistical conclusions can be drawn, and relationships between attributes can be elucidated using knowledge discovery techniques, which are well known in the art.

[00475] In one or more embodiments, the term "memory" can mean a "memory medium" and/or "computer readable memory medium" which is intended to include various types of memory or storage, including an installation medium, e.g., a CD-ROM, or floppy disks, a random access memory or computer system memory such as DRAM, SRAM, EDO RAM, Rambus RAM, NVRAM, EPROM, EEPROM, flash memory etc., and/or a non-volatile memory such as a magnetic media, e.g., a hard drive, and/or optical storage. The memory medium can include other types of memory as well, or combinations thereof. In one or more

embodiments, the memory medium can be and/or include an article of manufacture and/or a software product. For example, an article of manufacture and/or a software product can include a memory medium that includes instructions executable by a processor to perform one or more portions of one or more methods and/or processes described herein.

[00476] In addition, the memory medium can be located in a first computer in which the programs are executed, or can be located in a second different computer and/or hardware memory device that connects to the first computer over a network. In one or more embodiments, the second computer provides the program instructions to the first computer for execution. The memory medium can also be a distributed memory medium, e.g., for security reasons, where a portion of the data is stored on one memory medium and the remaining portion of the data can be stored on a different memory medium. Also, the memory medium can include one of the networks to which the current network is coupled, e.g., a SAN (Storage Area Network).

[00477] In one or more embodiments, each of the systems described herein may take various forms, including a personal computer system, server computer system, workstation, network appliance, Internet appliance, wearable computing device, personal digital assistant (PDA), tablet computing device, laptop, mobile telephone, mobile multimedia device, embedded computer system, television system, and/or other device. In general, the terms “computing device”, “computer”, and/or "computer system" can be broadly defined to encompass any device having a processor which executes instructions from a memory medium.

[00478] It is noted that, in one or more embodiments, one or more of the method elements described herein and/or one or more portions of an implementation of a method element can be performed in varying orders, can be repeated, can be performed concurrently with one or more of the other method elements and/or one or more portions of an implementation of a method element, or can be omitted. Additional and/or duplicated method elements can be performed as desired. For example, a process and/or method can perform one or more described method elements concurrently with duplicates of the one or more described method elements. For instance, multiple methods, processes, and/or threads can be implemented using same described method elements.

[00479] In one or more embodiments, concurrently can mean simultaneously. In one or more embodiments, concurrently can mean apparently simultaneously according to some

metric. For example, two or more method elements and/or two or more portions of an implementation of a method element can be performed such that they appear to be simultaneous to a human. It is also noted that, in one or more embodiments, one or more of the system elements described herein may be omitted and additional system elements can be added as desired.

[00480] Further modifications and alternative embodiments of various aspects of the invention may be apparent to those skilled in the art in view of this description. Accordingly, this description is to be construed as illustrative only and is for the purpose of teaching those skilled in the art the general manner of carrying out the invention. It is to be understood that the forms of the invention shown and described herein are to be taken as embodiments. Elements and materials may be substituted for those illustrated and described herein, parts and processes may be reversed, and certain features of the invention may be utilized independently, all as would be apparent to one skilled in the art after having the benefit of this description of the invention. Changes may be made in the elements described herein without departing from the spirit and scope of the invention as described in the following claims.

WHAT IS CLAIMED IS:

1. A method, comprising:

receiving identification information associated with a telephony device operated by a user;

providing, to the telephony device, media associated with a plurality of call states;

determining a message for the user based on at least one of the identification information associated with the telephony device and a current call state associated with the telephony device from among the plurality of call states;

receiving, from the telephony device, user input information indicating a request for or an acceptance of the message; and

in response to receiving, from the telephony device, the user input information indicating the request for or the acceptance of the message, sending the message to the telephony device.

2. A system, comprising:

a processor; and

a memory medium coupled to the processor;

wherein the memory medium includes instructions, which when executed by the processor, cause the system to perform:

receiving identification information associated with a telephony device operated by a user;

querying, based on the identification information associated with the telephony device, a database for profile information of the user;

providing a plurality of menu options to the telephony device based on the profile information of the user;

receiving, from the telephony device, user input information indicating one or more selections from the plurality of menu options;

determining a message for the user based on the user input information indicating the one or more selections from the plurality of menu options; and

sending the message to the telephony device.

3. A method, comprising:

exchanging a plurality of media presentations of a plurality of call portions for compensation;

receiving a plurality of calls from a plurality of callers; and

providing the plurality of media presentations of the plurality of call portions to the callers via a call management system.

4. A system, comprising:

an auction server, coupled to a network, configured to exchange a plurality of media presentations of a plurality of call portions for compensation; and

a call management system, coupled to the network, configured to:

receive a plurality of calls from a plurality of callers; and

provide the plurality of media presentations of the plurality of call portions to the callers via a call management system.

5. A method, comprising:

receiving a call from a telephony device operated by a user;

determining identification information associated with the telephony device;

querying, based on the identification information associated with the telephony device, a database for profile information of the user;

building a menu that includes a plurality of menu options, wherein the plurality of menu options are based on the identification information associated with the telephony device and the profile information;

providing the menu to the telephony device;

receiving, from the telephony device, user input information indicating one or more selections from the plurality of menu options.

6. A system, comprising:

a processor;

a memory medium coupled to the processor;

a router coupled to the processor via a first network connection;

a database coupled to processor via a second network connection;

a voice gateway coupled to processor via a network connection;

wherein the router is configured to receive call from telephony devices;

wherein the voice gateway is configured to receive input from the telephony devices and provide the user input to the processor via the second network connection; and

wherein the memory medium includes instructions, which when executed by the processor, cause the system to perform:

receiving, via the router, a call from a telephony device operated by a user;

determining identification information associated with the telephony device;

querying, based on the identification information associated with the telephony device, the database for profile information of the user;

building a menu that includes a plurality of menu options, wherein the plurality of menu options are based on the identification information associated with the telephony device and the profile information;

providing the menu to the telephony device;

receiving, from the telephony device and via the voice gateway, user input information indicating one or more selections from the plurality of menu options.

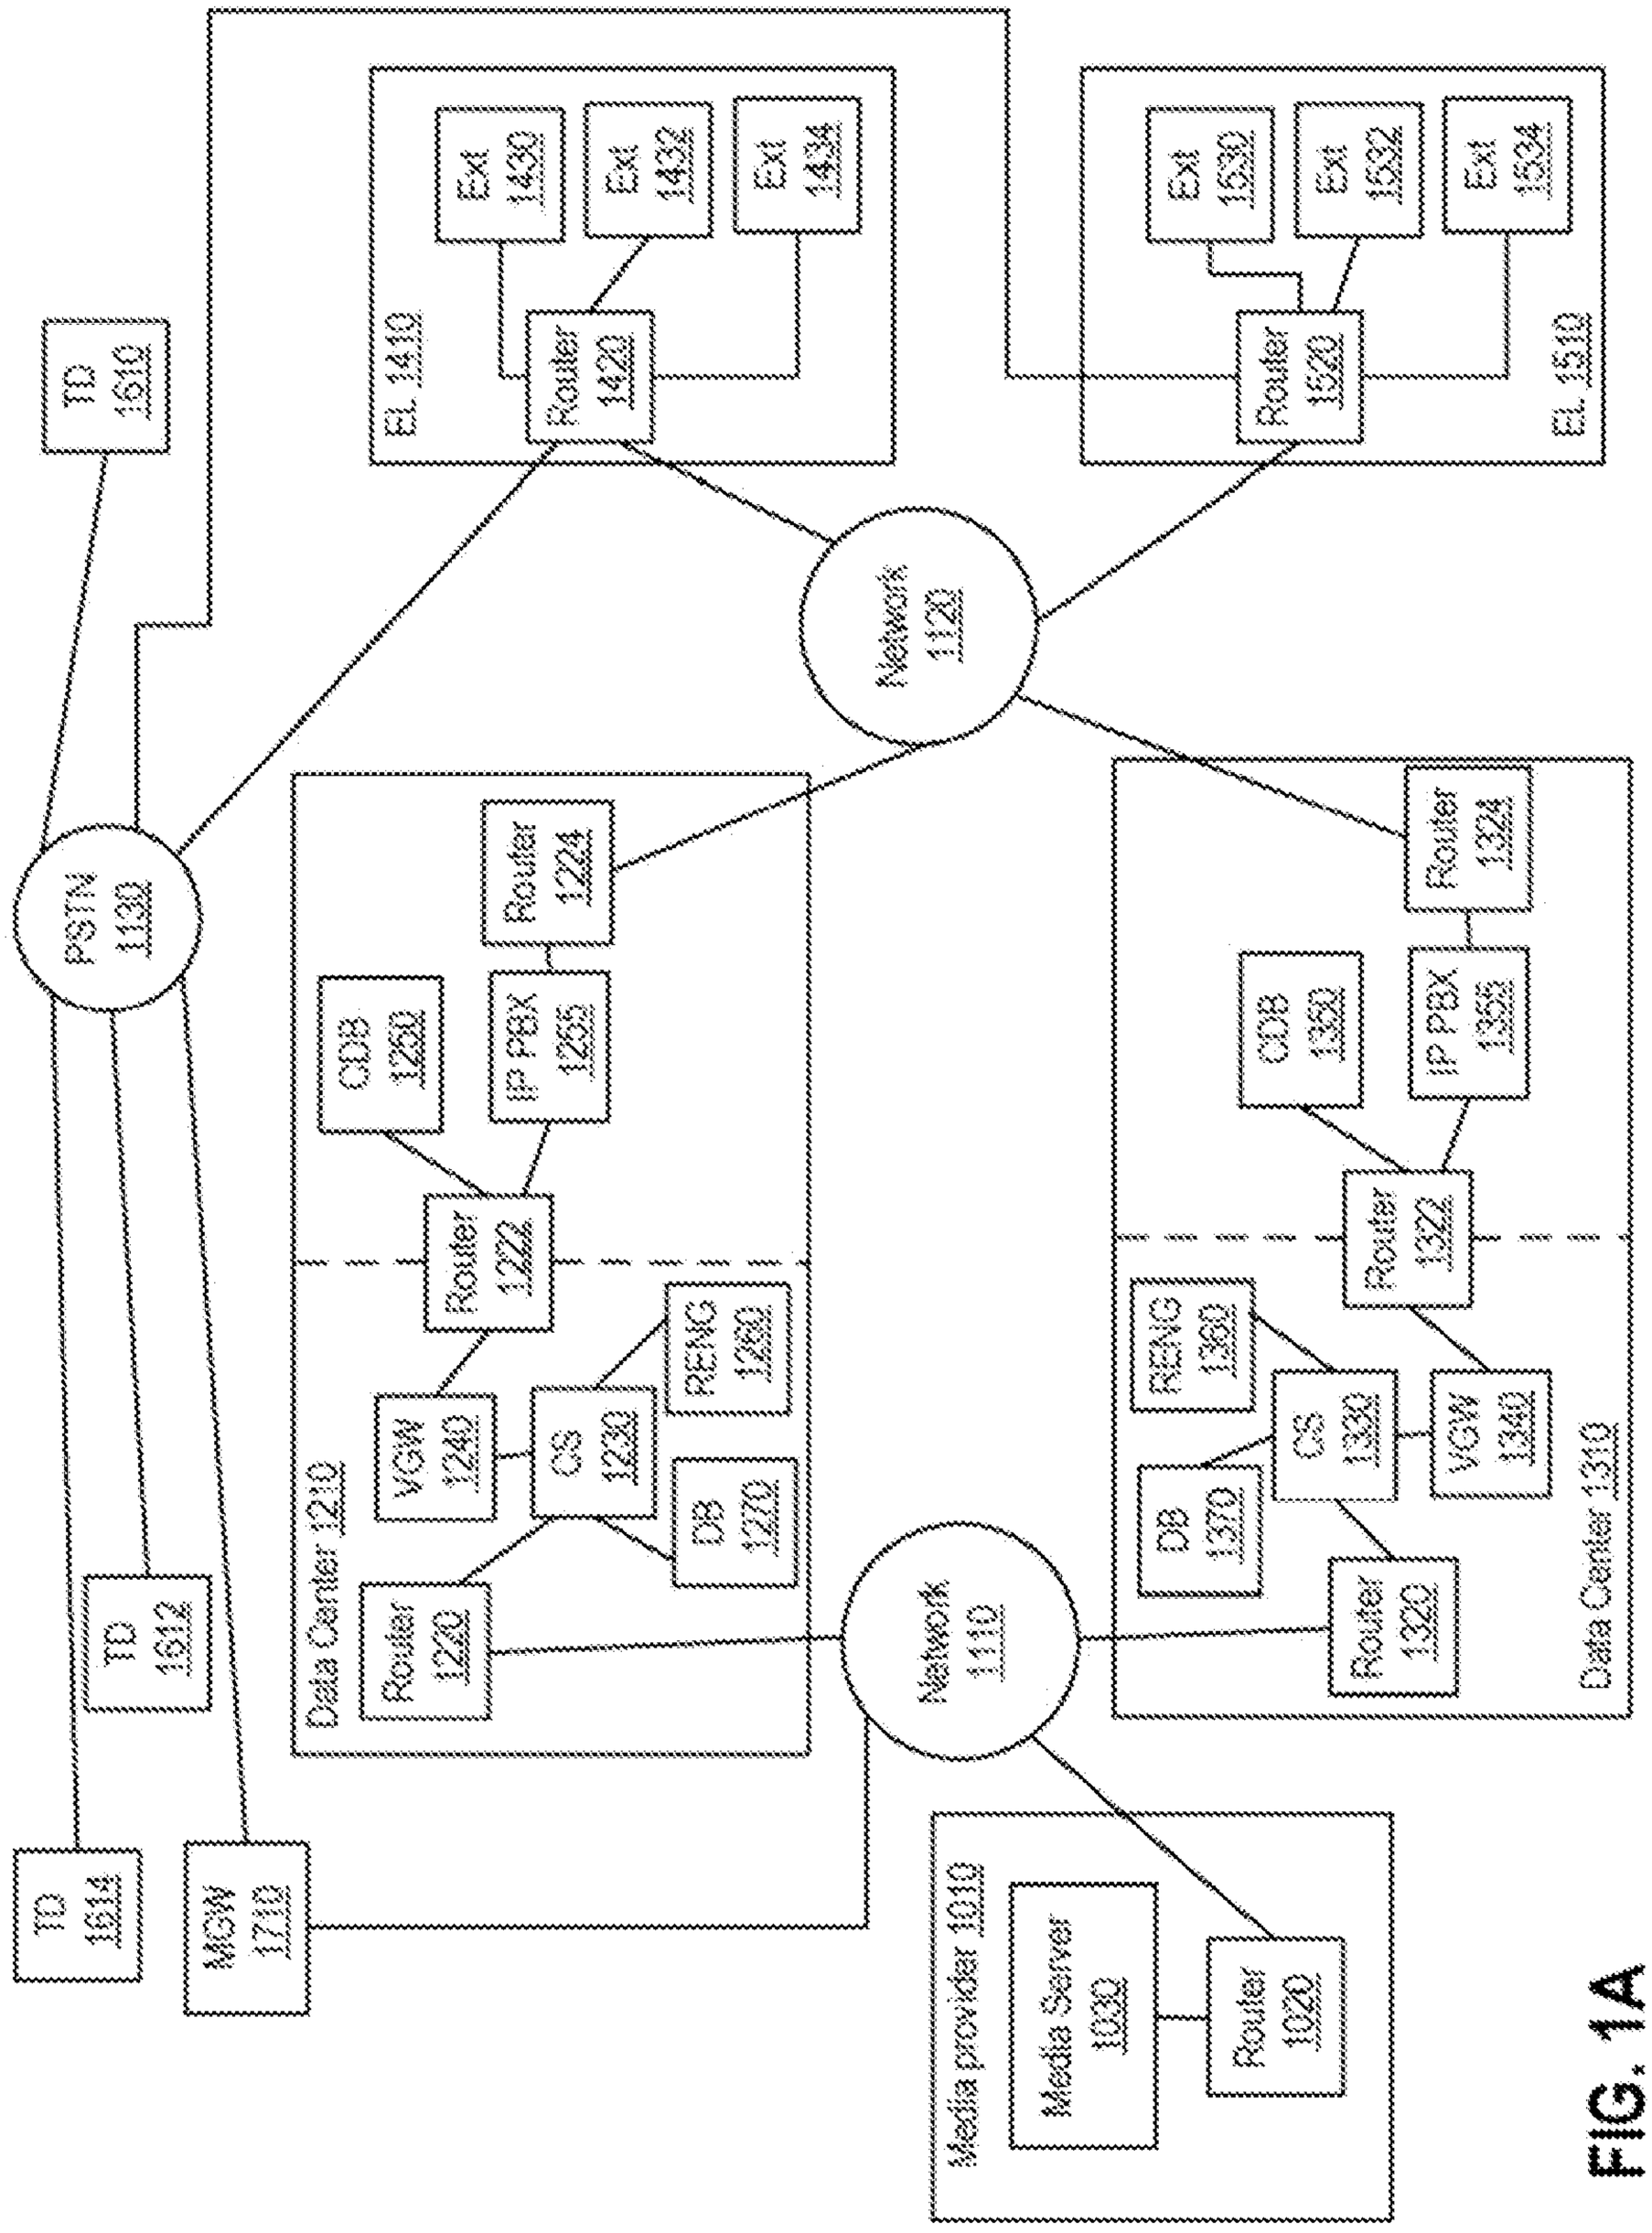
TD 1610
EL 1410
Ext 1430
Ext 1432
Ext 1434
Router 1420
EL 1510
Ext 1530
Ext 1532
Ext 1534
Router 1520
Network 1120
PSTN 1130
TD 1612
TD 1614
Data Center 1210
Router 1224
CDB 1250
IP PBX 1255
Router 1222
VGW 1240
CS 1230
RENG 1260
DB 1270
Router 1220
MGW 1710
Network 1110
Data Center 1310
Router 1324
CDB 1350
IP PBX 1355
Router 1322
RENG 1360
DB 1370
CS 1330
VGW 1340
Router 1320
Media provider 1010
Media Server 1030
Router 1020

FIG. 1A

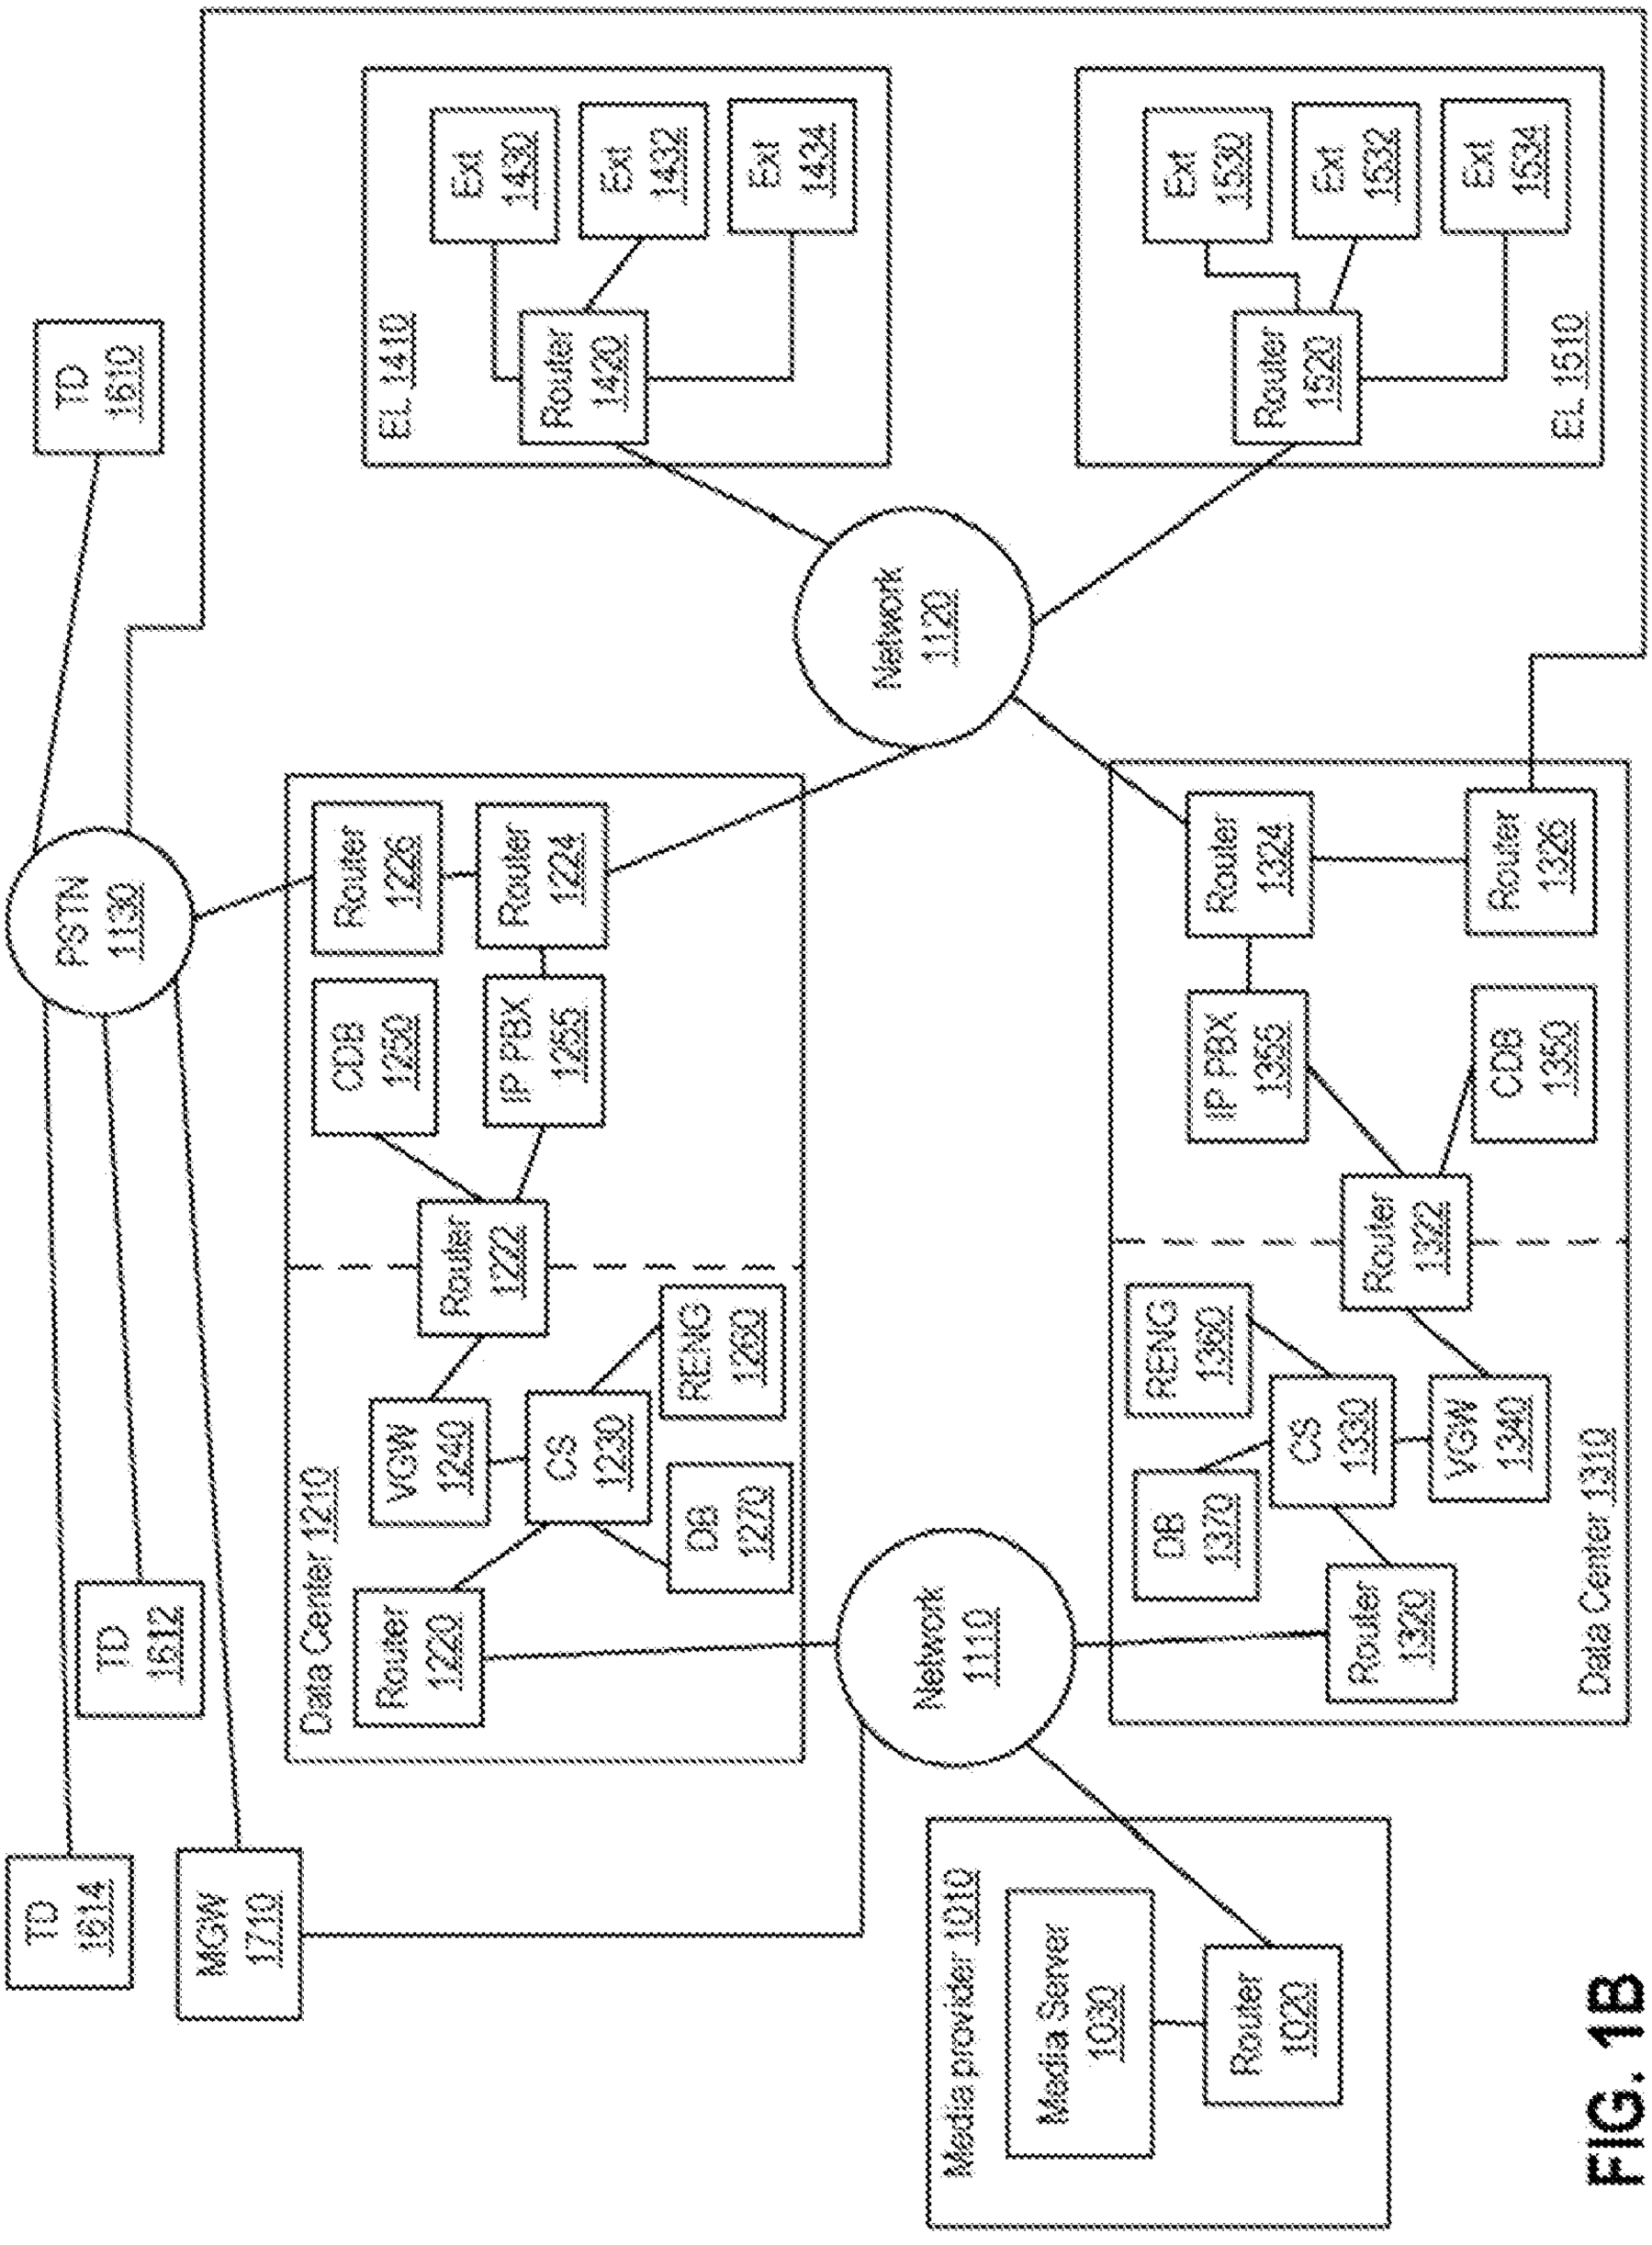
TD 1610
EL 1410
Ext 1430
Ext 1432
Ext 1434
Router 1420
EL 1510
Ext 1530
Ext 1532
Ext 1534
Router 1520
Network 1120
PSTN 1130
Router 1226
Router 1224
CDB 1250
IP PBX 1255
Router 1222
VGW 1240
CS 1230
RENG 1260
DB 1270
Router 1220
Data Center 1210
Router 1324
Router 1326
IP PBX 1355
CDB 1350
Router 1322
RENG 1360
CS 1330
VGW 1340
DB 1370
Router 1320
Data Center 1310
Network 1110
TD 1612
TD 1614
MGW 1710
Media provider 1010
Media Server 1030
Router 1020

FIG. 1B

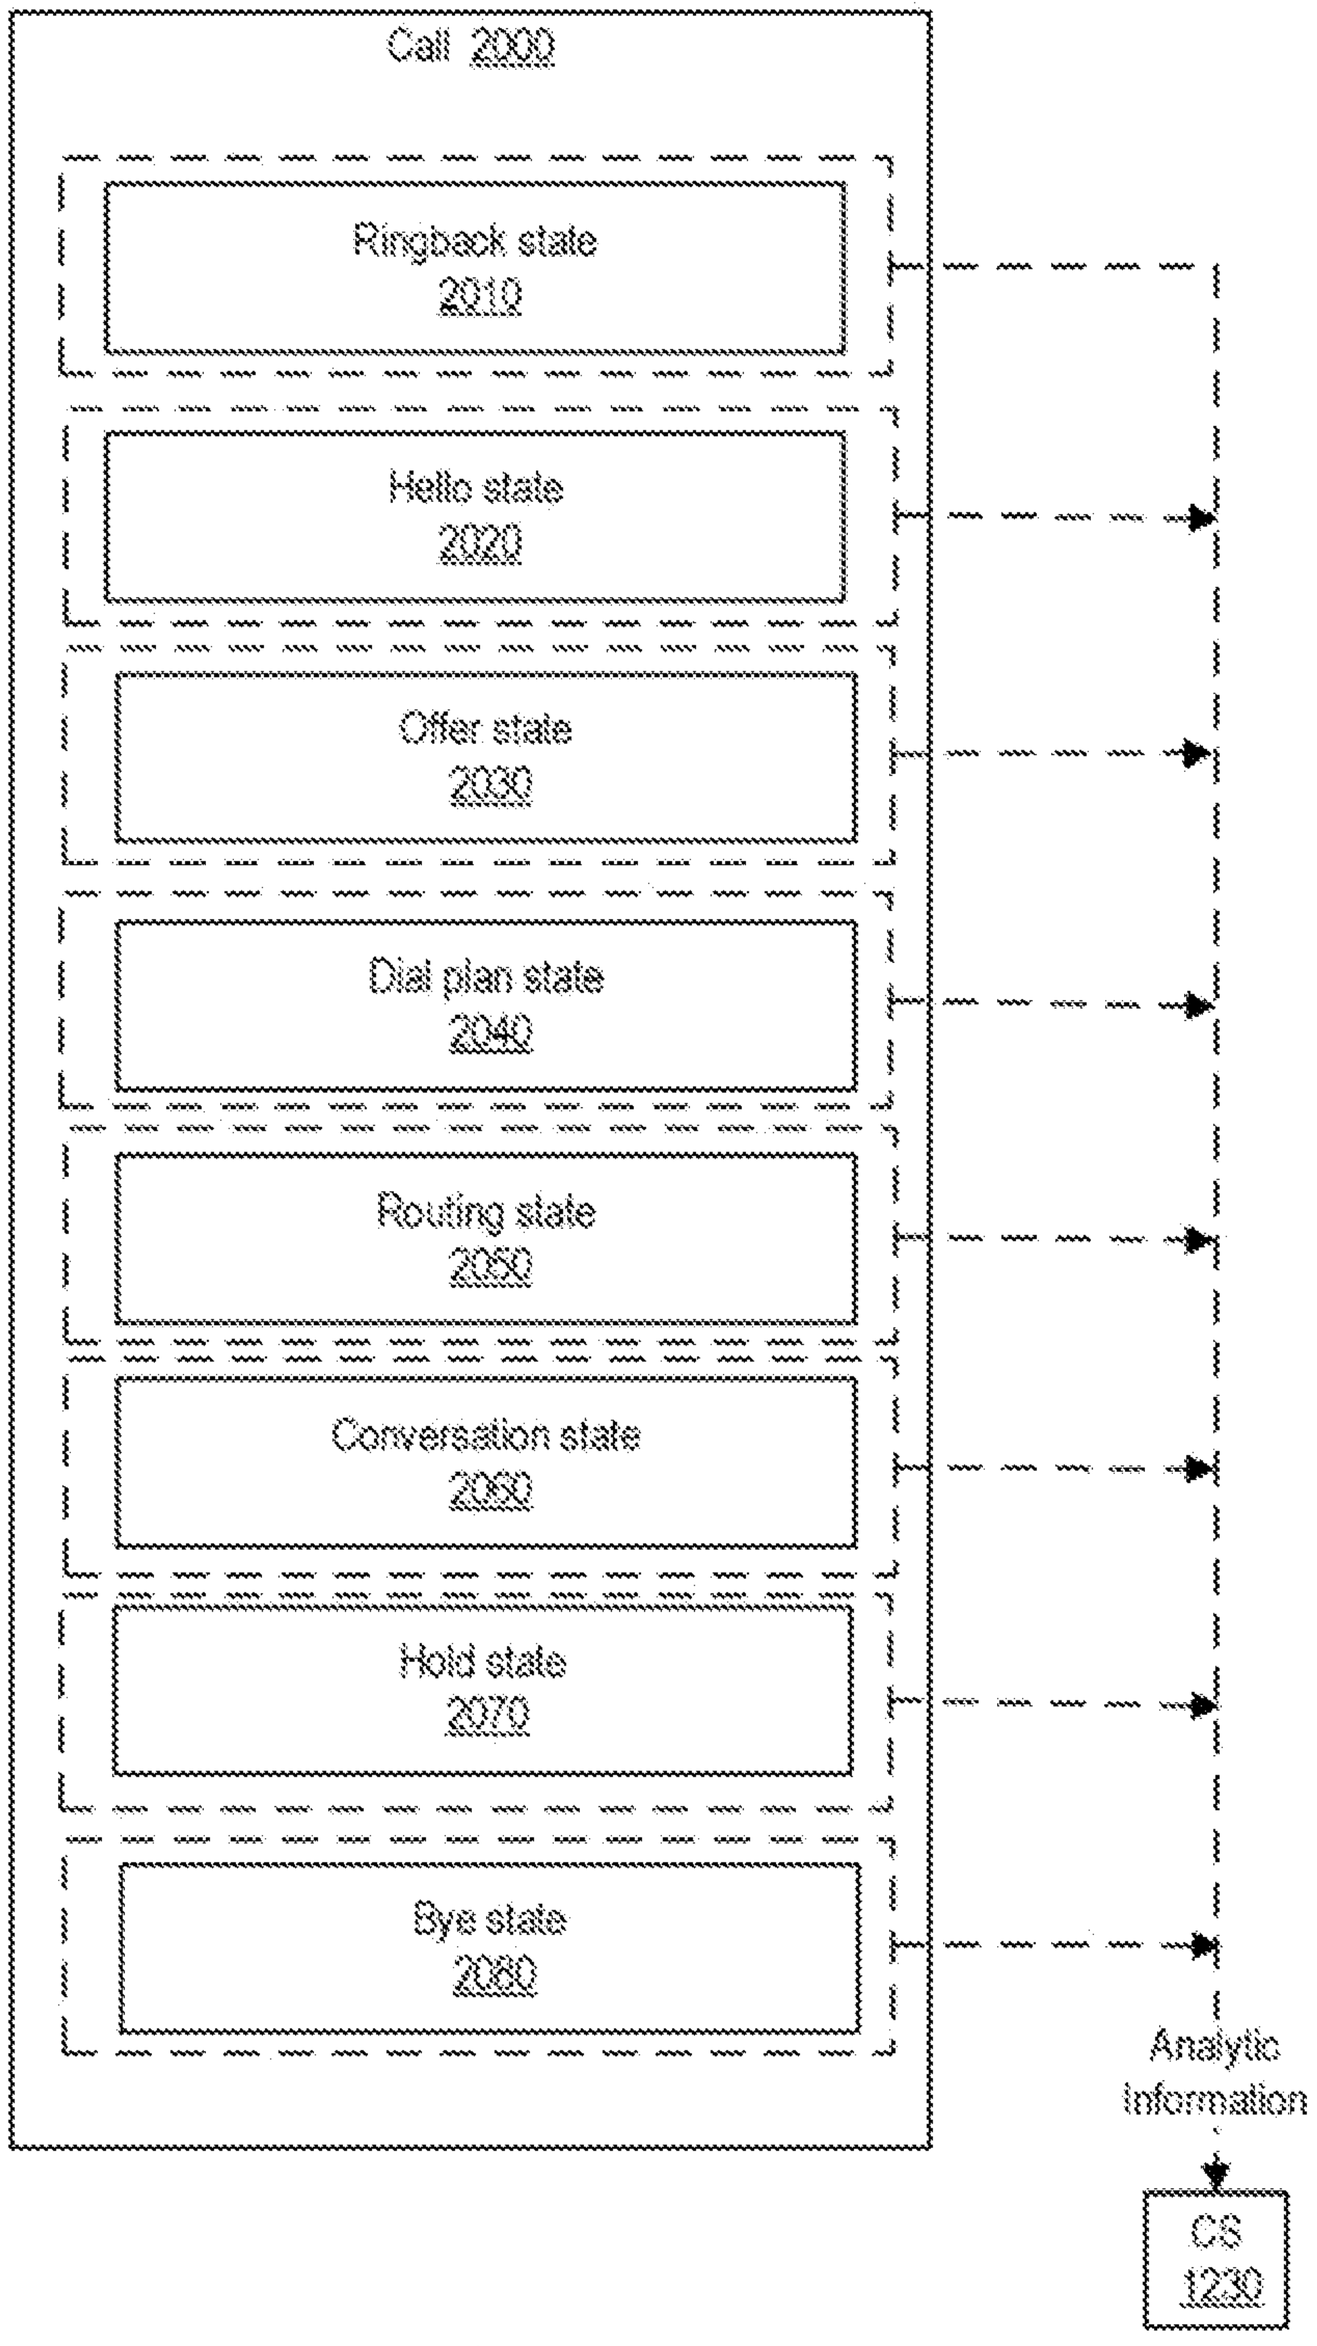
Call 2000
Ringback state 2010
Hello state 2020
Offer state 2030
Dial plan state 2040
Routing state 2050
Conversation state 2060
Hold state 2070
Bye state 2080
Analytic Information
CS 1230

