

(12) STANDARD PATENT
(19) AUSTRALIAN PATENT OFFICE

(11) Application No. **AU 2004222762 B2**

(54) Title
Architecture for an extensible real-time collaboration system

(51) International Patent Classification(s)
G06F 9/54 (2006.01) **G06F 15/00** (2006.01)
G06F 9/38 (2006.01) **H04L 29/06** (2006.01)
G06F 9/46 (2006.01)

(21) Application No: **2004222762** (22) Date of Filing: **2004.10.19**

(30) Priority Data

(31) Number	(32) Date	(33) Country
60599807	2004.08.06	US
10918855	2004.08.14	US
60513790	2003.10.23	US
10918333	2004.08.13	US

(43) Publication Date: **2005.05.12**

(43) Publication Journal Date: **2005.05.12**

(44) Accepted Journal Date: **2010.03.04**

(71) Applicant(s)
Microsoft Corporation

(72) Inventor(s)
Ganesan, Krishnamurthy;Chitturi, Ajay P.;Osborne, Robert J.;Potra, Adrian;Barkley, Warren Vincent;Bobde, Nikhil P.;Ostergren, Brian H.;Han, Mu

(74) Agent / Attorney
Davies Collison Cave, 1 Nicholson Street, Melbourne, VIC, 3000

(56) Related Art
US 5539886
US 5724508
US 20030014488

ABSTRACT OF THE DISCLOSURE

An architecture for an extensible real-time collaboration system is provided. The architecture presents a unified application program interface for writing application programs that use communications protocols. The architecture has activity objects, endpoint objects, and multiple media stacks. These objects may use various communications protocols, such as Session Initiation Protocol or Real-Time Transport Protocol to send and receive messages. The activity objects, endpoint objects, and multiple media stacks may each have one or more APIs that an application developer can use to access or provide collaboration-related functionality. These objects map the API to the underlying implementation provided by other objects. Using the activity objects enables a developer to provide less application logic than would otherwise be necessary to provide complex collaboration services.

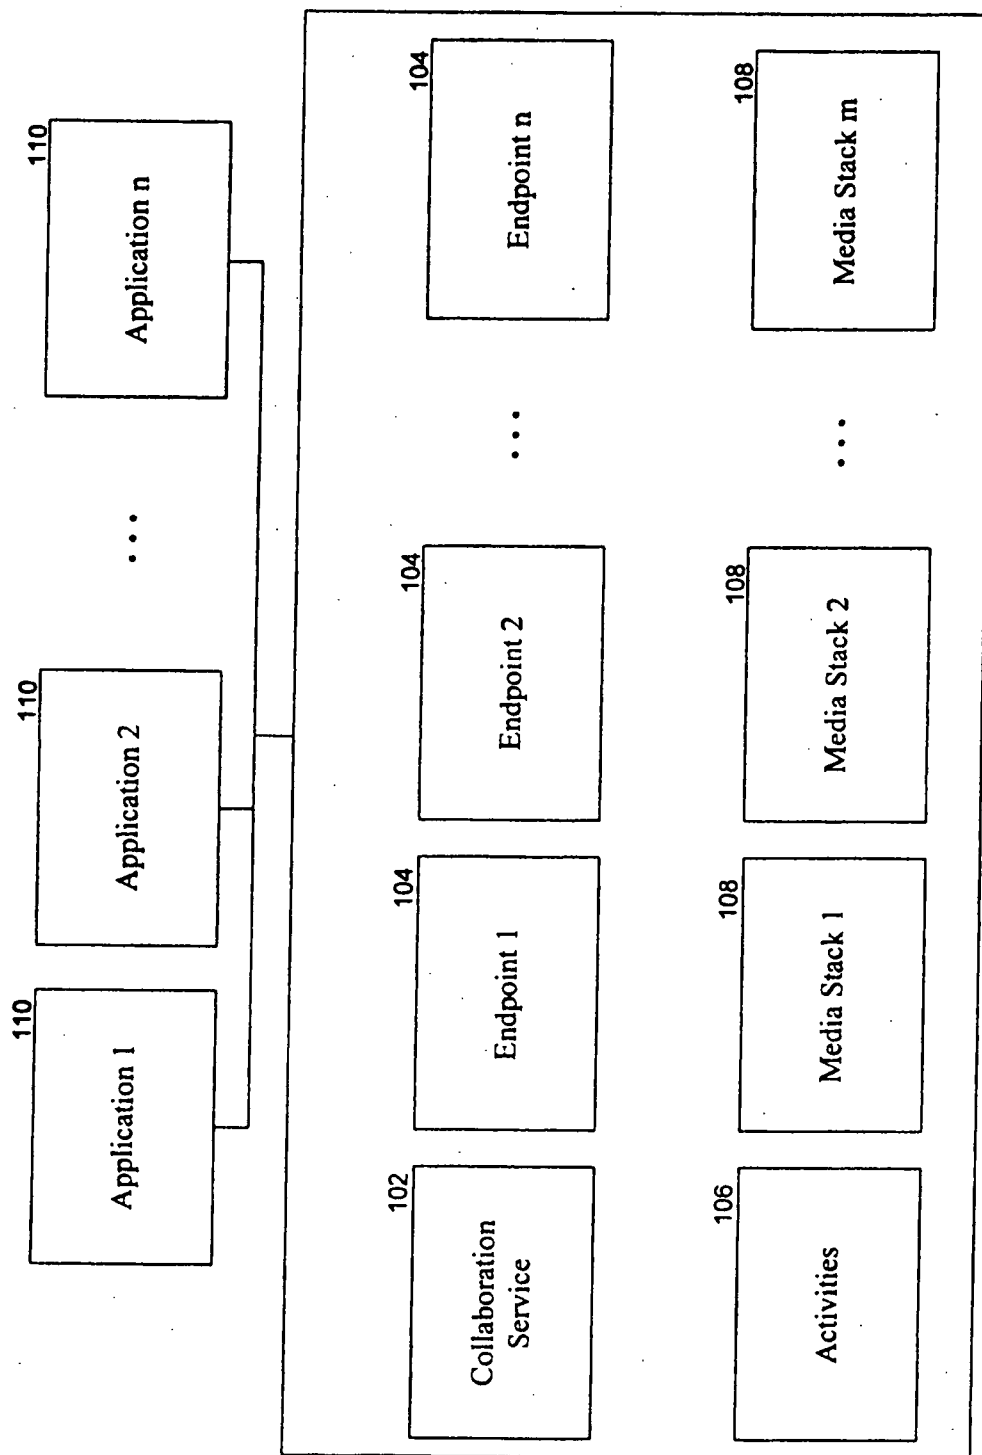


FIG. 1

2004222762 19 Oct 2004

AUSTRALIA
PATENTS ACT 1990
COMPLETE SPECIFICATION

NAME OF APPLICANT(S)::

