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

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
12 January 2012 (12.01.2012)

PCT

(10) International Publication Number
WO 2012/006237 A2

(51) **International Patent Classification:**
H04N 21/2668 (201 1.0 1) ***H04N 21/266*** (20 11.0 1)

(21) **International Application Number:**
PCT/US201 1/042786

(22) **International Filing Date:**
1 July 201 1 (01 .07.201 1)

(25) **Filing Language:** English

(26) **Publication Language:** English

(30) **Priority Data:**
12/833,035 9 July 2010 (09.07.2010) US

(71) **Applicant** *(for all designated States except US):* **INTEL CORPORATION** [US/US]; 2200 Mission College Boulevard, Santa Clara, California 95052 (US).

(72) **Inventors; and**
(75) **Inventors/ Applicants** *(for US only):* **YARVIS, Mark D.** [US/US]; 2839 NW 126th Avenue, Portland, Oregon 97229 (US). **GARG, Sharad** [US/US]; 15 117 NW Dane Lane, Portland, Oregon 97229 (US).

(74) **Agents: GARRETT, Patrick** et al; GARRETT IP, LLC, c/o CPA GLOBAL, P.O. Box 52050, Minneapolis, Minnesota 55402 (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:

— *without international search report and to be republished upon receipt of that report (Rule 48.2(g))*

(54) **Title:** SYSTEM AND METHOD FOR PRIVACY-PRESERVING ADVERTISEMENT SELECTION

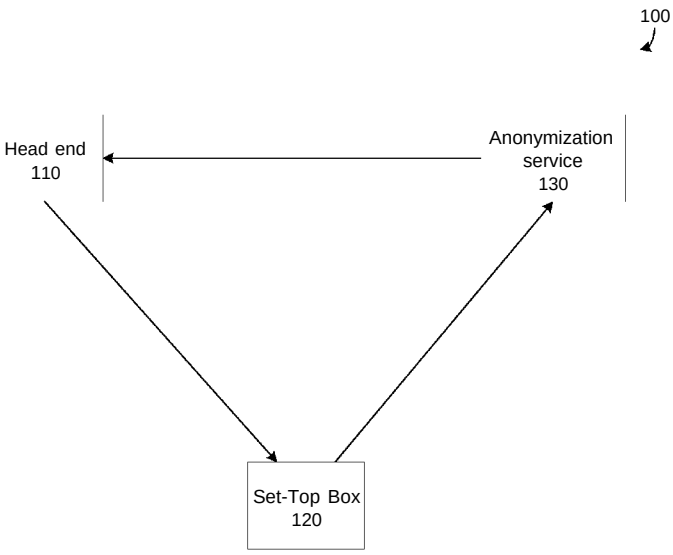

(57) **Abstract:** Methods and systems to provide targeted advertising in a subscription television setting, wherein the households that receive the targeted advertisement may retain anonymity. This may be achieved by marrying an anonymization protocol and infrastructure with a mechanism for STB selection and reporting. This anonymizes the set of households that are selected to play an ad, by ensuring that when an STB meets defined criteria and sends a 'criteria met' message (during the selection process) or 'ad shown' message (during a validation or counting process) to a head end, an STB or household identifier is not revealed, thus anonymizing the household information from the head end. In addition, this allows the head end to control the ad campaign process by advertising only to those households that meet specified criteria, and to count the number of actual instances in which an ad was displayed. The head end is able to count the number of times an ad was displayed, without knowing which particular household(s) viewed the ad.

WO 2012/006237 A2

SYSTEM AND METHOD FOR PRIVACY-PRESERVING ADVERTISEMENT SELECTION

BACKGROUND

Advertisements are sometimes intended for a particular group or demographic of potential consumers. Targeted advertisement in the context of subscription television (e.g., cable or satellite) is often inserted by a user's set-top box (STB) into advertising slots in a live Moving Picture Experts Group (MPEG) transport stream. Often a head end may control which set-top box plays a particular ad. Protocols often claim to be "privacy preserving" because the set-top box makes the selection decision, and the private information used to make that decision may be contained in the STB, rather than at the cable/broadcast head end. However, in such an arrangement private information may be inferred from the communication between the STB and the head end during the processes of (1) selecting households to play an ad and (2) counting the number of times the ad was presented by these STBs.

One protocol for targeted advertising, in which households are selected by the head end to play an ad, is as follows. An ad campaign is provided to the cable/broadcast head end from some external source (such as media buyer or advertiser). That ad campaign has an associated set of criteria that describes the desired viewers (e.g., households with net income greater than $100,000, or families with preschool children) and a total number of households to reach. Each STB contains information about its household (e.g., net income). When an opportunity to show an ad is approaching, the head end sends an "ad opportunity" message to all STBs, where the message contains an identifier for the advertisement (the contents of the ad may be pushed to the STB separately) and also contains the criteria for showing that ad.

The STBs use the household's private information to evaluate the criteria for a match. If there is a match, the STB sends a "criteria met" message to the head end, identifying that STB. The head end selects a subset of the households from those that responded (the subset corresponding to the number of households that need to be reached), and sends a "show ad" message to each selected STB, requesting that the ad be shown. The ads will be locally stored on the STB hard disk (HD) in advance, and shown locally from the HD. To minimize overhead, a distributor of the ad can push ad files onto the STB during low-traffic periods (e.g., in the middle of the night) when bandwidth is least

utilized. After the ad is shown, an "ad shown" message is sent from the STB to the head end, allowing the individual showing of the advertisement to be counted. This is necessary in order to verify that the ad has been shown to the required number of households.

While the above protocol does not require the user's specific private information to leave the STB, the "criteria met" message does inform the head end (any other party who can access this message) that a specific STB meets specific criteria. While explicit private information (e.g., the exact household net income) is not sent from STB to the head end, by sending the "criteria met" message, the STB is providing private information about its specific household to the head end (e.g., by stating that the household meets the criteria, the STB is indicating to the head end that this household's net income is greater than $100,000, for example). Therefore the conventional protocol still reveals sensitive information.

BRIEF DESCRIPTION OF THE DRAWINGS/FIGURES

FIG. 1 is a block diagram showing the topology of a system that may implement an embodiment.

FIG. 2 is a block diagram showing an alternative topology of a system that may implement an embodiment.

FIG. 3A and 3B show a flowchart that illustrates an anonymization protocol, according to an embodiment.

FIG. 4 is a block diagram of illustrating the use of the anonymization protocol in an exemplary topology, according to an embodiment.

FIG. 5 is a flowchart illustrating overall processing of an embodiment.

FIG. 6 is a flowchart illustrating the initial processing by the head end of a request to play an ad, according to an embodiment.

FIG. 7 is a flowchart illustrating the processing by the head end of STB responses, according to an embodiment.

FIG. 8 is a flowchart illustrating the processing by an STB to determine whether to show an ad, according to an embodiment.

FIG. 9 schematically illustrates the use of the anonymization protocol according to an embodiment.

FIG. 10 is a block diagram illustrating a software or firmware embodiment of logic that is executed at a head end.

FIG. 11 is a block diagram illustrating a software or firmware embodiment of logic that is executed at an anonymization node.

FIG. 12 is a block diagram illustrating a software or firmware embodiment of logic that is executed at an STB.

In the drawings, the leftmost digit(s) of a reference number identifies the drawing in which the reference number first appears.

DETAILED DESCRIPTION

A preferred embodiment is now described with reference to the figures, where like reference numbers indicate identical or functionally similar elements. Also in the figures, the leftmost digit of each reference number corresponds to the figure in which the reference number is first used. While specific configurations and arrangements are discussed, it should be understood that this is done for illustrative purposes only. A person skilled in the relevant art will recognize that other configurations and arrangements can be used without departing from the spirit and scope of the description. It will be apparent to a person skilled in the relevant art that this can also be employed in a variety of other systems and applications other than what is described herein.

Disclosed herein are methods and systems to provide targeted advertising in a multimedia setting, such as, for example, subscription television (cable or satellite) or internet content delivery (e.g., hulu.com or youtube.com). The households that receive the targeted advertisement may retain anonymity and prevent the exposure of private information. This may be achieved by using an anonymization protocol and infrastructure with a mechanism for STB selection and reporting. This may anonymize the set of households that are selected to play an ad by ensuring that when an STB meets defined criteria and sends a 'criteria-met' message (during the selection process) or 'ad shown' message (during a validation or counting process) to a head end, an STB or household identifier is not revealed, thus anonymizing the household information from the head end or other parties. In addition, this may simultaneously allow the head end to control the ad campaign process by advertising only to those households that meet specified criteria, and to count the number of actual instances in which an ad was displayed. The head end may also be able to count the number of times an ad is displayed, without knowing which individual households viewed the ad.

An anonymization service may utilize the anonymization protocol to allow the head end to collect anonymous data from STBs during the selection process. A selection

protocol may allow the head end to control which STBs may play a specific ad. The anonymization service may help determine the answer to two questions: (1) How many households meet the criteria for displaying the ad? and (2) How many STBs played a specific ad?

FIG. 1 is a block diagram illustrating components that may be used in an embodiment. A head end 110 is shown in communication with an STB 120. This connection may allow the head end 110 to inform the STB 120 that an opportunity exists to display an advertisement. The head end 110 may also identify the ad and specify the criteria that must be met before STB 120 can show the ad. This channel may also allow the head end to request that the STB 120 show the ad.

If the STB 120 determines that its household meets the stated criteria, then STB 120 may communicate this determination to the head end 110, via an anonymization service 130. The anonymization service 130 may allow the STB 120 to tell the head end 110 that the household of STB 120 meets the criteria, in a manner that preserves the anonymity of the household of STB 120. The process by which the anonymity is preserved will be described in greater detail below.

After the advertisement is shown by STB 120, the fact that the ad was shown may be communicated by the STB 120 to the head end 110, via the anonymization service 130. Because this communication would otherwise reveal that the household of STB 120 meets the stated criteria, the use of service 130 may conceal this fact, as will be described in greater detail below.

While a single STB 130 is shown in figure 1, in embodiments of the invention, one or more STBs may be present in a single household. All such STBs would be in communication with head end 110 and anonymization service 130, in the same manner as STB 120. Moreover, more than one household will typically have STBs; again, each of these STBs, in their respective households, maybe in communication with head end 110 and anonymization service 130 in the same manner as STB 120.

The structure of the head end 110 and the STB 120 may vary in various embodiments, as would be understood by a person of ordinary skill in the art. The anonymization service 130 may be embodied in a network node that includes a server or other device having one or more programmable processors, or otherwise having the circuitry required for connectivity to the head end 110 and the STB 120. Such a network node may be viewed as an anonymization node. As will be described in greater detail

below, anonymization service 130 may also include functionality that implements an anonymization protocol in hardware, software, firmware, or some combination thereof.

In an embodiment of the invention, communications between a head end and an STB may pass through a content distribution network. This is illustrated in the embodiment of FIG. 2. Here, a head end 210 is shown in communication with an STB 220, via a content distribution network 215. In such an embodiment, an anonymization service 230 may be in communication with head end 210. Moreover, the anonymization service 230 may communicate with the STB 230 via the head end 210 and its content distribution network 215. As will be described below, an anonymization protocol may be used to conceal private information of households from the head end 210.

An anonymization protocol is illustrated in FIGs. 3A and 3B, according to an embodiment. Note that in the following description and figures, $E(K_i, x, ...)$ denotes encryption of data field(s) (x,...) using a key K_i. Analogously, $D(K_j, y, ...)$ denotes decryption of data field(s) y using a key K_j.

The protocol as illustrated in FIGs. 3A and 3B may be used to preserve the privacy of households during a targeted advertising process. The steps of the illustrated protocol are shown for the transmission of generic data from STBs to a head end. As will be described below, this data may be an indication as to whether a household meets certain marketing criteria, or may be an indication that an ad was shown in a household, for example.

At 305, an STB may perform a public key encryption process to encrypt data. To perform this encryption, the STB may use the public component of a public key pair, where the corresponding private component is held by the head end. This public encryption key associated with the head end is referred to herein as K_{PuH}, and the associated private component held at the head end is referred to as K_{prH}. The encryption of the data is therefore shown at 305 as

$$E(K_{PuH}, data)$$

i.e., the encryption of data using the key K_{PuH}. Note that all encryptions of data using public keys described below can also be implemented as encryption of a symmetric key using a public key, where the symmetric key is used in turn to encrypt the data.

At 310, the STB may combine this encrypted result with an identifier for the STB (shown as "ID_{STB}"), a timestamp ("time"), and authentication information ("auth_info") in an embodiment. The timestamp may be an indication of the time and date, or may be the

value of a periodically incrementing counter. The authentication information may be a system- wide or local password, a password specific to the STB, or a hash of the data field signed using a private key known to the STB (where the public component of this key is known to the anonymization service). The timestamp and the authentication information may be used at the recipient (the anonymization service) to verify the legitimacy of the communication, as will be described below. The combination of $E(K_{PUH}, data)$, IDSTB, the timestamp, and the authentication information may be represented as a concatenation of the binary representations of these values in an embodiment.