FIG. 2

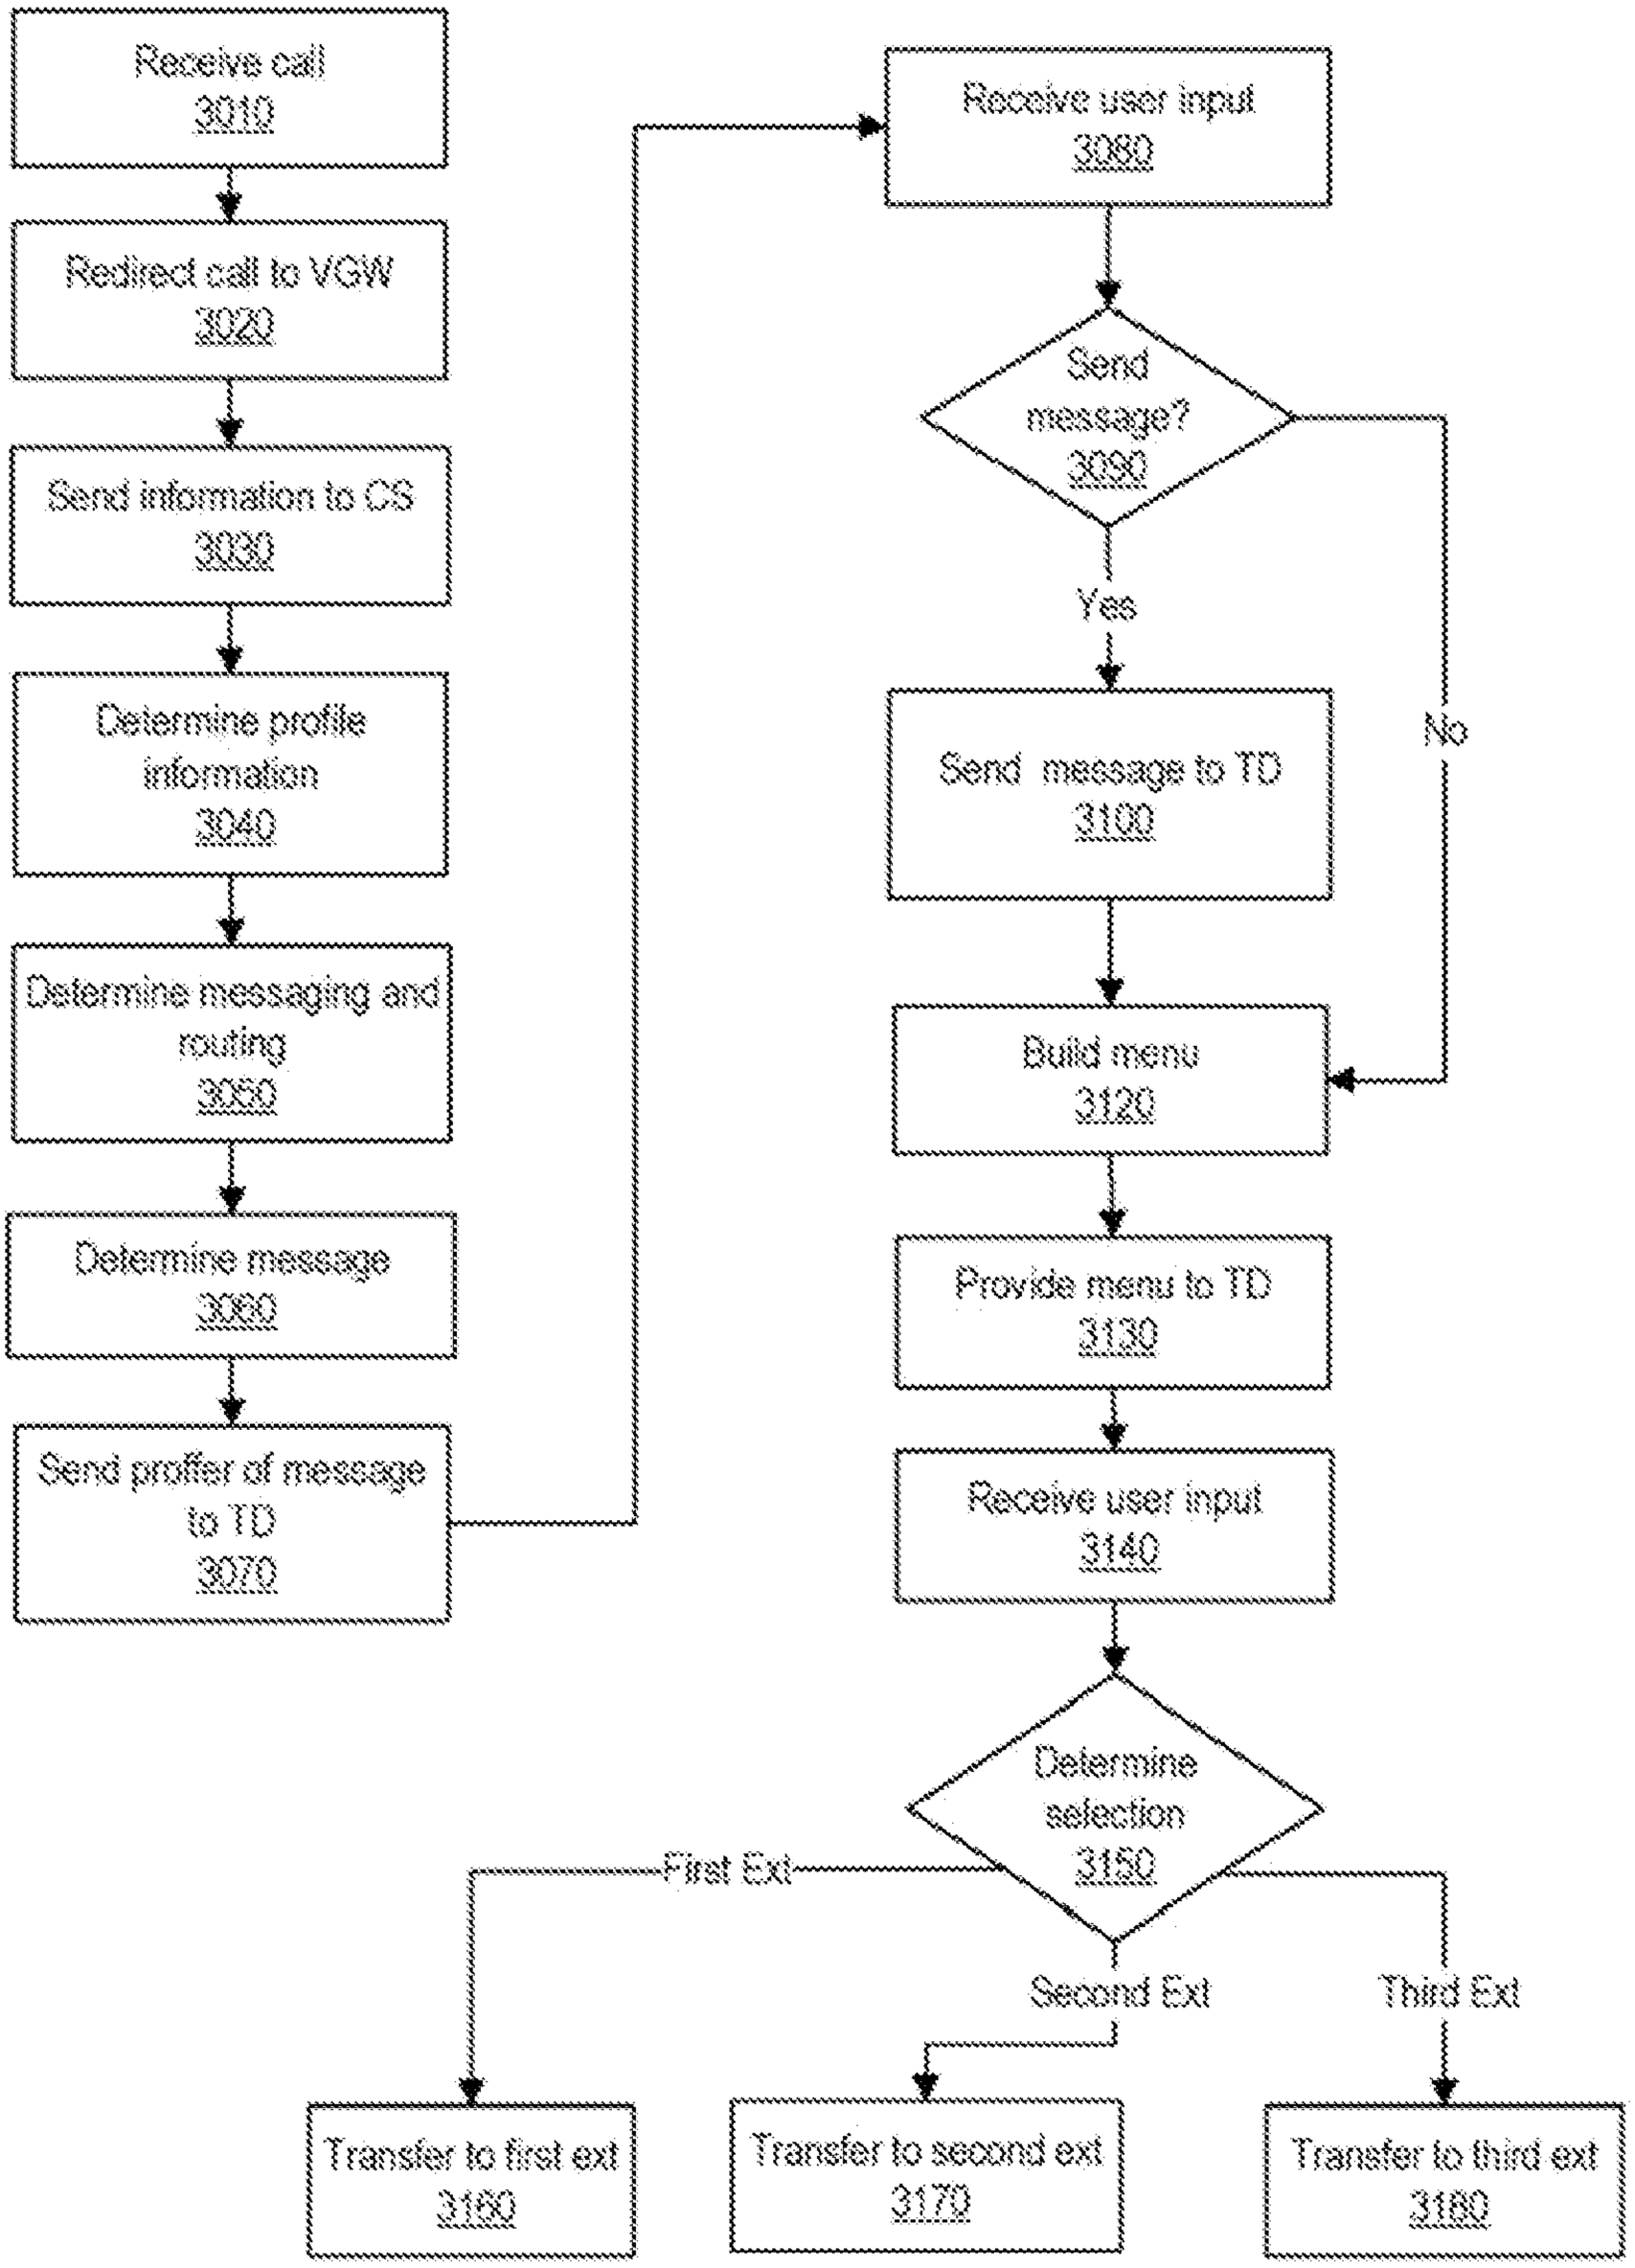
Receive call
3010
Redirect call to VGW
3020
Send information to CS
3030
Determine profile information
3040
Determine messaging and routing
3050
Determine message
3060
Send proffer of message to TD
3070
Receive user input
3080
Send message?
3090
Yes
No
Send message to TD
3100
Build menu
3120
Provide menu to TD
3130
Receive user input
3140
Determine selection
3150
First Ext
Second Ext
Third Ext
Transfer to first ext
3160
Transfer to second ext
3170
Transfer to third ext
3180

FIG. 3

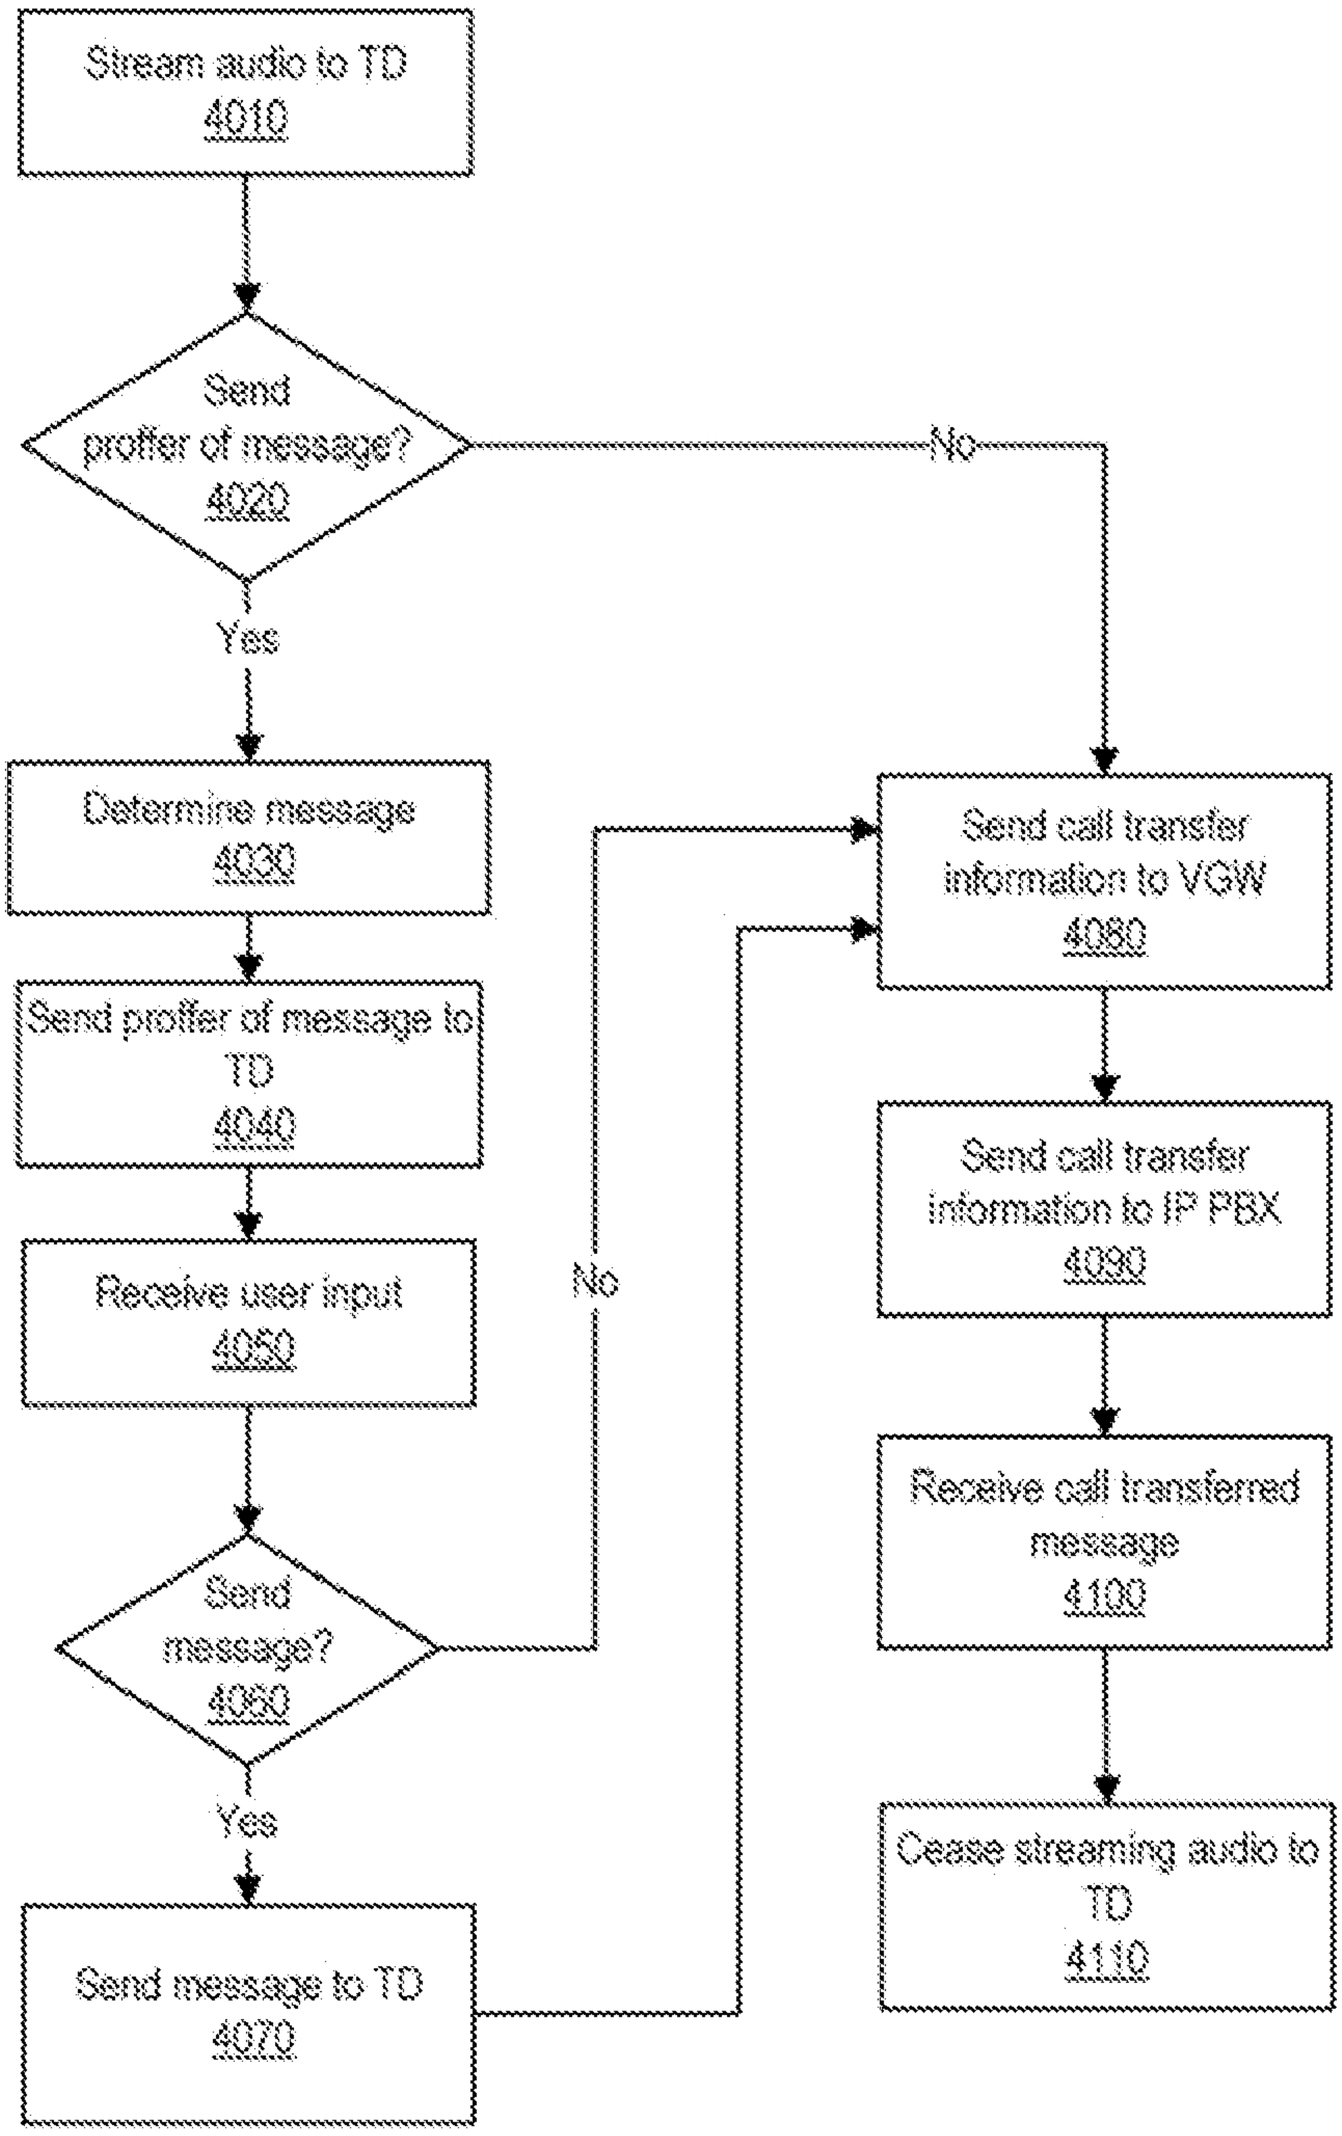
Stream audio to TD
4010
Send proffer of message?
4020
No
Yes
Determine message
4030
Send proffer of message to TD
4040
Receive user input
4050
Send message?
4060
No
Yes
Send message to TD
4070
Send call transfer information to VGW
4080
Send call transfer information to IP PBX
4090
Receive call transferred message
4100
Cease streaming audio to TD
4110

FIG. 4

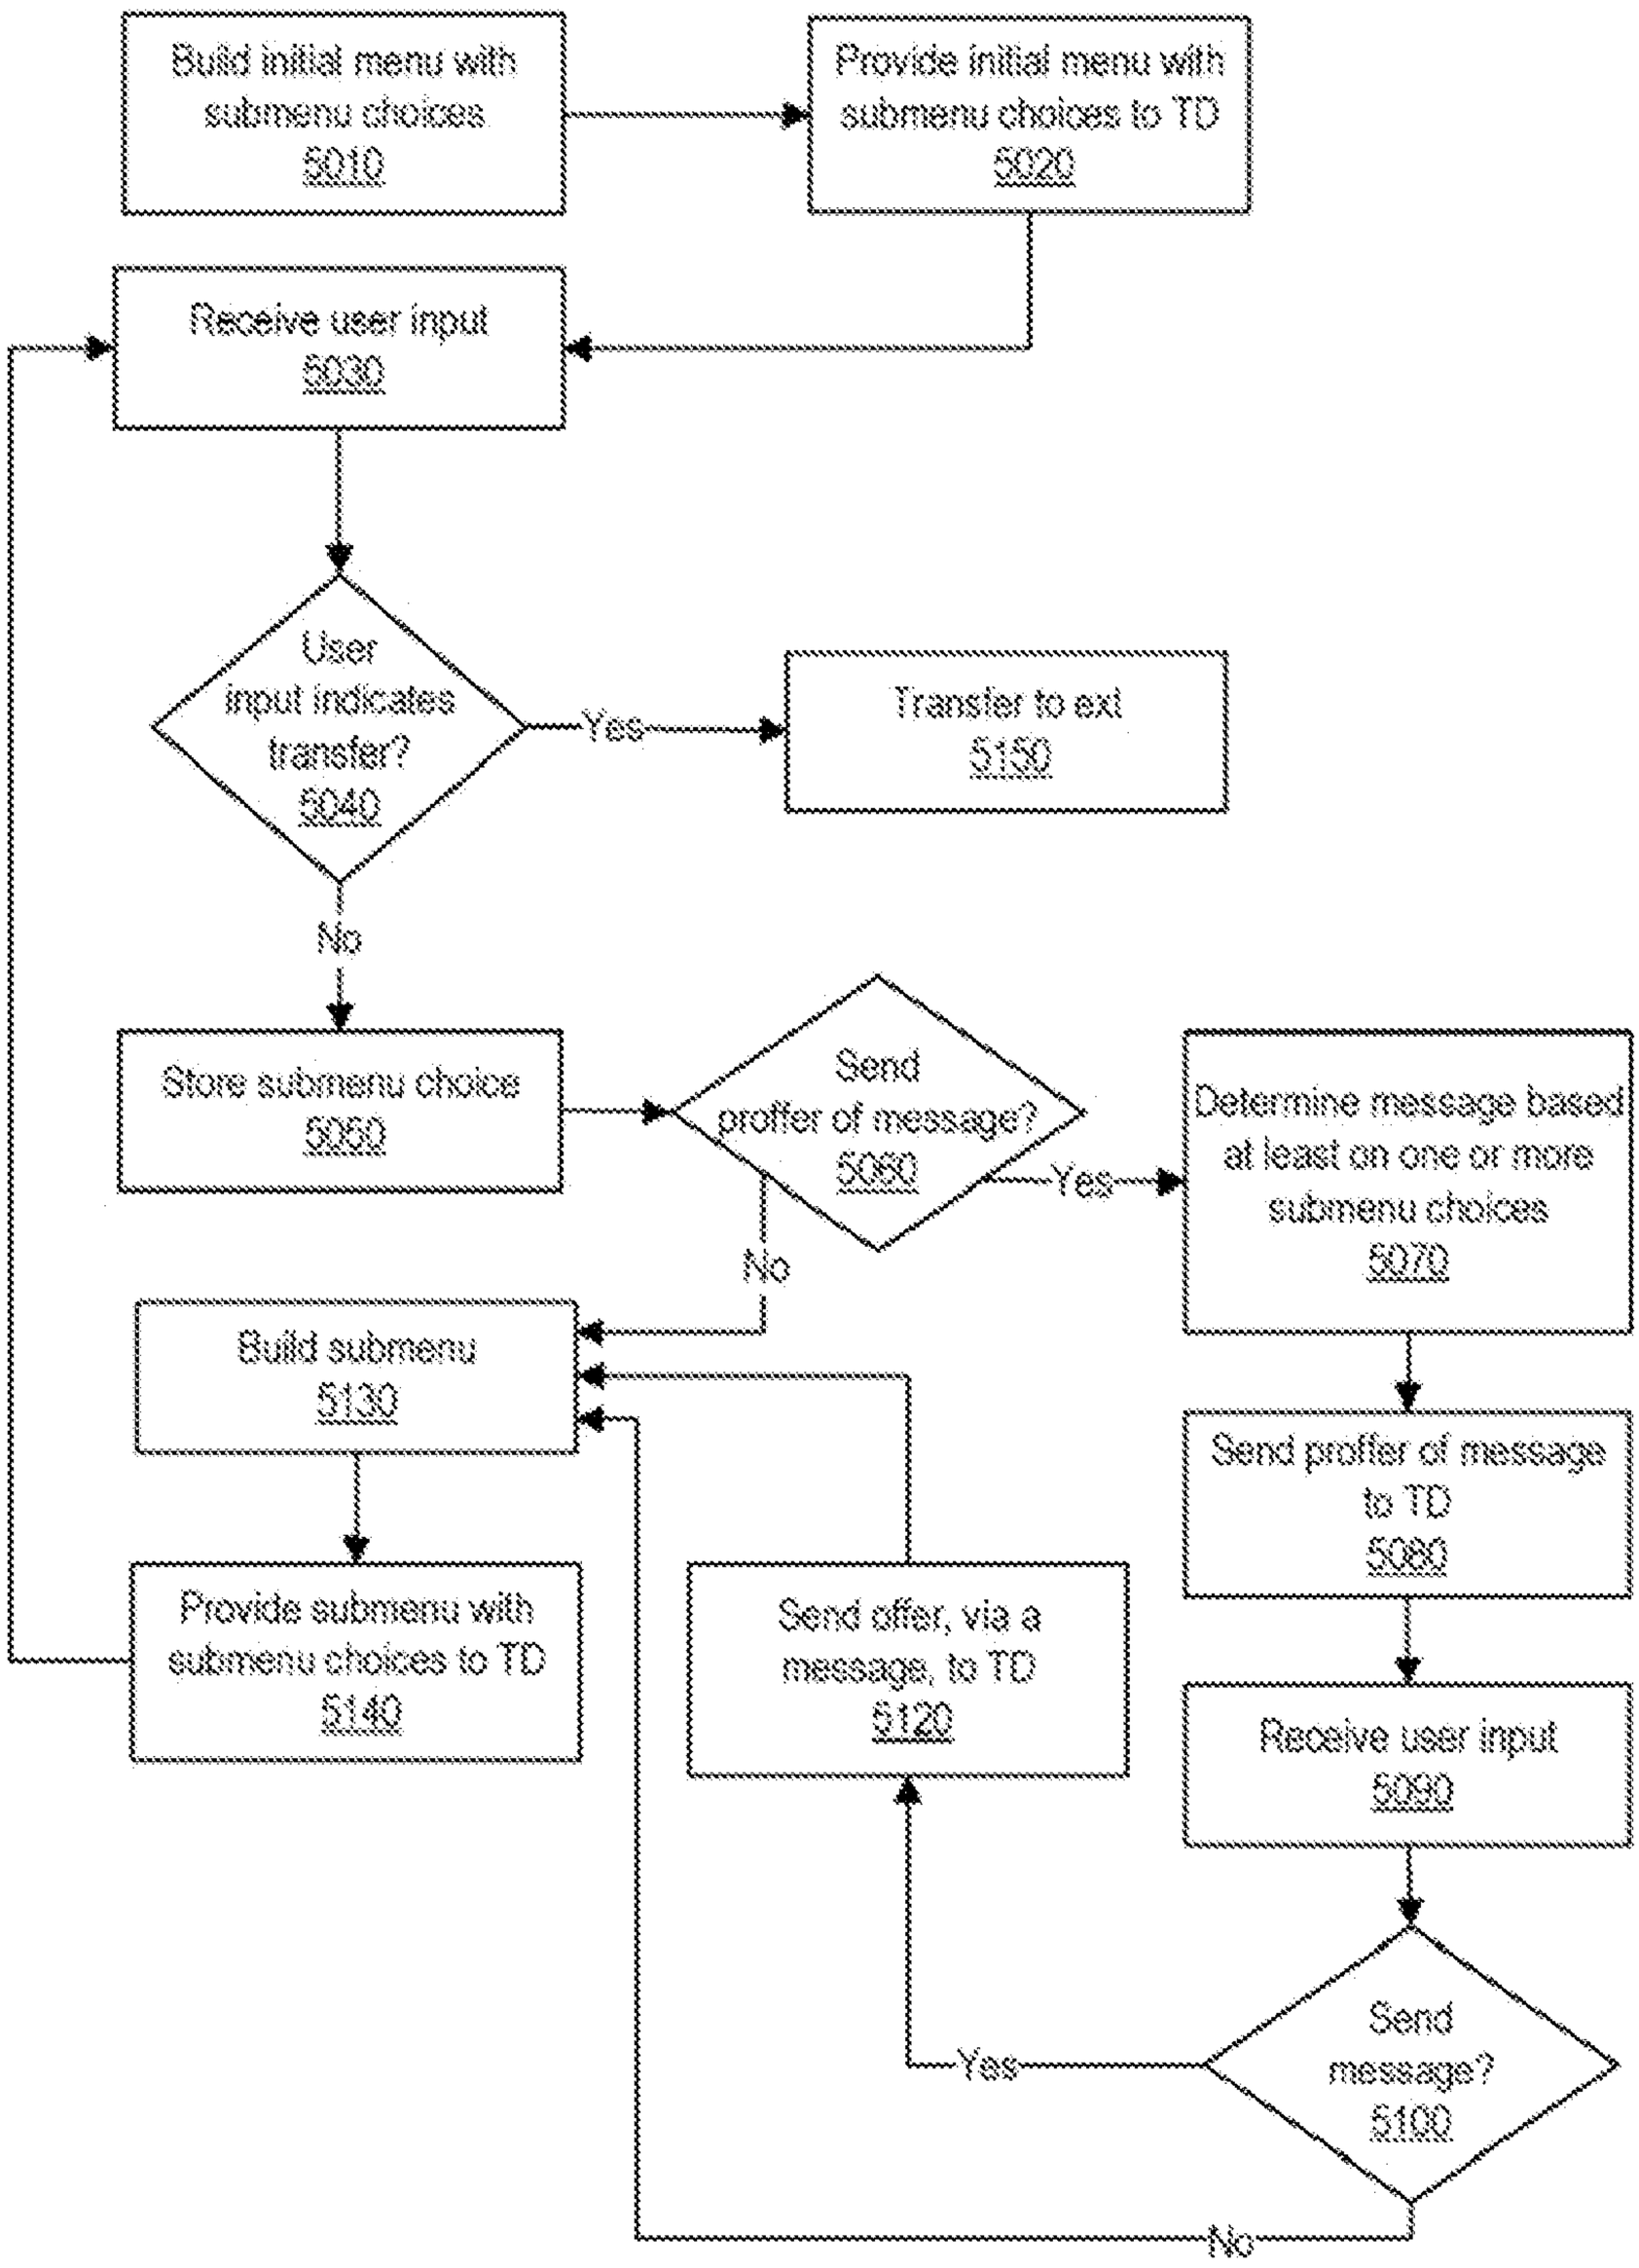
Build initial menu with submenu choices 5010
Provide initial menu with submenu choices to TD 5020
Receive user input 5030
User input indicates transfer? 5040
Yes
Transfer to ext 5150
No
Store submenu choice 5050
Send proffer of message? 5060
Yes
Determine message based at least on one or more submenu choices 5070
No
Build submenu 5130
Send proffer of message to TD 5080
Provide submenu with submenu choices to TD 5140
Send offer, via a message, to TD 5120
Receive user input 5090
Send message? 5100
Yes
No

FIG. 5

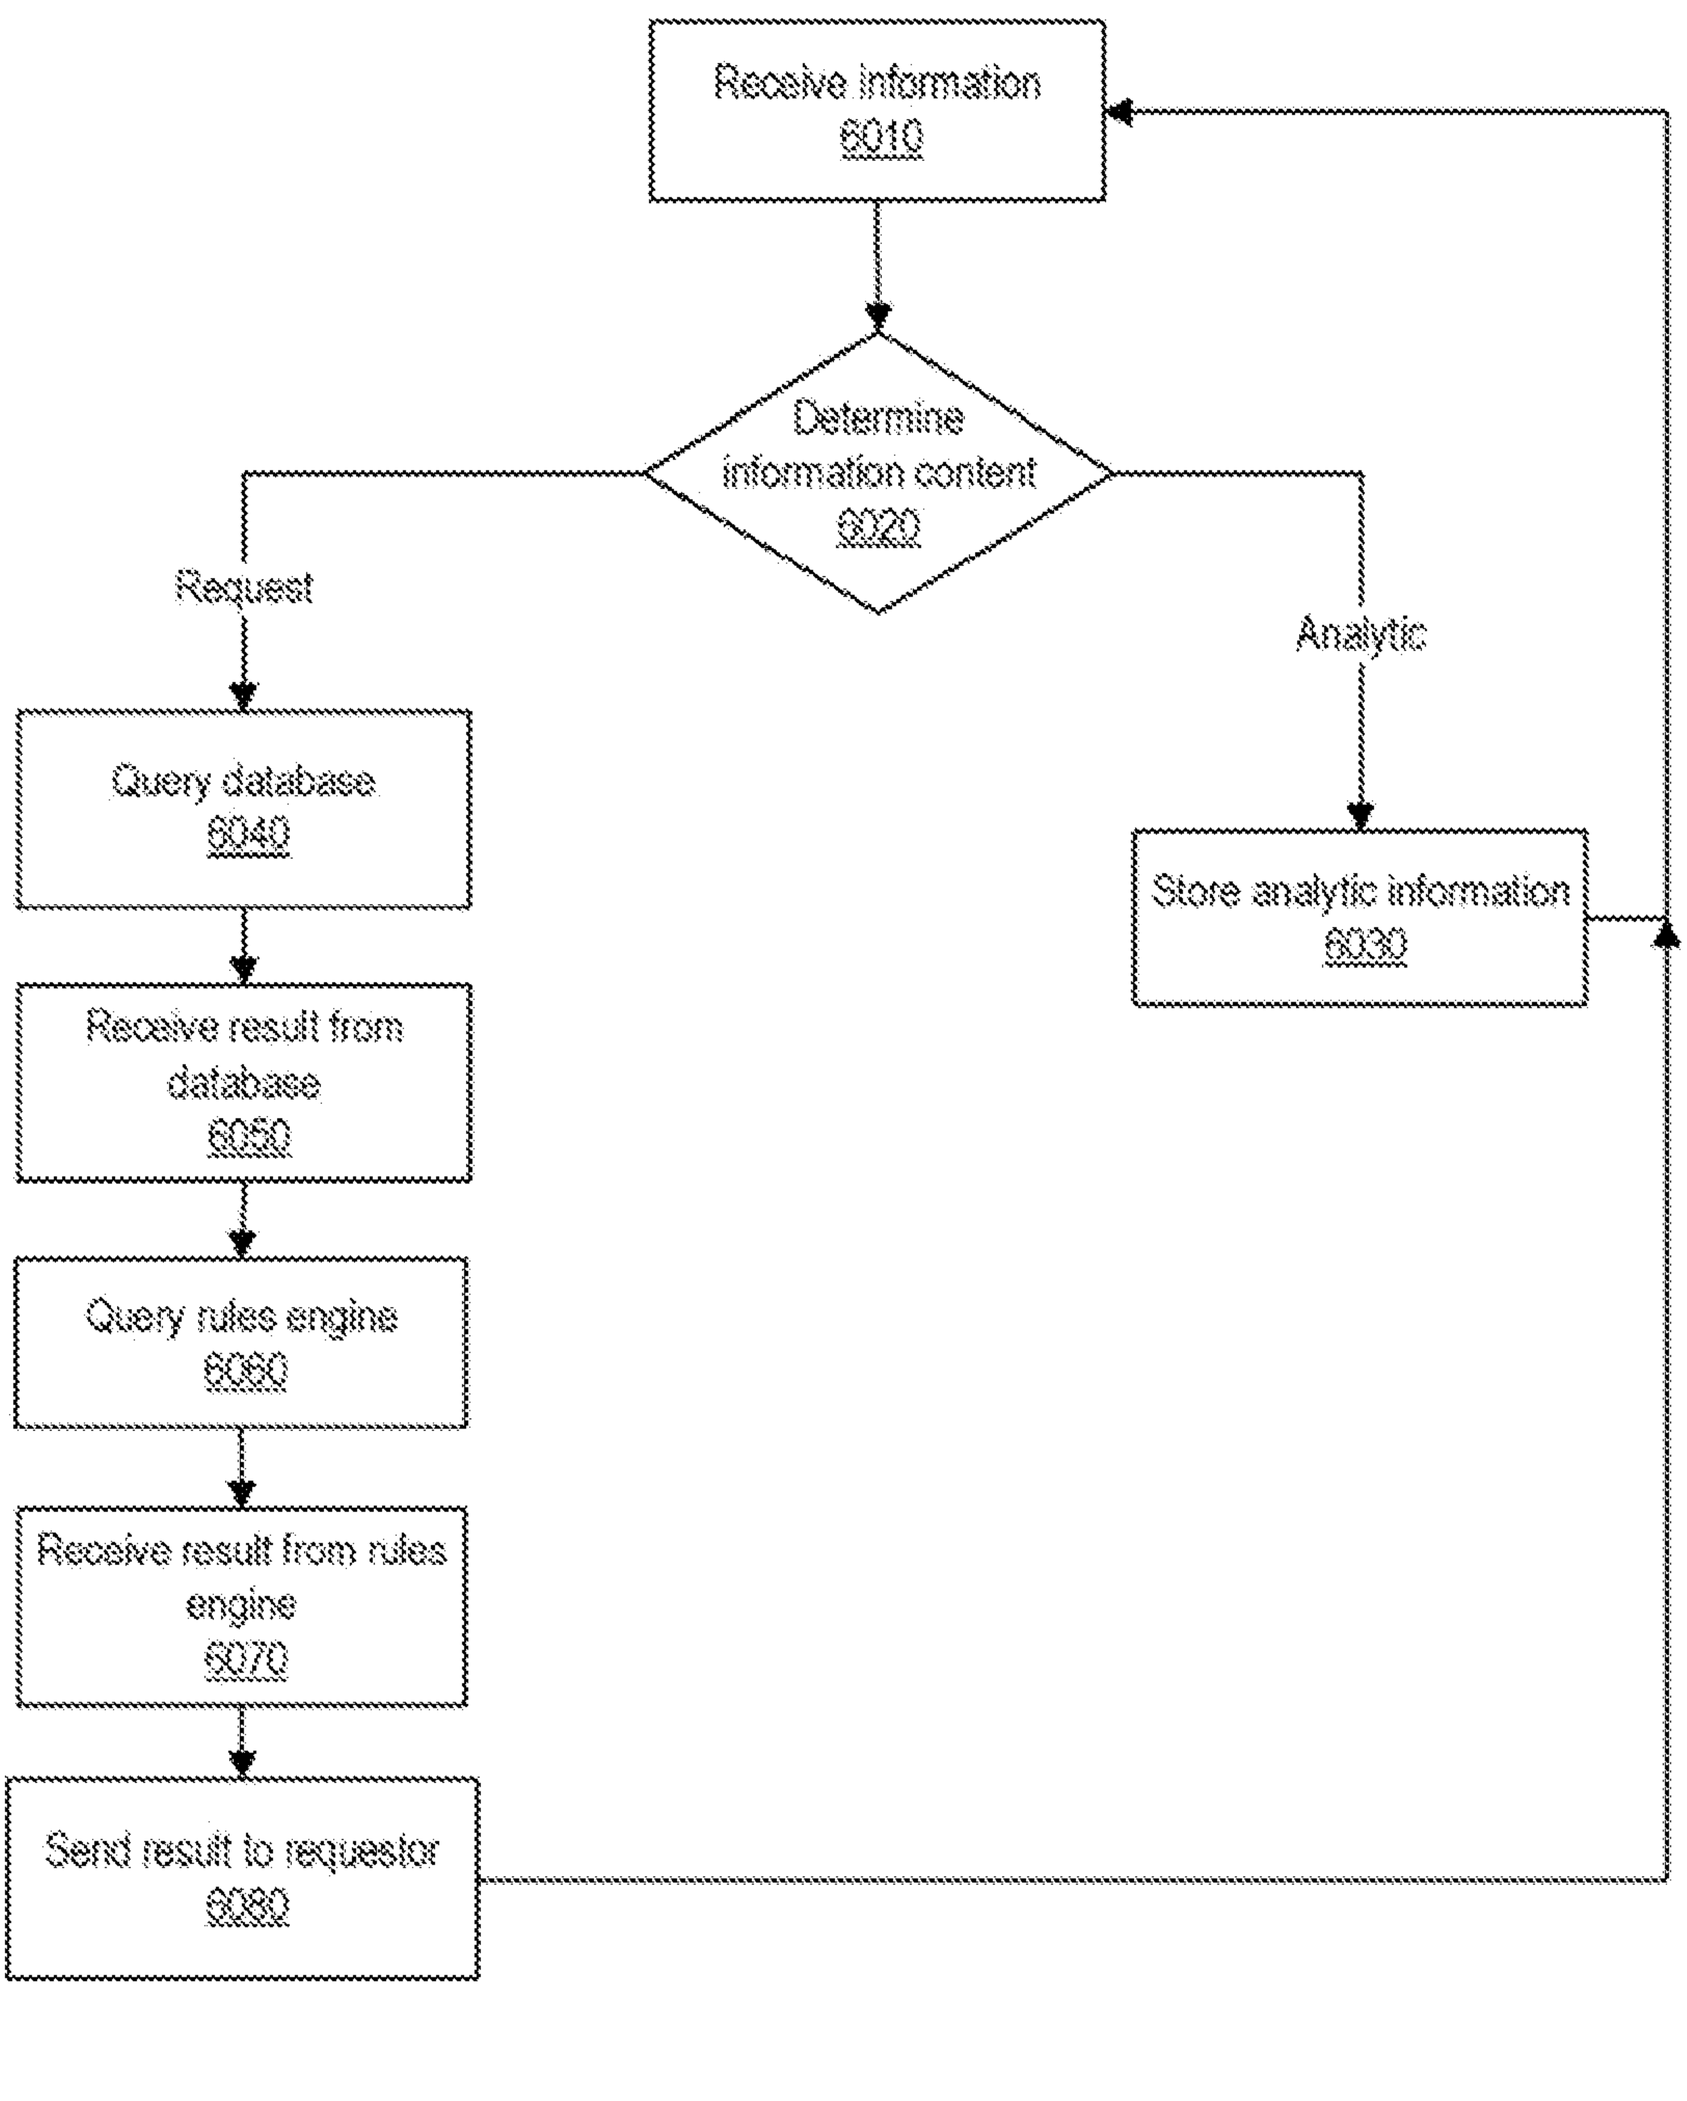
Receive information
6010
Determine information content
6020
Request
Analytic
Query database
6040
Receive result from database
6050
Query rules engine
6060
Receive result from rules engine
6070
Send result to requestor
6080
Store analytic information
6030

