

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

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
4 February 2010 (04.02.2010)

PCT

(10) International Publication Number
WO 2010/014831 A1

- (51) International Patent Classification:
H04J 3/16 (2006.01)
 - (21) International Application Number:
PCT/US2009/052283
 - (22) International Filing Date:
30 July 2009 (30.07.2009)
 - (25) Filing Language: English
 - (26) Publication Language: English
 - (30) Priority Data:
61/084,739 30 July 2008 (30.07.2008) US
 - (71) Applicant (for all designated States except US): **VISIBLE WORLD, INC.** [US/US]; 527 West 34th Street, New York, NY 10001 (US).
 - (72) Inventors; and
 - (75) Inventors/Applicants (for US only): **HABERMAN, Seth** [US/US]; 1 West 85th Street, Apt. 6a, New York, NY 10024 (US). **NIEMEIJER, Gerrit** [US/US]; 23 Arcularius Terrace, Maplewood, NJ 07040 (US). **JANSEN, Alex** [US/US]; 101 West End Street, Apt. 5J, New York, NY 10023 (US). **VAN DE POL, Erik** [US/US]; 527 West 34th Street, New York, NY 10001 (US). **FIFE, Andrew** [US/US]; 527 West 34th Street, New York, NY 10001 (US).
 - (74) Agents: **BRODOWSKI, Michael, H.** et al.; K & L Gates LLP, State Street Financial Center, One Lincoln Street, Boston, MA 0211-2950 (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, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.
 - (84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).
- Published:**
— with international search report (Art. 21(3))

(54) Title: SYSTEMS AND METHODS FOR DYNAMIC BANDWIDTH ALLOCATION FOR ADDRESSABLE CONTENT

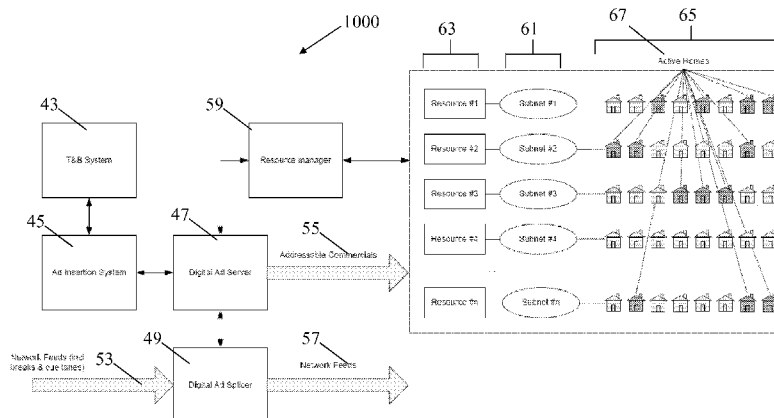


Fig. 10

(57) Abstract: Systems and methods for push-based dynamic bandwidth allocation deliver addressable, advertising content in a digital network. Bandwidth is allocated on a push basis in response to receiving a cue tone from a broadcast stream. The cue tone contains data indicating an addressable break. A portion of bandwidth is then allocated to an addressable content stream based on the data of the cue tone. The addressable content stream is then streamed to a receiver during the addressable break and the receiver is tuned from the broadcast stream to the addressable content stream for the duration of the addressable break. The additional bandwidth is de-allocated at the end of the addressable break.



WO 2010/014831 A1

55113-3

SYSTEMS AND METHODS FOR DYNAMIC
BANDWIDTH ALLOCATION FOR ADDRESSABLE CONTENT

FIELD OF INVENTION

[0002] The present invention generally relates to cable transmission, and more specifically to methods and systems for dynamic bandwidth allocation for addressable content.

BACKGROUND

[0003] Cable operators and vendors, such as Big Band, Motorola, Cisco, and Imagine Communications, have developed and are further refining systems to measure bandwidth demand or “pull” by cable consumers and allocate capacity within a cable plant node on a real-time dynamic basis. These types of systems are referred to as “switched digital video” systems. Switched digital video (“SDV”) systems are being implemented to resolve growing consumer demand of cable bit/bandwidth. With hundreds of possible television channels and an increasing number of high-definition (“HD”) channels, cable service providers are being stretched to the limits of their network capacity in order to provide uninterrupted, quality service to their subscribers. In addition to audio and video data transmitted for television services, many providers also package Internet, Video-on-Demand and digital telephone services to subscribers – all within the same cable infrastructure.

[0004] To accommodate the increased demand for bandwidth, cable providers limit the transmission of a particular channel until it is requested by a subscriber. For example, a certain channel is not constantly broadcast to a home, or neighborhood. When a subscriber tunes to that channel on the digital set-top box, a signal is sent to the cable provider to “turn-on” the channel. The provider then transmits the stream of data containing that channel’s

video and audio through the cable to the set-top box and on to the subscriber's television.

Should a second subscriber in the same service area call up that same channel, the stream is forwarded on to that set top box eliminating the need for a second stream of the same channel.

[0005] The transmission speeds of signals over the cable lines is fast enough that the subscriber is unaware that seconds before tuning to that station, the station was not being broadcast at all.

[0006] Current switched digital video systems are designed to alleviate and overcome bandwidth transmission limitations from the consumer on a "pull" basis (i.e., dynamically allocating bandwidth based upon subscriber usage and demand). For example, when a digital Cable TV subscriber starts a Video-On-Demand ("VOD") session to watch a movie, a traditional sequence of events includes following steps: the subscriber selects a particular piece of On-Demand content from the menu on the Digital Cable Set-top Box ("DSTB"); the DSTB initiates a request to the back-end VOD server to start a session; the VOD server performs some authentication and/or billing functions; and then the VOD server allocates the content. The VOD server then will attempt to allocate bandwidth for this session. If no bandwidth can be allocated at the time of the request, the VOD server will notify the DSTB which will then inform the user of the failed VOD session with an "Error: Please try again later" message.

[0007] Once bandwidth is allocated, the VOD server will begin streaming the content (via the newly allocated bandwidth) to the DSTB, where it is rendered to the subscriber. After the session is over (either because the subscriber actively stopped the session, or because the session timed out), the bandwidth is de-allocated and returned to the network.

[0008] Another example of the traditional “pull-based” model of bandwidth allocation occurs when a digital Cable TV subscriber changes channels in a cable system that uses SDV technology to save bandwidth. The subscriber requests the DSTB to tune from channel X to channel Y, either through a Program Guide, or by pressing a Channel-Up/Channel-Down button on the remote control, or by entering the channel number directly on the remote control. The DSTB initiates a request to the back-end SDV server conveying that the DSTB will leave channel X and tune to channel Y. The SDV server will first check if this was the only viewer in the service area that was still watching channel X; if that is the case, the SDV server will remove channel X from the active channel line-up for this particular cable system service areas, and de-allocate the bandwidth that the channel data stream station was occupying.

[0009] The SDV server will then check to see if channel Y is already available in the active channel line-up for this particular cable system service area (signifying that at least one other subscriber’s DSTB in the service area is tuned to this channel already); if it is not yet available, the SDV server will allocate bandwidth for it and add it to the active channel line-up. The SDV server will then send a message back to the DSTB to indicate the (new) position of channel Y in the active channel line-up, at a position ‘*n*’. Upon receipt of this message, the DSTB will tune to position *n*, and the viewer will start viewing channel Y.

[0010] There remains an untapped resource for cable and satellite service providers, as well as advertisers alike in adapting a dynamic bandwidth allocation protocol on a “push” basis to opportunistically exploit the gaps or holes in the available bandwidth efficiently and effectively. That is, dynamically allocating available bandwidth to certain channels for the inclusion of additional services, such as advanced advertising and content delivery, while minimizing bandwidth allocated to a channel but not which is being used.

SUMMARY OF THE INVENTION

[0010a] According to an aspect of the present disclosure, there is provided a method of allocating bandwidth comprising: receiving a cue tone from a broadcast stream, the cue tone comprising data indicating at least a start-time of an addressable break in the broadcast stream; allocating a portion of available bandwidth to an addressable content stream based on the data of the cue tone; streaming the addressable content stream to a receiver during the addressable break; switching the receiver from the broadcast stream to the addressable content stream for the duration of the addressable break; and de-allocating the portion of bandwidth at the end of the addressable break, wherein the portion of bandwidth is allocated at a first time prior to the start-time of the addressable break and is de-allocated at a second time after an end-time of the addressable break.

[0010b] According to another aspect of the present disclosure, there is provided a system for dynamically allocating bandwidth comprising: a server comprising addressable content; a splicer receiving a broadcast stream and inserting a stream of the addressable content from the server into the broadcast stream; and an edge device detecting a cue tone, the cue tone comprising data indicating at least a start-time of an addressable break, the edge device allocating a portion of available bandwidth to the stream of the addressable content based on the data of the cue tone for a duration, the edge device de-allocating the portion of bandwidth at the end of the duration of the addressable break, wherein the portion of bandwidth is allocated at a first time prior to the start-time of the addressable break and is de-allocated at a second time after an end-time of the addressable break.

[0010c] According to another aspect of the present disclosure, there is provided a system for dynamically allocating bandwidth comprising: a splicer receiving and transmitting a broadcast stream, the splicer detecting a cue tone, the cue tone comprising data indicating at least a start-time of an addressable break; and a server comprising addressable content, the server transmitting an addressable content stream, wherein the server allocates a portion of bandwidth for a duration to the addressable content stream based on the data of the cue tone and the server de-allocates the portion of bandwidth at an end of the duration of the addressable break, wherein the portion of bandwidth is allocated at a first time prior to the

start-time of the addressable break and is de-allocated at a second time after an end-time of the addressable break.

[0010d] According to another aspect of the present disclosure, there is provided a system for dynamically allocating bandwidth comprising: a splicer receiving and transmitting
5 a broadcast stream, the splicer detecting a cue tone, the cue tone comprising data indicating at least a start-time of an addressable break; and a server comprising addressable content, the server transmitting an addressable content stream, and a resource manager receiving a request for bandwidth from the server, the resource manager defining subnetworks of available subscribers, the resource manager determining a portion of the available subscribers in the
10 subnetwork selected to receive the addressable content, the resource manager allocating a portion of bandwidth for a duration to the addressable content stream based on the data of the cue tone and the portion of available subscribers in the subnetwork selected to receive the addressable content, the resource manager de-allocating the portion of bandwidth at the end of the duration of the addressable break, wherein the portion of bandwidth is allocated at a first
15 time prior to the start-time of the addressable break and is de-allocated at a second time after an end-time of the addressable break.

[0010e] According to another aspect of the present disclosure, there is provided a system for dynamically allocating bandwidth comprising: means for receiving a cue tone from a broadcast stream, the cue tone comprising data indicating at least a start-time of an
20 addressable break in the broadcast stream; means for allocating a portion of available bandwidth to an addressable content stream based on the data of the cue tone; means for streaming the addressable content stream to a receiver during the addressable break; and means for de-allocating the portion of bandwidth at the end of the addressable break, wherein the portion of bandwidth is allocated at a first time prior to the start-time of the addressable
25 break and is de-allocated at a second time after an end-time of the addressable break.

[0010f] According to another aspect of the present disclosure, there is provided a method of allocating bandwidth comprising: receiving a trigger from a content distribution stream, the trigger comprising data indicating at least a start-time of an addressable break in the content distribution stream; allocating a portion of available bandwidth to an addressable

content stream based on the data of the trigger; streaming the addressable content stream to a receiver during the addressable break; switching the receiver to the addressable content stream for the duration of the addressable break; and de-allocating the portion of bandwidth at the end of the addressable break, wherein the portion of bandwidth is allocated at a first time prior to the start-time of the addressable break and is de-allocated at a second time after an end-time of the addressable break.

