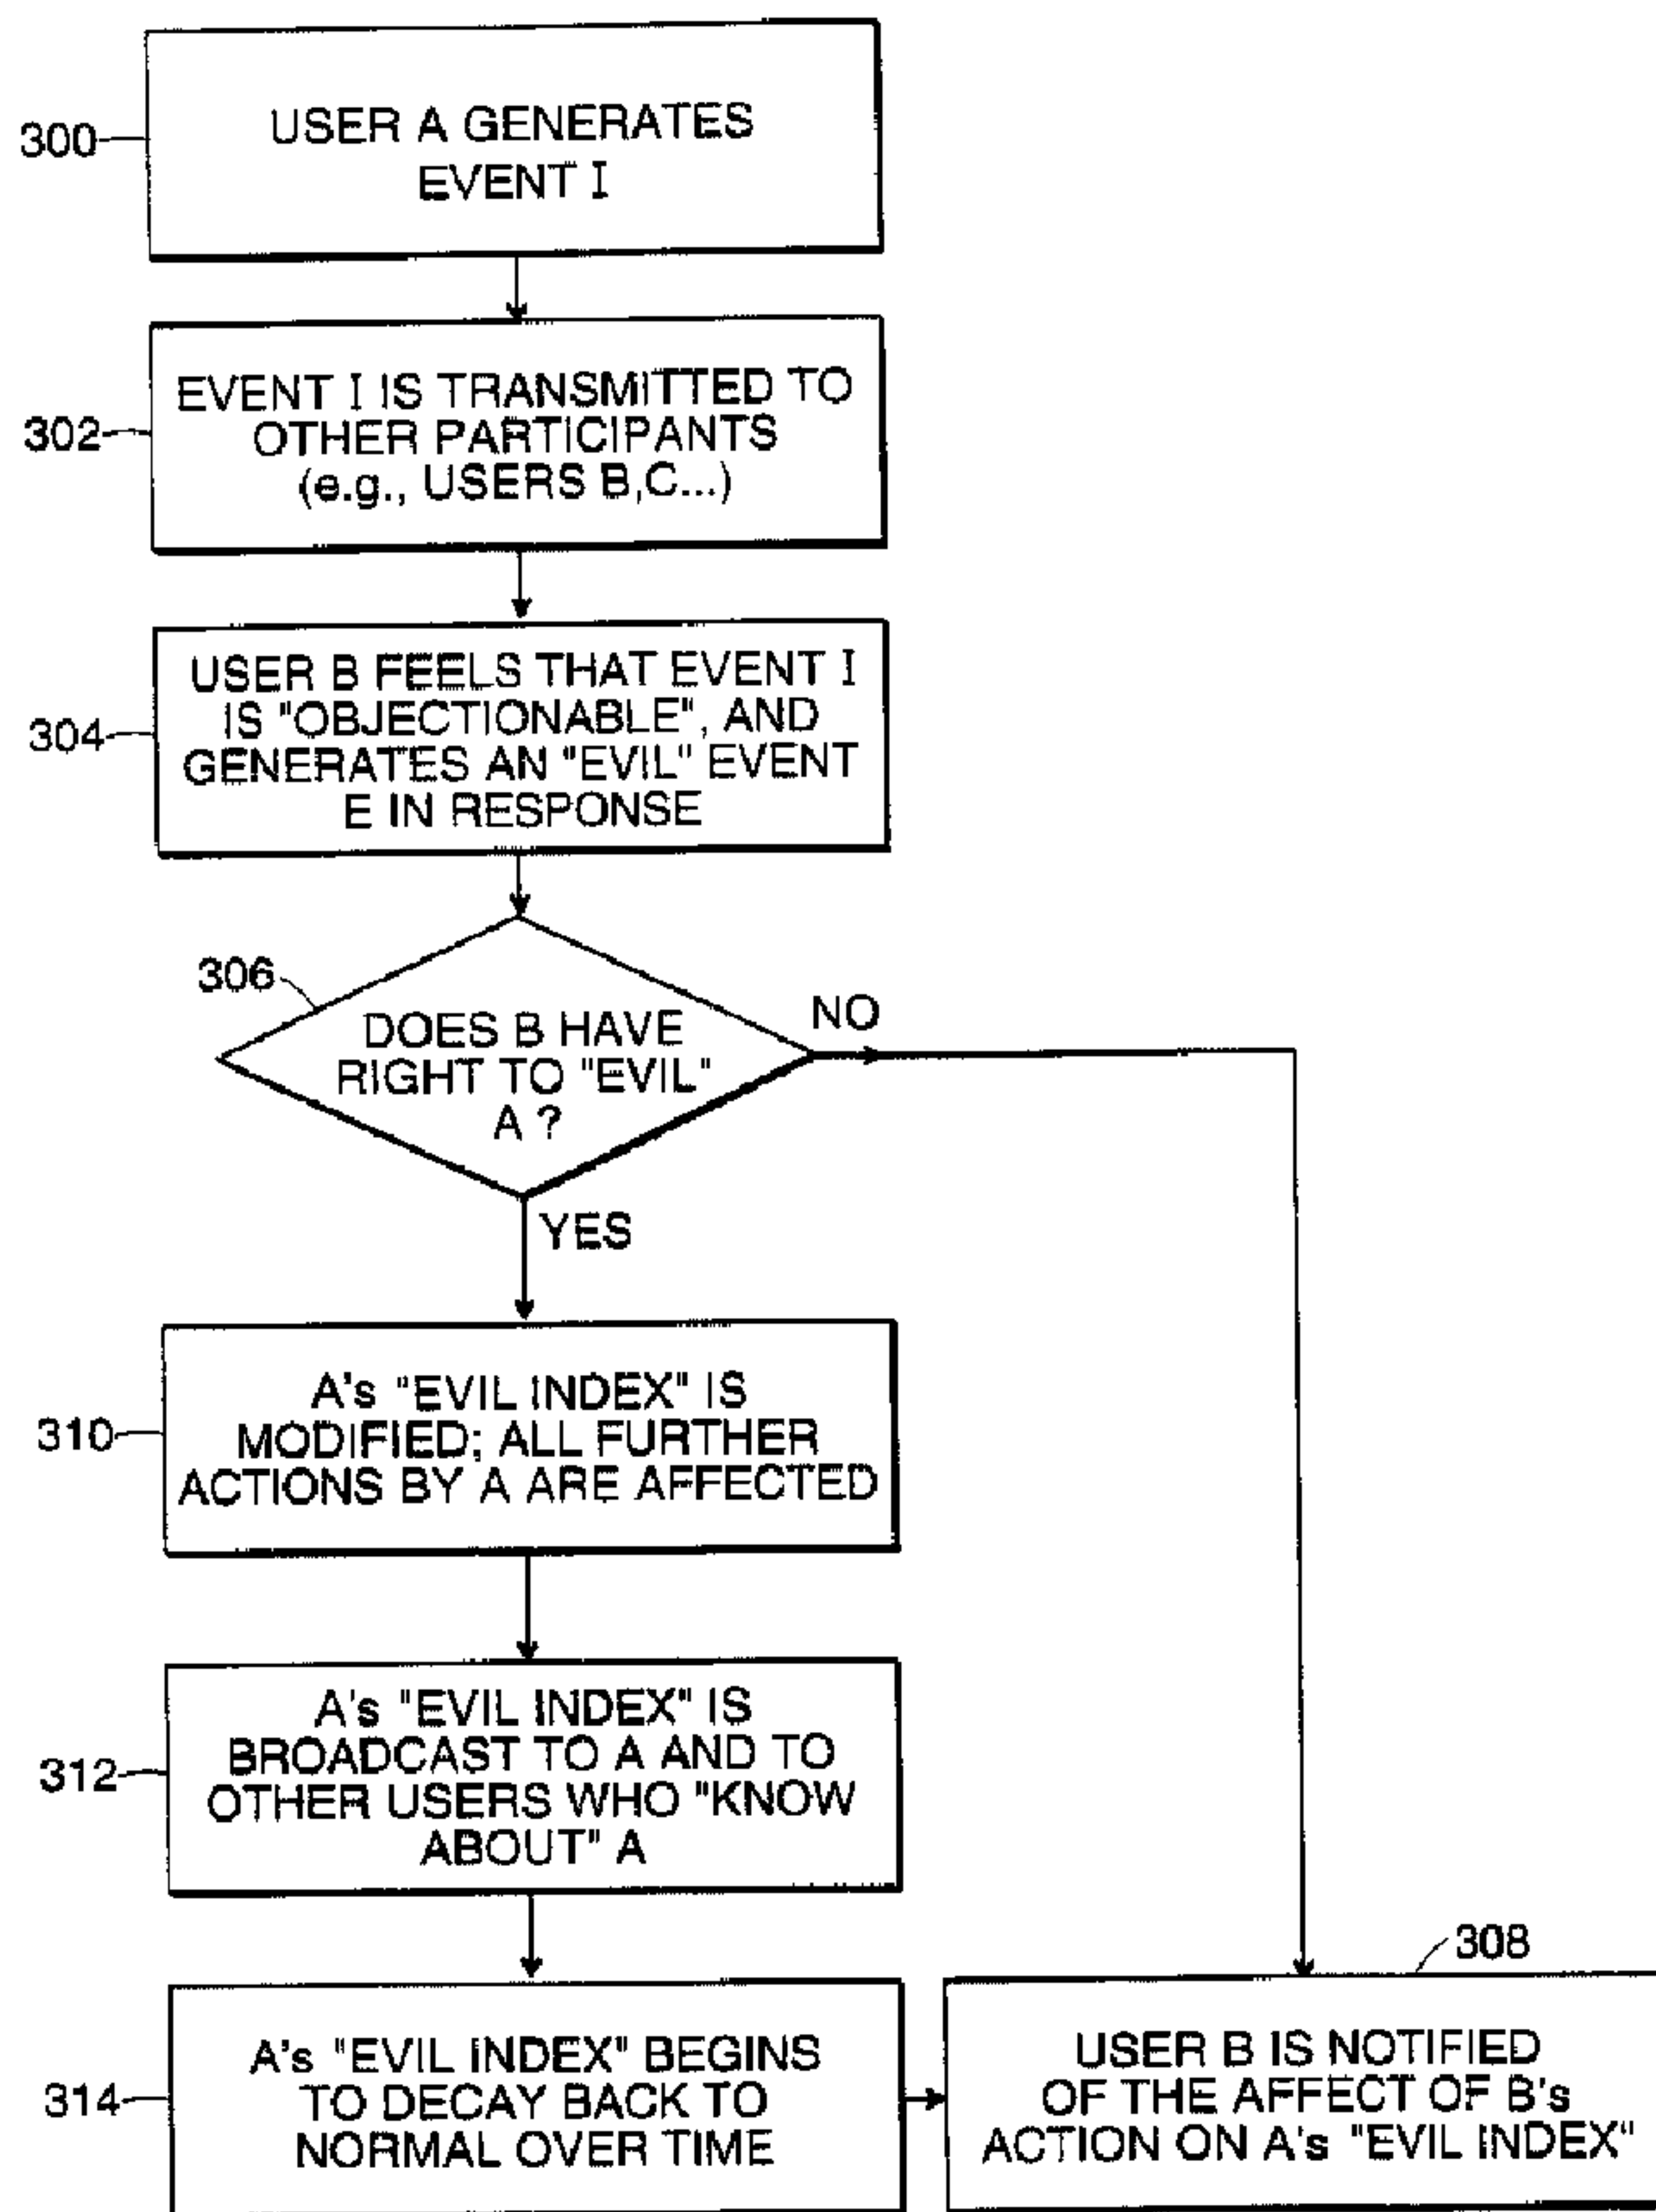




(86) Date de dépôt PCT/PCT Filing Date: 1998/05/15  
 (87) Date publication PCT/PCT Publication Date: 1998/11/26  
 (45) Date de délivrance/Issue Date: 2003/01/07  
 (85) Entrée phase nationale/National Entry: 1999/11/16  
 (86) N° demande PCT/PCT Application No.: US 1998/009969  
 (87) N° publication PCT/PCT Publication No.: 1998/053387  
 (30) Priorité/Priority: 1997/05/20 (60/047,235) US

(51) Cl.Int.<sup>7</sup>/Int.Cl.<sup>7</sup> G06F 1/00  
 (72) Inventeurs/Inventors:  
 MORRIS, HARRY W., US;  
 BOSCO, ERIC, US;  
 LIPPKE, DAVID LOWELL, US;  
 STEELE, COLIN ANTHONY, US  
 (73) Propriétaire/Owner:  
 AMERICA ONLINE, INC., US  
 (74) Agent: RICHES, MCKENZIE & HERBERT LLP

(54) Titre : FORUMS DE DISCUSSION EN LIGNE A AUTO-CONTROLE ET A LIMITATION AUTOMATIQUE DE LA FREQUENCE D'ENVOI  
 (54) Title: SELF-POLICING, RATE LIMITING ONLINE FORUMS



(57) Abrégé/Abstract:

A method and means for self-policing and automatically rate-limiting multiple-user online forums. The preferred embodiment of the invention includes a set of rules that permit users to censure other users. A censured user has one or more "privileges" (which may include access to the online computer system or the effective message rate of a communications connection) taken

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

away or diminished; the privilege is gradually restored if the censure user behaves. The censoring participant and other pertinent participants are notified of the effect of the censoring on the censored user. In another aspect of the invention, the online computer system automatically tracks the rate at which a user sends certain types of messages, and can message rate limit a user who uses too many system resources by sending a large number of messages in rapid succession.



