

(19) World Intellectual Property
Organization
International Bureau



(43) International Publication Date
22 July 2004 (22.07.2004)


PCT

(10) International Publication Number
WO 2004/061608 A2

- (51) International Patent Classification⁷: **G06F**
- (21) International Application Number:
PCT/US2003/041320
- (22) International Filing Date:
23 December 2003 (23.12.2003)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
10/334,291 31 December 2002 (31.12.2002) US
- (71) Applicant: **MOTOROLA, INC., A CORPORATION OF THE STATE OF DELAWARE** [US/US]; 1303 East Algonquin Road, Schaumburg, IL 60196 (US).
- (72) Inventors: **VASUDEVAN, Venugopal**,; 4741 Fairfax Avenue, Palatine, IL 60067 (US). **MYSORE, Jayanth P.**,; 59 Hummingbird Lane, Streamwood, IL 60107 (US). **HANNEEF, Anwar M.**,; 980 Countryside Drive, Apt. #103, Palatine, IL 60067 (US). **ALMAULA, Jay R.**,; 1121 Pine Valley Drive, #208, Schaumburg, IL 60173 (US).
- (74) Agents: **MAY, Steven A.**, et al.; MOTOROLA, INC., Intellectual Property Dept., 1303 East Algonquin Road, Schaumburg, IL 60196 (US).
- (81) Designated States (*national*): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, UZ, VC, VN, YU, ZA, ZM, ZW.
- (84) Designated States (*regional*): ARIPO patent (BW, GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IT, LU, MC, NL, PT, RO, SE, SI, SK, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).
- Published:**
— without international search report and to be republished upon receipt of that report

[Continued on next page]

(54) Title: SYSTEM AND METHOD FOR RENDERING CONTENT ON MULTIPLE DEVICES



(57) Abstract: A method for delivering content to be rendered by multiple devices (120) is provided. Indications of resources available to a user are received (356), the resources including rendering resources that are provided by the multiple devices (120), each of the multiple devices being coupled to a network (108, 116), wherein at least one of the multiple devices provides at least one rendering resource available for use of the user and provides at least one rendering resource available for simultaneous use of another user. Content requested by the user is received (162), the content including multiple content types. A mapping of content types to the multiple devices is determined (166), wherein the mapping is based on rendering resources provided by each of the multiple devices. Content types of the content requested by the user are then delivered (178) via the network to the multiple devices according to the mapping.



For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

SYSTEM AND METHOD FOR RENDERING CONTENT ON MULTIPLE DEVICES

5 Cross Reference to Related Applications

The present disclosure is related to commonly owned U.S. Patent Application No. 10/334,848, entitled "Method and Apparatus for Linking Multimedia Content Rendered Via Multiple Devices," filed on December 31, 2002 (Attorney Docket No. CS20693RL), which is herein incorporated by reference in its entirety for all purposes.

Technical Field

The present disclosure relates to rendering multimedia content, and more particularly to rendering multimedia content via multiple devices.

Background

Some mobile computing devices with limited display capabilities are also provided with web browsing capability. For example, a personal digital assistant (PDA) may include a wireless modem and web browsing software. As another example, a cellular phone may include a display and web browsing software. But because of typical limitations of these devices (e.g., power, memory, display resolution, screen size), they may not be capable of adequately rendering content from the internet that is intended for full-sized computers (e.g., desktop computers, laptop computers, etc.). For example, a display of a PDA or cell phone may be incapable of displaying anything other than text or simple icons because of its display resolution and screen size. Alternatively, the rendering of content by the device may be of too poor quality to be useful. For instance, a picture from a web page that is displayed on a mobile device may be unintelligible because of the display's low resolution, and/or small screen size.

Similarly, mobile computing devices typically may not be capable of adequately rendering audio or video streams that may be available to it. For example, a PDA may not include an audio system, or its audio system may be capable of generating only lower quality audio. Such a PDA may not be capable of adequately rendering, for example, an MP3 music file available from the internet.

In some instances, the communication link through which the mobile computing device is linked to the Internet may be a bottleneck. For instance, the bandwidth of a cellular link may not permit adequate rendering of a video stream.

In general, content available via the Internet may not be able to be rendered as intended on many mobile computing devices.

Brief Description of the Drawings

FIG. 1 is a block diagram of an example system that enables delivering of content to multiple devices for rendering on the multiple devices.

FIG. 2 is a flow diagram of an example method for delivering content to multiple devices for rendering on the multiple devices.

FIG. 3 is a block diagram of an example resource manager, and illustrates an example of its interaction with other systems.

FIG. 4 is a flow diagram of an example method that may be implemented by the resource manager client.

FIG. 5 is a flow diagram of an example method that may be implemented by the resource manager.

FIG. 6 is a flow diagram of an example method that may be implemented by the resource manager client of an owned rendering device.

FIG. 7 is a flow diagram of an example method that may be implemented by the resource manager.

FIG. 8 is a block diagram of an example multi-device proxy server, and illustrates an example of its interaction with other systems.

FIG. 9 is a flow diagram of an example method that may be implemented by the multi-device proxy server.

FIG. 10 is a flow diagram of another example method that may be implemented by the multi-device proxy server.

FIG. 11 is a block diagram of an example delivery control subsystem.

FIG. 12 is a flow diagram of an example method that may be implemented by the multi-device proxy server.

FIG. 13 is a flow diagram of an example method that may be implemented by
5 the multi-device proxy server.

FIG. 14 is a table illustrating one example format for representing user preferences related to preferred devices for rendering particular types of content.

FIG. 15 is a table illustrating one example format for representing user preferences related to preferred communication links for delivering content to
10 particular types of devices.

FIG. 16 is a table illustrating one example format for representing user preferences related to whether content should be delivered to the same device or different devices.

FIG. 17 is a flow diagram of an example method that may be implemented by
15 the delivery control subsystem.

FIG. 18 is a flow diagram of another example method that may be implemented by the delivery control subsystem.

FIG. 19 is a flow diagram of an example method that may be implemented by the delivery control subsystem.

FIG. 20 is a block diagram of an example subsystem for providing user interface mechanisms that may assist a user with interpreting and navigating content being rendered on multiple devices.
20

FIG. 21 is a flow diagram of an example method that may be implemented by the example subsystem of FIG. 20.

FIG. 22 is a flow diagram of an example method that may be implemented by
25 the example subsystem of FIG. 20.

Detailed Description

Example systems and methods are described in which content may be
30 rendered by a plurality of devices. By rendering content on a plurality devices, limitations of a single device may be mitigated. For example, a user may have a PDA with a low resolution display. Additionally, the user may have access to a laptop

computer having a display with much higher resolution and/or larger screen size. As described herein, the user may be able to render, for example, a web page that includes text, low resolution graphics, and high resolution graphics using the PDA and the laptop computer. For instance, text and low resolution graphics of the web
5 page can be delivered to, and rendered by, the PDA, and high resolution graphics of the web page can be delivered to, and rendered by, the laptop computer.

System Overview

FIG. 1 is a block diagram of an example environment 100 in which content
10 may be delivered to multiple devices for rendering on the multiple devices. A content server 104 may be coupled to a network 108. The content server 104 may provide content to users via the network 108. The content may include, for example, web pages, streaming video, streaming audio, etc. A single content server 104 may provide multiple types of content (e.g., web pages, video, audio), or different content
15 servers may provide different types of content (e.g., the environment 100 may include a web page server, a video server, and/or an audio server). The network 108 may comprise a wide area network (WAN), the Internet, an intranet, an extranet, a local area network (LAN), etc.

A multi-device proxy server 112 may also be coupled to the network 108 and
20 to a network 116, and rendering devices 120a and 120b are coupled with the network 116. The multi-device proxy server 112 will be described in more detail below. In general, the multi-device proxy server 112 may act as a gateway between the rendering devices 120a and 120b and the network 108. For example, if information is to be transmitted from a device coupled to the network 108, such as the content server
25 104, to the rendering device 120a, such information should pass through the multi-device proxy server 112. The network 116 may comprise a local area network, an internet, an intranet, an extranet, a local area network, etc. Although the network 116 is illustrated as being separate from the network 108, the network 116 may be, or include, a subset of the network 108.

30 A resource manager 124 may be coupled to the network 116. Although the resource manager 124 is illustrated as being separate from the multi-device proxy server 112, some or all of the resource manager 124 may be a component of the multi-

device proxy server 112. The resource manager 124 and rendering devices 120a and 120b may be located in a physical neighborhood that is indicated by the box 128. In general, the resource manager 124 may interact with a user device, for example rendering device 120a, to permit a user to access other devices in the neighborhood
5 128, for example rendering device 120b, for rendering content on those other devices.

One or more rendering devices 120c may also be included that may not be in the physical neighborhood 128 of other rendering devices (e.g., rendering devices 120a and 120b). Such rendering devices 120c may be coupled to the network 108. Although device 120c is termed a rendering device, the device 120c need not be used
10 to render content. For example, the device 120c could be used to store content for later rendering by another device. The rendering devices 120 may comprise a cell phone, a pager, a PDA, a laptop computer, a desk top computer, a workstation, a server, etc.