At 315, the STB may encrypt this combination using the public component of another public key pair, $K_{PU}A$. The private component of this key pair, $K_{PR}A$, may be held at the anonymization service. At 315, the STB therefore performs

$$E(K_{PUA}, E(Kp_{uH}, data), ID_{STB}, time, auth\_info).$$

At 320, the STB may send the result to the anonymization service. Note that each STB that received the "ad opportunity" message from the head end may perform the sequence 305-320.

At 325, the anonymization service may decrypt each received message $E(K_{PUA}, E(Kp_{uH}, data), IDSTB, time, auth\_info)$ using its private key $K_{PR}A$. There should be one such message received from every STB that received the "ad opportunity" message from the head end. The anonymization service may therefore perform the operation

$$D(Kp_{rA}, E(Kp_{UA}, E(Kp_{uH}, data), IDSTB, time, auth\_info))$$

for each received message. This may result in the values $E(K_{PUH}, data)$, ID_{STB}, time, auth_info for each STB. Referring now to FIG. 3B, at 360, for each message received from an STB, the anonymization service may then check the timestamp to verify that it is sufficiently recent and does not match the timestamp of any other previously received message. If the timestamp is too old or matches that of a previously received message, then the newly received message may be a duplicate or delayed message, perhaps sent by a hacker or other third party intending to compromise the protocol. The anonymization service may also check the authentication data to verify the legitimacy of the sending party (i.e., the STB).

At 362, the anonymization service may compute a checksum over all messages $E(Kp_{uH}, data)$ received from the respective STBs. The checksum may be a cyclic redundancy code (CRC) or other checksum value. In an alternative embodiment, the

checksum may be the output of a cryptographic function . At 365, the anonymization service may encrypt a new timestamp and and the checksum, using K_{PuH} :

$E(Kp_{uH}, time, checksum)$

At 370, the anonymization service may send $E(K_{PuH}, time, checksum)$ and each $E(Kp_{uH}, data)$ to the head end as a collective set of messages. By sending the full set of messages $E(Kp_{uH}, data)$, one from each STB, the anonymity of each individual STB may be preserved. At 375, the head end may decrypt these values. The head end may therefore perform the decryption

$D(Kp_{rH}, E(Kp_{uH}, time, checksum))$

and recover the time and checksum values. The head end may also perform

$D(K_{PrH}, E(K_{PuH}, data))$

for each $E(K_{PuH}, data)$ and recover the data.

As a result of this protocol, the head end may never see the identity of the source of the data (i.e., the STB and its household). The source may be identified at the anonymization service, but may not be revealed to the head end. Moreover, because the identity of the source may be encrypted between the STB and the anonymization service, an outside observer may likewise be unable to identify the source of the data. In addition, because the data remains encrypted until it is received at the head end, the anonymization service may not determine what information was sent from the STB to the head end.

The flow of information is illustrated in FIG. 4, according to an embodiment. Each of several STBs, STB1,. ..3 (also labeled as 421, 422, …, 423 respectively) may perform a public key encryption process to encrypt the data. To perform this encryption, the STBs may use the public component K_{PrH} of a public key pair, where the corresponding private component Kp_{uH} may be held by the head end 410. The result of the encryption of the data is shown as

$E(K_{PuH}, data_i)$.

Each STB may combine this encrypted result with an identifier for the STB (shown as "IDsTBi"), [a] timestamp, and authentication information in an embodiment. The timestamp and the authentication information may be used at the recipient (the anonymization service 430) to verify the legitimacy of the communication.

Each STB may encrypt this combination using the public component $K_{J>UA}$. The private component Kp_{rA} of this key pair may be held at the anonymization service 430. Each STBi may therefore perform

$$E(K_{PUA}, E(Kp_{uH}, data;), ID_{STBI}, time, auth\_info).$$

Each STB may then send its result to the anonymization service 430.

The anonymization service 430 may decrypt the received messages using its private key K_{prA}. This results in the values $E(K_{PuH}, data;)$, $ID_{sTBi,}$ time, and auth_info. The anonymization service 430 may then check the timestamp to verify that it does not match the timestamp of any other previously received message. The anonymization service 430 may also check the authentication data to verify the legitimacy of each sending party (i.e., STBi).

The anonymization service 430 may receive the results $E(Kp_{uH}, data;)$ from many nodes, one from each STBi. After a timeout, the anonymization service may gather up all of these results, and assemble them in a random order. The anonymization service 430 may also encrypt a new timestamp and a checksum of the list of results, using K_{PuH} :

$$E(Kp_{uH}, time, checksum)$$

The anonymization service 430 may send $E(Kp_{uH}, time, checksum)$ and the randomly ordered list containing each $E(Kp_{uH}, data;)$ to the head end 410. The head end 410 may then decrypt these values. The head end 410 may therefore perform the decryption

$$D(K_{prH}, E(Kp_{uH}, time, checksum))$$

and recover the time and checksum values. The head end 410 may check that the checksum is consistent with the data received, and that the timestamp does not match those of previously received messages. The head end may also perform

$$D(K_{prH}, E(Kp_{uH}, data_i))$$

and recover each segment data;.

In an embodiment, each STB may provide entertainment and advertising content to its respective household. These homes are shown in FIG. 4 as households 471,. .., 473. Generally, an STB may access or store information that requires privacy, e.g., household income, demographics of household members, geographical location of the home, etc. It is this household information that may be used to match the criteria for the ad. The ad may be targeted at households with a certain income level, or households having residents of a certain age group, for example. If the criteria is met, then the STB; may state in data; that its household meets the criteria and is eligible to view the ad.

In an embodiment, the head end 410 may be part of the infrastructure for a content distributor 460. Content distributor 460 may be a cable company, a satellite television company, or an internet content delivery provider, for example. The anonymization

service 430 may be managed by a third party 480, independent of the content distributor 460.

The anonymization protocol discussed above may be used as illustrated in FIG. 5, according to an embodiment. At 510, a request to distribute an advertisement may be received and processed by the head end. This request may be accompanied by criteria for display, such that the ad is targeted to households meeting these criteria, as well as the number of households that should be reached. The head end may send the criteria and an identifier for the ad to the STBs in an "ad opportunity" message. In an embodiment, the actual advertisement (e.g., and audio/visual file) may have previously been distributed to all STBs and stored there. The head end therefore can send the ID of the ad to the STBs, without having to send the actual ad.

At 520, all STBs may examine the criteria in the "ad opportunity" message, compare it to information about the STB's household stored on the STB, and determine whether the criteria is met. It then may reply to the head end via the anonymization service by sending a "criteria met" message. To do this, the anonymization protocol discussed above with respect to FIGs. 3A, 3B, and 4 may be used. In this case, the data transmitted by each STB to the head end may include an indication as to whether the criteria is met by the household of the STB.

At 530, the head end may process the anonymized "criteria met" responses and determine the number of STBs that have responded in the affirmative, i.e., indicating that their respective households meet the criteria. The head end may then broadcast, to all the STBs, a request to show the ad. In an embodiment, this message may be encrypted by the head end. In such an embodiment, the encryption may be performed using a private key held by the head end, where the corresponding public key, needed for decryption at each STB, is publicly available. Such an arrangement would prevent an unauthorized party from posing as a head end for purposes of this transmission.

At 540, each STB may make a determination as to whether the ad is to be shown. As will be described in greater detail below, this determination may include but is not limited to a determination as to whether the criteria is met. If the STB determines that the ad is to be shown, then the ad may be shown.

At 550, each STB may send the identifier of the ad and an indication of whether or not the ad was played to the head end in an "ad shown" message, again using the anonymization service and protocol described above. For each message sent by the STBs,

the data field contains an indication as to whether the ad was shown, instead of an indication of whether the criteria is met. At 560, the head end may count the number of households to which the advertisement was shown.

The processing of a request to distribute an advertisement (510 of FIG. 5) is illustrated in greater detail in FIG. 6, according to an embodiment. At 610, a request may be received from an external source at the head end, asking that an advertisement be played. In an embodiment, this request may include an identifier for the ad ("ad ID"), the criteria for households where the ad is to be played (e.g., particular household demographics or income level), and the number of households for which the ad is requested to be played ("N").

In some circumstances, it may be desirable to play the ad in as many households as possible. In this case, N will have an unlimited value. At 620, a determination may be made as to whether N has an unlimited value. If N is not unlimited and has a finite value, then at 640, the criteria, N, and the ad identifier may be cached at the head end, and at 650 a message may be sent to all STBs, where the message includes the ad identifier and the criteria. This message may state that there is an opportunity to show the identified advertisement for households meeting the stated criteria.

If N has an unlimited value, then at 630 the head end may send a message to all STBs telling them to show the ad if they meet the criteria, specifying the criteria, and defining the value $\mathbf{P_{show}} = 1$. The use of the value $\mathbf{P_{show}}$ will be described in greater detail below. In this context, the setting of this value to 1 effectively tells each STB that it is eligible to show the ad if it meets the criteria.

The processing, by the head end, of STB responses (530 of FIG. 5) is illustrated in greater detail in FIG. 7, according to an embodiment. At 710, the head end may receive one or more messages from respective STBs indicating that the criteria is met by their respective households. Note that while the head end receives these messages, it never knows the identities of the specific households or STBs that have responded. At 720, the head end may decrypt these messages using the private component of its public key pair, as described above. At 730, the head end may count the number of affirmative responses to learn the total number M of households that meet the criteria. Negative responses may be ignored.

At 740, a determination may be made as to whether $M < N$. If not, then at 750 the head end may define $\mathbf{P_{show}} = N/M$, where N is the number of households are required or

desired by the advertiser to see the ad. This value represents the proportion of eligible households (i.e., those that meet the criteria) that may be required to see the ad, in order for the goal of N to be reached. If the condition of 740 is true, then the number of households required to see the ad meets or exceeds the number of households satisfying the criteria. In this case, all eligible households may be shown the ad. This is made possible by setting $P_{show}=1$, as will be described below.

At 770, a message may be sent to all STBs, telling them that if their respective households meet the criteria, then they are eligible to show the ad, pending further processing at the STB involving P_{show}, as will be described below. In an embodiment, this message may include an identifier for the ad, the criteria, and P_{show}. Further, in an embodiment, these values may be sent in encrypted form, encrypted using the private component of the key pair of the head end, K_{PrH}.

The determination of whether to show the advertisement (540 of FIG. 5) is illustrated in greater detail in FIG. 8, according to an embodiment. At 805, the STB may receive a "show ad" message from the head end including an add ID, criteria, and P_{show}. At 810, a determination may be made at the STB as to whether the household meets the criteria. If not, then at 820 the STB may prepare a negative "ad shown" message to be sent to the head end via the anonymization protocol; this message may be sent to the head end via the anonymization protocol at 550. Otherwise, at 830, a random number P may be generated, where $0 < P < 1$. In an embodiment, the random number P may be produced using a noise source; alternatively, P may be generated using a deterministic pseudorandom data source. At 840, a determination may be made as to whether $P < P_{show}$. If so, then at 850 the advertisement may be shown. If not, then at 820 the STB may prepare a negative "ad shown" message to the head end via the anonymization protocol. At 860, a message may be prepared, saying that the ad was shown. This message may be sent to the head end via the anonymization protocol at 550.

Recall that if N is unlimited (meaning that as many households as meet the criteria are to be shown the ad), the head end defined $P_{show}=1$. If Pshow was defined by the head end to equal 1, then for every household meeting the criteria (at 810), the STB generates a random P (at 830) wherein $P < P_{show}$, and the ad may be shown in every such household at 850.

Where N is not unlimited, then $P_{show} = N/M$ as defined at the head end. By having each STB choose a random P (at 830) and determining if $P < P_{show}$ (at 840), the expected number of households showing the ad (at 850) will be equal to N.

Note that in certain embodiments, the value of P_{show} as derived at a head end may be varied somewhat from the calculation shown at 750. For example, if the head end wishes to decrease the chances that the ad is displayed to fewer than N households, then it may increase the value of P_{show} by a small amount. This may result in a larger set of STBs showing the ad. Analogously, if the head end wishes to decrease the chances that the ad is displayed to more than N households, then it may decrease the value of P_{show} by a small amount.

FIG. 9 schematically illustrates the anonymization protocol as it is applied above, according to an embodiment. At 910, the head end may issue a message to all STBs, announcing the opportunity to show an advertisement. This message may include the ID for the ad, plus the criteria for showing the ad. At 920, each STB may respond with a message that states whether the criteria is met by the particular household. These messages may be sent to the anonymization service. Each message may include an encrypted indication of whether the criteria is met, where the encryption may be performed using the public component Kp_{uH} of an asymmetric key pair. The private component of this key pair may be held only by the headend. This encrypted indication is shown as $E(K_{PuH}$, criteria_met_y/n), where criteria_met_y/n occupies the data field shown in FIG. 3A and B. This encrypted indication may be further encrypted, along with an ID of the STB (shown as STBi), a timestamp, and authentication information, using the public component Kp_{UA} of another asymmetric key pair. The private component of this latter public key pair may be held by the anonymization service. The message is therefore shown as $E(K_{PuA}$, $E(K_{PuH}$, criteria_met_y/n), $ID_{STBi,}$ time, auth_info), where such a message may be generated by each STBi for transmission to the anonymization service. While the figure shows three such messages for different STBs, it is to be understood that the total number of STBs (and associated messages) may be greater or fewer.