PCT

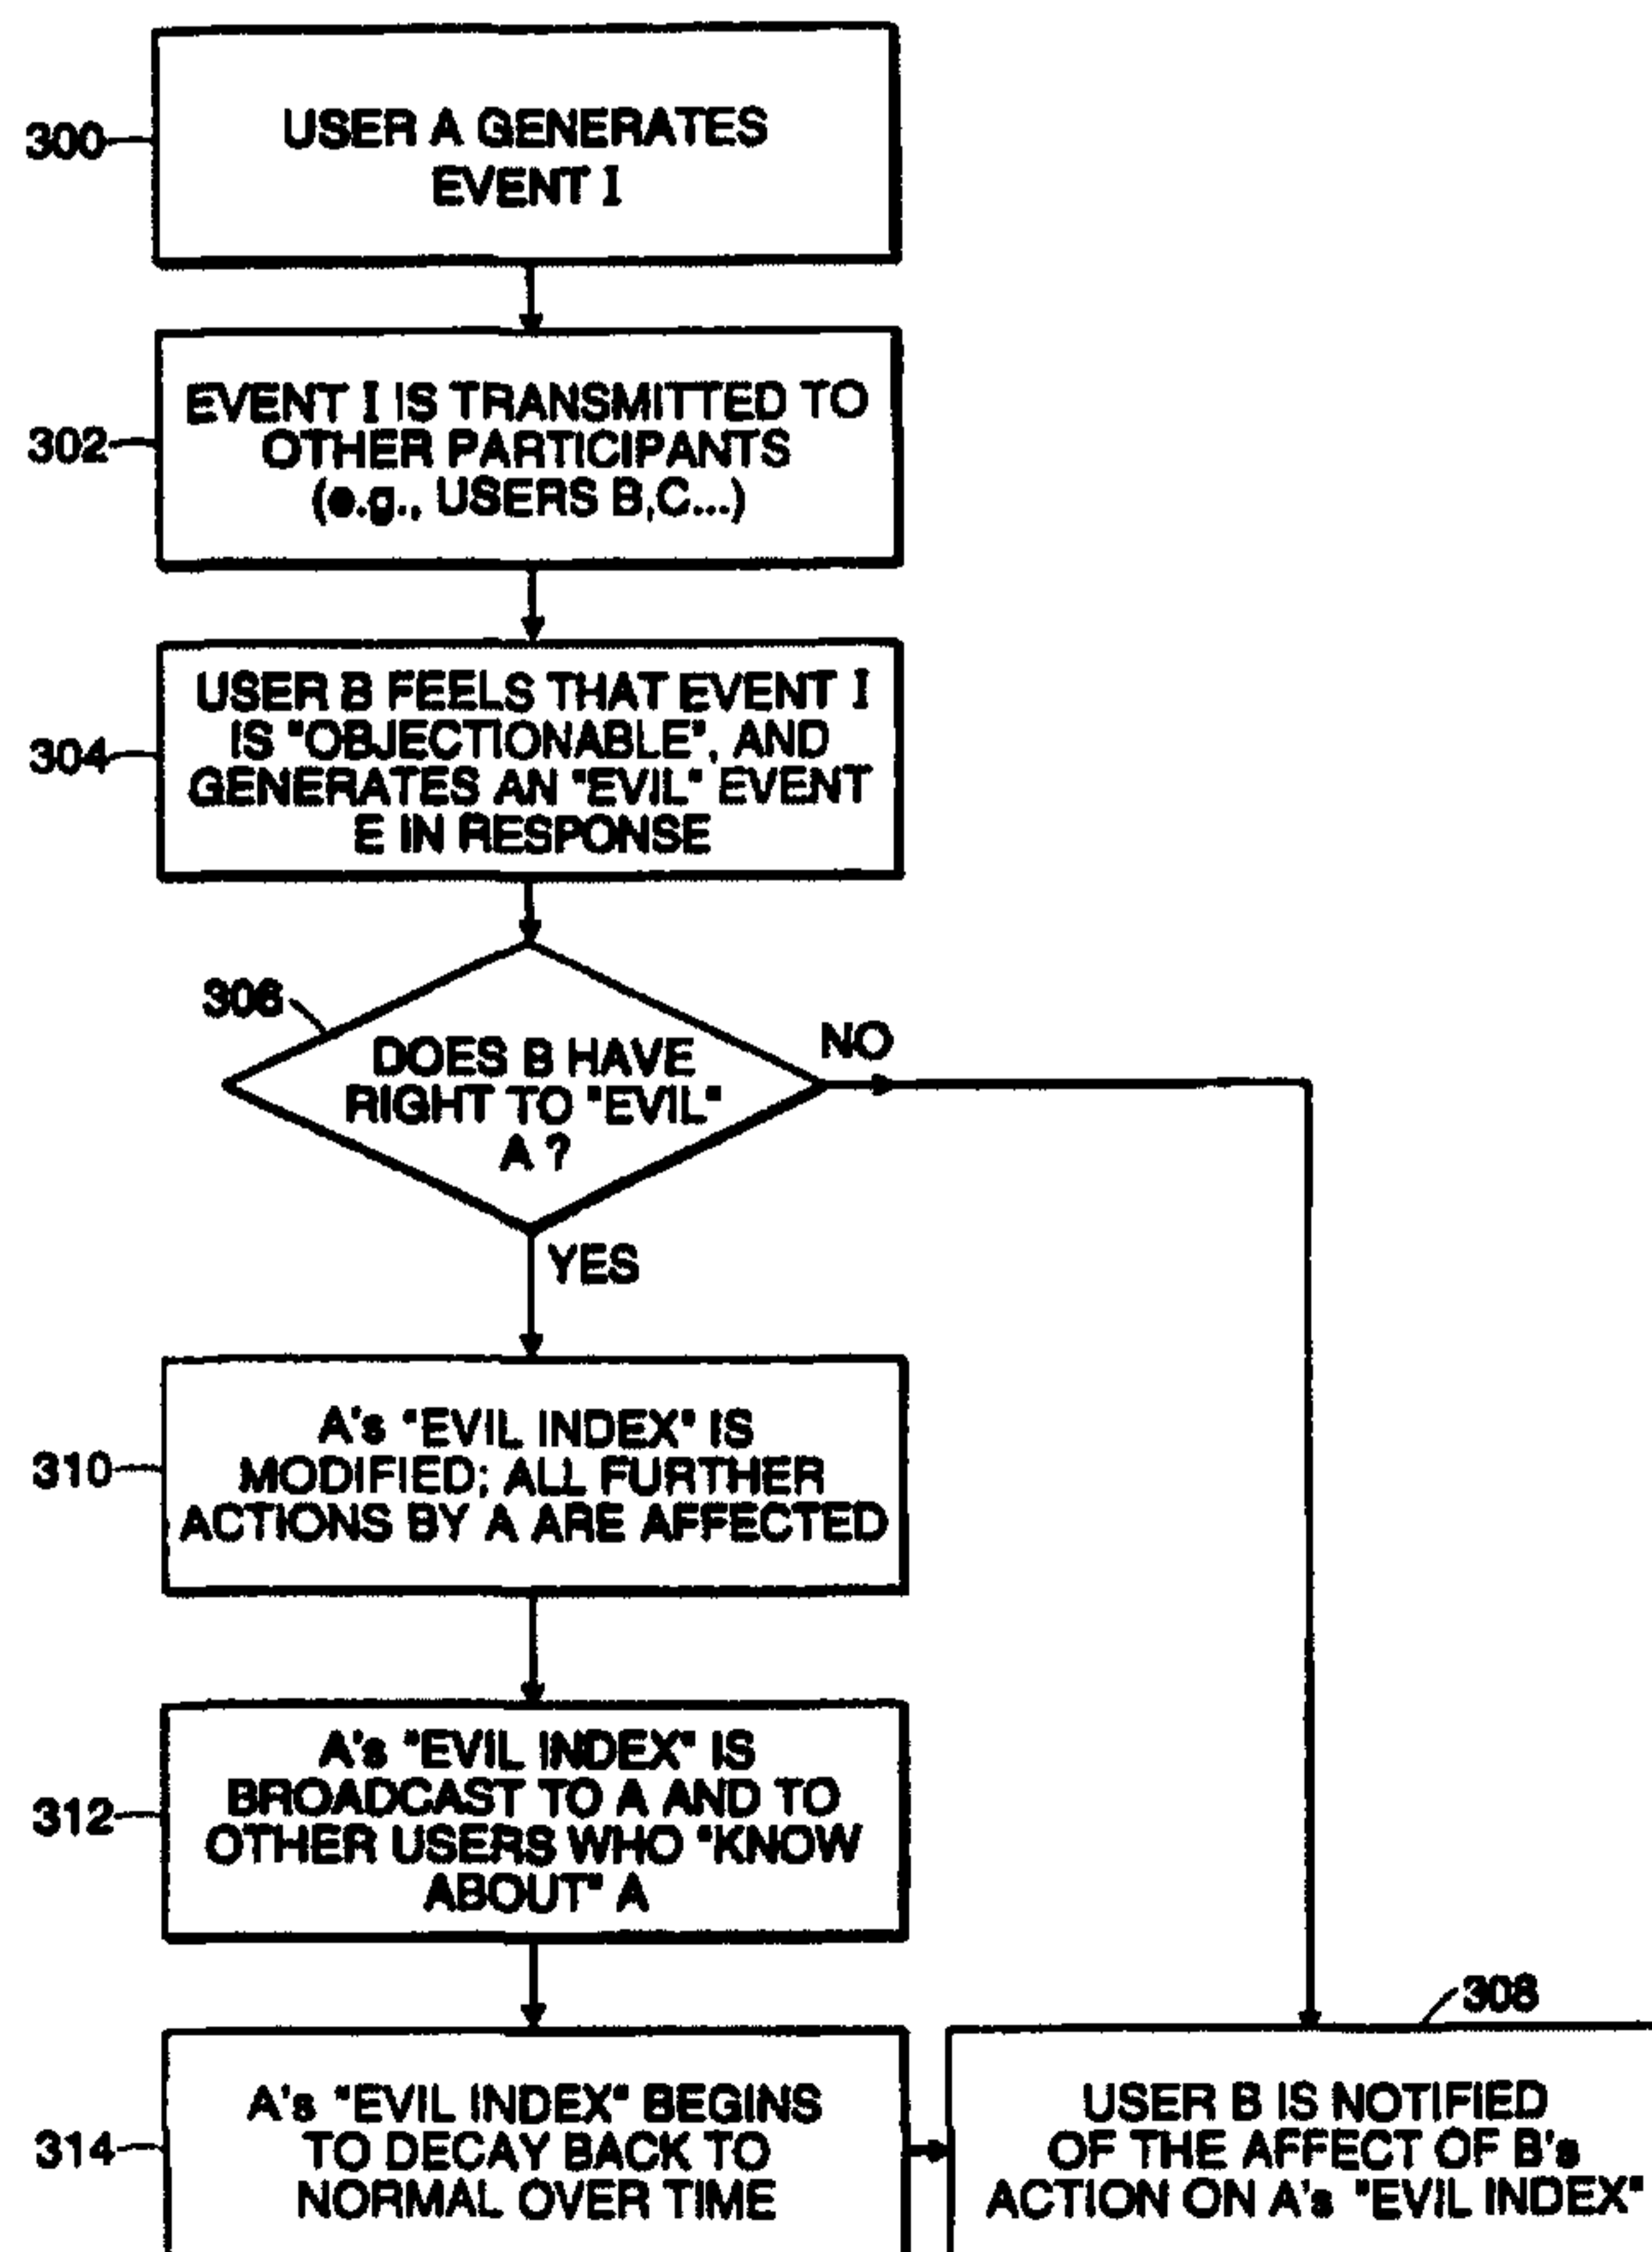
WORLD INTELLECTUAL PROPERTY ORGANIZATION  
International Bureau

INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

<b>(51) International Patent Classification <sup>6</sup> :</b>  <b>G06F 1/00</b>	<b>A1</b>	<b>(11) International Publication Number:</b> <b>WO 98/53387</b>  <b>(43) International Publication Date:</b> 26 November 1998 (26.11.98)
<b>(21) International Application Number:</b> PCT/US98/09969 <b>(22) International Filing Date:</b> 15 <sup>15</sup> 18 May 1998 (18.05.98)  <b>(30) Priority Data:</b> 60/047,235                      20 May 1997 (20.05.97)                      US  <b>(71) Applicant:</b> AMERICA ONLINE, INC. [US/US]; 22000 AOL Way, Dulles, VA 20166 (US).  <b>(72) Inventors:</b> MORRIS, Harry, W.; Apartment F, 1719 Ascot Way, Reston, VA 20190 (US). BOSCO, Eric; 2430 13th Court N., Arlington, VA 22201 (US). LIPPKE, David, Lowell; 1441 Kingsvale Circle, Herndon, VA 20170 (US). STEELE, Colin, Anthony; 635 Cobbler Terrace, Leesburg, VA 20175 (US).  <b>(74) Agent:</b> PHILLIPS, John, C.; Fish & Richardson P.C., 601 Thirteenth Street, N.W., Washington, DC 20005 (US).		<b>(81) Designated States:</b> AU, BR, CA, JP, MX, European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE).  <b>Published</b> <i>With international search report.</i>

**(54) Title:** SELF-POLICING, RATE LIMITING ONLINE FORUMS**(57) Abstract**

A method and means for self-policing and automatically rate-limiting multiple-user online forums. The preferred embodiment of the invention includes a set of rules that permit users to censure other users. A censored user has one or more "privileges" (which may include access to the online computer system or the effective message rate of a communications connection) taken away or diminished; the privilege is gradually restored if the censoring user behaves. The censoring participant and other pertinent participants are notified of the effect of the censoring on the censored user. In another aspect of the invention, the online computer system automatically tracks the rate at which a user sends certain types of messages, and can message rate limit a user who uses too many system resources by sending a large number of messages in rapid succession.



- 1 -

**SELF-POLICING, RATE LIMITING ONLINE FORUMS****BACKGROUND**

This application is a continuation of provisional application number 60/047,235 filed May 20, 1997.

**5 Technical Field**

This invention relates to distributed computer services, particularly computer services having online forums.

**Background Information**

An online forum is a communications interchange in which  
10 people may communicate with others through successive  
electronic transmissions between respective computer  
systems. An online forum, or any other type of  
distributed computer services, may be implemented on a  
distributed computer system such as that shown in FIG. 1.  
15 Forum participants (equivalently, users of the computer  
services) typically are scattered across a large  
geographical area and communicate with one or more  
central server systems 100 through respective client  
systems 102 (e.g., a personal or laptop computer). In  
20 practice, the server system 100 typically will not be a  
single monolithic entity but rather will be a network of  
interconnected server computers, possibly physically  
dispersed from each other, each dedicated to its own set  
of duties and/or to a particular geographical region. In  
25 such a case, the individual servers are interconnected by  
a network of communication links, in known fashion. One  
such server system is "America Online" from America  
Online Incorporated of Virginia.

Each client system 102 runs client software that allows  
30 it to communicate in a meaningful manner with  
corresponding software running on the server system 100.

- 2 -

The client systems 102 communicate with the server system 100 through various channels, such as a modem 104 connected to a telephone line 106 or a direct Internet connection using a transfer protocol such as TCP/IP  
5 (Transfer Control Protocol/Internet Protocol). The server system 100 is responsible for receiving input from the client systems 102, manipulating the collective body of input information (and possibly information from other sources) into a useful format, and retransmitting the  
10 formatted information back to one or more clients 102 for output on an output device, such as a display screen.

Referring to FIG. 2, one type of forum is a "chat room" 200, in which the various participants 204 (e.g., "Allens9," "JOSHUA ALEX," etc.) may enter text which  
15 appears in a scrolling text window 202 on each participant's computer display screen. In the example in FIG. 2, the chat room 200 has 22 participants whose identities (or "screen names") are listed in a scrolling window 210. A participant 204 may respond to the comment  
20 of another participant 204 by entering a line of text in an edit box 206 and activating (e.g., by clicking with a pointer device, such as a mouse) a SEND button 208. In response, the text in the scrolling text window 202 scrolls upwards and the newly entered line of text is  
25 displayed at the bottom of the scrolling text window 202. In the illustrated example, the last participant to enter a comment was JOSHUA ALEX, who typed "TEXAS."

The chat room 200 shown in FIG. 2 is "public", meaning that it has multiple participants who were placed in the  
30 chat room by the computer-service provider and who most likely never have met or conversed with one another before. A comment by a participant in a public forum may be seen by all of the participants of the chat room. If a participant desires some privacy, that participant may

- 3 -

"open" and enter a "private" chat room (for example, by clicking on a SETUP button 212), and thereafter invite one or more other participants to enter the private chat room. Once in a private forum, participants may  
5 communicate with one another without fear that uninvited participants will be able to see their comments.

When a participant in a forum, whether public or private, makes a comment that others in the forum regard as offensive, in poor taste, wildly incorrect, or otherwise  
10 objectionable, the offending participant most likely will be "flamed" by one or more of the other participants. A "flame" is a reprimand or other stringent response directed at the offending party. One purpose behind flaming another participant is to dissuade the offender,  
15 through embarrassment or intimidation, from making further objectionable comments. In this manner, if the offending user chooses to curb his or her behavior in response to the flaming, a forum may be crudely regulated or "policed" by the forum's participants. However, the  
20 offending participant may continue to behave in an objectionable manner. Further, a participant who overly "flames" other participants may also be objectionable. Accordingly, participant policing of forums does not always work well. In such cases, offended participants  
25 may drop out of "flame-filled" forums, and/or the online service must devote resources to actively police problematic participants.

Other objectionable behavior includes sending one or more messages to "spoof" other users as to the sender's  
30 identity in order to try to get confidential information (e.g., credit card numbers or passwords) sent in response (sometimes called "password fishing").

- 4 -

Another problem that can arise in online systems is "resource hogging", where a participant uses features such as broadcast or multi-cast messaging to send a large number of messages to other users in a short period of  
5 time (sometimes called "spamming"). Such resource hogging deprives other users of server resources, and can slow an online system response time to undesirable levels.

Accordingly, the inventor has determined that there is a need for a better way to police recalcitrant participants  
10 in online forums and to reduce spamming. The present invention provides a method and means for accomplishing this goal.

**SUMMARY**

The invention provides a method and means for self-  
15 policing and automatically rate-limiting multiple-user online forums. The preferred embodiment of the invention includes a set of rules that permit users to censure other users. A censored user has one or more "privileges" (which may include access to the online computer system  
20 or the effective message rate of a communications connection) taken away or diminished; the privilege is gradually restored if the censored user behaves. The censoring participant and other pertinent participants are notified of the effect of the censoring on the  
25 censored user. In another aspect of the invention, the online computer system automatically tracks the rate at which a user sends certain types of messages, and can message rate limit a user who uses too many system resources by sending a large number of messages in rapid  
30 succession. The amount of rate limiting may be a function of the amount of censure that has been applied to the user being limited.

- 5 -

Advantages of this invention may include one or more of the following. The techniques described here enable a multiple-user online service (e.g., a chat room or other forum) to be self-policing. Access to the computer  
5 service for a particular user is automatically regulated by the computer based on input from other users concerning the conduct of the user under consideration, or based on the message rate of a user. Users of a computer-based system have the ability to sanction a  
10 misbehaving user and thereby have the offending user's access to the system denied or curtailed. Unlike the conventional "flaming" approach to policing, which typically fails because it relies on the misbehaving user curbing his or her own behavior or results in retaliatory  
15 "flaming", the policing techniques described here are based on predetermined rules and operate automatically in response to votes cast by other users of the computer system. The offending user has no choice in the matter once he or she chooses to misbehave. Because the  
20 operation of these policing techniques are automatic, the computer-based service provider need not expend personnel time and resources to police public forums, direct user communications, and the like.

The details of one or more embodiments of the invention  
25 are set forth in the accompanying drawings and the description below. Other features, objects, and advantages of the invention will be apparent from the description and drawings, and from the claims.

#### DESCRIPTION OF DRAWINGS

30 FIG. 1 shows a prior art distributed computer system of the type used for providing online computer services.

FIG. 2 is a screen shot showing an example of a prior art online computer forum.

- 6 -

FIG. 3 is a flowchart of a basic embodiment of the self-policing aspect of the invention.

FIG. 4 is a flowchart of a basic embodiment of the rate-limiting aspect of the invention.

5 Like reference numbers and designations in the various drawings indicate like elements.

#### **DETAILED DESCRIPTION**

The prospect of millions of concurrent users of an online computer system puts great pressure on the ability of the  
10 system provider to police abusive participants. Accordingly, the invention provides a self-policing online environment that makes participants responsible for their behavior. That is, other participants can express an opinion about that behavior, and have that opinion affect  
15 the offending user in a negative way and be "visible" to other participants. Providing such a self-policing user feedback mechanism lowers the cost of providing online services to users and provides a more "user-friendly" environment for participants.

20 In another aspect of the invention, the online computer system automatically tracks the rate at which a user sends certain types of messages, and can "rate limit" a user who "hogs" too many system resources by sending a large number of messages (or messages of selected types)  
25 in rapid succession.

#### *Participant Self-Policing*

FIG. 3 is a flowchart of a basic embodiment of the self-policing aspect of the invention. Initially, user A generates an event I, such as a message (STEP 300). Event  
30 I is transmitted to some number of other participants, such as users B and C in a public forum on an online

- 7 -

computer system (STEP 302). However, event I may be a direct communication between user A to user B, such as by use of the Instant Message™ feature of the America Online computer service. User B may feel that event I is  
5 "objectionable" or "evil" (a subjective decision by B), and therefore generates a "vote" against the behavior of user A by sending a special type of response message - an "evil" event E (STEP 304). In the preferred embodiment, a user cannot "evil" another participant except in response  
10 to a message from the participant, or otherwise in cases where the participant's actions effect the user doing the "eviling." In order to reduce retaliation, users cannot "evil" other participants directly in response to being "eviled".