Also, a transcoder 132 may be coupled to the network 108. The transcoder
15 132 may be employed to transcode content from the content server 104 that may be in a format that cannot be rendered by one of the rendering devices 120 into a format that can be rendered.

FIG. 2 is a flow diagram of an example method 150 for delivering content to multiple devices, and will be described with reference to FIG. 1.

20 At block 152, a user, via the rendering device 120a, for example, may discover the resource manager 124 in a neighborhood in which the user is located. This discovery may be accomplished using various techniques including standardized discovery mechanisms such as the Service Location Protocol (SLP), the BLUETOOTH™ standard, Sun Microsystems' Jini technology, technology from the
25 Salutation Consortium, etc.

At block 154, the user may reserve resources for rendering content on devices other than the user's main device. The resources may be software resources or hardware resources. For example, if a user's main device is rendering device 120a, the user may reserve resources on another device in the neighborhood 128, such as
30 rendering device 120b. The user may interact with the resource manager 124, using the rendering device 120a, to reserve resources on the rendering device 120b for

rendering content. For example, if the rendering device 120b includes an audio system, the user may reserve the audio system of the device 120b.

Additionally, the user may reserve resources on, or request that content be delivered to, a remote device, such as device 120c. For instance, the user may instruct
5 the multi-device proxy server 112 to deliver certain types of content to the remote device 120c, or to deliver certain specific content (e.g., a file, audio stream, video stream, etc.) to the remote device 120c. In one example, the user may reserve resources on the remote device 120c, or request delivery to the remote device 120c, without interacting with a resource manager 124.

10 At block 158, the user may request content from the content server 104. For example, the user may request content using the rendering device 120a.

At block 162, content from the content server 104 that was requested by the rendering device 120a may be received by the multi-device proxy server 112. At block 166, the multi-device proxy server 112 may determine a mapping of the
15 requested content to the devices to which the user has access. For example, the mapping may indicate text and graphics should go to device 120a while audio should go to device 120b or device 120c.

In some instances, the multi-device proxy server 112 may download software to a rendering device 120 if it is determined that the rendering device 120 cannot
20 render the content in the format provided by the content server 104 (block 168). The downloaded software may provide the rendering device 120 with the capability to render the content in the format provided by the content server 104.

In some instances, the multi-device proxy server 112 may utilize the transcoder 132 if it is determined that a rendering device cannot render content in the
25 format provided by the content server 104. In such cases, the multi-device proxy server 112 may transmit a portion of the content to the transcoder 132 at block 170. As an example, the transcoder may transcode stereo audio data into mono audio data. Then, at block 174, the multi-device proxy server 112 may receive the transcoded portion of the content from the transcoder 132. Although the transcoder 132 is
30 illustrated as a separate device in FIG. 1, the multi-device proxy server 112 may also include a transcoder as an alternative, or in addition, to the transcoder 132.

At block 178, the multi-device proxy server 112 may deliver the content, or a portion of the content, to the multiple devices according to the mapping determined at block 166.

Also, in some instances, a user may wish to have content delivered to a remote device, such as rendering device 120c. In these cases, the user may instruct the multi-device proxy server 112 to deliver, for example, certain types of content to the remote device 120c. The multi-device proxy server 112 then determines a mapping of the requested content to the devices that to which the user has access, including the remote device 120c. For example, the mapping may indicate text and graphics should go to device 120a while audio should go to device 120c.

Next, the multi-device proxy server 112 delivers the content to the appropriate devices according to the mapping. For example, text and graphics could be delivered to a user's PDA, and audio could be delivered to a desktop, located near the user (e.g. in a conference room in which the user is located), with an audio system. Further, the audio could also be delivered to the user's desktop computer at the user's home for storage on the computer's hard drive, a CD-ROM, a DVD, etc., and/or for processing by software on the user's desktop computer.

Resource Manager

The resource manager 124 may be located within, nearby, or remote from the physical neighborhood 128. As discussed above, some or all of the resource manager 124 may be a component of the multi-device proxy server 112. Alternatively, the resource manager 124 may be entirely separate from the multi-device proxy server 112. The resource manager 124 may be implemented using one or more computers such as a personal computer, workstation, server, mainframe, set-top box, cellular base station computer, etc. In some examples, one instance of a resource manager 124 may correspond to each neighborhood, whereas in other examples, one instance of a resource manager 124 may correspond to several neighborhoods.

FIG. 3 is a block diagram of an example resource manager 124, and illustrates an example of its interaction with other systems, including rendering devices 120. The resource manager 124 may include a resource allocation subsystem 204, a user interaction subsystem 206, a resource monitoring subsystem 208, and a shared

resources database 212. The resource manager 124 may also include a communication agent 216 for communicating with other systems such as rendering devices 120 and the multi-device proxy server 112. Any type of suitable protocol may be used to communicate with other systems. For example, protocols such as the
5 Session Invitation Protocol (SIP), Hyper Text Transfer Protocol (HTTP), Simple Mail Transfer Protocol (SMTP), etc. or a similar proprietary protocol may be used.

Communication with other systems may be, for example, via the network 108, the network 116, and/or direct communication with rendering devices proximate to the resource manager 124. Such communication may include communication via a
10 wired communication link such as a LAN, a telephone line, a cable line, a fiber optic line, etc., or a wireless link such as a wireless LAN (e.g., the IEEE 802.11x standards), a wireless communication link according to the BLUETOOTH™ standard, a cellular link, a two-way paging link, etc.

The resource monitoring subsystem 208 may communicate with devices in a
15 neighborhood to determine what resources are available to users. For example, a conference room may include a computer having a full-size display and an audio system. The resource monitoring subsystem 208 may communicate with this computer to ascertain that the computer has resources available for use, such as its display and its audio system. The resource manager 124 may treat a device as an
20 aggregate of component resources. Thus, multiple users may be able to share a single device.

The user interaction subsystem 206 may interact with a user to permit the user to reserve resources on devices in the neighborhood that are managed by the resource manager 124. The user interaction subsystem 206 may advertise services of the
25 resource manager to users in the neighborhood. It may advertise using various techniques including standardized mechanisms such as the SLP, the BLUETOOTH™ standard, Sun Microsystems' Jini technology, technology from the Salutation Consortium, etc.

For example, a user in a conference room may, using a PDA, communicate
30 with the user interaction subsystem 206 to determine what resources are available in the conference room, and to reserve some or all of those resources.

The user interaction subsystem 206 may communicate with the shared resources database 212 to store/retrieve information relating to, for example, total resources in a neighborhood, reserved resources in the neighborhood, available resources in the neighborhood, etc. It may additionally communicate with the resource allocation subsystem to enforce access control policies, charging policies, etc. set by the owner(s) of the resources that the user wishes to reserve.

The resource manager 124 may communicate with devices in a neighborhood 128, such as rendering devices 120a and 120b. Rendering device 120a may be, for example, a device having resources that a user does not want to make available to others. Such a device will be referred to herein as an "owned device." Rendering device 120b may be, for example, a device having resources which may be made available for use by others. Such a device will be referred to herein as a "rentable device." The rendering devices 120a and 120b may each include resource manager clients 230a and 230b respectively. The resource manager clients 230a and 230b may be the same, or they may differ, for example, depending on whether its corresponding rendering device is an owned device or a rentable device. The resource manager clients 230 may interact with the resource manager 124 to request resources for a user, notify the resource manager 124 regarding the device's resources, update the resource manager 124 about the current state of a resource, etc.

FIG. 4 is a flow diagram of an example method that may be implemented by the resource manager client 230. The method 250 may be implemented by a processor configured by software on a tangible medium such as a CD-ROM, a floppy disk, a hard drive, a digital versatile disk (DVD), or a memory associated with the processor. Persons of ordinary skill in the art will readily appreciate that the entire method 250 or parts thereof could alternatively be executed by a device other than a processor, and/or embodied in firmware and/or dedicated hardware in a well known manner. Further, although the example method is described with reference to the flow diagram in FIG. 4, persons of ordinary skill in the art will readily appreciate that many other methods may alternatively be used. For example, the order of execution of the blocks may be changed, and/or the blocks may be changed, eliminated, or combined.

The method 250 may be used to inform the resource manager 124 and/or the multi-device proxy server 112 of resource capabilities and/or availability of a rendering resource.

At block 254, capabilities of the rendering device may be determined. For example, device capabilities may be determined by querying system software, reading from a list, registry, database, etc. Device capabilities may include information relating to the device type, the device's processor (e.g., type, speed, etc.), memory/storage (e.g., type, size, etc.), communication link (e.g., type, proxy server for the particular link, maximum bandwidth, etc.), display (e.g., display size, display resolution, whether the display is a color display, bits per pixel, etc.), sound system (whether present, type, etc.), etc.

At block 258, the device capabilities may be registered with the resource manager 124 and/or the multi-device proxy server 112. For example, an owned device may register device capabilities with the multi-device proxy server 112, whereas a rentable device may register device capabilities with a resource manager 124 associated with a neighborhood in which the rentable device is located. The device capabilities may be expressed using a variety of mechanisms including standardized mechanisms such as the Composite Capability/Preference Profiles (CC/PP) framework, a proprietary protocol, etc.