[0010g] According to another aspect of the present disclosure, there is provided an apparatus comprising: one or more processors; and memory storing instructions that, when executed by the one or more processors, cause the apparatus to: receive a trigger from a content distribution stream, the trigger comprising data indicating at least a start-time of an addressable break in the content distribution stream; allocate a portion of available bandwidth to an addressable content stream based on the data of the trigger; stream the addressable content stream to a receiver during the addressable break; switch the receiver to the addressable content stream for the duration of the addressable break; and de-allocate the portion of bandwidth at the end of the addressable break, wherein the portion of bandwidth is allocated at a first time prior to the start-time of the addressable break and de-allocated at a second time after an end-time of the addressable break.

[0010h] According to another aspect of the present disclosure, there is provided a system for dynamically allocating bandwidth comprising: a splicer for receiving and transmitting a content distribution stream, the splicer being configured to detect a trigger, the trigger comprising data indicating a start-time of an addressable break; and a server comprising addressable content, the server being configured to transmit an addressable content stream, wherein the server allocates a portion of bandwidth for a duration to the addressable content stream based on the data of the trigger and the server de-allocates the portion of bandwidth at the end of the duration of the addressable break, wherein the server allocates the bandwidth at a first time prior to the start-time of the addressable break, and de-allocates the portion of bandwidth at a second time after an end-time of the addressable break.

[0010i] According to another aspect of the present disclosure, there is provided a method of allocating bandwidth comprising: receiving, at a splicer, a content distribution

stream; inserting, by the splicer, a stream of addressable content into the content distribution stream, the addressable content being stored on a server; detecting, at an edge device, a trigger from the content distribution stream, the trigger comprising data indicating at least a start-time of an addressable break in the content distribution stream; allocating, by the edge device, a portion of available bandwidth to the stream based on the data of the trigger for a duration; streaming the addressable content stream to a receiver during the addressable break; switching the receiver to the addressable content stream for the duration of the addressable break; and de-allocating the portion of bandwidth at the end of the addressable break, wherein the edge device allocates the bandwidth at a first time prior to the start-time of the addressable break, and de-allocates the portion of bandwidth at a second time after an end-time of the addressable break.

[0010j] According to another aspect of the present disclosure, there is provided a system for dynamically allocating bandwidth comprising: a server comprising addressable content; a splicer for receiving a content distribution stream and inserting a stream of the addressable content from the server into the content distribution stream; and an edge device for detecting a trigger, the trigger comprising data indicating at least a start-time of an addressable break, the edge device allocating a portion of available bandwidth to the stream of the addressable content based on the data of the trigger for a duration, the edge device de-allocating the portion of bandwidth at the end of the duration of the addressable break, wherein the edge device allocates the bandwidth at a first time prior to the start-time of the addressable break, and de-allocates the portion of bandwidth at a second time after the end-time of the addressable break.

[0010k] According to another aspect of the present disclosure, there is provided a system for dynamically allocating bandwidth comprising: means for receiving a trigger from a content distribution stream, the trigger comprising data indicating at least a start-time of an addressable break in the content distribution stream; means for allocating a portion of available bandwidth to an addressable content stream at a first time prior to the start-time of the addressable break; means for streaming the addressable content stream to a receiver during the addressable break; and means for de-allocating the portion of bandwidth at a second time after the end-time of the addressable break.

[0010l] According to another aspect of the present disclosure, there is provided a method comprising: receiving, at a splicer, a content distribution stream; inserting, by the splicer, a stream of addressable content into the content distribution stream, the addressable content being stored on a server; detecting, at an edge device, a trigger from the content distribution stream, the trigger comprising data indicating at least a start time of an addressable break in the content distribution stream; allocating, by the edge device, a portion of available bandwidth to an addressable content stream prior to the start time of the addressable break and for a duration; streaming, by the edge device, the addressable content stream to a receiver during the addressable break; switching, by the edge device, the receiver from the content distribution stream to the addressable content stream for the duration of the addressable break; and de-allocating, by the edge device, the portion of bandwidth at an end time of the addressable break.

[0010m] According to another aspect of the present disclosure, there is provided a system comprising: a server comprising addressable content; a splicer for receiving a content distribution stream and inserting a stream of the addressable content from the server into the content distribution stream; and an edge device for detecting a trigger, the trigger comprising data indicating at least a start time of an addressable break, the edge device allocating a portion of available bandwidth to the stream of the addressable content prior to the start time of the addressable break and for a duration, the edge device de-allocating the portion of bandwidth at an end time of the addressable break.

[0010n] According to another aspect of the present disclosure, there is provided a system comprising: means for receiving a trigger from a content distribution stream, the trigger comprising data indicating at least a start time of an addressable break in the content distribution stream; means for allocating a portion of available bandwidth to an addressable content stream prior to the start time of the addressable break and for a duration; means for streaming the addressable content stream to a receiver during the addressable break; and means for de-allocating the portion of bandwidth at an end time of the addressable break.

[0010o] According to another aspect of the present disclosure, there is provided a method comprising: receiving a content stream, at least a portion of the content stream

comprising an addressable content stream; detecting, in the content stream, a trigger comprising data indicating at least a start time of an addressable break in the content stream; allocating, based on the data associated with the trigger and prior to the start time of the addressable break, a portion of available bandwidth to the addressable content stream; streaming, to a receiver and for a duration of the addressable break, the addressable content stream; and de-allocating, after an end time of the addressable break, the portion of available bandwidth.

[0010p] According to another aspect of the present disclosure, there is provided an apparatus comprising a processor and a memory, the memory storing computer-executable instructions which, when executed by the processor, cause the apparatus to perform operations comprising: receiving a content stream, at least a portion of the content stream comprising an addressable content stream; detecting, in the content stream, a trigger comprising data indicating at least a start time of an addressable break in the content stream; allocating, based on the data associated with the trigger and prior to the start time of the addressable break, a portion of available bandwidth to the addressable content stream; streaming, to a receiver and for a duration of the addressable break, the addressable content stream; and de-allocating, after an end time of the addressable break, the portion of available bandwidth.

[0010q] According to another aspect of the present disclosure, there is provided a computer-readable storage medium comprising computer-executable instructions which, when executed by a device, cause the device to perform operations comprising: receiving a content stream, at least a portion of the content stream comprising an addressable content stream; detecting, in the content stream, a trigger comprising data indicating at least a start time of an addressable break in the content stream; allocating, based on the data associated with the trigger and prior to the start time of the addressable break, a portion of available bandwidth to the addressable content stream; streaming, to a receiver and for a duration of the addressable break, the addressable content stream; and de-allocating, after an end time of the addressable break, the portion of available bandwidth.

[0010r] According to another aspect of the present disclosure, there is provided a method comprising: receiving data indicating at least a start time of an addressable break in a

content stream; allocating, based on the data, prior to the start time of the addressable break, a portion of available bandwidth to an addressable content stream; transmitting, to a receiver and for an entire duration of the addressable break starting at the start time, the addressable content stream; and de-allocating, after the entire duration of the addressable break, the portion of available bandwidth.

[0010s] According to another aspect of the present disclosure, there is provided an apparatus comprising: a processor; and a memory storing instructions that, when executed by the processor, cause the apparatus to: receive data indicating at least a start time of an addressable break in a content stream; allocate, based on the data, prior to the start time of the addressable break, a portion of available bandwidth to an addressable content stream; transmit, to a receiver and for an entire duration of the addressable break starting at the start time, the addressable content stream; and de-allocate, after the entire duration of the addressable break, the portion of available bandwidth.

[0010t] According to another aspect of the present disclosure, there is provided a non-transitory computer-readable storage medium storing instructions that, when executed by a processor of an apparatus, cause the apparatus to: receive data indicating at least a start time of an addressable break in a content stream; allocate, based on the data, prior to the start time of the addressable break, a portion of available bandwidth to an addressable content stream; transmit, to a receiver and for an entire duration of the addressable break starting at the start time, the addressable content stream; and de-allocate, after the entire duration of the addressable break, the portion of available bandwidth.

[0011] Embodiments of the invention include systems and methods for dynamic bandwidth allocation to deliver addressable, advertising content in a digital network to users, using allocation techniques driven by asynchronous events related to the advertising content, instead of being requested by an end-user (in other words, it is based on a “push model” instead of on a “pull model”). Embodiments of the system are described herein in the context of addressable content in a multi-cast multimedia network (e.g. a digital cable TV system, or a Direct To Home satellite TV system), however one skilled on the art should recognize the

81633727

invention is equally applicable to other content-initiated bandwidth allocation systems as well (for example it could be used to allocate bandwidth to send personalized news or sports sequences to TV viewers).

[0012] According to one embodiment, a method of allocating bandwidth on a push basis includes receiving a cue tone from a broadcast stream. The cue tone contains data indicating an addressable break. A portion of additional bandwidth is then allocated to an addressable content stream based on the data of the cue tone. The addressable content stream is then streamed to a receiver during the addressable break and the receiver is tuned from the broadcast stream to the addressable content stream for the duration of the addressable break. The additional bandwidth is de-allocated at the end of the addressable break.

[0013] Another embodiment of the invention includes a system for dynamically allocating bandwidth. The system includes a storage server containing addressable content and a splicer that is capable of receiving a broadcast stream inserting a stream of the addressable content from the storage server into the broadcast stream. An edge device is included which is capable of detecting a cue tone in the broadcast stream. The cue tone contains data indicating an addressable break. The edge device is further capable of allocating a portion of bandwidth

to the stream of the addressable content based on the data of the cue tone for a duration. The edge device is also capable of de-allocating the portion of bandwidth at the end of the duration of the addressable break.

[0014] Another embodiment of the invention includes a system for dynamically allocating bandwidth having a splicer capable of receiving and transmitting a broadcast stream. The splicer detects a cue tone in the broadcast stream. The cue tone contains data indicating an addressable break. The system also includes a storage server containing addressable content. The storage server is capable of transmitting an addressable content stream and allocates a portion of bandwidth for a duration to the addressable content stream based on the data of the cue tone. The storage server de-allocates the portion of bandwidth at the end of the duration of the addressable break.

[0015] Another embodiment of the invention includes a system for dynamically allocating bandwidth including a splicer for receiving and transmitting a broadcast stream. The splicer detects a cue tone in the broadcast stream which contains data indicating an addressable break. A storage server contains addressable content and transmits an addressable content stream. A resource manager receives a request for additional bandwidth from the storage server. The resource manager defines a subnetwork of available subscribers and determines a portion of the available subscribers in the subnetwork selected to receive the addressable content. The resource manager allocates a portion of the bandwidth for a duration to the addressable content stream based on the data of the cue tone and the portion of available subscribers in the subnetwork selected to receive the addressable content. The resource manager then de-allocates the portion of bandwidth at the end of the duration of the addressable break.

[0016] In yet another embodiment of the invention, a system for dynamically allocating bandwidth includes components for receiving a cue tone from a broadcast stream. The cue tone includes data indicating at least the beginning of an addressable break in the broadcast stream. The system also includes components for allocating a portion of available bandwidth to an addressable content stream based on the data of the cue tone and components for streaming the addressable content stream to a receiver during the addressable break. Components for tuning the receiver from the broadcast stream to the addressable content stream for the duration of the addressable break are also included. Lastly, the system includes components for de-allocating the portion of bandwidth at the end of the addressable break.

BRIEF DESCRIPTION OF THE DRAWINGS

[0017] These embodiments and other aspects of this invention will be readily apparent from the detailed description below and the appended drawings, which are meant to illustrate and not to limit the invention, and in which:

FIG. 1 is a diagram of frequency allocations of a broadcast type cable delivery streaming system including addressable advertisements;

FIG. 2 is a timeline depiction of bandwidth allocation in a network stream as a function of time;

FIG. 3 is a timeline depiction of bandwidth allocation in a network stream as a function of time in accordance with an embodiment of the invention;

FIG. 4 is a timeline depiction of bandwidth allocation in a network stream as a function of time in accordance with an embodiment of the invention;

FIG. 5 depicts a method of dynamic bandwidth allocation in accordance with an embodiment of the invention;

FIG. 6 is a block diagram of a dynamic bandwidth allocation system in accordance with an embodiment of the invention;

FIG. 7 is a block diagram of a dynamic bandwidth allocation system in accordance with an embodiment of the invention;

FIG. 8 is a block diagram of a dynamic bandwidth allocation system in accordance with an embodiment of the invention;

FIG. 9 is a block diagram of a dynamic bandwidth allocation system including a resource manager in accordance with an embodiment of the invention; and

FIG. 10 is a block diagram of a dynamic bandwidth allocation system including a resource manager in accordance with an embodiment of the invention.