FIG. 6

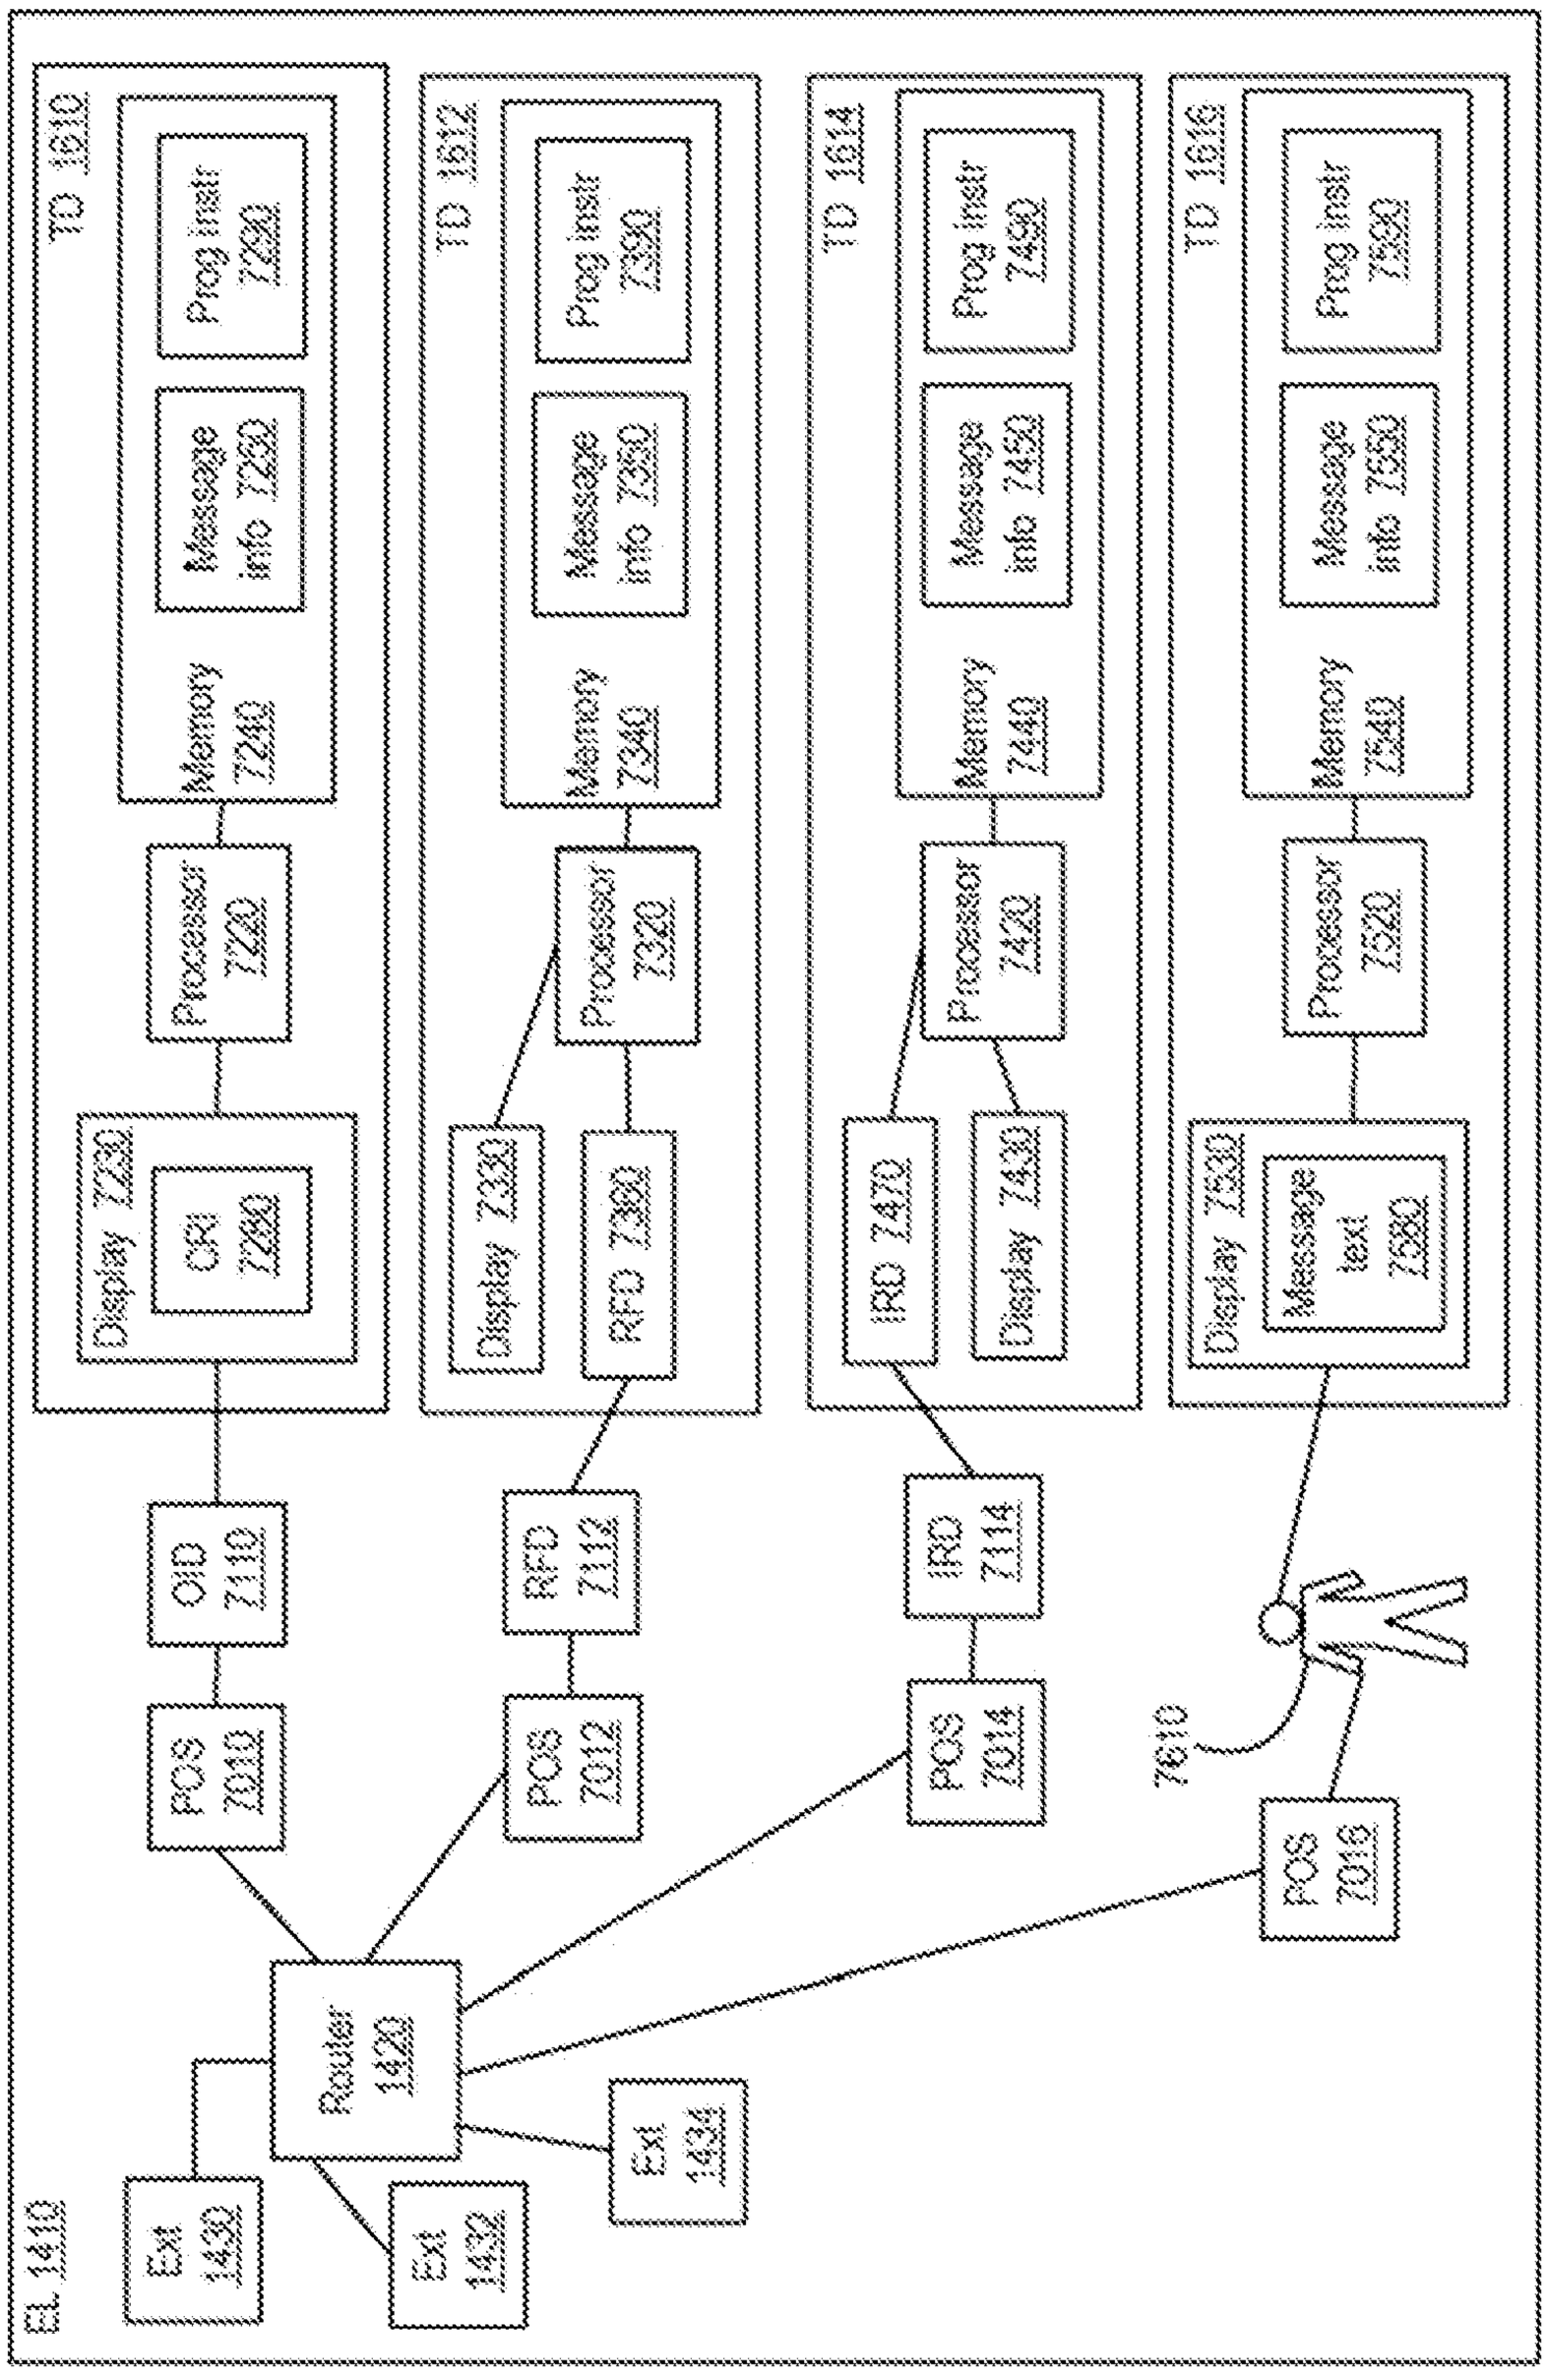
EL 1410
Router 1420
Ext 1430
Ext 1432
Ext 1434
OID 7110
RFD 7112
IRD 7114
POS 7010
POS 7012
POS 7014
POS 7016
7e10
TD 1610
Processor 7220
Memory 7240
Message info 7260
Prog instr 7280
Display 7230
ORI 7290
TD 1612
Processor 7320
Memory 7340
Message info 7360
Prog instr 7380
Display 7330
RFD 7390
TD 1614
Processor 7420
Memory 7440
Message info 7460
Prog instr 7480
IRD 7470
Display 7430
TD 1616
Processor 7520
Memory 7540
Message info 7560
Prog instr 7580
Display 7530
Message text 7590

FIG. 7

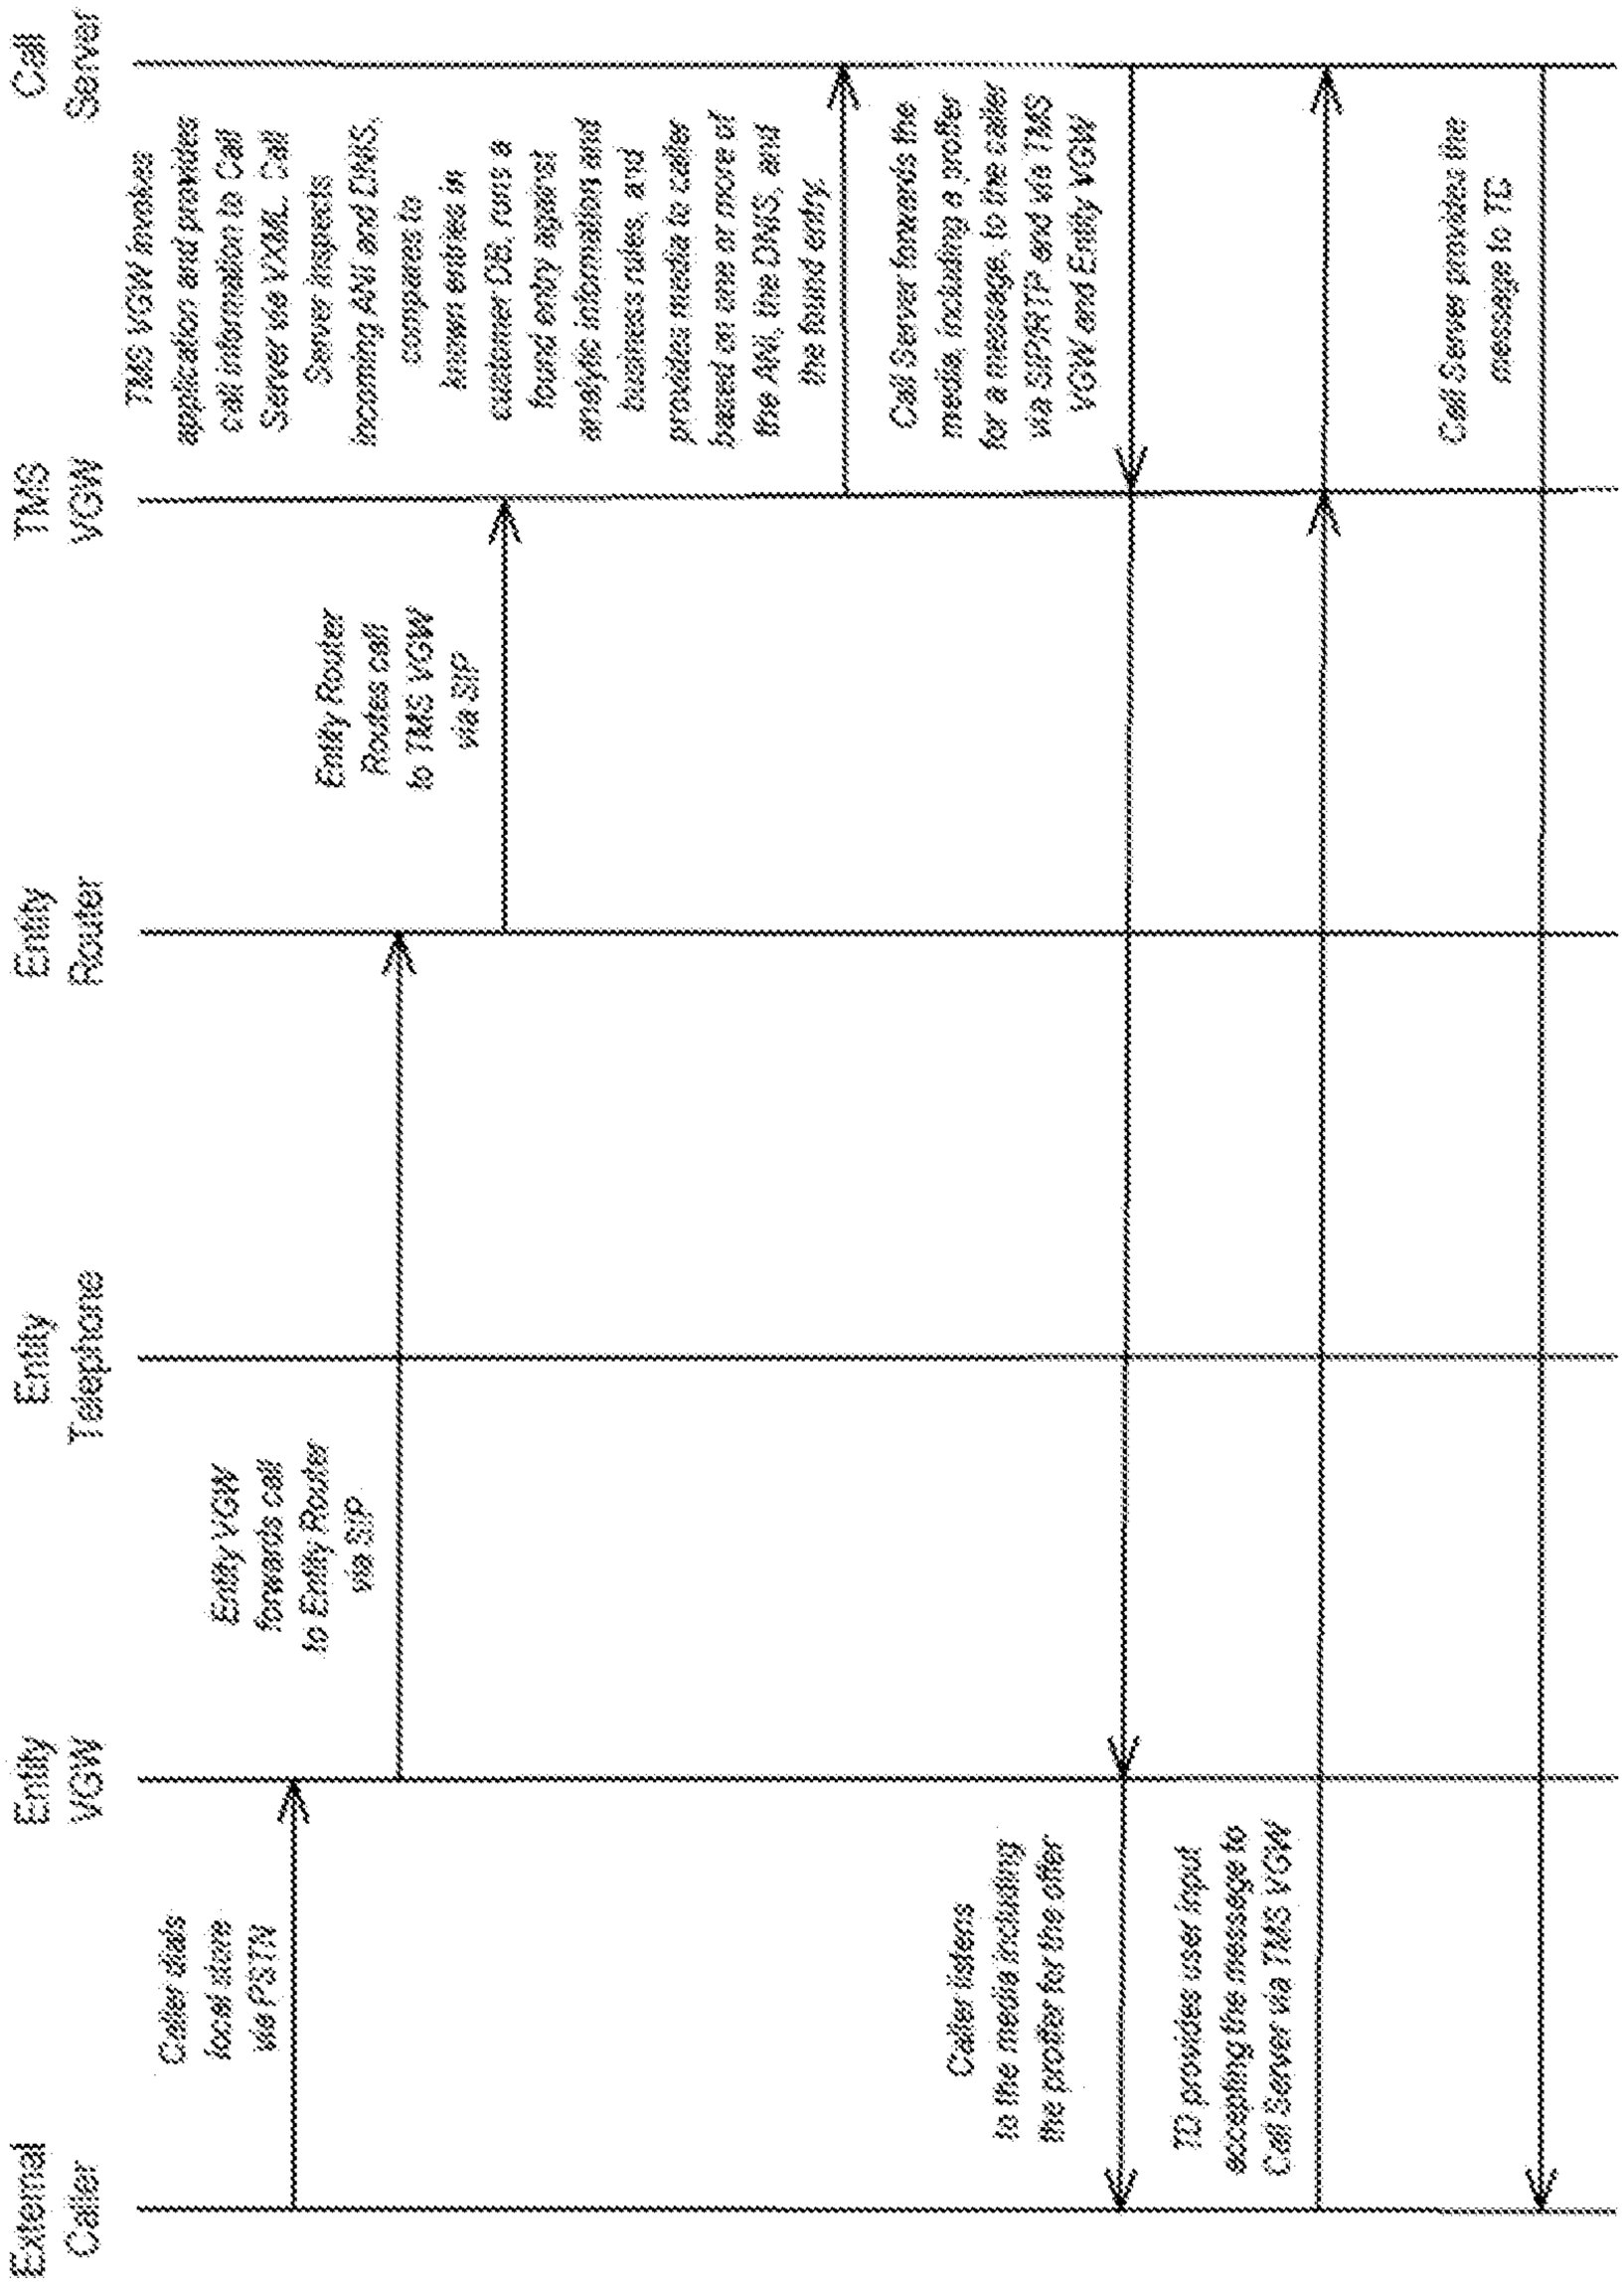
External Caller
Entity VGW
Entity Telephone
Entity Router
TMS VGW
Call Server
Caller dials local store via PSTN
Entity VGW forwards call to Entity Router via SIP
Entity Router Routes call to TMS VGW via SIP
TMS VGW invokes application and provides call information to Call Server via VXML. Call Server requests incoming ANI and DNIS, compares to known entries in customer DB, runs a found entry against analytic information and business rules, and provides media to caller based on one or more of the ANI, the DNIS, and the found entry
Call Server forwards the media, including a prefer for a message, to the caller via SIP/RTP and via TMS VGW and Entity VGW
Call Server provides the message to TG
Caller listens to the media including the prefer for the offer
TG provides user input accepting the message to Call Server via TMS VGW

FIG. 8A

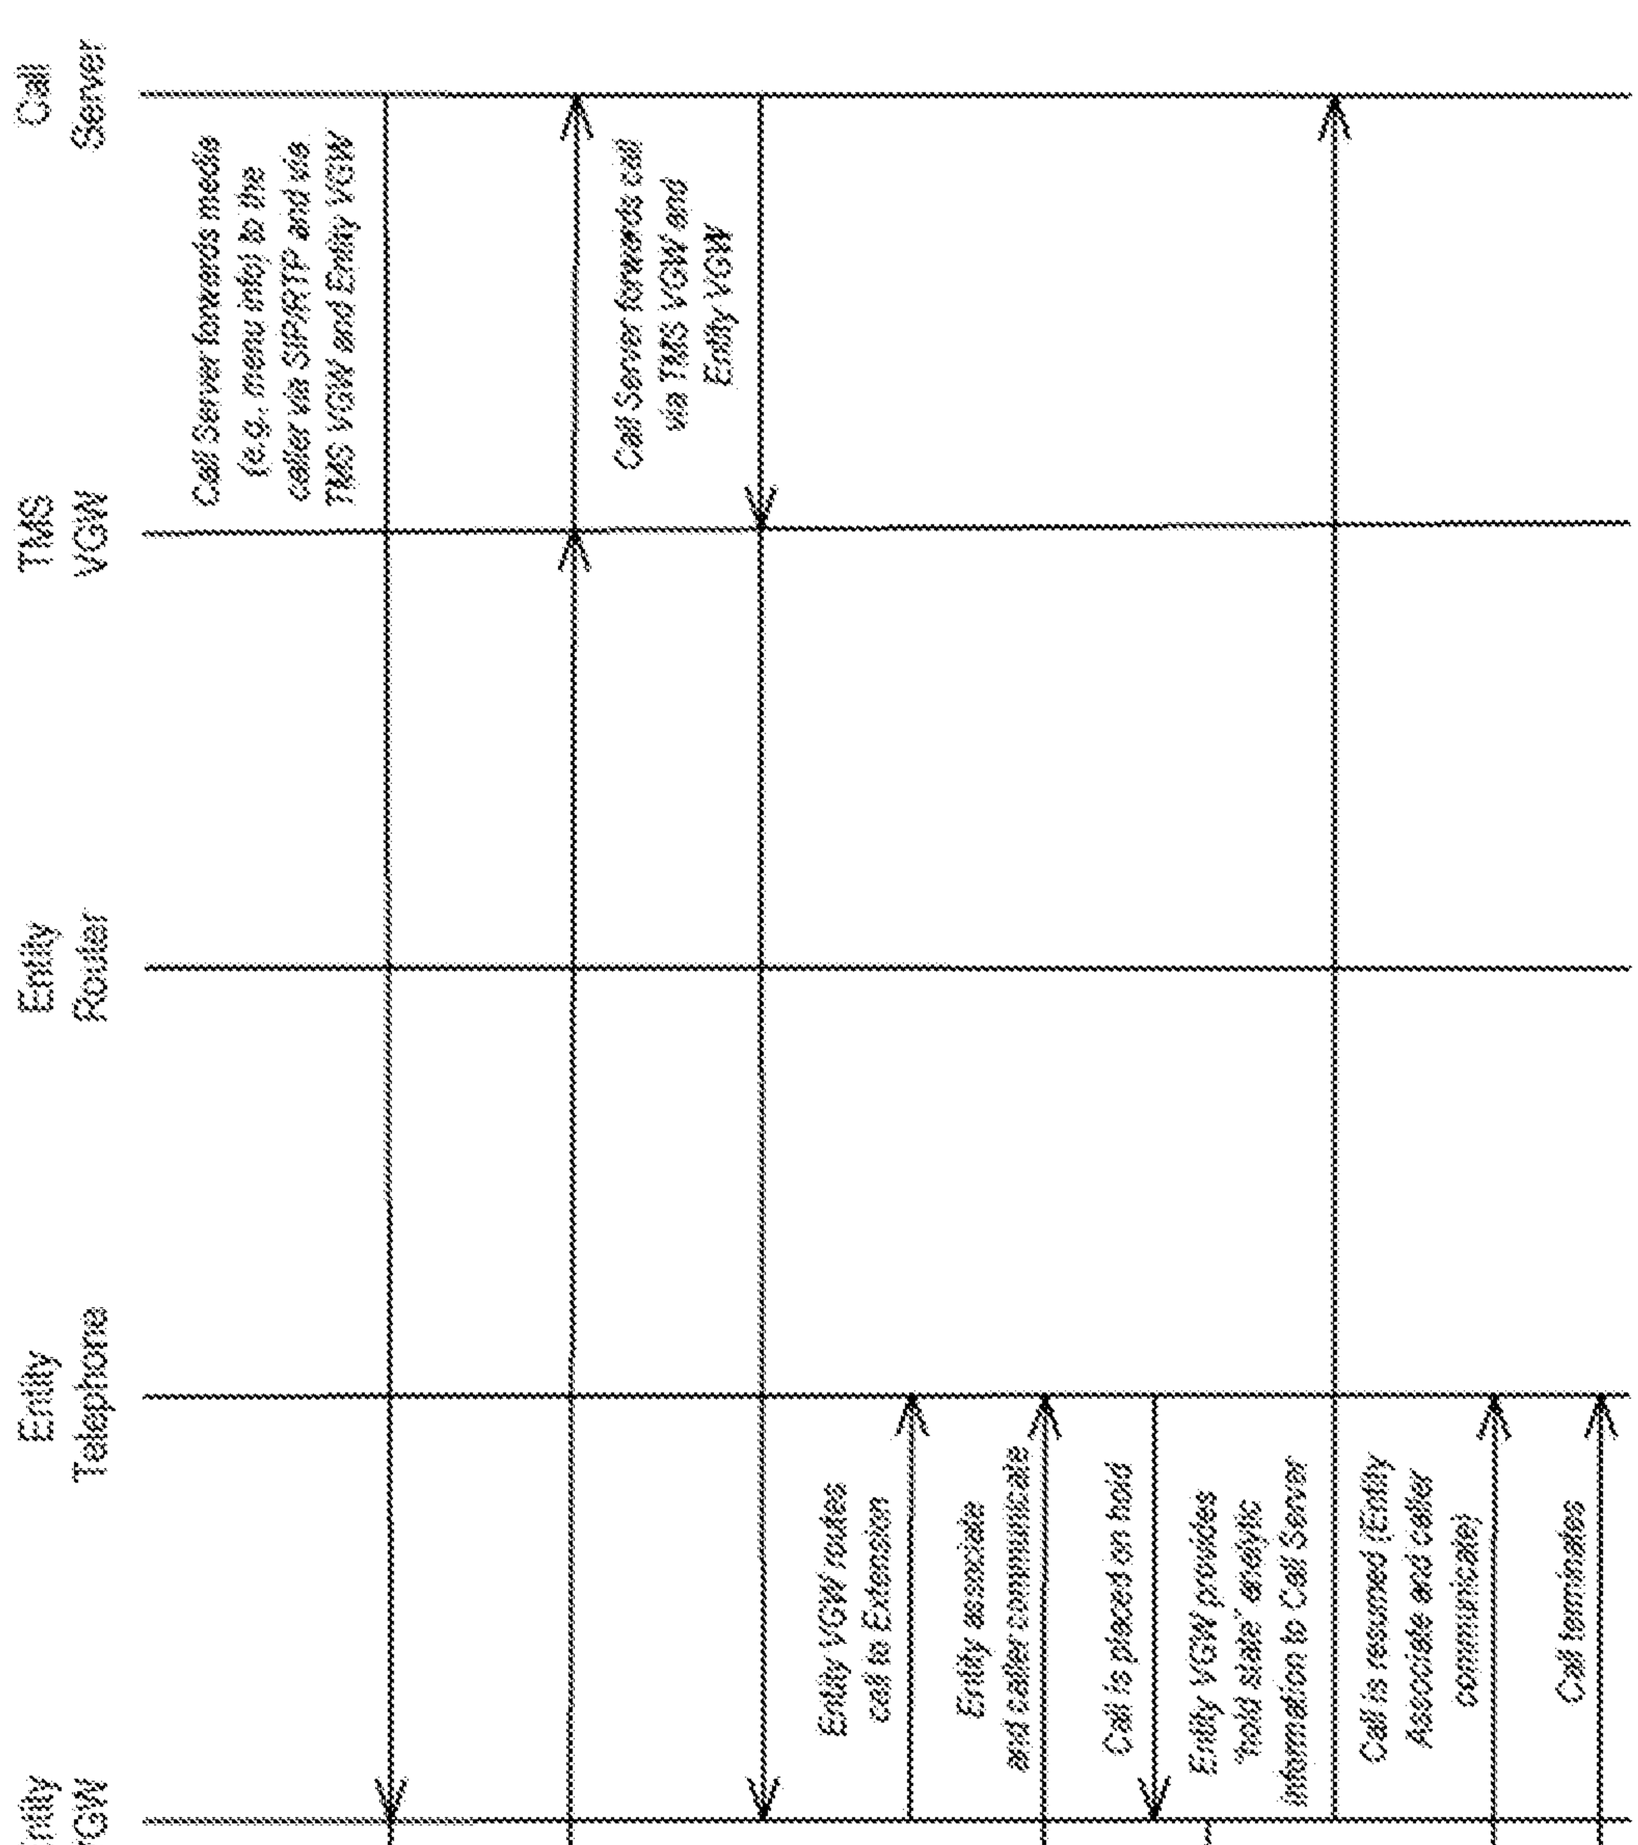
External Caller
Entity VGW
Entity Telephone
Entity Router
TMS VGW
Call Server
Caller listens to the media
ID provides user input to Call Server via TMS VGW
Entity VGW routes call to Extension
Entity associate and caller communicate
Call is placed on hold
Entity VGW provides "hold state" analytic information to Call Server
Entity VGW provides appropriate "On Hold" media to ID
Call is resumed (Entity Associate and caller communicate)
Call terminates
Call Server forwards call via TMS VGW and Entity VGW
Call Server forwards media (e.g. menu info) to the caller via SIP/RTP and via TMS VGW and Entity VGW

FIG. 8B

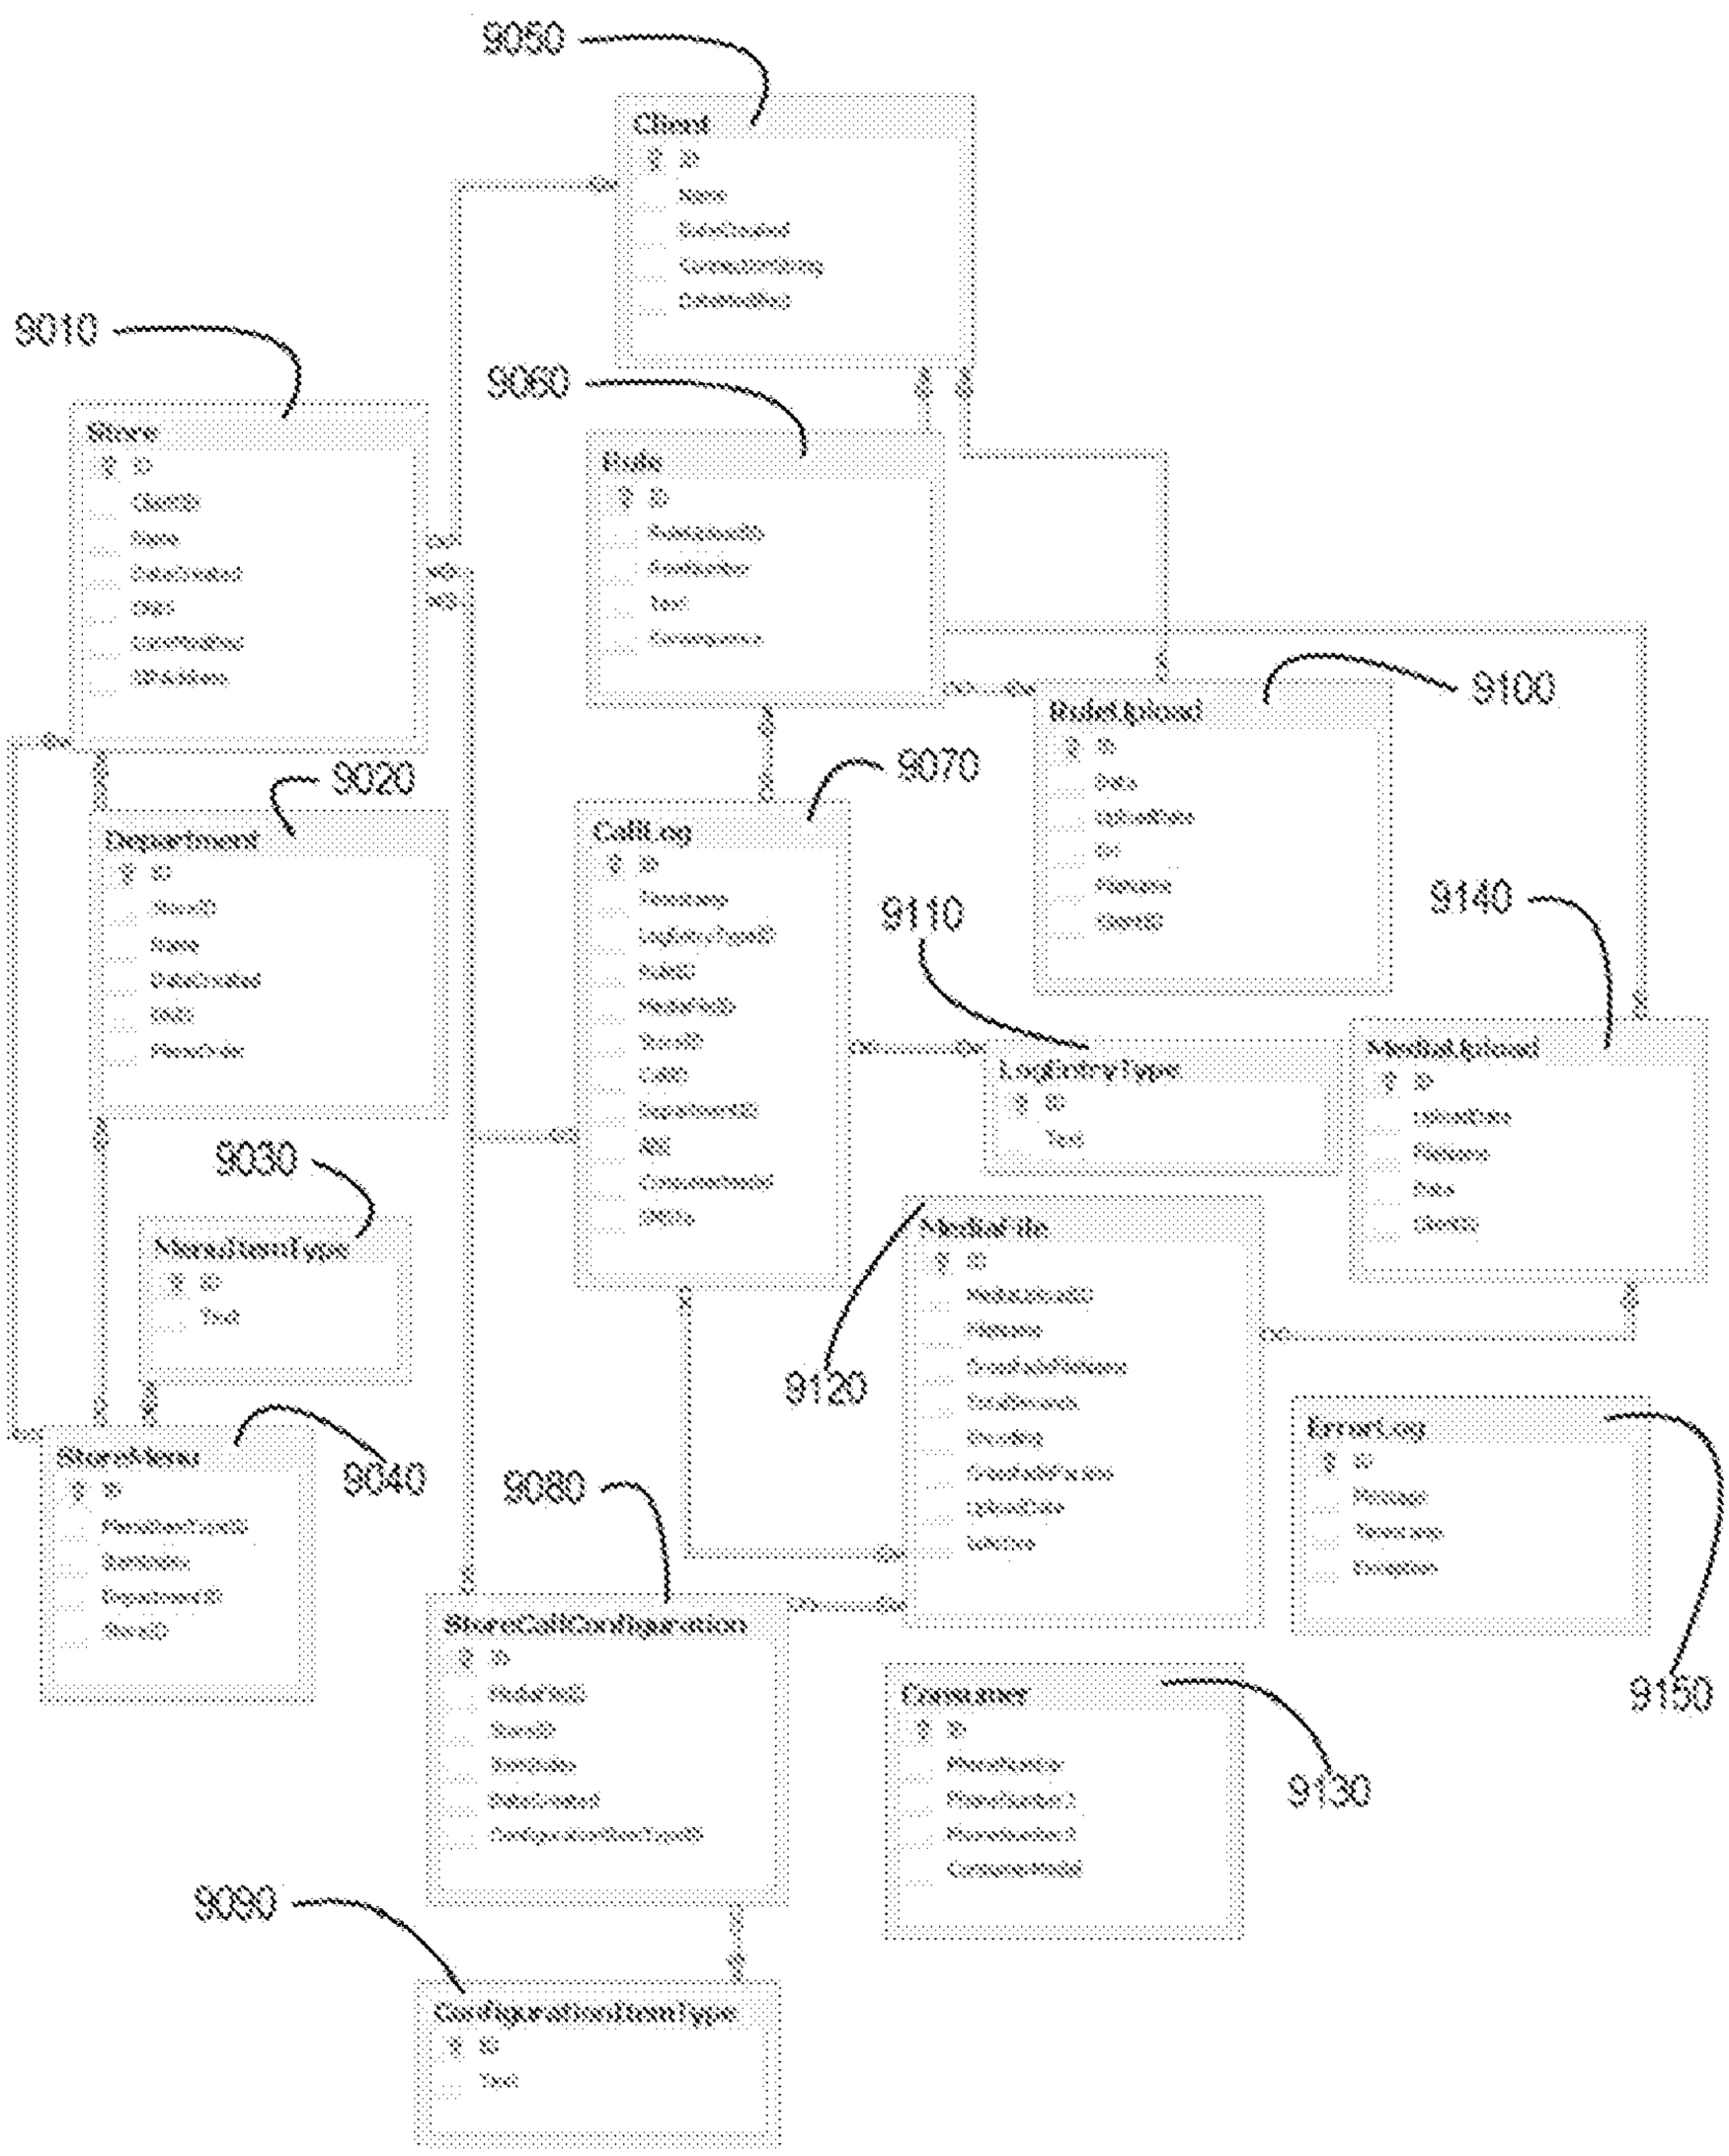
9050
9010
9060
9100
9070
9020
9110
9140
9030
9120
9040
9080
9150
9130
9090

FIG. 9

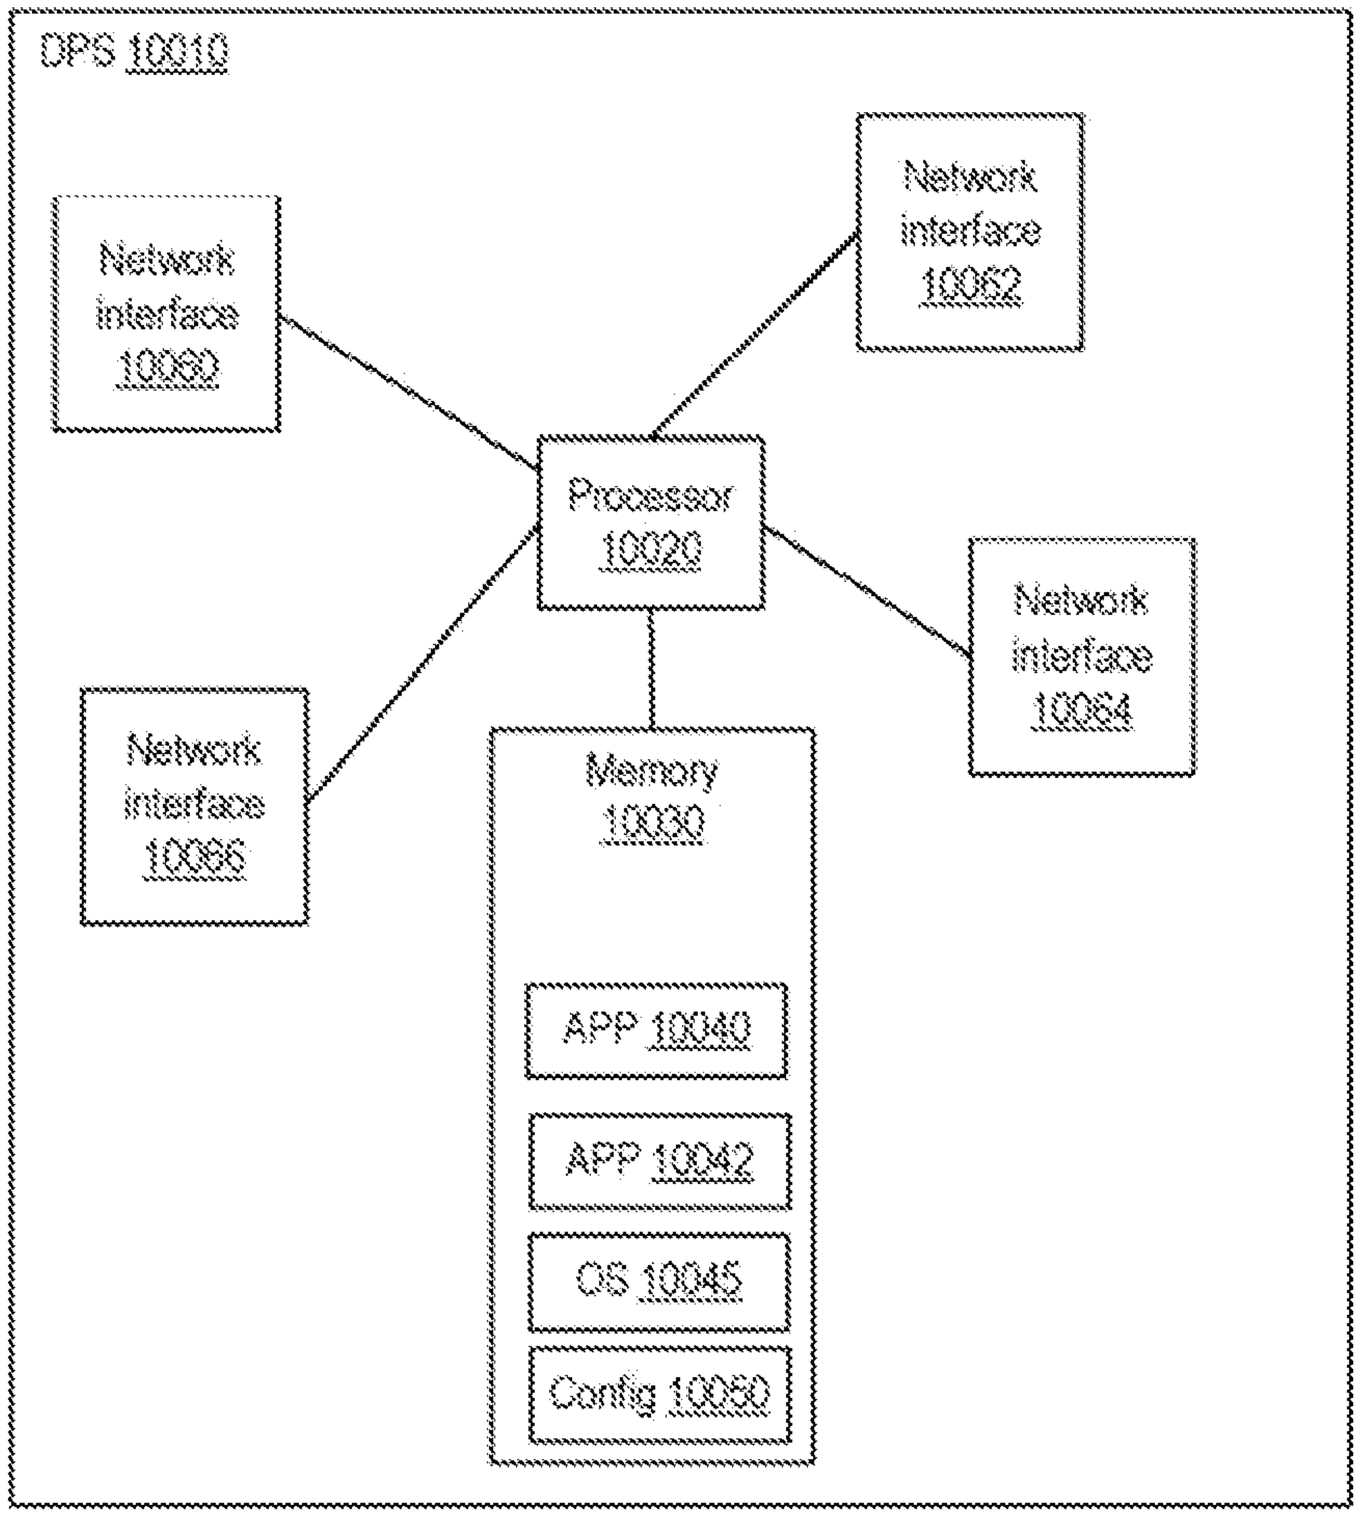
OPS 10010
Network interface 10060
Network interface 10062
Processor 10020
Network interface 10064
Network interface 10066
Memory 10030
APP 10040
APP 10042
OS 10045
Config 10050