The anonymization service may decrypt each such message using its private key, to recover the ID of the STBi, the time stamp, and the authentication information. The latter two values may be checked by the anonymization service as described above. This decryption may also recover the encrypted indication as to whether the criteria is met by STB;, $E(Kp_{uH}$, criteria_met_y/n). This encrypted indication may then be combined with

the indication received from each STBi and forwarded by the anonymization service, as a set, to the head end at 930. In an embodiment, the ordering of the set of encrypted indications may be randomized before sending to the head end. The anonymization service may also send the ID of the advertisement, a timestamp, and a checksum that may be a function of the encrypted indications. These latter three values may be encrypted using K_{puH}.

The head end may then decrypt the received messages using its private key, to identify the ad and to count the number of STBs at which the criteria is met. The head end never knows the identities of the specific STBs which meet the criteria. The checksum may also check the received timestamp and checksum values. The head end may calculate a value $\mathbf{P_{show}}$ as described above, based in part on the number of households indicating that the criteria are met. At 940, the head end may then send a message to each STB, instructing it to show the ad if the criteria is met and if the values P (generated at the STBs as described above) permit. As described above, this message may be encrypted by the head end using K_{PrH}. In this case, each STB would decrypt this message using K_{PuH}.

Some STBs will then show the ad while others may not. At each STB an encrypted indication may be generated, identifying the ad and stating whether the ad was shown. This indication may be encrypted using K_{PuH}, and is shown as $E(K_{PuH}$, ad ID, ad_shown_y/n). Here, the fields ad ID and ad_shown_y/n collectively represent the field of data illustrated in FIG. 3A and B. The particular STB is not identified in this indication. This indication may then be further encrypted using K_{PUA}, along with the identifier of the STB, a timestamp, and authentication information. This message is therefore shown as $E(K_{PUA}, E(K_{PuH}$, ad ID, ad_shown_y/n), IDS_{TBI}, time, auth_info). At 950, each STB may send such a message to the anonymization service, incorporating the appropriate indication of whether the ad was shown by the STB.

The anonymization service may then decrypt the messages to recover the ID of the STBi, the time stamp, and the authentication information. The latter two values may be checked as described above. This decryption may also yield the encrypted indication as to whether the ad was shown, $E(K_{PuH}$, ad_shown_y/n). This encrypted indication may then be combined with indications from each STBi and forwarded by the anonymization service collectively to the head end at 960. The anonymization service may also send the ID of the advertisement, a timestamp, and a checksum that may be a function of the encrypted indications. These latter three values may be encrypted using K_{PuH}.

The head end may then decrypt the received messages using its private key, to identify the ad and to process the checksum and timestamp. The head end can then count the number of STBs that showed the advertisement.

One or more features disclosed herein may be implemented in hardware, software, firmware, or combinations thereof, including discrete and integrated circuit logic, application specific integrated circuit (ASIC) logic, programmable gate arrays, and/or microcontrollers, or may be implemented as part of a domain-specific integrated circuit package, or a combination of integrated circuit packages. The term software, as used herein, refers to a computer program product including a computer readable medium having computer program logic stored therein to cause a computer system to perform one or more features and/or combinations of features disclosed herein.

Software or firmware embodiments are illustrated in the context of computing systems shown in FIGs. 10-12. In FIG. 10, system 1000 may be located at a head end, and may include a processor 1020 and a body of memory 1010 that may include one or more computer readable media that may store computer program logic 1040. Memory 1010 may be implemented as a hard disk and drive, a removable media such as a compact disk and drive, a read-only memory (ROM) or random access memory (RAM) device, for example. Processor 1020 and memory 1010 may be in communication using any of several technologies known to one of ordinary skill in the art, such as a bus. Computer program logic 1040 contained in memory 1010 may be read and executed by processor 1020. One or more I/O ports and/or I/O devices, shown collectively as I/O 1030, may also be connected to processor 1020 and memory 1010.

Computer program logic 1040 may include initial request processing logic 1050. This logic may be responsible for processing an initial request for the playing of an advertisement, as received from an advertiser or content provider at the head end. As noted above, such an initial request may also include an identifier for the specific advertisement, criteria for households that may play the ad, and a value N representing the number of households that are to view the ad. Initial request processing logic 1050 may be responsible for instructing STBs to show the ad when N is unlimited, and otherwise informing STBs that there is an opportunity to show this ad.

Computer program logic 1040 may also comprise STB response processing logic 1060. This body of logic may be responsible for receiving messages from STBs indicating that criteria have been met, decrypting these messages, counting the number of affirmative

responses, computing P_{show} as necessary, and instructing STBs to show the advertisement. Note that in alternative embodiments, the decryption process may be controlled by logic 1060, where the actual decryption may be performed by one or more hardware components.

Computer program logic 1040 may also comprise statistics collection logic 1070. Logic 1070 may be responsible for receiving messages from STBs indicating whether the ad has been shown, and counting the total number of households that have been shown the ad.

In alternative embodiments, the computer program logic 1040 may be organized differently in order to implement the processing described herein. Different logic modules may be used instead of those shown in FIG. 10. Moreover, in alternative embodiments, more or fewer logic modules may be used, as would be understood by those of ordinary skill in the art.

Referring to FIG. 11, system 1100 may be located at an anonymization service, and may include a processor 1120 and a body of memory 1110 that may include one or more computer readable media that may store computer program logic 1140. Memory 1110 may be implemented as a hard disk and drive, a removable media such as a compact disk and drive, or a read-only memory (ROM) or random access memory (RAM) device, for example. Processor 1120 and memory 1110 may be in communication using any of several technologies known to one of ordinary skill in the art, such as a bus. Computer program logic 1140 contained in memory 1110 may be read and executed by processor 1120. One or more I/O ports and/or I/O devices, shown collectively as I/O 1130, may also be connected to processor 1120 and memory 1110.

Computer program logic 1140 may include decryption logic 1150. This logic may be responsible for decrypting the encrypted messages coming from the STBs, as described above. Computer program logic 1140 may also include verification logic 1160, which may be responsible for verifying the data included in those messages, such as the authentication data, and checking the timestamp as discussed above. Computer program logic 1140 may also comprise encryption logic 1170. Logic 1170 may be responsible for encrypting messages to be sent to the head end, as described above.

In alternative embodiments, the computer program logic 1140 may be organized differently in order to implement the processing described herein. Different logic modules may be used instead of those shown in FIG. 11. Moreover, in alternative embodiments,

more or fewer logic modules may be used, as would be understood by those of ordinary skill in the art.

Moreover, while the embodiment of FIG. 11 suggests that encryption and decryption may be performed in software or firmware, in alternative embodiments encryption and decryption may be performed in hardware. In such embodiments, the control of the encryption and decryption operations may be performed by processor 1120 as directed by computer program logic 1140.

Referring to FIG. 12, system 1200 may be located at an STB, and may include a processor 1220 and a body of memory 1210 that may include one or more computer readable media that may store computer program logic 1240. Memory 1210 may be implemented as a hard disk and drive, a removable media such as a compact disk and drive, or a read-only memory (ROM) or random access memory (RAM) device, for example. Processor 1220 and memory 1210 may be in communication using any of several technologies known to one of ordinary skill in the art, such as a bus. Computer program logic 1240 contained in memory 1210 may be read and executed by processor 1220. One or more I/O ports and/or I/O devices, shown collectively as I/O 1230, may also be connected to processor 1220 and memory 1210.

Computer program logic 1240 may include household evaluation logic 1250. This logic may be responsible for comparing information about the household to the criteria for showing the advertisement, in order to determine if the criteria are met. Computer program logic 1240 may also include encryption logic 1260, which may be responsible for encrypting messages indicating whether the criteria are met and messages indicating whether an ad has been shown, for example, as described above. Computer program logic 1240 may also comprise decryption logic 1270. Logic 1270 may be responsible for decrypting messages, such as the message from the head end telling the STB to show the advertisement.

Computer program logic 1240 may also include randomization logic 1280, which may be responsible for determining a random value between 0 and 1, as described above. Alternatively, in embodiments the randomization process may use hardware for generating the random value, in which case logic 1280 may be responsible for controlling or sampling the randomization hardware. Computer program logic 1240 may also include display decision logic 1290, which may be responsible for determining whether to show an ad, given the random value P and the received value $\mathbf{P_{show}}$, as described above.

In alternative embodiments, the computer program logic 1240 may be organized differently in order to implement the processing described herein. Different logic modules may be used instead of those shown in FIG. 12. Moreover, in alternative embodiments, more or fewer logic modules may be used, as would be understood by those of ordinary skill in the art.

Moreover, while the embodiment of FIG. 12 suggests that encryption and decryption may be performed in software or firmware, in alternative embodiments encryption and decryption may be performed in hardware. In such embodiments, the control of the encryption and decryption operations may be performed by processor 1220 as directed by computer program logic 1240.

Methods and systems are disclosed herein with the aid of functional building blocks illustrating the functions, features, and relationships thereof. At least some of the boundaries of these functional building blocks have been arbitrarily defined herein for the convenience of the description. Alternate boundaries may be defined so long as the specified functions and relationships thereof are appropriately performed.

While various embodiments are disclosed herein, it should be understood that they have been presented by way of example only, and not limitation. It will be apparent to persons skilled in the relevant art that various changes in form and detail may be made therein without departing from the spirit and scope of the methods and systems disclosed herein. Thus, the breadth and scope of the claims should not be limited by any of the exemplary embodiments disclosed herein.

CLAIMS

WHAT IS CLAIMED IS:

1. A method, comprising:

receiving, at an anonymizing service node, information from a plurality of set top boxes (STBs) encrypted by each STB using a public component of a first public key pair, the information from each STB including

an identifier of the STB,

an identifier of an advertisement, and

an indication as to whether one or more criteria for display of the ad are met, wherein the indication is encrypted with the public component of a second public key pair;

decrypting the encrypted information from each STB, to recover the identifier of each STB;

encrypting the identifier of the advertisement using the public component of the second public key pair; and

forwarding all encrypted indications and the encrypted identifier of the advertisement together to a head end, enabling the head end to count the number M of STBs at which the criteria is met,

wherein the private component of the first public key pair is held at the anonymizing service node and is not available to the head end and the private component of the second public key pair is held at the head end and not available to the anonymizing service.

2. The method of claim 1, wherein the information from each STB further includes a first timestamp and authentication data, such that the first timestamp and authentication data are also encrypted using the public component of the first public key pair; and

wherein the method of claim 1 further comprises:

after decrypting the encrypted information from each STB, verifying the authentication data and determining whether the first timestamp matches any previously received timestamps.

3. The method of claim 1, further comprising:

receiving, at the anonymizing service node, further information from each STB, encrypted using the public component of the first public key pair, the further information including

the identifier of the STB, and

an indication of whether the advertisement was displayed at the STB,

encrypted with the public component of the second key pair ,

decrypting the encrypted further information from each STB, to recover the identifier of each STB, said decryption performed using the private component of the first public key pair;

encrypting the identifier of the advertisement with the public component of the second public key pair; and

forwarding all encrypted indications of whether the advertisement was displayed and the encrypted identifier of the advertisement together to the head end, enabling the head end to count the number of STBs that displayed the advertisement.

4. The method of claim 3, wherein the further information includes authentication data and a second timestamp, such that the authentication data and the second timestamp were also encrypted using the public component of the first public key pair; and

wherein the method of claim 3 further comprises:

after decrypting the encrypted further information, verifying the authentication data and determining whether the first timestamp matches any previously received timestamp.

5. The method of claim 4, further comprising:

calculating a checksum on the basis of the STBs' respective indications of whether the advertisement was displayed;

encrypting the checksum and said second timestamp using the public component of the second public key pair; and

forwarding the encrypted checksum and second timestamp to the head end along with each of the STBs' respective indications of whether the advertisement was displayed.

6. The method of claim 3, wherein the indication of whether the advertisement was displayed is generated at each STB,

wherein a decision regarding whether to display the advertisement is made at each STB in response to a command issued to each STB from the head end, wherein the command includes a parameter $P_{show} = N/M$,

where N is the number of households desired to see the advertisement and M is the number of households indicating that the criteria is met, as counted by the headend;

wherein after receiving the command, each STB creates a random number between 0 and 1 and displays the ad only if the random number is less than or equal to P_{show}.

7. The method of claim 1, wherein the STB creates the indication that the criteria is met, in response to a broadcast from the head end identifying an opportunity to display the advertisement and specifying the criteria for displaying the advertisement.