DETAILED DESCRIPTION

[0018] The invention will be more completely understood through the following detailed description, which should be read in conjunction with the attached drawings. Detailed embodiments of the invention are disclosed herein, however, it is to be understood that the disclosed embodiments are merely exemplary of the invention, which may be embodied in various forms. Therefore, specific functional details disclosed herein are not to be interpreted as limiting, but merely as a basis for the claims and as a representative basis for teaching one skilled in the field to variously employ the invention in virtually any appropriately detailed embodiment.

[0019] Embodiments of the invention include systems and methods for dynamic bandwidth allocation to deliver addressable, advertising content in a digital network to users, using allocation techniques driven by asynchronous events related to the content, instead of being requested by an end-user (in other words, it is based on a “push model” instead of on a “pull model”).

[0020] Turning now to Figure 1, frequency allocations 100 for a system for addressable advertising in a digital cable system is depicted in accordance with an embodiment of the invention. As depicted in the illustrative figure, the system includes, without limitation, a number of linear local ad-supported cable networks 5 (shown as a function of Time): twenty broadcast networks transmitted in two distinct transport streams broadcast at 6 MHz intervals. Each network may have its own local addressable breaks 7 at pre-defined times in regularly scheduled programming for advertisements or other addressable content. During these breaks, addressable commercials are streamed through additional bandwidth 9 and the DSTBs that are tuned to the regularly scheduled program are forcibly tuned 11 away from the program to a particular advertisement 15. The DSTBs will tune back to the network program 13 at the end of the commercial. The illustrative system can tailor advertising content within the linear broadcast network content by matching demographic profiles of the subscribers with a particular selection of addressable commercial content suited for the subscriber. For example, if a subscriber's credit score is above a certain threshold and maintains a substantially high household income then addressable content for a premium credit card may, such as a gold or platinum card be may be selected to present to the subscriber during a break.

[0021] Traditional approaches to bandwidth allocation for such an addressable advertising systems have been based on either static bandwidth allocation (i.e., allocate the bandwidth at

all times), or schedule-based bandwidth allocation (i.e., allocate the bandwidth when an addressable break is scheduled). Local advertising breaks follow a schedule that is unique per network. For example on ESPN's broadcast channel, advertising breaks occurring between 5:30 A.M. and 6:30 A.M. on a particular day may be scheduled as follows:

- Tuesday Jan 1, 2009, 5:45am, a :30 second break
- Tuesday Jan 1, 2009, 6:12am, a :60 second break
- Tuesday Jan 1, 2009, 6:43am, a :60 second break
- Tuesday Jan 1, 2009, 7:14am, a :90 second break

Other networks may each have their own unique schedule tailored to their respective broadcasted programs.

[0022] In practice, however, the scheduled time for an addressable break may be only indicative. That is, the actual time when the break occurs may differ slightly (or sometimes significantly) from the scheduled time. When broadcasting live events or alternative programming due to scheduling conflicts or blackout restrictions, for example, broadcast programmers cannot predict or ensure when programming breaks will occur. Holding the break time as an indicator offers flexibility to the network programmer to shift commercial breaks based on unpredictable programming schedules. To remedy the inaccuracy in predicted scheduling times, broadcast programmers and service providers use a window concept. A window provides a time interval for which a scheduled break is valid. A break may be scheduled for 5:10 A.M., with a window open-time of 5:00 A.M. and a window close time of 5:20 A.M., which means that the break is scheduled for 5:10 A.M., but could occur as early as 5:00 A.M. and as late as 5:20 A.M. Any break that occurs during that window is considered to be the 5:10 A.M. break, even if it did not actually occur at that precise time.

[0023] Bandwidth allocation for addressable breaks can not be based on scheduled time, as described above, because it is inaccurate. Thus, bandwidth (and other resources) for addressable breaks is allocated based on the break window as described above. In traditional allocation systems, bandwidth is allocated for the full duration of the break window, because that is the defined time interval when the scheduled break might occur. Figure 2 depicts the frequency allocations 200 for a traditional window-based bandwidth allocation system. As the broadcast stream progresses through time, a window 21 opens at time, T_0 , which is the time the additional content bandwidth 9 is allocated for that stream. If the actual break interval 17 occurs at a different time than the scheduled break interval 19, the bandwidth 9 is still available because the bandwidth 9 is allocated for the entire window period 21, ending at time T_C .

[0024] This schedule-based system or, more accurately, window-based system, of bandwidth allocation is functional, but is also inefficient. Bandwidth is a scarce resource in many digital broadcast systems, and the window-based allocation approach locks up bandwidth for the full duration of a window, while it is only effectively being used for the duration of the break, a fraction of the entire window. For example, if a window size of twenty minutes is selected for a sixty-second break, then the bandwidth is allocated for twenty minutes; however the bandwidth is only actually used for the sixty seconds, so the time-bandwidth allocation, in this example is twenty times higher than the actually needed time-bandwidth to broadcast the addressable content.

[0025] One embodiment of the invention, as depicted in Figure 3, includes a frequency bandwidth illustration 300 for a system using real-time break programming data to allocate bandwidth 9 only when an actual break 17 occurs. When a broadcast system knows when an addressable break 17 is going to occur, a form of pre-roll notification called a "cue tone" is

transmitted. Traditionally these cue tones were audio tones, however with the proliferation of digital signal processing and broadcast transmission, digital cue tones, as described in the American National Standards Institute (“ANSI”) and Society of Cable and Telecommunications Engineers (“SCTE”) Standard 35, are being used increasingly. Other (more proprietary) cue tones are in existence as well such as Internet-Protocol (“IP”) -based cue tones. Cue tones may sound or trigger up to several seconds before the actual break occurs, and contain precise information identifying the actual start-time and stop-time of the break.

[0026] One embodiment of the invention uses these tones to create an efficient, real-time dynamic bandwidth allocation system for addressable advertising as illustrated in Figure 4. A broadcast stream 23 including the addressable content stream (shown here overlapping the broadcast stream) is shown as a function of time T , defining the events associated with the bandwidth allocation of the embodiment. At time T_0 no additional bandwidth is allocated to the addressable content, and regularly scheduled programming of the network broadcast is being transmitted the subscriber on the tuned channel. At time T_1 , an addressable break window is opened, as traditional broadcast systems would use, however in this embodiment, no bandwidth is yet allocated to the stream 23. A cue tone 25 occurs at time T_2 , specifying time, T_3 , for the start of actual addressable break, but still no bandwidth is allocated. When time T_3 arrives as the actual start of addressable break 17, additional bandwidth 9 is now allocated to the addressable content stream and spliced into the broadcast stream nearly instantaneously. Time T_4 occurs during the actual break 17 and additional bandwidth 9 remains allocated for the broadcast of the addressable break content. The additional bandwidth 9 is de-allocated immediately at Time T_5 , the end of the break, as dictated by the cue tone 25 at Time T_2 , tuning the receiver back to the broadcast stream. At time T_6 , the addressable break 17 has already ended, but the break window remains open, however no

bandwidth is allocated. The break window is closed at a later time T_7 . It should be noted that, according to one embodiment, the window open and window close events shown in Figure 4 are irrelevant to allocation of additional bandwidth.

[0027] The embodiment of the dynamic bandwidth allocation system described above is more efficient than the schedule-based allocation method. For the example of a twenty minute window with a sixty-second break, previously mentioned, the schedule based system would allocate twenty times more bandwidth for the same break than the dynamic system of the present embodiment (i.e., twenty minutes of costly allocated bandwidth versus sixty seconds), with the same effective bandwidth used. Alternatively, even in an embodiment of the invention in which bandwidth would be allocated immediately upon receipt of the cue tone (at Time T_2) instead of at the exact break-start (T_3), to provide a buffer, or safety period to ensure available bandwidth, significant efficiencies are still maintained using the dynamic allocation system of the embodiment.

[0028] One embodiment of the invention includes a method 500 for dynamically allocating bandwidth for addressable content, as depicted in Figure 5. As described above, each station may have distinct schedule breaks for addressable content to be broadcast to subscribers. The break schedule 27 may be created and managed in a traffic and billing system, and driven by an advertisement insertion module, both described in further detail below. Using the time scale outlined above in connection with Figure 4, an addressable break window may open 29 at Time T_0 , indicating the possibility of an upcoming break. At Time T_1 a cue tone is received and processed 31 for the start-time, duration, end-time or other data 33. At Time T_3 , the additional bandwidth is allocated 35 at or near the same time as the actual break begins as contained in the cue tone. At and during Time T_4 the additional bandwidth is allocated and used to stream the addressable content 37 to subscriber's receiver. When the addressable

break ends, at Time T_5 , the additional bandwidth is de-allocated 39 and returned to a pool of available bandwidth to be used for other applications or requests. The addressable break window may remain open for a period, T_6 after the end of the actual break until the window is scheduled to close at Time T_7 41.

[0029] Turning now to Figure 6, an implementation of a dynamic bandwidth allocation system 600 is depicted in accordance with an embodiment of the invention. According to the embodiment, a traffic and billing (“T&B”) system module 43 manages the schedules for the local advertising breaks. An advertising insertion system module 45 manages insertion of local commercials in the local breaks. The local commercials inserted in this illustrated embodiment are non-addressable or non-targeted advertisements broadcast to all transmission receivers. A digital ad server 47 manages local commercials for one or more sets of particular ad insertion zones. A digital ad splicer 49 actually splices local digital commercials into the digital network feeds 53. The splicer 49 interacts with the digital ad server 47 by inserting the commercials obtained from the ad server 47. An edge device module 51 manages the addressable commercials by managing and allocating the required bandwidth needed to stream the addressable commercials. When the edge device module 51 detects an addressable break, through a digital cue tone or other signal as described above, the edge device module 51 dynamically allocates the necessary bandwidth for the duration of the break. The edge device module 51 then uses the allocated bandwidth to stream out the commercials 55 in addition to the network feeds 57 and then de-allocates the bandwidth again at the end-time obtained from the cue tone. If multiple, simultaneous, addressable breaks should occur, the edge device module 51 attempts to allocate bandwidth for all of the breaks. Once the edge device module 51 runs out of available bandwidth, it will stop servicing addressable advertising breaks until enough bandwidth is returned. For example, if two or more broadcast channels happen to send cue tones to open up addressable content breaks at

the same time, or if the breaks will overlap, the edge device module 51 will attempt to allocate the necessary bandwidth to both, or all channels, having simultaneous or overlapping addressable breaks, until the spectrum of available bandwidth is diminished beyond the capacity to carry further signals.

[0030] In an alternative embodiment, the functionality of the dynamic bandwidth allocation system derived from the edge device, may be built into digital ad servers and/or splicers. An exemplary embodiment 700 is depicted in Figure 7. A T&B system 43 manages the schedules for the local advertising breaks. The advertising insertion system module 45 manages insertion of local commercials in the local breaks together with a digital ad splicer 49. These are regular local commercials, not the addressable commercials. The ad server 47 manages bandwidth for addressable commercials when the splicer 49 detects and notifies the ad server 47 that a digital cue tone has occurred. The ad server 47 will then allocate the necessary bandwidth for the exact duration of the break, as detailed in the cue tone, and use the allocated bandwidth to stream the commercials 55 out to the DSTBs. When the end of the break occurs, the ad server 47 will de-allocate the bandwidth freeing it up for subsequent allocation, and will instruct the splicer 49 to return to the network feed 57. The illustrative embodiment may also handle multiple, simultaneous, addressable advertising breaks as described above in connection with the system of Figure 6.

