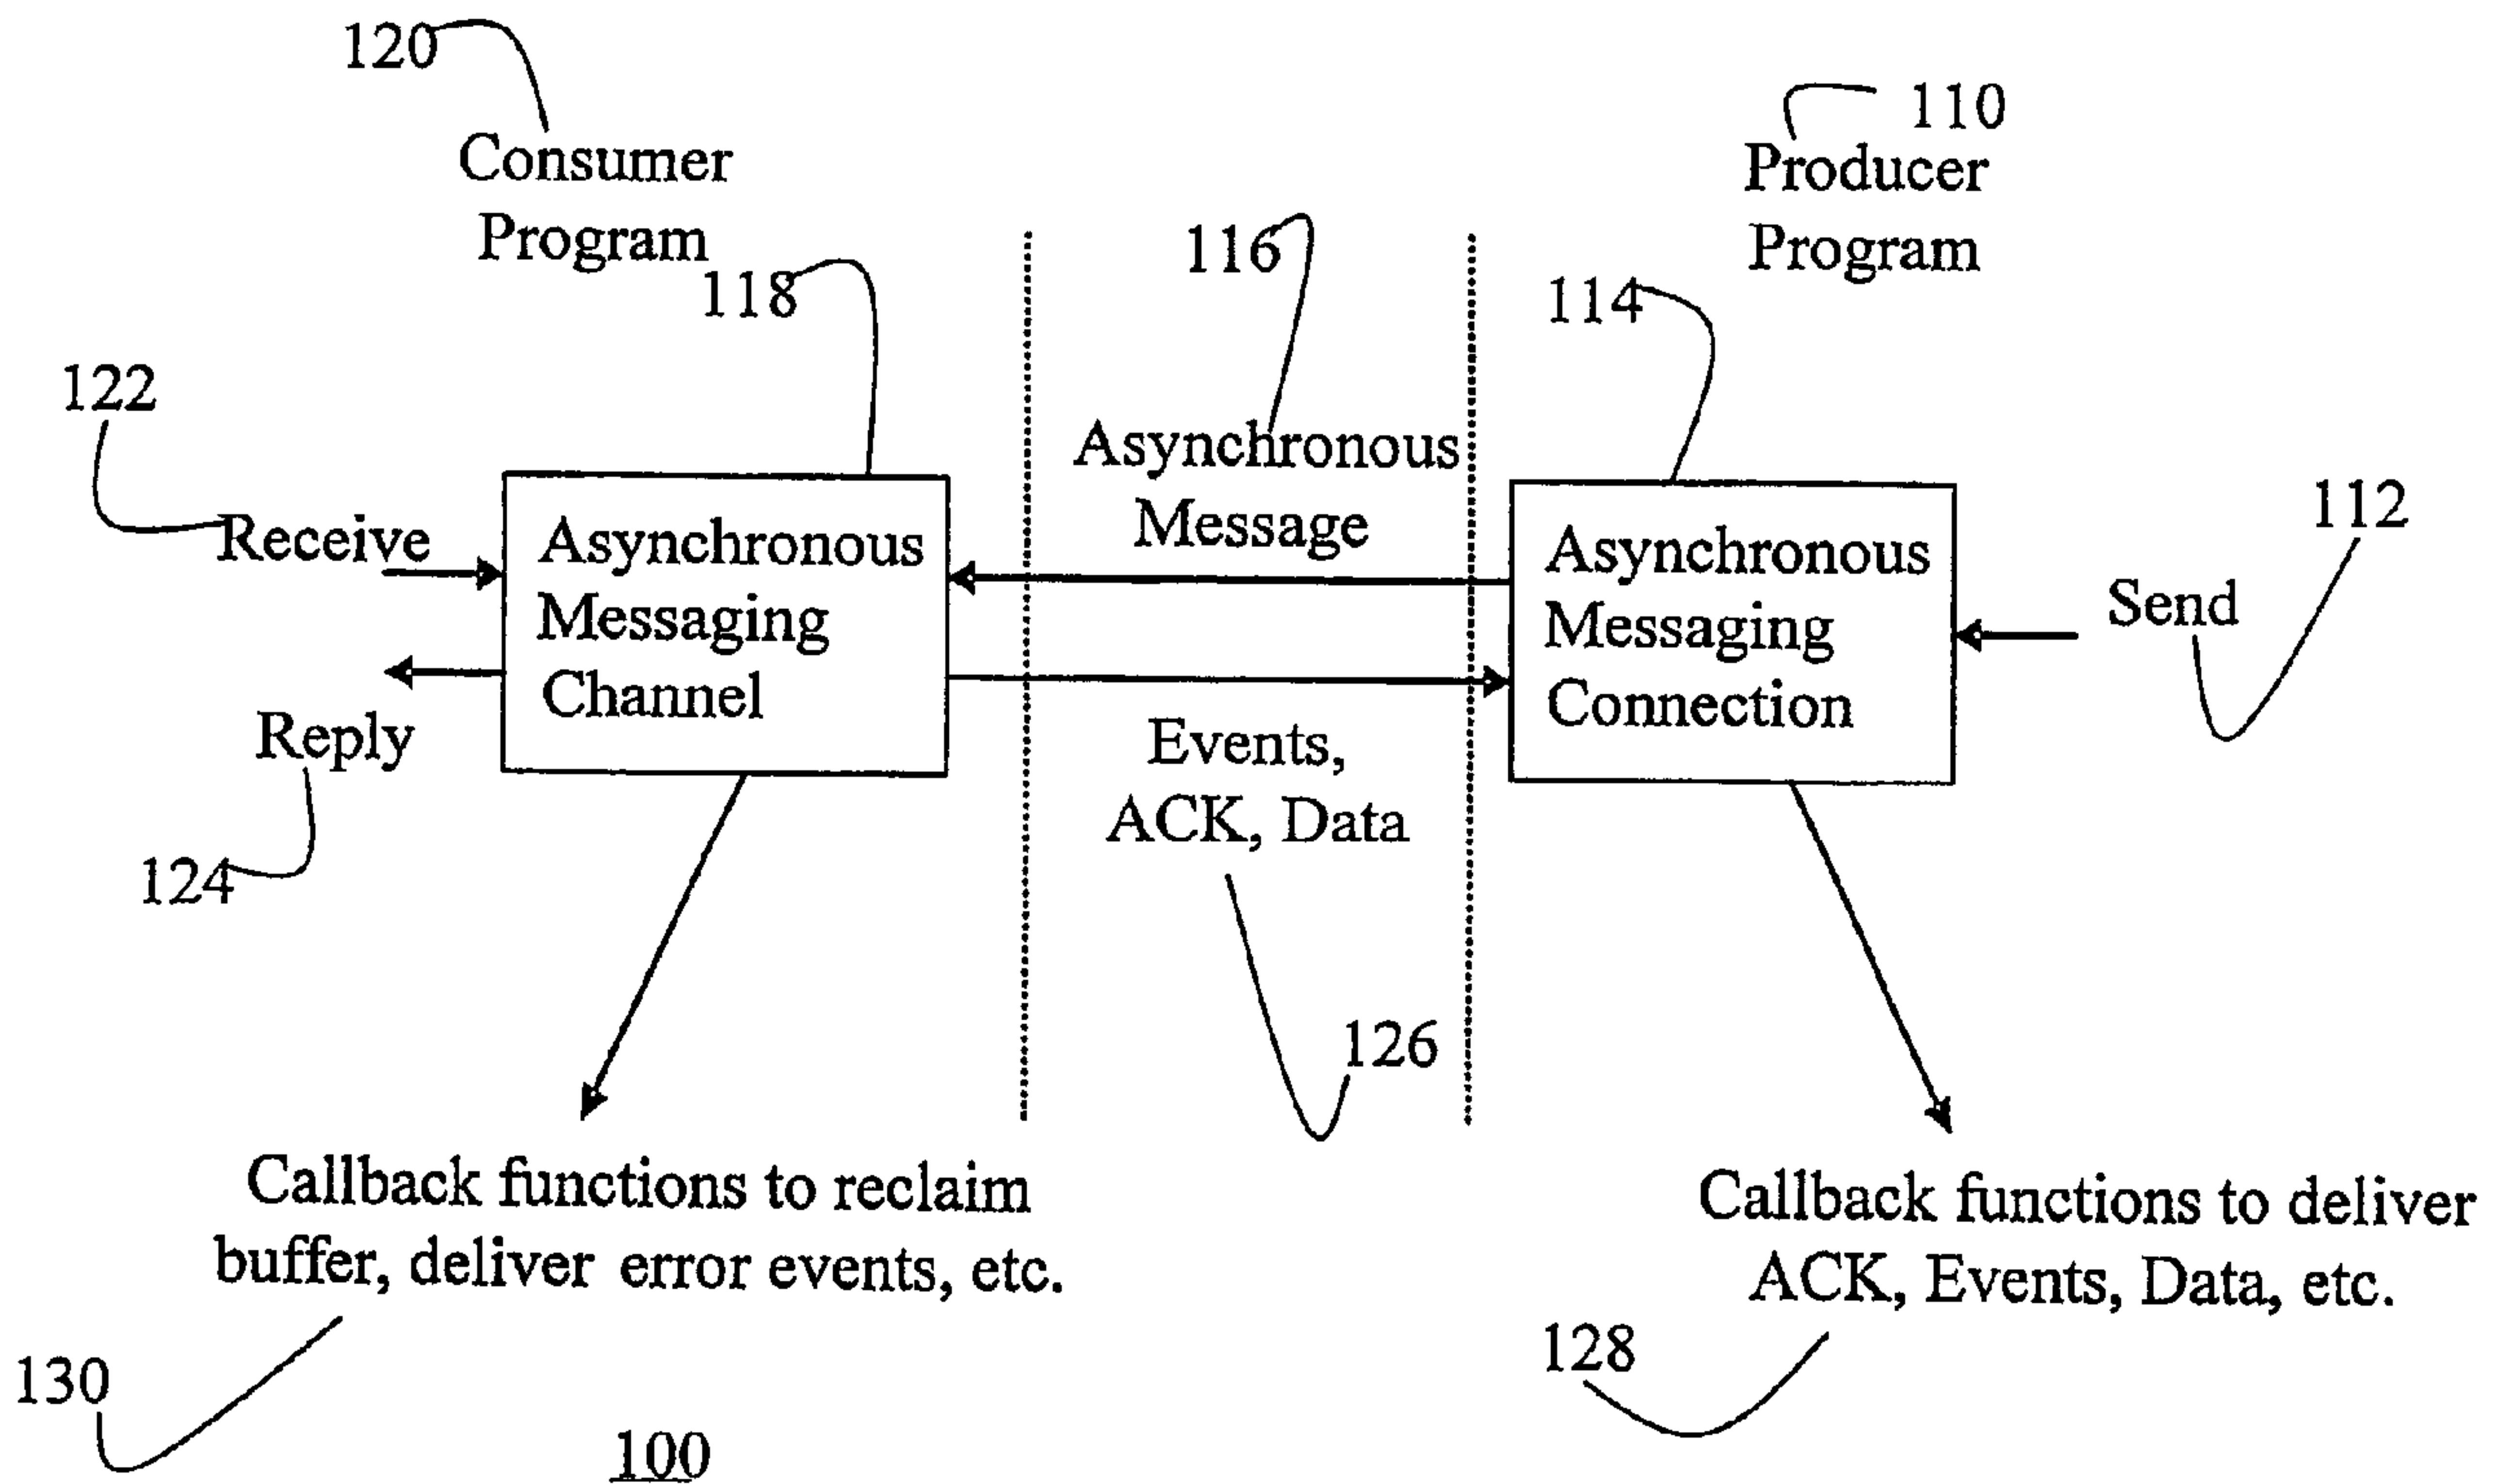




(22) Date de dépôt/Filing Date: 2006/02/10
 (41) Mise à la disp. pub./Open to Public Insp.: 2006/08/14
 (30) Priorités/Priorities: 2005/02/14 (US60/652929);
 2005/06/03 (US11/145105)

(51) Cl.Int./Int.Cl. *G06F 9/54* (2006.01)
 (71) Demandeur/Applicant:
 QNX SOFTWARE SYSTEMS, CA
 (72) Inventeurs/Inventors:
 ZHOU, HAO, CA;
 MARINEAU-MES, SEBASTIEN, CA;
 VAN DER VEEN, PETER, CA;
 KATHAIL, PARDEEP, US;
 BELAIR, STEVE, US
 (74) Agent: OYEN WIGGS GREEN & MUTALA LLP

(54) Titre : SYSTEME DE MESSAGERIE ASYNCHRONE RAPIDE A PROTECTION DE MEMOIRE DANS UN ENVIRONNEMENT A PROCESSUS MULTIPLES ET FILES MULTIPLES
 (54) Title: FAST AND MEMORY PROTECTED ASYNCHRONOUS MESSAGE SCHEME IN A MULTI-PROCESS AND MULTI-THREAD ENVIRONMENT



(57) Abrégé/Abstract:

An asynchronous message passing mechanism that allows for multiple messages to be batched for delivery between processes, while allowing for full memory protection during data transfers and a lockless mechanism for speeding up queue operation and queuing and delivering messages simultaneously.

ABSTRACT

An asynchronous message passing mechanism that allows for multiple messages to be batched for delivery between processes, while allowing for full memory protection during data transfers and a lockless mechanism for speeding up queue operation and queuing and
5 delivering messages simultaneously.

**FAST AND MEMORY PROTECTED ASYNCHRONOUS MESSAGE SCHEME IN A
MULTI-PROCESS AND MULTI-THREAD ENVIRONMENT**

INVENTORS:

5

**Hao Zhou
Sebastien Marineau-Mes
Peter van der Veen
Pardeep Kathail
Steve Belair**

10

BACKGROUND OF THE INVENTION

1. Technical Field.

[0001] The present application relates to computer operating systems, and more specifically, to inter-process communications in multi-process and/or multi-threaded environments.

2. Related Art.

15

[0002] Inter-process communication (IPC), which generally refers the exchange of data between programs either within the same computer or over a network, has become vital in today's real time distributed operations systems. IPC may be implemented using various data transfer methodologies, and is typically provided by the kernel module of the operating system. The kernel, which provides various user level computer programs with secure access to the computer's hardware, may provide IPC to allow coordination of processing among various processes and threads running on the system. As known in the art, a thread is a conveniently sized collection of programming steps that are scheduled and executed as a group, while a process may act as a "container" of threads. Processes may define the address space within which threads will execute. A process may contain at least one thread.

20

25

30

[0003] Message passing may be implemented to provide IPC throughout the entire system. In general, a message may be a packet of bytes passed from one process to another with no special meaning attached to the content of the message. The data in a message may have meaning for the sender of the message and for its receiver, but for no one else. Message passing not only allows processes to pass data to each other, but also provides a means of synchronizing the execution of several processes. As they send, receive, and reply to messages, processes undergo various "changes of state" that affect when and for how long,

they may run. Knowing their states and priorities, the operating system may schedule processes as efficiently as possible to optimize the available processor resources.

5 [0004] To manage these changes of state and avoid deadlock situations that may occur due to communications taking place in the wrong state, operating systems employ synchronous message passing systems. Synchronous message passing systems are those that require coordination among the sending, receiving, and replying to of messages between the threads or processes. While these synchronous systems are ideal for enforcing in-order processing of messages, they are prone to out of state and deadlock conditions, and do not provide for high level of data throughput as messages must be sent individually. Moreover, these problems
10 become exacerbated as the number of intercommunication processes or threads increase, limiting their effectiveness in today's data intensive processing.

[0005] To accommodate these processing needs, asynchronous systems have been developed to transfer messages independently of the coordination between communicating threads or processes. While these systems do provide great benefits, they still suffer from various