At block 262, dynamic state information for the device may be determined. For example, dynamic state information for the device may be determined by querying system software, reading from a list, registry, database, etc. Dynamic state information may include information relating to the device's processor (e.g., current load, load available to a user), memory (e.g., utilization, amount available), communication link (e.g., bandwidth available), display (e.g., percentage used, maximum window size available, etc.), availability of the sound system (whether currently in use), etc.

At block 266, it may be determined whether the dynamic state information determined at block 262 should be transmitted to the resource manager 124 and/or the multi-device proxy server 112. For example, dynamic state information may be transmitted upon power-up of the device, periodically, upon a change in the dynamic state information, etc.

If dynamic state information need not be transmitted, the flow may proceed to block 262. If dynamic state information should be transmitted, the flow may proceed to block 270. At block 270, dynamic state information may be transmitted to the resource manager 124 and/or the multi-device proxy server 112. For example, an owned device may transmit dynamic state information to the multi-device proxy server 112, whereas a rentable device may transmit dynamic state information to a resource manager 124 associated with a neighborhood in which the rentable device is located. A proprietary mechanism may be used to express dynamic state. An example extension to the CC/PP framework that provides dynamic state information for a device may be:

```

15      <?xml version="1.0"?>
      <rdf:RDF xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#">
        <rdf:Description rdf:about="CPU">
          <Load>0.3</Load>
          <User>0.1</User>
        </rdf:Description>
        <rdf:Description rdf:about="Memory">
          <Utilization>0.4</Utilization>
20      <Available>100K</Available>
        </rdf:Description>
        <rdf:Description rdf:about="Bandwidth">
          <Available>15Kbps</Available>
        </rdf:Description>
25      <rdf:Description rdf:about="Display">
          <Available>0.6</Available>
          <Max-rectangle>200x150</Max-rectangle>
        </rdf:Description>
        <rdf:Description rdf:about="Sound">
30      <Status>ON</Status>
        </rdf:Description></rdf:RDF>

```

where "CPU ... Load" may be a number between 0 and 1 that indicates the current percentage load of the central processing unit (CPU), "CPU ... User" may be a number between 0 and 1 that indicates percentage load of the CPU available to a user, "Memory ... Utilization" may be a number between 0 and 1 that indicates the current utilization of the memory, "Memory ... Available" may be a number indicating the amount of memory available, "Bandwidth" may be a number indicating the current bandwidth available to the device on the communication link, "Display ... Available"

may be a number between 0 and 1 indicating the percentage of the display available, "Display ... Max-rectangle" may be provide the dimensions in pixels of a maximum rectangle on the display available to a user, and "Sound ... Status" may be a text value indicating whether the sound system is in use.

5 A complete list of dynamic state information need not be sent at block 270. For example, if only particular dynamic state information has changed since the last transmission of dynamic state information, only the changed dynamic state information may be transmitted.

10 After the dynamic state information is transmitted, the flow may return to block 262.

FIG. 5 is a flow diagram of an example method 300 that may be implemented by the resource manager 124. In particular, the method 300 may be implemented by the resource monitor subsystem 208. The method 300 may be implemented by a processor configured by software on a tangible medium such as a CD-ROM, a floppy
15 disk, a hard drive, a digital versatile disk (DVD), or a memory associated with the processor. Persons of ordinary skill in the art will readily appreciate that the entire method 300 or parts thereof could alternatively be executed by a device other than a processor, and/or embodied in firmware and/or dedicated hardware in a well known manner. Further, although the example method is described with reference to the
20 flow diagram in FIG. 5, persons of ordinary skill in the art will readily appreciate that many other methods may alternatively be used. For example, the order of execution of the blocks may be changed, and/or the blocks may be changed, eliminated, or combined.

25 The method 300 may be used to process device capabilities and dynamic state information received from rentable rendering devices in a particular neighborhood.

At block 304, it may be determined whether information received from the rendering device 120 corresponds to a registration of a new device. If it is a registration of a new device, the flow may proceed to block 308. At block 308, the device capabilities of the new device may be converted into one or more resources.
30 At block 312, the one or more resources determined at block 308 may be stored in a resource registry. The resource registry may indicate, for example, resources in a neighborhood that may be reserved and used/shared by one or more users. Referring

back to FIG. 3, the resource registry may be included in the shared resources database 212.

If at block 304 it is determined that the received information is not a registration of a new device, the flow may proceed to block 316. At block 316, it may be determined whether the received information corresponds to dynamic state information for a rendering device 120. If it does correspond to dynamic state information, the flow may proceed to block 320. At block 320, the dynamic state information may be used to update corresponding resource information in the resource registry.

FIG. 6 is a flow diagram of an example method 350 that may be implemented by the resource manager client 230 of an owned rendering device that a user uses to interact with the resource manager. The method 350 may be implemented by a processor configured by software on a tangible medium such as a CD-ROM, a floppy disk, a hard drive, a digital versatile disk (DVD), or a memory associated with the processor. Persons of ordinary skill in the art will readily appreciate that the entire method 350 or parts thereof could alternatively be executed by a device other than a processor, and/or embodied in firmware and/or dedicated hardware in a well known manner. Further, although the example method is described with reference to the flow diagram in FIG. 6, persons of ordinary skill in the art will readily appreciate that many other methods may alternatively be used. For example, the order of execution of the blocks may be changed, and/or the blocks may be changed, eliminated, or combined.

The method 350 may be used to reserve and/or unreserve resources of rendering devices 120 in a neighborhood using the resource manager 124.

Reserving/unreserving a resource may include reserving/unreserving the entire resource or a portion of the resource. For example, if a user has reserved 50% of a display screen, the user may unreserved half of that resource, thus giving the user 25% of the display screen.

At block 352, it may be determined whether the user desires to reserve additional resources, or whether the user desires to unreserve resources. If the user desires to reserve additional resources, the flow may proceed to block 354.

At block 354, a request for an indication of available resources in the neighborhood may be transmitted to the resource manager 124. At block 356, an indication of available resources in the neighborhood may be received from the resource manager 124. The indication of available resources may include, for example, a list of resources. Additionally, the indication of available resources may also provide indications of the devices associated with the available resources.

At block 358, the user may select one or more resources that the user wishes to reserve. For example, the user may use a touch screen, keypad, stylus, mouse, etc. to select particular resources from a list of resources. At block 360, the selected resources may be transmitted to the resource manager 124.

At block 362, it may be determined whether the resource manager 124 confirmed that the user reserved the selected resources sent at block 360. If a confirmation is received, the flow may proceed to block 364. At block 364, the rendering device 120 that implements the method 350 may send a notification to the multi-device proxy server 112 indicating that the additional resources selected at block 360 are to be added to an available resources pool associated with the user. As an alternative, and as will be described subsequently, the resource manager 124 may send such a notification.

If at block 362, it is determined that the resource manager 124 did not confirm that the user reserved the selected resources, the flow may proceed back to block 358.

If at block 352, it is determined that the user desires to unreserve resources, the flow may proceed to block 380. At block 380, a request for an indication of resources in the neighborhood reserved by the user may be transmitted to the resource manager 124. At block 382, an indication of reserved resources may be received from the resource manager 124. The indication of reserved resources may include, for example, a list of resources. Additionally, the indication of reserved resources may also provide indications of the devices associated with the reserved resources.

At block 384, the user may select one or more resources that the user wishes to unreserve. For example, the user may use a touch screen, keypad, stylus, mouse, etc. to select particular resources from a list of resources. At block 386, the selected resources may be transmitted to the resource manager 124.

At block 388, it may be determined whether the resource manager 124 confirmed that the user unreserved the selected resources sent at block 386. If a confirmation is received, the flow may proceed to block 390. At block 390, the rendering device 120 that implements the method 350 may send a notification to the multi-device proxy server 112 indicating that the resources selected at block 384 are to be removed from the available resources pool associated with the user. As an alternative, and as will be described subsequently, the resource manager 124 may send such a notification.

If at block 388, it is determined that the resource manager 124 did not confirm that the user unreserved the selected resources, the flow may proceed back to block 384.

FIG. 7 is a flow diagram of an example method 400 that may be implemented by the resource manager 124. In particular, the method 400 may be implemented by the user interaction subsystem 206. The method 400 may be implemented by a processor configured by software on a tangible medium such as a CD-ROM, a floppy disk, a hard drive, a digital versatile disk (DVD), or a memory associated with the processor. Persons of ordinary skill in the art will readily appreciate that the entire method 400 or parts thereof could alternatively be executed by a device other than a processor, and/or embodied in firmware and/or dedicated hardware in a well known manner. Further, although the example method is described with reference to the flow diagram in FIG. 7, persons of ordinary skill in the art will readily appreciate that many other methods may alternatively be used. For example, the order of execution of the blocks may be changed, and/or the blocks may be changed, eliminated, or combined.