FIG. 10

Figure 1P

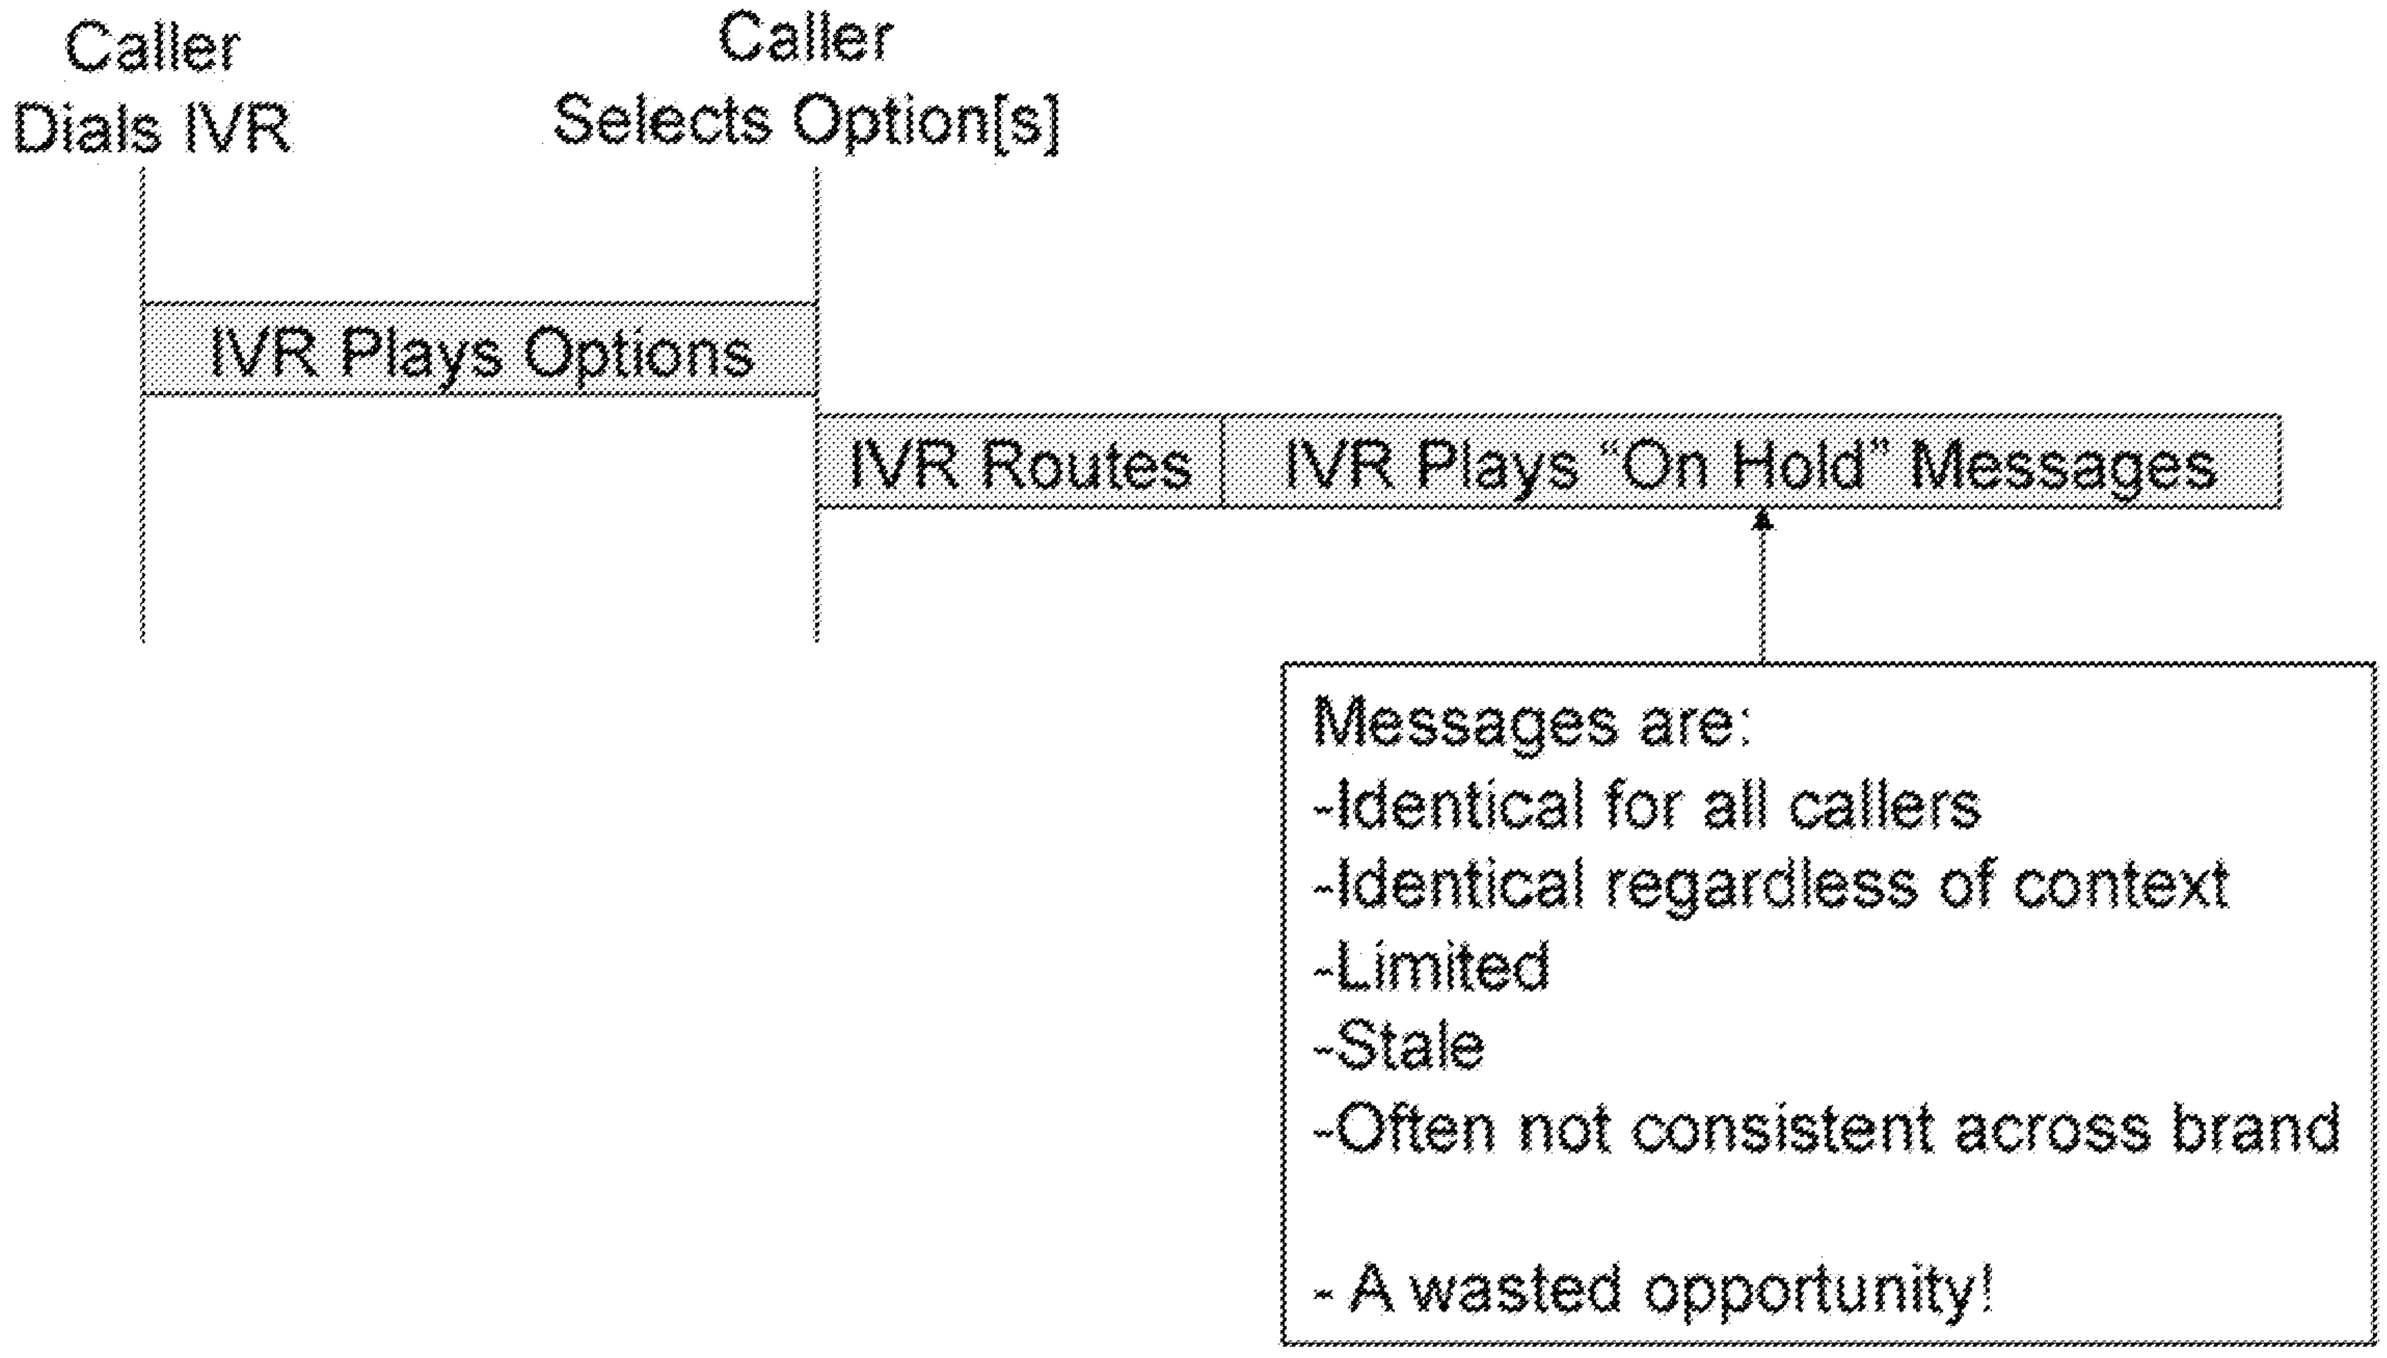
Caller
Dials IVR
Caller
Selects Option[s]
IVR Plays Options
IVR Routes
IVR Plays "On Hold" Messages
Messages are:
-Identical for all callers
-Identical regardless of context
-Limited
-Stale
-Often not consistent across brand
- A wasted opportunity!

Figure 2P

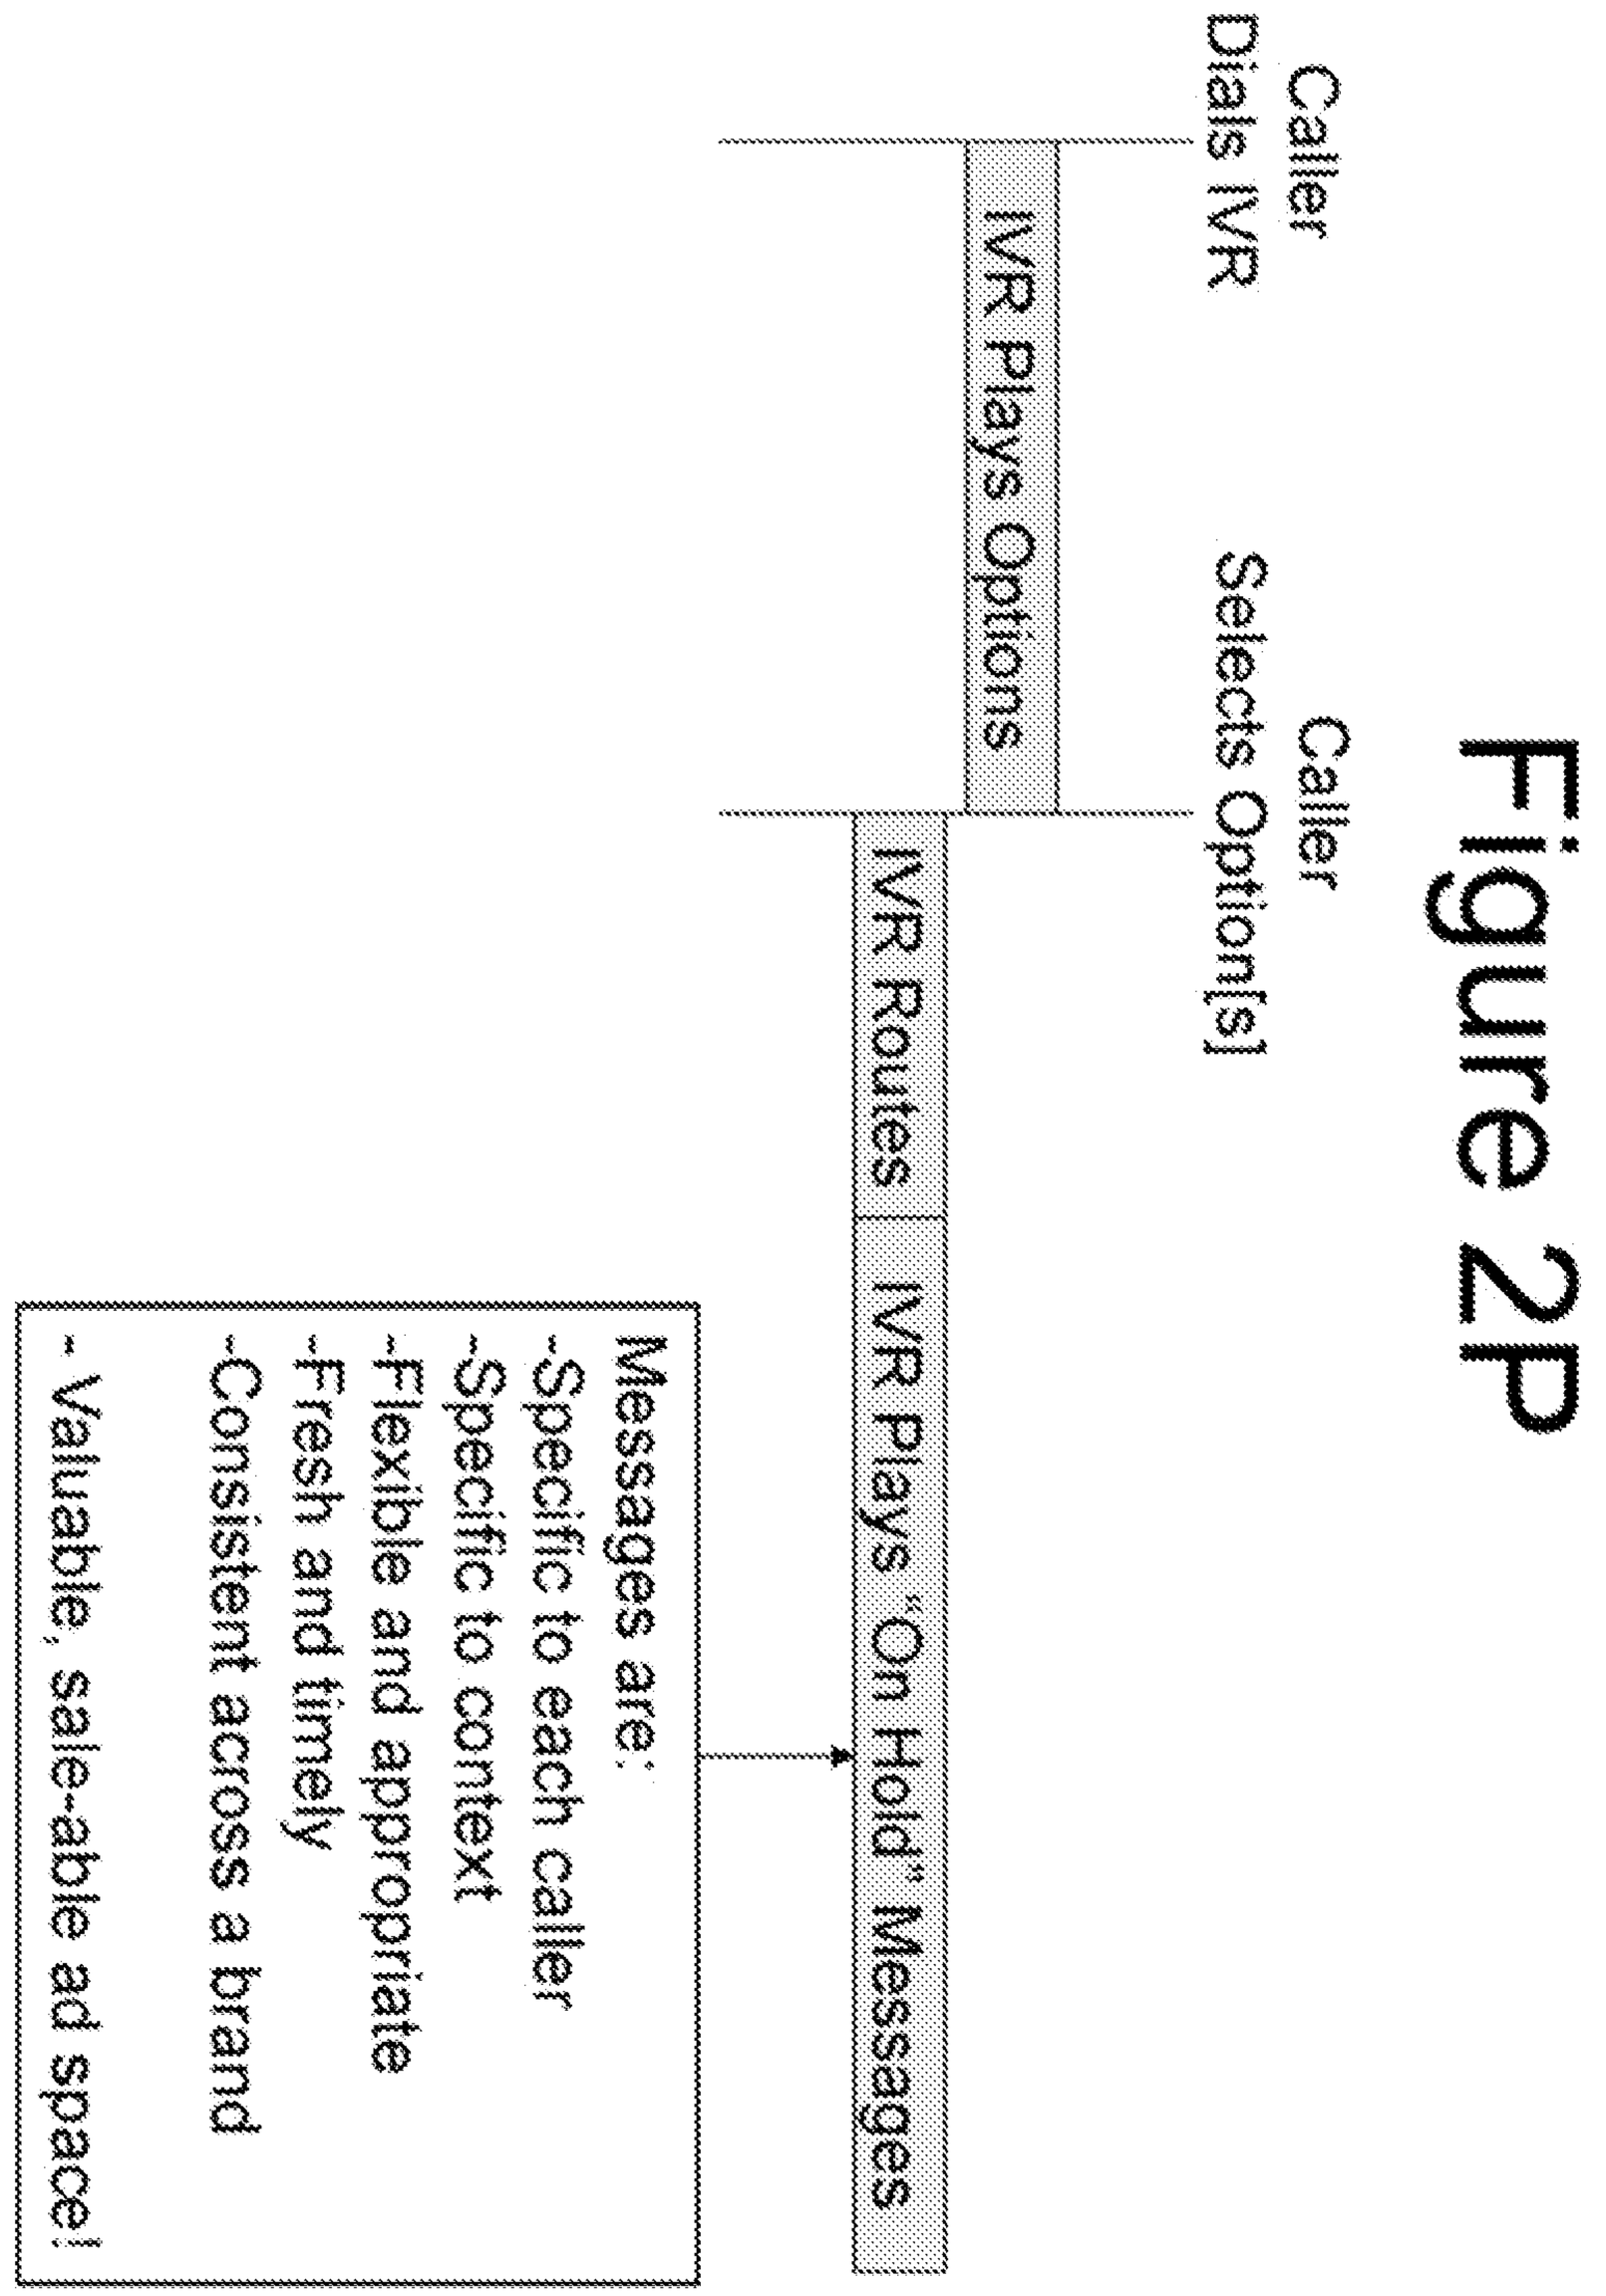
Caller
Selects Option[s]
IVR Plays Options
IVR Routes
IVR Plays "On Hold" Messages
Messages are:
-Specific to each caller
-Specific to context
-Flexible and appropriate
-Fresh and timely
-Consistent across a brand
- Valuable, sale-able ad spaces

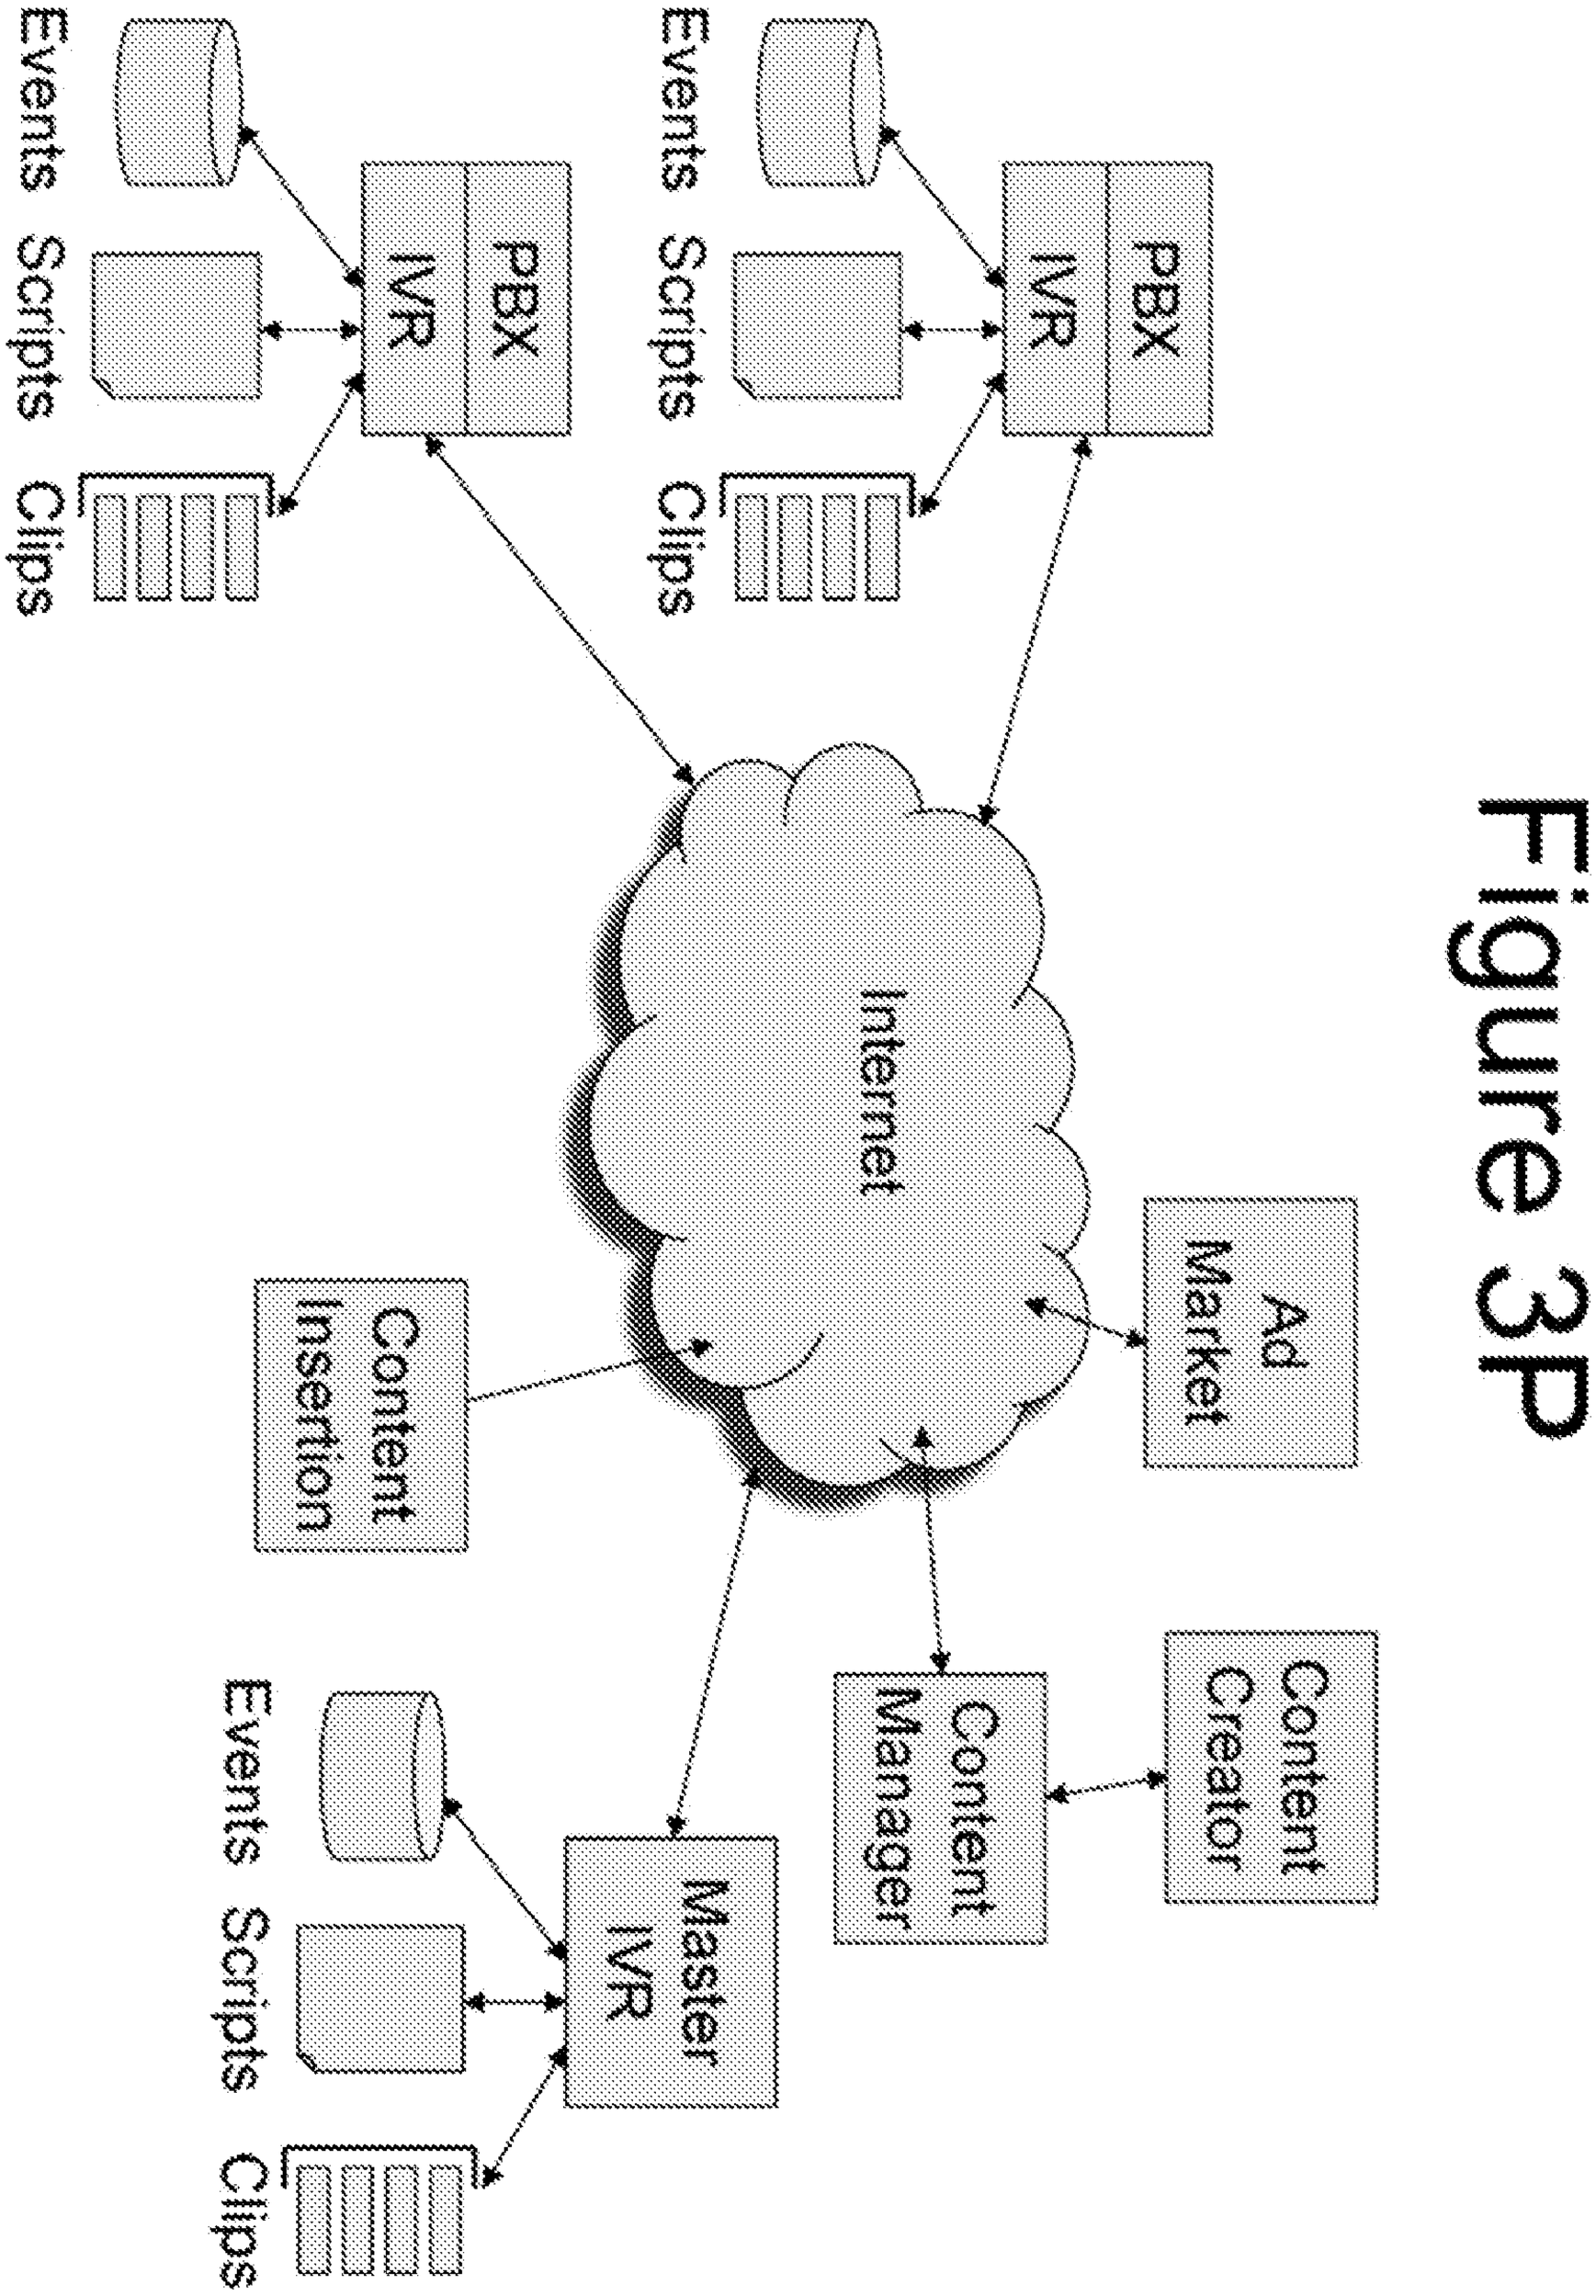
Events
Scripts
Clips
PBX
IVR
Events
Scripts
Clips
PBX
IVR
Internet
Ad Market
Content Insertion
Content Manager
Content Creator
Master IVR
Events
Scripts
Clips

Figure 3P

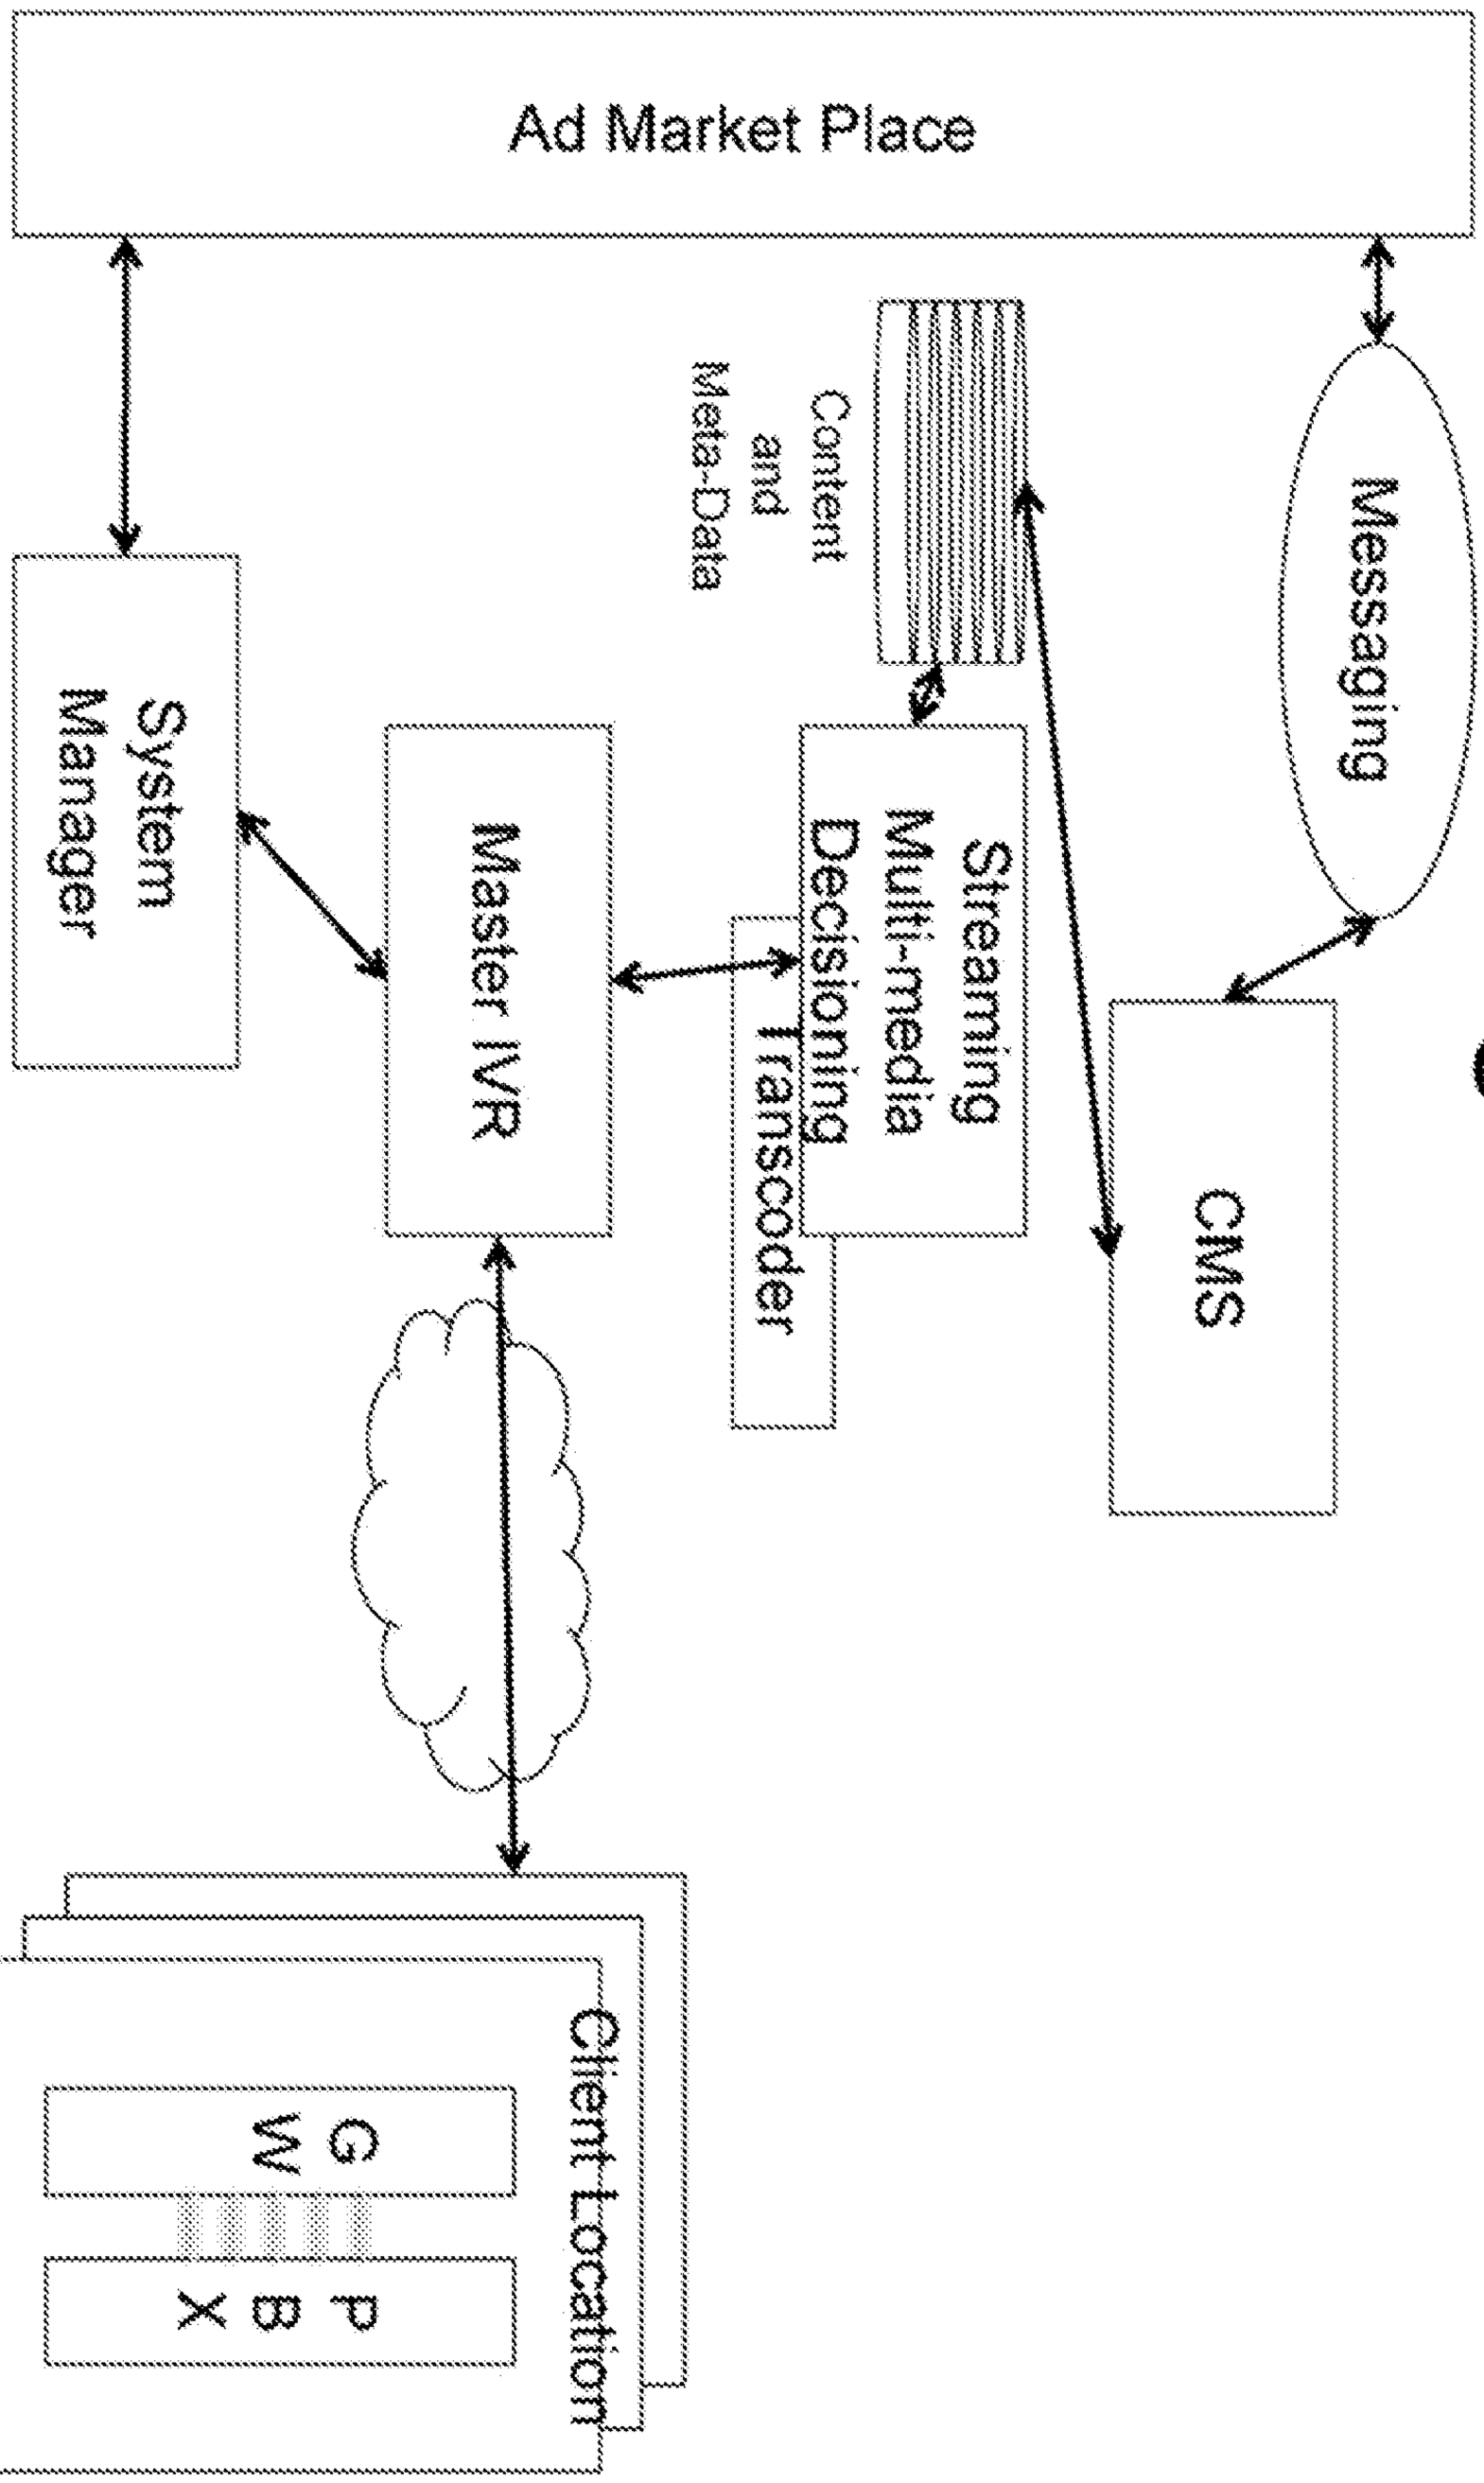
Ad Market Place
Messaging
CMS
Content and Meta-Data
Streaming Multi-media Decisioning
Transcoder
Master IVR
System Manager
Client Location
PDA
GW