Microsoft Corporation

ADDRESS FOR SERVICE:

DAVIES COLLISON CAVE
Patent Attorneys
1 Nicholson Street, Melbourne, 3000, Australia

INVENTION TITLE:

Architecture for an extensible real-time collaboration system

The following statement is a full description of this invention, including the best method of performing it known to me/us:-

CROSS-REFERENCE TO RELATED APPLICATION(S)

- [0001] This application claims the benefit of U.S. Provisional Application No. 60/513,790, filed on October 23, 2003, U.S. Provisional Application No. 60/599,807, filed on August 6, 2004, U.S. Parent Application 10/918,333, filed on August 13, 2004 and U.S. Parent Application 10/918,855 filed on August 14, 2004, which are incorporated herein in their entirety by reference.

TECHNICAL FIELD

- [0002] The described technology relates generally to data communications and, more particularly, to an architecture for an extensible real-time collaboration system.

BACKGROUND

- [0003] Various communications applications and protocols enable communications between software programs or users. As examples, real-time communications applications such as MICROSOFT WINDOWS MESSENGER and Voice over Internet Protocol ("VoIP") enable communications between users sending each other text, video, or voice data. These applications may use various protocols, such as Session Initiation Protocol ("SIP"), Real-Time Transport Protocol ("RTP"), and Real-Time Control Protocol ("RTCP"), to establish sessions and send communications-related information. SIP is an application-layer control protocol that devices can use to discover one another and to establish, modify, and terminate sessions between devices. RTP is a protocol for delivering audio and video over the Internet, and is frequently used in streaming media systems and videoconferencing systems in conjunction with other protocols such as RTCP and

H.323. RTCP is a protocol that enables a client application to monitor and control data sent or received using RTP, and is used with RTP. SIP and RTP/RTCP are Internet proposed standards. Their specifications, "RFC 3261" and "RFC 3550," and respectively, are available on the Internet at [www.ietf.org at /rfc/rfc3261.txt](http://www.ietf.org/rfc/rfc3261.txt) and [www.faqs.org at /rfcs/rfc3550.html](http://www.faqs.org/rfcs/rfc3550.html), respectively, and are incorporated herein in their entirety by reference.

[0004] Applications may additionally use other protocols. Applications may use enhanced versions of the protocols indicated above or altogether different protocols that are designed to carry specialized data. As an example, when a new or improved protocol for carrying videoconferencing information becomes available, an application developer creating or modifying an application may desire to use the new or improved protocol, e.g., to improve performance or offer additional features. To use the new or improved protocol, the application developer may need to modify portions of the application that interact with communications protocols because an improvement to a protocol or a new protocol may have a different interface than a protocol already being used. As an example, while a protocol may have a NewSession interface to create a new session, an improved version may have a StartSession method that creates and starts a session and accepts additional parameters. Because StartSession accepts additional parameters, its interface is different from NewSession, and so an application using NewSession may need to be modified to use StartSession. When a protocol has a different interface, application developers may need to learn the different interface and modify their applications to use this interface to use the protocol.

[0005] Application developers may need to become familiar with details of each of the many communications protocols they use in applications they develop. As an example, when an application developer uses SIP and RTP/RTCP, the application developer may need to be familiar with all three protocols to provide program logic relating to the protocols. An application developer not familiar with all three protocols may need additional training and time to become familiar with these protocols. Furthermore, when the application is to be modified to work with

2004222762 16 Feb 2010

- 3 -

additional or improved protocols, the application developer may need to revise or add programming logic so that the application can function with these protocols. This could lead to additional development expense and difficulty.

[0006] Furthermore, various protocols present a variety of complexities. For example, to provide a videoconferencing facility in an application, the application's developer would have to become familiar with a number of protocols and provide logic to coordinate these protocols to add videoconferencing capabilities. Adding other collaboration capabilities to an application, such as text messaging, voice messaging, etc., presents other similar complexities.

10 [0007] Thus, an architecture for an extensible real-time collaboration system that facilitates addition of collaboration features in an application without significant investment in developer training would have significant utility.

SUMMARY

15

[0008] An architecture for an extensible real-time collaboration system is described. The architecture presents a unified application program interface ("API") for writing application programs that use communications protocols. The architecture has activity objects, endpoint objects, and multiple media stacks.

20 These objects may use various communications protocols, such as Session Initiation Protocol or Real-Time Transport Protocol to send and receive messages containing information. The activity objects, endpoint objects, and multiple media stacks may each have one or more APIs that an application developer can use to access or provide collaboration-related functionality. These objects map the API to
25 the underlying implementation provided by other objects. Using the activity objects enables a developer to provide less application logic than would otherwise be necessary to provide complex collaboration services.

[0008a] Some embodiments relate to a method performed by a computing
30 system for providing real-time collaboration services in an application, comprising:
instantiating an activity object, the activity object providing the functionality

of an endpoint object and a media stack object both underlying the activity object, and having an application program interface for providing a collaboration service; and

invoking methods of the application program interface,

wherein the invoked methods provide management services of the endpoint object which provides an abstraction of a user at a computing device and is identified by a combination of an electronic address of the user and an endpoint identifier, the management services including providing presence information relating to whether the user is present at the computing device; and

wherein the invoked methods further provide content communications services of the media stack object between applications that are communicating information of the collaboration service.