15 performance issues. For example, some asynchronous message passing systems do not provide for full memory protected data transfers. As a result, programs may corrupt the address space of one another. Additionally, asynchronous message passing systems typically allow messages to be only buffered or sent at any point in time, blocking the thread from queuing new messages if previously queued messages are being transferred.

20 [0006] Accordingly, there is a need for an asynchronous message passing system that may provide mechanisms for sending multiple messages by message buffering, sending and/or receiving a batch of messages based on a triggering method. The system may also provide for full memory protection when passing messages between the threads or processes, and may also provide a lockless queuing mechanism that allows for sending and buffering
25 messages simultaneously.

SUMMARY

[0007] Systems and methods of managing asynchronous interprocess communications in distributed operating systems to accelerate existing mechanism in today's operating systems are described. Existing asynchronous messaging schemes make one kernel call to send or
30 receive each message and employ locking schemes that slow performance by introducing extra locking overhead and not allowing simultaneous messages queuing and delivery in a multiprocessor system. In addition, some do not provide full memory protection for these

communications. The systems and methods described here provide an asynchronous message passing mechanism that allows for multiple messages to be batched for delivery between processes while allowing for full memory protection during data transfers, and a lockless mechanism for speeding up queue operation and queuing and delivering messages simultaneously.

[0008] Other systems, methods, features and advantages of the invention will be, or will become, apparent to one with skill in the art upon examination of the following figures and detailed description. It is intended that all such additional systems, methods, features and advantages be included within this description, be within the scope of the invention, and be protected by the following claims.

BRIEF DESCRIPTION OF THE DRAWINGS

[0009] The invention can be better understood with reference to the following drawings and description. The components in the figures are not necessarily to scale, emphasis instead being placed upon illustrating the principles of the invention. Moreover, in the figures, like referenced numerals designate corresponding parts throughout the different views.

[0010] Figure 1 shows an exemplary bidirectional communication model for delivering asynchronous messages between threads of an operating system.

[0011] Figure 2 shows exemplary connection object data structures for passing messages asynchronously.

[0012] Figure 3 shows an exemplary flow chart depicting the operations of a kernel module and an asynchronous messaging library.

[0013] Figure 4 is a diagram depicting the relationship between exemplary functions of an asynchronous messaging library and associated kernel calls of a kernel module for passing messages asynchronously.

[0014] Figure 5 is an exemplary communication model to use asynchronous messages to deliver large batch data, and to use synchronous message to deliver control and status events.

[0015] Figure 6 is an exemplary flow chart depicting exemplary asynchronous message passing operations of a consumer program and producer program.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

[0016] The asynchronous message passing methods and systems described may be provided as software programs that are generally stored in an executable form on a computer readable medium such as a random access memory (RAM), read only memory (ROM), optical disk (CD-ROM), or magnetic storage medium (hard drive or portable diskette), or may be implemented by hardware means, or other means known in the art. These functionalities may be provided, for example, as combination of functions of an asynchronous messaging library and extensions to existing synchronous kernel calls of a real time operating system (RTOS). The asynchronous messaging library may include functions for both sending and receiving messages. The functionalities described also may be provided over multiple libraries.