Figure 4P

Figure A1

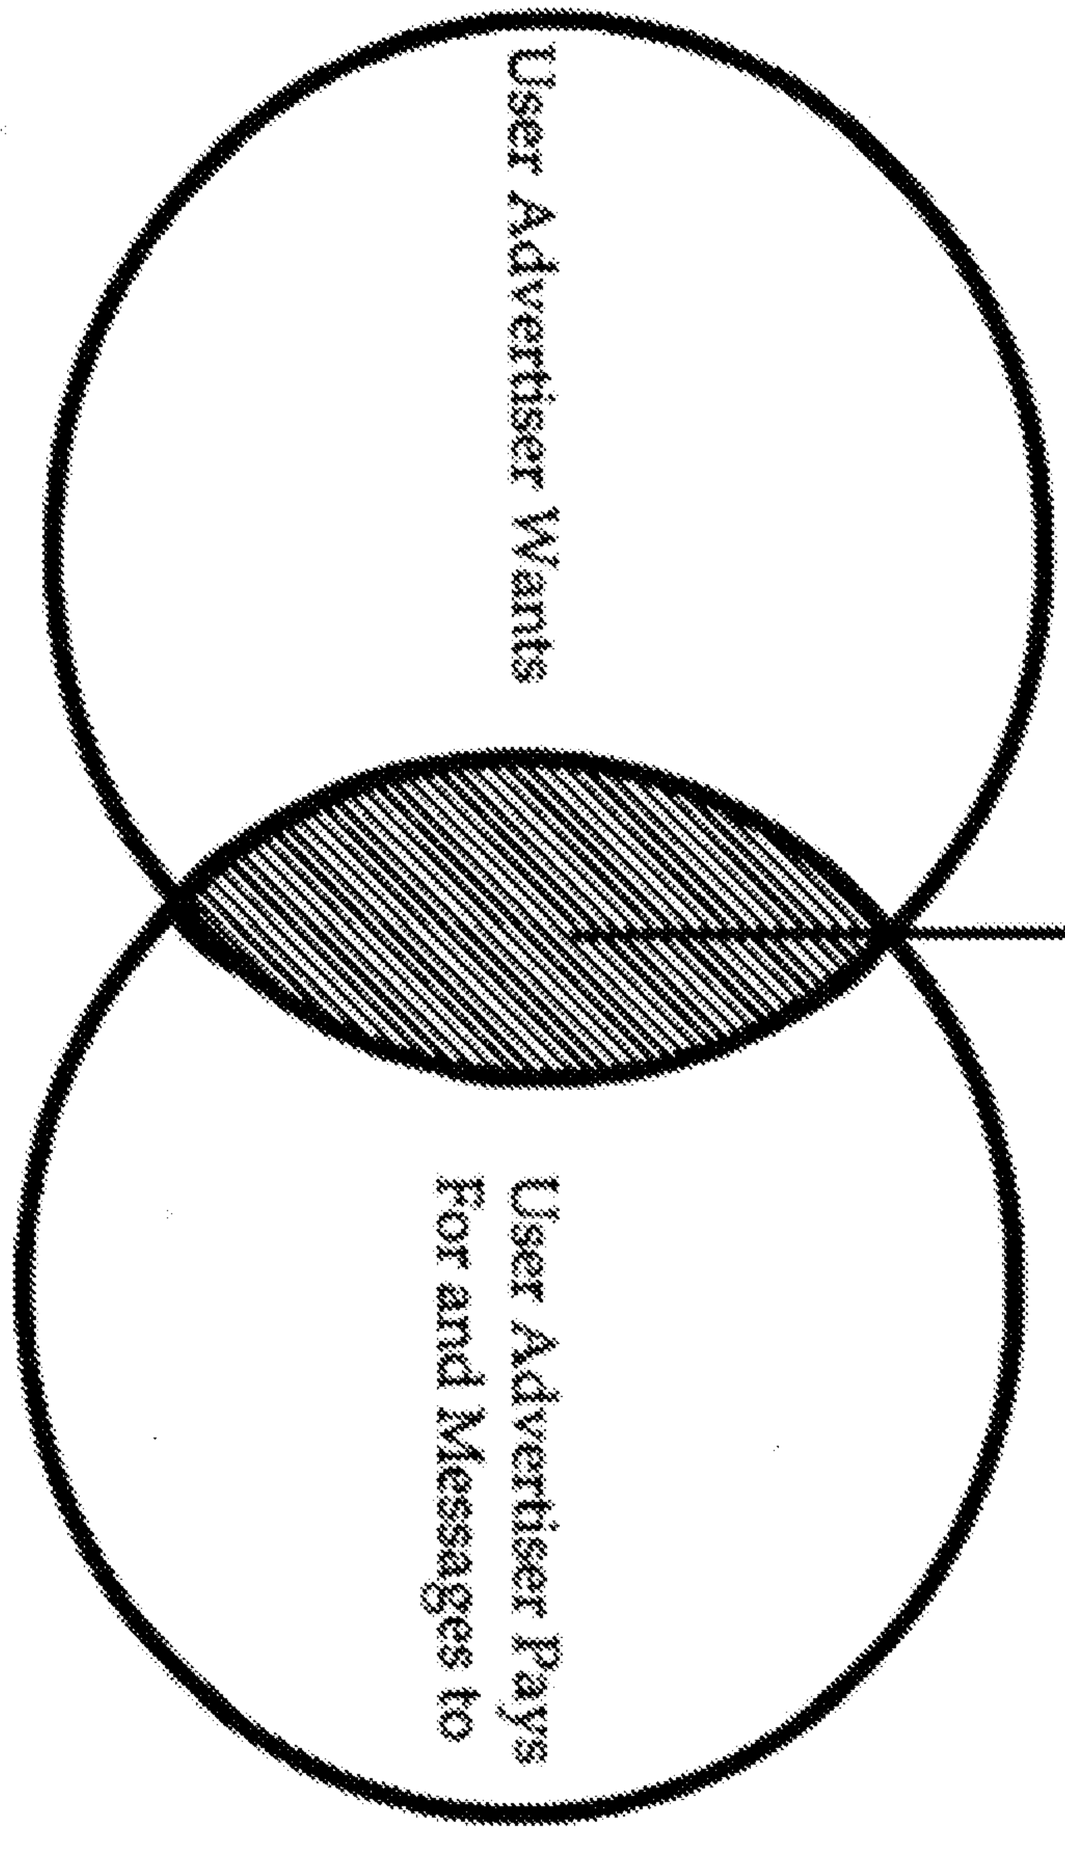
User Advertiser Wants
User Advertiser Wants and Messages To
User Advertiser Pays For and Messages to

Figure A2

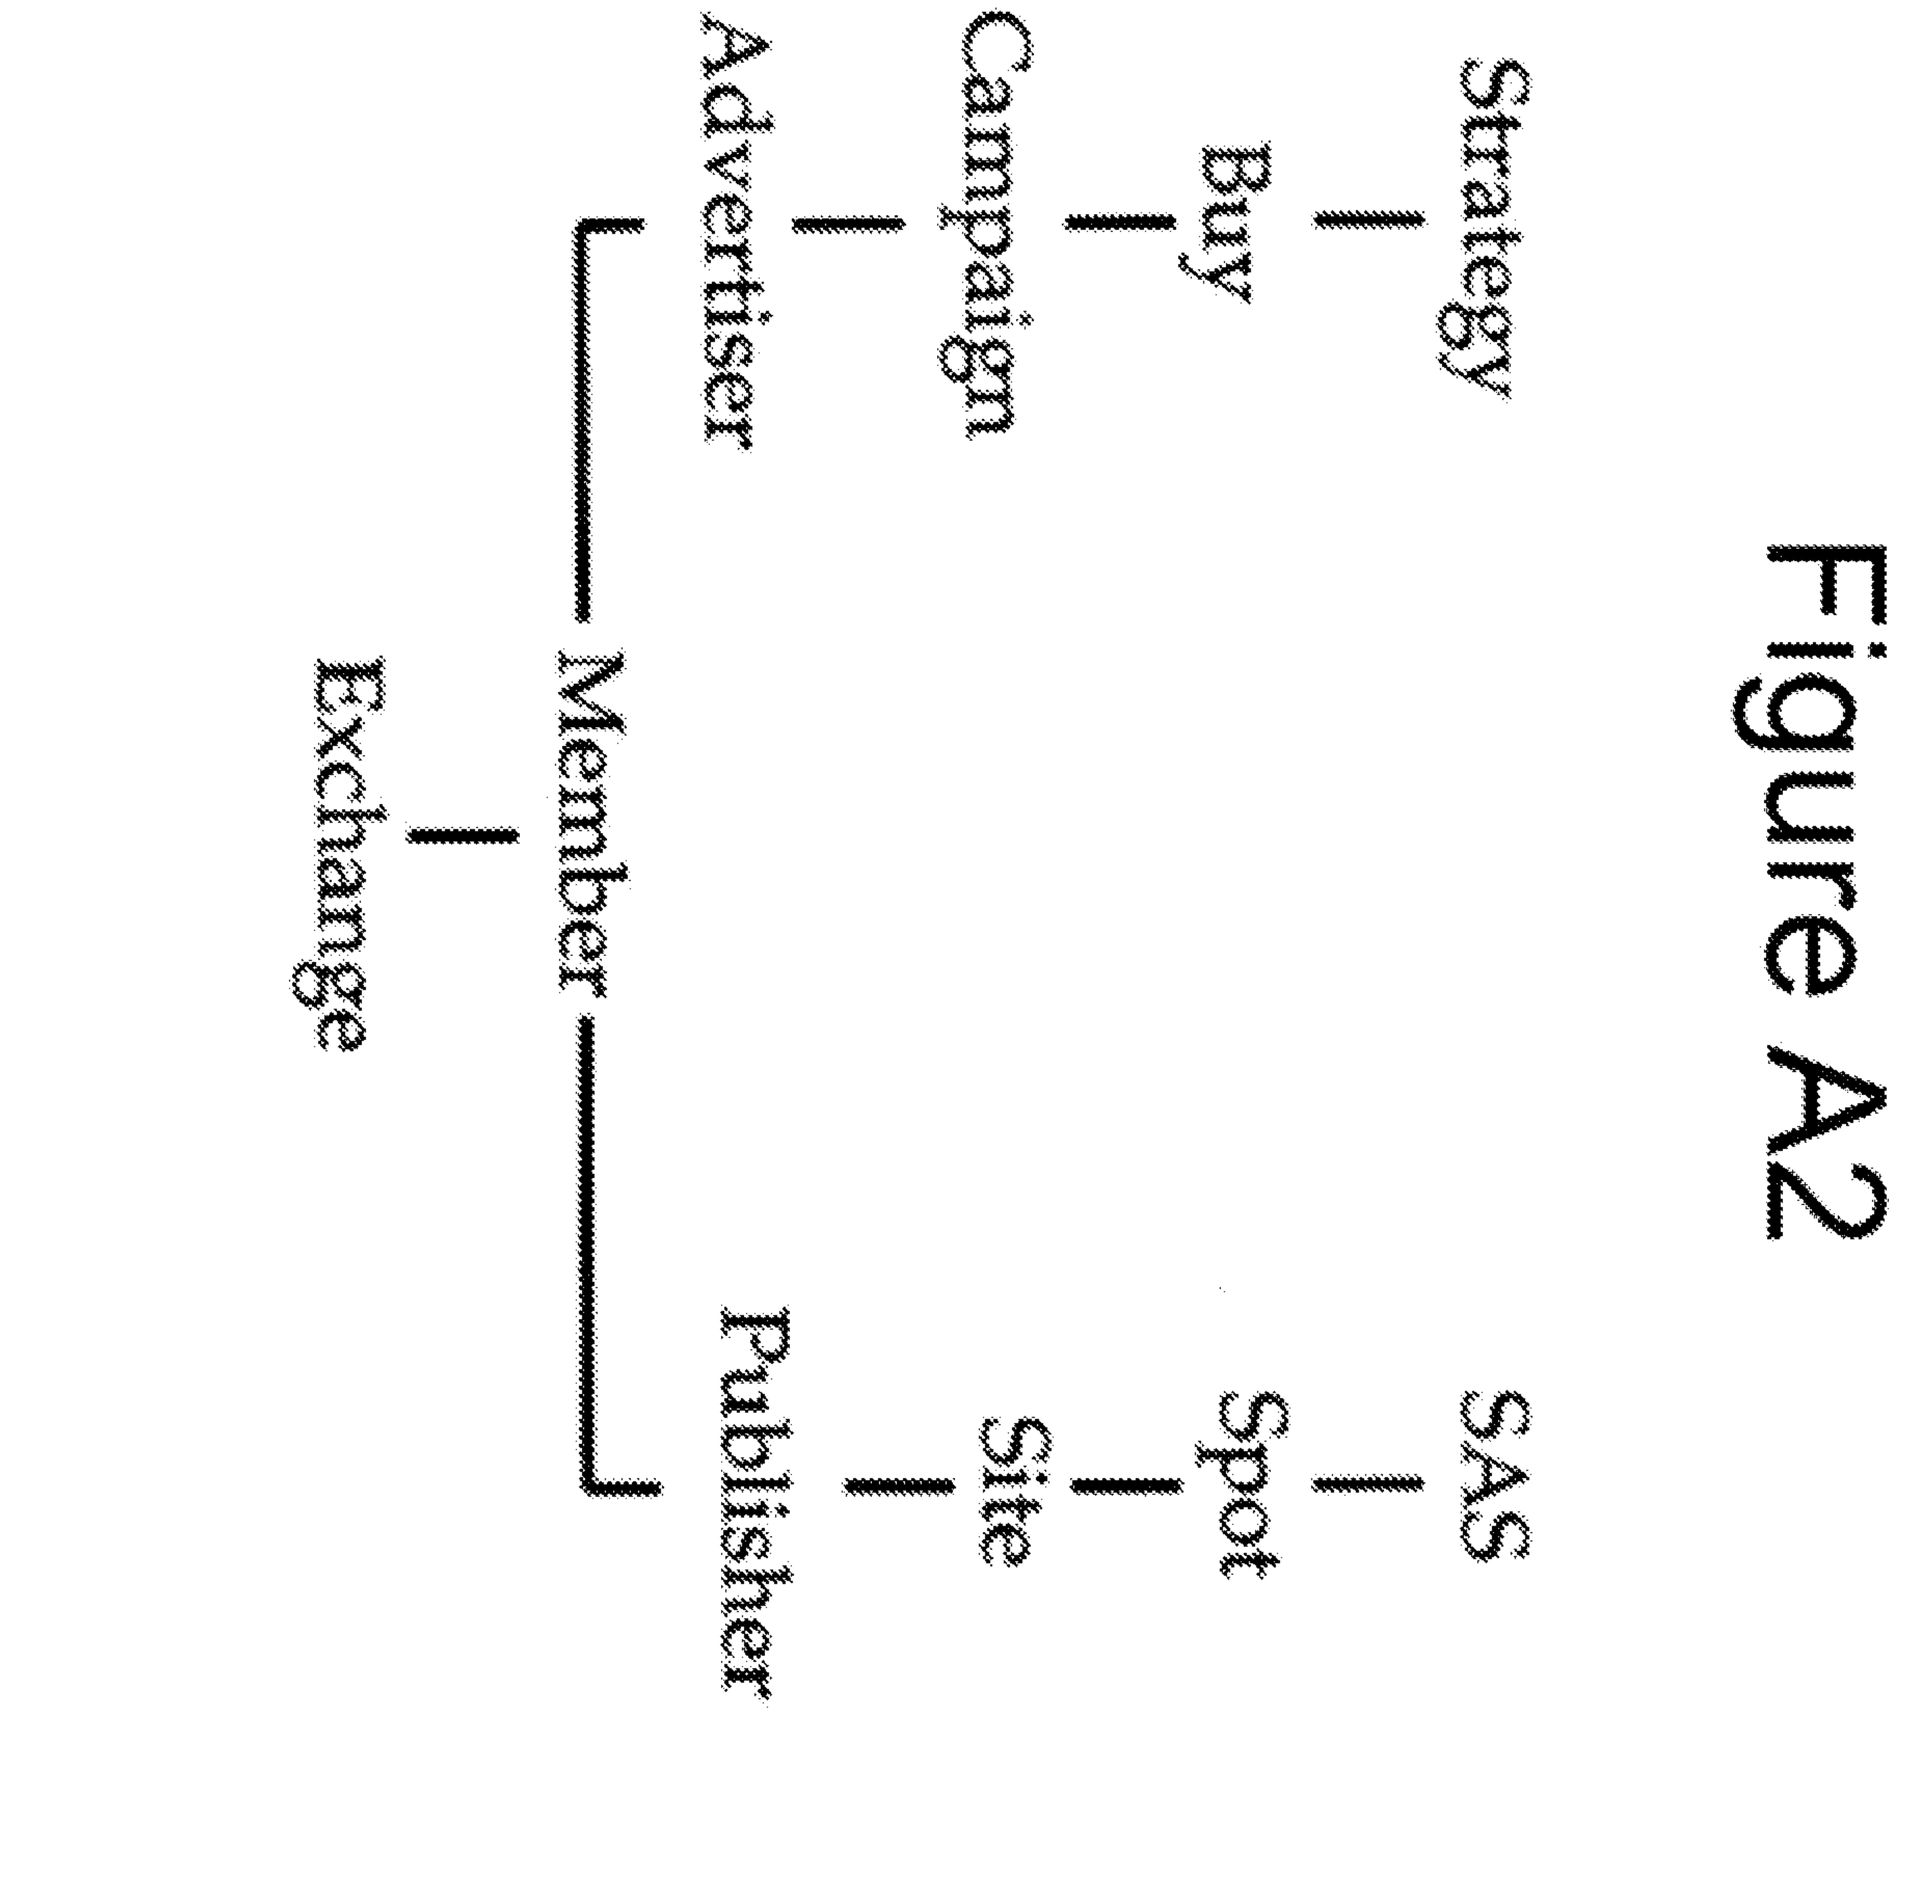
Strategy
Buy
Campaign
Advertiser
Member
Exchange
SAS
Spot
Site
Publisher

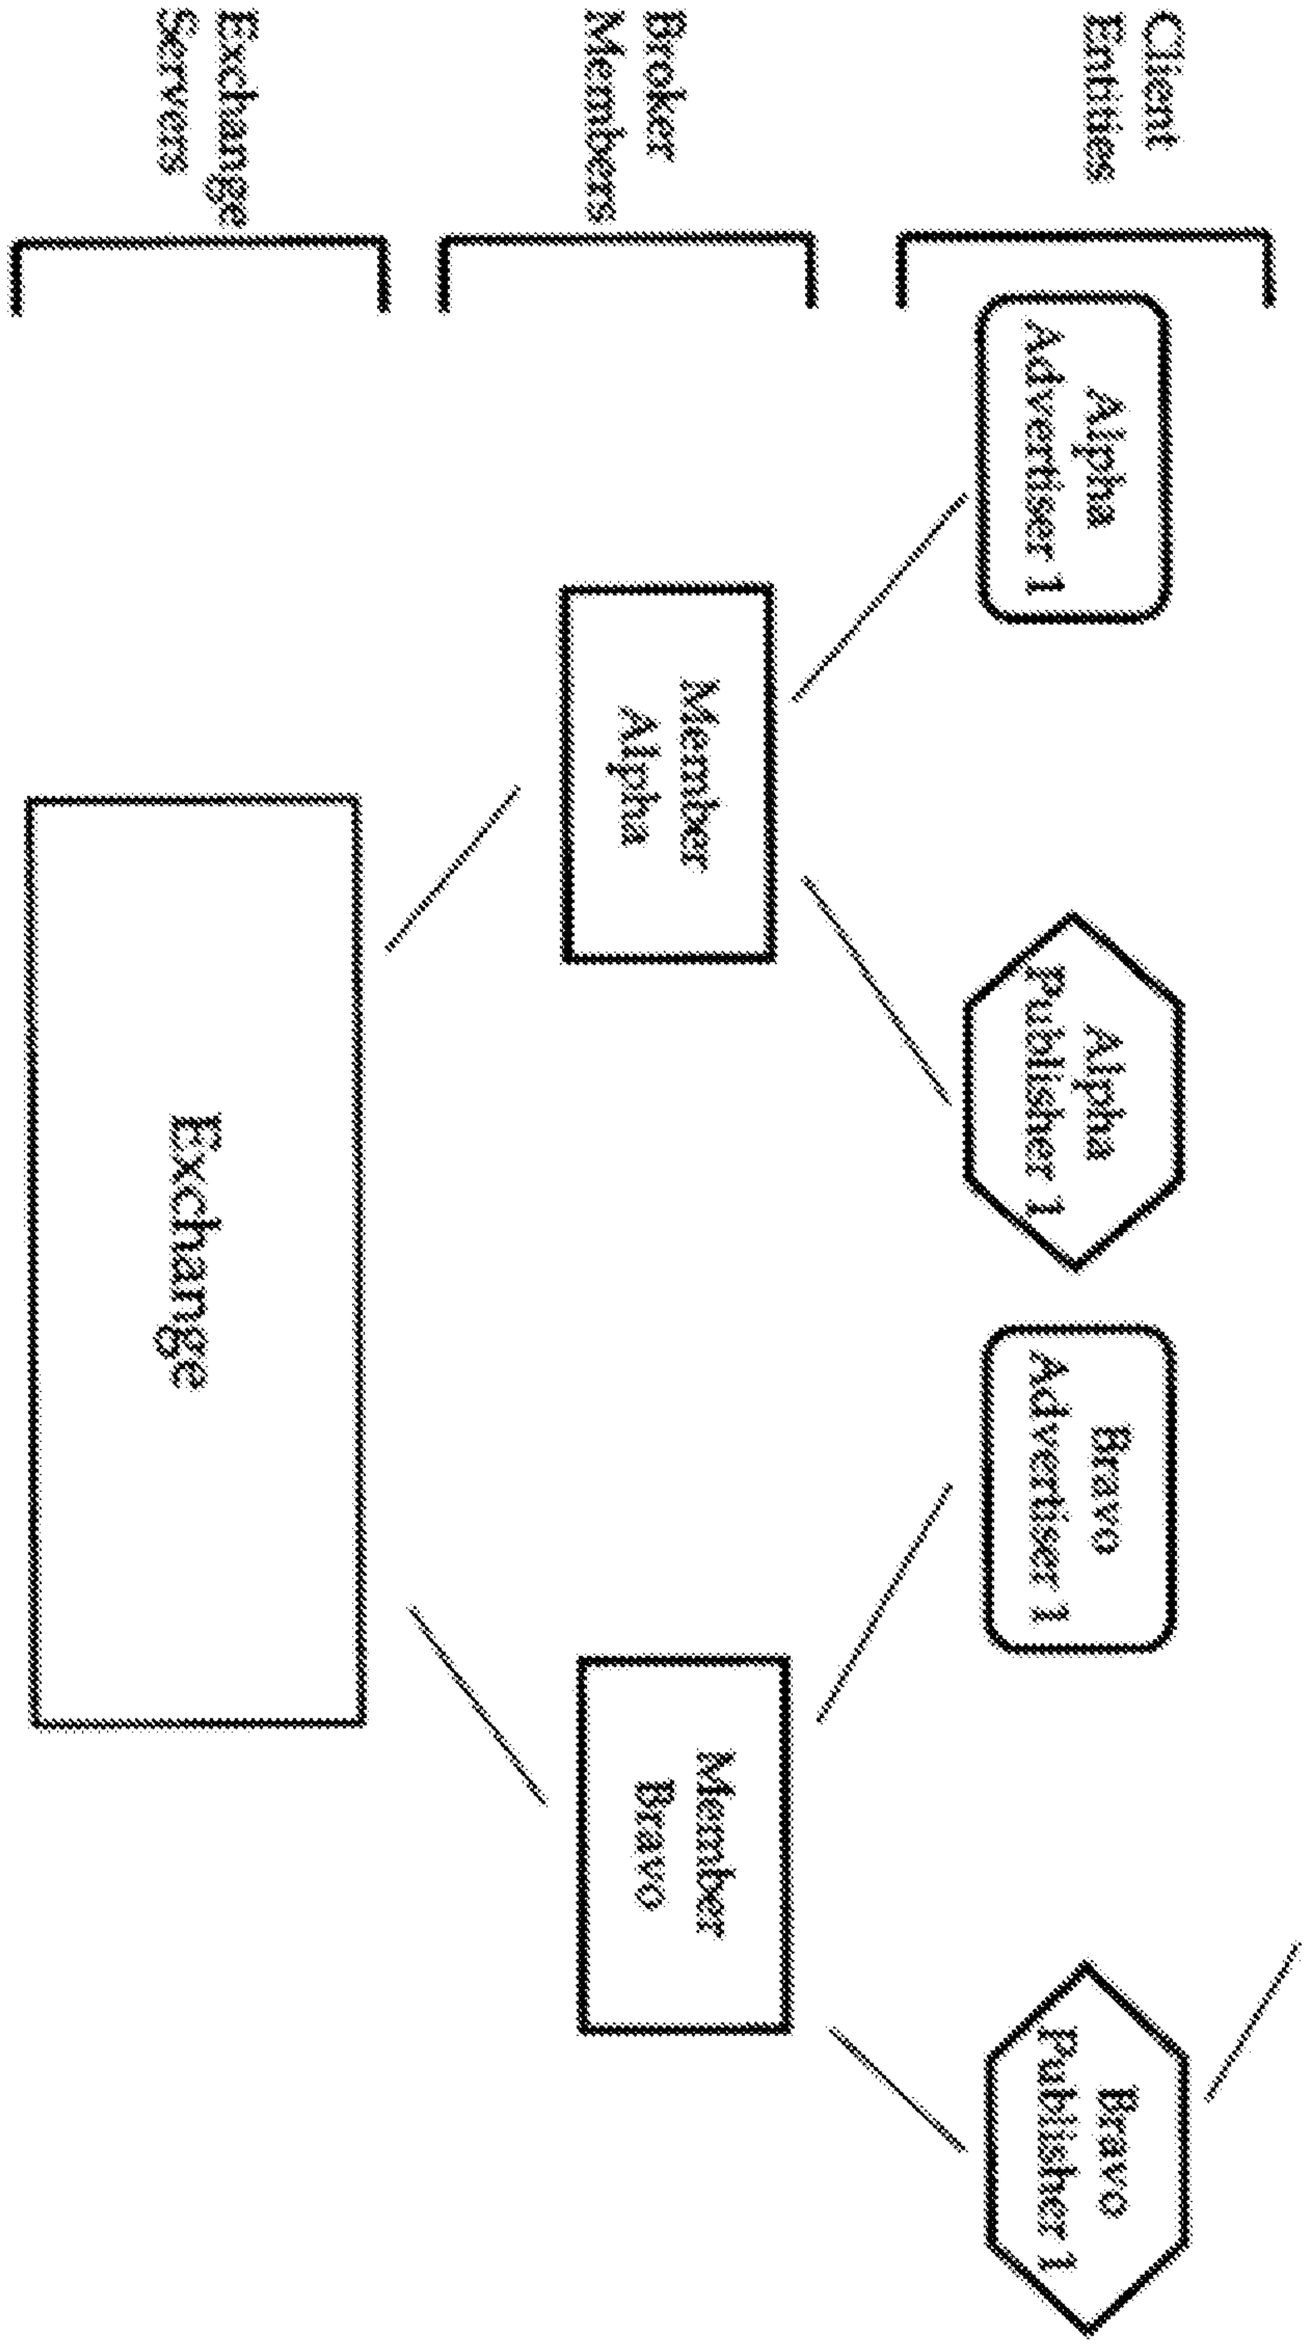
Exchange Servers
Broker Members
Client Entities
Alpha Advertiser 1
Member Alpha
Alpha Publisher 1
Exchange Image
Bravo Advertiser 1
Member Bravo
Bravo Publisher 1
End User/Site Visitor

Figure A3

Figure A4

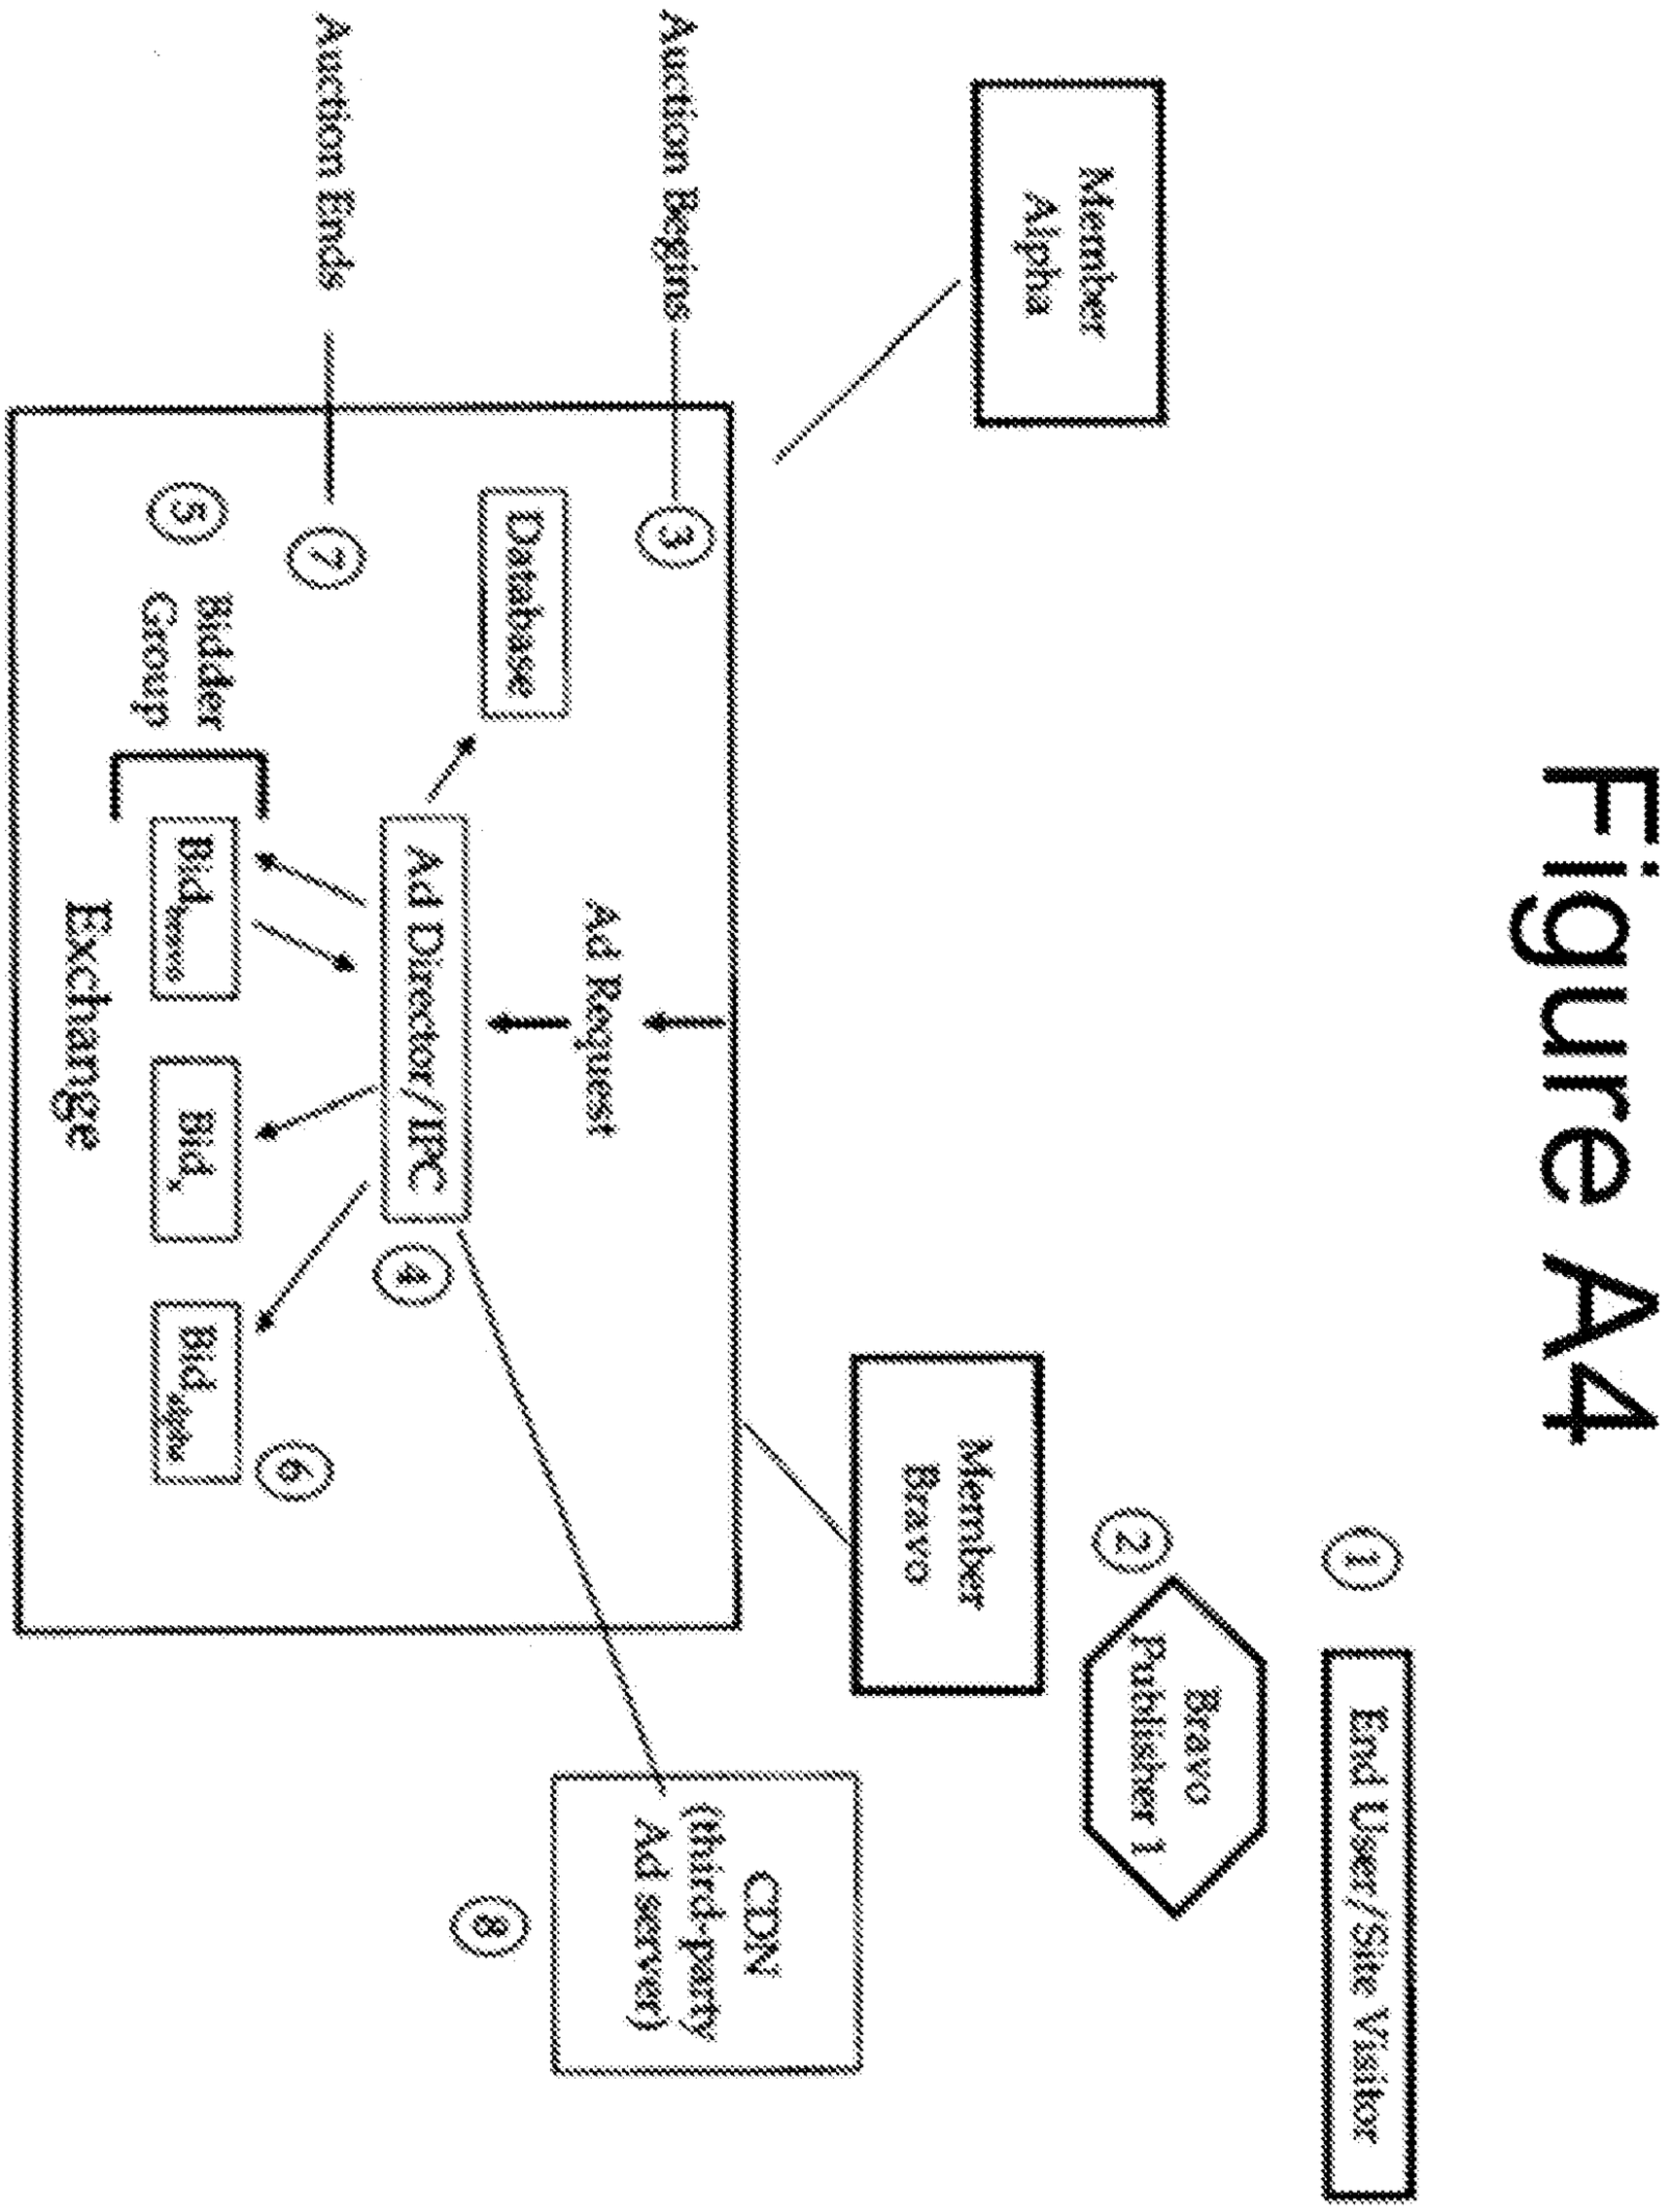
End User/Site Visitor
1
Bravo Publisher 1
2
Member Bravo
CDN (third-party Ad server)
8
Member Alpha
Auction Ends
Auction Begins
Database
Ad Request
Ad Director/IFC
Bidder Group
bid
bid
bid
bid
Exchange
3
4
5
6
7