8. A system, comprising:

a processor; and

a memory in communication with said processor, said memory for storing a plurality of processing instructions for directing said processor to:

receive information from a plurality of set top boxes (STBs) encrypted by each STB using a public component of a first public key pair, the information from each STB including

an identifier of the STB,

an identifier of an advertisement, and

an indication as to whether one or more criteria for display of the advertisement are met, wherein the indication is encrypted with the public component of a second public key pair;

decrypt the encrypted information from each STB, to recover the identifier of each STB;

encrypt the identifier of the advertisement using the public component of the second public key pair; and

forward all encrypted indications and the encrypted identifier of the advertisement together to a head end, enabling the head end to count the number M of STBs at which the criteria are met,

wherein the private component of the first public key pair is held at the anonymizing service node and is not available to the head end, and the private component of the second public key pair is held at the head end and not available to the anonymizing service.

9. The system of claim 8, wherein the information from each STB further includes a first timestamp and authentication data, such that the first timestamp and authentication data are also encrypted using the public component of the first public key pair; and

wherein said memory further stores processing instructions for directing said processor to:

after decrypting the encrypted information from each STB, verify the authentication data and determine whether the first timestamp matches any previously received timestamps.

10. The system of claim 8, wherein said memory further stores processing instructions for directing said processor to:

receive, at the anonymizing service node, further information from each STB, encrypted using the public component of the first public key pair, the further information including

the identifier of the STB, and

an indication of whether the advertisement was displayed at the STB, encrypted with the public component of the second key pair;

decrypt the encryptedfurther information from each STB, to recover the identifier of each STB, said decryption performed using the private component of the first public key pair;

encrypt the identifier of the advertisement with the public component of the second public key pair; and

forward all encrypted indications of whether the advertisement was displayed and the encrypted identifier of the advertisement together, enabling the head end to count the number of STBs that displayed the advertisement.

11. The system of claim 10, wherein the indication further information includes authentication data and a second timestamp, such that the authentication data and the second timestamp have also been encrypted using the public component of the first public key pair; and

wherein said memory further stores processing instructions for directing said processor to:

after decrypting the encrypted further information, verify the authentication data and determining whether the second timestamp matches any previously received timestamp.

12. The system of claim 11, wherein said memory further stores processing instructions for directing said processor to:

calculate a checksum on the basis of the STBs' respective indications of whether the advertisement was displayed;

encrypt the checksum and the second timestamp using the public component of the second public key pair; and

forward the encrypted checksum and second timestamp to the head end along with each of the STBs' respective indications of whether the advertisement was displayed.

13. The system of claim 10, wherein the indication of whether the advertisement was displayed is generated at each STB,

wherein a decision regarding whether to display the advertisement is made at each STB in response to a command issued to each STB from the head end, wherein the command includes a parameter $P_{show} = N/M$,

where N is the number of households desired to see the advertisement and

M is the number of households indicating that the criteria is met, as counted by the head end; and

wherein after receiving the command, each STB creates a random number between 0 and 1 and displays the advertisement only if the random number is less than or equal to $\mathbf{P_{sh0}}_{w}$.

14. The system of claim 8, wherein the STB creates the indication that the criteria is met, in response to a broadcast from the head end identifying an opportunity to display the advertisement and specifying the criteria for displaying the advertisement.

15. A computer program product comprising a computer useable medium having computer program logic stored thereon, the computer control logic comprising:

logic configured to cause a processor to receive information from a plurality of set top boxes (STBs) encrypted by each STB using a public component of a first public key pair, the information from each STB including

an identifier of the STB,

an identifier of an advertisement, and

an indication as to whether one or more criteria for display of the advertisement are met, wherein the indication is encrypted with the public component of a second public key pair;

logic configured to cause the processor to decrypt the encrypted information from each STB, to recover the identifier of each STB;

logic configured to cause the processor to encrypt the identifier of the advertisement using the public component of the second public key pair; and

logic configured to cause the processor to forward all encrypted indications and the encrypted identifier of the advertisement together to a head end, enabling the head end to count the number M of STBs at which the criteria are met,

wherein the private component of the first public key pair is held at the anonymizing service node and is not available to the head end and the private component of the second public key pair is held at the head end and not available to the anonymizing service.

16. The computer program product of claim 15, wherein the information from each STB further includes a first timestamp and authentication data, such that the first timestamp and authentication data are also encrypted using the public component of the first public key pair; and

wherein said computer control logic further comprises:

logic configured to cause the processor to, after decrypting the encrypted information from each STB, verify the authentication data and determine whether the first timestamp matches any previously received timestamps.

17. The computer program product of claim 15, wherein said computer control logic further comprises:

logic configured to cause the processor to receive, at the anonymizing service node, further information from each STB, encrypted using the public component of the first public key pair, the further information including

the identifier of the STB, and

an indication of whether the advertisement was displayed at the STB, encrypted with the public component of the second key pair,

wherein the indication of whether the advertisement was displayed was encrypted using the public component of the first public key pair;

logic configured to cause the processor to decrypt the encrypted further information from each STB, to recover the identifier of eachSTB, the decryption performed using the private component of the first public key pair; and

logic configured to cause the processor to forward all encrypted indications of whether the advertisement was displayed and the encrypted identifier of the advertisement together to the head end, enabling the head end to count the number of STBs that displayed the advertisement.

18. The computer program product of claim 17, wherein the further information includes authentication data and a second timestamp, such that the authentication data and the second timestamp have also been encrypted using the public component of the first public key pair; and

wherein said computer control logic further comprises:

logic configured to cause the processor to, after decrypting the encrypted further information, verify the authentication data and determining whether the second timestamp matches any previously received timestamp.

19. The computer program product of claim 18, wherein said computer control logic further comprises:

logic configured to cause the processor to calculate a checksum on the basis of the STBs' respective indications of whether the advertisement was displayed;

logic configured to cause the processor to encrypt the checksum and the second timestamp using the public component of the second public key pair; and

logic configured to cause the processor to forward the encrypted checksum and second timestamp to the head end along with each of the STBs' respective indications of whether the advertisement was displayed.

20. The computer program product of claim 17, wherein the indication of whether the advertisement was displayed is generated at each STB,

wherein a decision regarding whether to display the advertisement is made at each STB in response to a command issued to each STB from the head end, wherein the command includes a parameter $P_{show} = N/M$,

where N is the number of households desired to see the advertisement and

M is the number of households indicating that the criteria are met, as counted by the head end; and

wherein after receiving the command, each STB creates a random number between 0 and 1 and displays the advertisement only if the random number is less than or equal to $P_{sh0\ w}$.

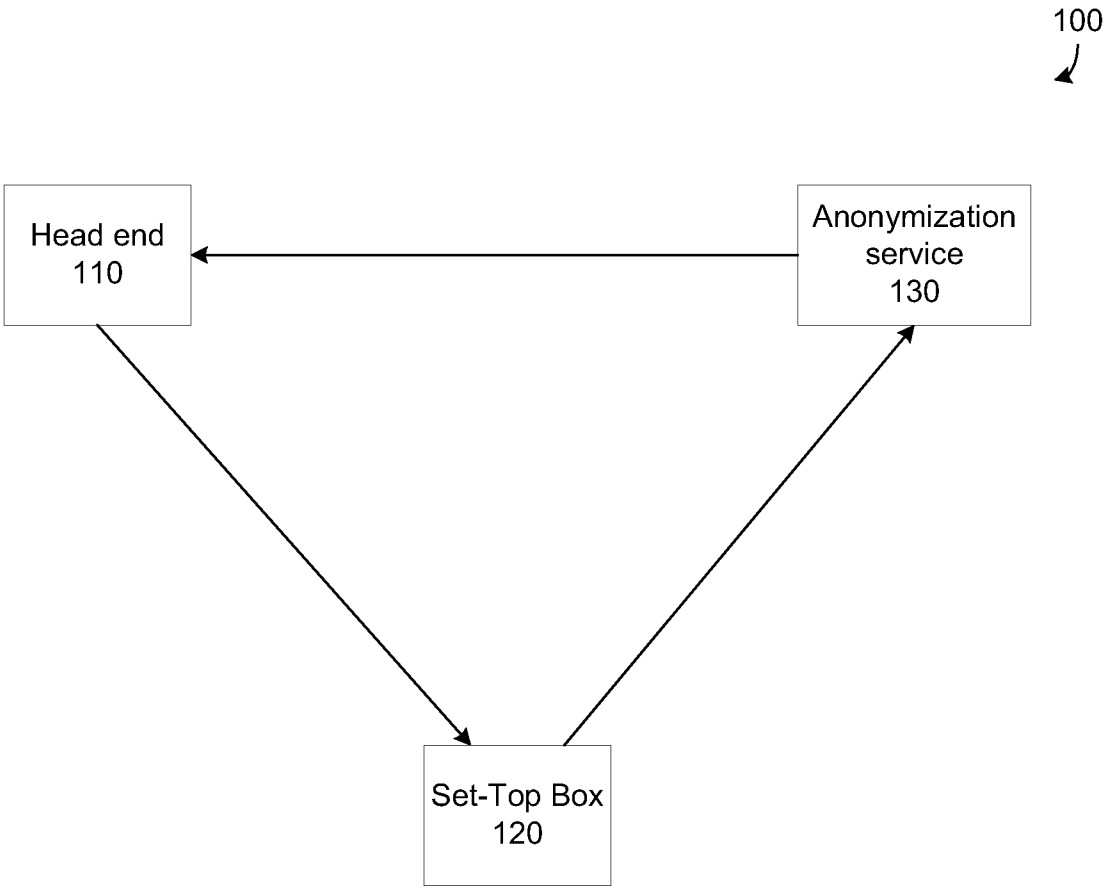

FIG. 1

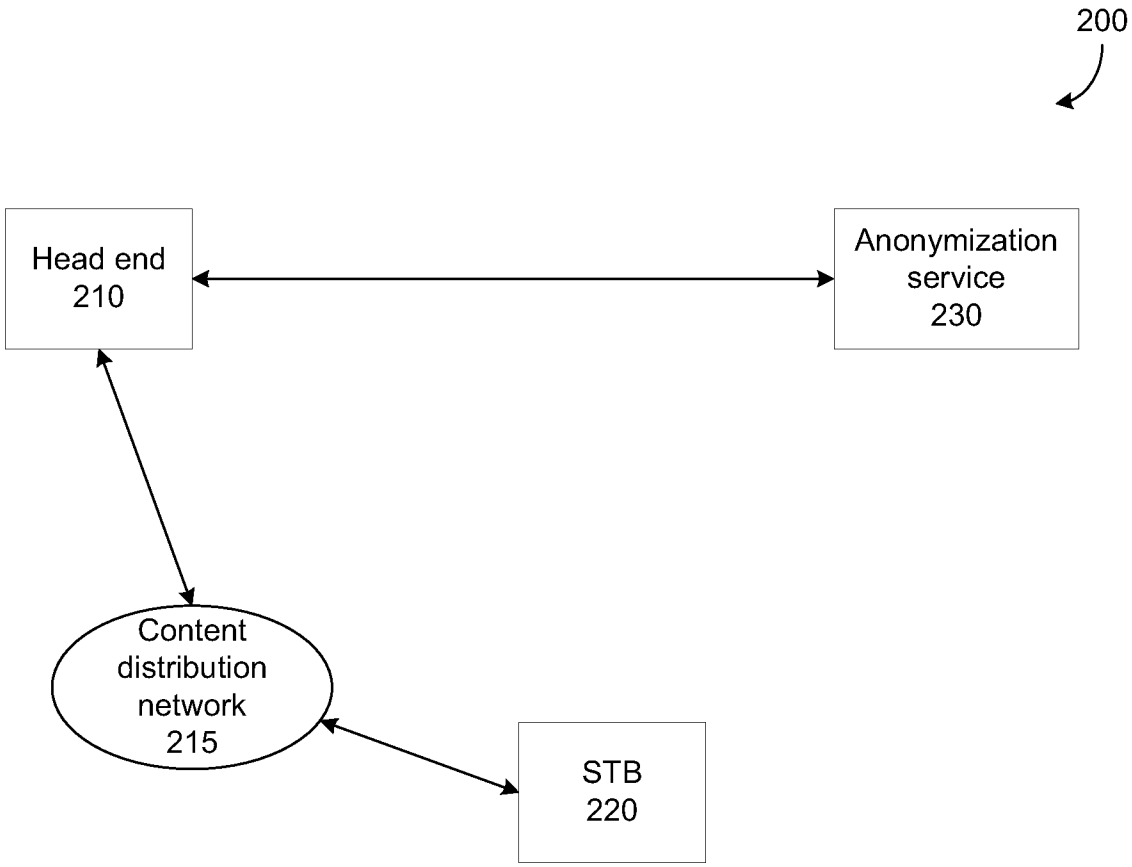
200
Head end
210
Anonymization
service
230
Content
distribution
network
215
STB
220