[0008b] Some embodiments relate to an extensible real-time collaboration system, comprising:

multiple media stack objects;

an endpoint object for providing or receiving signalling information, the endpoint object providing an abstraction of a user at a computing device and being identified by a combination of an electronic address of the user and an endpoint identifier; and

multiple activity objects, wherein the activity objects provide the functionality of the multiple media stack objects and endpoint object all underlying the multiple activity objects, provide an application program interface to applications and utilize the multiple media stack objects and endpoint object to provide collaboration services, wherein the collaboration services include providing presence information relating to whether the user is present at the computing device, further wherein the application uses the application program interface and does not need to provide logic to coordinate the media stack objects and endpoint object.

[0008c] Some embodiments relate to computer-readable storage storing computer-executable instructions for performing the methods described herein or

2004222762 16 Feb 2010

- 4A -

for implementing the systems described herein.

BRIEF DESCRIPTION OF THE DRAWINGS

5 [0009] Figure 1 is a block diagram illustrating components of an architecture for an extensible real-time collaboration system.

[0010] Figure 2 is a block diagram illustrating components of an endpoint object of the extensible real-time collaboration system.

[0011] Figure 3 is a block diagram illustrating activity objects of the
10 extensible real-time collaboration system.

[0012] Figure 4 is a flow diagram illustrating a create_server_endpoint routine.

[0013] Figure 5 is an architecture diagram illustrating the architecture for an extensible real-time collaboration system.

15

DETAILED DESCRIPTION

[0014] In some embodiments, an architecture for an extensible real-time collaboration system is provided. The architecture provides a high-level application
20 program interface ("API") for writing application programs that use communications protocols to provide collaboration services. An application developer can add collaboration services to an application by using the API without needing to learn complexities associated with several underlying protocols that implement the collaboration services.

25 [0015] The architecture comprises activities objects, endpoint objects, and multiple media stacks. These objects may use various communications protocols, such as SIP or RTP/RTCP, to send and receive messages. The activities objects, endpoint objects, and multiple media stacks may each have one or more APIs that an application developer can use to access or provide functionality provided by the
30 objects. The application developer can choose to provide application logic that utilizes the APIs provided by endpoint objects and media stacks, or can choose to

2004222762 16 Feb 2010

- 4B -

provide application logic that utilizes the API provided by an activity object. By utilizing the APIs provided by the endpoint objects and media stacks, the application developer may be able to exercise a high degree of flexibility, but may have to provide significantly more application logic than if only an activity object's API is used. An application developer may choose to use the activity object's API for several reasons. The API of the activity objects provides a higher-level interface than the APIs of the endpoint object and media stacks. Furthermore, the

2004222762 16 Feb 2010

activity objects coordinate the endpoint object and the media stack, and so application logic may not need to be provided to perform the coordination.

[0016] Activity objects provide services to applications and other objects. Activity objects are objects that enable an application to participate in a variety of specific activities, and may include, as examples, instant messaging, teleconferencing, videoconferencing, application sharing, and other activities. Activity objects may be considered to be "wrapper" objects that provide the functionality of underlying objects, including endpoint objects and media stacks. In particular, activity objects coordinate endpoint objects and media stacks to provide a seamless and integrated session to other objects using the activity objects, such as applications.

[0017] An example of the benefit of using an activity object is provided by the following example. An application developer may desire to provide videoconferencing capabilities in an application. To do so, the application developer would first have to become familiar with a signaling protocol, such as SIP, and a media protocol, such as RTP/RTCP. Next, the application developer may have to provide application logic to create a session, determine whether a contact with whom videoconferencing is desired is presently online, send an invitation to join a videoconference, negotiate various parameters relating to the videoconference, capture audio and video from sound and video capturing hardware, and finally exchange audio/video data using RTP/RTCP. In contrast, by using a videoconferencing activity object of the architecture, many of these steps are eliminated because the videoconference activity object is especially designed to consolidate this application program logic into a few higher-level interfaces. The architecture has similar activity objects for a variety of other collaboration activities. Furthermore, the architecture provides support for additional activity objects to be added in the future.

[0018] An endpoint object provides management services, such as signaling. An endpoint object comprises profile, publishing/subscribing, signaling, and protocol stack components. The profile component may provide an abstraction of a user through an API. The publishing/subscribing component provides interfaces to track presence and availability information relating to users. The signaling

component may be used to provide or receive infrequent, transactional messages relating to establishing or controlling sessions. The signaling component may also be used for negotiating, e.g., media parameters such as bandwidth limit. The protocol stack component is used by the profile, publishing/subscribing, and signaling components to send or receive data, and supports various protocols, including SIP.

[0019] As previously described, the activity object provides a facility for application developers to add collaboration services to applications easily. As an example, an application developer desiring to add videoconferencing in an application may provide logic to create (or "instantiate") a videoconferencing activity object. The videoconferencing activity object may then instantiate an endpoint object for signaling (such as a SIP endpoint object) and a messaging media stack for carrying audio/video data. During subsequent videoconferencing activities (e.g., when sending or receiving an audio/video stream), the activity object may coordinate the objects it created or instantiated. In an embodiment, the application may create the objects and, optionally, provide an indication of the objects to the activity object. In such a case, the activity object may not need to create these objects.

[0020] Media stack objects provide content communications services, such as handling audio/video data streams. As an example, a media stack object may use RTP/RTCP to send or receive audiovisual information relating to a videoconference.

[0021] Turning now to the figures, Figure 1 is a block diagram illustrating components of an architecture for an extensible real-time collaboration system in an embodiment. The architecture for an extensible real-time collaboration system comprises a collaboration service object 102, multiple endpoint objects 104, activity objects 106, and multiple media stacks 108. One or more applications 110 may utilize the architecture by accessing various methods, properties, and events relating to the architecture. An application developer writing an application may utilize the architecture by using a unified API instead of having to learn and

implement a different API for each media stack, protocol, or other component that the application or architecture may use.

[0022] The collaboration service object 102 provides a facility for applications to share multiple endpoint objects and may provide a consistent API across a number of endpoint objects. As an example, if endpoint object 1 provides an interface relating to receiving (or sending) information, and endpoint object 2 similarly provides an interface relating to receiving (or sending) information, but the two interfaces use different names yet perform similar functions, the collaboration service object may provide a common name for both interfaces. When an application developer uses this common name in an application, the application developer may not need to revise the application when a new or modified object that provides an interface with a different name is used with the collaboration service object.