Figure A5

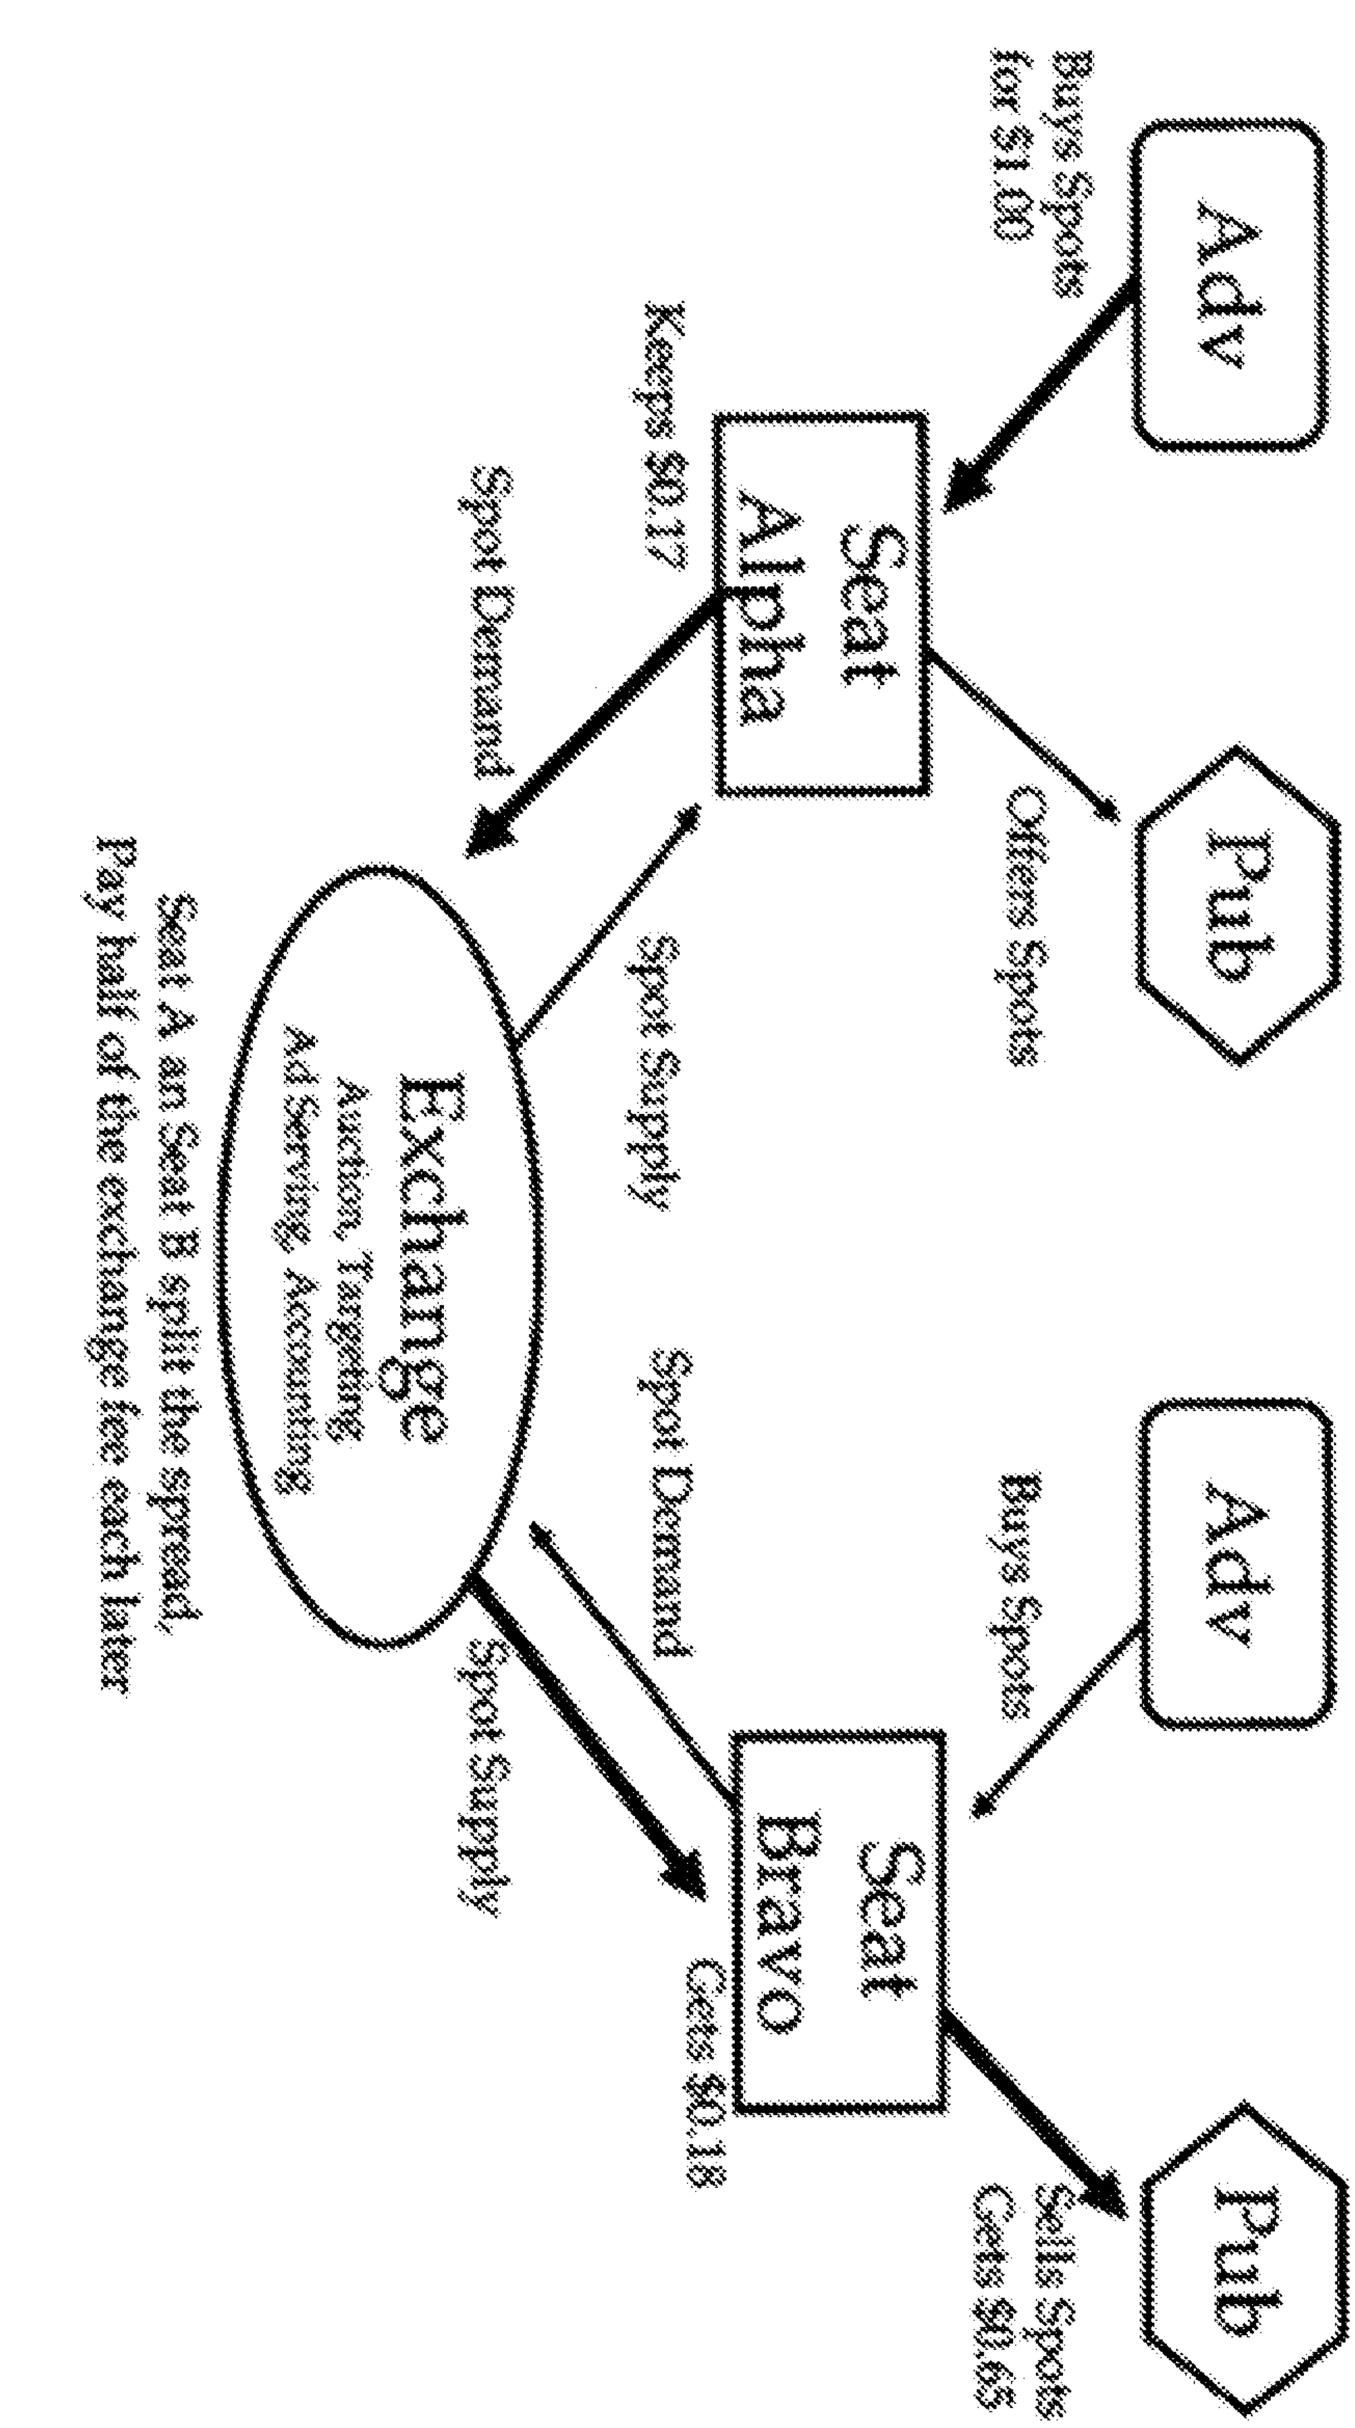
Adv
buys spots for $100
Seat Alpha
Keeps $0.17
Spot Demand
Offers Spots
Pub
Exchange
Auction, Trading
Ad Serving, Accounting
Spot Supply
Spot Supply
Spot Demand
Seat A an Seat B split the spread,
Pay half of the exchange fee each later
Adv
buys spots
Seat Bravo
Gets $0.18
Sells spots
Gets $0.65
Pub

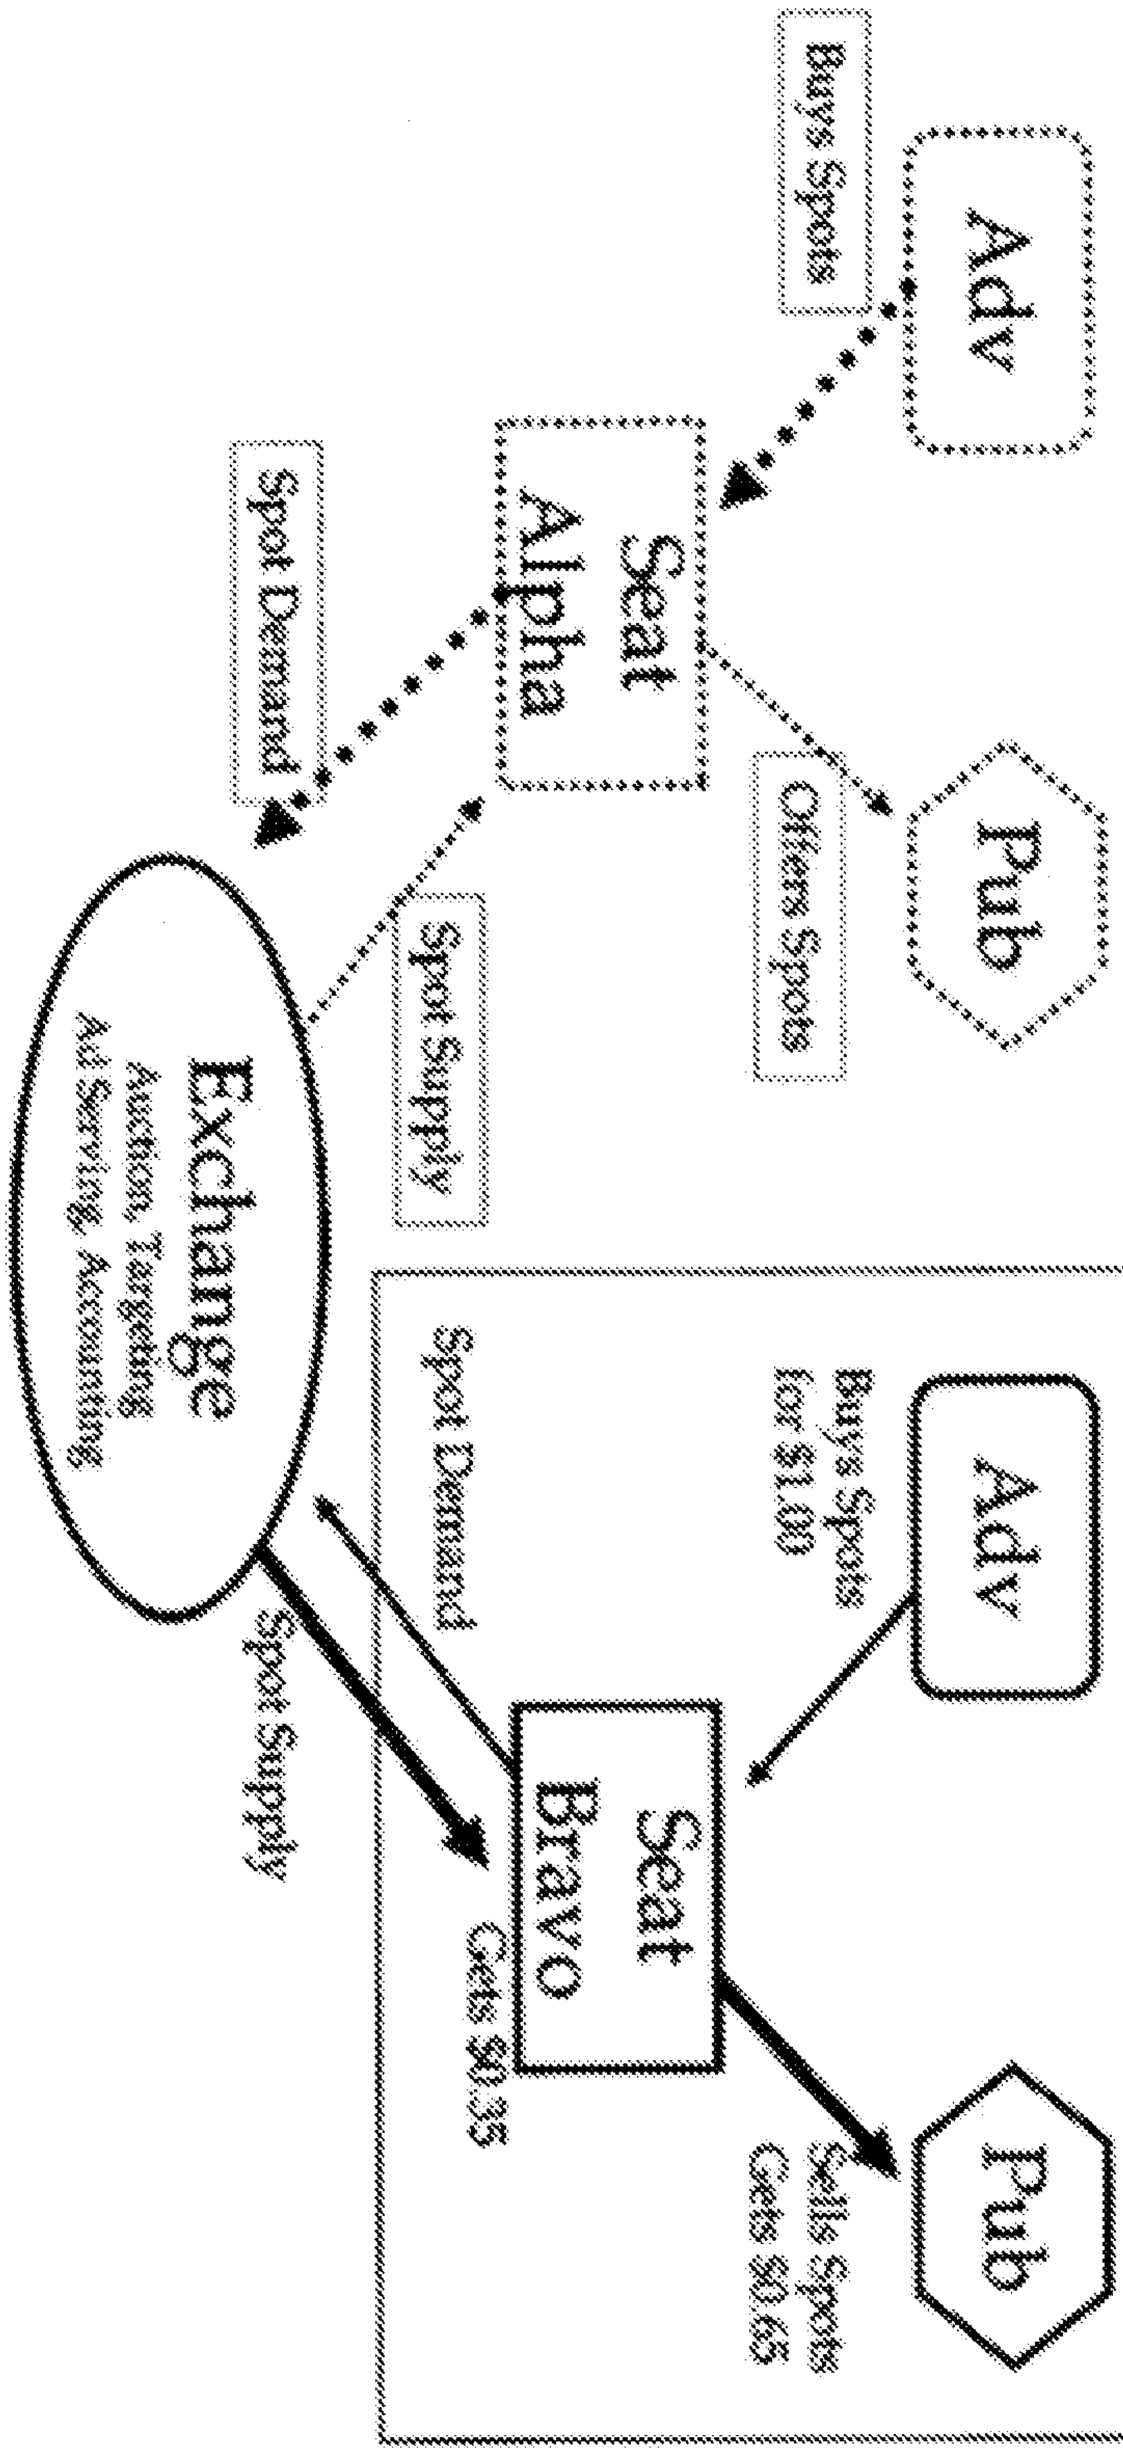
Buys Spots
Adv
Seat Alpha
Spot Demand
Offers Spots
Pub
Spot Supply
Exchange
Auction, Targeting
Ad Serving, Accounting
Private Network
Adv
Buys Spots for $100
Seat Bravo
Spot Demand
Gets $0.35
Sells Spots
Gets $0.65
Pub
Spot Supply

Figure A6

Figure A7

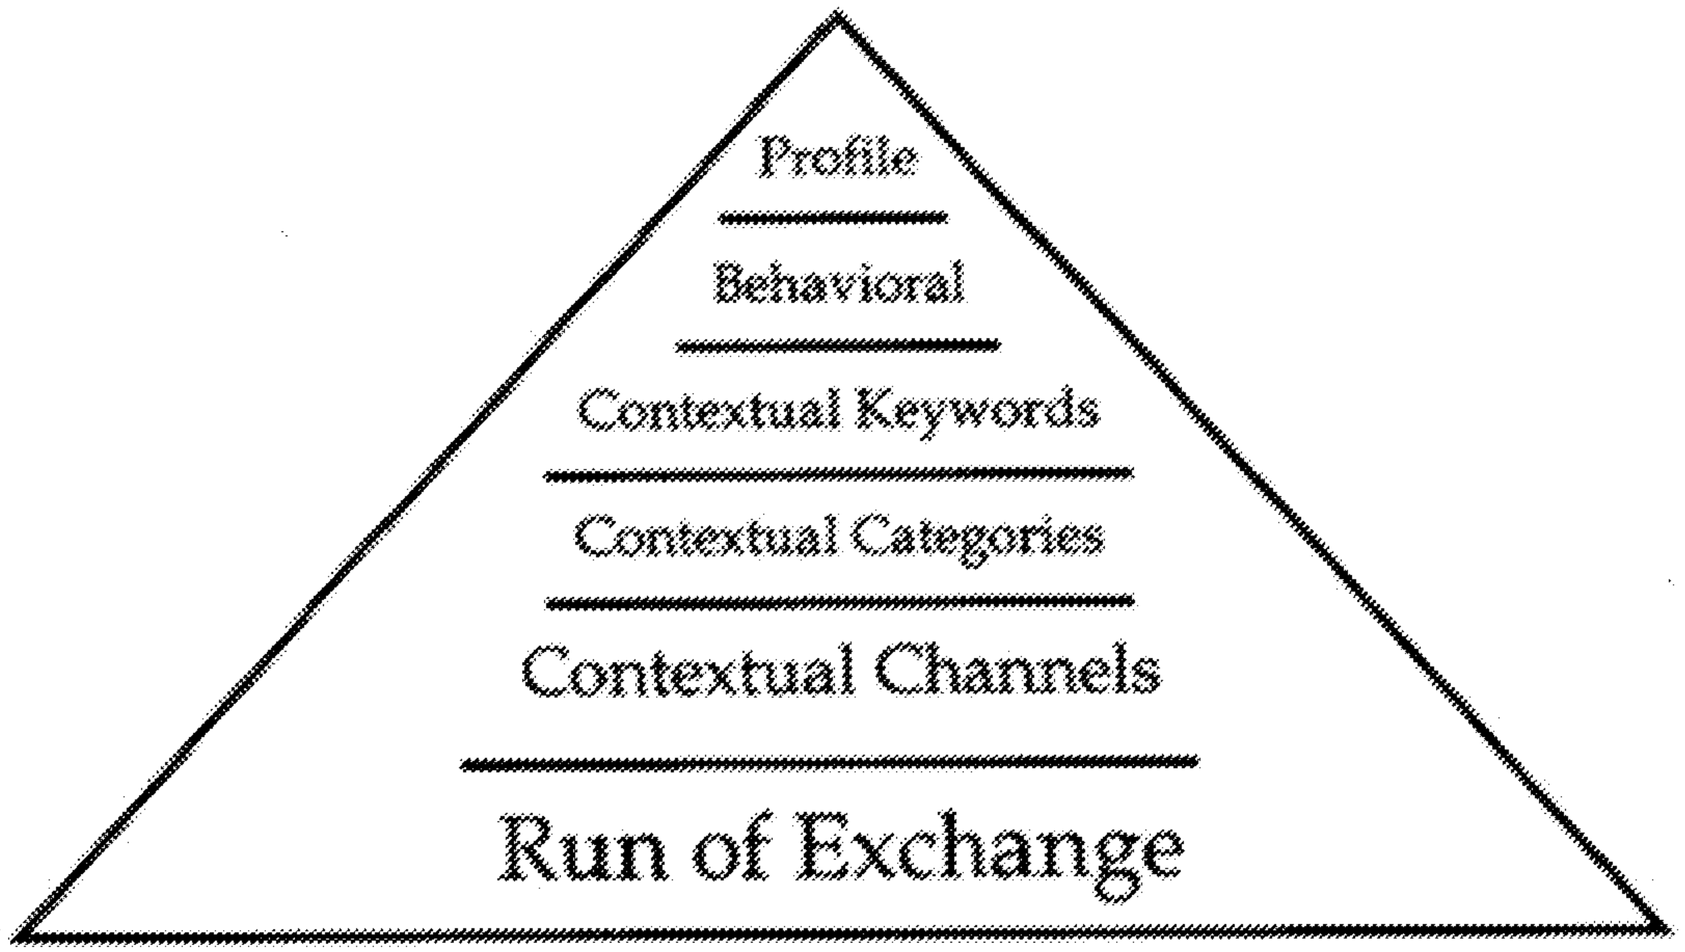
Profile
Behavioral
Contextual Keywords
Contextual Categories
Contextual Channels
Run of Exchange

Figure A8

1 Million Impressions

price	
1.00	100,000 IMP's
.55	400,000 IMP's
.35	200,000 IMP's
Unsold Or undersold	300,000 IMP's

Fixed Pate Pricing

PUBLISHER REVENUE $390

1 Million Impressions

price	
.97	100,000 IMP's
.71	135,000 IMP's
.70	135,000 IMP's
.60	135,000 IMP's
.57	215,000 IMP's
.52	100,000 IMP's
.44	100,000 IMP's
.22	80,000 IMP's

Value-Based Pricing

$604

Figure 1B

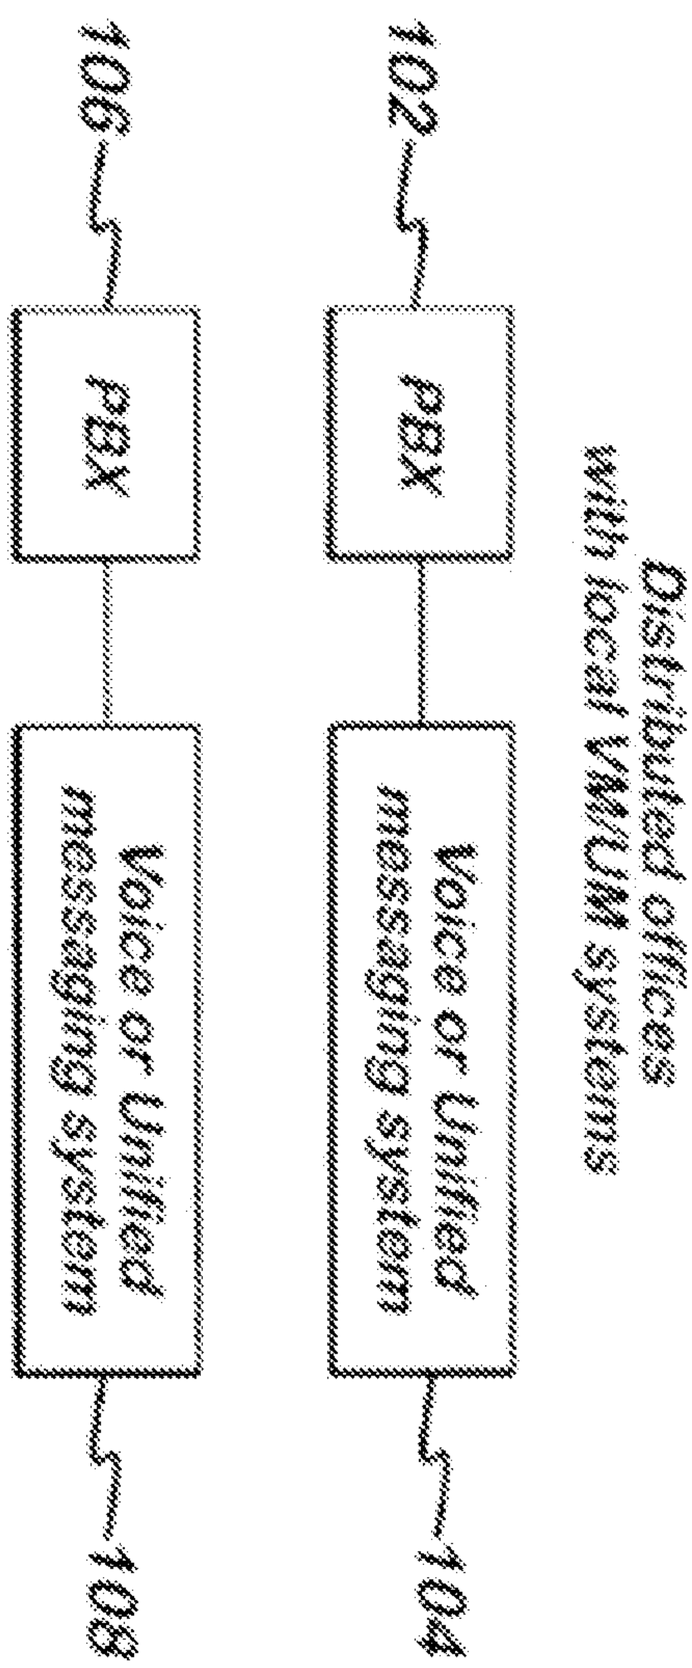
Distributed office
with local VM/UM systems
PBX
102
Voice or unified
messaging system
104
PBX
106
Voice or unified
messaging system
108

Figure 2B2

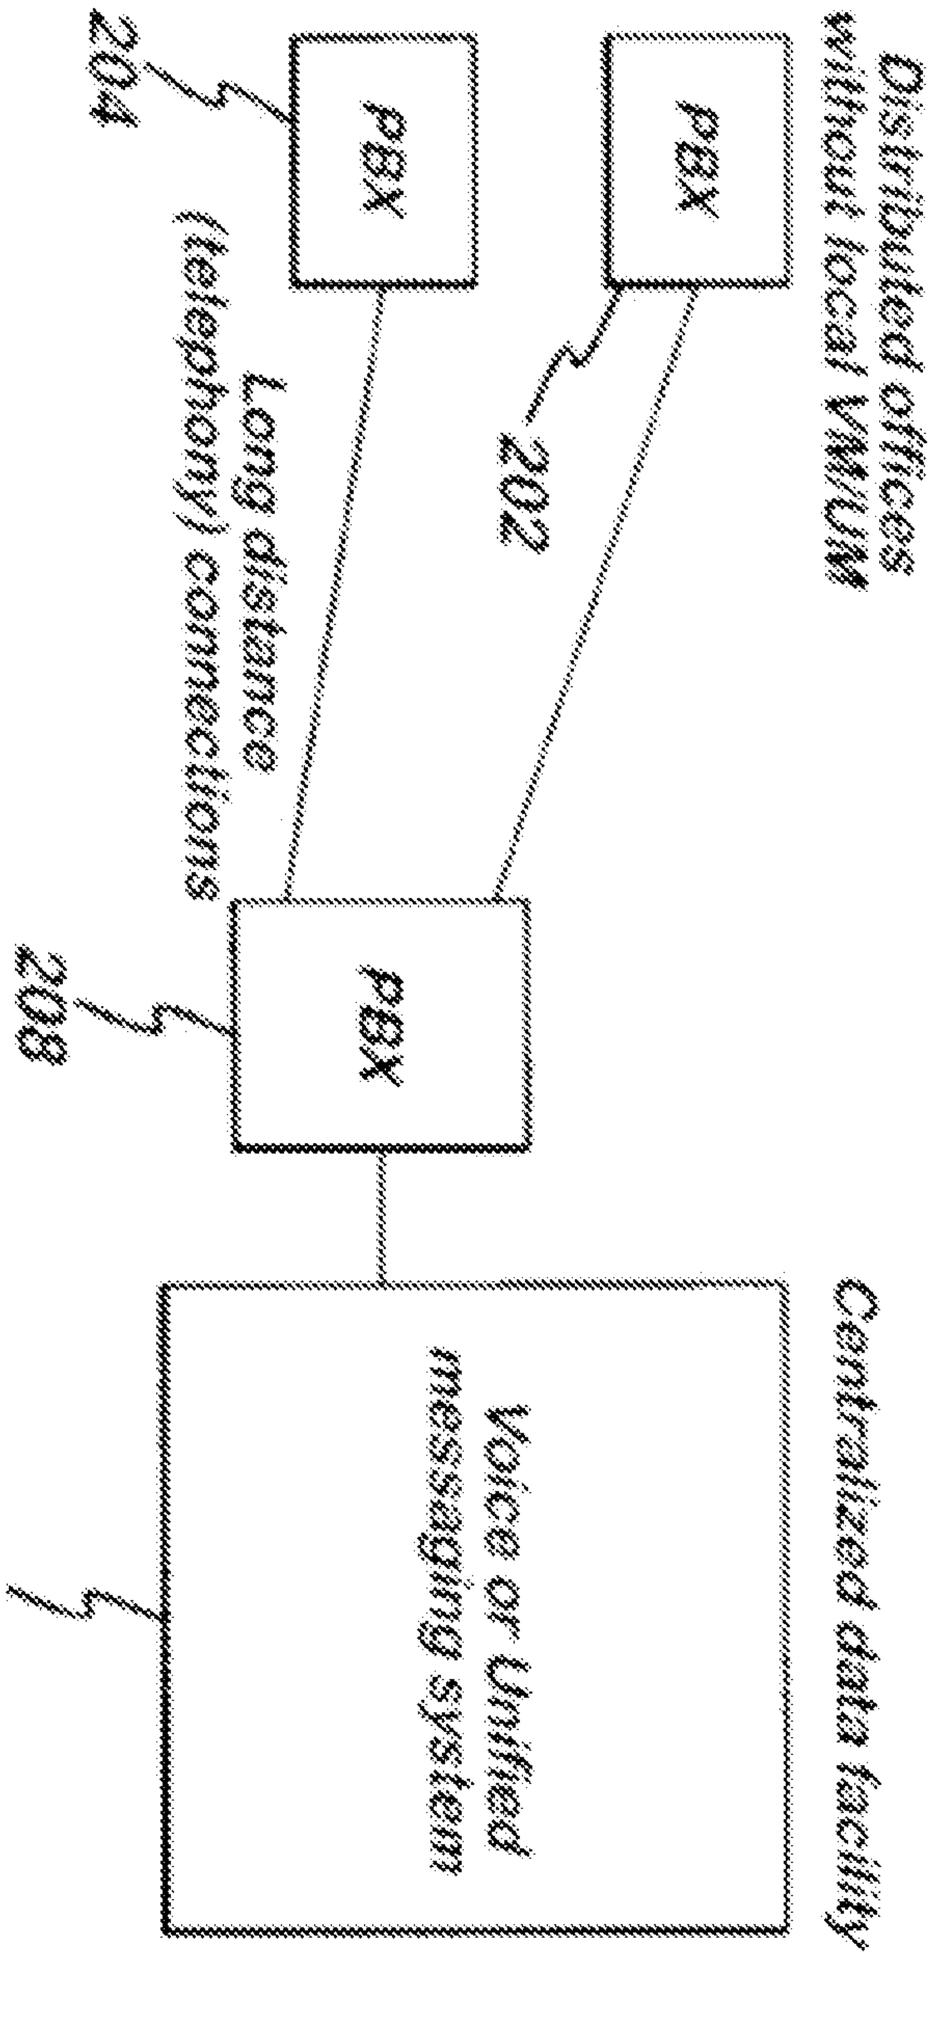
Distributed offices without local VM/UM
PBX
202
PBX
204
Long distance (telephony) connections
PBX
208
Centralized data facility
Voice or Unified messaging system
206

Figure B3

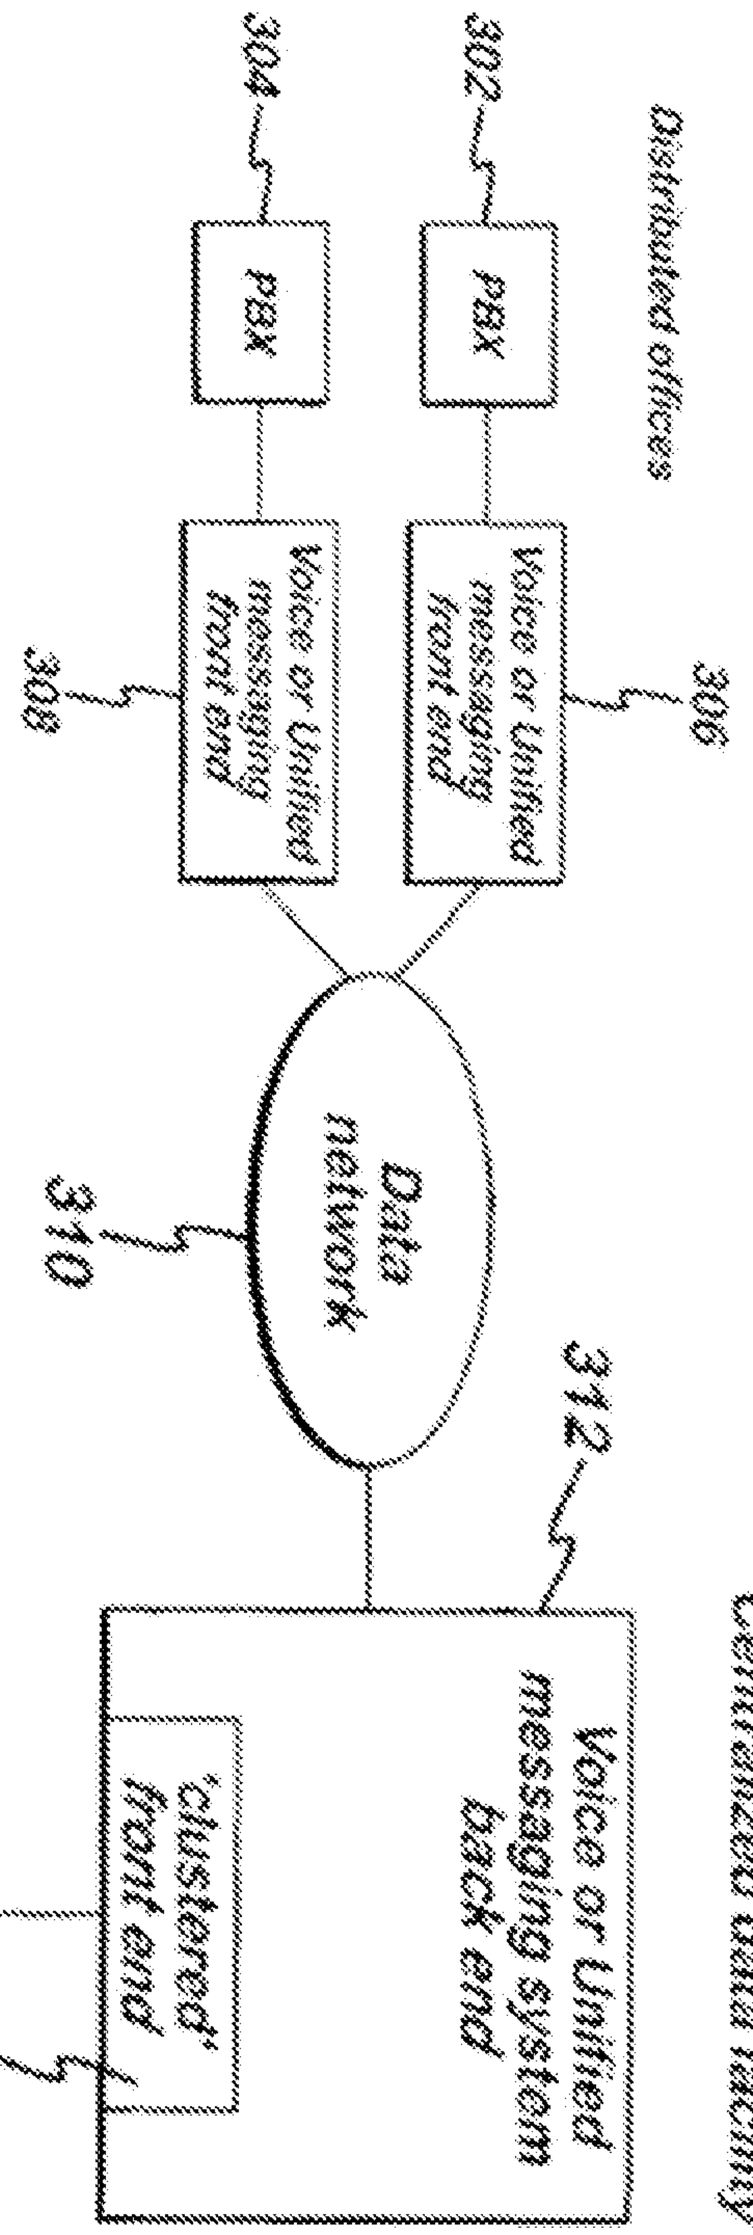
Centralized data facility
Voice or Unified messaging system back end
306
Voice or Unified messaging front end
308
Distributed offices
PBX
304
PBX
302
Cluster, front end
316
Data network
310
312
PBX
314
PSTN lines (Toll-free)

Figure B4

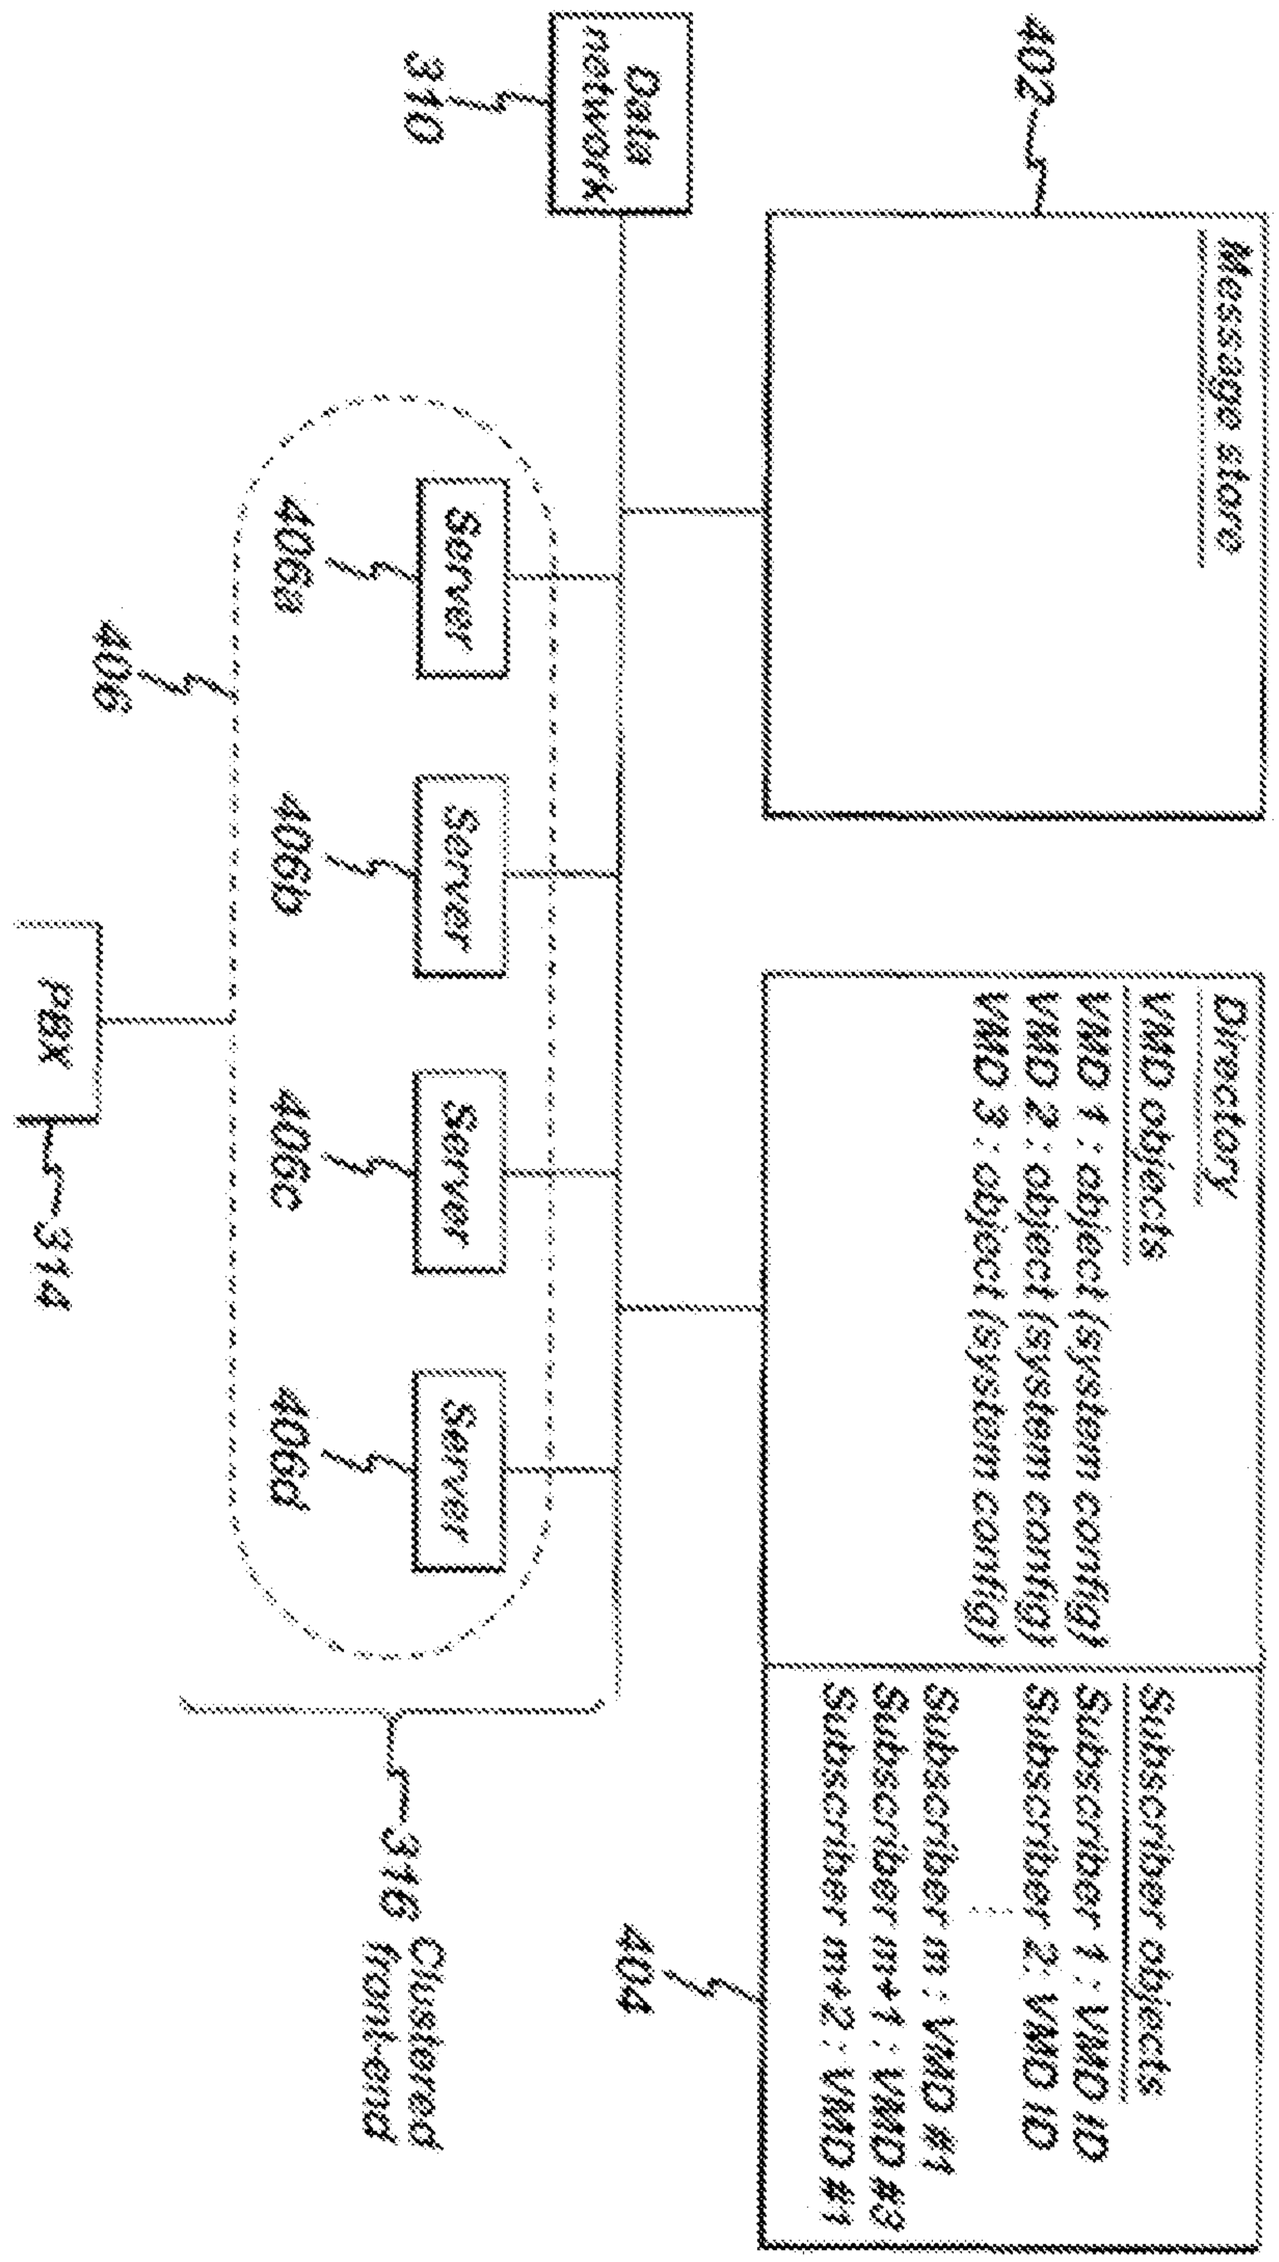
Data network
310
Message store
Directory
VMO objects
VMO 1 : object (system config)
VMO 2 : object (system config)
VMO 3 : object (system config)
Subscriber objects
Subscriber 1 : VMO ID
Subscriber 2: VMO ID
Subscriber m : VMO #4
Subscriber m+1 : VMO #3
Subscriber m+2 : VMO #1
402
404
406
408a
408b
408c
408d
Server
Server
Server
Server
PBX
314
316
Clustered front-end