FIG. 2

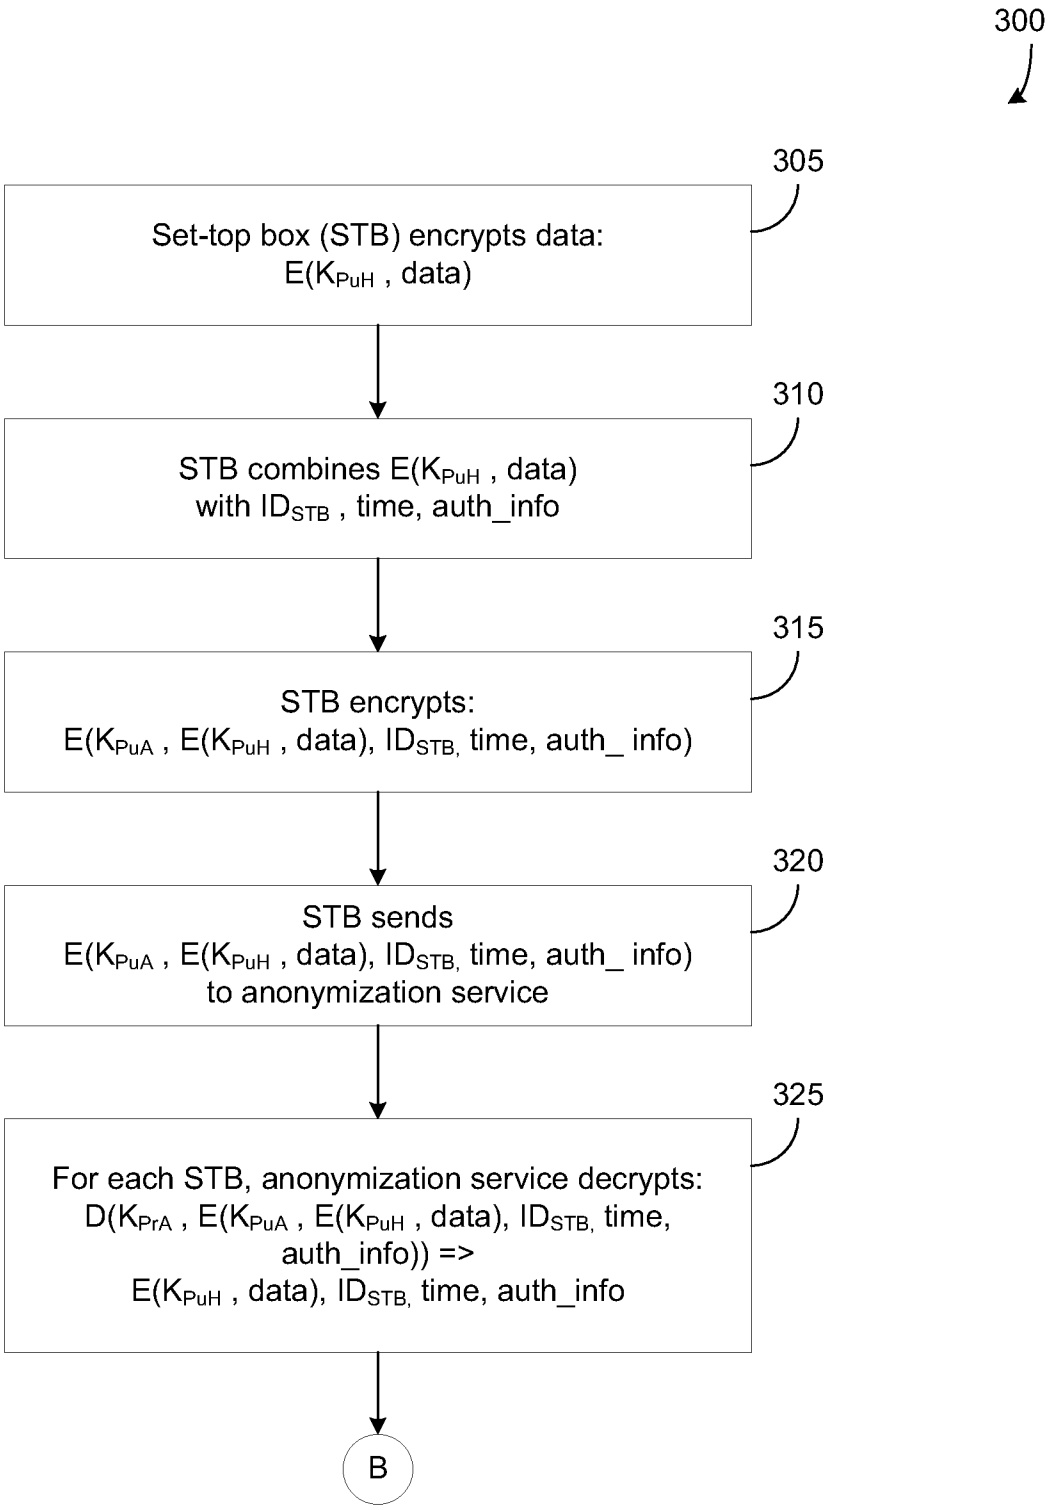
300
305
Set-top box (STB) encrypts data:
E(K_PuH , data)
310
STB combines E(K_PuH , data)
with ID_STB , time, auth_info
315
STB encrypts:
E(K_PuA , E(K_PuH , data), ID_STB, time, auth_ info)
320
STB sends
E(K_PuA , E(K_PuH , data), ID_STB, time, auth_ info)
to anonymization service
325
For each STB, anonymization service decrypts:
D(K_PrA , E(K_PuA , E(K_PuH , data), ID_STB, time,
auth_info)) =>
E(K_PuH , data), ID_STB, time, auth_info
B

FIG. 3A

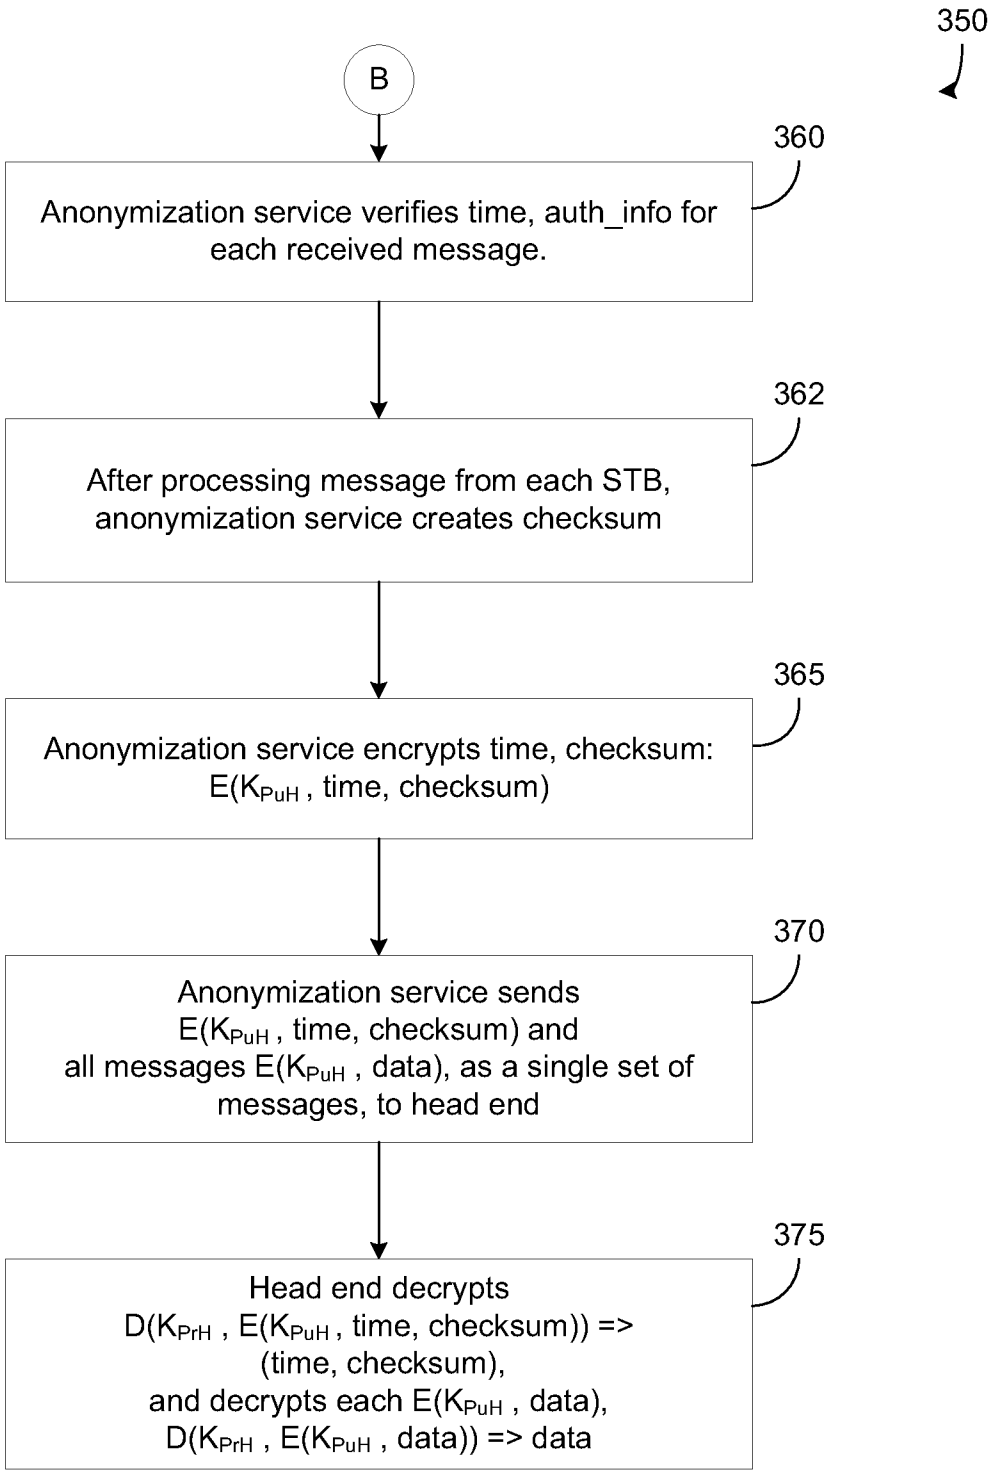
350
B
360
Anonymization service verifies time, auth_info for each received message.
362
After processing message from each STB, anonymization service creates checksum
365
Anonymization service encrypts time, checksum: E(K_PuH , time, checksum)
370
Anonymization service sends E(K_PuH , time, checksum) and all messages E(K_PuH , data), as a single set of messages, to head end
375
Head end decrypts D(K_PrH , E(K_PuH , time, checksum)) => (time, checksum), and decrypts each E(K_PuH , data), D(K_PrH , E(K_PuH , data)) => data