[0017] An exemplary bidirectional communication model 100 for delivering asynchronous messages between programs of an operating system is shown in FIG. 1. The model 100 includes a producer thread 110 that sends data to a consumer program 120. The programs 110 and 120 may be threads, processes, and the like operating on the same node of a distributed operating system. Alternatively, the programs 110 and 120 may operate on different nodes of the system. To pass asynchronous messages between the programs 110 and 120, the producer program 110 may call a send function 112 to forward messages to an asynchronous messaging connection 114 associated with the producer program 110. The connection 114 may provide a mechanism for sending asynchronous messages 116 to an asynchronous messaging channel 118 associated with the consumer program 120. For example, each connection 114 may buffer messages from the producer thread 110 until a certain trigger event has occurred. Messages may be buffered until, for example, a certain number of messages are written to the buffer or a certain amount of time has elapsed since the last passing of messages. Messages may be passed by joining the connection 114 and channel 118 at the kernel level, and copying data directly from memory space associated with the producer program to that associated with the consumer program 120. Alternatively, or additionally, triggering and data passing may be performed using any known method. The consumer thread 120 associated with this channel 118 may call a receive function 122 to collect the message 116, and a reply function 124 to send a reply or acknowledgement message 126 to the producer program 110. Call back functions 128, which may automatically be called upon the occurrence of a particular event, may be provided to claim the reply or acknowledgement messages 126, handle buffer errors, and the like. Additionally,

call back functions 130 may also be provided on the consumer program 120 side to handle errors, such as message delivery and send buffer errors.

[0018] The asynchronous messaging connection 114 and channel 118 may provide a level of abstraction that allows user programs to target particular ports 114 and 118, rather than specific threads, processes, and the like. The connection 114 and channel 118 may be bidirectional to both send and receive messages, or multiple one-way connections 114 and channels 118 may be used to send and receive messages among programs, threads, and the like. To receive asynchronous messages, programs such as the consumer program 120 may create an asynchronous messaging port. Asynchronous messaging ports designed to receive messages may be referred to as a channel 118. After a channel 118 is established, any program may attach to this channel to begin sending messages. Multiple threads, processes and the like may all connect to the same channel 118.

[0019] Each channel 118 may have a variety of attributes, such as a channel ID used to identify the channel 118. Additionally, each channel 118 may also include a queue for received messages. An additional queue for a free buffer list may also be included. Channels 118 may be defined by channel objects that store these channel attributes and buffer management data. For example, an asynchronous messaging channel object may include attributes that define a linked list of channels, a channel ID, flags for the channel, access permissions, buffer size, maximum number of buffers allowed, an event to be set for notification, a queue for received messages, a queue for a free buffer list, and a mutex, or mutual exclusion lock, used to protect the integrity of the received message queue. These channel objects may be stored in a hash table.

[0020] Another asynchronous messaging port, referred to as a connection 114, may be created by the producer program 110 for passing asynchronous messages 116 to a consumer program 120. In order to pass asynchronous messages to a consumer program 120, a connection 114 may be created and attached to a channel 118. Like a channel 118, the connection 114 may be defined by a connection object used to store connection attributes and buffer data management. If multiple threads of a process all attach to the same channel 118, a single connection 114 may be shared between the threads. Alternatively, or additionally, multiple connections may be created that attached to the same channel 118. For example, each connection 114 may include attributes that define callback functions, a message buffer size, a maximum number buffers allowed for a connection, and triggering information. Triggering information may include a number of messages at which point messages may be passed to the channel 118, an amount of time between message passing events, flags for

enabling and/or disabling triggering methods, and the like. Exemplary callback functions include error notification callback functions, and callback functions that may allow for the reclamation of buffer space upon message delivery. This information may be stored in a connection attribute object.

5 [0021] Each connection 114 may be shared among the kernel module and an asynchronous messaging library. The connection 114 may include a queue, linked list or the like of messages to be sent to the channel 118. Connections 114 may be defined by connection objects, which may map to a file descriptor to allow programs, such as the producer thread 110, to send messages directly to the connection via the file descriptor, which may be a
 10 connection identifier. Alternatively, other mechanisms may be used to pass messages from programs to connections 114. Referring also to Figure 2, the connection object may include such elements as an array of asynchronous message headers for messages to be sent to the channel 118, flags for the asynchronous connection, pointers to the start, head 204, and tail 206 of the message queue 220, the size of the list of headers, a pointer 202 to the start of free space in the message queue, error status of the connection, events, a number of pending
 15 messages, a timer that maybe used, for example, to trigger message passing 216, a connection attributes object, described above, and the like. The connection may be shared among the kernel and the asynchronous messaging library, for example, by allowing only the kernel to adjust the head pointer while allowing only the asynchronous messaging library to edit the
 20 tail and free pointers. This sharing mechanism may provide lockless access to the send message queue so that new messages may be queued while previously queued messages are passed to the channel 118. The asynchronous message may include such information as the error status of the message, a pointer to the input/output vector (iov) array that contains the body of the message, the size of the iov array, and a handler used to identify the
 25 asynchronous messages. Handlers may be used in both the producer 110 and consumer 120, and a map function may be used to establish a one-to-one relationship between the handlers used each sides. Alternatively, or additionally, the connection object may also include a condition variable and mutex, or mutual exclusion lock, which together can be used to protect integrity of the message queue by blocking other threads from executing when the queue is
 30 full in a known manner.

[0022] Exemplary connection object data structures are shown in Figure 2. As illustrated, a connection object may include attributes 210, a queue of message headers 220, and a list of sent messages 240 that may be waiting for a reply. The attributes 210 may include a pointer to the start 202 of free headers 226 in the message queue 220 and pointers to the start 204 and

end 206 of the messages list in the queue 220, which together may define the sent 222 and queued messages 224. The connection object may also include call back functions 208, a condition variable 212, mutex 214, and trigger criteria 216. After messages are passed to the channel, the sent messages 242 may be added to a linked list 240 of sent message headers. Each message may include a header 260 that may include one or more flags 262, a handle 264, a call back function 266, a send buffer 268, a receive buffer 270, and any additional information 272 as needed or desired.

[0023] As described above, channel 118 and connection 114 may be joined by at the kernel level. To manage this union, a union object may be created. The union object may include such information as flags for the channel, access permission, the size of the kernel buffer, the maximum number of buffers allowed, an event to be set for notification and an associated handler for the event, the number of messages in the channel send queue, credential information, and the like.

[0024] Referring also to Figure 3, an exemplary flow chart showing the operations of a kernel module 340 and an asynchronous messaging library 320 are shown. In operation the kernel 340 may check the message queue 220 upon the triggering event at step 302. As it passes messages to a channel 114 at step 304, the kernel 340 may move the head pointer 204, which points to the start message queue, to the appropriate message in the message queue 220. The asynchronous messaging library 320 may then move the sent messages 242 to the list of sent messages 240 at step 306. A handle 264, such as the address of the message header, may be used to generate a receive identifier for the consumer program 120. After the received message has been processed, a reply function may be called to send back a reply or acknowledgement message at step 308. This reply message may include, for example, the receiver identifier. The kernel 340 may then notify the asynchronous messaging library 320 that messages have been delivered by sending a pulse, which may be a fixed-size non-blocking message, setting a flag 262 in the message header 260, and the like at step 310. The asynchronous messaging library 320 may then deliver the reply or acknowledgement to the producer program 110, such as via a callback function and the like, at step 312. The reply messages may be buffered in the producer program's 110 asynchronous messaging library until a certain criteria is met, at which point a kernel call may be made to deliver multiple messages to the producer program 110.

[0025] A diagram 400 showing the relationship between exemplary functions 402, 404, 406, 408, 410, 452, 454, 456, 458, 460 and 462 of an asynchronous messaging library 420 and associated kernel calls 422, 424, 426, 432, 434, and 436 of a kernel module 440 is shown in

Figure 4. The asynchronous messaging library 420 may include functions such as an AsyncMsg_ConnectAttach function 402, an AsyncMsg_ConnectDetach function 404, AsyncMsg_Put function 406, an AsyncMsg_Flush function 408, call back functions 410, an AsyncMsg_ChannelCreate function 452, an AsyncMsg_ChannelDestroy function 454, an AsyncMsg_Get function 456, events 458, an AsyncMsg_Malloc function 460, and an AsyncMsg_Free function 462. The kernel 440 may include calls such as ConnecAttach 422, ConnectDetach 424, MsgSendAsync 426, Event 428, ChannelProperties 430, ChannelCreate 432, ChannelDestroy 434, and MsgReceiveAsync 436.

[0026] Referring also to Figure 6, an exemplary flowchart depicting asynchronous message passing operations of a consumer program and producer program. The AsyncMsg_ChannelCreate function 452 may be called initially by a consumer program 120 to create an asynchronous messaging channel 118 at step 602. The function 452 may return a channel identifier, such as an integer, used to identify the channel. This channel identifier is then broadcast to the producer 110 at step 604, which may call the AsyncMsg_ConnectAttach function 402 to establish a connection 114 at step 606. The producer program 110 may allocate a buffer by calling the AsyncMsg_Malloc function 460, and calling the AsyncMsg_Put function 406 to send a message at step 608. Alternatively, or additionally, the producer program 110 may use its own buffers and using a call back function established with the connection to reclaim its buffers. Errors in the sending process may be handled using associated call back functions 410.