15 The online system receives event E and checks a database to see whether user B has the right to "evil" user A (STEP 306). For example, this step can limit users who themselves have been restricted from "eviling" other users.

20 If user B does not have the right to "evil" user A, then user B is notified of the affect of user B's action on user A's "evil index" - a value of how restricted user A is with respect to use of resources on the online system (STEP 308). In this case, user B's action has no affect  
25 on user A.

If user B does have the right to "evil" user A, then user A's "evil index" is modified by a suitable amount, which will affect further actions by user A (STEP 310), as described below. The amount of modification can be based  
30 on the past behavior of user A, or of users A and B. The amount of modification may also depend on the type of "eviling" asserted by user B. For example, user B may wish to assert an "eviling" event E anonymously rather

- 8 -

than be identified. Anonymous "eviling" may be accorded lesser weight. That is, an anonymous "eviling" response may count as fewer evil "votes" than it would if the eviling user's identity was revealed. In one embodiment, 5 an "eviling" user may set up an automatic "eviling" response to all communications from another specific user. In this case, the effect of the "eviling", which is intended to discourage communication from the specific user, is given very low weight (equivalently, counts as 10 very few evil "votes") because of its automatic invocation.

User A's "evil index" is "broadcast" to other users who "know about" user A (STEP 312). For example, such users might be all of the participants in a private chat room, 15 or all visitors to a public chat room. In one embodiment, a modified "evil index" is reported to each user in a current chat room's user list and to each user who has the "eviled" users on his or her "buddy list". Thus, an "eviled" user is publicly chastised.

20

Thereafter, user A's "evil index" begins to gradually "decay" back to normal over time (STEP 314). Such decay 25 may be linear, exponential, step-wise, or some other function. Meanwhile, user B is notified of the affect of user B's action on user A's "evil index" (STEP 308). In this case, user B's action has had an affect on user A.

Some of the steps noted above may be done in different 30 order without substantially changing the effect of the process. For example, STEPS 312, and 314, and 308 may be performed in different order.

- 9 -

A basic "penalty" for having a non-normal "evil index" is denial of access to a forum or the online service until the user's "evil index" has decayed back to normal. In a more refined embodiment, a user's "evil index" affects a rate limit which governs a user's ability to send (and/or receive) messages. This feature allows other participants to "evil" a user who "flames" or "spams" them, and thus reduce the rate at which the recalcitrant user can send and/or receive messages. A description of rate limiting is set forth below.

A server database (which may be centralized or distributed) stores a user's "evil index". A user's "evil index" can be maintained in a user-record as a global total, or by forum, or both. The value of each "evil index" can be used to control the user's ability to log on to the online computer system or access selected forums, and/or the effective rate of message or file transmissions.

Other rules may applied in different embodiments. For example, the following rules can be implemented:

A user must be in a forum (e.g., a chat room, but including direct communication between users, such as the Instant Message™ feature of the America Online computer service) for a specified amount of time before being allowed to "evil" another user in that forum; this reduces "hit-and-run flaming". The minimum amount of time may vary from forum to forum, and from user to user (e.g., a forum "sysop" may be immune to "eviling"). A user's user-record in the server database thus would record a time-of-entry for a forum. For example, a user's time-of-entry to a forum would be compared to the current time in STEP 306 of FIG. 3 to

- 10 -

determine if the user had the right to "evil"  
another participant.

5 A user must be in a forum for a specified amount  
of time before being allowed to "evil" another  
user who has been in that forum for a longer  
period of time. For example, a user's time-of-  
entry to a forum would be compared to the time-of-  
entry of another participant in STEP 306 of FIG. 3  
10 to determine if the user had the right to "evil"  
that other participant. The specified amount of  
time may vary from forum to forum, and from user  
to user.

15 A user's eviling response may be accorded  
increased weight (equivalently, counted as extra  
"evil" votes) based on the "seniority" of the  
"eviling" user. Each additional unit of time  
spent in a forum could enhance a user's  
"seniority," thereby allowing long term user's  
more "eviling" power than new-comers. That is,  
20 being "eviled" by an "old-timer" can have a  
greater than normal affect on modification of a  
chastised user's "evil index" in STEP 310 of FIG.  
3. A user's user-record in the server database  
thus would record a total-time-of-access for each  
25 forum, which could be a "lifetime" total (such  
that users who had belonged to the online computer  
service for longer periods would have greater  
"eviling" power than relatively new subscribers to  
the service) , a running total for a selected  
30 period of time, or a session total. The time  
period for accumulating extra votes and the  
"eviling" value of extra votes may vary from forum  
to forum, and from user to user.

- 11 -

5 A user may be limited in the number of "eviling" votes that can be cast in any one online session or time period (e.g., a day or week). A user's user-record in the server database thus would record the number of "eviling" votes cast globally or by forum. For example, a user's number of previous "eviling" votes cast could be examined in STEP 306 of FIG. 3 to determine if the user had the right to "evil" another participant.

10 The amount of modification of a user's "evil index" in STEP 310 of FIG. 3 after being "eviled" is preferably non-linear, so that the first few times a user is "eviled" has little effect (and possibly no effect until some threshold level of  
15 accumulated "eviling" votes from several participants are asserted against the user) on the user's "evil index" (everyone makes a few mistakes). However, for subsequent times that a user is "eviled", the user's "evil index" preferably is modified to a much greater degree, increasing the penalty for recidivism. The  
20 determination of how much to modify a user's "evil index" can be based on the user's global "evil index" (i.e., a "rap sheet" evaluation) or forum "evil index" (i.e., a "fresh start" approach). A user's user-record in the server database would thus record the number of times the user has been "eviled", either globally or by forum, for  
25 example. The recorded number might be a "lifetime" total or a running total for a selected period of  
30 time. The amount of "evil index" modification may also vary from forum to forum, and from user to user.

- 12 -

5 As an example, each user may be given an initial "evil index" of 100. A first "offense" may reduce that value to 95; a second "offense" reduces the value to 85; third and subsequent offenses reduce the current value by 15 units. Alternatively, each user is given an initial "evil index" of 0 and "eviling" increases that value. If a range of 100 is used, a user's "evil index" may be regarded as a "percentage of evil", up to 100% "evil".

10 The decay rate for returning a user's modified "evil index" to normal can vary from forum to forum. For example, the decay in chat rooms (where most flaming occurs) may be less than in other forums. Using the example from immediately above,  
15 the user's "evil index" may "decay" back to "normal" at the rate of 2 units per hour in a chat room, but at the rate of 5 units per hour in other forums.

20 A user may query a forum to determine the rate of "eviling" within the forum as a measure of how much members of the forum "evil" one another. A high rate may indicate that one or more members are misbehaving. The system keeps track of all users participating in a forum, so the current  
25 "eviling" rate is a time-weighted average of the number of "eviling" votes cast by the current participants. In an alternative embodiment, a user may query a forum to determine the accumulated value of the "evil index" of all current users.  
30 This measures how often the users have misbehaved in the past, and can serve as an estimate of the likelihood that the current users will misbehave in the future. The system keeps track of all users participating in a forum, so the total "evil

- 13 -

index" is the sum of the pertinent forum "evil index" for each participant.

5 In some embodiments, a user who has been "eviled" has a lessened ability to "evil" other users (thus reducing retaliation). However, some online systems implement message types, such as broadcast or multi-cast messages or self-repeating messages, that are more frequently used than other message types to flame or spam other participants. In some  
10 embodiments, it may be desirable to allow all users (even those with a modified "evil index", and thus possibly restricted in usage rights) to "evil" the originator of such message types. Such a determination would be made in STEP 306 of FIG.  
15 3 by examining the message type of event I sent by the originating user. The amount of modification of such an originator's "evil index" can be set to be greater than for flaming to discourage use of such message types for spamming.

20 When a user's "evil index" is modified, the user is notified that privileges, such as rate of messaging, have been limited. In one embodiment, a message is sent from the server that contains the user's current "evil index" to the "eviled" user, and optionally includes the current  
25 rate of decay back to normal. Such information allows a wide range of feedback to be presented to the user about his or her ability to interact with the online computer system. For example, a graphical "power meter" or "power bar graph" can be used to indicate the "evil index" of  
30 the user. For instance, a color-coded bar graph can be divided into (1) a green zone to represent that the user's "evil index" is normal; (2) a yellow zone to represent that the user's "evil index" has been modified slightly; (3) a red zone to represent that the user's

- 14 -

"evil index" has been modified significantly; and (4) a black zone to represent that access or message privileges have been suspended for a time. However, other methods of informing a recalcitrant user of his or her "evil index" can be used.

In one embodiment, a user can "rehabilitate" his or her "evil index" by visiting advertisements displayed on the online computer system; "visits" can be determined by requiring the user to "click" on an advertisement. The user is rewarded with more "power" by adjusting the value of the user's "evil index" more towards normal.

#### *Automatic Rate Limiting*

In one embodiment, both input and output messaging rates of a user are limited based on the behavior of the user and/or available system resources. Such rate limiting can stop malicious users and errant client computers from consuming more than a fair share of online system resources. However, preferably the rate limiting system is weighted to tolerate brief bursts of messaging activity while penalizing unacceptably large rates of messaging. Rate limiting can also limit aggregate input to a server to a level at which the system is reasonably loaded under normal conditions. Rate limiting can also be combined with "eviling" by automatically adjusting a users rate limit parameters based on their "evil index."

In one embodiment, input rate limiting - from user to server - is conducted on a per user connection basis, and within a user connection on a per-message type bases (that is, rate limits for different types of messages may be set to different values). In one embodiment, rate limiting for a user is achieved in accordance with the following algorithm:

- 15 -

- (1) Define  $A$  as the running average of inter-message time gaps for the last  $N$  messages of selected types that the user has attempted to send; a system selected value  $I$  is used as the first value for  $A$ . Calculation of  $A$  can be done, for example, on a forum basis (accounting only for messages sent in the current forum), session basis (accounting for messages sent in the user's current online session), or message-count basis (accounting for the last  $N$  messages sent at any time in any forum).
- (2) If  $A$  is below a warning threshold  $W$  (indicating that the user is approaching a point of sending messages too frequently), when the user attempts to send a message, send the user a warning message but transmit the user's message.
- (3) If  $A$  is below a rate limit threshold  $R$  (indicating that the user is sending messages too frequently) when the user attempts to send a message, send the user a warning message and drop the user's message.
- (4) Repeat the above steps until  $A$  rises above a clear threshold  $C$  (indicating that the user is not sending messages too frequently), at which time the rate limiting condition is considered cleared.
- (5) If at any time  $A$  drops below a disconnect threshold  $D$ , disconnect the user.

FIG. 4 is a flowchart of a basic embodiment of the rate-limiting aspect of the invention, showing a slightly different order for the steps described above. A user attempts to send a message (STEP 400). Average  $A$  is computed (STEP 402). If  $A$  is less than a disconnect threshold  $D$  (STEP 404), the user is disconnected (STEP 406). Otherwise, if  $A$  is less than a rate limit threshold  $R$  (STEP 408), the message is dropped, the user is warned

- 16 -

that the rate limited has been exceeded, and the user is flagged as "rate limited" (STEP 410).

Otherwise, if  $A$  is less than a warning threshold  $W$  (STEP 412), a determination is made as to whether the user is  
5 rate limited (STEP 414). If not, the message is sent, but the user is warned that the rate limit is being approached (STEP 416). Otherwise, the message is dropped and the user is warned that the rate limited has been exceeded (STEP 418).