FIG. 3B

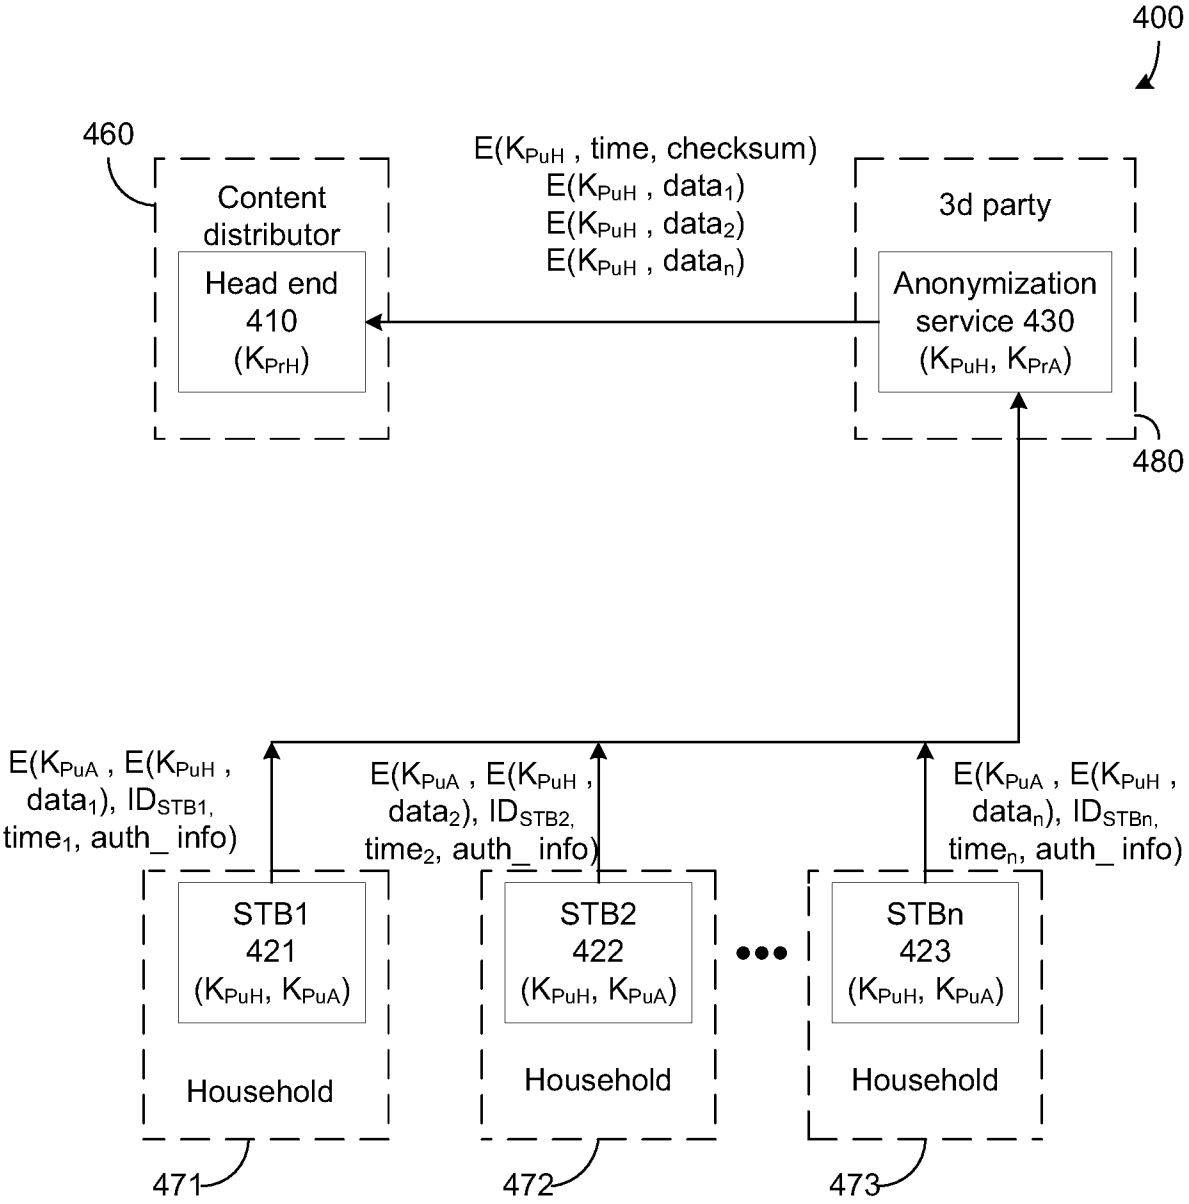
400
460
Content distributor
Head end 410 (K_{PrH})
$E(K_{PuH}$, time, checksum)
$E(K_{PuH}$, $data_1)$
$E(K_{PuH}$, $data_2)$
$E(K_{PuH}$, $data_n)$
3d party
Anonymization service 430 (K_{PuH}, K_{PrA})
480
$E(K_{PuA}$, $E(K_{PuH}$, $data_1)$, $ID_{STB1,}$ $time_1$, auth_ info)
$E(K_{PuA}$, $E(K_{PuH}$, $data_2)$, $ID_{STB2,}$ $time_2$, auth_ info)
$E(K_{PuA}$, $E(K_{PuH}$, $data_n)$, $ID_{STBn,}$ $time_n$, auth_ info)
STB1 421 (K_{PuH}, K_{PuA})
STB2 422 (K_{PuH}, K_{PuA})
STBn 423 (K_{PuH}, K_{PuA})
Household
Household
Household
471
472
473

FIG. 4

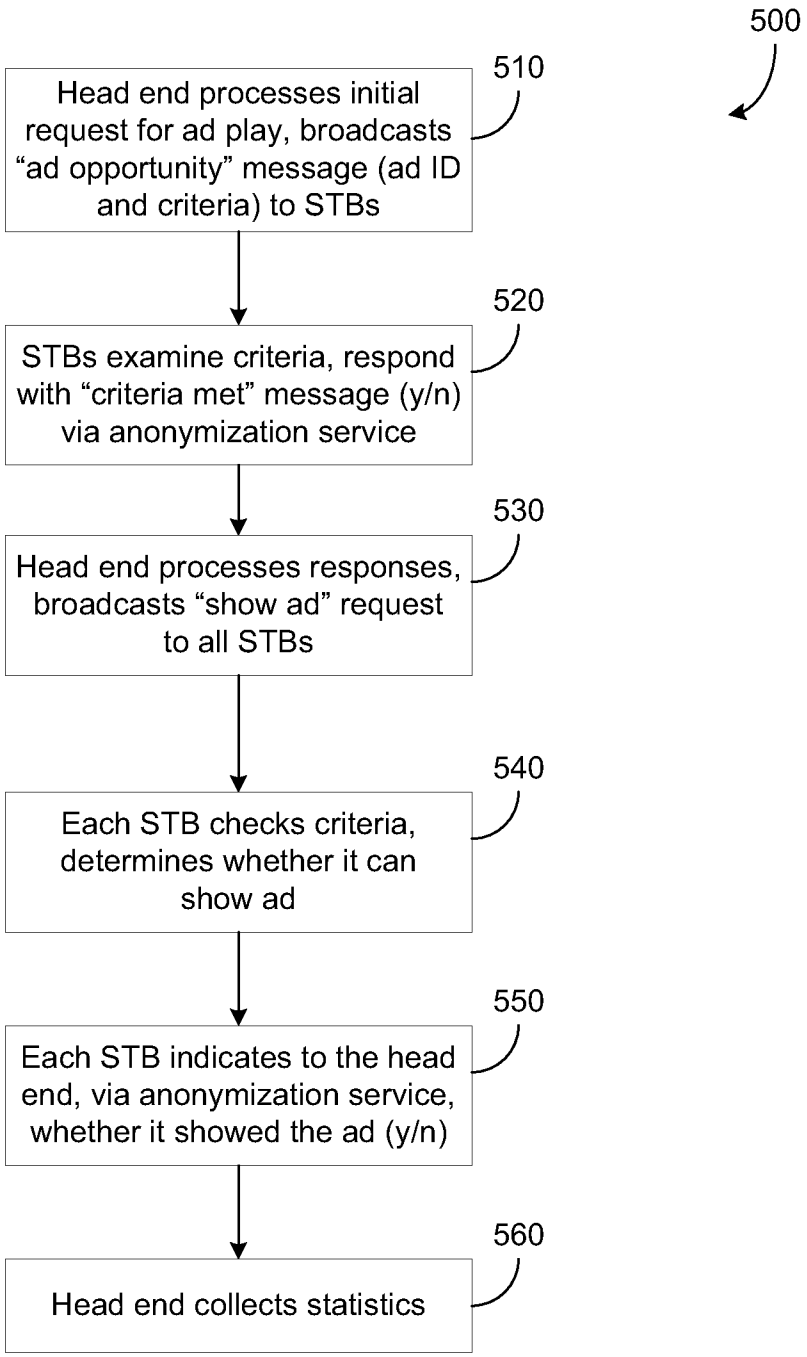
500
510
Head end processes initial request for ad play, broadcasts "ad opportunity" message (ad ID and criteria) to STBs
520
STBs examine criteria, respond with "criteria met" message (y/n) via anonymization service
530
Head end processes responses, broadcasts "show ad" request to all STBs
540
Each STB checks criteria, determines whether it can show ad
550
Each STB indicates to the head end, via anonymization service, whether it showed the ad (y/n)
560
Head end collects statistics

FIG. 5

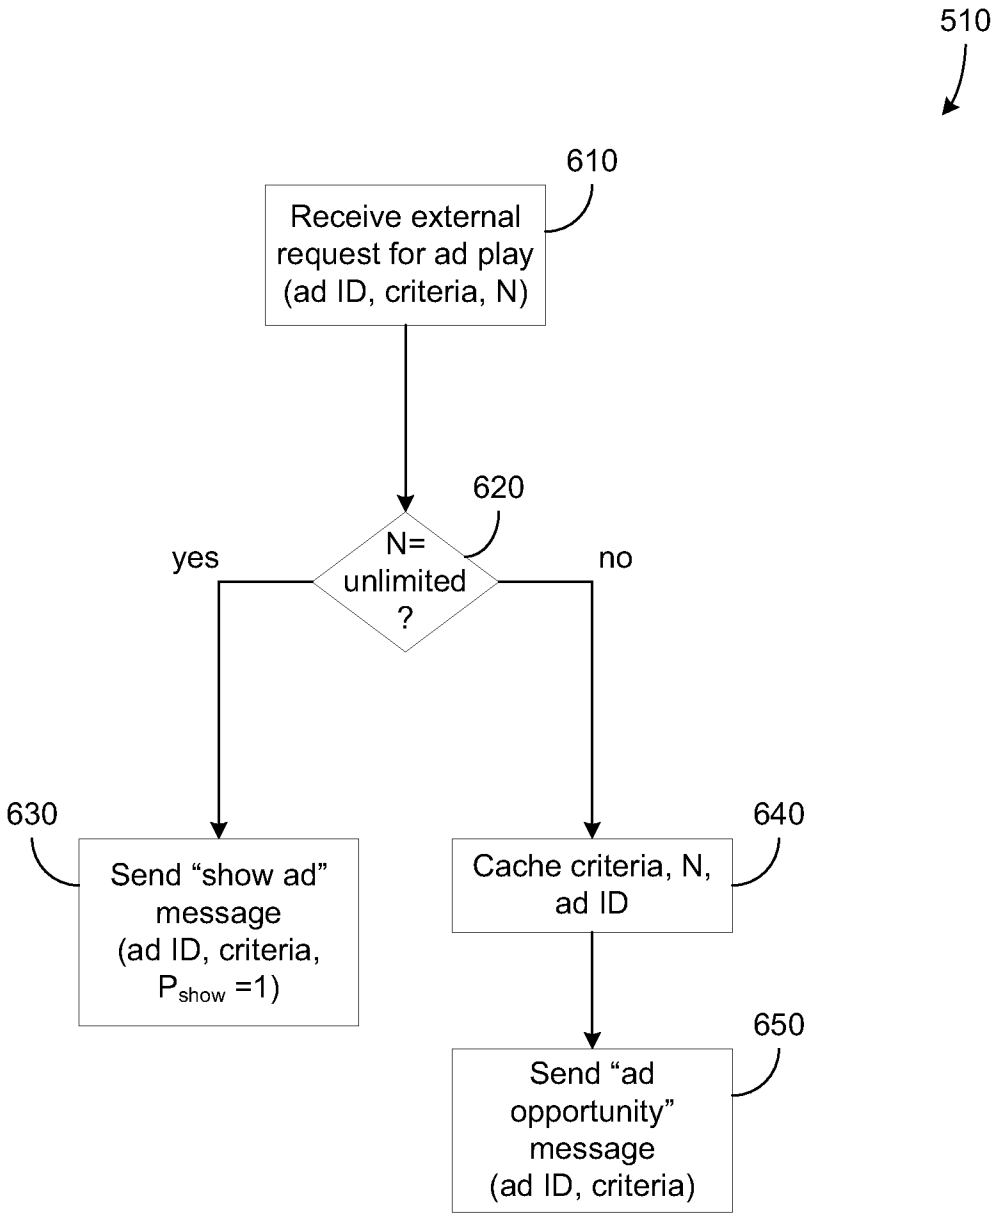
510
610
Receive external request for ad play (ad ID, criteria, N)
620
N= unlimited ?
yes
no
630
Send “show ad” message (ad ID, criteria, P_{show} =1)
640
Cache criteria, N, ad ID
650
Send “ad opportunity” message (ad ID, criteria)