[0027] After a message has been passed to the channel 118 at step 608, the consumer program 120 may call the AsyncMsg_Get function 456 to retrieve the message at step 610. This function may be blocking or non-blocking, and may use the library's 420 internal buffer to receive a buffer, which may then be freed using the AsyncMsg_Free function 462. The internal buffer space may be configured to automatically align the buffers for page swapping and copy-on-write optimization for large messages, which are typical candidates for asynchronous transfer. When the message passing is complete, a channel 118 may be destroyed at step 612 via the AsyncMsg_ChannelDestroy function 454, which in turn calls the ChannelDestroy 434 kernel call to destroy the channel 118.

[0028] The consumer program 120 may establish various channel attributes when creating a channel 118. For example, the consumer program 120 may set the channel 118 as blocking or non-blocking, as well as the data transfer type. The channel 118 may transfer data in packets, or a stream of data may be used. Channel permissions, a buffer size, maximum number of buffers, notification event, and buffer allocation callback function may also be

established. Channel permissions may be used to control who can make a connection to the channel 118. The buffer allocation callback may allocate receive buffers if the consumer program 120 wants to use its own buffers and free buffers when the channel 118 is destroyed. The number of buffers to be actually allocated buffers may be passed to the consumer program 120. When freeing buffer space, an error code may be returned to the consumer program 120. Additionally, or alternatively, any of the channel attributes may be set by the consumer 120 via the AsyncMsg_ChannelCreate function 452. The AsyncMsg_ChannelCreate function 452 may pass these parameters in turn to a ChannelCreate kernel call 432. The kernel call may create a channel having the attributes established above.

[0029] As described above, the producer program 110 may connect to the channel 118 at step 606 using the AsyncMsg_ConnectAttach function 402, which may connect a process identified by a process identifier with a channel associated with its own identifier and return a connection identifier, such as an integer, to the producer program 402 by calling the ConenctAttach kernel call 424. Flags may also be set to create a connection 114 to the channel 118 using library buffers or program buffers, and to designate the connection 114 as blocking. If library buffers are to be used, the AsyncMsg_Malloc function 460 may be called to allocate the buffers. The buffers may be freed via the AsyncMsg_Free function 462. Error handling call back functions may also be established set for the connection 114. The producer may detach from the channel 118 at step 614 by calling the AsyncMsg_ConnectDetach function 404, which in turn calls the ConnectDetach 424 kernel call to destroy the connection, when message passing to the consumer 120 is finished. The queued messages in the connection may be discarded or, alternatively, sent to the consumer 120 before destroying the connection.

[0030] Once the connection 118 has been created, the AsyncMsg_Put function 406 may be called to add messages to the queue 220 at step 608. When putting a message on the queue 220, the user may override the default handler associated with a particular message, so that customized functionality may be implemented on a per message basis. When the trigger criteria has been met, or whenever the user calls the AsyncMsg_Flush function 408, messages may be passed to the channel 118 by the MsgSendAsync 426 kernel call, which may notify the kernel that some asynchronous messages have been queued for delivery. The queued messages may be delivered upon the receipt of the MsgReceiveAsync 436 kernel call, which may be called by the AsyncMsg_Get function 456. Alternatively, the messages may be transferred before the MsgReceiveAsync 436 kernel call is made. The MsgReceiveAsync

436 kernel call may be blocking or non-blocking. Additionally, an AsyncMsg_ConnectAttr function 464 may be provided to allow a user to set or retrieve connection and/or channel properties for connections or channels specified by a connection or channel identifier. The properties may be set or retrieved via a ChannelProperties kernel call 230.

5 [0031] Optionally, a synchronous communication path may be provided in conjunction with an asynchronous channel 118 and connection 114 pair to provide for synchronous communications between the threads in parallel to the asynchronous message passing. The synchronous communication path may be provided using any known manner. For example, an exemplary model 500 for using asynchronous messages to deliver large batch data, and to
10 use synchronous message to deliver control and status events is shown in Figure 5. The synchronous communications may include events and synchronous messages that require in-order processing. A producer process 510 may collect data from some input devices, such as sensors, hard disks, and the like, buffers the data 516, and send it to another consumer process 520 for processing. In order to obtain a high throughput, the system 500 may use the
15 asynchronous messaging to deliver the data 516 from producer program 510 to consumer because of its 'batch delivery' feature that greatly reduces the overhead. In the consumer process 520, an asynchronous messaging channel 518 may be created to receive the data 516. A separate synchronous channel (event channel 538) may be created to receive synchronous messages and pulses. A worker thread 540 may be blocked on the synchronous channel 538
20 to wait for service requests and events.

[0032] In the producer program 510, an asynchronous connection 514 is established to the consumer program's 520 asynchronous message channel 518 to send data 516. A synchronous channel 528 may also be created to receive error events 526, interrupt events 532, which may be generated by input devices 534, and synchronous messages 536. A
25 worker thread 530 may be blocked on the channel 528 to receive the error events 526 and interrupt events 532 and messages. The consumer process 520 may also create a connection 524 to this channel 528 to deliver various events 536, such as commands, acknowledgements, status events (*e.g.*, server is idle), and the like.

[0033] While various embodiments of the invention have been described, it will be apparent
30 to those of ordinary skill in the art that many more embodiments and implementations are possible within the scope of the invention. Accordingly, the invention is not to be restricted except in light of the attached claims and their equivalents.

CLAIMS

We claim:

1. A method of asynchronously passing messages in an operating system, the method comprising:
 - 5 generating, by a consumer process in communication with the operating system, a channel for receiving asynchronous messages;
 - generating, by a producer process, a connection for sending asynchronous messages;
 - sending messages, by the producer process, to the connection;
 - joining the channel and the connection;
 - 10 determining the occurrence of a triggering condition; and
 - passing messages from the connection to the channel after determining the occurrence of the triggering condition.
2. The method of claim 1, where the operating system includes a kernel module, the method further comprising joining the channel and the connection within the kernel module.
- 15 3. The method of claim 2, where messages are passed, by the kernel module, directly from a sender memory space associated with the producer process to a receiver memory space associated with the consumer process
4. The method of claim 2, where sending messages to the connection comprises adding the messages to a send messages queue.
- 20 5. The method of claim 4, where the send messages queue is shared between the kernel module and the producer process.
6. The method of claim 5, where the queue is shared in a lockless manner such that the producer process may send messages to the connection while the kernel module is passing messages to the channel.
- 25 7. The method of claim 6, where the lockless method comprises:
 - allowing only the kernel module to modify a head pointer to the beginning of a list of messages to be sent from the send messages queue; and

allowing only the producer program to modify a tail pointer to the end of the list of messages to be sent from of the send messages queue and a free pointer to the start of free buffer space associated with the send messages queue.

8. The method of claim 1, further comprising:

5 establishing a synchronous communication path between the producer process and the consumer process.

9. The method of claim 1, where the triggering condition is a predetermined number of messages in send messages queue.

10. The method of claim 1, where the triggering condition is an expiration of a time
10 period.

11. An operating system operable to asynchronously pass messages from a producer program to a consumer program, the operating system comprising:

15 an asynchronous messaging library in communication with the producer program and the consumer program and operable to provide a channel to the consumer program and a connection to the producer program, the connection operable to receive messages from the producer program and add them to a messages queue, the channel operable to queue
15 messages for the consumer program; and

a kernel module in communication with the asynchronous messaging library and operable to pass messages from the producer program to the consumer program.

20 12. The operating system of claim 11, where messages are queued in the connection until a triggering event occurs.

13. The operating system of claim 12, where the triggering event is a predetermined number of messages in send messages queue

25 14. The operating system of claim 12, where the triggering condition is an expiration of a time period.

15. The operating system of claim 11, where the messages are passed directly from a sender memory space associated with the producer process to a receiver memory space associated with the consumer process by the kernel module.

16. The operating system of claim 11, the messages queue is shared in a lockless manner such that the producer process may send messages to the connection while the kernel module is passing messages to the channel.

17. An asynchronous message passing architecture in communication with an operating system including a kernel module, the asynchronous message passing architecture operable to asynchronously pass messages from a sender to a receiver, the operating system comprising:

an asynchronous messaging library in communication with the sender and the receiver and operable to provide a channel to the receiver and a connection to the sender, the connection operable to receive messages from the sender and add them to a messages queue, the channel operable to queue messages for the receiver; and

where the kernel module is in communication with the asynchronous messaging library and operable to pass messages directly from the memory space associated with the sender to memory space associated with the receiver upon the occurrence of a triggering event.

18. The asynchronous message passing architecture of claim 17, the messages queue is shared in a lockless manner such that the sender may send messages to the connection while the kernel module is passing messages to the receiver.

19. The asynchronous message passing architecture of claim 18, where the lockless manner comprises:

allowing only the kernel module to modify a head pointer to the beginning of a list of messages to be sent from the messages queue.