10 If  $A$  is not less than the warning threshold  $W$  (STEP 412), and  $A$  is less than a clear threshold  $C$  (STEP 420), a determination is made as to whether the user is rate limited (STEP 421). If not, the message is sent (STEP  
15 423). Otherwise, the message is dropped and the user is warned that the rate limited has been exceeded (STEP 418).

Finally, if  $A$  is not less than the clear threshold  $C$  (STEP 420), the rate limit flag for the user is cleared (STEP 422) and the message is sent (STEP 423).

20 The rate limiting algorithm supports several "tunable" parameters:

- The running average of inter-message time gaps -  $A$
- The number of message receptions over which  $A$  is calculated -  $N$
- 25 • An initial average -  $I$
- A clear threshold -  $C$
- A warning threshold -  $W$
- A rate limit threshold -  $R$
- A disconnect threshold -  $D$

30 In one embodiment, the values for  $C$ ,  $W$ ,  $R$ , and  $D$  are selected such that  $C > W > R > D$ . The initial average

- 17 -

rate  $I$  can be weighted to increase the algorithm's tolerance of bursts of activity, such as "chatty startups" when a user joins a forum and sends several messages in succession. The threshold rates can be set  
5 globally for all user's, or "tuned" for each user.

In one embodiment, the difference between the clear threshold  $C$  and the rate limit threshold  $R$  can be "tuned" by the online computer system to alter the interval between commencement of rate limiting for a user and the  
10 resumption of normal activity. The difference between  $C$  and  $R$  may be automatically increased, for example, if the user sends excessive "spamming" type messages at too great a rate. Further, the threshold values  $C$ ,  $W$ ,  $R$ , and  $D$  can be dynamically modified by the online system as a  
15 way of limiting system resource usage if too many users are simultaneously using the system.

As noted above, when a user's message rate is limited, the user is notified. In one embodiment, a message is sent from the server to a rate limited user that contains  
20 values for each of the parameters outlined above, and another message is sent the next time that the server will accept messages from the user without restriction. Such information allows a wide range of feedback to be presented to the user about his or her ability to  
25 interact with the online computer system. For example, a graphical "power meter" or "power bar graph" can be used to indicate the "health" or "power" of the user. For instance, a color-coded bar graph can be divided into (1)  
30 effect; (2) a yellow zone to represent that the user's message rate is near the point where rate limiting would take effect; (3) a red zone to represent that message rate limiting is in effect; and (4) a black zone to represent that access privileges have been suspended for

- 18 -

a time. However, other methods can be used to inform a recalcitrant user of his or her message rate limit status.

In one embodiment, a user can "rehabilitate" his or her message rate limit status by visiting advertisements displayed on the online computer system. The user is rewarded with more "power" by adjusting the parameters.

In one embodiment, output rate limiting - from server to user - is performed in a similar manner, with the following differences:

- 6) The system server audits the rate of outgoing error messages of selected types (e.g., RATE\_TO\_HOST), and uses the above rate limiting algorithm with  $R=D$ , to prevent excessive load due to handling users that are generating rate errors.
- 7) The system server can limit the rate of outgoing messages of selected types (e.g., those types that may be used for spamming) using the above rate limiting algorithm when  $A < R$ . In this case, a notice message can be sent to the user, to the user's intended recipient, and to a server storing the user's "evil index". The intended recipient can "evil" the sending user, or the system server can automatically modify the user's "evil index" based on the sending of excessive "spamming" type messages at too great a rate.

The algorithms described above prevent abuse of online system resources without causing undue "pain" to users, including "power users" and users who are subject to delays and input/output timing changes by non-ideal networks (which can cause otherwise normal user activity to appear to be unacceptable activity).

- 19 -

*Rate Limiting and "Eviling"*

The concepts of automatic rate limiting and modification of a user's "evil index" can be combined to provide a refined self-policing, automatic rate limiting system that can regulate an "eviled" user's ability to participate in forums without requiring total exclusion from the forum. Further, a user's current global or forum "evil index" can be used to modify the parameters used to effect automatic (i.e., "non-eviled" basis) rate limiting for the user. For example, a user's ability to interact in a forum can be rate limited by modifying the values for *A*, *C*, *W*, *R*, and/or *D* as a function of whether the user's "evil index" has been modified, or as a function of the current value of the user's "evil index".

Conversely, the amount of modification of a user's "evil index" can be a function of the current value of *A*. Thus, conceptually, a user's permitted message rate  $R_p$  is a function of the user's "evil index" *EI* plus the user's rate of attempted message activity *A*:  $R_p = f(EI) + g(A)$

Conceptually, each user's user-record thus may look like the following table:

Forum ID	"evil index"	decay rate	# times "eviled"	# times "eviled" others	time-of-entry	total-time of-access	<i>A</i>	<i>C</i>	<i>W</i>	<i>R</i>	<i>D</i>
Global value											
Forum 1 value											
Forum 2 value											
<i>etc.</i>											

Separate "lifetime" and "session" records may be kept where useful. Of course, other or different information may be recorded for each user, and other ways of organizing such data may be used. Further, users may be

- 20 -

assigned a "class" designation (such as "sysop" (system operator), corporate users, paying members, non-paying members, etc.) which can be given weight in making a determination of right to "evil" or absolute or relative immunity from being "eviled" in STEP 306 of FIG. 3.

In setting rate parameter values for a user, a global rate table can be used which is indexed by a user's "evil index"; that is, message rate parameters are a function of the user's global or forum "evil index". A conceptual example of one such table might be the following ( $\Delta A$  represents an optional adjustment to be applied to the calculated value of A):

"evil index" value	$\Delta A$	C	W	R	D
0-20					
21-40					
41-60					
61-80					
81-90					
91-95					
96-100					

The values for A, C, W, R, and D in this table can be defined globally or by forum. If by forum, the values can be defined in absolute terms or as relative offsets to a global table. Multiple tables of this form also can be defined, indexed by message type, so that misuse of certain types of messages are punished more than misuse of other message types.

Alternatively, a non-tabular implementation could be used instead of the global rate table described above. For example, a user's rate parameter values ( $\Delta A$ , C, W, R, D)

- 21 -

could be defined as a function describing a relation between these quantities.

*Implementation*

The methods and mechanisms described here are not limited  
5 to any particular hardware or software configuration, but rather they may find applicability in any computing or processing environment used in connection with online computer services.

The invention may be implemented in hardware or software,  
10 or a combination of both. However, preferably, the invention is implemented in computer programs executing on programmable computers each comprising at least one processor, at least one data storage system (including volatile and non-volatile memory and/or storage  
15 elements), at least one input device, and at least one output device. Program code is applied to input data to perform the functions described herein and generate output information. The output information is applied to one or more output devices, in known fashion.

20 Each program is preferably implemented in a high level procedural or object oriented programming language to communicate with a computer system. However, the programs can be implemented in assembly or machine language, if desired. In any case, the language may be a compiled or  
25 interpreted language.

Each such computer program is preferably stored on a storage media or device (e.g., ROM or magnetic diskette) readable by a general or special purpose programmable computer, for configuring and operating the computer when  
30 the storage media or device is read by the computer to perform the procedures described herein. The inventive system may also be considered to be implemented as a

- 22 -

computer-readable storage medium, configured with a computer program, where the storage medium so configured causes a computer to operate in a specific and predefined manner to perform the functions described herein.

- 5 A number of embodiments of the present invention have been described. Nevertheless, it will be understood that various modifications may be made without departing from the spirit and scope of the invention. Accordingly, other embodiments are within the scope of the following claims.

- 23 -

**CLAIMS**

What is claimed is:

1. A method, performed on a computer system, of regulating a user's access to a computer-based service, the method comprising the steps of:
  - (a) receiving input about a first user from at least one other user of the computer-based service;
  - (b) evaluating the received input; and
  - (c) modifying the first user's ability to access the computer-based service based on a result of such evaluating.
2. The method of claim 1, in which the step of receiving input comprises gathering a vote from at least one other user of the computer-based service, the vote being cast in response to action taken by the first user.
3. The method of claim 2, in which the step of evaluating comprises applying predetermined access limiting criteria to the cast vote.
4. The method of claim 2, in which the step of evaluating further includes determining whether each other user is permitted to cast a vote against the first user.
5. The method of claim 4, in which the step of determining further includes evaluating whether a voting user has been accessing at least a portion of the computer-based service for a minimum amount of time.

- 24 -

6. The method of claim 4, in which the step of determining further includes evaluating whether a voting user has been accessing at least a portion of the computer-based service for a longer amount of time than the first user.
7. The method of claim 4, in which the step of determining further includes limiting the number of votes a user may cast.
8. The method of claim 2, in which the step of evaluating further includes giving extra weight to votes cast by another user against the first user if the other user has been accessing at least a portion of the computer-based service for a selected amount of time.
9. The method of claim 2, in which the step of modifying comprises denying the first user full access to the computer-based service based on the result of the evaluating.
10. The method of claim 9, further comprising the step of allowing the first user to regain full access to the computer-based service based on predetermined access resumption criteria.
11. The method of claim 1, further comprising the step of notifying selected other users that the first user's ability to access the computer-based service has been modified.
12. The method of claim 1, further comprising the step of providing an indicator of an amount of modification by all users accessing at least a portion of the computer-based service.

- 25 -

13. The method of claim 1, in which the computer-based service comprises an online public forum.
  
14. The method of claim 1, further comprising the step of permitting the first user to vote against  
5 another user based on a message type sent by such other user after the first user's ability to access the computer-based service has been modified.

- 5 15. A computer program, residing on a computer-readable medium, for regulating a user's access to a computer-based service, comprising instructions for causing a computer to:
- 10 (a) receive input about a first user from at least one other user of the computer-based service;
- (b) evaluate the received input; and
- 15 (c) modify the first user's ability to access the computer-based service based on a result of such evaluating.
16. The computer program of claim 15, in which the instructions for causing a computer to receive input further comprises instructions for causing a computer to gather a vote from at least one other user of the computer-based service, the vote being cast in response to action taken by the first user.
- 20
17. The computer program of claim 16, in which the instructions for causing a computer to evaluate further comprises instructions for causing a computer to apply predetermined access limiting criteria to the cast vote.
- 25
18. The computer program of claim 16, in which the instructions for causing a computer to evaluate further comprises instructions for causing a computer to determine whether each other user is permitted to cast a vote against the first user.
- 30

- 27 -

19. The computer program of claim 18, in which the instructions for causing a computer to determine further comprises instructions for causing a computer to evaluate whether a voting user has been accessing at least a portion of the computer-based service for a minimum amount of time.
- 5
20. The computer program of claim 18, in which the instructions for causing a computer to determine further comprises instructions for causing a computer to evaluate whether a voting user has been accessing at least a portion of the computer-based service for a longer amount of time than the first user.
- 10
21. The computer program of claim 18, in which the instructions for causing a computer to determine further comprises instructions for causing a computer to limit the number of votes a user may cast.
- 15
22. The computer program of claim 16, in which the instructions for causing a computer to evaluate further comprises instructions for causing a computer to give extra weight to votes cast by another user against the first user if the other user has been accessing at least a portion of the computer-based service for a selected amount of time.
- 20
- 25
23. The computer program of claim 16, in which the instructions for causing a computer to modify further comprises instructions for causing a computer to deny the first user full access to the computer-based service based on the result of the evaluation.
- 30

- 28 -

- 24 . The computer program of claim 23, further  
comprising instructions for causing a computer to  
allow the first user to regain full access to the  
computer-based service depending on predetermined  
5 access resumption criteria.
- 25 . The computer program of claim 15, further  
comprising instructions for causing a computer to  
notify selected other users of modification of the  
first user's ability to access the computer-based  
10 service.
- 26 . The computer program of claim 15, further  
comprising instructions for causing a computer to  
provide an indicator of an amount of modification  
by all users accessing at least a portion of the  
15 computer-based service.
- 27 . The computer program of claim 15, in which the  
computer-based service comprises an online public  
forum.
- 28 . The computer program of claim 15, further  
20 comprising instructions for causing a computer to  
permit the first user to vote against another user  
based on message type sent by such other user  
after the first user's ability to access the  
computer-based service has been modified.