[0023] Endpoint objects 104 provide a facility for signaling other objects. Signaling may be used between two endpoint objects having a session, e.g., so that an endpoint object can invite or request the other endpoint object to conduct an activity or exchange information relating to the session. As examples, an endpoint object may invite the other endpoint object of the session to an instant messaging conversation and subsequently may send text messages relating to the conversation. Endpoint objects are further described below in greater detail in relation to Figure 2.

[0024] Activity objects are components that enable an application to participate in a variety of collaboration-related activities. These components provide an API that an application developer can use to coordinate endpoint objects and media stacks. Activity objects 106 are further described below in greater detail in relation to Figure 3.

[0025] A media stack object 108 provides content communications services, such as handling data streams, and provides an API for other objects to send or receive the data. The architecture is capable of supporting virtually an infinite number of media stacks by virtue of the fact that the architecture does not need to distinguish between the data or media types. As a result, new media stacks may be added or

media stacks may be modified as requirements change. An example of a media stack is RTP/RTCP. This media stack may be used to send audiovisual information.

[0026] Figure 2 is a block diagram illustrating components of an endpoint object of the extensible real-time collaboration system in an embodiment. Endpoint objects provide management services, such as a facility for signaling other objects. Signaling may be used between two endpoint objects having a session, e.g., so that an endpoint object can invite or request the other endpoint object to conduct an activity or exchange information relating to the session. As examples, an endpoint object may invite the other endpoint object of the session to an instant messaging conversation and subsequently may send text messages associated with the conversation.

[0027] The architecture may support several varieties of endpoint objects, and each variety of endpoint object may be instantiated multiple times. As an example, there may be an endpoint object relating to a user's personal Internet service provider account (e.g., MSN.COM) and another endpoint object relating to the user's corporate Internet account (e.g., MICROSOFT.COM). The user may be logged in to service providers using the personal account on multiple devices (e.g., a handheld computing device and a desktop computing device) and may also be logged in using the corporate account on some of the devices (e.g., the desktop computing device). Thus, there may be two instances relating to the URI associated with the personal account. Individual instances of endpoint objects may then be uniquely identified by a combination of a uniform resource locator ("URI") and an endpoint identifier ("EID"). As an example, an endpoint object may be identified by the URI user@MSN.COM and by the EID "1234." As previously described, the EID may be used to particularly distinguish an instance of an endpoint object from another instance of an endpoint object that is associated with the same URI.

[0028] An endpoint object may provide a "data" interface and a "signaling" interface. The data interface may include methods, events, and properties relating to data that is published or subscribed to by the endpoint object. By publishing or

subscribing to data, an application may provide data or be notified of changes to data. The signaling interface may provide methods, events, and properties relating to controlling signaling of the endpoint object. As examples, signaling characteristics may include creating or joining sessions, entering or leaving conversations, accepting or declining invitations, and other signals.

[0029] The illustrated endpoint object 200 comprises a profile component 201, publishing and subscribing component 202, signaling component 204, and protocol stack component 206.

[0030] The profile component may provide an abstraction of a user through an API. It maintains service-related information for the user, such as an electronic address (e.g., URI), credentials used for the service provider, the service provider's status, capability, and policies. Users may have multiple profiles in relation to different service providers. The user could also have more than one profile per service provider. The profile component may be used when creating sessions with other users, e.g., to provide the user's URI. The profile component may provide methods to log on or off a user in relation to a service provider.

[0031] The publishing/subscribing component provides interfaces to track presence and availability information relating to users. Presence information relates to whether a user is present at a particular computing device. Availability information relates to whether the present user is available to receive a message or willing to do so. As an example, a user of a cellular telephone may be present when the cellular telephone is turned on, but may be unavailable for a messaging session when the user is participating in a phone call. Similarly, a user who has set a "busy" indication in MICROSOFT WINDOWS MESSENGER may be present but unavailable for messaging.

[0032] As further examples, the presence object may provide information relating to a user who is present and available to participate in a MICROSOFT WINDOWS MESSENGER conversation using a computing device and is also available to participate in a teleconference using a cellular telephone. When the user is no longer logged into MICROSOFT WINDOWS MESSENGER, the presence object may update this information so that an application using the presence object is

able to determine that the user is no longer present or available to participate in the MICROSOFT WINDOWS MESSENGER conversation. Thus, presence information indicates whether users or other objects are present. Various service providers or protocols may use different mechanisms to produce or provide presence information. So that an application developer does not need to be aware of the multiple ways of producing or providing presence information, an application developer may use the endpoint object to produce or use presence information.

[0033] The publishing/subscribing component provides a subscribe interface to create a subscription to another object's publication, a publish interface to provide subscriptions to other objects, and a notify interface to receive notifications relating to services whose publications have been subscribed to. These interfaces enable an application to use the component to provide, receive, or track presence information. As an example, when a user participates in a MICROSOFT WINDOWS MESSENGER conversation using a personal computer and participates in a teleconference using a cellular telephone, the publishing/subscribing component may detect and report the user's presence at both locations. A URI and EID may together uniquely identify instances of endpoint objects. Because a user may be present at multiple locations simultaneously, the user's URI may be indicated as present at these multiple locations. The addition of an EID in relation to a given URI provides a mechanism to uniquely identify a particular instance of presence.

[0034] Notifications may be provided relating to various information. As examples, notifications may be provided relating to whether a user is online, busy, available, out for lunch, etc. Notifications may also be provided relating to geographic location of a user (e.g., as provided by a global positioning system or "GPS") contact information, calendar information, out of office messages, audio/video capabilities, etc.

[0035] The signaling component may also be used to provide or receive infrequent, transactional messages relating to establishing or controlling sessions. This component may also be used for negotiating, e.g., media parameters such as frames per second.

[0036] The protocol stack object is responsible for sending and receiving information using a protocol. As an example, SIP may be used to send or receive signaling information. In various embodiments, other protocols can equally be used. In an embodiment, an endpoint object may be compatible with multiple protocols. In such a case, the endpoint object may be able to use, e.g., multiple protocols, as necessary, to send or receive information. Alternatively, the architecture may support multiple endpoint object-protocol combinations as separate endpoint objects. In such a case, one endpoint object may be used for SIP and another for some other protocol.