FIG. 6

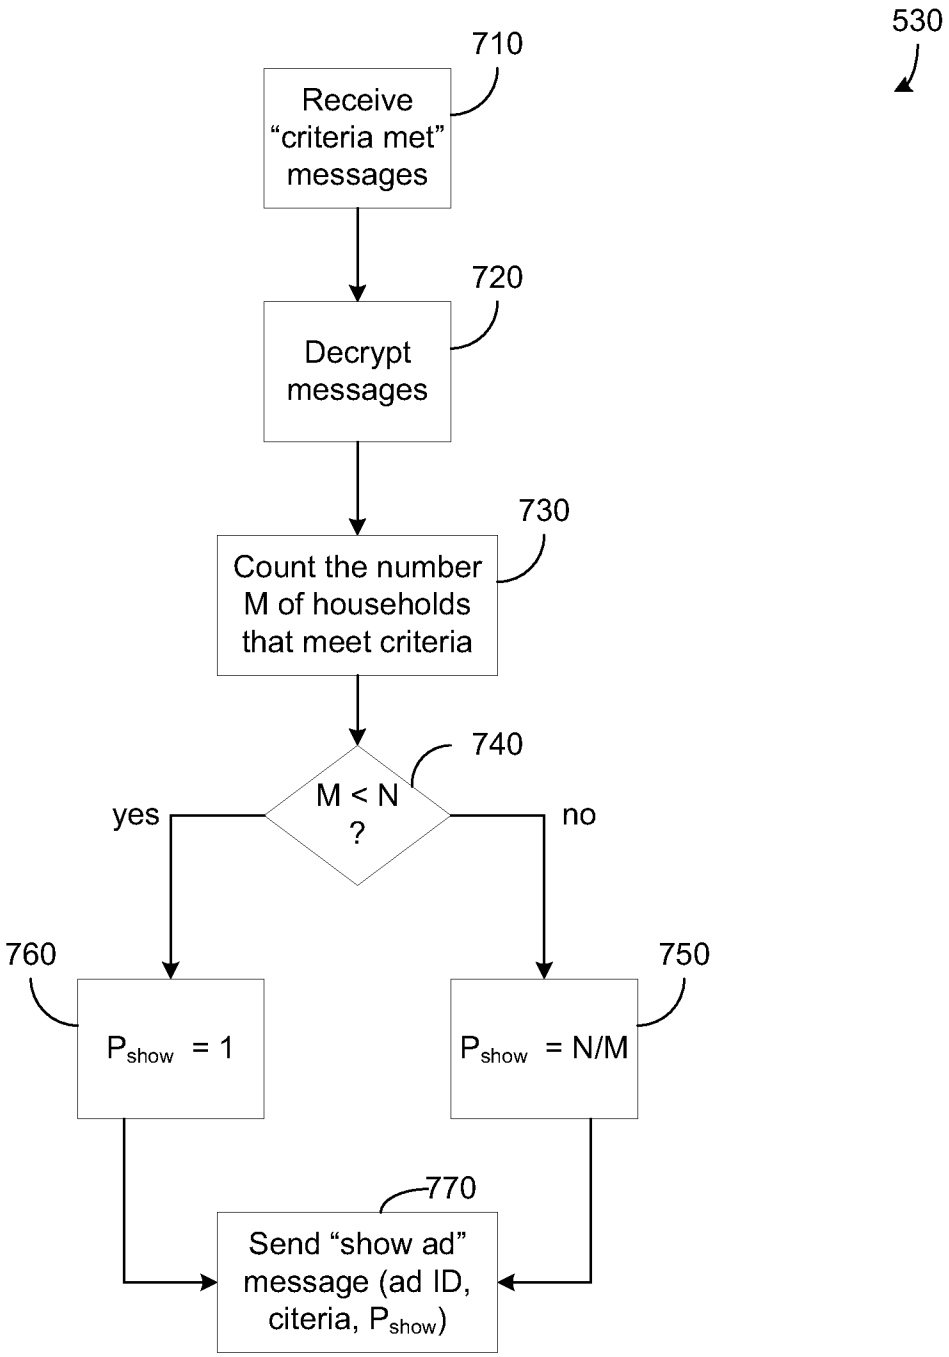
530
710
Receive "criteria met" messages
720
Decrypt messages
730
Count the number M of households that meet criteria
740
M < N ?
yes
no
760
P_{show} = 1
750
P_{show} = N/M
770
Send "show ad" message (ad ID, citeria, P_{show})

FIG. 7

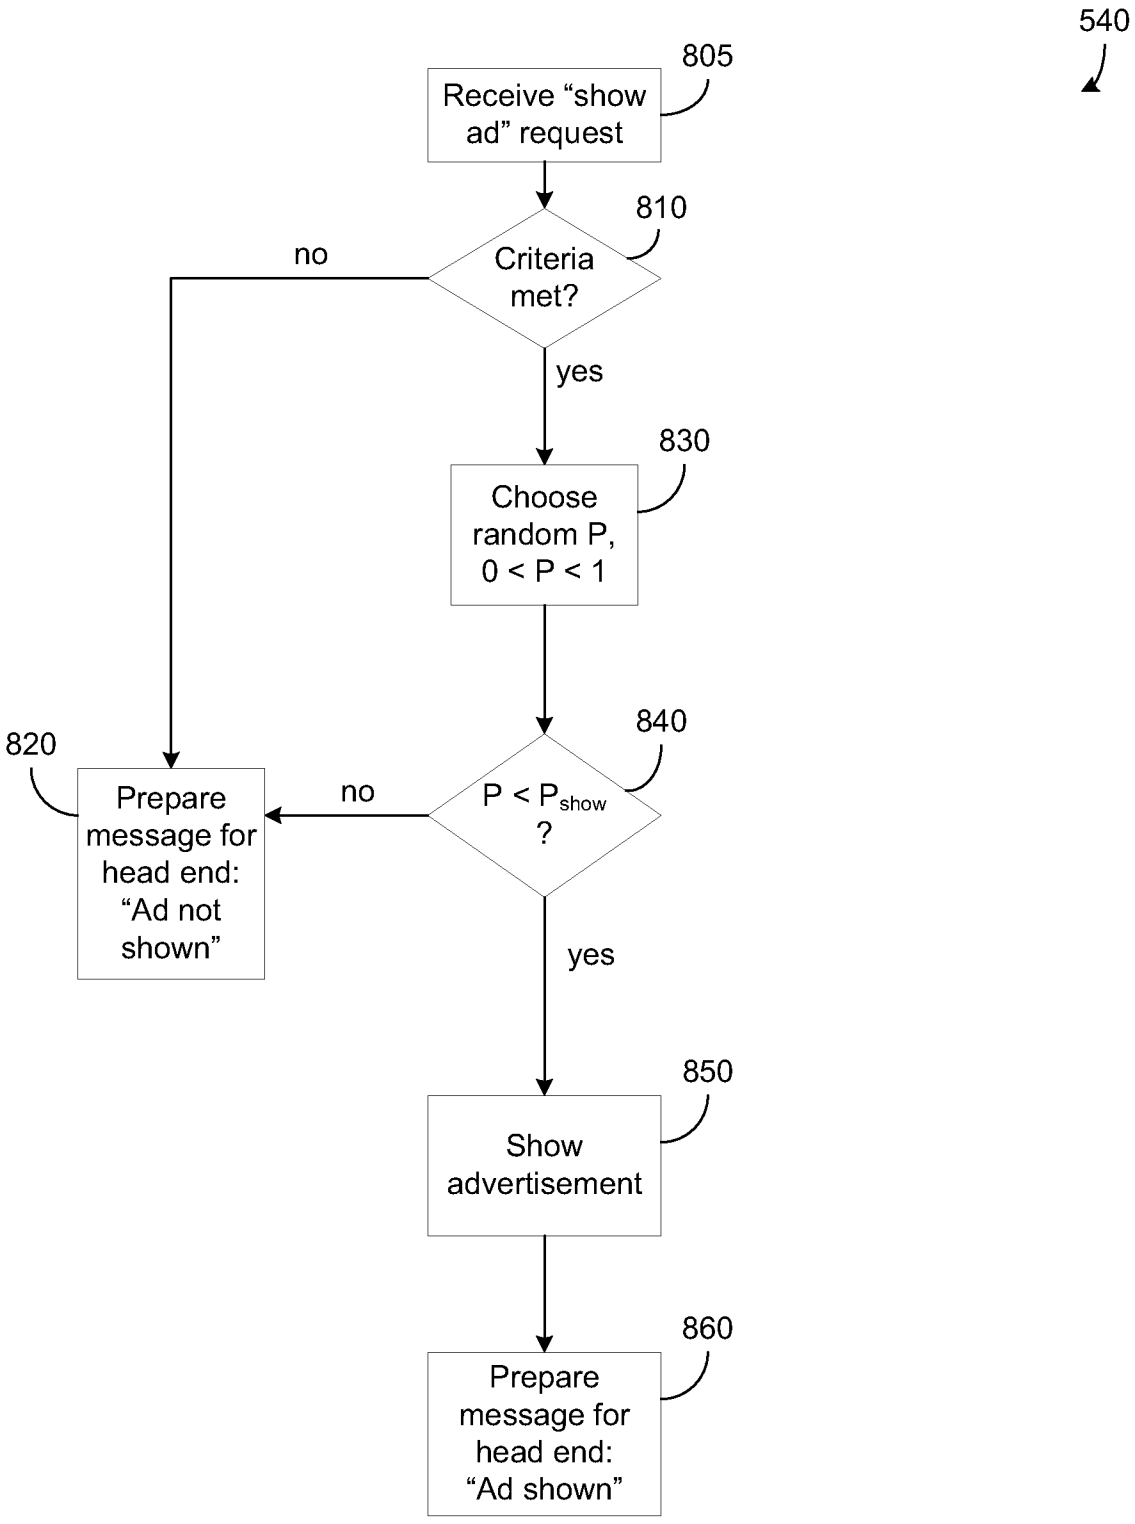
540
805
Receive "show ad" request
810
Criteria met?
no
yes
830
Choose random P, 0 < P < 1
840
P < P show ?
820
Prepare message for head end: "Ad not shown"
no
yes
850
Show advertisement
860
Prepare message for head end: "Ad shown"

FIG. 8

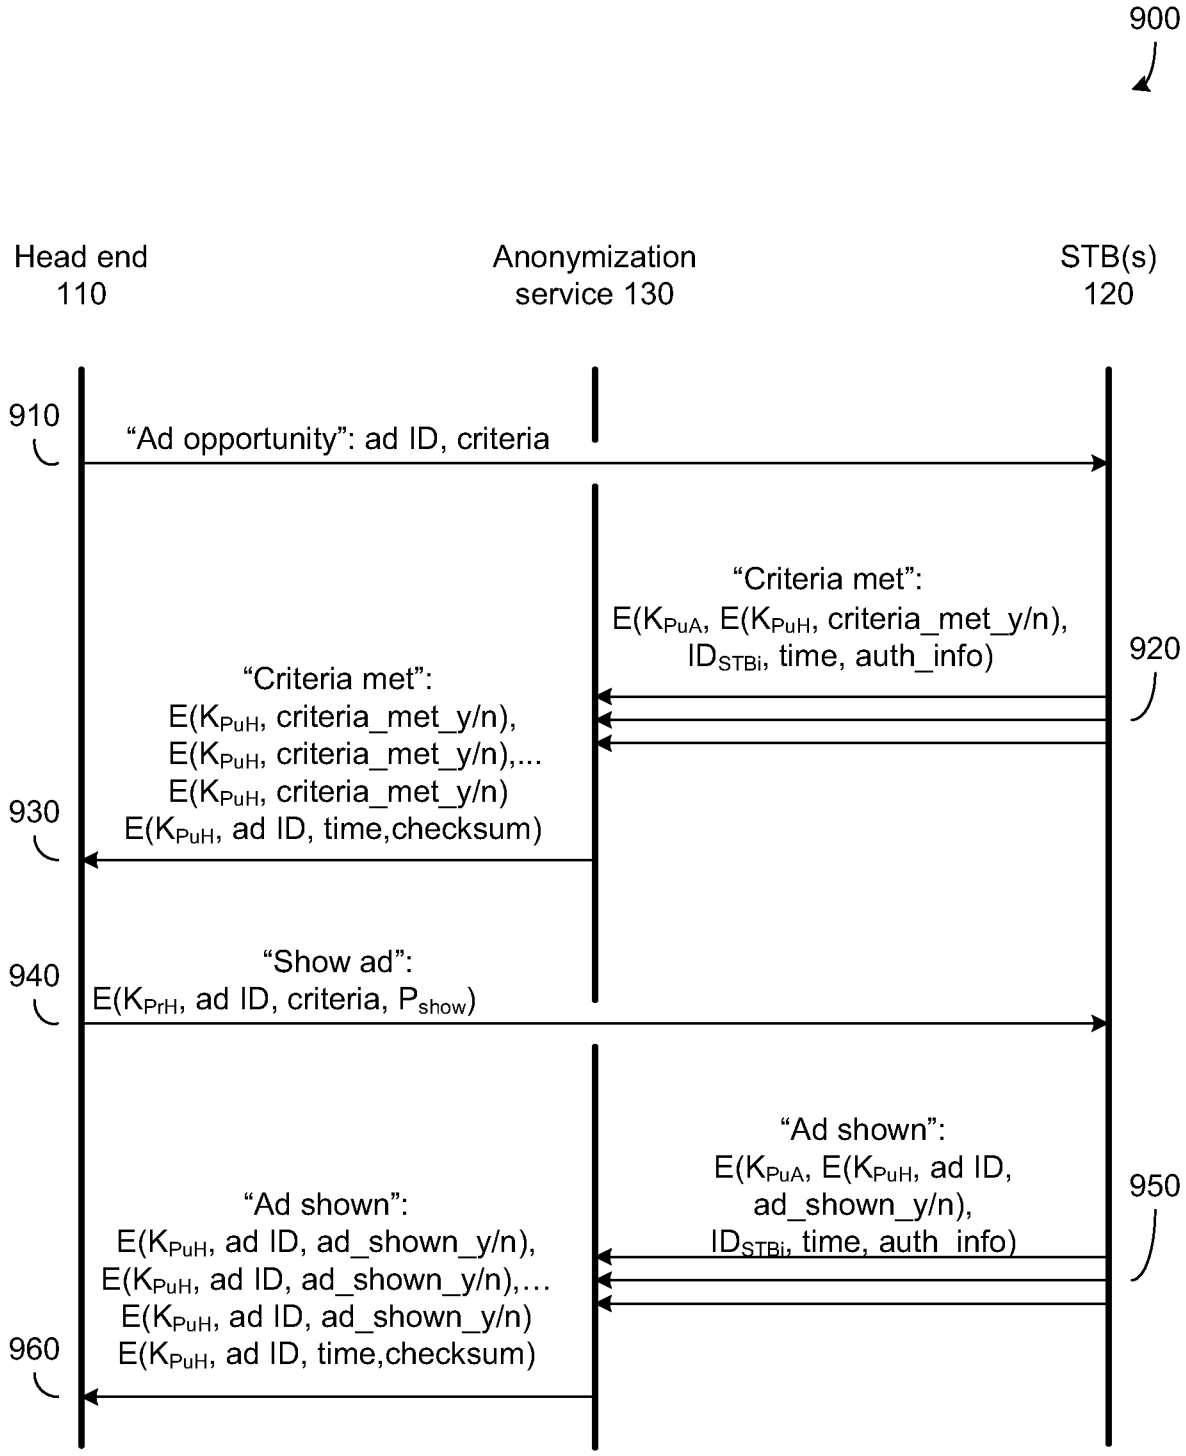
900
Head end 110
Anonymization service 130
STB(s) 120
910
"Ad opportunity": ad ID, criteria
"Criteria met": E(K_PuA, E(K_PuH, criteria_met_y/n), ID_STBi, time, auth_info)
920
"Criteria met": E(K_PuH, criteria_met_y/n), E(K_PuH, criteria_met_y/n),... E(K_PuH, criteria_met_y/n) E(K_PuH, ad ID, time,checksum)
930
940
"Show ad": E(K_PrH, ad ID, criteria, P_show)
"Ad shown": E(K_PuA, E(K_PuH, ad ID, ad_shown_y/n), ID_STBi, time, auth_info)
950
"Ad shown": E(K_PuH, ad ID, ad_shown_y/n), E(K_PuH, ad ID, ad_shown_y/n),… E(K_PuH, ad ID, ad_shown_y/n) E(K_PuH, ad ID, time,checksum)
960