Figure B5

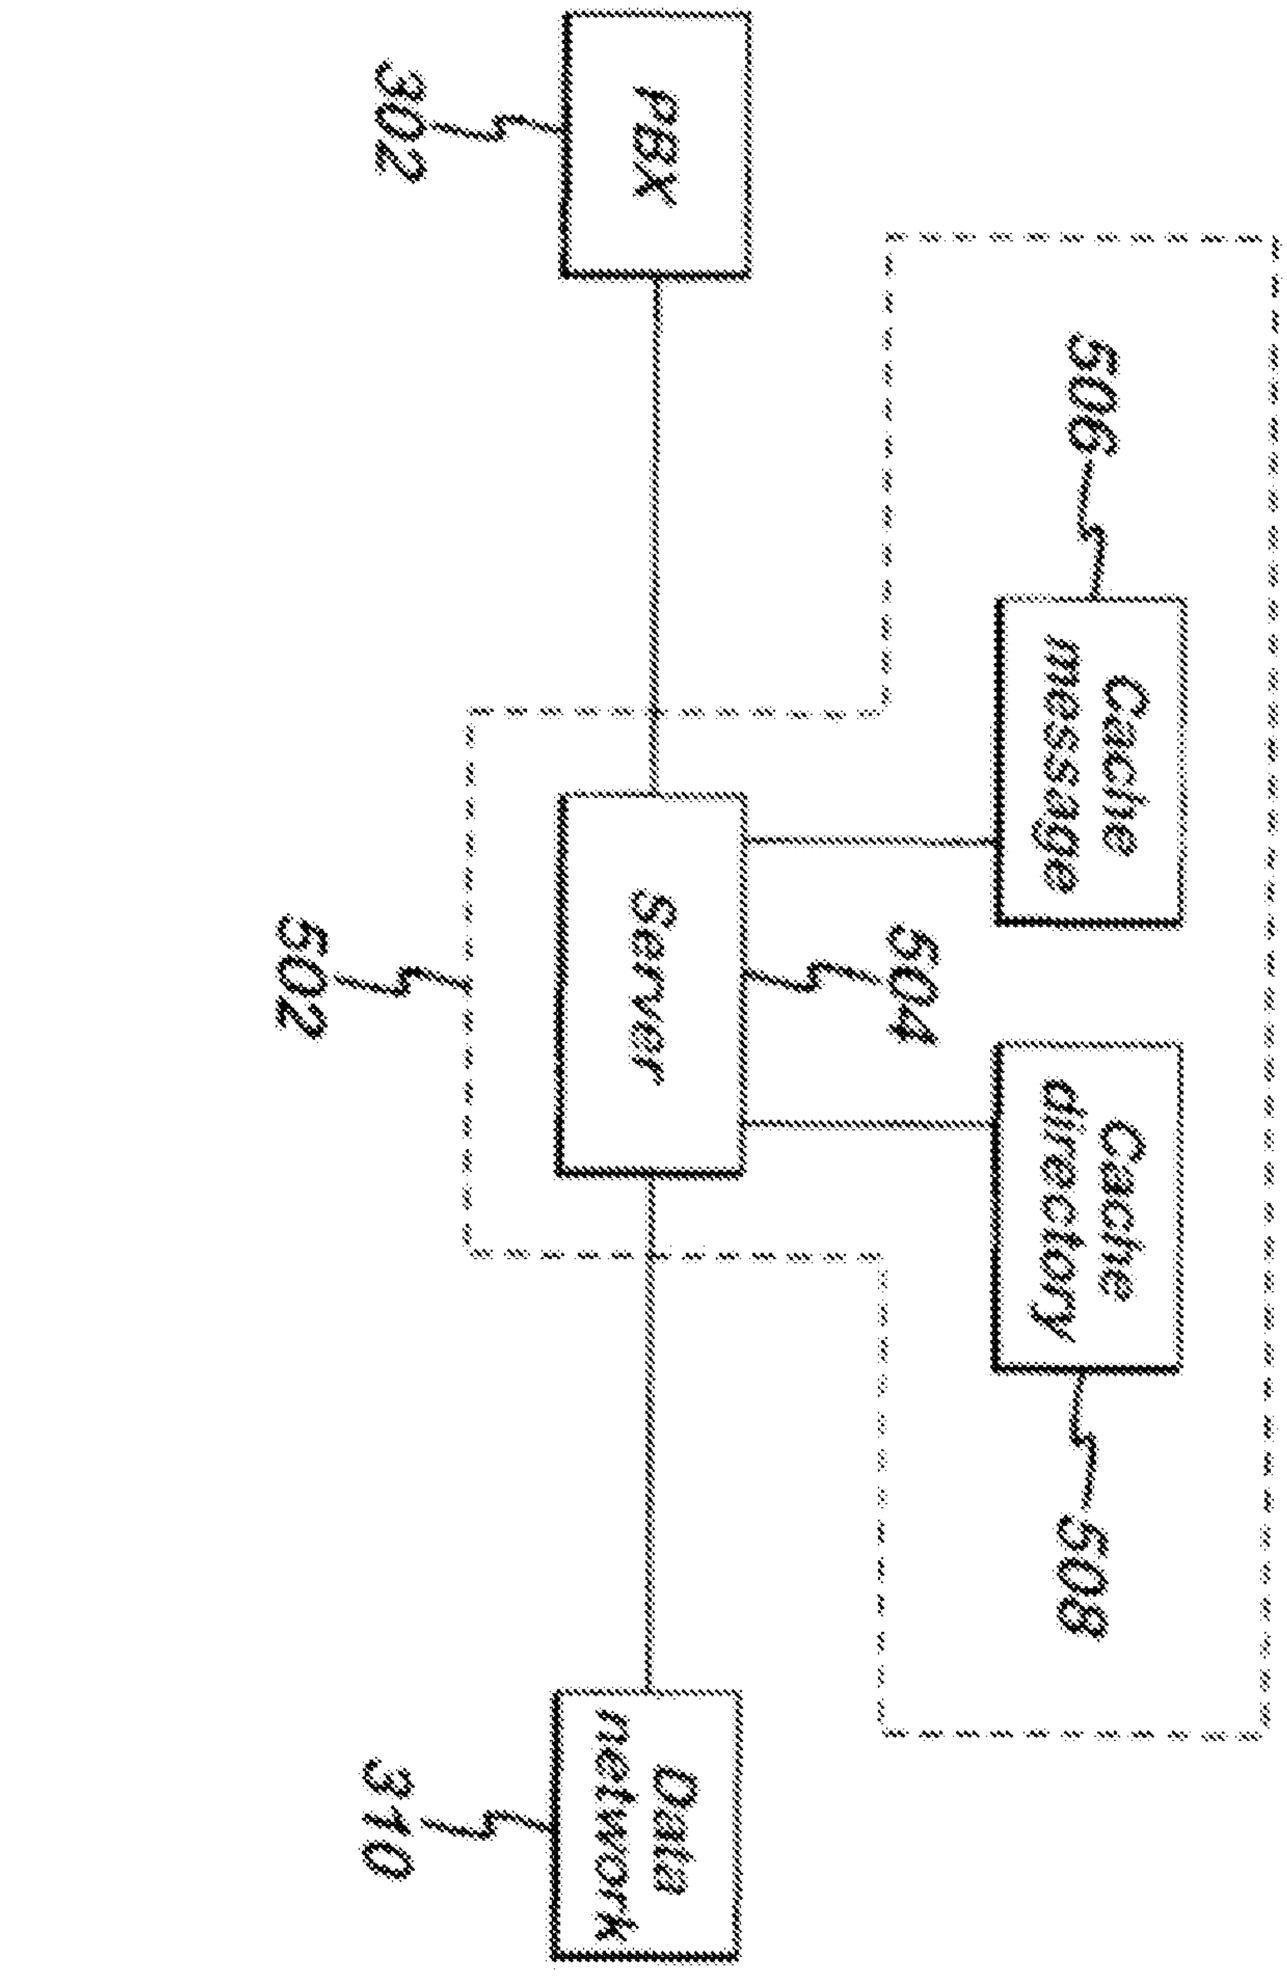
302
PBX
306
Cache message
304
Server
Cache directory
308
310
Data network

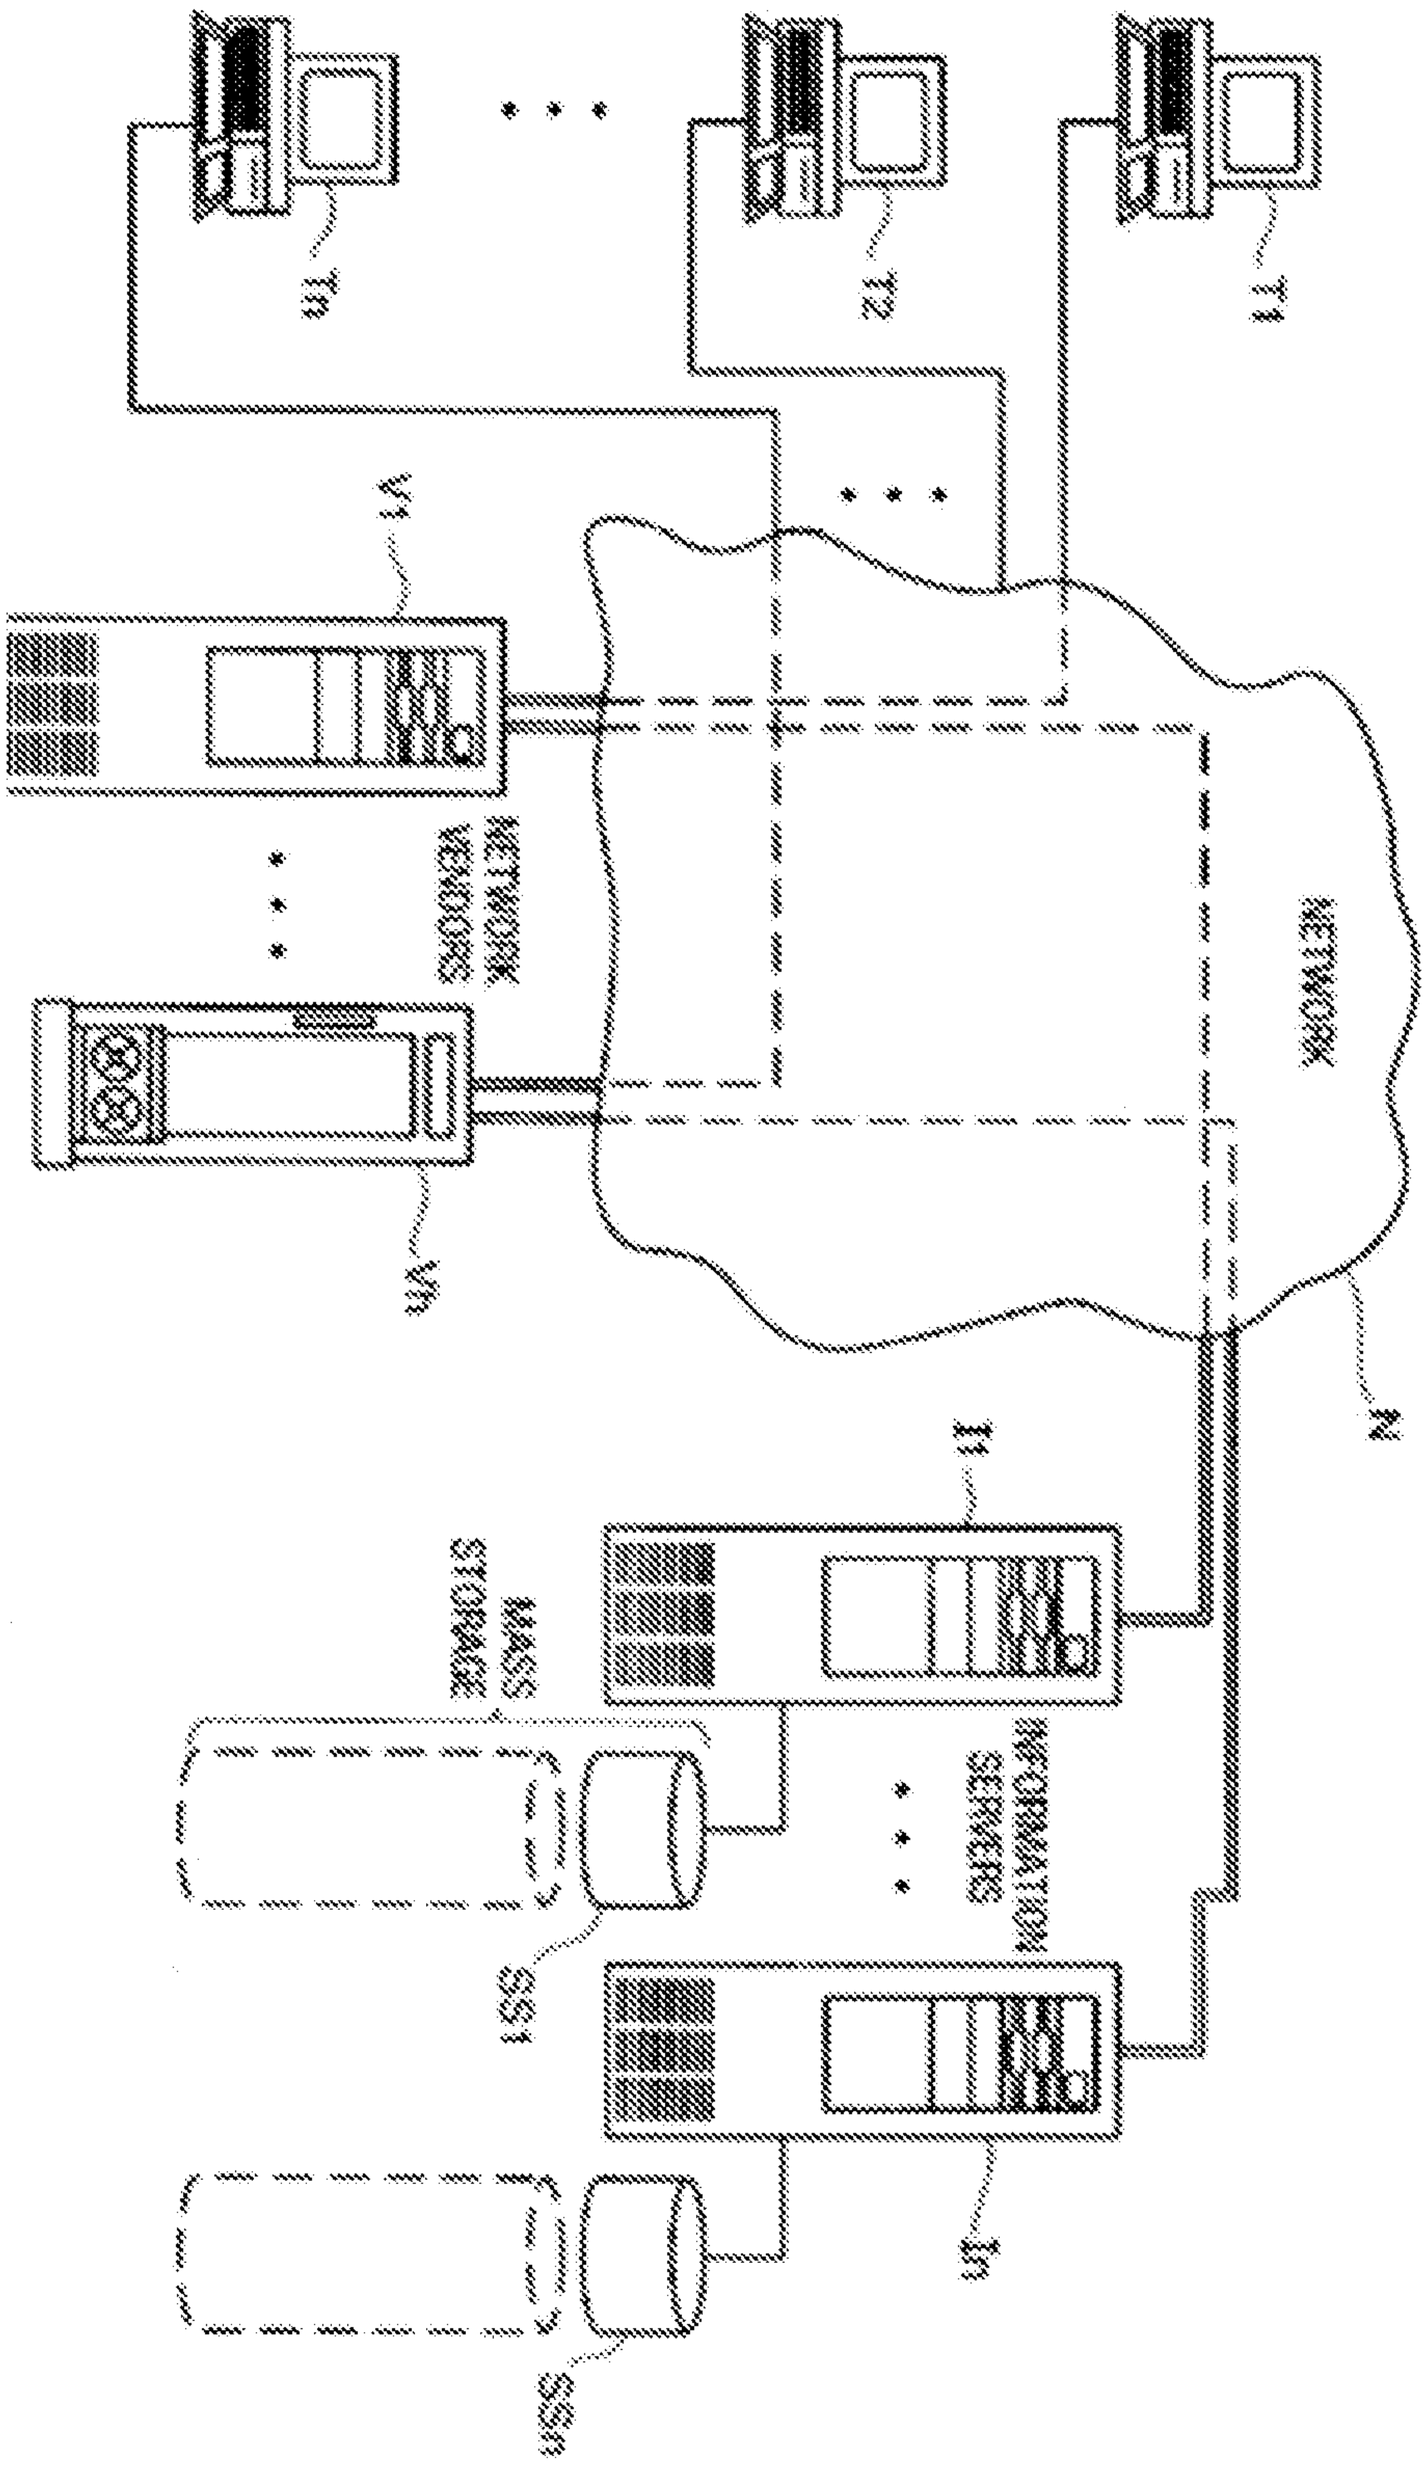

Figure C1

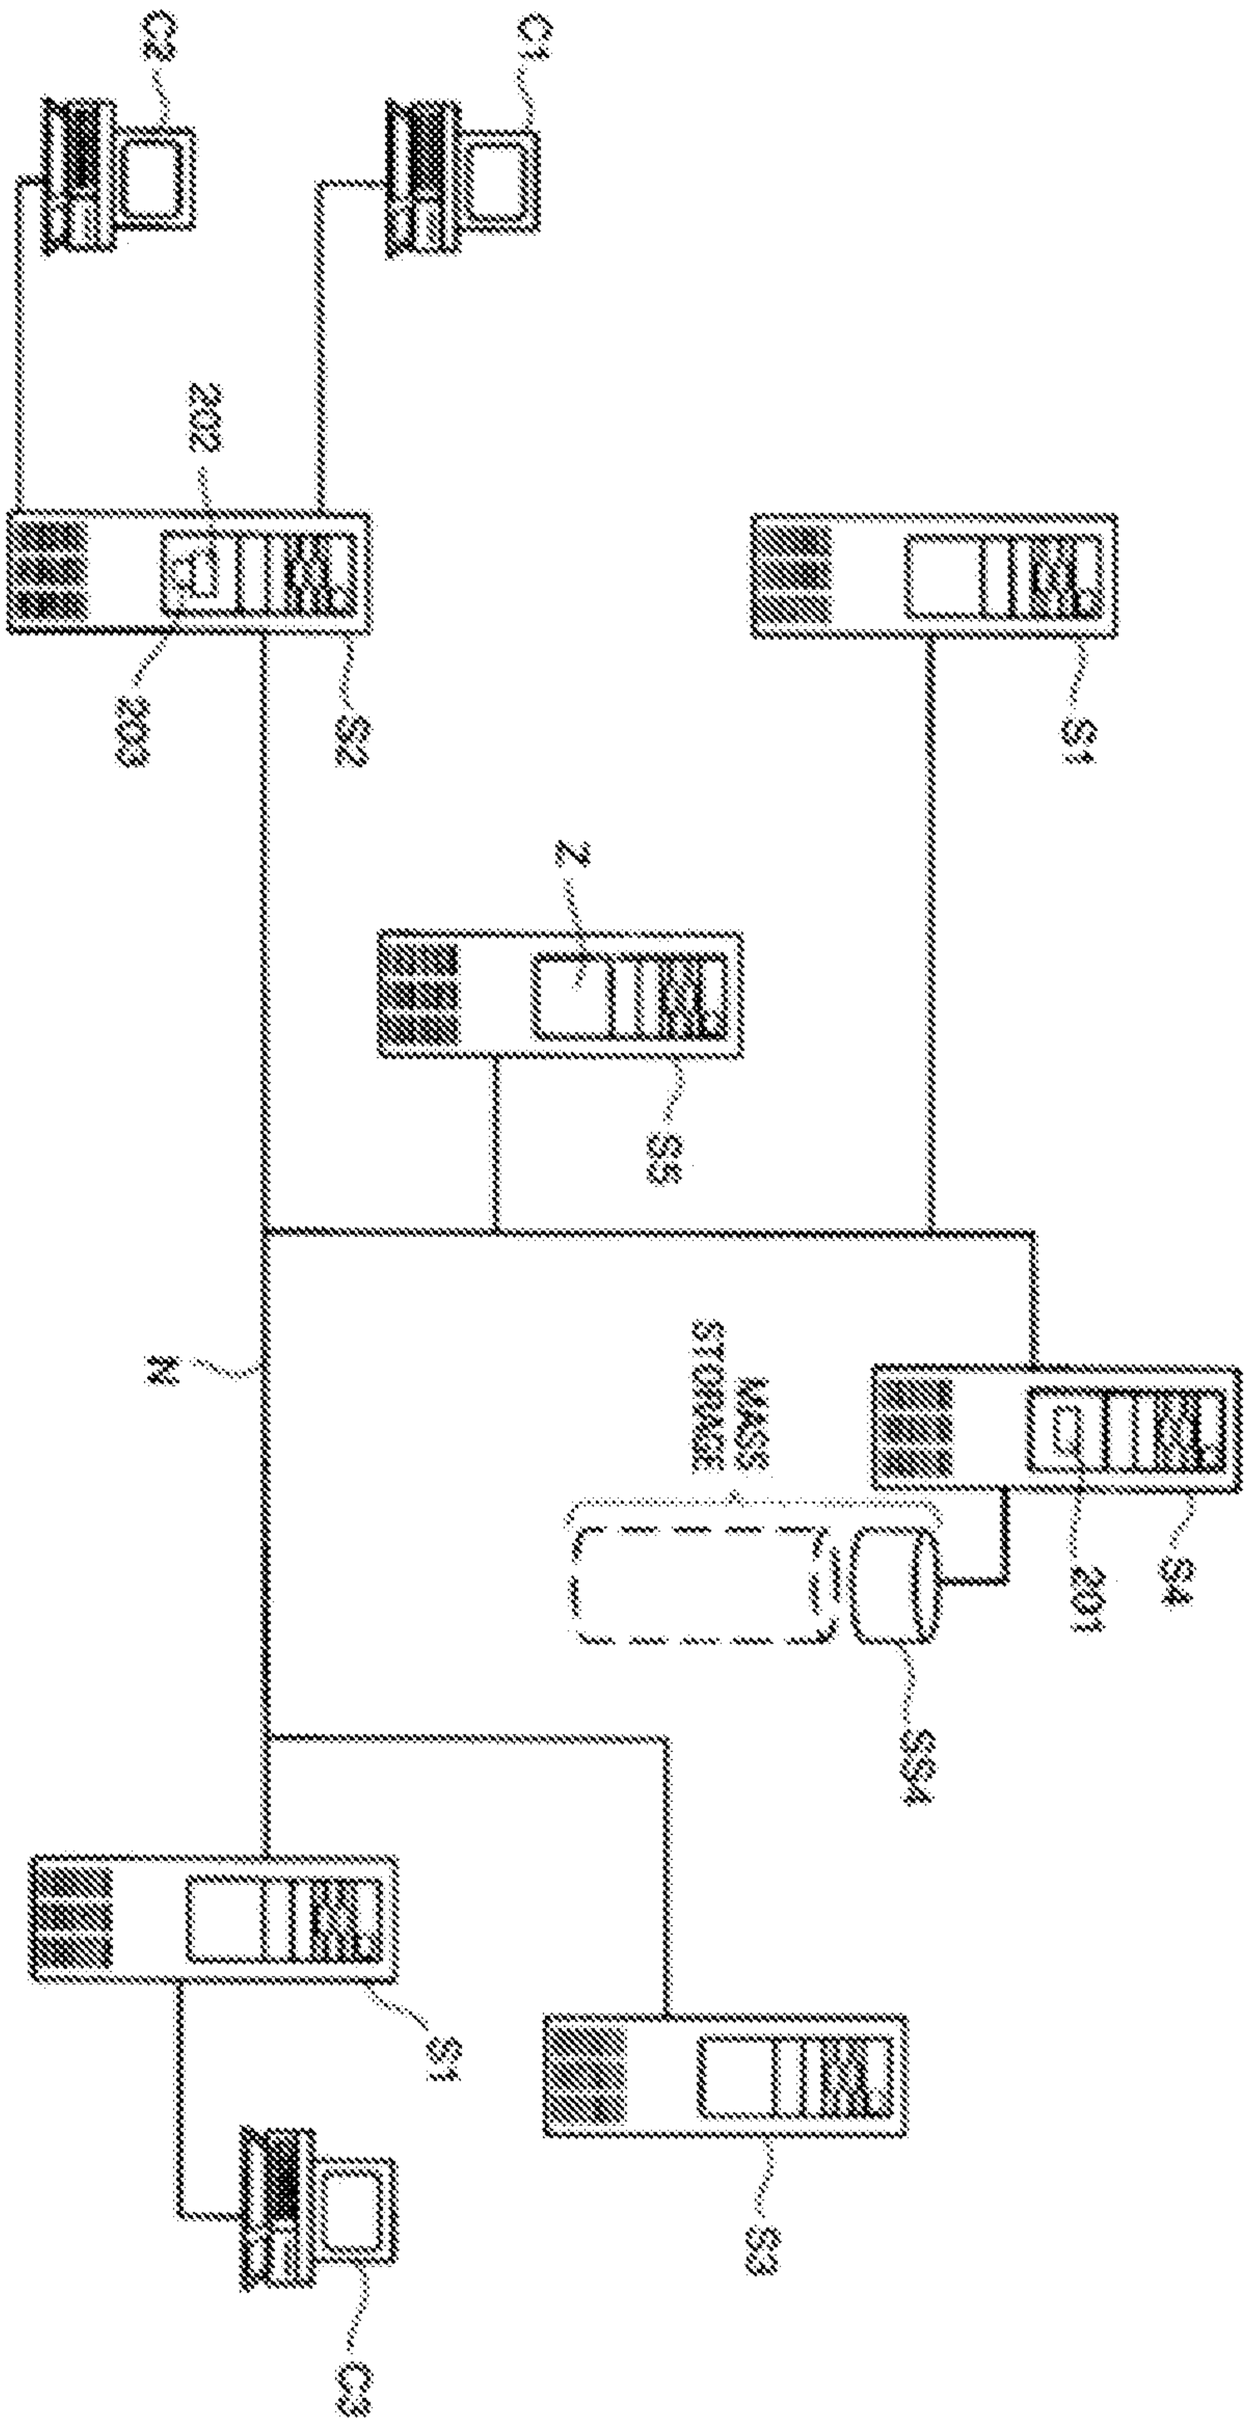

Fig. 2

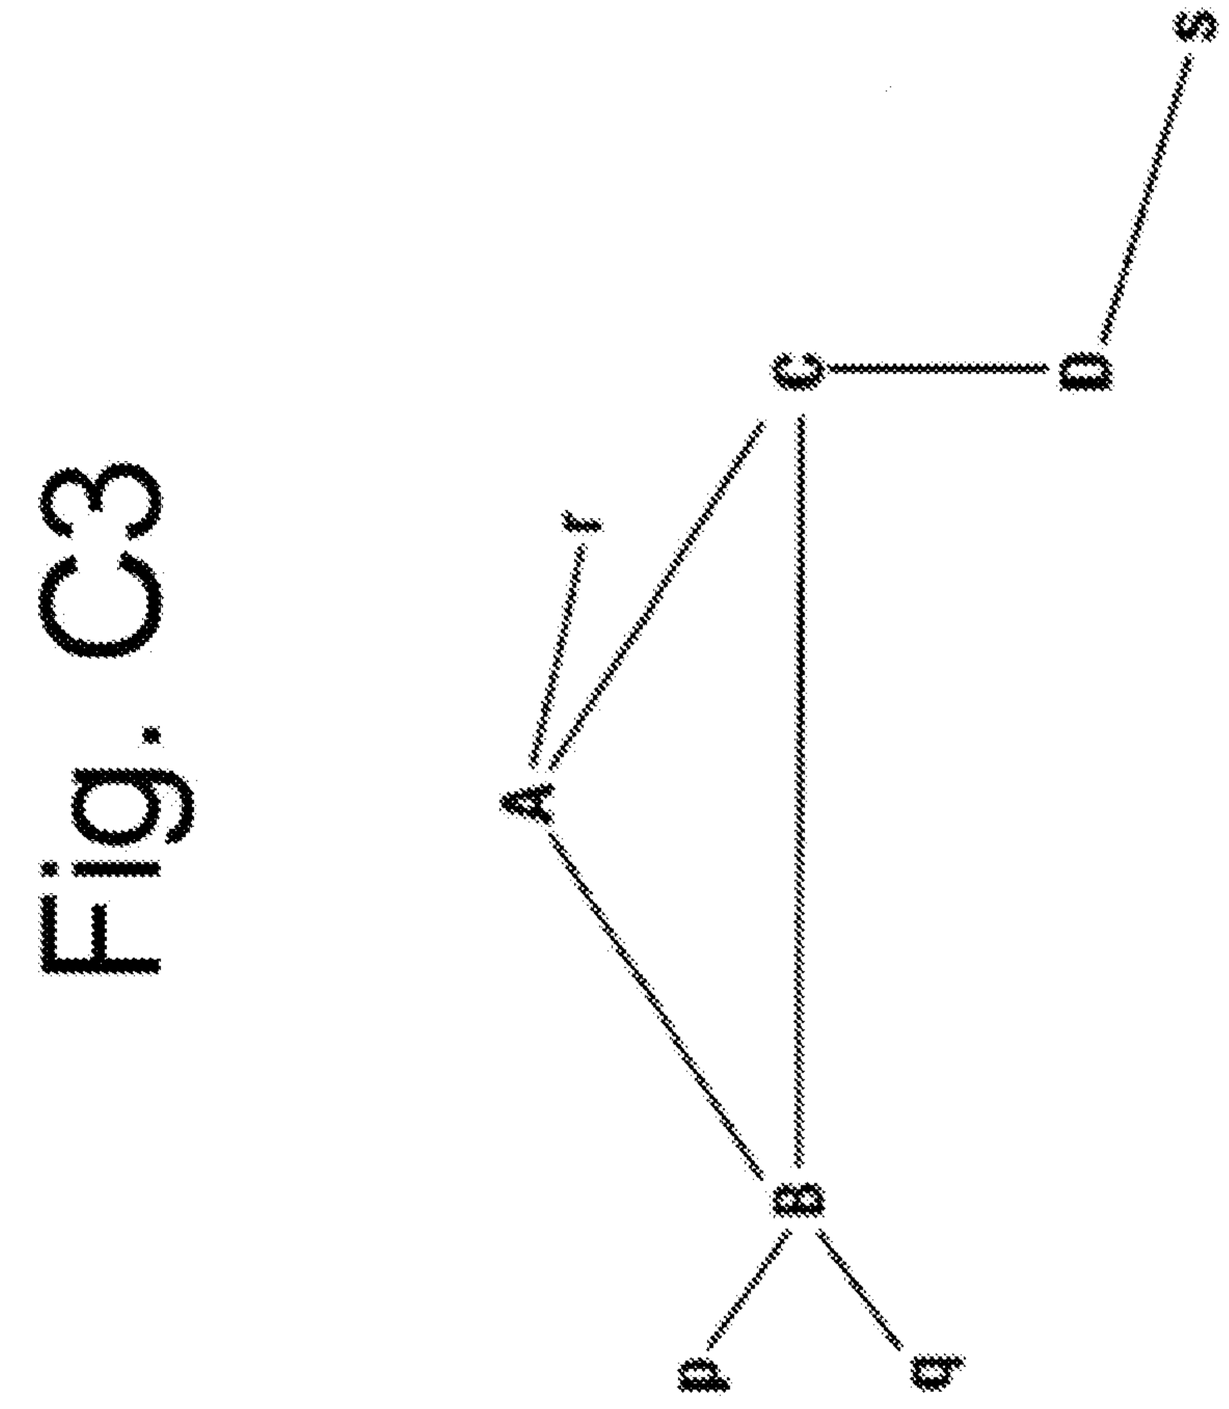

Fig. 33

Figure C4

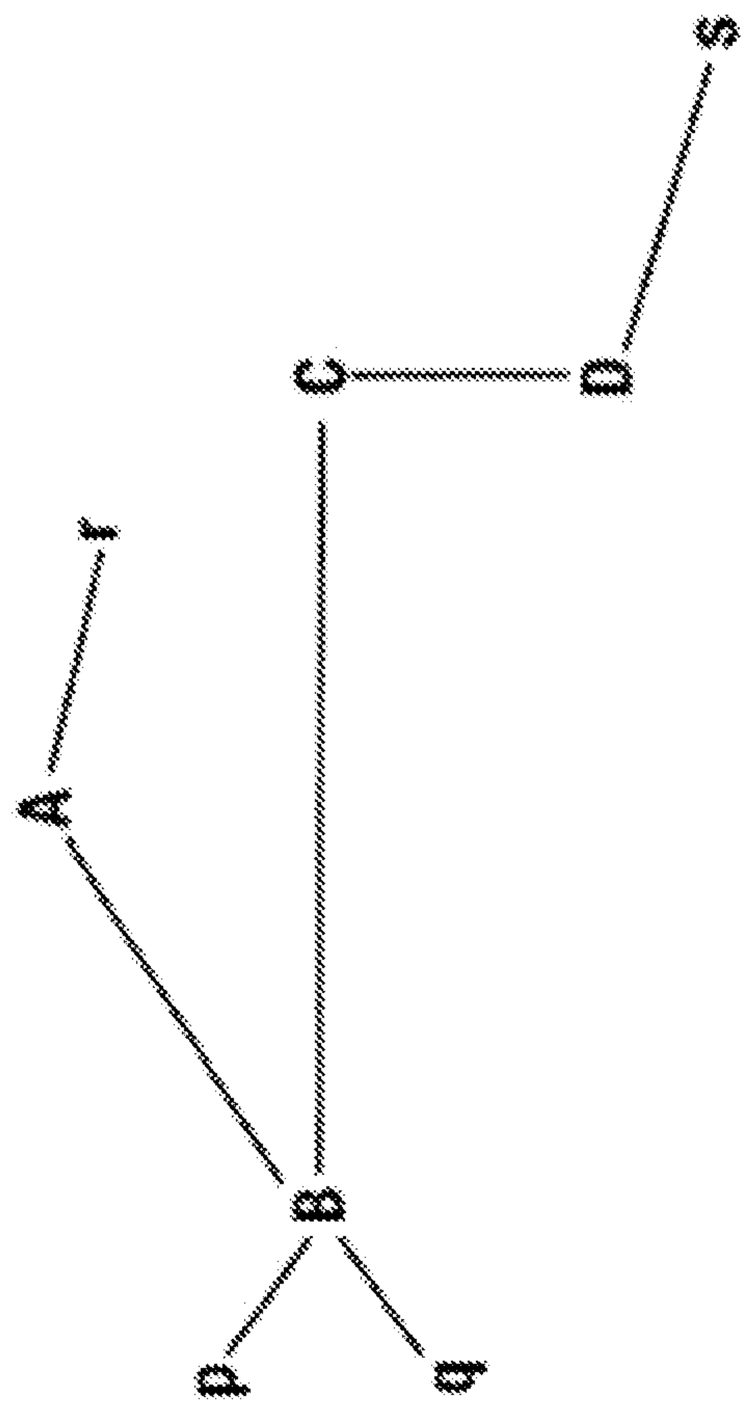

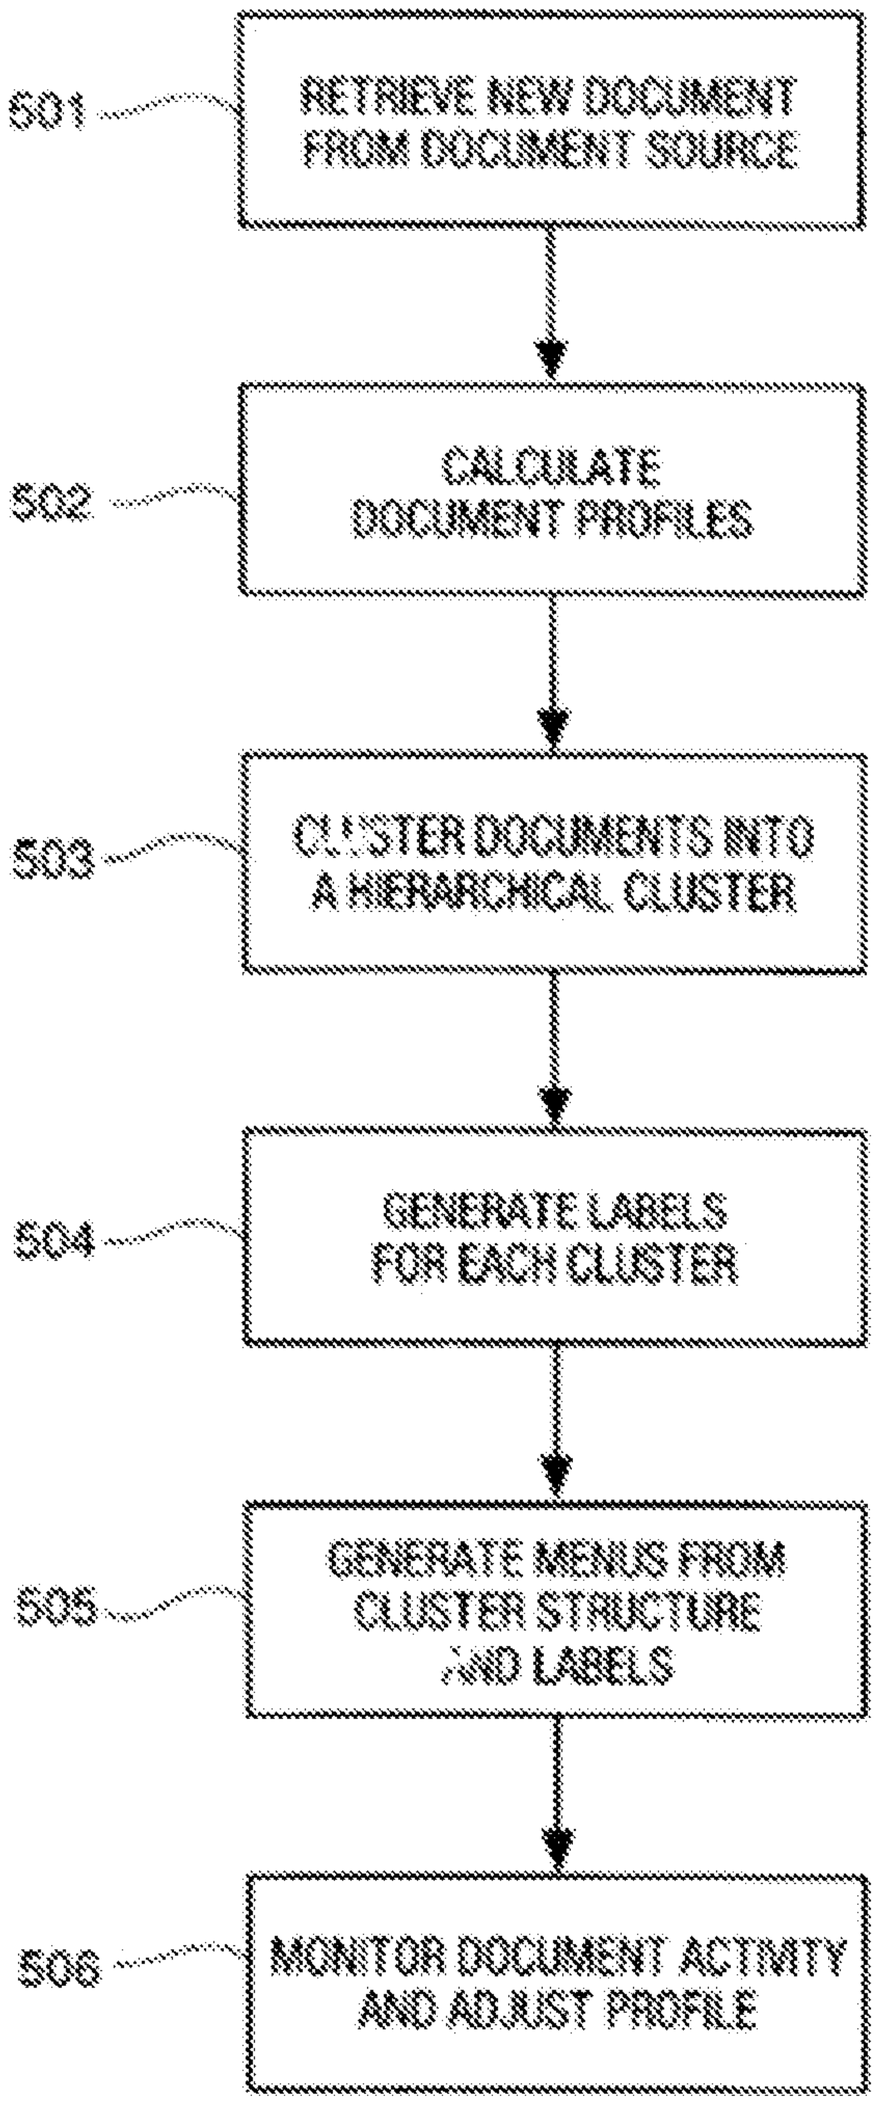
501
RETRIEVE NEW DOCUMENT FROM DOCUMENT SOURCE
502
CALCULATE DOCUMENT PROFILES
503
CLUSTER DOCUMENTS INTO A HIERARCHICAL CLUSTER
504
GENERATE LABELS FOR EACH CLUSTER
505
GENERATE MENUS FROM CLUSTER STRUCTURE AND LABELS
506
MONITOR DOCUMENT ACTIVITY AND ADJUST PROFILE