[0037] Figure 3 is a block diagram illustrating activity objects of the extensible real-time collaboration system in an embodiment. Activity objects provide services to applications and other objects. Activity objects are objects that enable an application to participate in a variety of specific activities, and may include, as examples, instant messaging, teleconferencing, videoconferencing, application sharing, and other activities. Activity objects may be considered to be "wrapper" objects that provide the functionality of underlying objects, including endpoint objects and media stacks. In particular, activity objects coordinate endpoint objects and media stacks to provide a seamless and integrated session to objects using the activity objects, such as applications.

[0038] Each activity object provides functionality that enables an application using the object to participate in a variety of activities. As an example, an activity object may encapsulate signaling and media information. In doing so, the activity object may serve as a bridge between signaling and media, so that an application using the activity object is provided an integrated view. For example, an application developer would merely request a videoconference rather than having to set up separate SIP and RTP/RTCP connections and sessions. The illustrated examples of activities are described below.

[0039] A collaboration session activity object 304 enables an application to provide collaboration services to the application's users. As an example, the collaboration session activity object may enable two users using two different computers to collaborate by authoring a shared document.

[0040] An application sharing activity object 310 enables an application using the application sharing activity object to provide functionality relating to the sharing of applications between users using different computing devices. As an example, two users may share a "whiteboard" application, e.g., using which a user can provide visual information that the other user can view and manipulate.

[0041] A messaging activity object 306 provides an ability for applications to provide messaging capabilities to the application's users. As an example, an application developer may want to enable messaging within an application the developer is creating. As a specific example, when two users are collaborating on a document by simultaneously editing a document, the users may want to be able to send text messages to one another or participate in a videoconference during the collaboration. In such a case, an application developer may use a collaboration session activity object as well as messaging or videoconferencing activity objects to enable users to collaborate and exchange messages with one another during the collaboration.

[0042] A videoconferencing activity object 308 provides videoconferencing facilities to an application. Videoconferencing may involve sending and receiving audiovisual information.

[0043] Additional activity objects are also contemplated, and are represented as activity objects 312.

[0044] An application developer may use objects of the architecture described above (and those not listed or described) by utilizing APIs provided by the objects. These objects may provide an easy-to-use API so that an application developer may not need to reference APIs provided by underlying components that provide the services associated with the activity objects. As an example, a messaging service provider may provide an API that a developer could use. To do so, the developer may need to spend time to learn the API, which can be quite complicated. Instead, the developer may wish to use a simpler API provided by an object of the architecture. Furthermore, the object may encapsulate steps that may be required to use multiple different objects. As an example, an application developer wishing to exchange messages between two computers may need to

utilize an API provided by the SIP as well as an API exposed by another low-level object that provides messaging services. In contrast, the application developer would only need to use the messaging activity object, and thereby be able to add messaging functionality to an application much more easily. Furthermore, the framework may operate to coordinate multiple objects, thereby requiring less programming logic from the application developer.

[0045] In an embodiment, a collaboration session contains and uses activity objects.

[0046] Figure 4 is a flow diagram illustrating a `create_server_endpoint` routine in an embodiment. The routine is called by an application to create an endpoint object that connects to a server. When an endpoint is created that connects to a server, information it publishes may be available to subscribing objects even after the created endpoint is no longer operating. Thus, an endpoint connected to a server may provide "per-URI" information meaning that the information remains available after the object's lifetime.

[0047] The routine begins at block 402. At block 404, the routine creates a new endpoint object, and indicates that the endpoint is related to an application. The indicated application may be provided as a parameter to a create function that operates to create the endpoint. When creating an endpoint, a "friendly" name may be provided so that the endpoint may be referenced by the friendly name. Alternatively, the newly created endpoint may be referenced by a unique identifier associated with the endpoint. This unique identifier may be generated by the system when the object is created.

[0048] At block 406, upon creating the endpoint, the application may register the newly created endpoint object to the server to allow the server to route messages to this endpoint. Upon receiving the register request from the endpoint object, the server may issue a challenge to the endpoint. The challenge may contain a "realm" used by the server. A realm may indicate a domain name associated with the server. As an example, the server may issue a challenge with the realm "MICROSOFT.com."

- [0049] At block 408, the routine responds to the challenge by providing credentials (e.g., user id and password) associated with the application. These credentials may be supplied by a user or automatically. The server may validate the credentials that the routine supplies. The credentials may be associated with the realm. For example, if the application provides credentials that are not associated with the server's realm ("MICROSOFT.com"), the server may not authenticate the application.
- [0050] If the registration is successful, the endpoint may be ready to receive messages from other objects. Otherwise, the endpoint may be unable to receive messages.
- [0051] The routine returns to its caller at block 412.
- [0052] In an embodiment, the server may enable an endpoint that is not successfully registered to send messages but not receive messages. Alternatively, in a weaker security model, the server may enable any endpoint to send or receive messages.
- [0053] In an embodiment, the server may challenge an endpoint when the endpoint attempts to send a message but not when it receives messages.
- [0054] An alternate routine (not shown) may create a peer-to-peer endpoint. A peer-to-peer endpoint is one that is not associated with a server. When an endpoint is not associated with a server, information the endpoint publishes may not be available once the endpoint stops operating.
- [0055] Figure 5 is an architecture diagram illustrating the architecture for an extensible real-time collaboration system in an embodiment. The architecture comprises an activity object 502, endpoint 504, and multiple media stack objects 506. These objects have been discussed in extensive detail. The architecture diagram indicates a relationship between the activity object, endpoint object, and media stack objects in an embodiment. Specifically, the architecture diagram indicates that functionality provided by an activity object includes functionality provided by the endpoint and media stack objects.
- [0056] In an embodiment, multiple activity objects may be used with corresponding media stack objects.