The method 400 may be used to reserve and/or unreserve resources of rendering devices 120 in a neighborhood for a user.

At block 402, it may be determined whether the user desires to reserve additional resources, or whether the user desires to unreserve resources. If the user desires to reserve additional resources, the flow may proceed to block 404.

At block 404, it may be determined what resources are available in a particular neighborhood. This may include determining what resources are in the neighborhood, and then determining which of those resources are available. The user interaction

subsystem 206 may determine resources available in the neighborhood by querying the shared resources database 212. Additionally, if the resource is a rentable resource, the user interaction subsystem 206 may determine which of the resources are available based on access control policies, charging policies, etc. enforced by the resource allocation subsystem 204.

At block 406, an indication of available resources in the neighborhood may be sent to the user's owned rendering device 120. The indication of available resources may include, for example, a list of resources. Additionally, the indication of available resources may also provide indications of the devices associated with the available resources. At block 408, an indication of resources selected by the user may be received.

At block 410, it may be determined whether the selected resources are still available. For example, the user interaction subsystem 206 may determine whether the selected resources are still available by querying the shared resources database 212. If the selected resources are no longer available, the flow may proceed back to block 406. If the selected resources are still available, the flow may proceed to block 412.

At block 412, the selected resources may be reserved for the user. For example, the user interaction subsystem 206 may modify the shared resources database 212 to indicate that the selected resources are reserved by the user.

At block 414, a confirmation may be sent to the user, confirming that the selected resources have been reserved. At block 416, the resource manager 124 may send a notification to the multi-device proxy server 112 indicating that the additional resources reserved at block 412 are to be added to an available resources pool associated with the user. As an alternative, and as was described previously, a rendering device may send such a notification.

If at block 402, it is determined that the user desires to unreserve resources, the flow may proceed to block 430. At block 430, it may be determined what resources are currently reserved by the user in the neighborhood. For example, the resource allocation subsystem 204 may determine resources reserved by the user in the neighborhood by querying the shared resources database 212.

At block 432, an indication of reserved resources in the neighborhood may be sent to the user's owned rendering device 120. The indication of reserved resources may include, for example, a list of resources. Additionally, the indication of reserved resources may also provide indications of the devices associated with the reserved
5 resources. At block 434, an indication of resources selected by the user may be received.

At block 436, the selected resources may be unreserved. For example, the resource allocation subsystem 204 may modify the shared resources database 212 to indicate that the selected resources are unreserved by the user and may be made
10 available to other users.

At block 440, a confirmation may be sent to the user, confirming that the selected resources have been unreserved. At block 442, the resource manager 124 may send a notification to the multi-device proxy server 112 indicating that the resources unreserved at block 436 are to be removed from an available resources pool
15 associated with the user. As an alternative, and as was described previously, a rendering device may send such a notification.

Multi-Device Proxy Server

As described previously, the multi-device proxy server 112 typically acts as a
20 proxy to a user in a multi-device communication environment. As just one specific example, the multi-device proxy server 112 may be implemented as an extension to a wireless application protocol (WAP) proxy server. The multi-device proxy server 112 may comprise one or more computers such as a personal computers, workstations, servers, mainframes, etc.

25 FIG. 8 is a block diagram of an example multi-device proxy server 112, and illustrates an example of its interaction with other systems, including rendering devices 120. FIG. 8 will be described with reference to FIG. 1.

The multi-device proxy server 112 may comprise a delivery control subsystem 504. The deliver control subsystem may communicate with other systems such as
30 rendering devices 120 and the resource manager 124 via a communication agent 508. The communication agent 508 may be the same as, or similar to, the communication agent 216 of the resource manager 124.

The multi-device proxy server 112 may also comprise a user preferences database 512, a user resources database 516, and a session state database 520. The delivery control subsystem 504 may store data to, and read data from, each of these databases.

5 The user preferences database 512 may include information relating to preferences of a user for rendering content on multiple devices. The user preferences database 512 may include information relating to, for example, preferred devices for rendering certain types of content (e.g., preferable that video be rendered on computer with a full size display rather than a cell phone), preferred splitting of content (e.g.,
10 prefer that two different streams of audio be routed to different devices rather than same device), preferred network for delivery (e.g., prefer that data delivered to cell phone be delivered via cell phone network rather than wireless LAN network).

 The user resources database 516 may include information relating to the resources available to a particular user. For example, information indicating that a
15 particular user has an owned PDA and has reserved a display of desktop computer may be stored in the user resources database 516.

 The session state database 520 may include information relating to the state of a particular user's session. This information may include, for example, the types of content being, or to be, delivered (e.g., text, video, audio). Content type information
20 may also include information relating to the format of the content (e.g., a video data rate, stereo audio vs. mono audio, etc.). The session state database 520 may also include information relating to the transport layer such as the addresses of the remote devices used, the port numbers used on those devices, etc.

 The multi-device proxy server 112 may additionally comprise an application
25 downloader 524 that may download to rendering devices applications, for example, to facilitate the rendering of content. The application downloader 524 may be controlled by the delivery control subsystem 504.

 The multi-device proxy server 112 may further comprise a transcoder interface 528 that may provide the deliver control subsystem 504 with an interface to a local
30 transcoder 532 (optional) and one or more external transcoders. The transcoder interface 528 may communicate with external transcoders via a variety of protocols including standardized mechanisms for invoking remote procedures such as the

Simple Object Access Protocol (SOAP), the Common Gateway Interface (CGI), Sun Microsystems' Remote Method Invocation (RMI) technology, the Internet Content Adaptation Protocol (ICAP), etc..

5 The content transfer protocol proxy 536 may include one or more different proxies corresponding to different protocols. For example, the content transfer protocol proxy 536 may include an HTTP proxy, a real-time protocol (RTP) proxy, a real time streaming protocol (RTSP) proxy, etc. The content transfer protocol proxy 536 may also include a parser for parsing content according to content types. The content types may include, for example, multipurpose internet mail extensions
10 (MIME) types. The content types may also include types defined, for example, by the author of the content using, for example, an annotation file. For example, content may be authored so that certain of the content can be viewed by certain persons, certain devices, etc.

In general, the content transfer protocol proxy 536 receives content destined
15 for a user, parses the content into content types, and then delivers the content to one or more rendering devices 120 under the control of the delivery control subsystem 504. For example, the content transfer protocol proxy 536 may deliver text and graphics to rendering device 120a, and deliver audio to rendering device 120b. Additionally, the content transfer protocol proxy 536 may send content to, and receive content from, a
20 transcoder (e.g., local transcoder 532 or an external transcoder) under control of the transcoder interface 528. Further, the content transfer protocol proxy 536 may store information regarding the state of a session in the session state database 520. For example, the content transfer protocol proxy 536 may store information relating to the different content types in a content session requested by a user.

25 The multi-device proxy server 112 may communicate with rendering devices, such as rendering devices 120a, 120b, and 120c (not shown in FIG. 8). Rendering devices 120a and 120c may be, for example, owned devices, and rendering device 120b may be, for example, a rentable device. The rendering devices 120a, 120b, and 120c may each include respective multi-device proxy clients 550. The multi-device
30 proxy clients 550 may be the same for each rendering device, or they may differ, for example, depending on whether its corresponding rendering device is an owned device or a rentable device. The multi-device proxy clients 550 on an owned device

may interact with the multi-device proxy server 112 to, for example, notify the multi-device proxy server 112 of resources available to a user, notify the multi-device proxy server 112 of user preferences regarding delivering content to particular devices available to the user, coordinate downloading of applications to facilitate rendering of content, etc. Additionally, the multi-device proxy client 550 may launch applications on the rendering devices 120, such as applications 562, to facilitate the rendering of content. For example, the multi-device proxy client 550 may launch a video player application on the rendering device 120 upon being becoming aware that video content will be, or has begun being, delivered to the rendering device 120.

FIG. 9 is a flow diagram of an example method 570 that may be implemented by the multi-device proxy server 112. The method 570 may be implemented by a processor configured by software on a tangible medium such as a CD-ROM, a floppy disk, a hard drive, a digital versatile disk (DVD), or a memory associated with the processor. Persons of ordinary skill in the art will readily appreciate that the entire method 570 or parts thereof could alternatively be executed by a device other than a processor, and/or embodied in firmware and/or dedicated hardware in a well known manner. Further, although the example method is described with reference to the flow diagram in FIG. 9, persons of ordinary skill in the art will readily appreciate that many other methods may alternatively be used. For example, the order of execution of the blocks may be changed, and/or the blocks may be changed, eliminated, or combined.

At block 572, a request for content may be received from a user. The request may be destined for a content server. At block 574, an entry in a list may be created, where the entry corresponds to the user that sent the request. The entry may include information so that when the multi-device proxy server 112 receives the requested content, the user corresponding to the content can be determined. Examples of such information include a session identifier that may uniquely identify a user's request, the port number on which the request was received, a user identifier that may uniquely identify the user, etc.

At block 576, the request may be reformatted so that the content server sends the requested content to the multi-device proxy server 112. For example, the requestor's identity may be changed from the user's device to that of the multi-device

proxy server 112. Then, at block 578, the reformatted request may be sent to the content server.