[0031] According to another embodiment of the invention, as depicted in Figure 8, a dynamic bandwidth allocation system 800 includes a digital ad server 47 to allocate bandwidth by transmitting a call to a resource manager module 59. A T&B system 43 manages the schedules for the local advertising breaks. The advertising insertion system module 45 manages insertion of local commercials in the local breaks together with a digital ad splicer 49. These are regular local commercials, not the addressable commercials. The ad

server 47 further manages bandwidth for addressable commercials when the splicer 49 detects and notifies the ad server 47 that a digital cue tone has occurred. Specifically, the ad server 47 may request bandwidth from the resource manager 59 whenever a cue tone occurs. The resource manager module 59 manages a pool of available bandwidth that may or may not be shared with other bandwidth-intensive applications such as VOD, SDV or other applications. The resource manager module 59 receives requests for bandwidth from the digital ad server and can serve these requests based on a prioritization policy, such as a first-come-first served policy. Other more sophisticated, non-linear prioritization policies may also be implemented without deviating from the scope of the invention. The digital ad server will then obtain the necessary bandwidth for the exact duration of the break, as detailed in the cue tone, and use the allocated bandwidth to stream the commercials 55 out to the DSTBs. When the end of the break occurs, the ad server will return the bandwidth back to the resource manager 59.

[0032] According to another embodiment of the invention, as depicted in Figure 9, a dynamic bandwidth allocation system 900 includes a T&B system 43 that manages the schedules for the local advertising breaks. Whenever an ad splicer 49 receives a digital cue tone, it will notify the ad server 47. The ad server 47 then requests bandwidth for addressable commercials from resource manager 59. The resource manager 59 manages a pool of available network bandwidth. In fact the network is divided up in several sub-networks 61, where each subnetwork 61 has its own bandwidth resource 63. Each subnetwork 61 also represents a number of connected (subscriber) homes 65 within the cable system, direct-to-home satellite system, or IPTV system. An example of such a subnetwork 61 in a cable system is a service group, where depending on the system each service group may comprise several hundred to few thousand subscriber homes 65. When the resource manager 59 receives a request for addressable commercials (A, B, C, D, E) from the ad server 47, the

resource manager 59 may allocate that bandwidth in every one of its subnetworks 61. The resource manager 59 knows which individual homes 65 are in each subnetwork 61. The resource manager 59 also knows which home is profiled or targeted with each commercial (A, B, C, D, E). Rather than allocating additional bandwidth for all five commercials (A, B, C, D, E) in all subnetworks 61, the resource manager 59 will only allocate the additional bandwidth for a commercial in a subnetwork 61 if there is at least one home in that subnetwork 61 that will be targeted with that particular commercial. The result is a more efficient bandwidth allocation as illustrated below:

Subnet#	# Homes	Home Profiles	Bandwidth Allocated
1	7	A,B,A,B,B,E,A	3 x
2	13	D,B,A,C,A,A,A,A,D,C,C,C,B	4 x
3	10	E,A,E,E,A,E,B,E,E,E	3 x
4	18	C,C,B,A,C,E,D,A,A,B,C,B,E,E,A,B,C,C	5 x
5	4	A,C,C,B	3 x

For example, Subnet #1 may include seven homes in its defined network. According to the resource manager's profile, those seven homes are targeted to receive commercials A, B and E (three home receiveing commercial A, three home receiving commercial B and one home receiving commercial E). Accordingly, the resource manager may only allocate the additional bandwidth required to stream those three commericals to the homes of assigned to Subnet #1, instead of allocating the additional bandwidth for the full five addressable advertisements.

[0033] According to yet another embodiment of the invention, as depicted in Figure 10, a dynamic bandwidth allocation system 1000 includes a T&B system 43 that manages the schedules for the local advertising breaks. Whenever an ad splicer 49 receives a digital cue tone, it will notify the ad server 47. The ad server 47 then requests bandwidth for addressable commercials from resource manager 59. The resource manager 59 manages a pool of available network bandwidth. In fact the network is divided up in several sub-networks 61,

where each subnetwork 61 has its own bandwidth resource 63. Each subnetwork 61 also represents a number of connected (subscriber) homes 65 within the system. When the resource manager 59 receives a request for addressable commercials (A, B, C, D, E) from the ad server 47, the resource manager 59 may allocate that bandwidth in every one of its subnetworks 61. The resource manager 59 knows which individual homes 65 are in each subnetwork 61. The resource manager 59 also knows which home is profiled or targeted with each commercial (A, B, C, D, E). In addition, the resource manager 59 also knows which of the homes is currently active. A home is active if it is ready to receive one of the addressable commercials (A, B, C, D, E). An example of an active home 67, is a home that is tuned to a channel that is ready to receive an addressable break. As the home is tuned to an appropriate channel, the home is ready to receive one of the addressable commercials at an upcoming addressable break. Rather than allocating bandwidth for all five commercials in all subnetworks 61, the resource manager 59 will only allocate bandwidth for a commercial in a subnetwork, if there is at least one active home in that subnetwork that will be targeted with that commercial. The result is an even more efficient bandwidth allocation as illustrated below:

Subnet#	# Homes	Active Home Profiles	Bandwidth Allocated
1	7	A, A, B	2 x
2	13	D, A, C, B	4 x
3	10		0 x
4	18	C, C, B, B, C, C	2 x
5	4	A	1 x

For example, Subnet #1 may include seven homes in its defined network. Of those seven homes, only three homes are currently active (i.e. currently tuned to a channel ready to receive addressable content). Of those three active homes, two are targeted to receive commercial A and one home is targeted to receive commercial B. Because only three homes are currently active, and of those three home, only two addressable commercials are targeted

55113-3

to those homes, bandwidth need only be allocated for two distinct addressable content streams.

[0034] While the certain embodiments described herein include a resource manager in system configurations like those shown in Figures 8-10, one skilled in the art should recognize that the resource manager may be implemented in alternative system configurations, such as the system of Figure 6, without deviating from the spirit and scope of the invention.

[0035] While the embodiments described herein are depicted as modular systems with defined functionalities, one skilled in the art should recognize that the present invention is not limited to the exemplary embodiments and other system architectures using different combinations of functionality in digital ad servers, digital ad splicers, edge devices, and resource managers may be implemented without deviating from the scope of the invention.

[0036] While the invention has been described with reference to illustrative embodiments, it will be understood by those skilled in the art that various other changes, omissions and/or additions may be made and substantial equivalents may be substituted for elements thereof without departing from the scope of the invention. In addition, many modifications may be made to adapt a particular situation or material to the teachings of the invention without departing from the scope thereof. Therefore, it is intended that the invention not be limited to the particular embodiment disclosed for carrying out this invention, but that the invention will include all embodiments falling within the scope of the appended claims. Moreover, unless specifically stated, any use of the terms first, second, etc. do not denote any order or importance, but rather the terms first, second, etc. are used to distinguish one element from another.

CLAIMS:

1. A method of allocating bandwidth comprising:
 - receiving a cue tone from a broadcast stream, the cue tone comprising data indicating at least a start-time of an addressable break in the broadcast stream;
 - allocating a portion of available bandwidth to an addressable content stream based on the data of the cue tone;
 - streaming the addressable content stream to a receiver during the addressable break;
 - switching the receiver from the broadcast stream to the addressable content stream for the duration of the addressable break; and
 - de-allocating the portion of bandwidth at the end of the addressable break, wherein the portion of bandwidth is allocated at a first time prior to the start-time of the addressable break and is de-allocated at a second time after an end-time of the addressable break.
2. A system for dynamically allocating bandwidth comprising:
 - a server comprising addressable content;
 - a splicer receiving a broadcast stream and inserting a stream of the addressable content from the server into the broadcast stream; and
 - an edge device detecting a cue tone, the cue tone comprising data indicating at least a start-time of an addressable break, the edge device allocating a portion of available bandwidth to the stream of the addressable content based on the data of the cue tone for a duration, the edge device de-allocating the portion of bandwidth at the end of the duration of the addressable break,

wherein the portion of bandwidth is allocated at a first time prior to the start-time of the addressable break and is de-allocated at a second time after an end-time of the addressable break.

3. The system of claim 2 further comprising a resource manager managing a pool of available bandwidth and allocating a portion of the available bandwidth to the edge device.

4. The system of claim 3 wherein the resource manager defines a plurality of subnetworks, the resource manager allocating the portion of available bandwidth based on the targeted subscribers in each subnetwork.

5. The system of claim 4 wherein the resource manager allocates the portion of bandwidth in each subnetwork based on active targeted subscribers in the subnetwork.

6. The system of claim 2 further comprising a traffic and billing module configured to manage a schedule for addressable breaks.

7. The system of claim 2 further comprising an insertion module configured to insert non-addressable content into the broadcast stream.

8. A system for dynamically allocating bandwidth comprising:

a splicer receiving and transmitting a broadcast stream, the splicer detecting a cue tone, the cue tone comprising data indicating at least a start-time of an addressable break; and

a server comprising addressable content, the server transmitting an addressable content stream, wherein the server allocates a portion of bandwidth for a duration to the addressable content stream based on the data of the cue tone and the server de-allocates the portion of bandwidth at an end of the duration of the addressable break,

wherein the portion of bandwidth is allocated at a first time prior to the start-time of the addressable break and is de-allocated at a second time after an end-time of the addressable break.

9. The system of claim 8 further comprising a resource manager managing a pool of available bandwidth and allocating a portion of the available bandwidth to the server.

10. The system of claim 9 wherein the resource manager defines a plurality of subnetworks, the resource manager serving the portion of available bandwidth based on targeted subscribers in each subnetwork.

11. The system of claim 10 wherein the resource manager allocates the portion of bandwidth based on the active targeted subscribers in each subnetwork.

12. The system of claim 8 further comprising a traffic and billing module configured to manage a schedule for addressable breaks.

13. The system of claim 8 further comprising an insertion module configured to insert non-addressable content into the broadcast stream.

14. A system for dynamically allocating bandwidth comprising:

a splicer receiving and transmitting a broadcast stream, the splicer detecting a cue tone, the cue tone comprising data indicating at least a start-time of an addressable break; and

a server comprising addressable content, the server transmitting an addressable content stream, and a resource manager receiving a request for bandwidth from the server, the resource manager defining subnetworks of available subscribers, the resource manager determining a portion of the available subscribers in the subnetwork selected to receive the addressable content, the resource manager allocating a portion of bandwidth for a duration to the addressable content stream based on the data of the cue tone and the portion of available subscribers in the subnetwork selected to receive the addressable content, the resource manager de-allocating the portion of bandwidth at the end of the duration of the addressable break,

wherein the portion of bandwidth is allocated at a first time prior to the start-time of the addressable break and is de-allocated at a second time after an end-time of the addressable break.

15. The system of claim 14 wherein the resource manager determines the available subscribers selected to receive the addressable content based on subscribers that are active to receive the broadcast stream.

16. The system of claim 14 wherein the resource manager determines the available subscribers selected to receive the addressable content based on subscribers tuned to the broadcast stream.

17. The system of claim 14 wherein the portion of the available subscribers selected to receive the addressable content is selected by at least one demographic characteristic.

18. A system for dynamically allocating bandwidth comprising:

means for receiving a cue tone from a broadcast stream, the cue tone comprising data indicating at least a start-time of an addressable break in the broadcast stream;

means for allocating a portion of available bandwidth to an addressable content stream based on the data of the cue tone;

means for streaming the addressable content stream to a receiver during the addressable break; and

means for de-allocating the portion of bandwidth at the end of the addressable break,

wherein the portion of bandwidth is allocated at a first time prior to the start-time of the addressable break and is de-allocated at a second time after an end-time of the addressable break.

19. The system of claim 18 further comprising means for splicing an advertisement into the addressable content stream to the receiver.