- [0057] Applications may call methods of objects of the architecture, or may directly call methods of underlying objects that implement functionality. By calling methods of the architecture's objects, application developers may need to provide less logic, and may not need to revise the application logic when the underlying components change.
- [0058] The following presents some APIs provided by the architecture.
- [0059] An application can create an endpoint object by creating a new "CollaborationEndpoint." The application can provide the following parameters: a URI relating to the endpoint object, a server associated with the endpoint object, and an indication of network credentials.
- [0060] A NetworkCredential method provides the indication of network credentials. This method accepts as parameters an indication of a user account, password, and domain with which the account is associated.
- [0061] An Enter method registers the endpoint, and provides an indication of a success or failure. The method does not require parameters.
- [0062] A Publish method publishes presence information. As examples, an application can indicate that a user is online, offline, busy, on the phone, etc. The architecture is flexible enough to provide virtually an unlimited number of presence indications. For example, an application may choose to publish a user's GPS location.
- [0063] Conversely, a Subscribe method subscribes to an endpoint object's publication.
- [0064] An Invite method invites a user to a collaboration session. The method receives an indication of a URI that should be invited.
- [0065] An Accept method accepts an invitation. Conversely, a Decline method declines the invitation.
- [0066] A messaging activity object can be created from an InstantMessagingActivity class. This activity object supports various methods including, e.g., a method to send a message, SendMessage.
- [0067] The SendMessage method sends a message. It receives a message string as a parameter.

[0068] The methods may be provided by various objects of the architecture. The architecture may additionally provide methods, properties, and events relating to retrieving a user's properties, adding a user to a list of contacts, adding or removing participants from a conversation, and so on.

[0069] The computing device on which the architecture is implemented may include a central processing unit, memory, input devices (e.g., keyboard and pointing devices), output devices (e.g., display devices), and storage devices (e.g., disk drives). The memory and storage devices are computer-readable media that may contain instructions that implement the system. In addition, the data structures and message structures may be stored or transmitted via a data transmission medium, such as a signal on a communications link. Various communications links may be used, such as the Internet, a local area network, a wide area network, or a point-to-point dial-up connection.

[0070] The architecture may be implemented in a variety of operating environments, including computing devices running a MICROSOFT WINDOWS operating system. This operating environment is only one example of a suitable operating environment and is not intended to suggest any limitation as to the scope of use or functionality of the system. Other well-known computing systems, environments, and configurations that may be suitable for use include personal computers, server computers, hand-held or laptop devices including "smart" cellular telephones, multiprocessor systems, microprocessor-based systems, programmable consumer electronics, network PCs, minicomputers, mainframe computers, distributed computing environments that include any of the above systems or devices, and the like.

[0071] The architecture may be described in the general context of computer-executable instructions, such as program modules, executed by one or more computers or other devices. Generally, program modules include routines, programs, objects, components, data structures, etc., that perform particular tasks or implement particular abstract data types. Typically, the functionality of the program modules may be combined or distributed as desired in various embodiments.

From the foregoing, it will be appreciated that specific embodiments of the invention have been described herein for purposes of illustration, but that various modifications may be made without deviating from the spirit and scope of the invention. Accordingly, the invention is not limited except as by the appended claims.

Throughout this specification and the claims which follow, unless the context requires otherwise, the word "comprise", and variations such as "comprises" or "comprising", will be understood to imply the inclusion of a stated integer or step or group of integers or steps but not the exclusion of any other integer or step or group of integers or steps.

The reference to any prior art in this specification is not, and should not be taken as, an acknowledgement or any form of suggestion that that prior art forms part of the common general knowledge in Australia.

[0072]

THE CLAIMS DEFINING THE INVENTION ARE AS FOLLOWS:

1. A method performed by a computing system for providing real-time collaboration services in an application, comprising:
- 5 instantiating an activity object, the activity object providing the functionality of an endpoint object and a media stack object both underlying the activity object, and having an application program interface for providing a collaboration service; and
- invoking methods of the application program interface,
- 10 wherein the invoked methods provide management services of the endpoint object which provides an abstraction of a user at a computing device and is identified by a combination of an electronic address of the user and an endpoint identifier, the management services including providing presence information relating to whether the user is present at the computing device; and
- 15 wherein the invoked methods further provide content communications services of the media stack object between applications that are communicating information of the collaboration service.
2. The method of claim 1, wherein the endpoint object uses a Session
- 20 Initiation Protocol.
3. The method of claim 2, wherein the media stack object uses Real-Time Transport Protocol.
- 25 4. The method of claim 2, wherein the media stack object uses Real-Time Control Protocol.
5. The method of any one of claims 1 to 4, wherein providing the content communications services includes providing media.

2004222762 16 Feb 2010

- 19 -

6. An extensible real-time collaboration system, comprising:
multiple media stack objects;
an endpoint object for providing or receiving signalling information, the
endpoint object providing an abstraction of a user at a computing device and being
5 identified by a combination of an electronic address of the user and an endpoint
identifier; and
multiple activity objects, wherein the activity objects provide the functionality
of the multiple media stack objects and endpoint object all underlying the multiple
activity objects, provide an application program interface to applications and utilize
10 the multiple media stack objects and endpoint object to provide collaboration
services, wherein the collaboration services include providing presence
information relating to whether the user is present at the computing device, further
wherein the application uses the application program interface and does not need
to provide logic to coordinate the media stack objects and endpoint object.
- 15 7. The extensible real-time collaboration system of claim 6, wherein the
application program interface provides a consistent interface name even when two
underlying objects actually providing associated functionality use different interface
names.
- 20 8. The extensible real-time collaboration system of claim 6 or 7, wherein one
of the collaboration services is messaging.
9. The extensible real-time collaboration system of any one of claims 6 to 8,
25 wherein one of the collaboration services is videoconferencing.
10. The extensible real-time collaboration system of any one of claims 6 to 9,
wherein one of the collaboration services is application sharing.
- 30 11. The extensible real-time collaboration system of any one of claims 6 to 10,
adding media stack objects.