FIG. 1  
(PRIOR ART)

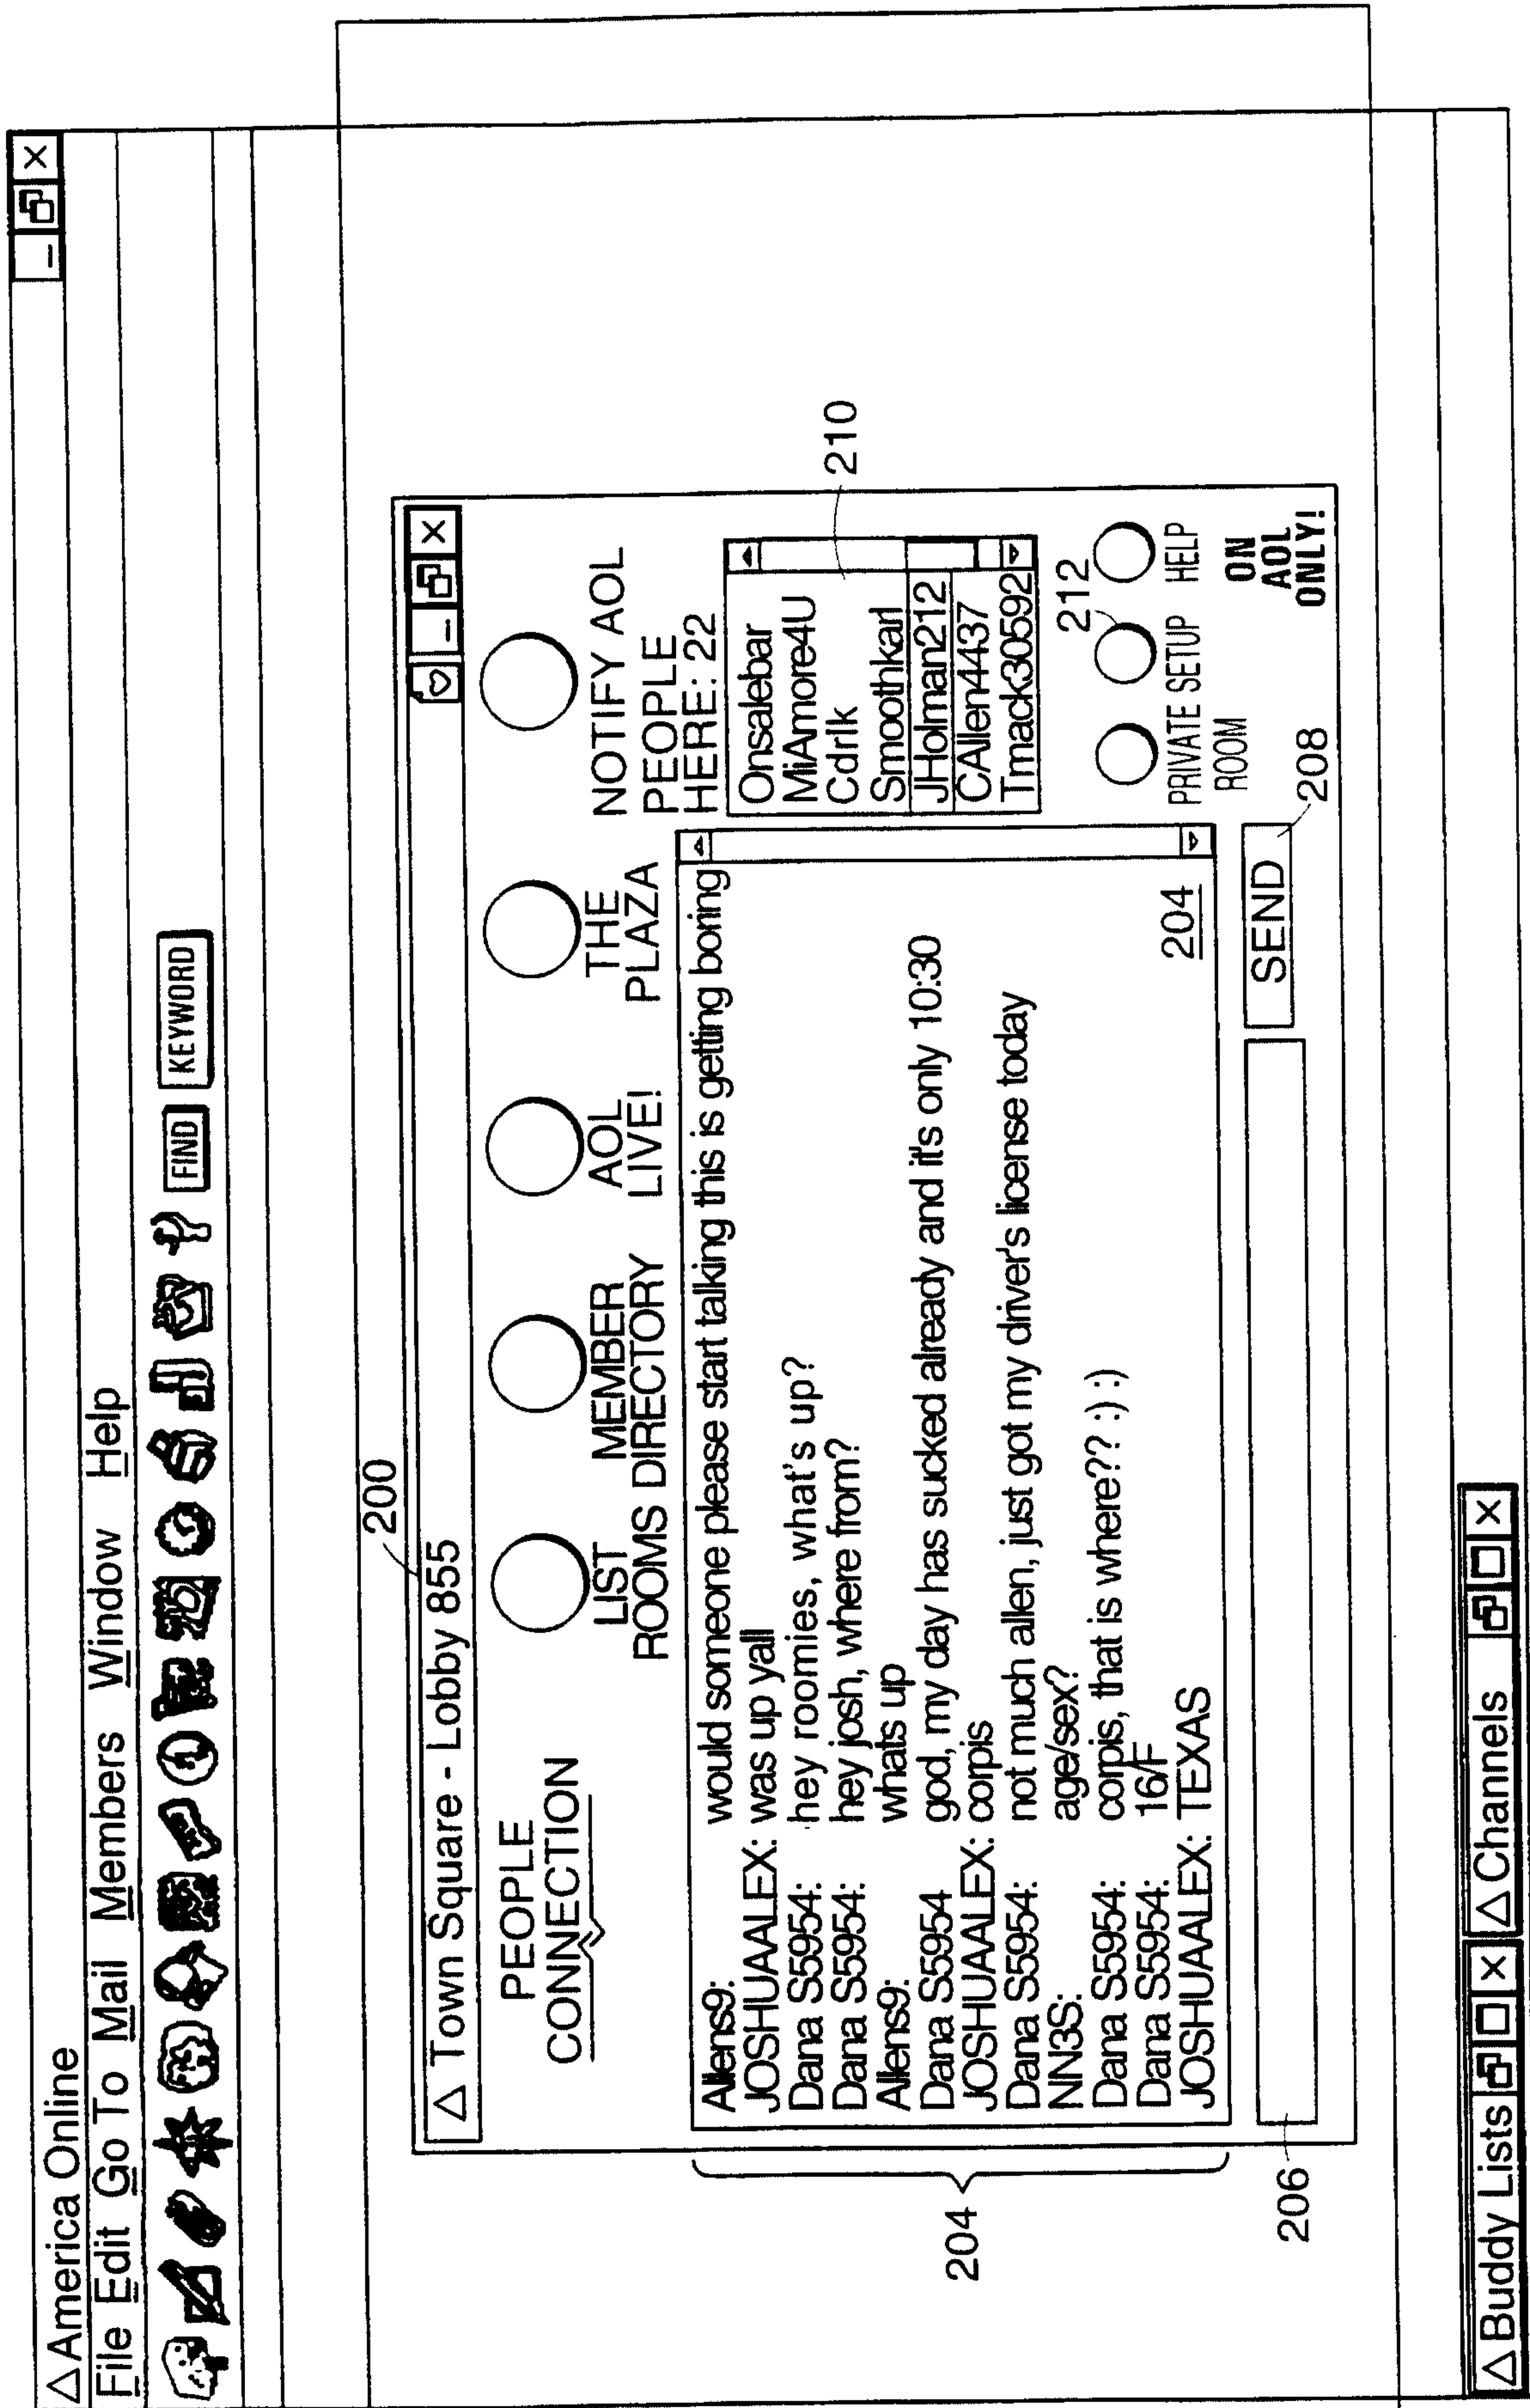


FIG. 2  
(PRIOR ART)