20. The system of claim 19 further comprising means for targeting the advertisement to a subscriber.

21. A method of allocating bandwidth comprising:

receiving a trigger from a content distribution stream, the trigger comprising data indicating at least a start-time of an addressable break in the content distribution stream;

allocating a portion of available bandwidth to an addressable content stream based on the data of the trigger;

streaming the addressable content stream to a receiver during the addressable break;

switching the receiver to the addressable content stream for the duration of the addressable break; and

de-allocating the portion of bandwidth at the end of the addressable break,

wherein the portion of bandwidth is allocated at a first time prior to the start-time of the addressable break and is de-allocated at a second time after an end-time of the addressable break.

22. The method of claim 21, wherein the portion of the available bandwidth is allocated via a resource manager.

23. The method of claim 22, wherein the resource manager is configured to define a plurality of subnetworks and to allocate the portion of available bandwidth based on targeted subscribers in each subnetwork.

24. The method of claim 22, wherein the resource manager is configured to define a plurality of subnetworks and to allocate the portion of bandwidth in each subnetwork based on active targeted subscribers in the subnetwork.

25. An apparatus comprising:

one or more processors; and memory storing instructions that, when executed by the one or more processors, cause the apparatus to:

receive a trigger from a content distribution stream, the trigger comprising data indicating at least a start-time of an addressable break in the content distribution stream;

allocate a portion of available bandwidth to an addressable content stream based on the data of the trigger;

stream the addressable content stream to a receiver during the addressable break;

switch the receiver to the addressable content stream for the duration of the addressable break; and

de-allocate the portion of bandwidth at the end of the addressable break,

wherein the portion of bandwidth is allocated at a first time prior to the start-time of the addressable break and de-allocated at a second time after an end-time of the addressable break.

26. The apparatus of claim 25, wherein the portion of the available bandwidth is allocated via a resource manager.

27. The apparatus of claim 26, wherein the resource manager is configured to define a plurality of subnetworks and to allocate the portion of available bandwidth based on targeted subscribers in each subnetwork.

28. The apparatus of claim 26, wherein the resource manager is configured to define a plurality of subnetworks and to allocate the portion of bandwidth in each subnetwork based on active targeted subscribers in the subnetwork.

29. A system for dynamically allocating bandwidth comprising:

a splicer for receiving and transmitting a content distribution stream, the splicer being configured to detect a trigger, the trigger comprising data indicating a start-time of an addressable break; and

a server comprising addressable content, the server being configured to transmit an addressable content stream, wherein the server allocates a portion of bandwidth for a duration to the addressable content stream based on the data of the trigger and the server de-allocates the portion of bandwidth at the end of the duration of the addressable break,

wherein the server allocates the bandwidth at a first time prior to the start-time of the addressable break, and de-allocates the portion of bandwidth at a second time after an end-time of the addressable break.

30. The system of claim 29, further comprising a resource manager configured to manage a pool of available bandwidth and to allocate the portion of the available bandwidth.

31. The system of claim 30, wherein the resource manager is configured to define a plurality of subnetworks and to allocate the portion of bandwidth based on targeted subscribers in each subnetwork.

32. The system of claim 30, wherein the resource manager is configured to define a plurality of subnetworks and to allocate the portion of bandwidth based on active targeted subscribers in the subnetwork.

33. The system of any one of claims 29 to 32, further comprising a traffic and billing module configured to manage a schedule for addressable breaks.

34. The system of any one of claims 29 to 33, further comprising an insertion module configured to insert non-addressable content into the content distribution stream.

35. A method of allocating bandwidth comprising:
- receiving, at a splicer, a content distribution stream;
 - inserting, by the splicer, a stream of addressable content into the content distribution stream, the addressable content being stored on a server;
 - detecting, at an edge device, a trigger from the content distribution stream, the trigger comprising data indicating at least a start-time of an addressable break in the content distribution stream;
 - allocating, by the edge device, a portion of available bandwidth to the stream based on the data of the trigger for a duration;
 - streaming the addressable content stream to a receiver during the addressable break;
 - switching the receiver to the addressable content stream for the duration of the addressable break; and
 - de-allocating the portion of bandwidth at the end of the addressable break,
- wherein the edge device allocates the bandwidth at a first time prior to the start-time of the addressable break, and de-allocates the portion of bandwidth at a second time after an end-time of the addressable break.
36. The method of claim 35, wherein the data of the trigger comprises a start-time and an end-time of the addressable break.
37. The method of claim 35, or 36, further comprising allocating the portion of bandwidth substantially at the start-time of the addressable break and de-allocating the portion of bandwidth substantially at an end-time of the addressable break.
38. A system for dynamically allocating bandwidth comprising:
- a server comprising addressable content;

a splicer for receiving a content distribution stream and inserting a stream of the addressable content from the server into the content distribution stream; and

an edge device for detecting a trigger, the trigger comprising data indicating at least a start-time of an addressable break, the edge device allocating a portion of available bandwidth to the stream of the addressable content based on the data of the trigger for a duration, the edge device de-allocating the portion of bandwidth at the end of the duration of the addressable break,

wherein the edge device allocates the bandwidth at a first time prior to the start-time of the addressable break, and de-allocates the portion of bandwidth at a second time after the end-time of the addressable break.

39. The system of claim 38, wherein the data of the trigger comprises a start-time and an end-time of the addressable break.

40. The system of claim 39, wherein the edge device allocates the portion of bandwidth substantially at the start-time of the addressable break, and de-allocates the portion of bandwidth substantially at the end-time of the addressable break.

41. The system of any one of claims 38 to 40, further comprising a resource manager that manages a pool of available bandwidth and allocates a portion of the available bandwidth to the edge device.

42. The system of claim 41, wherein the resource manager defines a plurality of subnetworks, the resource manager allocating the portion of available bandwidth based on targeted subscribers in each subnetwork.

43. The system of claim 42, wherein the resource manager allocates the portion of bandwidth in each subnetwork based on active targeted subscribers in the subnetwork.

44. The system of any one of claims 38 to 43, further comprising a traffic and billing module that manages a schedule for addressable breaks.

45. The system of any one of claims 38 to 44, further comprising an insertion module configured to insert non-addressable content into the content distribution stream.

46. A system for dynamically allocating bandwidth comprising:

means for receiving a trigger from a content distribution stream, the trigger comprising data indicating at least a start-time of an addressable break in the content distribution stream;

means for allocating a portion of available bandwidth to an addressable content stream at a first time prior to the start-time of the addressable break;

means for streaming the addressable content stream to a receiver during the addressable break; and

means for de-allocating the portion of bandwidth at a second time after the end-time of the addressable break.

47. The system of claim 46 further comprising means for allocating the portion of bandwidth substantially at the start-time of the addressable break and de-allocating the portion of bandwidth substantially at an end-time of the addressable break.

48. The system of claim 46 or 47, further comprising means for splicing an advertisement into the addressable content stream to the receiver.

49. The system of claim 48 further comprising means for targeting the advertisement to a subscriber.

50. A method comprising:

receiving, at a splicer, a content distribution stream;

inserting, by the splicer, a stream of addressable content into the content distribution stream, the addressable content being stored on a server;

detecting, at an edge device, a trigger from the content distribution stream, the trigger comprising data indicating at least a start time of an addressable break in the content distribution stream;

allocating, by the edge device, a portion of available bandwidth to an addressable content stream prior to the start time of the addressable break and for a duration;

streaming, by the edge device, the addressable content stream to a receiver during the addressable break;

switching, by the edge device, the receiver from the content distribution stream to the addressable content stream for the duration of the addressable break; and

de-allocating, by the edge device, the portion of bandwidth at an end time of the addressable break.

51. The method of claim 50, wherein the data of the trigger further comprises the end time of the addressable break.

52. The method of claim 50 or 51, further comprising managing, by a resource manager, a pool of available bandwidth and allocating, by the resource manager, a portion of the available bandwidth to the edge device.

53. The method of claim 52, further comprising defining, at the resource manager, a plurality of subnetworks, wherein the resource manager allocates the portion of available bandwidth based on targeted subscribers in each subnetwork.

54. The method of claim 53, wherein the resource manager allocates the portion of bandwidth in each subnetwork based on active targeted subscribers in the subnetwork.

55. The method of any one of claims 50 to 54, further comprising managing, at a traffic and billing module, a schedule for addressable breaks.

56. The method of any one of claims 50 to 55, further comprising inserting, at an insertion module, non-addressable content into the content distribution stream.

57. A system comprising:

a server comprising addressable content;

a splicer for receiving a content distribution stream and inserting a stream of the addressable content from the server into the content distribution stream; and

an edge device for detecting a trigger, the trigger comprising data indicating at least a start time of an addressable break, the edge device allocating a portion of available bandwidth to the stream of the addressable content prior to the start time of the addressable break and for a duration, the edge device de-allocating the portion of bandwidth at an end time of the addressable break.

58. The system of claim 57, wherein the data of the trigger further comprises an end time of the addressable break.

59. The system of claim 57 or 58, further comprising a resource manager for managing a pool of available bandwidth and allocating a portion of the available bandwidth to the edge device.

60. The system of claim 59, wherein the resource manager defines a plurality of subnetworks, the resource manager allocating the portion of available bandwidth based on targeted subscribers in each subnetwork.

61. The system of claim 60, wherein the resource manager allocates the portion of bandwidth in each subnetwork based on active targeted subscribers in the subnetwork.

62. The system of any one of claims 57 to 61, further comprising a traffic and billing module for managing a schedule for addressable breaks.

63. The system of any one of claims 57 to 62, further comprising an insertion module for inserting non-addressable content into the content distribution stream.

64. A system comprising:

means for receiving a trigger from a content distribution stream, the trigger comprising data indicating at least a start time of an addressable break in the content distribution stream;

means for allocating a portion of available bandwidth to an addressable content stream prior to the start time of the addressable break and for a duration;

means for streaming the addressable content stream to a receiver during the addressable break; and

means for de-allocating the portion of bandwidth at an end time of the addressable break.

65. The system of claim 64, further comprising means for splicing an advertisement into the addressable content stream to the receiver.

66. The system of claim 65 or 66, further comprising means for targeting the advertisement to a subscriber.

67. The system of any one of claims 64 to 66, further comprising means for managing a pool of available bandwidth and allocating a portion of the available bandwidth.

68. The system of any one of claims 64 to 67, further comprising means for managing a schedule for addressable breaks.

69. The system of any one of claims 64 to 68, further comprising means for inserting non-addressable content into the content distribution stream.

70. A method comprising:

receiving a content stream, at least a portion of the content stream comprising an addressable content stream;

detecting, in the content stream, a trigger comprising data indicating at least a start time of an addressable break in the content stream;

allocating, based on the data associated with the trigger and prior to the start time of the addressable break, a portion of available bandwidth to the addressable content stream;

streaming, to a receiver and for a duration of the addressable break, the addressable content stream; and

de-allocating, after an end time of the addressable break, the portion of available bandwidth.

71. The method of claim 70, further comprising switching, for the duration of the addressable break, the receiver to the addressable content stream.

72. The method of claim 70 or 71, wherein:

allocating, prior to the start time of the addressable break, the portion of available bandwidth comprises allocating, substantially at the start time of the addressable break, the portion of available bandwidth; and

de-allocating, after the end time of the addressable break, the portion of available bandwidth comprises de-allocating, substantially at the end time of the addressable break, the portion of available bandwidth.

73. The method of any one of claims 70 to 72, wherein the addressable content comprises a targeted advertisement.

74. The method of any one of claims 70 to 73, wherein the trigger comprises a cue tone.

75. The method of any one of claims 70 to 74, wherein the content stream further comprises non-addressable content.

76. The method of any one of claims 70 to 75, wherein the data further comprises the end time of the addressable break.

77. An apparatus comprising a processor and a memory, the memory storing computer-executable instructions which, when executed by the processor, cause the apparatus to perform operations comprising:

receiving a content stream, at least a portion of the content stream comprising an addressable content stream;

detecting, in the content stream, a trigger comprising data indicating at least a start time of an addressable break in the content stream;

allocating, based on the data associated with the trigger and prior to the start time of the addressable break, a portion of available bandwidth to the addressable content stream;

streaming, to a receiver and for a duration of the addressable break, the addressable content stream; and

de-allocating, after an end time of the addressable break, the portion of available bandwidth.