2004222762 16 Feb 2010

- 20 -

12. The extensible real-time collaboration system of any one of claims 6 to 11, wherein the endpoint object comprises a profile component providing the abstraction of the user at the computing device, a publishing and subscribing component providing interfaces to provide the presence information, a signaling component adapted to provide infrequent, transactional messages relating to establishing or controlling a real-time collaboration session, and a protocol stack component responsible for sending and receiving information using a protocol.
13. The extensible real-time collaboration system of claim 12, wherein the protocol stack component uses a Session Initiation Protocol.
14. The extensible real-time collaboration system of claim 12, wherein the protocol stack component uses a signaling protocol.
15. The extensible real-time collaboration system of any one of claims 6 to 14, wherein one of the multiple media stack objects uses Real-Time Transport Protocol.
16. The extensible real-time collaboration system of claim 15, wherein the activity objects coordinate the multiple media stack objects and the endpoint object under direction of an application.
17. A method substantially as hereinbefore described with reference to the accompanying drawings.
18. A system substantially as hereinbefore described with reference to the accompanying drawings.
19. Computer-readable storage storing computer-executable instructions for performing the method of any one of claims 1 to 5 and 17 or for implementing the system of any one of claims 6 to 16 and 18.

2004222762 16 Feb 2010

- 21 -

20. A system comprising means for performing the method of any one of claims 1 to 5 or comprising the computer-readable storage of claim 19.