FIG. 10 is a flow diagram on another example method 584 that may be implemented by the multi-device proxy server 112. The method 584 will be described with reference to FIG. 9. The method 584 may be implemented by a processor configured by software on a tangible medium such as a CD-ROM, a floppy disk, a hard drive, a DVD, or a memory associated with the processor. Persons of ordinary skill in the art will readily appreciate that the entire method 584 or parts thereof could alternatively be executed by a device other than a processor, and/or embodied in firmware and/or dedicated hardware in a well known manner. Further, although the example method is described with reference to the flow diagram in FIG. 10, persons of ordinary skill in the art will readily appreciate that many other methods may alternatively be used. For example, the order of execution of the blocks may be changed, and/or the blocks may be changed, eliminated, or combined.

At block 586, content that was requested by a user may be received from a content server. For example, the content may be in response to the reformatted request sent to the content server at block 578 of FIG. 9. At block 588, the user that requested the content may be determined. For example, the user may be determined by examining the list in which an entry for the user was created at block 574 of FIG. 9.

At block 590, the requested content may be processed to determine the different content types (e.g., text, graphics, video, audio, etc.) included in the content.

At block 592, a mapping of the content types determined at block 590 to resources available to the user may be determined. The resources available may be on multiple devices. Additionally, the multiple devices may include devices in a neighborhood as well as remote devices. At block 594, the content types may be sent to the multiple devices according to the mapping determined at block 592.

FIG. 11 is a block diagram of an example delivery control subsystem 504. FIG. 11 will be described with reference to FIG. 8. The delivery control subsystem may include a user resources manager 604, a mapping generator 610, and a session orchestrator 620. In general, the user resources manager 604 may receive information from an owned rendering device 120 or a resource manager 124, via the

communication agent 508, regarding resources that are available to the user. This information may be used to update the user resources database 516. Additionally, the user resources manager 604 may receive information regarding user preferences from an owned rendering device 120, and store this information in the user preferences
5 database 512.

Generally, the mapping generator 610 may receive user preferences information from the user preferences database 512, information relating to resources available to the user from the user resources database 516, and information relating to the content requested by the user from the session state database 520. Based on this
10 information, the mapping generator 610 may generate a mapping of content types to devices available to the user. This device mapping may be provided to the session orchestrator 620.

In general, the session orchestrator 620 may receive the device mapping from the mapping generator 610, and content type information from the session state
15 database 520. Based on this information, the session orchestrator 620 may generate control information for controlling the application downloader 524, the transcoder interface 528, and the content transfer protocol proxy 536.

User Resources Manager

20 FIG. 12 is a flow diagram of an example method that may be implemented by the multi-device proxy server 112. In particular, the method 650 may be implemented by the user resources manager 604. The method 650 may be implemented by a processor configured by software on a tangible medium such as a CD-ROM, a floppy disk, a hard drive, a digital versatile disk (DVD), or a memory associated with the
25 processor. Persons of ordinary skill in the art will readily appreciate that the entire method 650 or parts thereof could alternatively be executed by a device other than a processor, and/or embodied in firmware and/or dedicated hardware in a well known manner. Further, although the example method is described with reference to the flow diagram in FIG. 12, persons of ordinary skill in the art will readily appreciate
30 that many other methods may alternatively be used. For example, the order of execution of the blocks may be changed, and/or the blocks may be changed, eliminated, or combined.

The method 650 may be implemented upon receiving information from an owned rendering device. For instance, the multi-device proxy server 112 may receive information from an owned rendering device indicating resources and capabilities of the device as described with reference to FIG. 4. The method 650 is similar to the
5 method 300 described with reference to FIG. 5.

At block 652, the identity of the user may be determined. Information related to the user's identity (e.g., name, user identification, login name, etc.) may be included with the received information from the owned rendering device. Additionally, the user resources manager 604 may send a message to the owned device 120 prompting
10 the user to submit information related to the user's identity (e.g., name, user identification, login name, password, etc.).

At block 654, it may be determined whether information received from the owned rendering device 120 corresponds to a registration of a new device. If it is a registration of a new device, the flow may proceed to block 658. At block 658, the
15 device capabilities of the new device may be converted into one or more resources. At block 662, the one or more resources determined at block 658 may be stored in a resource registry associated with the user. The resource registry may indicate, for example, resources available to the user. Referring back to FIG. 8, the resource registry may be included in the user resources database 516.

20 If at block 654 it is determined that the received information is not a registration of a new device, the flow may proceed to block 668. At block 668, it may be determined whether the received information corresponds to dynamic state information for an owned rendering device 120. If it does correspond to dynamic state information, the flow may proceed to block 670. At block 670, the dynamic
25 state information may be used to update corresponding resource information in the resource registry associated with the user.

FIG. 13 is a flow diagram of an example method that may be implemented by the multi-device proxy server 112. In particular, the method 700 may be implemented by the user resources manager 604. The method 700 may be implemented by a
30 processor configured by software on a tangible medium such as a CD-ROM, a floppy disk, a hard drive, a digital versatile disk (DVD), or a memory associated with the processor. Persons of ordinary skill in the art will readily appreciate that the entire

method 700 or parts thereof could alternatively be executed by a device other than a processor, and/or embodied in firmware and/or dedicated hardware in a well known manner. Further, although the example method is described with reference to the flow diagram in FIG. 13, persons of ordinary skill in the art will readily appreciate that many other methods may alternatively be used. For example, the order of execution of the blocks may be changed, and/or the blocks may be changed, eliminated, or combined.

The method 700 may be implemented upon receiving information relating to the addition or removal of devices available to the user. For instance, the multi-device proxy server 112 may receive information from an owned rendering device or the resource manager 124 indicating that additional resources are available or that resources are no longer available as described with reference to FIGs. 6 and 7.

At block 704, the identity of the user may be determined. Block 704 may be implemented in a manner similar to block 652 of FIG. 12. At block 708, it may be determined whether additional resources are available to the user. If additional resources are available to the user, the flow may proceed to block 712. At block 712, the additional resources may be added to the resource registry associated with the user.

If at block 708, it is determined that additional resources are not available, it may be determined at block 716 if resources that were available to the user are no longer available. If resources are no longer available, the flow may proceed to block 720. At block 720, the resources that are no longer available to the user may be removed from the resource registry associated with the user.

User Preferences

Referring again to FIGs. 8 and 11, information from the user preferences database 512 may be used by the mapping generator 610 to generate a mapping of content types to devices. The user preference information may be assigned to a user based on one or more default configurations, or may be obtained from the user, for example, when the user registers for use of the multi-device proxy server 112, upon requesting content via the multi-device proxy server 112, etc. The user preference information may be modified, for example, upon the user's request, upon receiving a

type of content for which the user has not yet provided preference information, upon using a particular device or device type for which the user has not yet provided preference information, etc.

5 The user preference information stored in the database 512 may be represented in a variety of formats. For example, the user preference information may be represented in a utility function format. Examples of user preference information in a utility function format will now be described with reference to FIGs. 14, 15, and 16. For ease of explanation, particular examples will be described, where the particular examples are illustrated in tabular format.

10 FIG. 14 is a table illustrating one example format for representing user preferences related to preferred devices for rendering particular types of content. In particular, FIG. 14 illustrates an example table 800 for representing user preferences related to delivering types of content to particular devices.

15 Rows of the table 800 may correspond to particular devices or device types, and columns of the table 800 may correspond to content types (e.g., text, graphics, video, audio, etc.). The list of devices may include devices currently available to the user, a list of devices that the user has used in the past, a list of devices specified by the user, a default list of devices, etc.

20 A utility of a particular type of content delivered to a particular device may be represented as a number in a range. In the example of FIG. 14, the range may be between 0 and 10, inclusive, where a value of 0 indicates lowest utility, and a value of 10 indicates the highest utility. The example table 800 indicates that delivering text content to a PDA provides high utility, whereas delivering graphics to the PDA provides a lower utility, and delivering video or audio provides the lowest utility. For
25 example, the display of the PDA may be adequate for displaying text, but video on the PDA may be of poor quality. Similarly, the PDA may not include an audio system.

FIG. 15 is a table illustrating one example format for representing user preferences related to preferred communication links for delivering content to particular types of devices. In particular, FIG. 15 illustrates an example table 840 for
30 representing user preferences related to preferred communication links for delivering content to particular devices.

Rows of the table 840 may correspond to particular devices or device types, and columns of the table 840 may correspond to particular communication links or communication link types (e.g., a wireless LAN (WLAN), a general packet radio service (GPRS) network, etc.). The list of devices may include devices currently
5 available to the user, a list of devices that the user has used in the past, a list of devices specified by the user, a default list of devices, etc. The list of links may include links currently being utilized by the user, a list of links that the user has used in the past, a list of links specified by the user, a default list of links, etc.