20. The asynchronous message passing architecture of claim 19, where the lockless manner further comprises:

allowing only the asynchronous messaging library to modify a tail pointer to the end of the list of messages to be sent from the messages queue and a free pointer to the start of free buffer space associated with the messages queue.

21. A method for asynchronously passing messages between processes of a processing system, the method comprising:

generating a channel that facilitates receipt of asynchronous messages by a first process;

5 generating a connection that facilitates sending of asynchronous messages by a second process;

queuing a plurality of asynchronous messages at the connection using the second process;

detecting a triggering condition; and

10 transferring a batch of the plurality of asynchronous messages from the connection to the channel in response to the occurrence of the triggering condition.

22. The method of claim 1, where the processing system comprises a kernel, and where the method comprises transferring the batch of the plurality of asynchronous messages from the connection to the channel through one or more kernel processes.

15 23. The method of claim 1 or 2, where the processing system comprises an asynchronous messaging library, and where the method comprises generation of the channel and the connection using one or more calls to the asynchronous messaging library.

24. The method of claim 1, where the method further comprises allowing generally concurrent queuing and transfer of messages at the connection.

20 25. The method of claim 1, where the processing system comprises a kernel and an asynchronous messaging library, and where the method comprises:

queuing the plurality of asynchronous messages in a message queue associated with the connection;

25 providing a head pointer to the beginning of a list of messages that are to be transferred from the message queue;

providing a tail pointer to the end of the list of messages that are to be transferred from the message queue;

providing a free pointer to the start of free buffer space of the message queue;

30 restricting modification of a head pointer to processes of the kernel module, where the head pointer provides an index to the beginning of a list of messages that are to be sent from

the message queue; and

restricting modification of a tail pointer to one or more operations executed by the second process, where the tail pointer provides an index to the end of the list of messages that are to be sent from of the message queue.

5 26. The method of claim 25, where the method further comprises providing a free pointer serving as an index to the start of free buffer space of the message queue. .

27. The method of claim 1, further comprising:
establishing a synchronous communication path between the first process and the second process.

10 28. The method of claim 1, where the triggering condition comprises the queuing of a predetermined number of messages in the message queue.