3/4

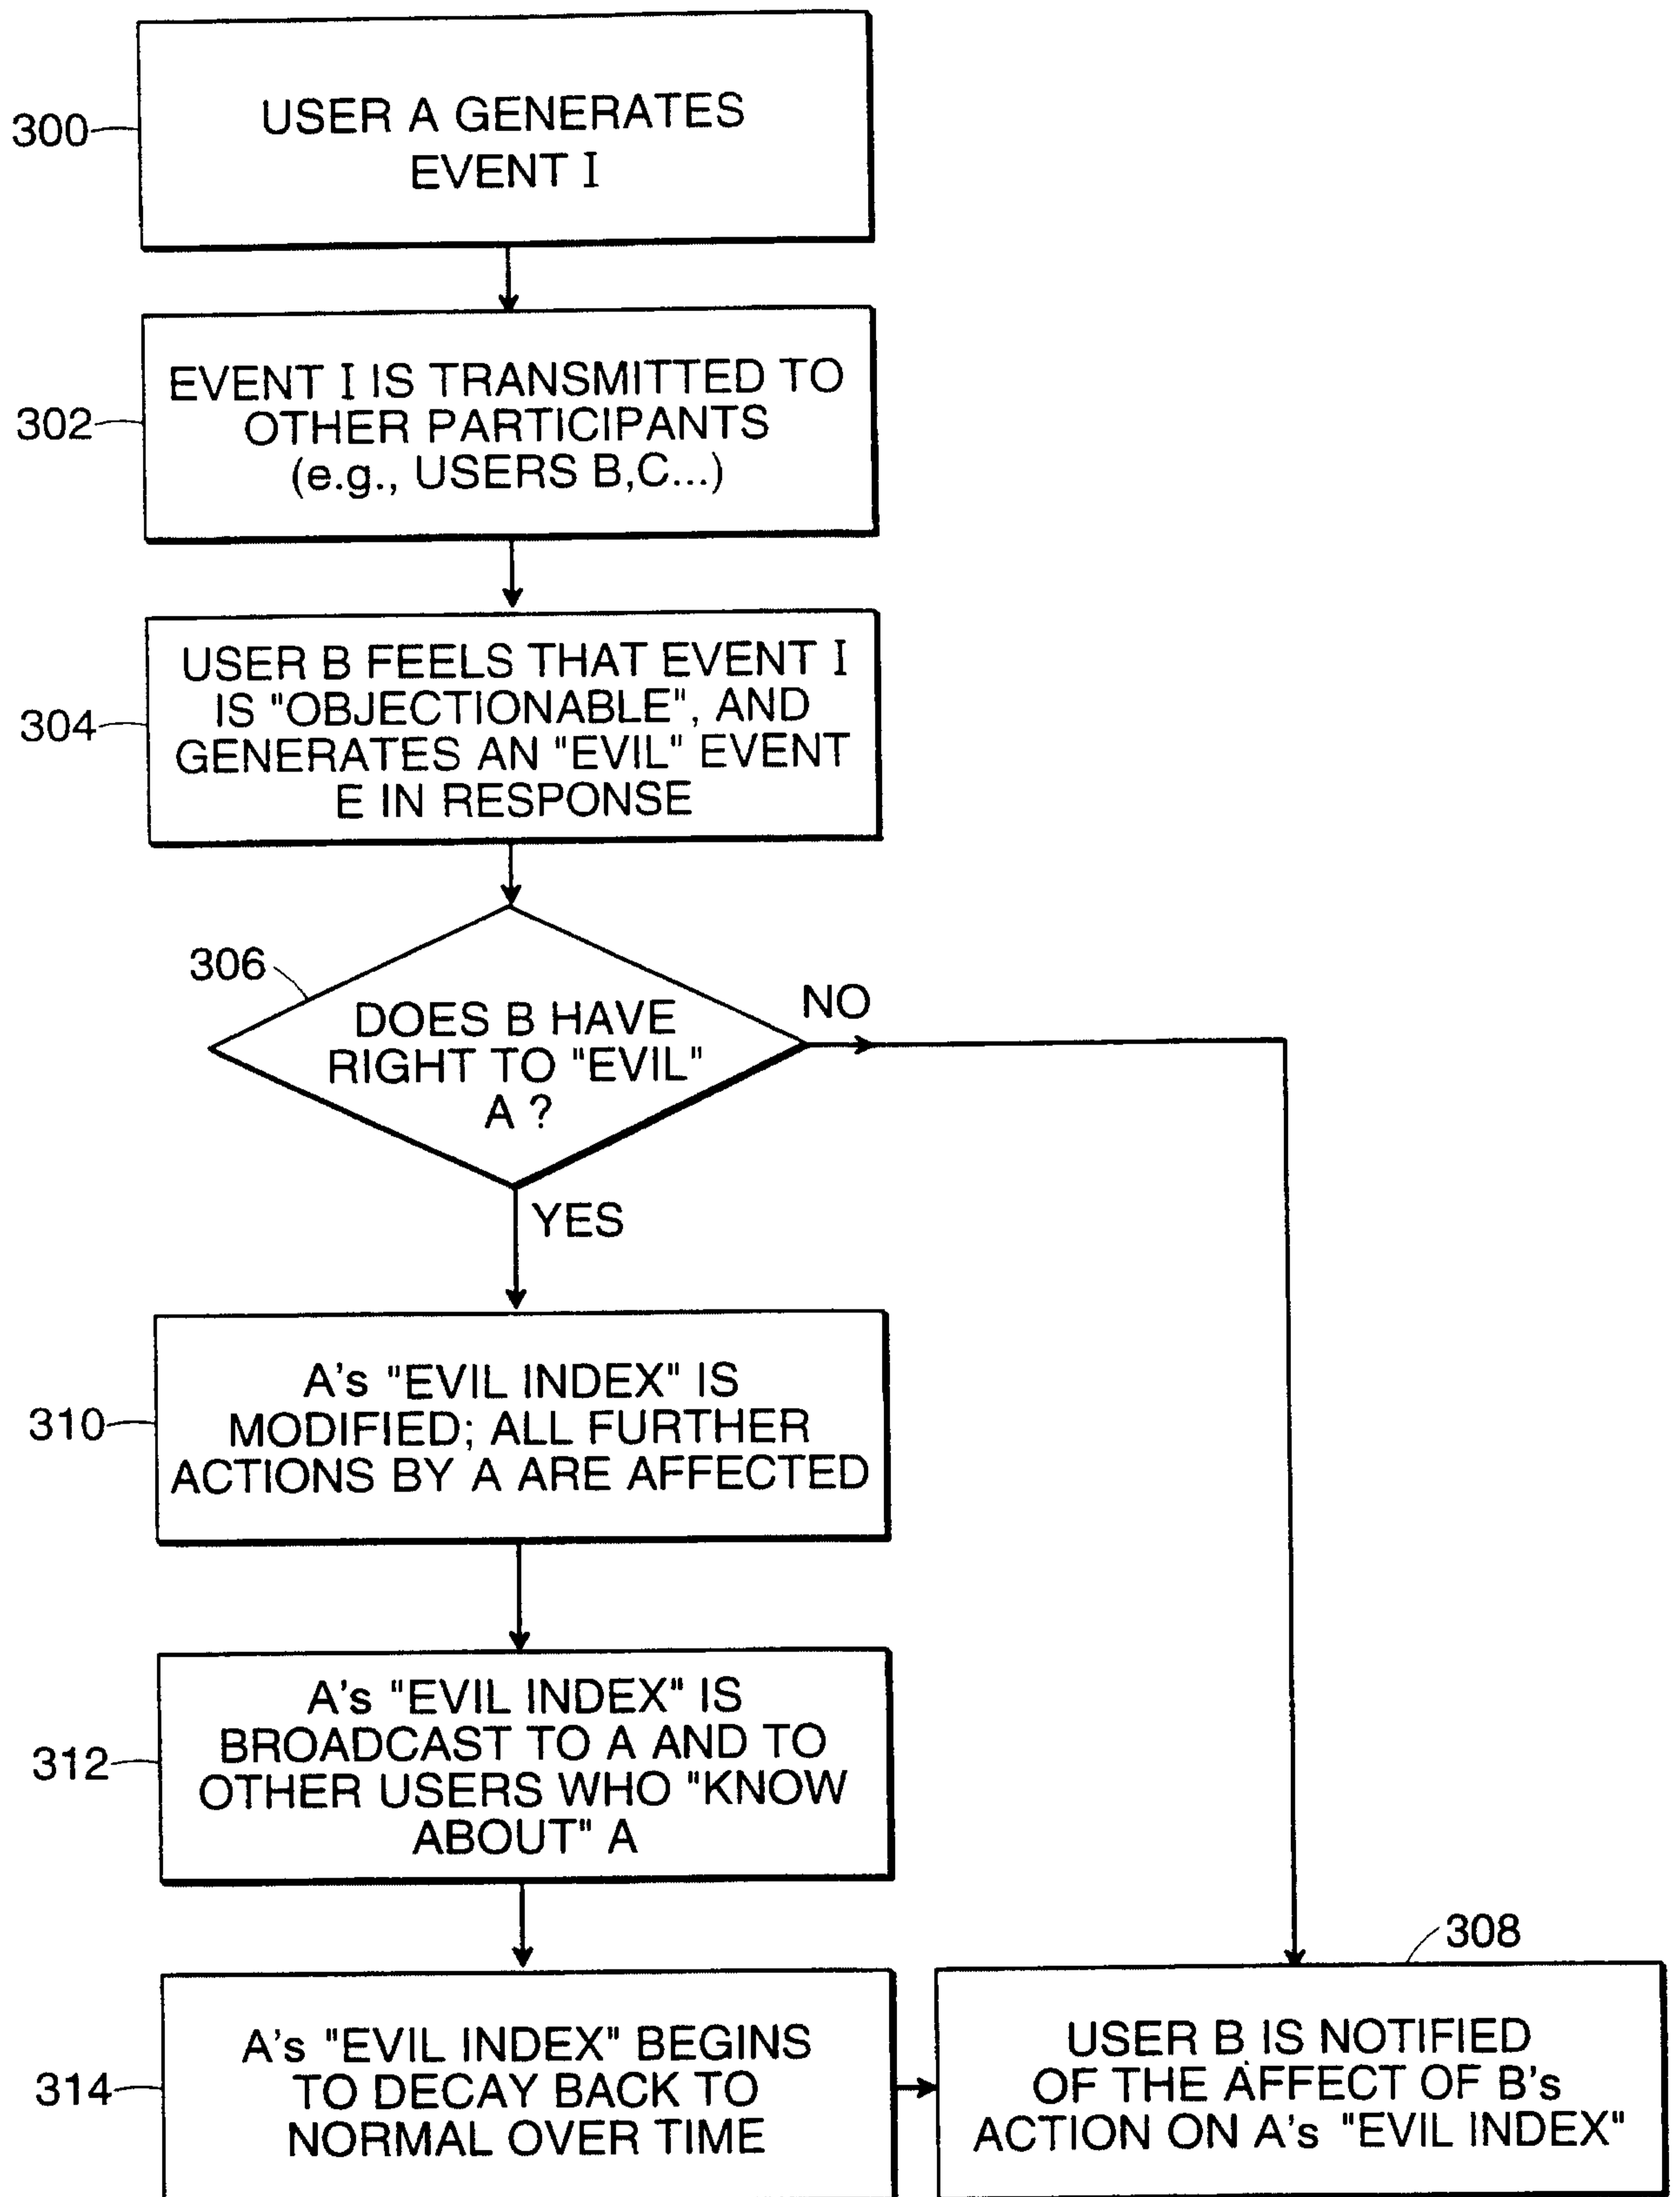


FIG. 3