A utility of using a particular communication link to deliver content to a
10 particular device may be represented as a number in a range. In the example of FIG. 15, the range may be between 0 and 10, inclusive, where a value of 0 indicates lowest utility, and a value of 10 indicates the highest utility. The example table 840 indicates that delivering content to a cellular phone via a GPRS link provides higher utility than via a WLAN link. Additionally, the table 840 indicates that delivering content to a
15 PDA or laptop computer via a WLAN link provides higher utility than via a GPRS link.

FIG. 16 is a table illustrating one example format for representing user preferences related to whether content should be delivered to the same device or different devices. In particular, FIG. 16 illustrates an example table 860 for
20 representing user preferences related to whether components of one or more content sessions should be delivered to the same or different devices. For example, it may be easier to synchronize video and audio corresponding to a movie via a single device rather than trying to synchronize video and audio split across two devices, or having to view the audio and video on two devices with a slight time delay between the video
25 and audio.

The rows and columns of the table 860 may correspond to particular components of one or more content sessions. In the example of FIG. 16, the suffix indicates the session to which a component belongs (i.e., "Caption 1," "Audio 1," and "Video 1" are components of a first content session, and "Audio 2" is a component of
30 a second content session, different from the first content session). The list of components may include components of sessions that the user has requested to be delivered, a list of possible components, a default list of components, etc.

A utility of a component being delivered via the same device as another component may be represented as a number in a range. In the example of FIG. 16, the range may be between 0 and 100, inclusive, where a value of 0 indicates lowest utility, and a value of 100 indicates the highest utility. Additionally, a large negative number (e.g., -100) may be used to represent a strong preference for the components not to be delivered via the same device.

The example table 860 indicates a high utility where the audio and video components of a first session (i.e., Audio 1 and Video 1) are delivered to the same device. Additionally, the table 860 indicates a lower utility where the caption and the video components of the first session (i.e., Caption 1 and Video 1) are delivered to the same device. Similarly, the table 860 indicates a low utility where the caption and the audio components of the first session (i.e., Caption 1 and Audio 1) are delivered to the same device. Further, the table 860 indicates strong preference that an audio component of a second session (i.e., Audio 2) not be delivered to the same device as the device or devices to which the components of the first session are being delivered.

Although the above examples of utility function representations are described in the format of tables, those of ordinary skill in the art will appreciate that utility function representations may be provided in a variety of formats. Additionally, the above examples describe particular numerical values and ranges, those of ordinary skill in the art will appreciate that a variety of values and ranges may be utilized in different implementations.

Further, although user preferences are described above in utility function representations, those of ordinary skill in the art will appreciate that user preferences may be represented in other ways as well. Additionally, other types of user preferences may be employed as well. For example, preferences may be based on device capabilities. For instance, a user may prefer rendering audio on a device having stereo audio capability rather than a device having only mono audio capability. Additionally, preferences may be based on the nature of the information included in the content. For example, it may be preferable to display large amounts of text on a desktop computer rather than a PDA. Additionally, it may be preferable to display certain information that the user does not want others to view on the user's PDA,

rather than on a shared display, even if the shared display could render the content much more adequately than the PDA.

Mapping Generator

5 Referring again to FIGs. 8 and 11, the mapping generator 610 generally may receive user preferences information from the user preferences database 512, information relating to resources available to the user from the user resources database 516, and information relating to the content requested by the user from the session state database 520. Based on this information, the mapping generator 610
10 may generate a mapping of content types to devices available to the user. The mapping may also include information for delivering content to the devices via particular communication links (e.g., a WLAN link vs. a GPRS link).

FIG. 17 is a flow diagram of an example method that may be implemented by the delivery control subsystem 504. In particular, the method 900 may be
15 implemented by the mapping generator 610. The method 900 will be described with reference to FIGs. 8 and 11. The method 900 may be implemented by a processor configured by software on a tangible medium such as a CD-ROM, a floppy disk, a hard drive, a digital versatile disk (DVD), or a memory associated with the processor. Persons of ordinary skill in the art will readily appreciate that the entire method 900
20 or parts thereof could alternatively be executed by a device other than a processor, and/or embodied in firmware and/or dedicated hardware in a well known manner. Further, although the example method is described with reference to the flow diagram in FIG. 17, persons of ordinary skill in the art will readily appreciate that many other methods may alternatively be used. For example, the order of execution of the blocks
25 may be changed, and/or the blocks may be changed, eliminated, or combined.

At block 904, content type information for a particular session may be received. For example, the content type information may be received from session state database 520. The content type information may indicate, for example, that a particular session may include text, video, and audio.

30 At block 908, an indication or indications of resources available to the user associated with the session may be received. For instance, a list of available resources may be received from the user resources database 516. As an example, the list of

resources may indicate that the available resources includes an audio system and a display system on a cell phone, and a display system on a laptop computer. Additionally, the list of resources may indicate WLAN and GPRS links available on both the cell phone and the laptop computer.

5 At block 912, user preferences information may be received, for example, from the user preferences database 512. As described with reference to FIGs. 14-16, the user preferences information may be represented as one or more utility function representations.

 At block 916, a best utility value may be initialized. For example, the best
10 utility value may be initialized to reflect a very low utility. At block 920, multiple possible mappings of content types in the session to devices and communication links available to the user may be generated. For example, all of the possible mappings may be generated. The content types in the session may be indicated by the information received at block 904. The devices and communication links available to
15 the user may be determined by the information received at block 908. The multiple possible mappings may be generated, for example, as a list of possible mappings.

 At block 924, one of the possible mappings generated at block 920 may be selected. For example, if a list of possible mappings is generated at block 920, the first mapping in the list may be selected. As alternatives, one of the mappings may be
20 randomly or pseudo-randomly selected, the last mapping in the list may be selected, etc.

 At block 928, a utility value of the selected mapping may be generated. The utility value may be generated, for example, according to one or more utility function representations received at block 912. At block 932, the utility value generated at
25 block 928 is compared with a best utility value. If at block 932, the utility value generated at block 928 indicates a higher utility than the current best utility value, the flow may proceed to block 936. At block 936, the utility value generated at block 928 may be stored as the new best utility value. Additionally, an indication of the mapping corresponding to the new best utility value may be stored. If at block 932
30 the utility value generated at block 928 does not indicate a higher utility than the current best utility value, the flow may proceed to block 940.

At block 940, it may be determined whether all of the mappings generated at block 920 have been tried. If all of the mappings have not been tried, then at block 944 a new mapping may be selected, and the flow may proceed back to block 928. If at block 940 it is determined that all of the mappings have been tried, the flow may proceed to block 948. At block 948, the best mapping may be provided to the session orchestrator 620.

FIG. 18 is a flow diagram of another example method that may be implemented by the delivery control subsystem 504. In particular, the method 960 may be implemented by the mapping generator 610. The method 960 will be described with reference to FIGs. 8 and 11. The method 960 may be implemented by a processor configured by software on a tangible medium such as a CD-ROM, a floppy disk, a hard drive, a DVD, or a memory associated with the processor. Persons of ordinary skill in the art will readily appreciate that the entire method 960 or parts thereof could alternatively be executed by a device other than a processor, and/or embodied in firmware and/or dedicated hardware in a well known manner. Further, although the example method is described with reference to the flow diagram in FIG. 18, persons of ordinary skill in the art will readily appreciate that many other methods may alternatively be used. For example, the order of execution of the blocks may be changed, and/or the blocks may be changed, eliminated, or combined.

The method 960 includes many of the same blocks as the method 900 of FIG. 17. For examples, blocks 904, 908, 912, 916, 928, 932, 936, and 948 may be the same.

At block 964, a possible mapping may be randomly or pseudo-randomly determined. At block 968, it may be determined if all mappings have been tried. As one example, it may be determined whether all of the possible mappings have been tried. As another example, it may be determined whether a predetermined number of mappings have been tried, or whether block 964 was invoked a predetermined number of times.

The mapping generator 610 (FIG. 11) may be invoked one or more times during a session. For example, the mapping generator 610 may be invoked in response to initially receiving content that was requested by the user. Additionally, the mapping generator 610 may be invoked during a session when additional

resources become available to the user, or when resources allocated to the user become unavailable. Additionally, the mapping generator 610 may be invoked when the user leaves a neighborhood or enters a new neighborhood. Entering or leaving a neighborhood may be detected manually, for example, by the user notifying the multi-
5 device proxy server 112 of the neighborhood change. Additionally, entering or leaving a neighborhood may be detected automatically, for example, via user device interaction with a resource manager 124, via a location detection system such as a wide area positioning system (e.g., a global positioning system (GPS), a Loran-C system, etc.), a local area positioning system (e.g., an in-building location system),
10 etc.

Session Orchestrator

Referring again to FIGs. 8 and 11, the session orchestrator 620 generally may receive user mapping information from the mapping generator 610 and information
15 relating to the types of content in the current session. Then, the session orchestrator 620 may generate control signals for controlling the content transfer protocol proxy so that the various types of content are delivered to the appropriate rendering devices, and via the appropriate communication links, according to the mapping information. As will be described in more detail below, the session orchestrator 620 may also
20 generate control signals for controlling the application downloader 524 and transcoders (local or external) via the transcoder interface 528.