Figure 5

Figure 6G

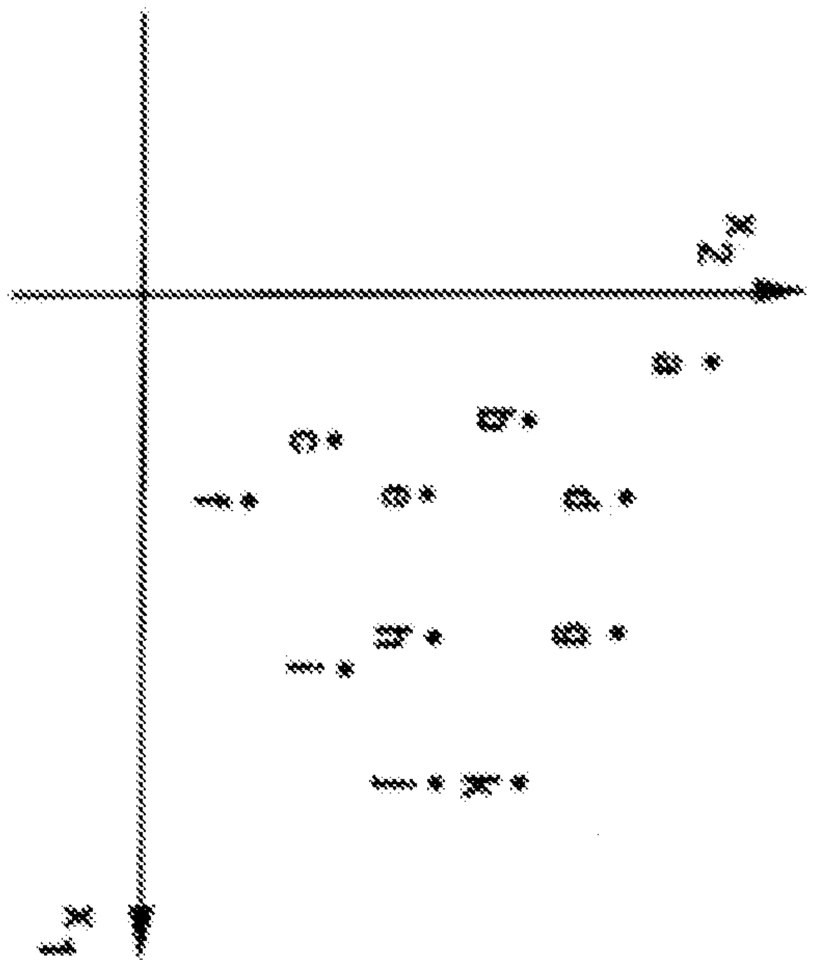

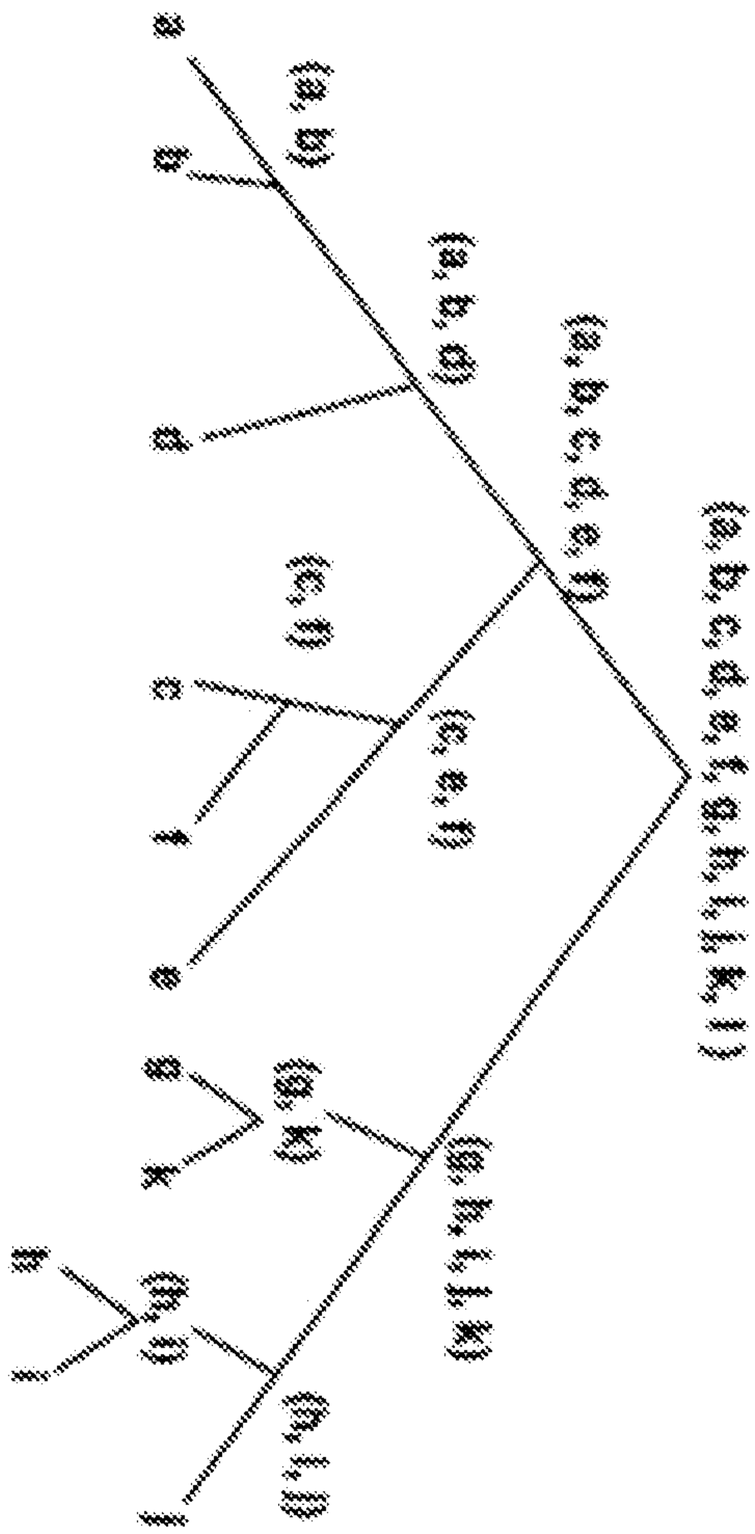
(a,b,c,d,e,f,g,h,i,j,k,l,m,n)
(b,c,d,e,f,g,h,i,j,k,l,m,n)
(a,b,c,d,e,f,g,h,i,j,k,l,m)
(a,b,c,d,e,f,g,h,i,j,k,l)
(a,b,c,d,e,f,g,h,i,j,k)
(a,b,c,d,e,f,g,h,i,j)
(a,b,c,d,e,f,g,h,i)
(a,b,c,d,e,f,g,h)
(a,b,c,d,e,f,g)
(a,b,c,d,e,f)
(a,b,c,d,e)
(a,b,c,d)
(a,b,c)
(a,b)
(a)
a
b
c
d
e
f
g
h
i
j
k
l
m
n

Figure 7

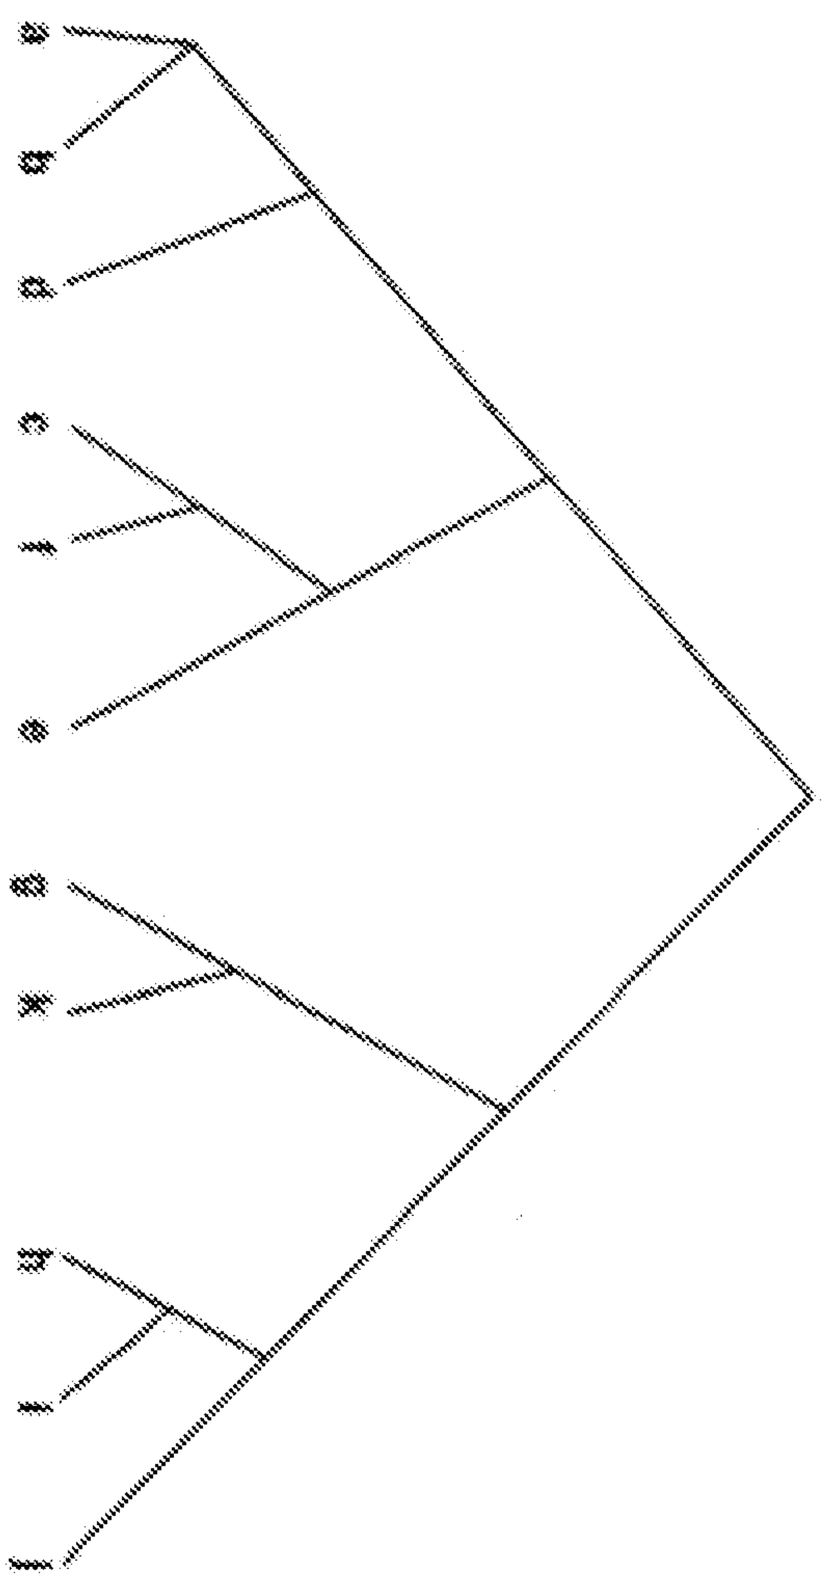

Figure 38

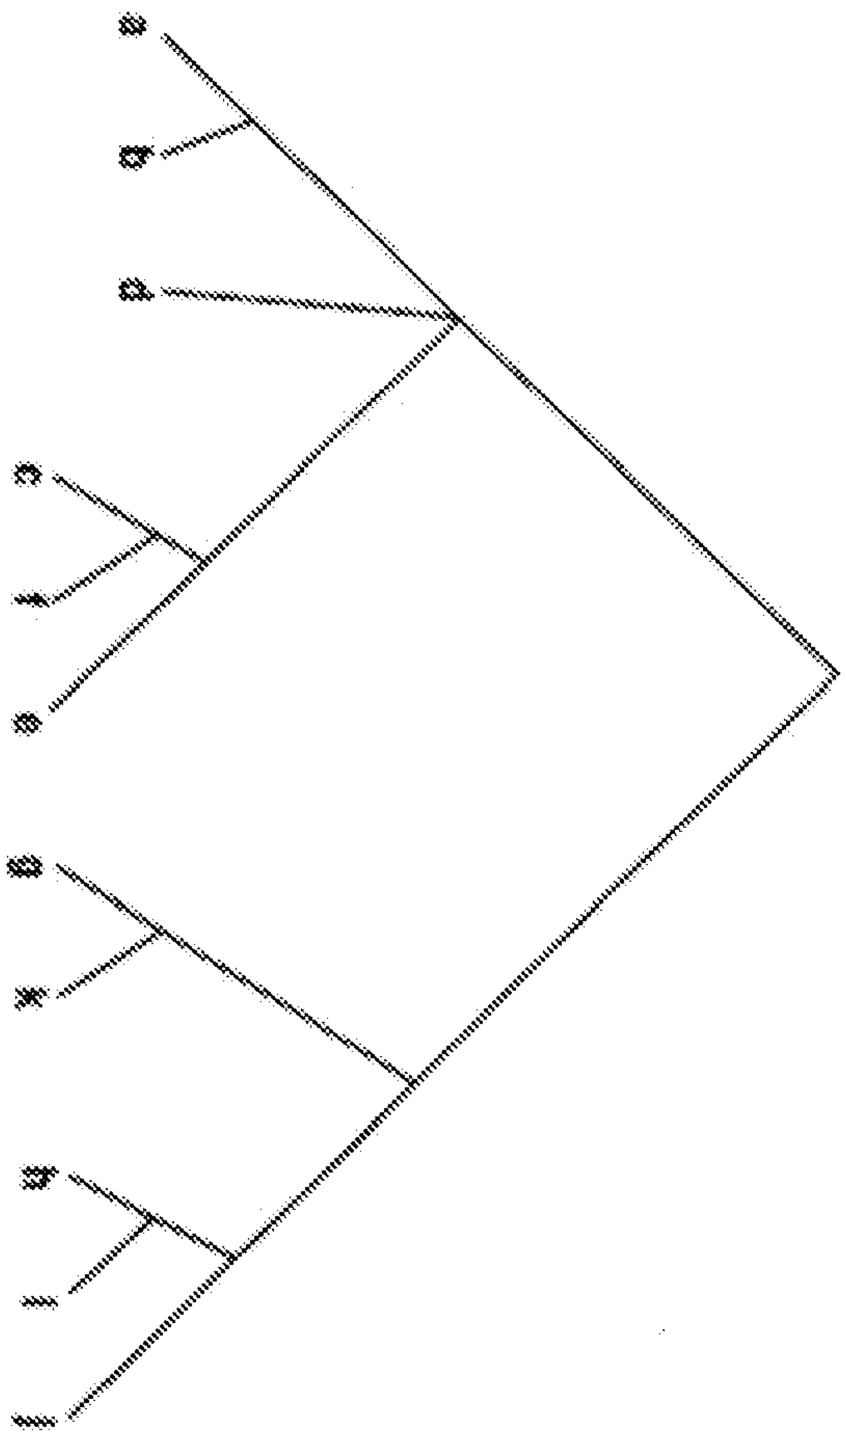

Figure 63

Figure 1C1

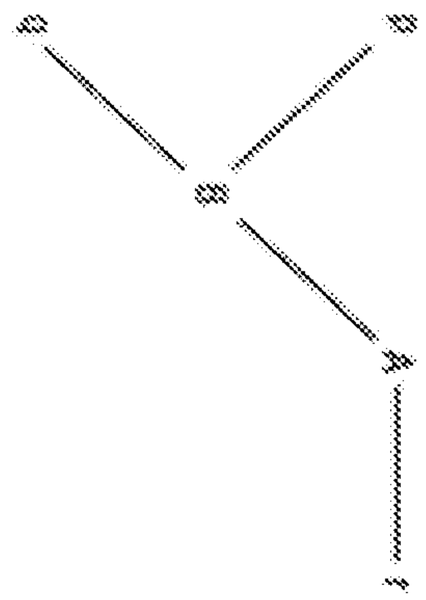

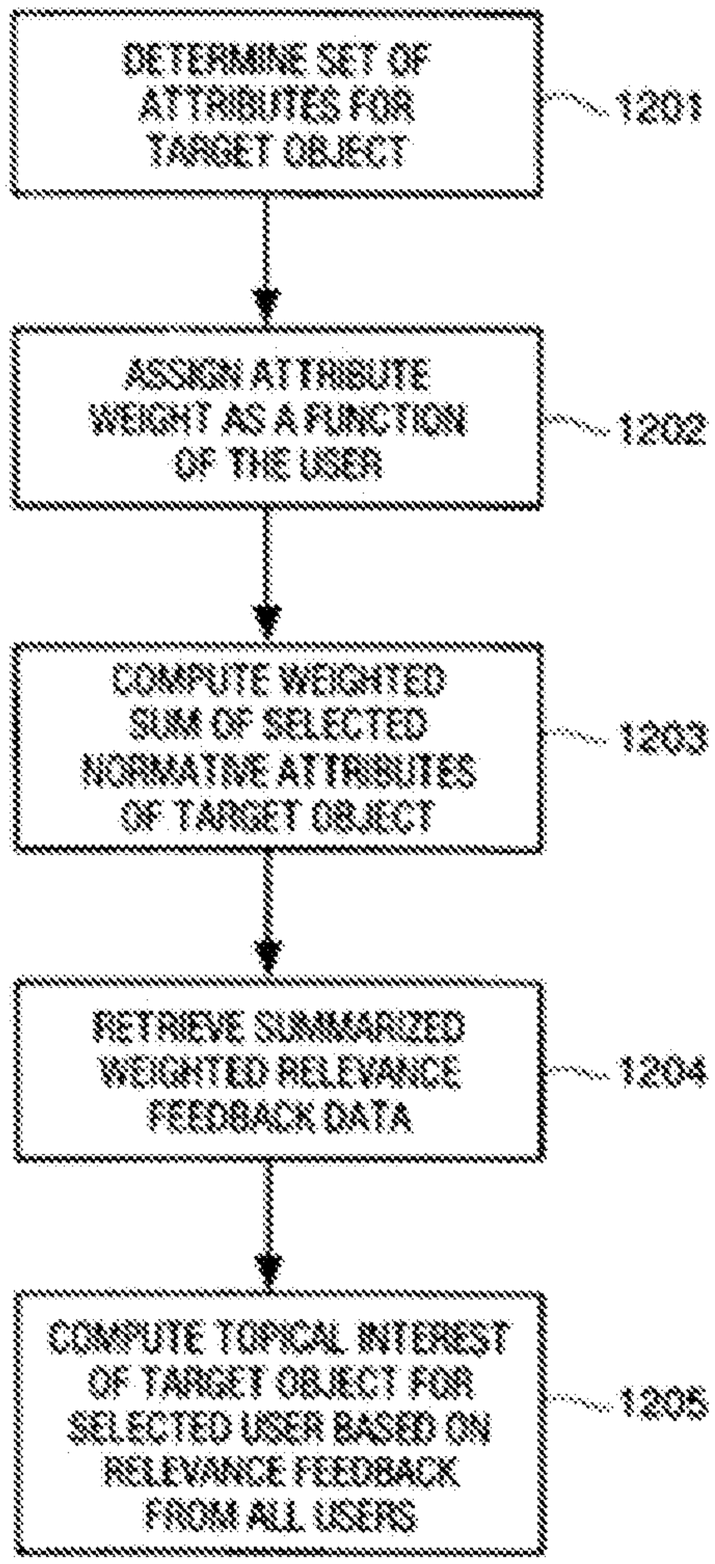
DETERMINE SET OF ATTRIBUTES FOR TARGET OBJECT
1201
ASSIGN ATTRIBUTE WEIGHT AS A FUNCTION OF THE USER
1202
COMPUTE WEIGHTED SUM OF SELECTED NORMATIVE ATTRIBUTES OF TARGET OBJECT
1203
RETRIEVE SUMMARIZED WEIGHTED RELEVANCE FEEDBACK DATA
1204
COMPUTE TOPICAL INTEREST OF TARGET OBJECT FOR SELECTED USER BASED ON RELEVANCE FEEDBACK FROM ALL USERS
1205

Figure 12

Figure C / 13A

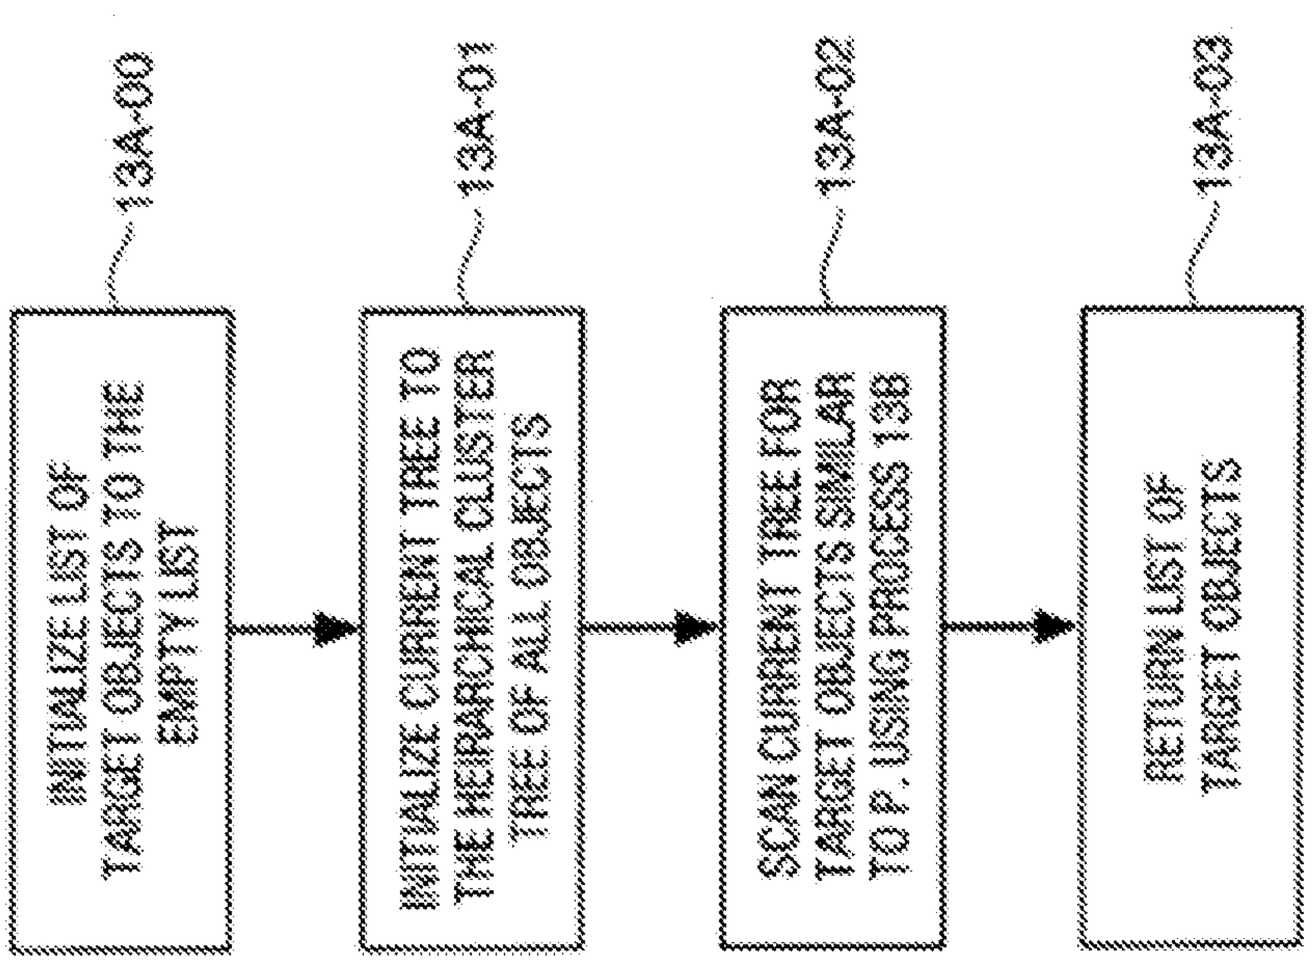
134-00
INITIALIZE LIST OF TARGET OBJECTS TO THE EMPTY LIST
134-01
INITIALIZE CURRENT TREE TO THE HIERARCHICAL CLUSTER TREE OF ALL OBJECTS
134-02
SCAN CURRENT TREE FOR TARGET OBJECTS SIMILAR TO P, USING PROCESS 13B
134-03
RETURN LIST OF TARGET OBJECTS

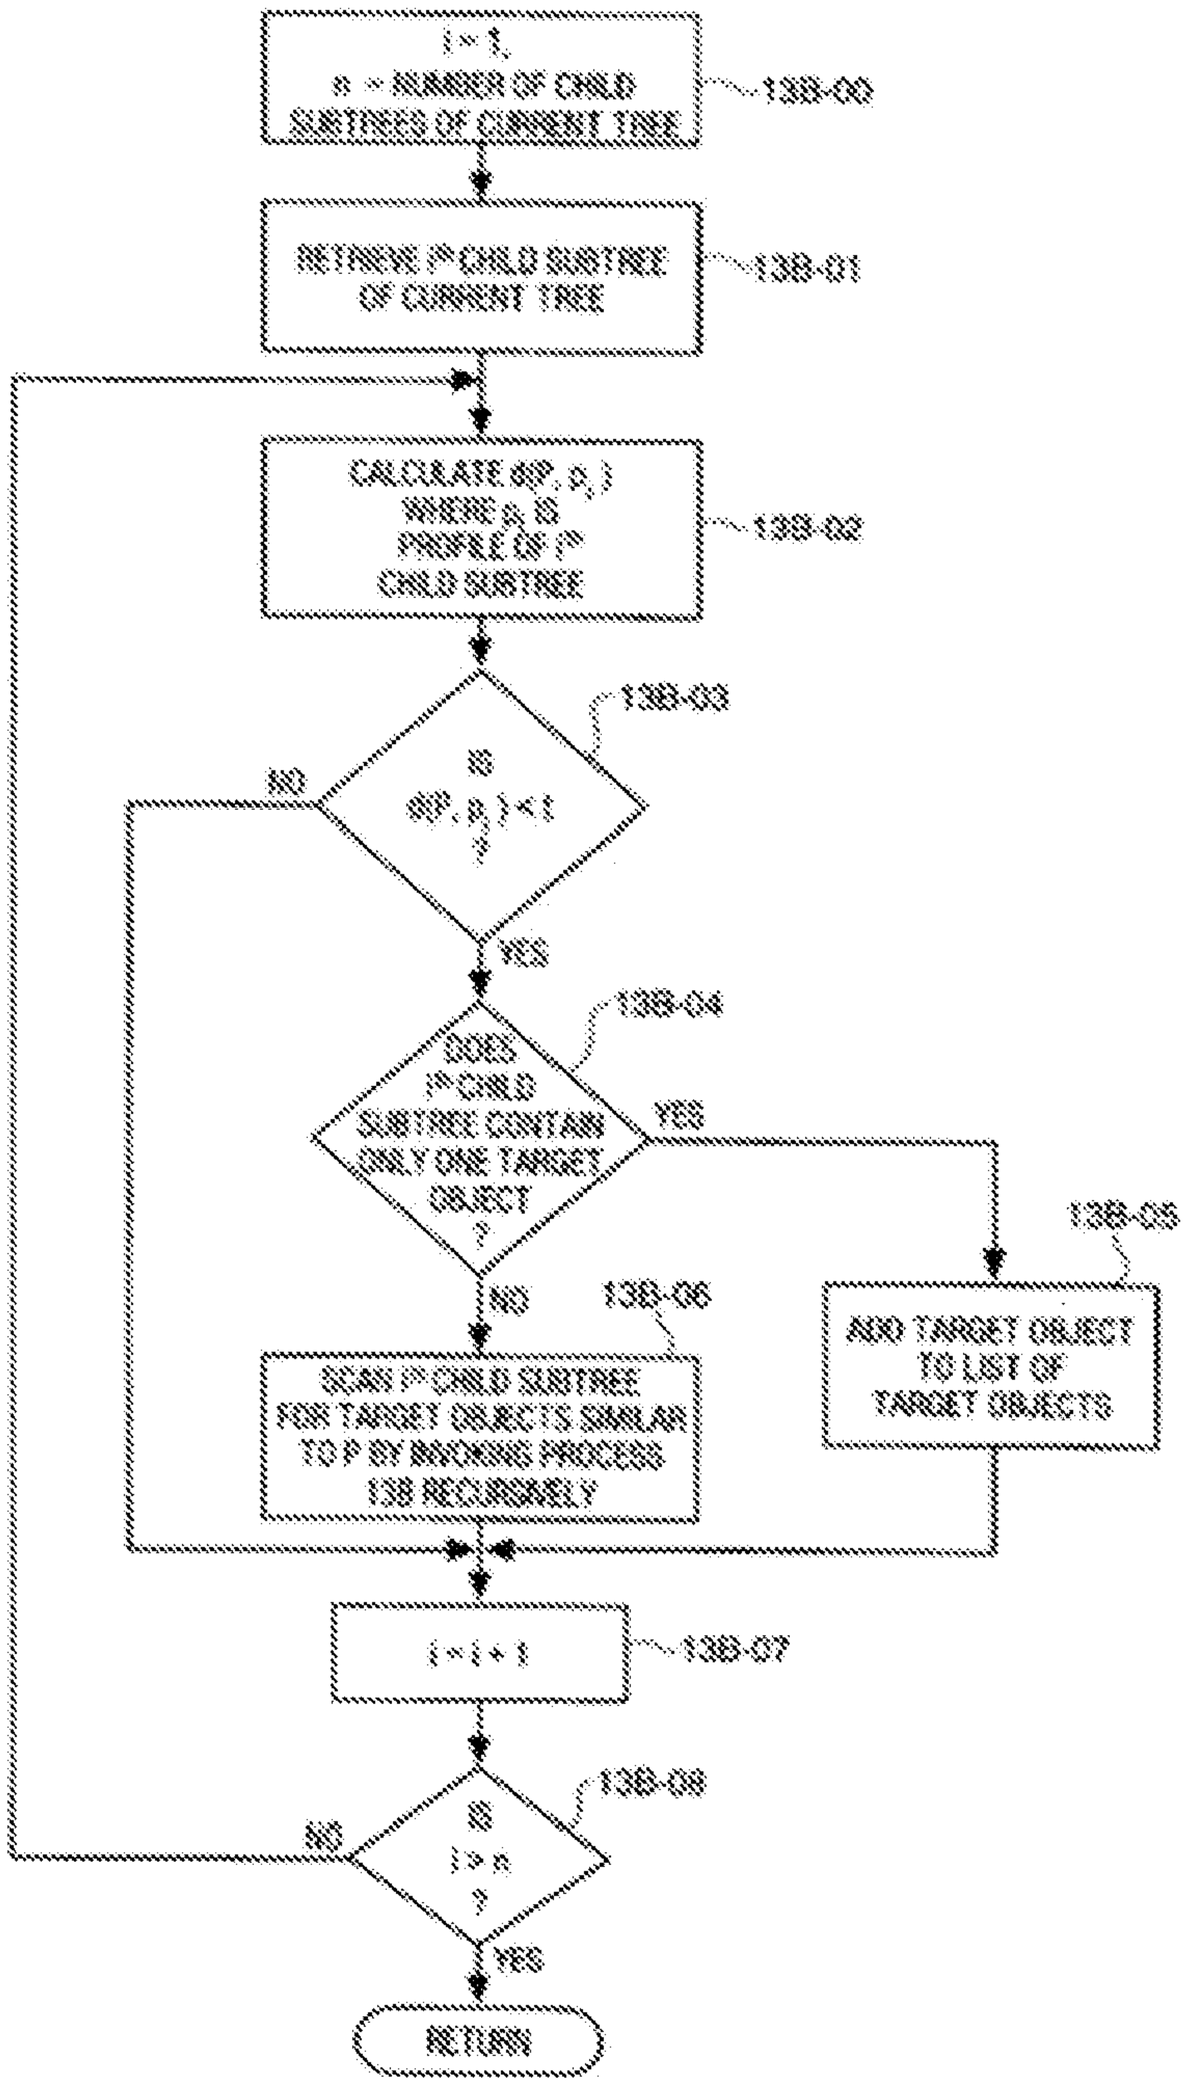
i = 1,
n = NUMBER OF CHILD SUBTREES OF CURRENT TREE
138-00
RETRIEVE iTH CHILD SUBTREE OF CURRENT TREE
138-01
CALCULATE θ(P, p_i) WHERE p_i IS PROFILE OF iTH CHILD SUBTREE
138-02
IS θ(P, p_i) < t ?
138-03
NO
YES
DOES iTH CHILD SUBTREE CONTAIN ONLY ONE TARGET OBJECT ?
138-04
YES
ADD TARGET OBJECT TO LIST OF TARGET OBJECTS
138-05
NO
SCAN iTH CHILD SUBTREE FOR TARGET OBJECTS SIMILAR TO P BY INVOKING PROCESS 138 RECURSIVELY
138-06
i = i + 1
138-07
IS i > n ?
138-08
NO
YES
RETURN

Figure 13B

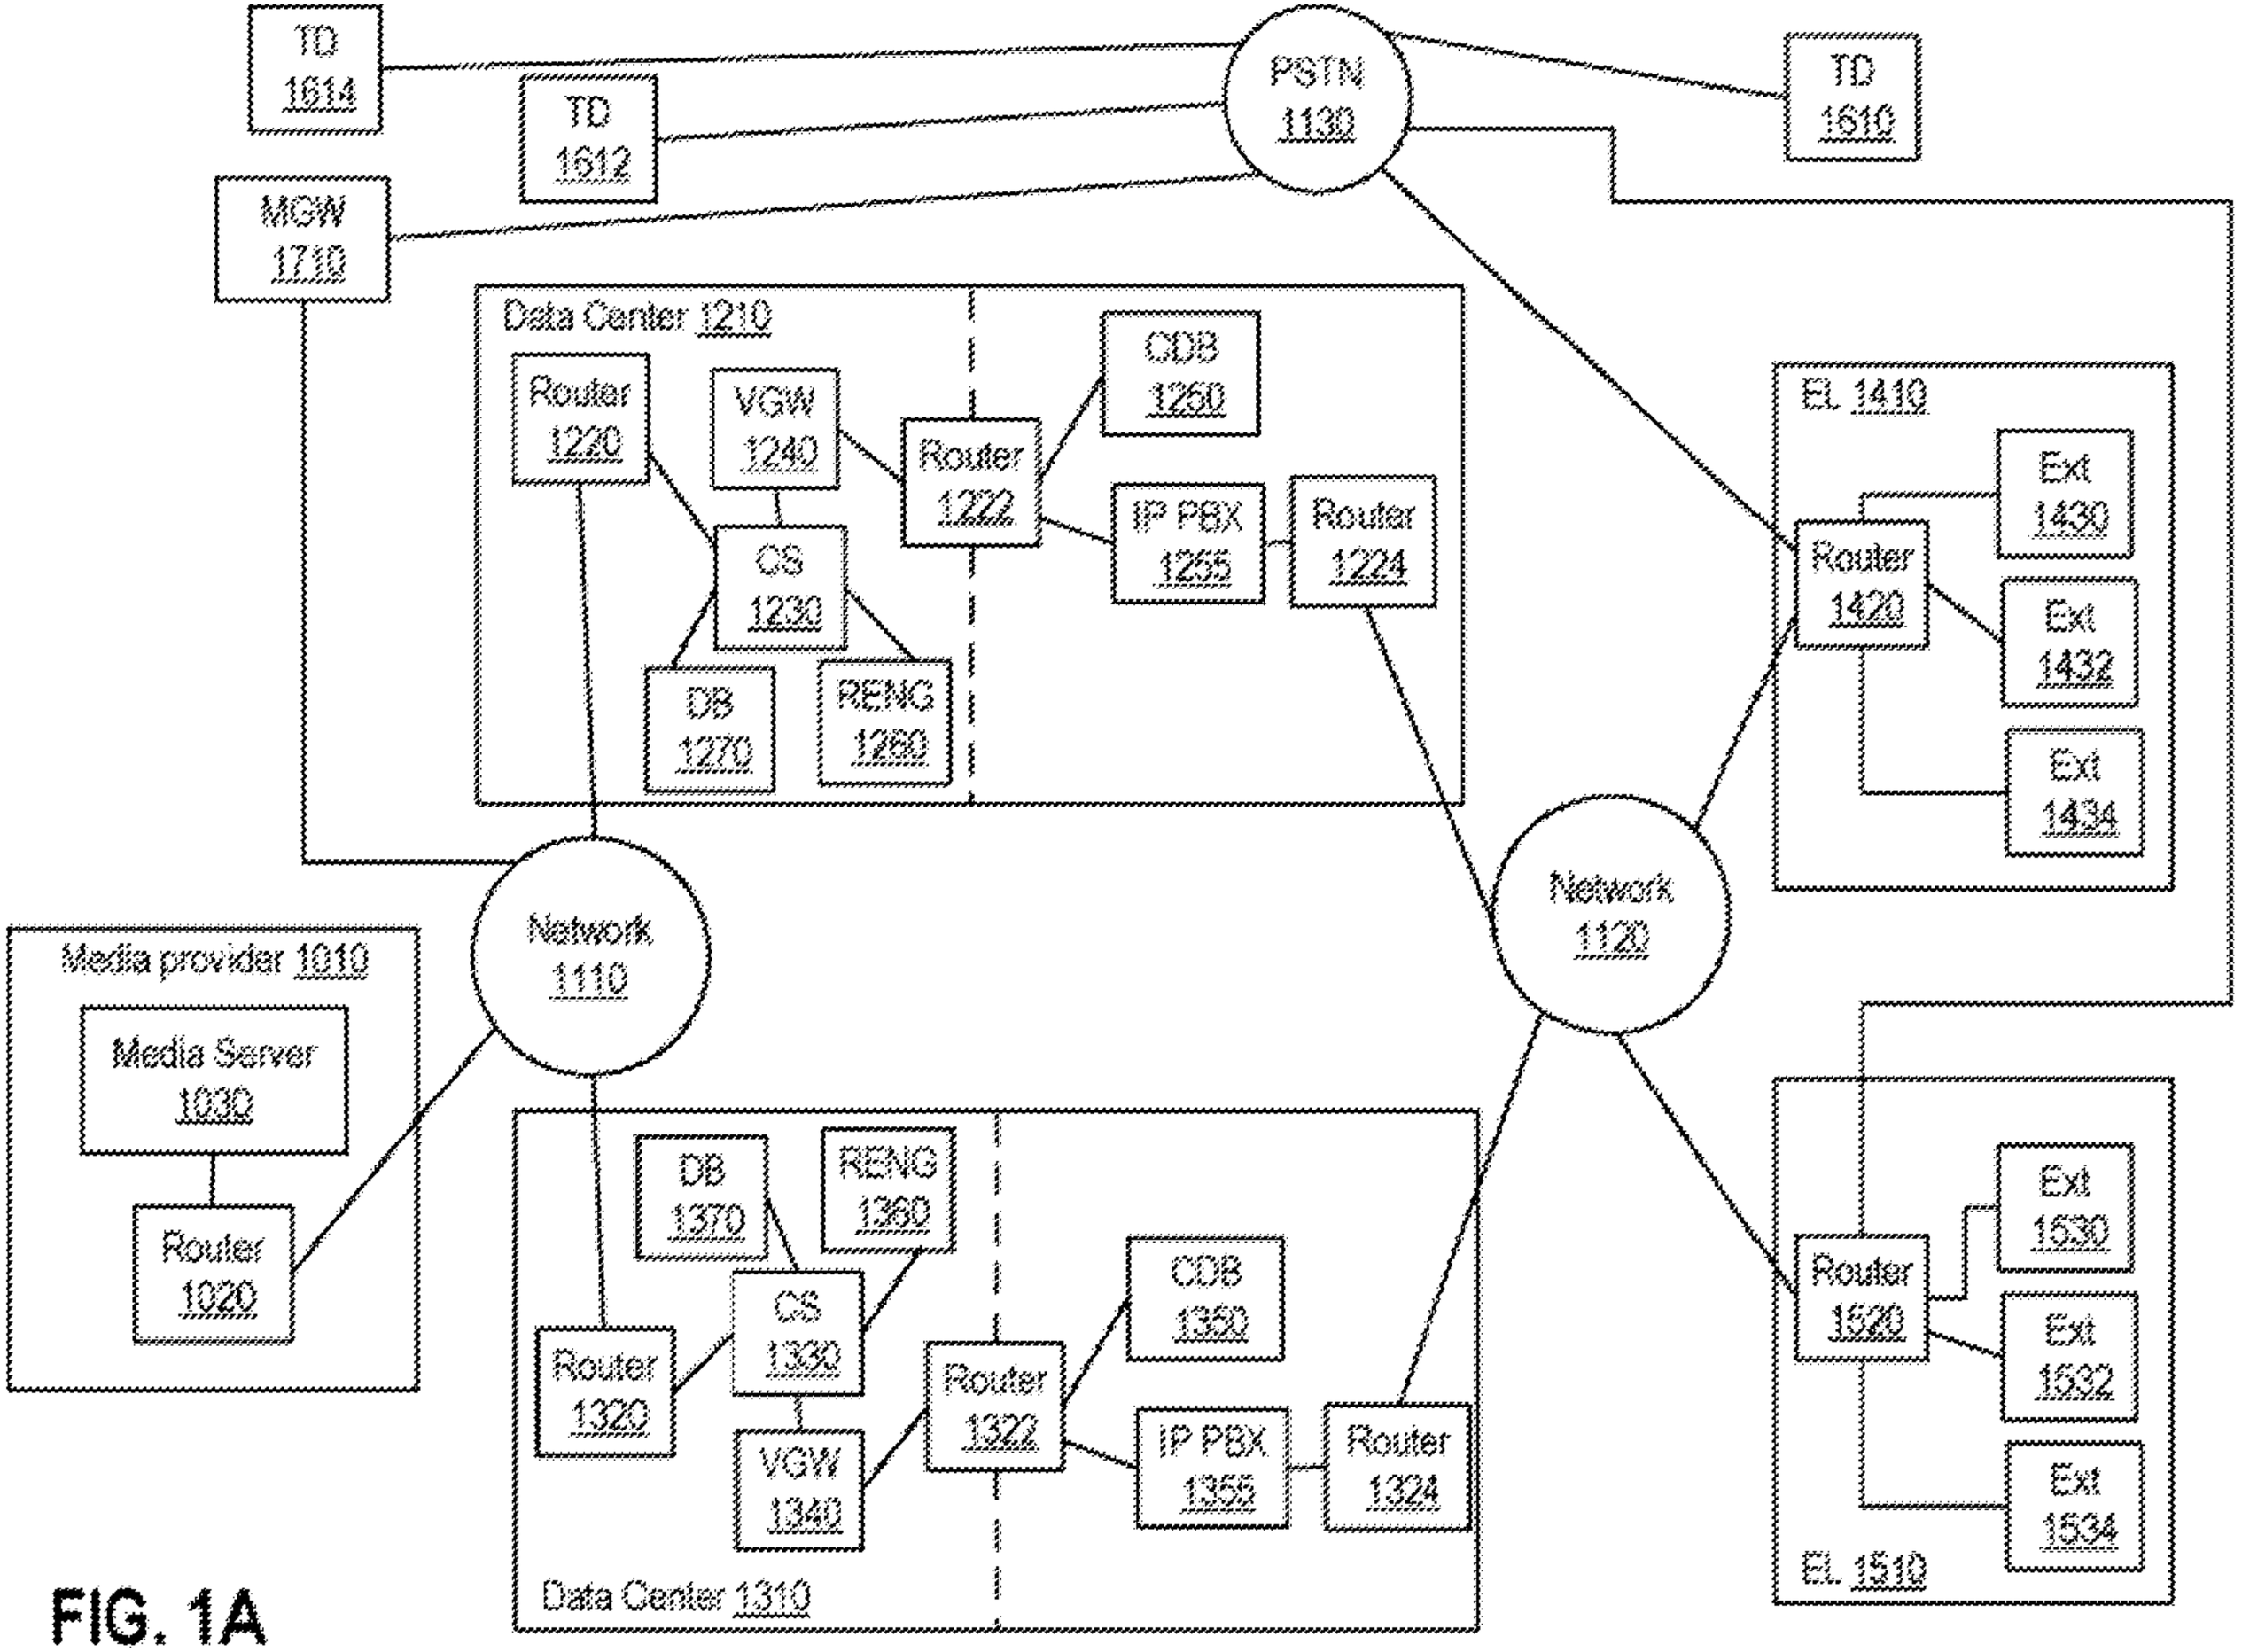
TD
1614
TD
1612
PSTN
1130
TD
1610
MGW
1710
Data Center 1210
Router
1220
VGW
1240
Router
1222
CDB
1260
IP PBX
1255
Router
1224
CS
1230
DB
1270
RENG
1280
EL 1410
Ext
1430
Router
1420
Ext
1432
Ext
1434
Network
1120
Network
1110
Media provider 1010
Media Server
1030
Router
1020
DB
1370
RENG
1360
CS
1330
CDB
1350
Router
1320
Router
1322
VGW
1340
IP PBX
1355
Router
1324
Data Center 1310
Ext
1530
Router
1520
Ext
1532
Ext
1534
EL 1510

FIG. 1A