78. The apparatus of claim 77, wherein the operations further comprise switching, for the duration of the addressable break, the receiver to the addressable content stream.

79. The apparatus of claim 77 or 78, wherein:

allocating, prior to the start time of the addressable break, the portion of available bandwidth comprises allocating, substantially at the start time of the addressable break, the portion of available bandwidth; and

de-allocating, after the end time of the addressable break, the portion of available bandwidth comprises de-allocating, substantially at the end time of the addressable break, the portion of available bandwidth.

80. The apparatus of any one of claims 77 to 79, wherein the addressable content comprises a targeted advertisement.

81. The apparatus of any one of claims 77 to 80, wherein the trigger comprises a cue tone.

82. The apparatus of any one of claims 77 to 81, wherein the content stream further comprises non-addressable content.

83. The apparatus of any one of claims 77 to 82, wherein the data further comprises the end time of the addressable break.

84. A computer-readable storage medium comprising computer-executable instructions which, when executed by a device, cause the device to perform operations comprising:

receiving a content stream, at least a portion of the content stream comprising an addressable content stream;

detecting, in the content stream, a trigger comprising data indicating at least a start time of an addressable break in the content stream;

allocating, based on the data associated with the trigger and prior to the start time of the addressable break, a portion of available bandwidth to the addressable content stream;

streaming, to a receiver and for a duration of the addressable break, the addressable content stream; and

de-allocating, after an end time of the addressable break, the portion of available bandwidth.

85. The computer-readable storage medium of claim 84, further comprising switching, for the duration of the addressable break, the receiver to the addressable content stream.

86. The computer-readable storage medium of claim 84 or 85, wherein:

allocating, prior to the start time of the addressable break, the portion of available bandwidth comprises allocating, substantially at the start time of the addressable break, the portion of available bandwidth; and

de-allocating, after the end time of the addressable break, the portion of available bandwidth comprises de-allocating, substantially at the end time of the addressable break, the portion of available bandwidth.

87. The computer-readable storage medium of any one of claims 84 to 86, wherein the addressable content comprises a targeted advertisement.

88. The computer-readable storage medium of any one of claims 84 to 87, wherein the trigger comprises a cue tone.

89. The computer-readable storage medium of any one of claims 84 to 88, wherein the content stream further comprises non-addressable content.

90. A method comprising:

receiving data indicating at least a start time of an addressable break in a content stream;

allocating, based on the data, prior to the start time of the addressable break, a portion of available bandwidth to an addressable content stream;

transmitting, to a receiver and for an entire duration of the addressable break starting at the start time, the addressable content stream; and

de-allocating, after the entire duration of the addressable break, the portion of available bandwidth.

91. The method of claim 90, further comprising switching, for the entire duration of the addressable break, the receiver to the addressable content stream.

92. The method of claim 90 or claim 91, wherein:

allocating the portion of the available bandwidth comprises allocating, substantially at the start time of the addressable break, the portion of the available bandwidth; and

de-allocating the portion of the available bandwidth comprises de-allocating, substantially at an end time of the addressable break, the portion of the available bandwidth.

93. The method of any one of claims 90 to 92, wherein the addressable content stream comprises a targeted advertisement.

94. The method of any one of claims 90 to 93, wherein the data comprises at least one of a trigger and a cue tone.

95. The method of any one of claims 90 to 94, wherein at least a portion of the content stream comprises non-addressable content.

96. The method of any one of claims 90 to 95, wherein the data further indicates an end time of the addressable break, the entire duration of the addressable break being defined between the start time and the end time.

97. An apparatus comprising:

a processor; and

a memory storing instructions that, when executed by the processor, cause the

apparatus to:

receive data indicating at least a start time of an addressable break in a content stream;

allocate, based on the data, prior to the start time of the addressable break, a portion of available bandwidth to an addressable content stream;

transmit, to a receiver and for an entire duration of the addressable break starting at the start time, the addressable content stream; and

de-allocate, after the entire duration of the addressable break, the portion of available bandwidth.

98. The apparatus of claim 97, wherein the instructions, when executed, further cause the apparatus to switch, for the entire duration of the addressable break, the receiver to the addressable content stream.

99. The apparatus of claim 97 or claim 98, wherein:
allocating the portion of the available bandwidth comprises allocating, substantially at the start time of the addressable break, the portion of the available bandwidth; and
de-allocating the portion of the available bandwidth comprises de-allocating, substantially at an end time of the addressable break, the portion of the available bandwidth.

100. The apparatus of any one of claims 97 to 99, wherein the addressable content stream comprises a targeted advertisement.

101. The apparatus of any one of claims 97 to 100, wherein the data comprises at least one of a trigger and a cue tone.

102. The apparatus of any one of claims 97 to 101, wherein at least a portion of the content stream comprises non-addressable content.

103. The apparatus of any one of claims 97 to 102, wherein the data further indicates an end time of the addressable break, the entire duration of the addressable break being defined between the start time and the end time.

104. A non-transitory computer-readable storage medium storing instructions that, when executed by a processor of an apparatus, cause the apparatus to:
receive data indicating at least a start time of an addressable break in a content stream;
allocate, based on the data, prior to the start time of the addressable break, a portion of available bandwidth to an addressable content stream;
transmit, to a receiver and for an entire duration of the addressable break starting at the start time, the addressable content stream; and

de-allocate, after the entire duration of the addressable break, the portion of available bandwidth.

105. The non-transitory computer-readable storage medium of claim 104, wherein the instructions, when executed, further cause the apparatus to switch, for the entire duration of the addressable break, the receiver to the addressable content stream.

106. The non-transitory computer-readable storage medium of claim 104 or claim 105, wherein:

allocating the portion of the available bandwidth comprises allocating, substantially at the start time of the addressable break, the portion of the available bandwidth; and

de-allocating the portion of the available bandwidth comprises de-allocating, substantially at an end time of the addressable break, the portion of the available bandwidth.

107. The non-transitory computer-readable storage medium of any one of claims 104 to 106, wherein the addressable content stream comprises a targeted advertisement.

108. The non-transitory computer-readable storage medium of any one of claims 104 to 107, wherein the data comprises at least one of a trigger and a cue tone.

109. The non-transitory computer-readable storage medium of any one of claims 104 to 108, wherein the data further indicates an end time of the addressable break, the entire duration of the addressable break being defined between the start time and the end time.




Figure 1.




Figure 2.




Figure 3.




Figure 4.

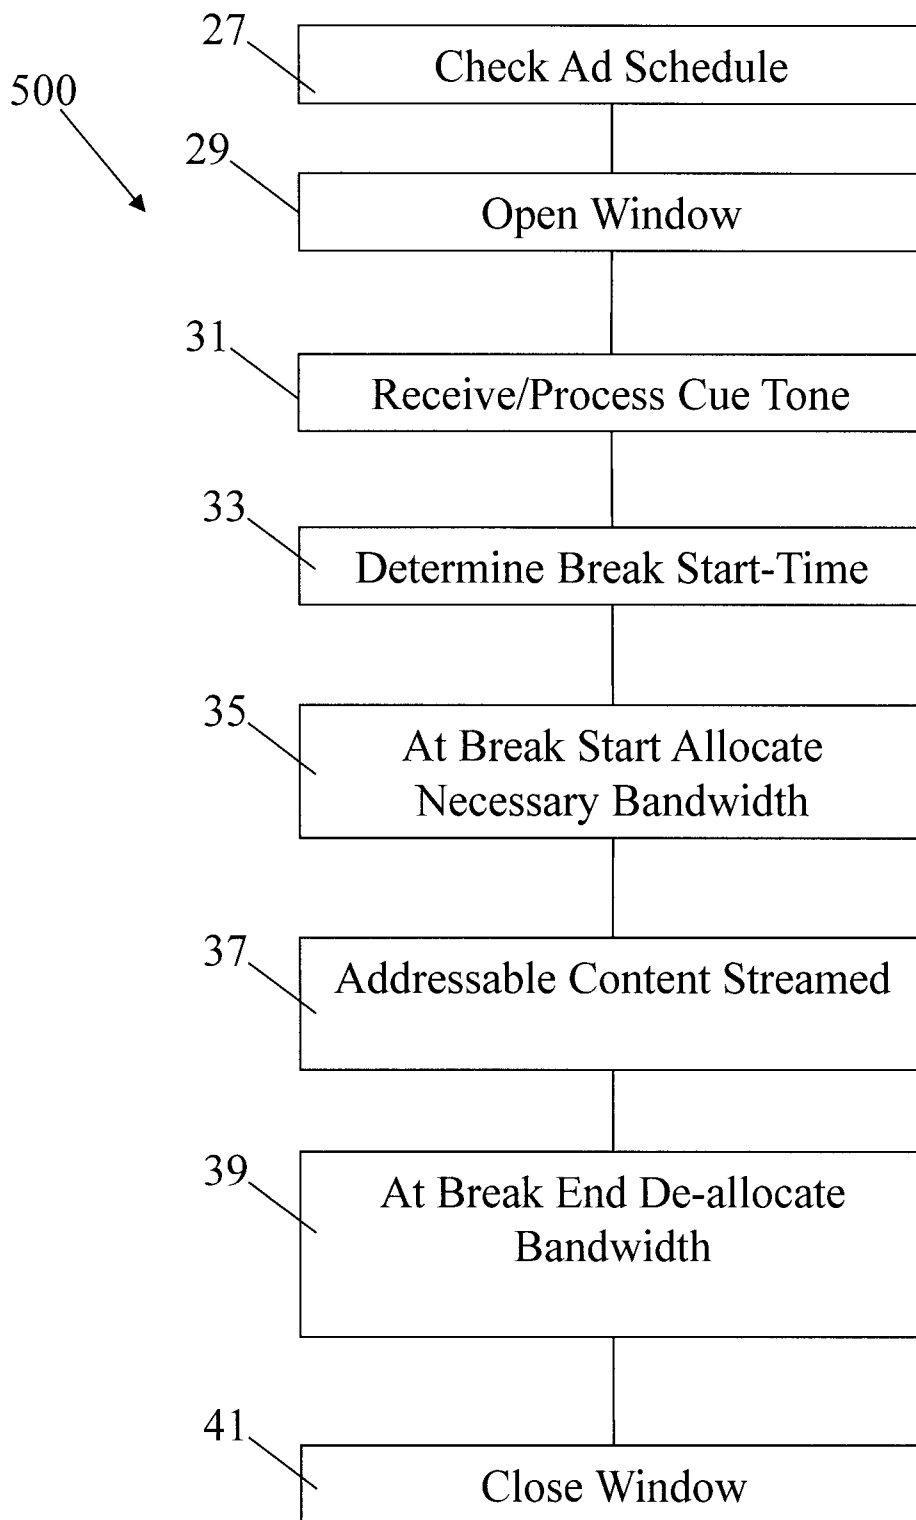


Figure 5.