FIG. 19 is a flow diagram of an example method that may be implemented by the delivery control subsystem 504. In particular, the method 1000 may be implemented by the session orchestrator 620. The method 1000 will be described
25 with reference to FIGs. 8 and 11. The method 1000 may be implemented by a processor configured by software on a tangible medium such as a CD-ROM, a floppy disk, a hard drive, a digital versatile disk (DVD), or a memory associated with the processor. Persons of ordinary skill in the art will readily appreciate that the entire method 1000 or parts thereof could alternatively be executed by a device other than a
30 processor, and/or embodied in firmware and/or dedicated hardware in a well known manner. Further, although the example method is described with reference to the flow diagram in FIG. 19, persons of ordinary skill in the art will readily appreciate

that many other methods may alternatively be used. For example, the order of execution of the blocks may be changed, and/or the blocks may be changed, eliminated, or combined.

At block 1004, the mapping, user resources, and content type information may be received from the mapping generator 610, the user resources database 516, and the session state database 520, respectively. At block 1008, a first content type from the content types of the session may be selected. At block 1012, it may be determined whether the mapping indicates there is a corresponding resource/device for the content type. If there is no corresponding resource/device, the flow may proceed to block 1016. At block 1016, the content of the content type for which a resource/device is not available may be stored. For example, the content may be temporarily stored until a resource for rendering the content type becomes available. Then, the flow may proceed to block 1020, where it may be determined if other content types of the session still remain. If all of the content types in the session have been handled, the flow may end. If other content types still remain, the flow may proceed to block 1024. At block 1024, a next content type in the session may be selected, and the flow may proceed back to block 1012.

If at block 1012 it is determined that there is a mapped device that corresponds to the content type, the flow may proceed to block 1028. At block 1028, it may be determined whether the mapped resource to which the content type corresponds can render the content type in its current format. For example, audio content may be in MP3 format, whereas the audio resource of the device selected by the mapping may be able to handle audio only in a different format. If the mapped resource can render the content type in its current format, the flow may proceed to block 1032. At block 1032, the content transfer protocol proxy 536 may be controlled to deliver the content type to the device indicated by the mapping information received at block 1004.

If at block 1028, it is determined that the mapped resource cannot render the content in its current format, the flow may proceed to block 1036. At block 1036, it may be determined whether the mapped device can receive a software download to permit it to render the content type in its current format. For example, the device may be incapable of receiving the download. Also, a user of a device may be unwilling to permit such downloads. Additionally, a software download for permitting the device

to render the content type in its current format may not be available. If the mapped device can receive the software download, the flow may pass to block 1040.

At block 1040, the application downloader 524 may be controlled to download the application software to the mapped device. Additionally, the user resources
5 database 516 may be updated to reflect the resources of the device have been changed via the downloaded application. Then, the flow may proceed to block 1032.

At block 1036, if it determined that the device cannot receive the application download, the flow may proceed to block 1044. At block 1044, the transcoder
10 interface 528 may be controlled to convert, using a transcoder, the current format of the content type into a new format that can be rendered by the mapped resource. The transcoder interface 528 may cause the content transfer protocol proxy to send the content type to a transcoder (local or external), and the content transfer protocol proxy may receive the content type in a new format from the transcoder. Then, the flow may proceed to block 1032.

15

User Interface Mechanisms for Multi-Device Rendering

In some examples, the multi-device proxy server 112 may modify the content with user interface mechanisms that assist a user with interpreting and navigating content being rendered on multiple devices. FIG. 20 is a block diagram of an
20 example subsystem 1100 for providing such user interface mechanisms. The particular user interface mechanisms that the example subsystem 1100 provides are coined "wormholes." Further details of systems for implementing such user interface mechanisms are provided in U.S. Patent Application No. 10/334,848, entitled
25 "Method and Apparatus for Linking Multimedia Content Rendered Via Multiple Devices," (Attorney Docket No. CS20693RL). In general, a wormhole may be an icon, or other indicator, within the content rendered on the user's device that indicates to the user that portions of the content have been removed and are being, or can be rendered, on a different device. The wormholes may permit the user to understand the content context in which this removed content originated. Additionally, the user may
30 use the wormhole mechanism to cause the removed content to be rendered, or to be indicated to the user. For example, a user viewing a web page with removed pictures

on a PDA could activate a wormhole on the PDA, causing the removed picture to flash on a shared display.

For ease of explanation, the wormhole subsystem 1100 will be described in the context of web pages. Additionally, the wormhole subsystem 1100 will be
5 described with reference to FIG. 8.

The wormhole subsystem 1100 may be, in whole or in part, a component of the multi-device proxy server 112, or may be implemented as a separate system. In general, the wormhole subsystem 1100 may receive content that has been transcoded for rendering on multiple devices, and then modifies the transcoded content to include
10 user interface mechanisms (wormholes). Additionally, the wormhole subsystem 1100 may receive indications that a wormhole was activated, and may instruct a wormhole client on a rendering device 120 to respond to the activation. For example, a user may be viewing a web page rendered such that text and wormholes appear on a PDA and images appear on a desk top computer. When a user activates a wormhole via the
15 PDA, the wormhole subsystem 1110 may instruct the desk top computer displaying an image associated with the wormhole to cause the image to blink.

The wormhole subsystem may include a wormhole inserter 1104, a wormhole database 1108, and a wormhole manager 1112. The wormhole inserter 1104 may be a component of the local transcoder 532. The wormhole inserter 1104 may receive a
20 first subset of the requested content, and may generate a modified first subset of the content that includes wormholes.

FIG. 21 is a flow diagram of an example method that may be implemented by the wormhole inserter 1104. The method 1140 may be implemented by a processor configured by software on a tangible medium such as a CD-ROM, a floppy disk, a
25 hard drive, a digital versatile disk (DVD), or a memory associated with the processor. Persons of ordinary skill in the art will readily appreciate that the entire method 1140 or parts thereof could alternatively be executed by a device other than a processor, and/or embodied in firmware and/or dedicated hardware in a well known manner. Further, although the example method is described with reference to the flow diagram
30 in FIG. 21, persons of ordinary skill in the art will readily appreciate that many other methods may alternatively be used. For example, the order of execution of the blocks may be changed, and/or the blocks may be changed, eliminated, or combined.

At block 1144, a first subset of the requested content may be received. For example, a web page requested by the user may have been partitioned into content types (e.g., text, graphics, etc.). The first subset of content may correspond to a first transcoded web page that includes text of the requested web page, but not many of the graphics. A second subset of the content may correspond to a second transcoded web page that includes the graphics missing from the first transcoded web page.

At block 1148, some or all of the originally requested content that is missing from the first subset may be determined. For example, the graphics missing from the first web page may be determined.

At block 1152, wormholes corresponding to the missing content may be inserted in the first subset of content. The wormholes may include an indication of the content to which they correspond. Additionally, the wormholes may include an indication of the rendering device on which the missing content is being rendered, or should be rendered. These indications may be included in the wormholes themselves, or may be stored in the wormhole database 1108.

For example, wormholes corresponding to the graphics missing from the first web page and included in the second web page may be inserted in the first web page. The wormholes may include indications of the content in the second web page to which they correspond, and the device on which the second web page is being rendered.

A wormhole may include, for example, a button, a link, an icon, etc. The wormhole may provide an indication of the content or the type of content that was removed. For example, if a picture was removed, the wormhole may include an icon that indicates a picture was removed, text that describes the picture, etc.

The wormhole inserter 1104 may also store information relating to the inserted wormholes in the wormhole database 1108. For example, the wormhole information may include an indication of the missing content to which it relates, on which device the missing content is being rendered, or should be rendered, etc. As an alternative, some or all of this information may be encoded in the wormhole itself.

When a user activates a wormhole on a rendering device 120 (e.g., by selecting, clicking on, etc. the wormhole), the rendering device 120 may transmit an indication of the wormhole selection to the multi-device proxy server 112. The

indication may include information that would enable the wormhole subsystem 1100 to determine the selected wormhole in the wormhole database 1108. As an alternative, the indication may include information to enable the wormhole subsystem 1100 to determine the missing content to which the selected wormhole relates, and the rendering device 120 on which the missing content is being rendered, or should be rendered.

The rendering device 120 may transmit the indication, for example, to the content transfer protocol proxy 536 or the communication agent 508 of the multi-device proxy server 112.

FIG. 22 is a flow diagram of an example method that may be implemented by the wormhole manager 1112. The method 1170 may be implemented by a processor configured by software on a tangible medium such as a CD-ROM, a floppy disk, a hard drive, a digital versatile disk (DVD), or a memory associated with the processor. Persons of ordinary skill in the art will readily appreciate that the entire method 1170 or parts thereof could alternatively be executed by a device other than a processor, and/or embodied in firmware and/or dedicated hardware in a well known manner. Further, although the example method is described with reference to the flow diagram in FIG. 22, persons of ordinary skill in the art will readily appreciate that many other methods may alternatively be used. For example, the order of execution of the blocks may be changed, and/or the blocks may be changed, eliminated, or combined.