29. The method of claim 1, where the triggering condition comprises an expiration of a predetermined time period.

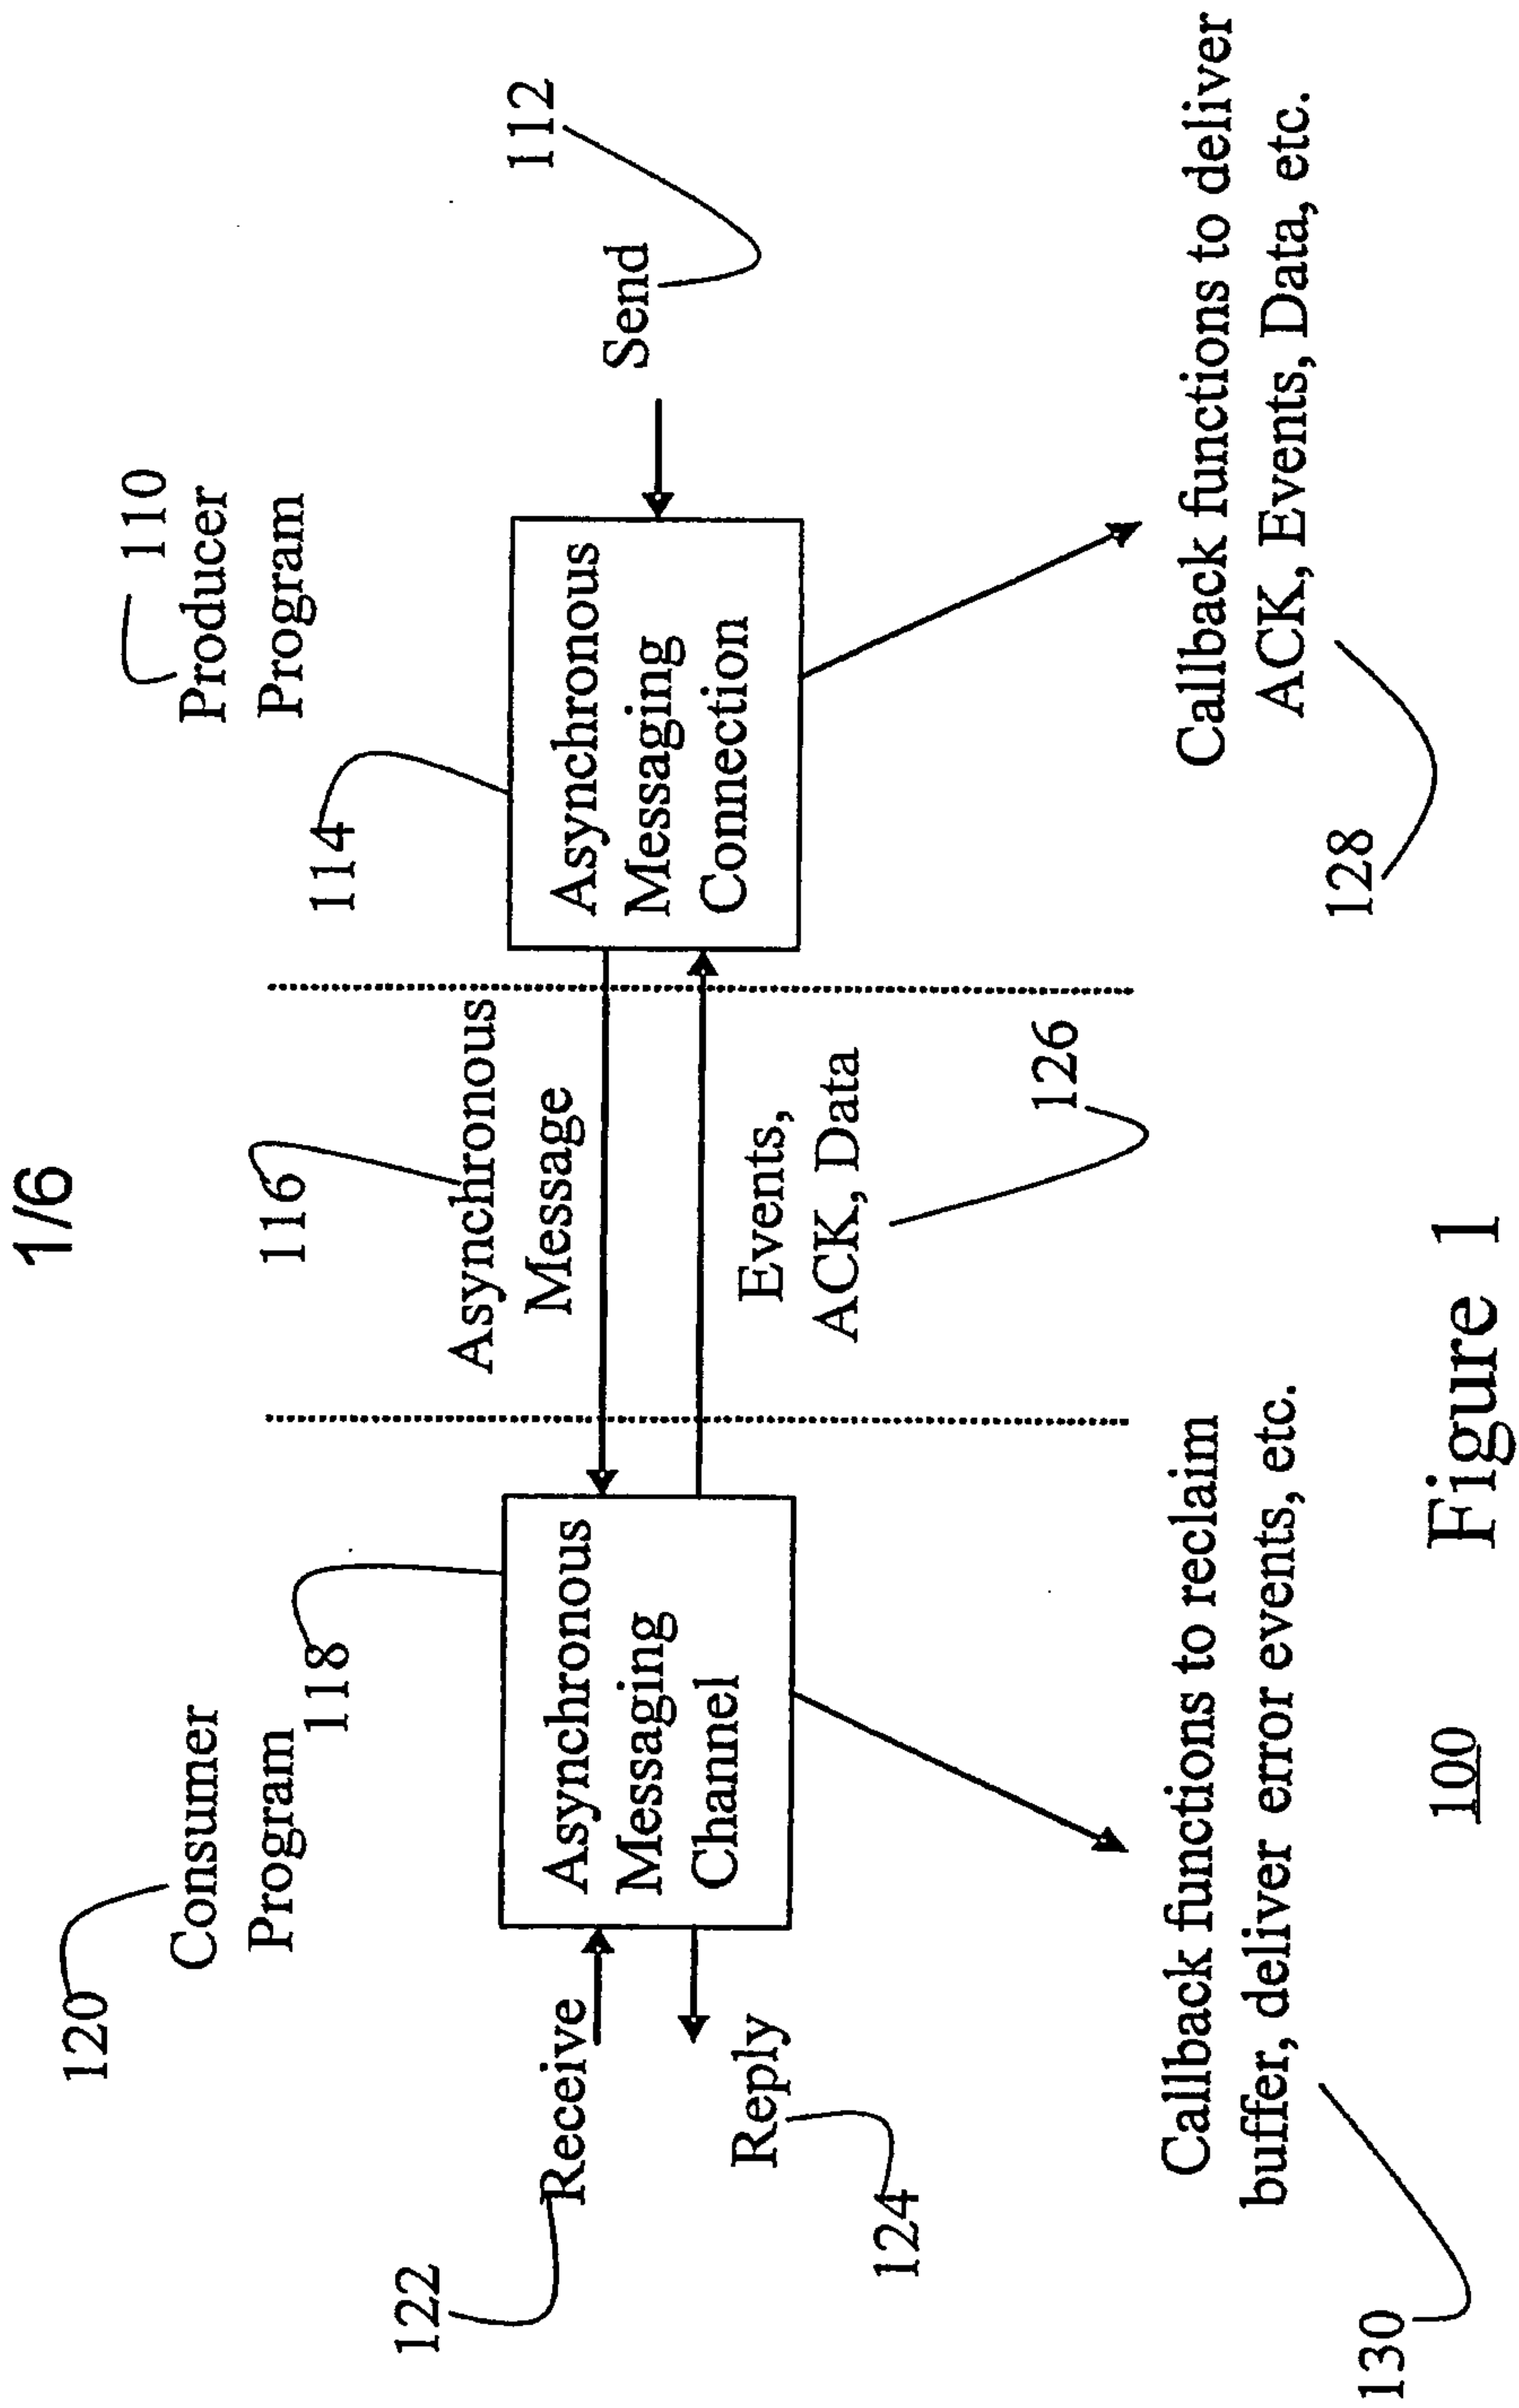


Figure 1