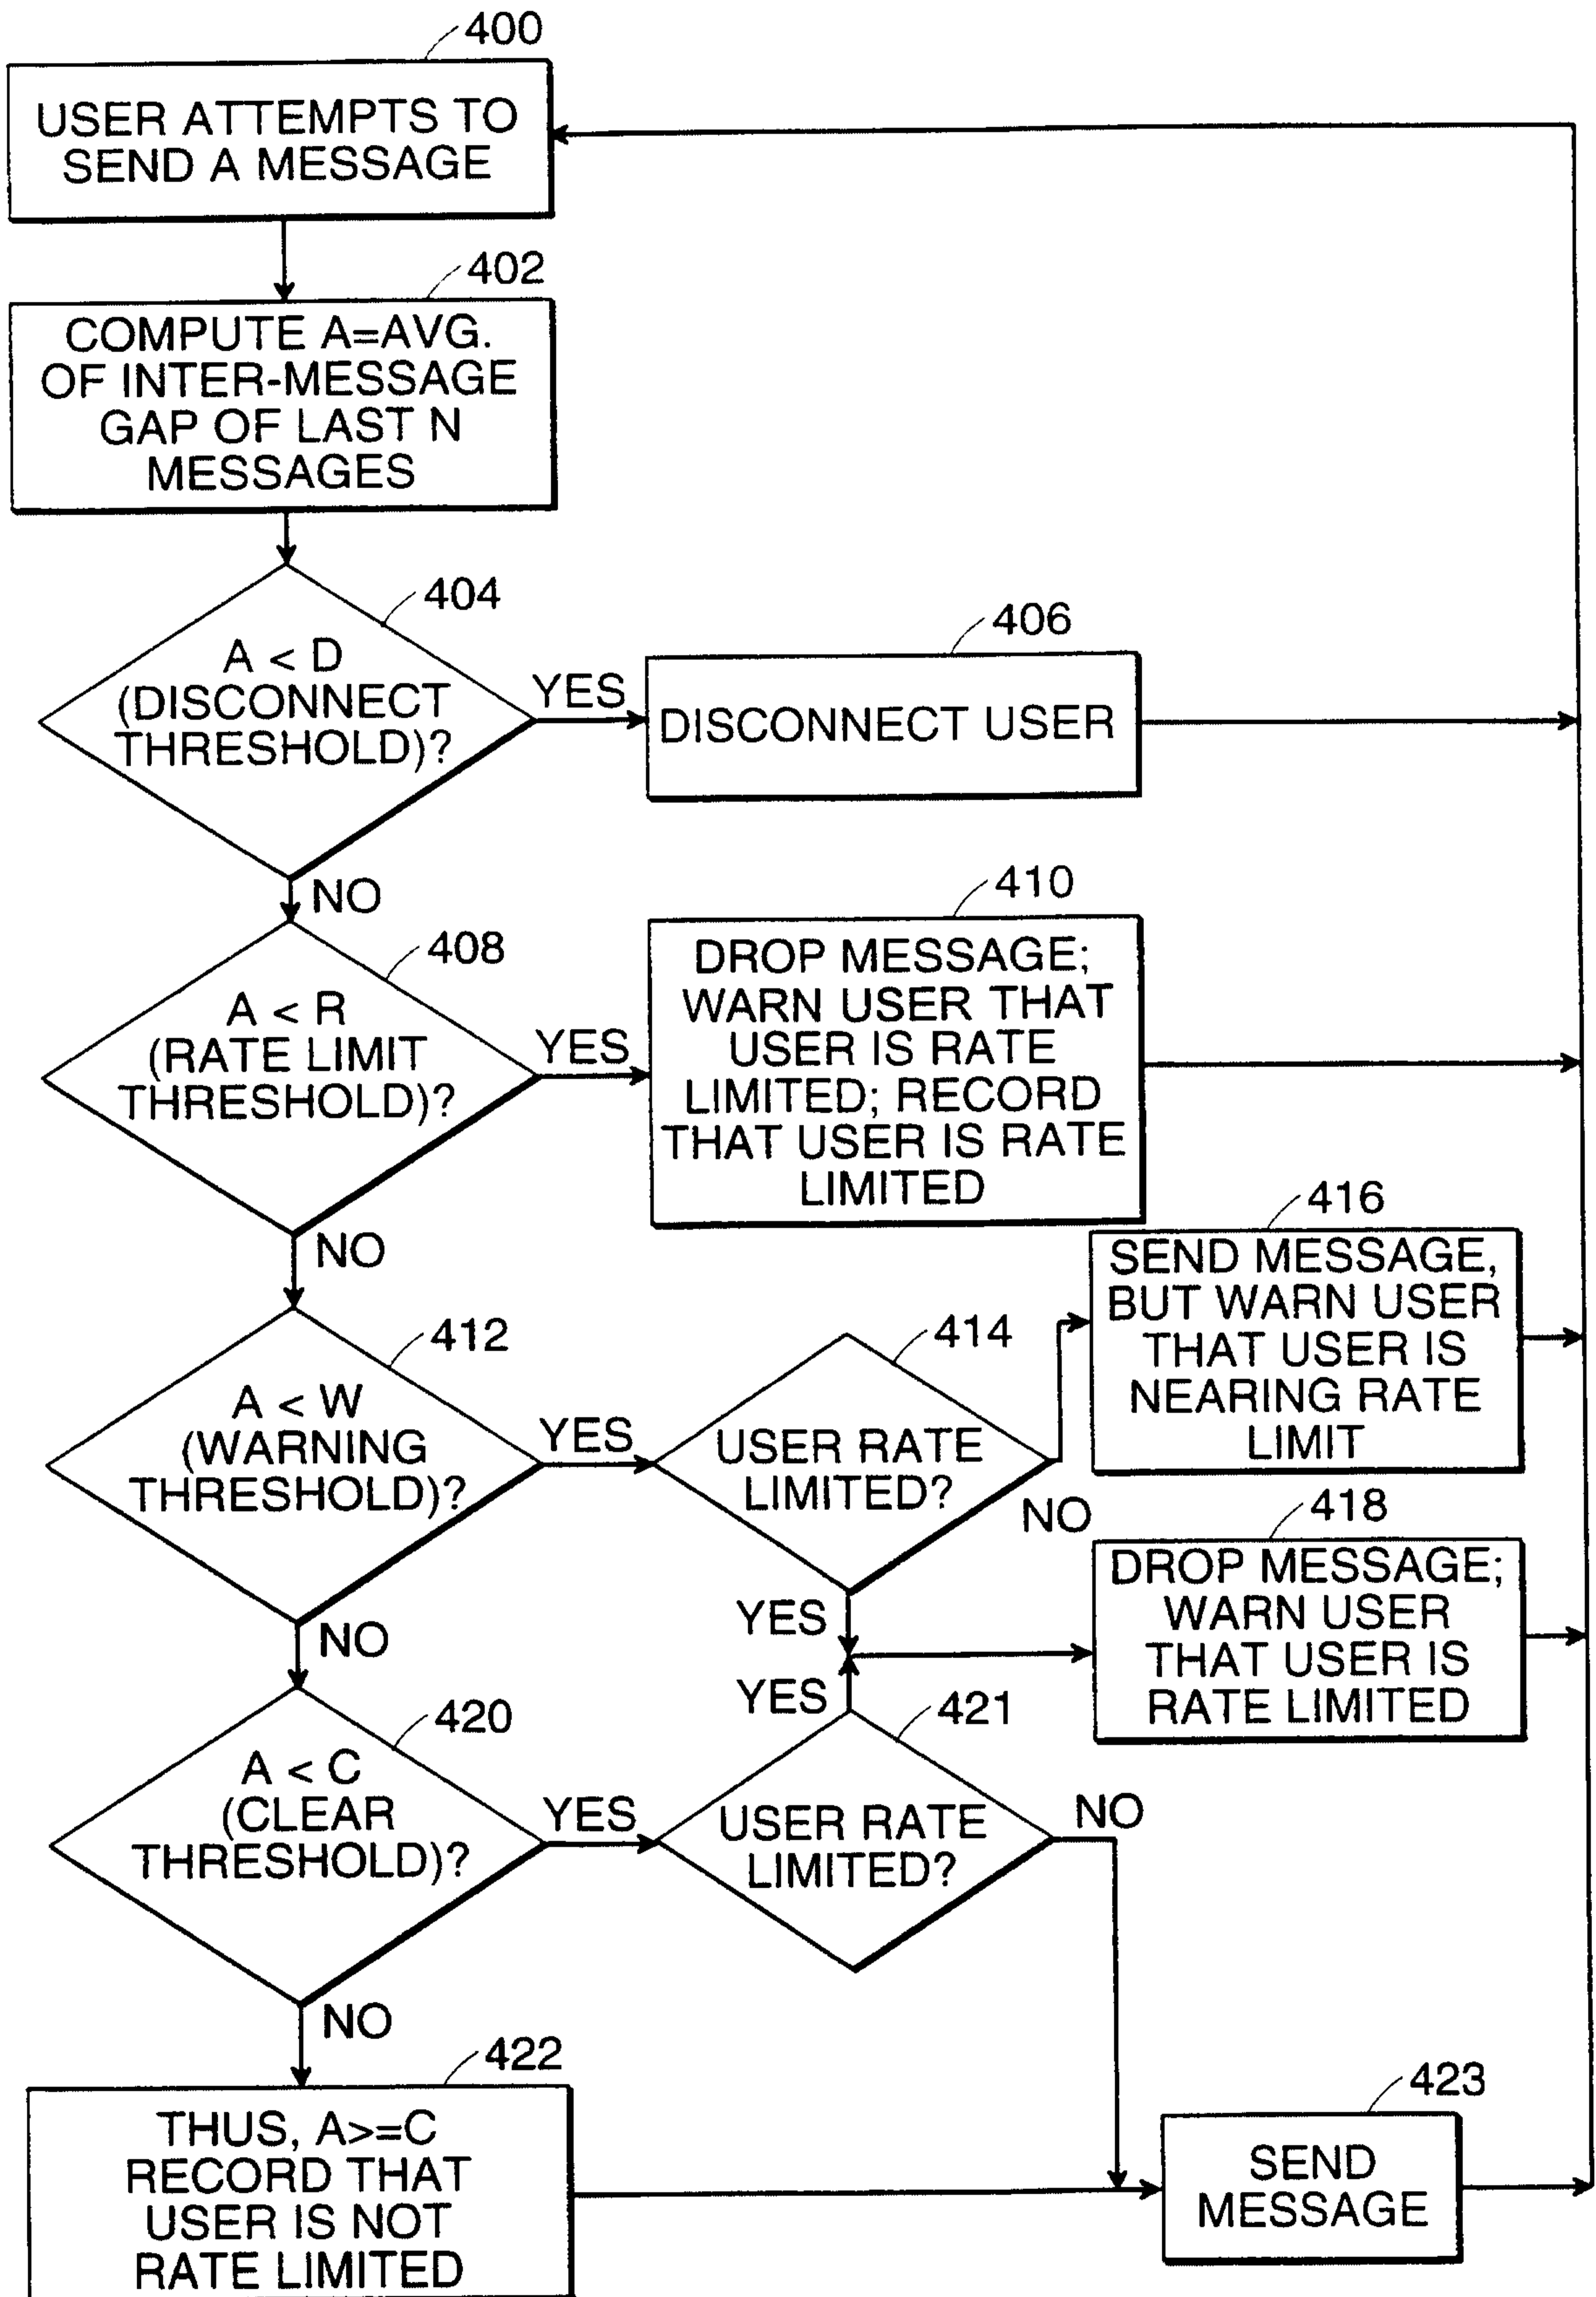


FIG. 4