FIG. 9

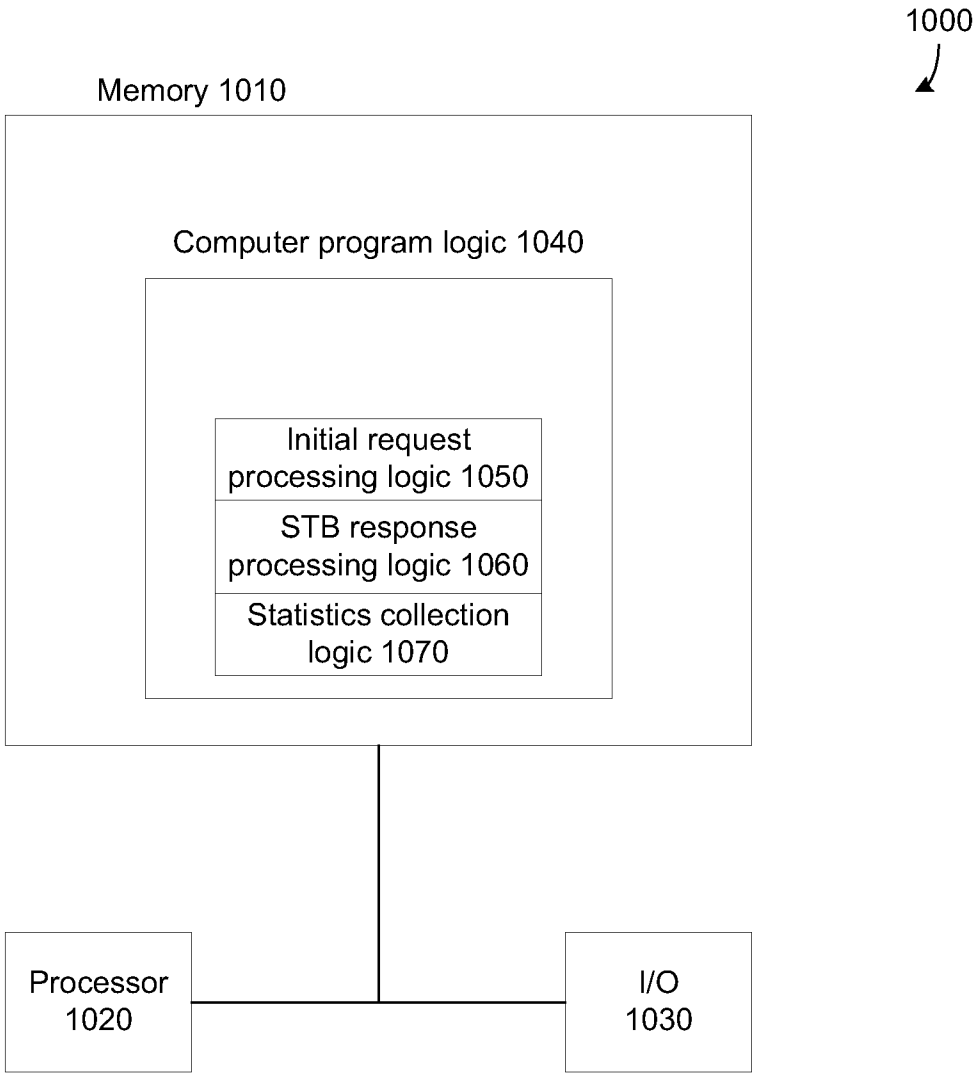
1000
Memory 1010
Computer program logic 1040
Initial request processing logic 1050
STB response processing logic 1060
Statistics collection logic 1070
Processor 1020
I/O 1030

FIG. 10

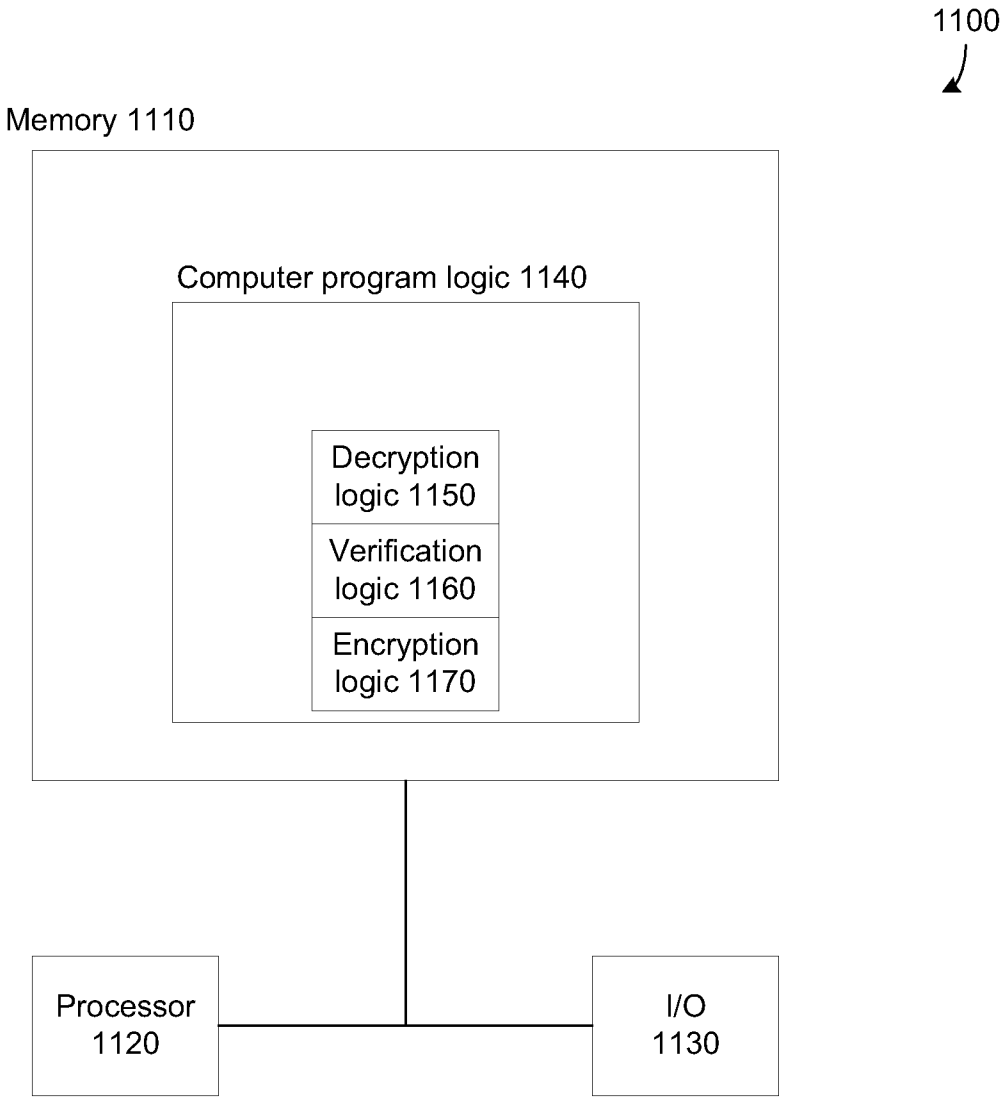
1100
Memory 1110
Computer program logic 1140
Decryption logic 1150
Verification logic 1160
Encryption logic 1170
Processor 1120
I/O 1130

FIG. 11

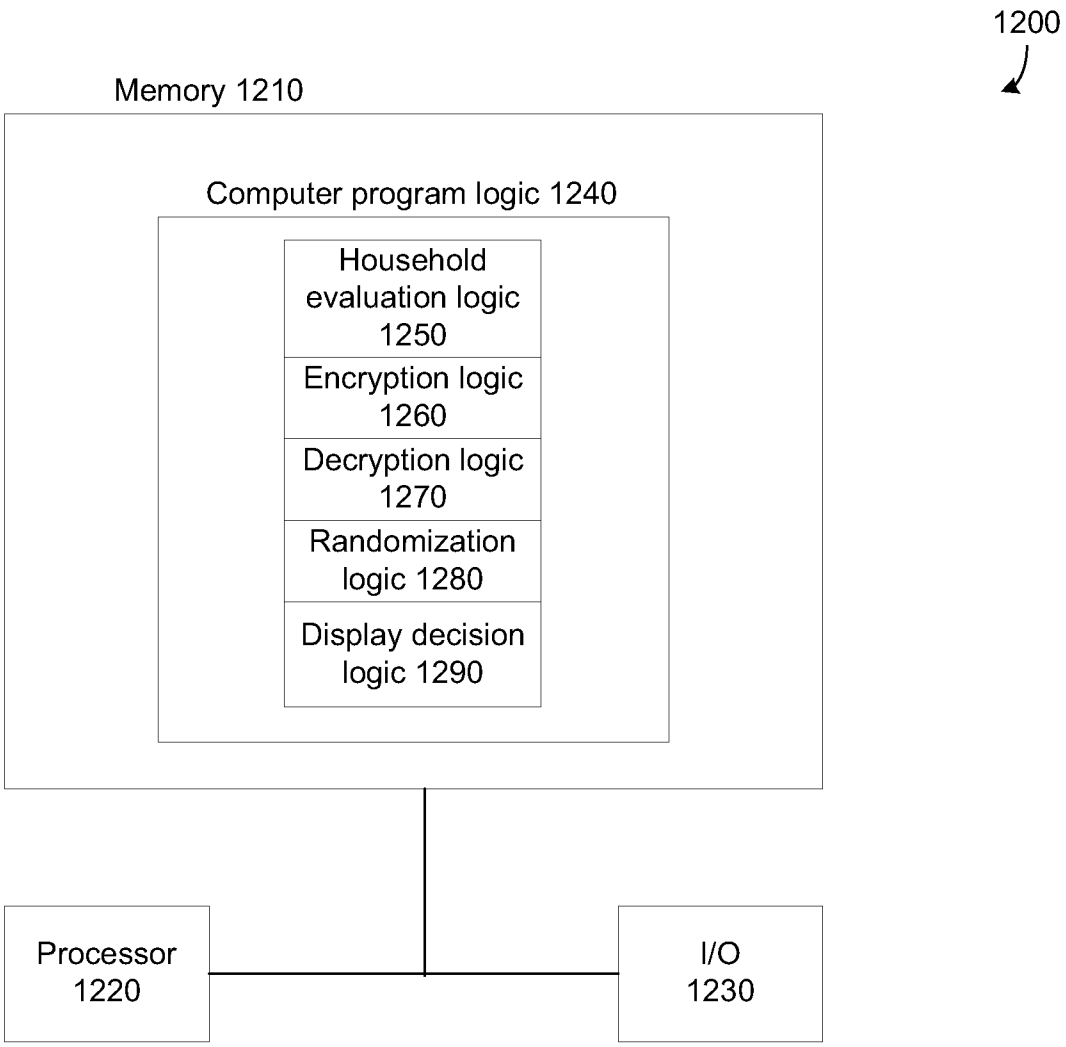
1200
Memory 1210
Computer program logic 1240
Household evaluation logic 1250
Encryption logic 1260
Decryption logic 1270
Randomization logic 1280
Display decision logic 1290
Processor 1220
I/O 1230

FIG. 12