600

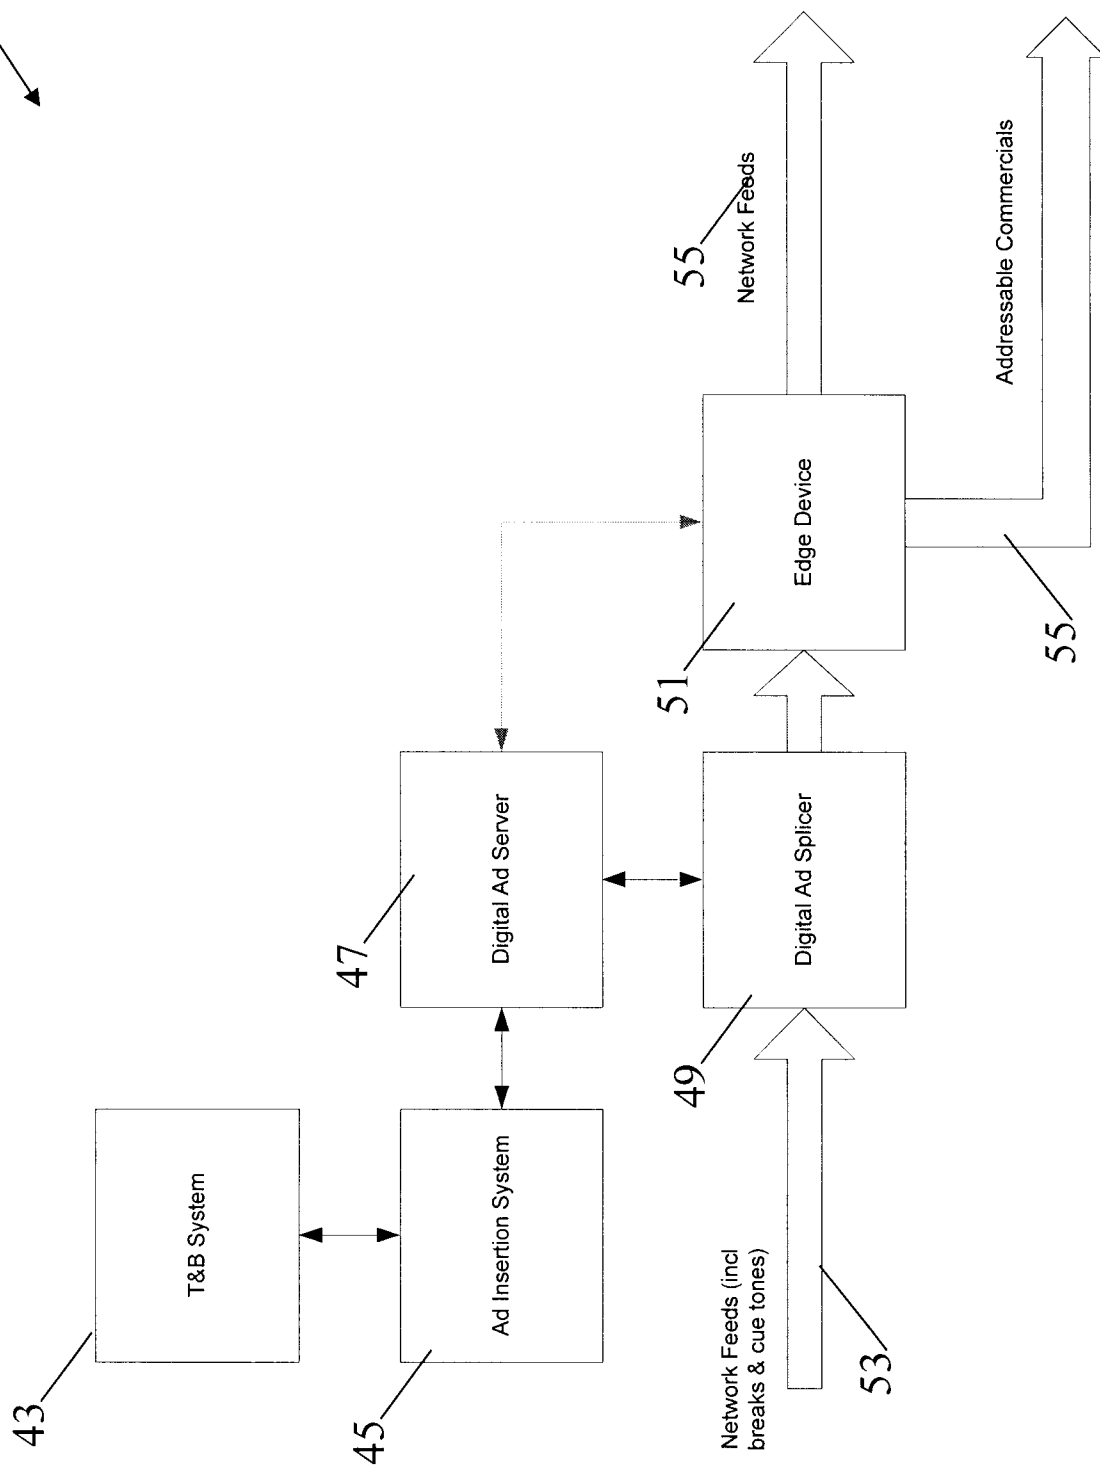


Figure 6.

700

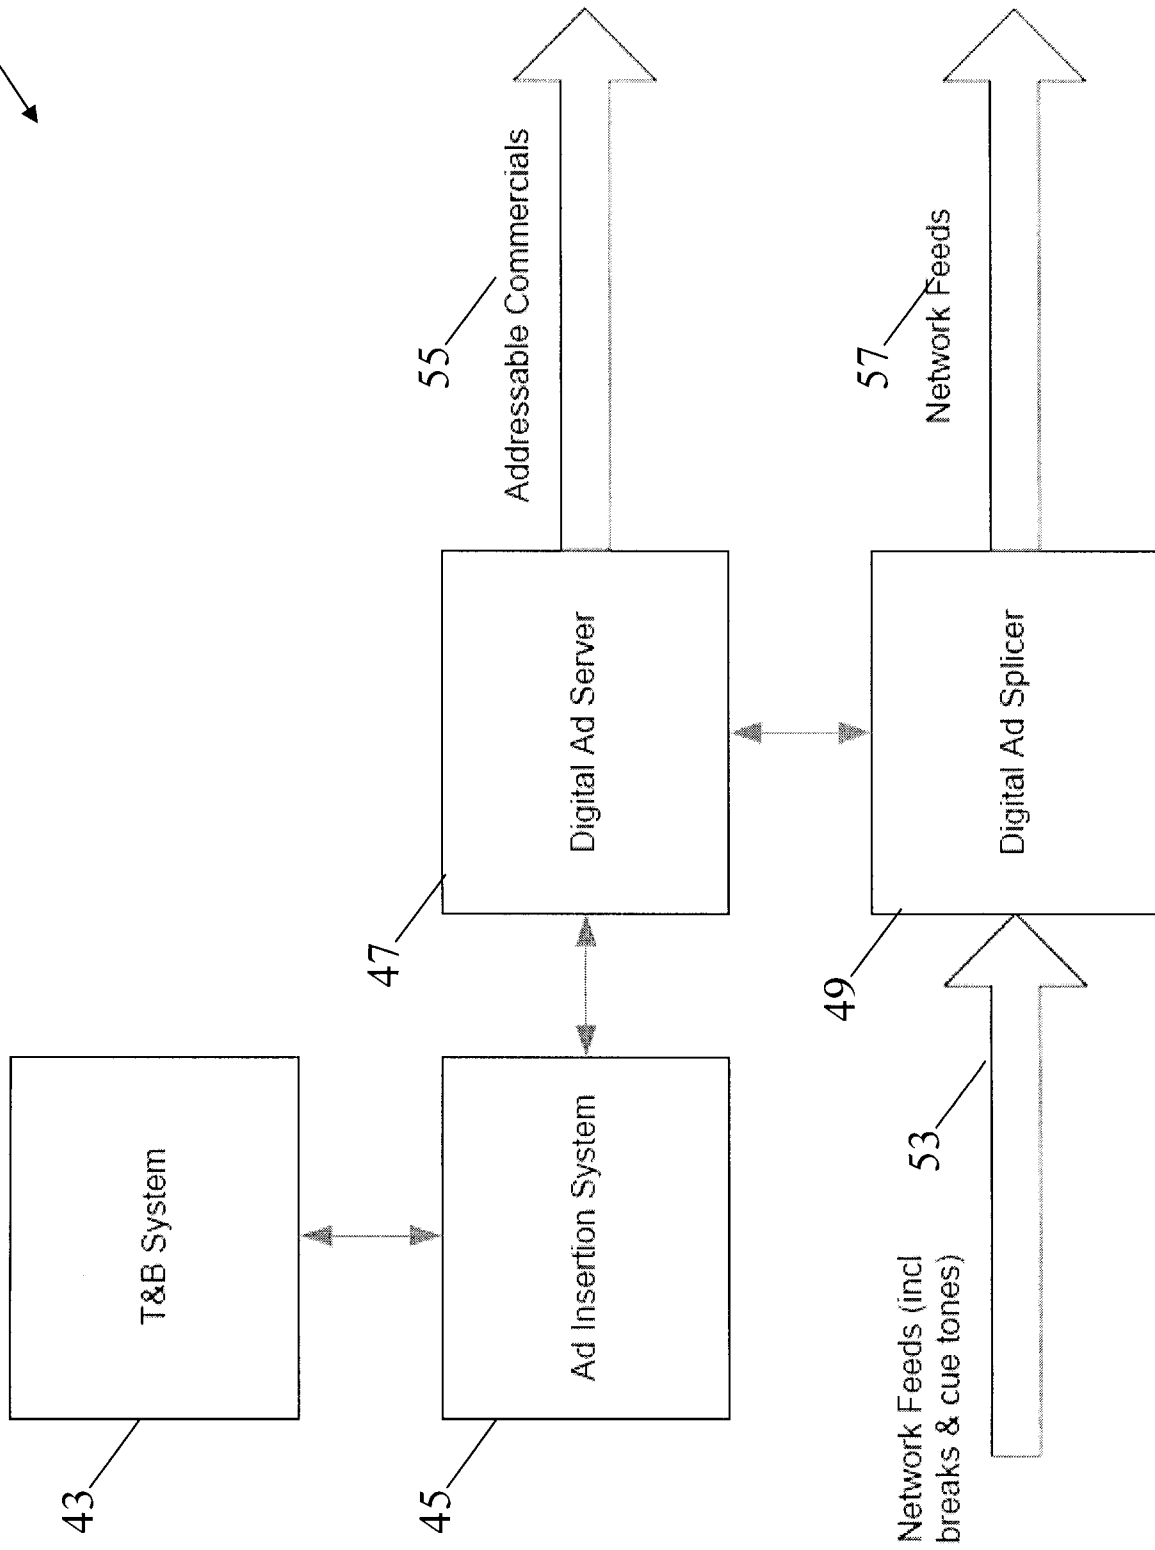


Figure 7.