At block 1174, an indication of a wormhole selection may be received. The indication may be received, for example, via the content transfer protocol proxy 536 or the communication agent 508 of the multi-device proxy server 112.

At block 1178, the content to which the wormhole corresponds may be determined. Additionally, the device that is rendering, or should render, the corresponding content may be determined. For example, the wormhole manager 1112 may determine, via the wormhole itself or via the wormhole database 1108, the missing content to which the wormhole corresponds and the rendering device 120 on which the missing content is to be rendered, or should be rendered.

At block 1182, the rendering device that is rendering, or should be rendering, the corresponding content, may be instructed to indicate the content to which the wormhole corresponds. For example, if the content is being rendered, the rendering

device 120 may be instructed to cause the content to blink, may generate an outline around the content, may cause a caption to appear below or over the content, may cause the caption to blink, etc. If the content is not yet rendered on the rendering device, the wormhole manager 1112 may cause the content to be sent to an
5 appropriate rendering device and rendered.

While the invention is susceptible to various modifications and alternative constructions, certain illustrative embodiments thereof have been shown in the drawings and are described in detail herein. It should be understood, however, that there is no intention to limit the disclosure to the specific forms disclosed, but on the
10 contrary, the intention is to cover all modifications, alternative constructions and equivalents falling within the spirit and scope of the disclosure as defined by the appended claims.

What is Claimed is:

1. A method for delivering content to be rendered by multiple devices, the
5 method comprising:
 - receiving indications of resources available to a user, the resources including rendering resources, the rendering resources provided by a plurality of devices, each device of the plurality of devices being coupled to a network, and wherein at least one device of the plurality of devices provides at least one rendering resource available for
10 use of the user and provides at least one rendering resource available for simultaneous use of another user;
 - receiving content requested by the user, the content including a plurality of content types;
 - determining a mapping of content types to the plurality of devices, wherein the
15 mapping is based on rendering resources provided by each of the plurality of devices; and
 - delivering content types of the content requested by the user to the plurality of devices according to the mapping, the content types delivered to the plurality of devices via the network.
- 20 2. A method as defined in claim 1, wherein determining the mapping is further based on preferences of the user.
3. A method as defined in claim 2, wherein the preferences include at least one of utility functions, preferences relating to whether to deliver content of a same session to a single device, preferences relating to whether to deliver content of
25 different sessions to a single device, and preferences relating to communication links for delivering content types to the plurality of devices.
4. A method as defined in claim 1, wherein one of the content types is in a first format and the method further comprises:

determining if the device to which the one content type is mapped is capable of processing the first format; and

downloading an application to the device to which the one content type is mapped, wherein the application is capable of processing the first format.

- 5 5. A tangible medium storing machine readable instructions comprising:
first set of prestored instructions to receive indications of rendering resources
available to a user, the rendering resources provided by a plurality of rendering
devices, each of the plurality of rendering devices being coupled to a network;
wherein at least one of the plurality of rendering devices provides at least one
10 rendering resource available for use of the user, and provides at least one rendering
resource available for simultaneous use of another user;
second set of prestored instructions to receive content requested by the user,
the content including a plurality of content types;
third set of prestored instructions to determine a mapping of content types to
15 the plurality of devices, wherein the mapping is based on rendering resources
provided by each of the plurality of devices; and
fourth set of prestored instructions to deliver content types of the content
requested by the user to the plurality of devices according to the mapping, the content
types delivered to the plurality of devices via the network.

- 20 6. A server for delivering content to be rendered by multiple devices, the server
comprising:
a computer operatively coupled to a content server via a first network, and
operatively coupled to a plurality of rendering devices via a second network, wherein
at least one of the plurality of rendering devices provides at least one rendering
25 resource available for use of a user, and provides at least one rendering resource
available for simultaneous use of another user, the computer comprising:
a memory;
a processor coupled to the memory, the processor configured to
receive indications of rendering resources available to a user,
30 the rendering resources provided by the plurality of rendering devices;

receive content requested by the user, the content including a plurality of content types, the content received from the content server via the first network,

5 determine a mapping of content types to the plurality of devices, wherein the mapping is based on rendering resources provided by each of the plurality of devices, and

 deliver content types of the content requested by the user to the plurality of devices according to the mapping, the content types delivered to the plurality of devices via the second network.

10 7. A method for allocating resources on at least one device for rendering content, the method comprising:

 providing a first device, via a network, with an indication of first and second resources available to a first user, wherein the first and second resources are for rendering content, and wherein the first and second resources are provided by a

15 second device coupled to the network;

 receiving from the first device a request to reserve the first resource;

 reserving the first resource for the first user;

 providing a third device, via a network, with an indication of the second resource available to a second user;

20 receiving from the third device a request to reserve the second resource; and
 reserving the second resource for the second user.

8. A method as defined in claim 7, further comprising:

 providing the first device, via the network, with an indication of a third resource available to the first user, wherein the third resource is for rendering content,

25 and wherein the third resource is provided by a fourth device;

 receiving from the first device a request to reserve the third resource; and

 reserving the third resource for the first user.

9. A method as defined in claim 7, further comprising:

receiving capability information from the second device, the capability information including information relating to capabilities for rendering content on the second device; and

- 5 determining a set of resources available on the second device based on the received capability information, the set of resources including the first and second resources for rendering content.

10. A tangible medium storing machine readable instructions comprising:
- first set of prestored instructions to provide a first device, via a network, with an indication of first and second resources available to a first user, wherein the first
- 10 and second resources are for rendering content, and wherein the first and second resources are provided by a second device coupled to the network;
- second set of prestored instructions to receive from the first device a request to reserve the first resource;
- third set of prestored instructions to reserve the first resource for the first user;
- 15 fourth set of prestored instructions to provide a third device, via a network, with an indication of the second resource available to a second user;
- fifth set of prestored instructions to receive from the third device a request to reserve the second resource; and
- sixth set of prestored instructions to reserve the second resource for the second
- 20 user.

11. A server for delivering content to be rendered by multiple devices, the server comprising:
- a computer operatively coupled to first, second, and third devices via a network, and operatively coupled to a plurality of rendering devices via a second
- 25 network, wherein at least one of the plurality of rendering devices provides at least one rendering resource available for use of a user, and provides at least one rendering resource available for simultaneous use of another user, the computer comprising:
- a memory;
- a processor coupled to the memory, the processor configured to

to provide the first device, via the network, with an indication of first and second resources available to a first user, wherein the first and second resources are for rendering content, and wherein the first and second resources are provided by the second device;

5 receive from the first device a request to reserve the first resource;

 reserve the first resource for the first user;

 provide the third device, via a network, with an indication of the second resource available to a second user;

10 receive from the third device a request to reserve the second resource; and

 reserve the second resource for the second user.

12. A method for determining a mapping for delivering content to multiple devices for rendering on the multiple devices, the method comprising:

15 receiving content requested by a user, the content comprising a plurality of content types;

 determining a set of resources available to the user for rendering content, the set of resources provided by a plurality of rendering devices;

 determining user preferences for delivering content, the user preferences including preference information relating to whether to deliver content to a single device or to multiple devices;

 determining a mapping of the plurality of content types to the plurality of rendering devices based on the set of resources available to the user and the user preferences.


25

13. A method as defined in claim 12, wherein the user preferences include at least one of utility functions, preference information relating to whether to deliver content of a same session to a single device or to multiple devices, preference information relating to whether to deliver content of different sessions to a single device or to multiple devices, and preference information relating to communication links for delivering content types to the plurality of devices.


30

14. A method for determining a mapping for delivering content to multiple devices for rendering on the multiple devices, the method comprising:
- receiving content requested by a user, the content comprising a plurality of content types;
- 5 determining a set of resources available to the user for rendering content, the set of resources provided by a plurality of rendering devices;
- determining user preferences for delivering content, the user preferences including preference information relating to preferred communication links for delivering content to the plurality of devices;
- 10 determining a mapping of the plurality of content types to the plurality of rendering devices based on the set of resources available to the user and the user preferences.



1/15




2/15

150**FIG. 2**


3/15

**FIG. 3****FIG. 4**


4/15




5/15

**FIG. 6**



6/15



**FIG. 7** 400

7/15


**FIG. 8**

8/15



**FIG. 11****FIG. 12**

10/15

700 **FIG. 13**

11/15

CONTENT TYPE DEVICE	TEXT	GRAPHICS	VIDEO	AUDIO
PDA	10	4	0	0
CELL PHONE	10	2	0	5
LAPTOP	10	10	10	10

800 **FIG. 14**


NETWORK DEVICE	WLAN	GPRS
PDA	10	5
CELL PHONE	5	10
LAPTOP	10	3

840 **FIG. 15**


CAPTION 1 AUDIO 1 VIDEO 1 AUDIO 2	CAPTION 1	AUDIO 1	VIDEO 1	AUDIO 2
CAPTION 1	0	20	50	-100
AUDIO 1	20	0	100	-100
VIDEO 1	50	100	0	-100
AUDIO 2	-100	-100	-100	0

860 **FIG. 16**


12/15



13/15

**FIG. 18**960

14/15

**FIG. 19**

1000