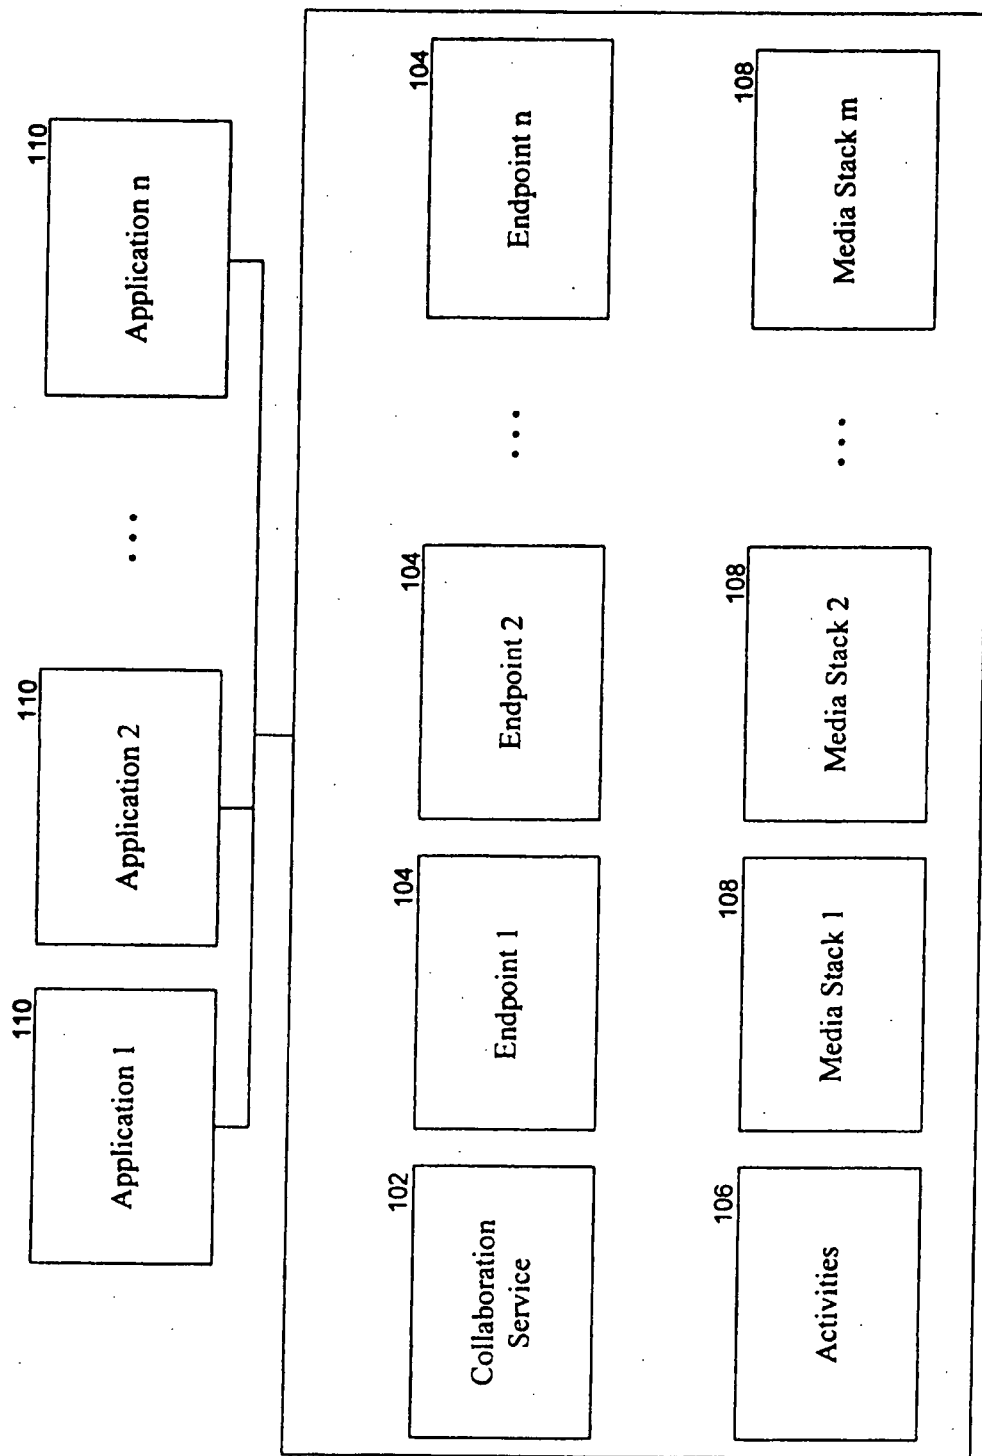


FIG. 1

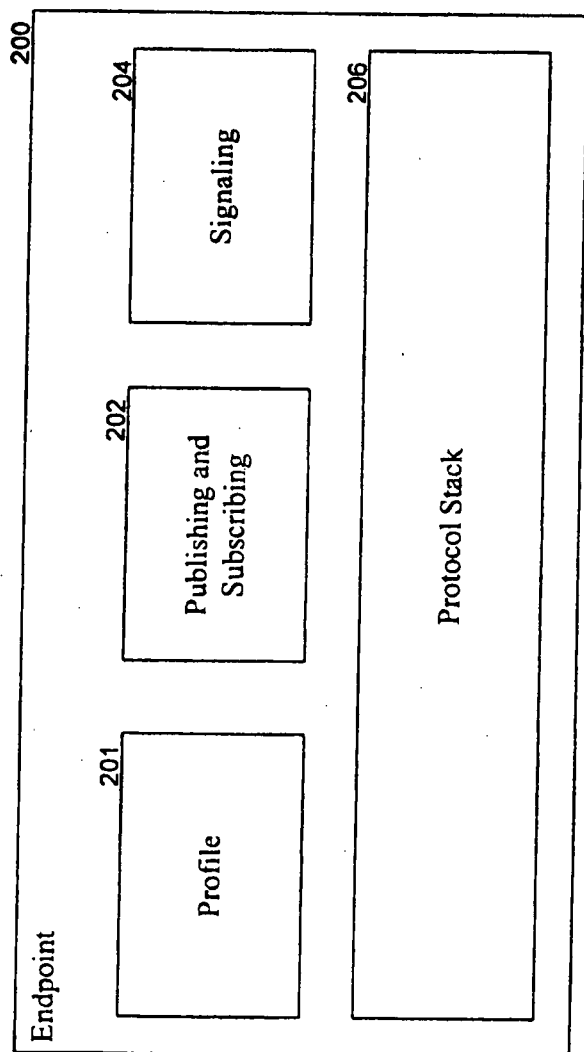


FIG. 2

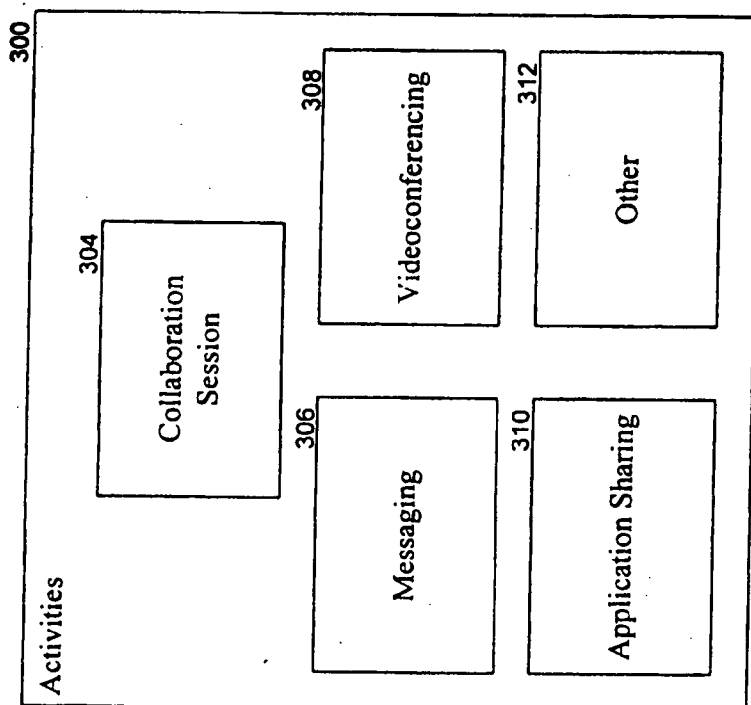


FIG. 3

Create_server_endpoint

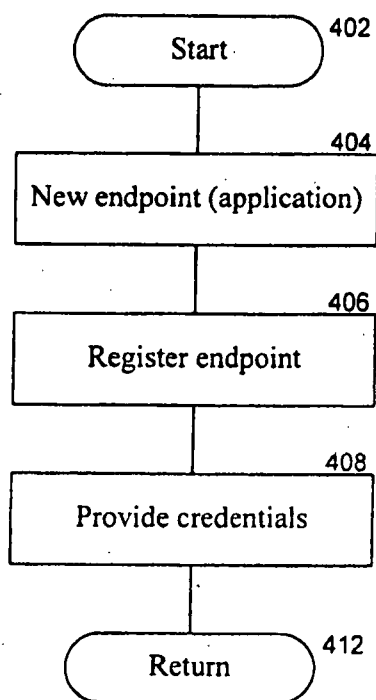


FIG. 4

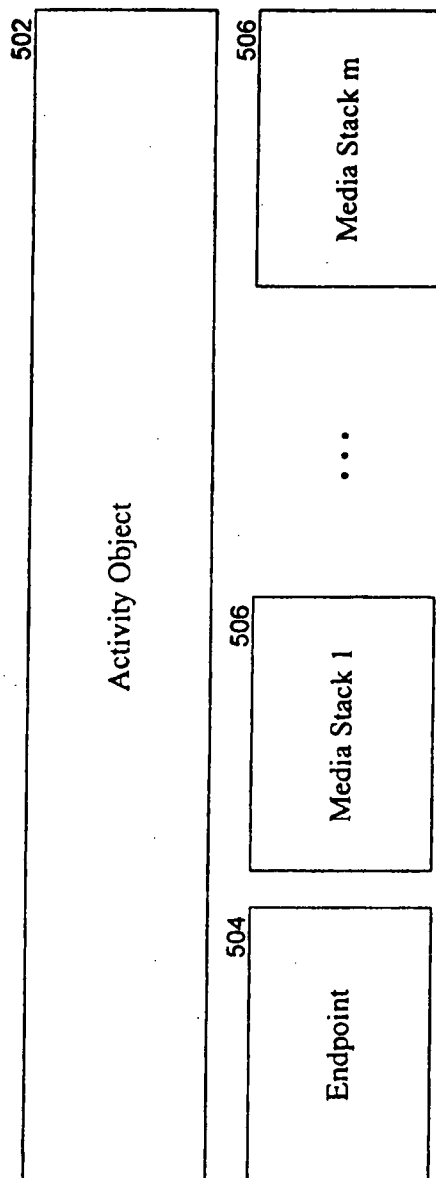


FIG. 5