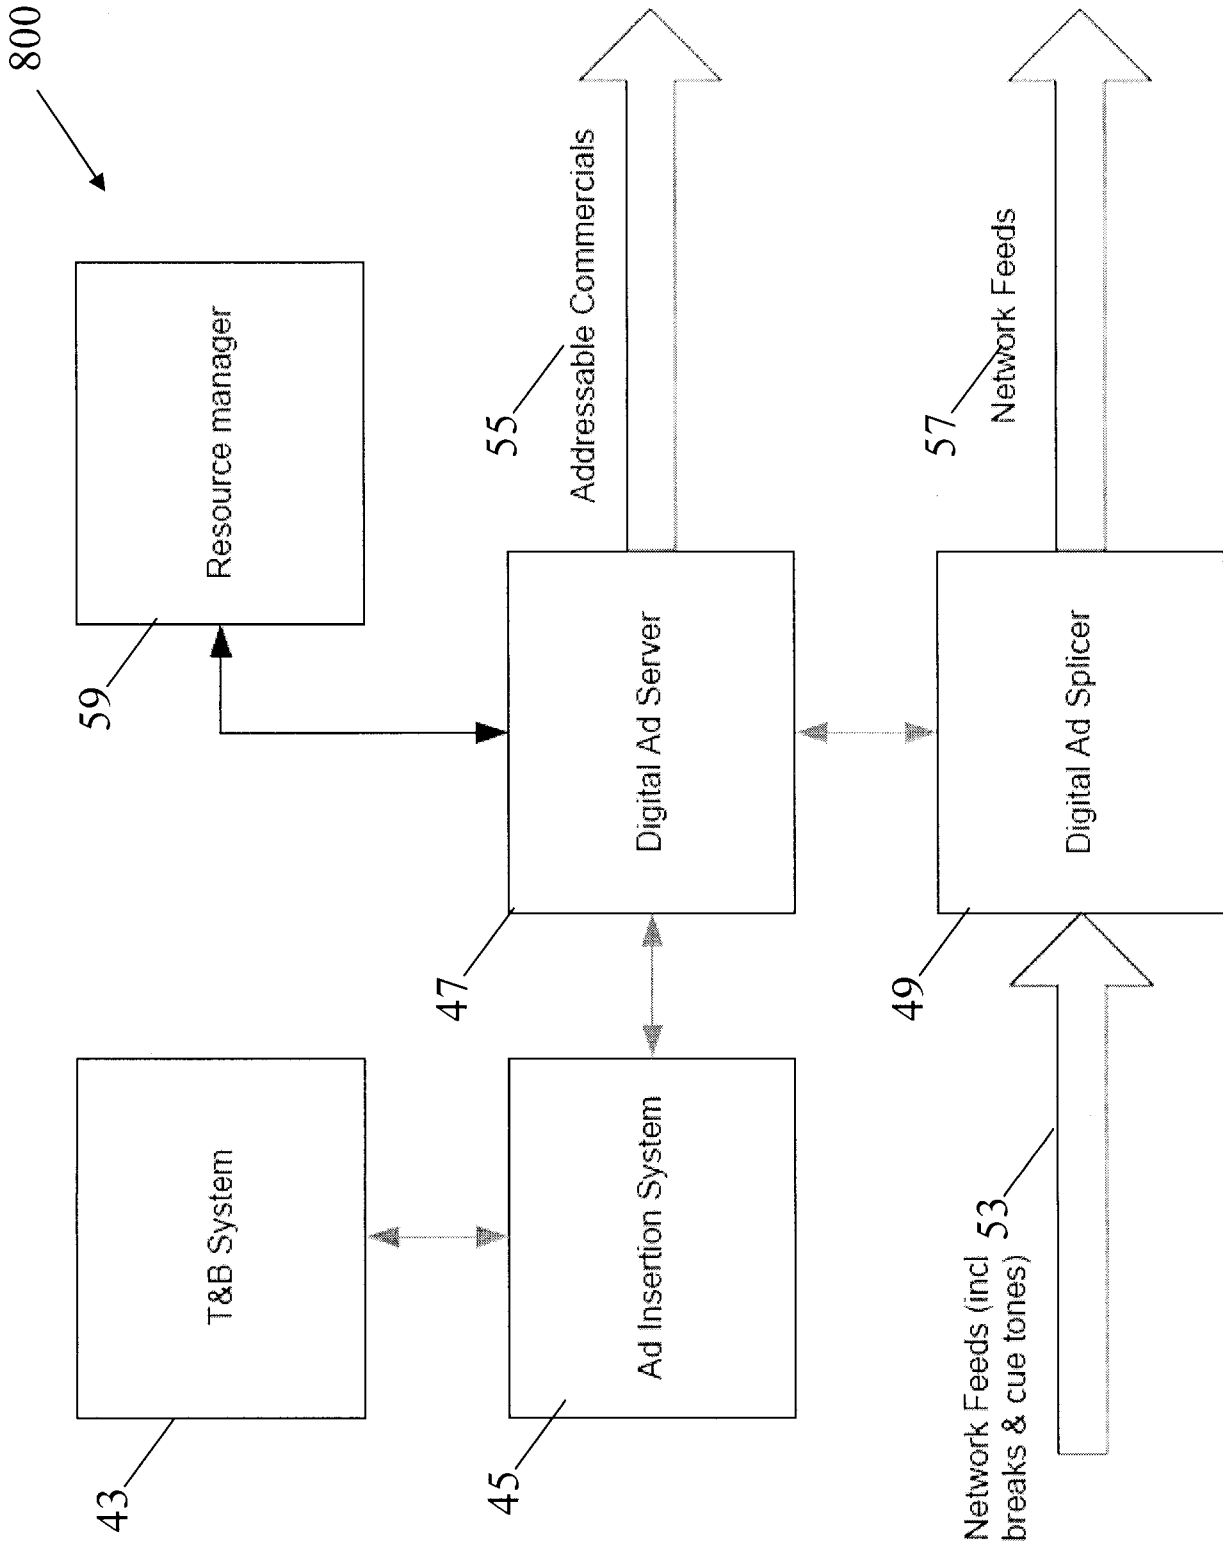


Figure 8.

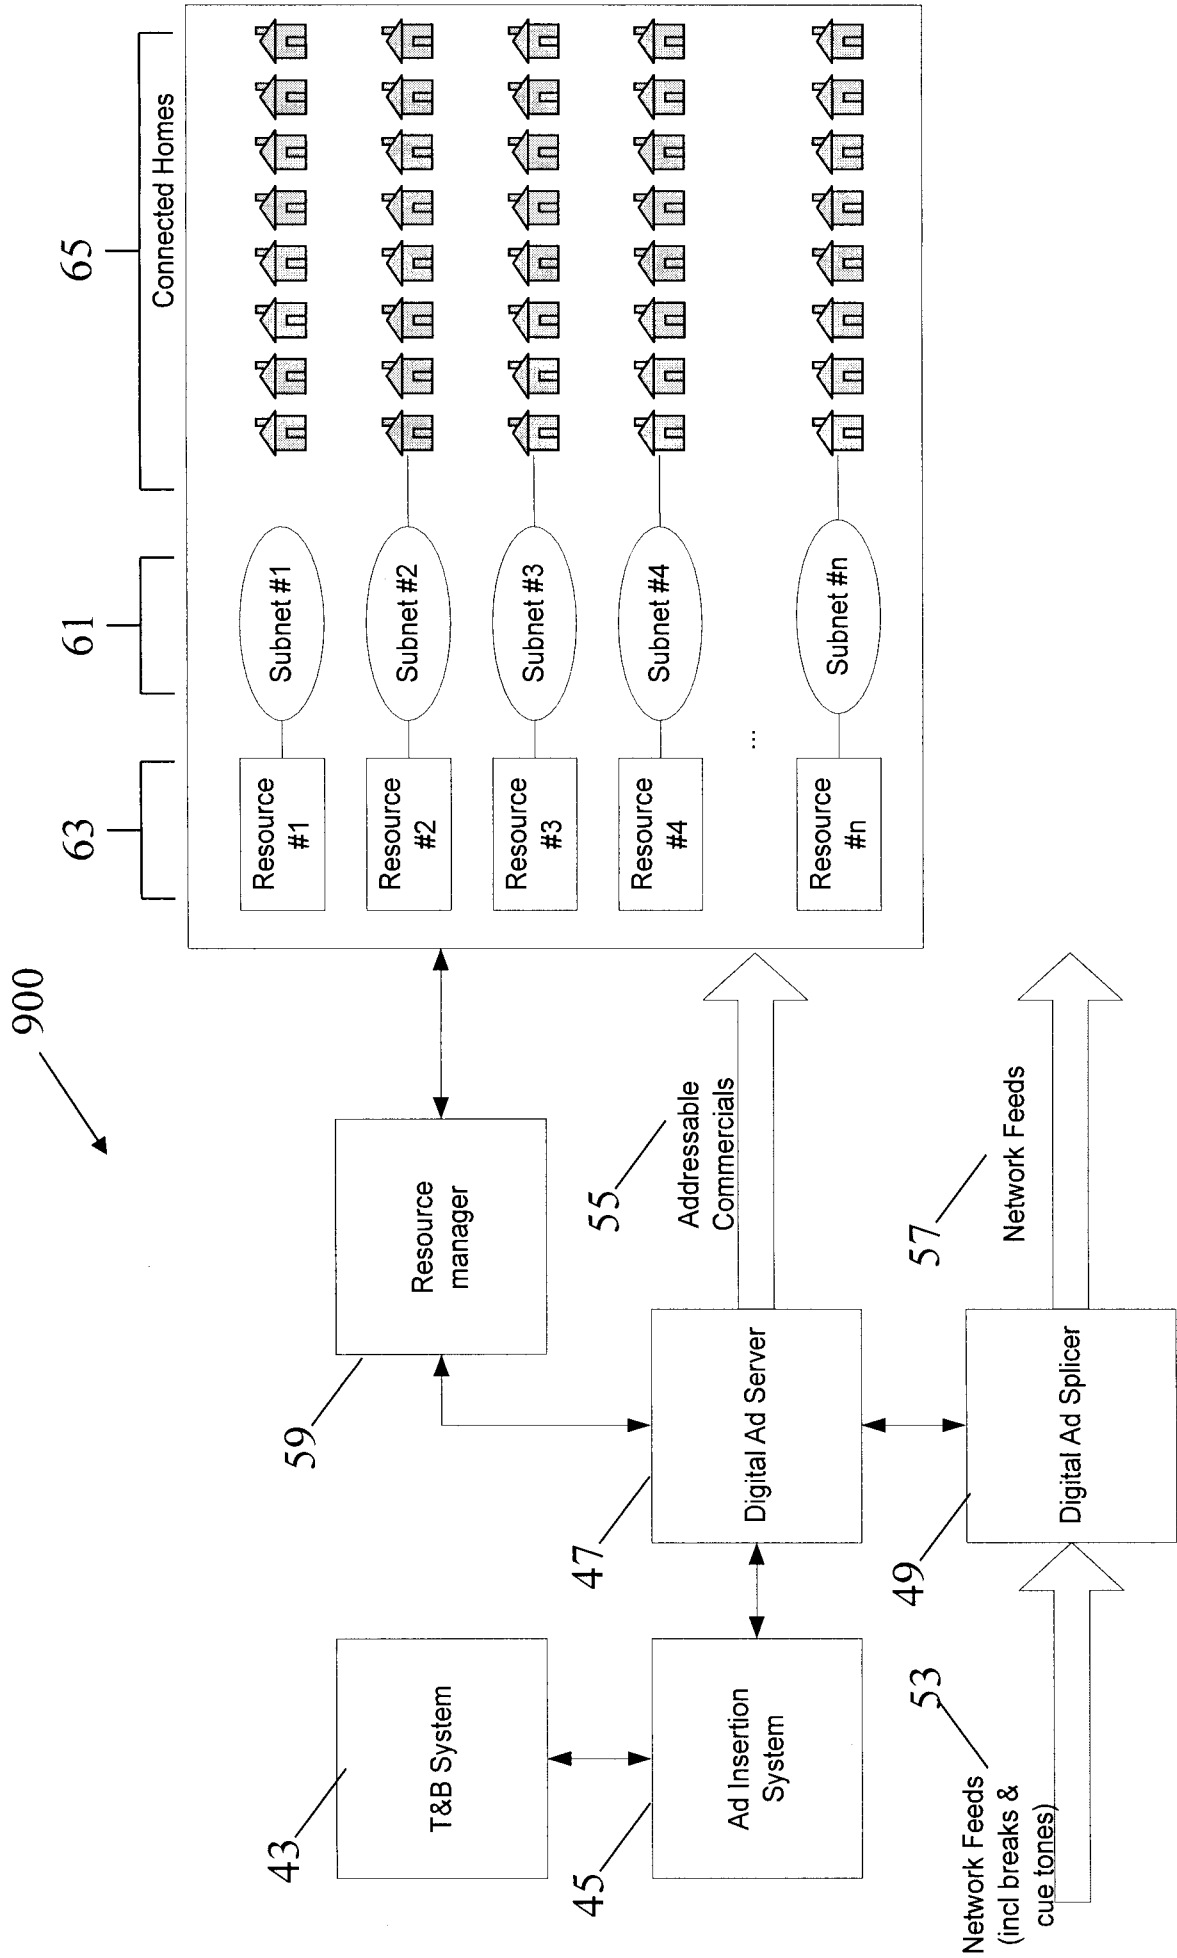


Figure 9.




Figure 10.