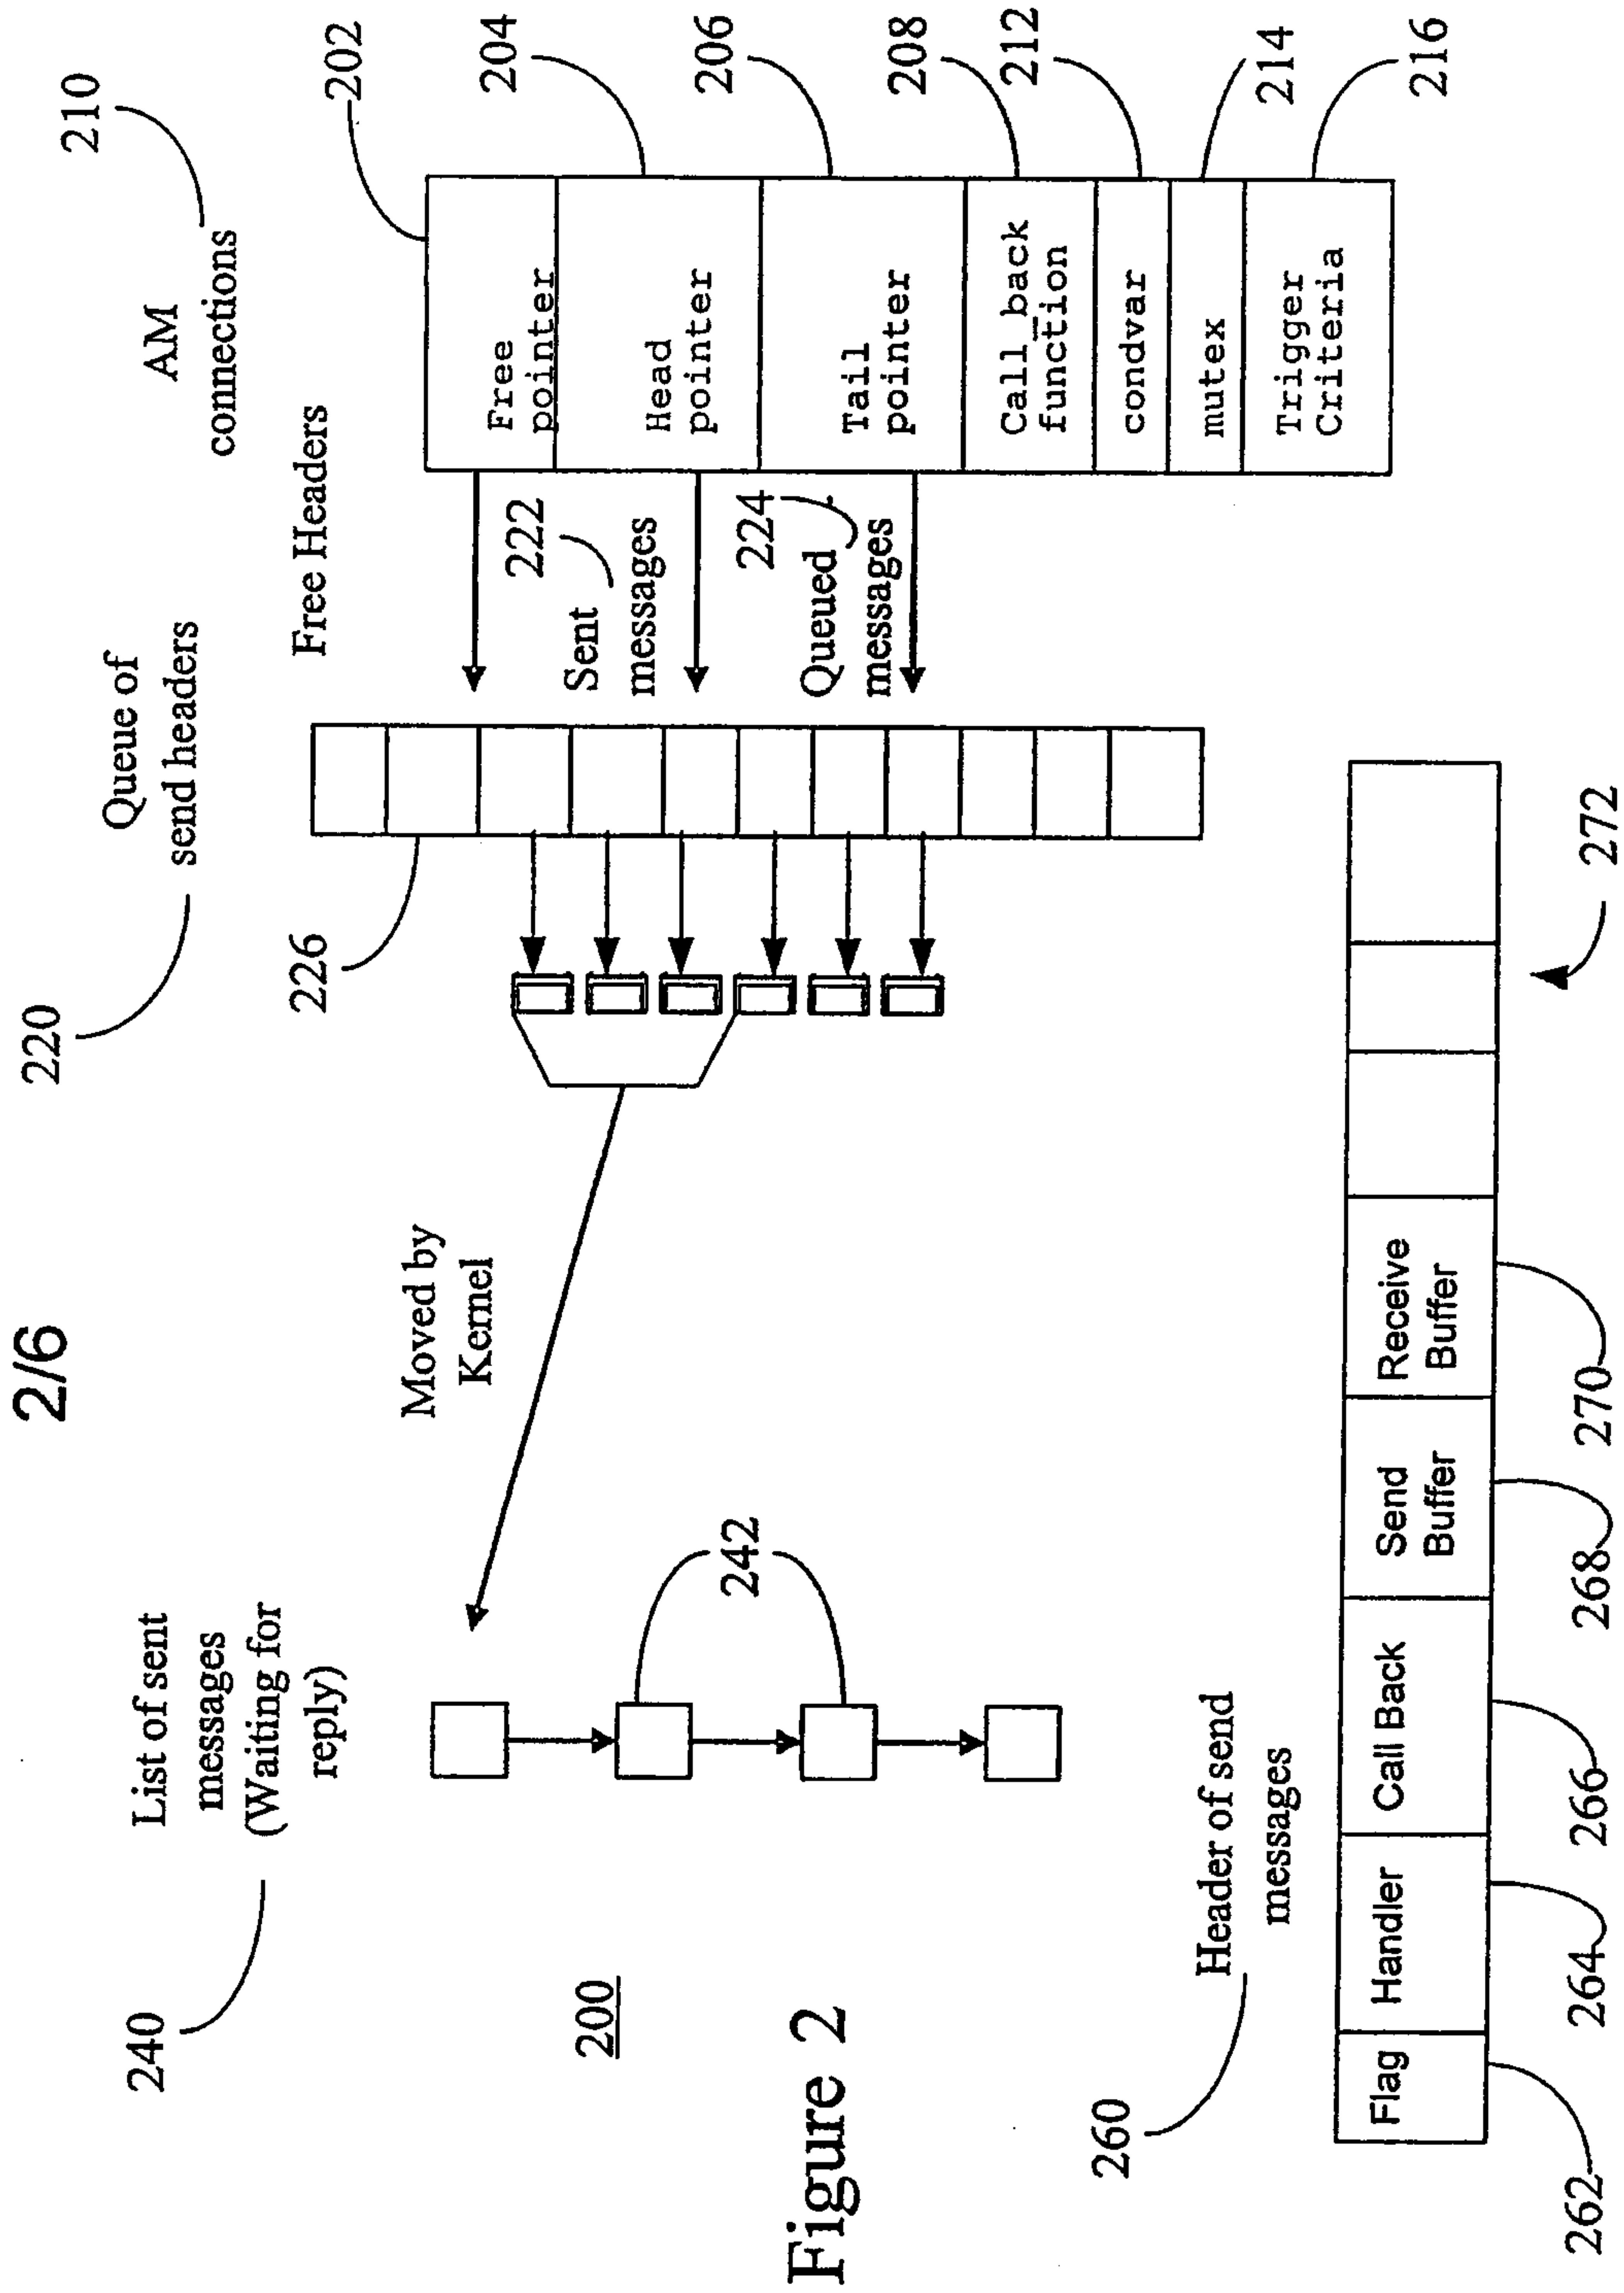


Figure 2

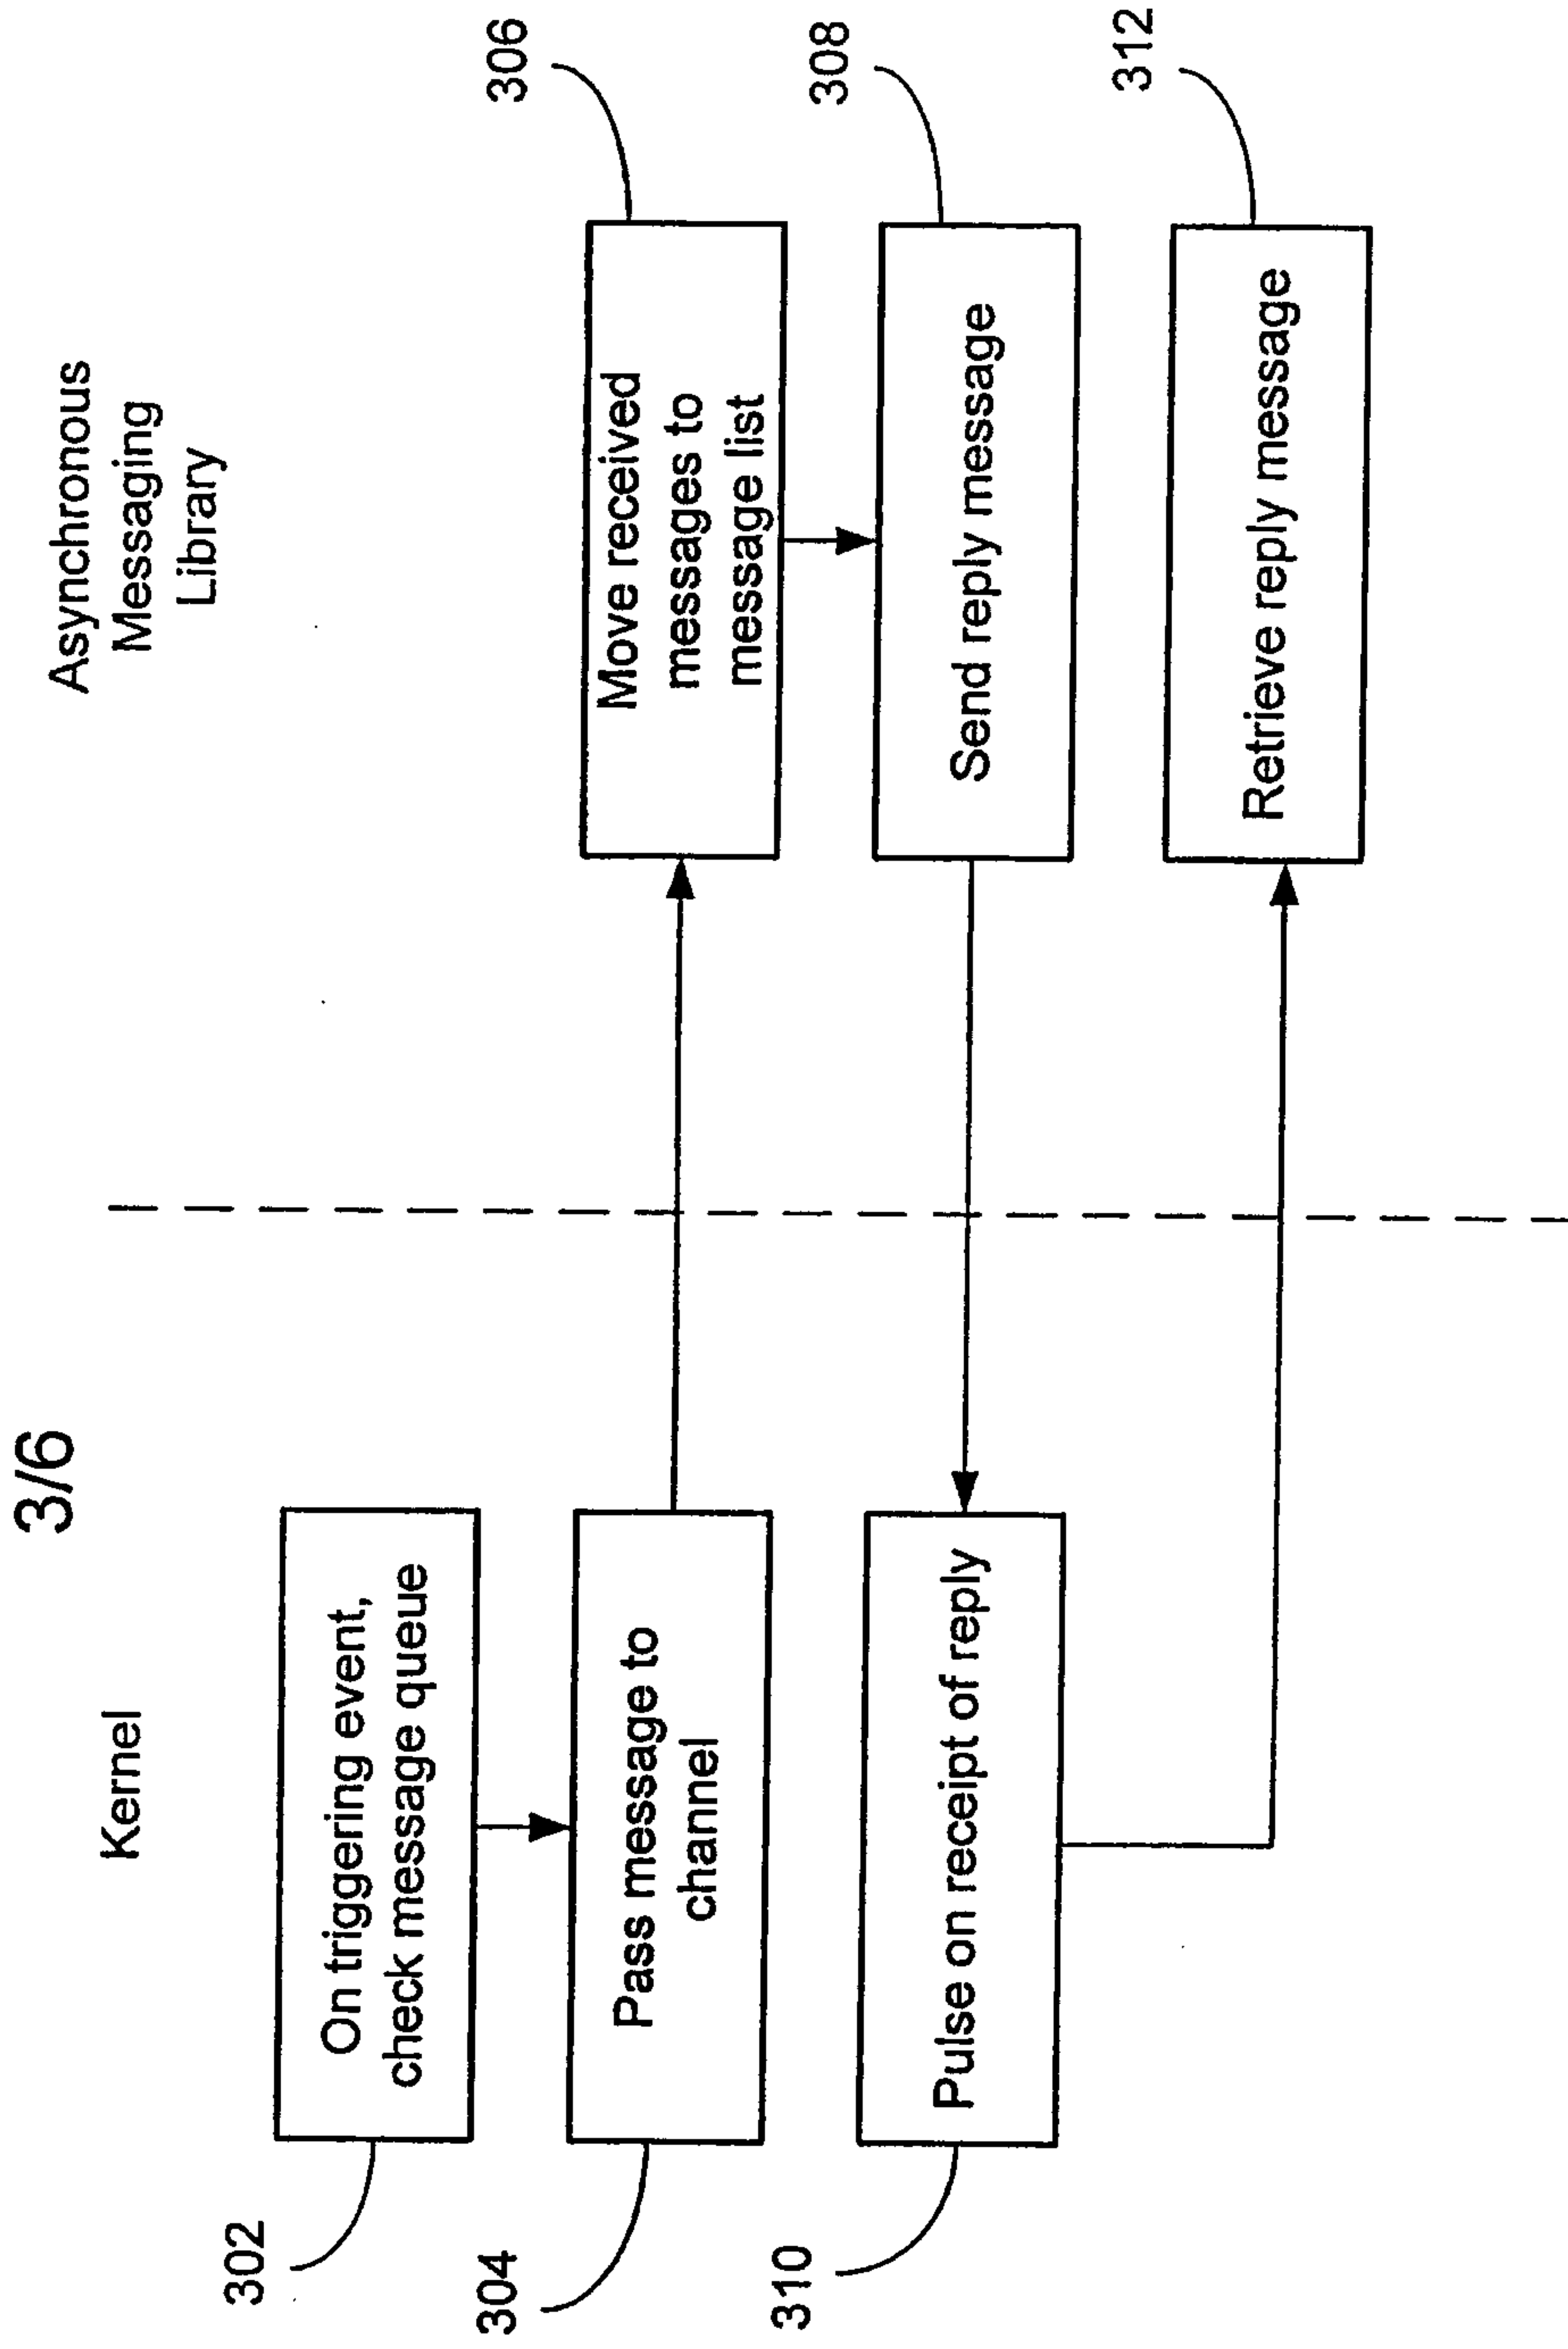


Figure 3

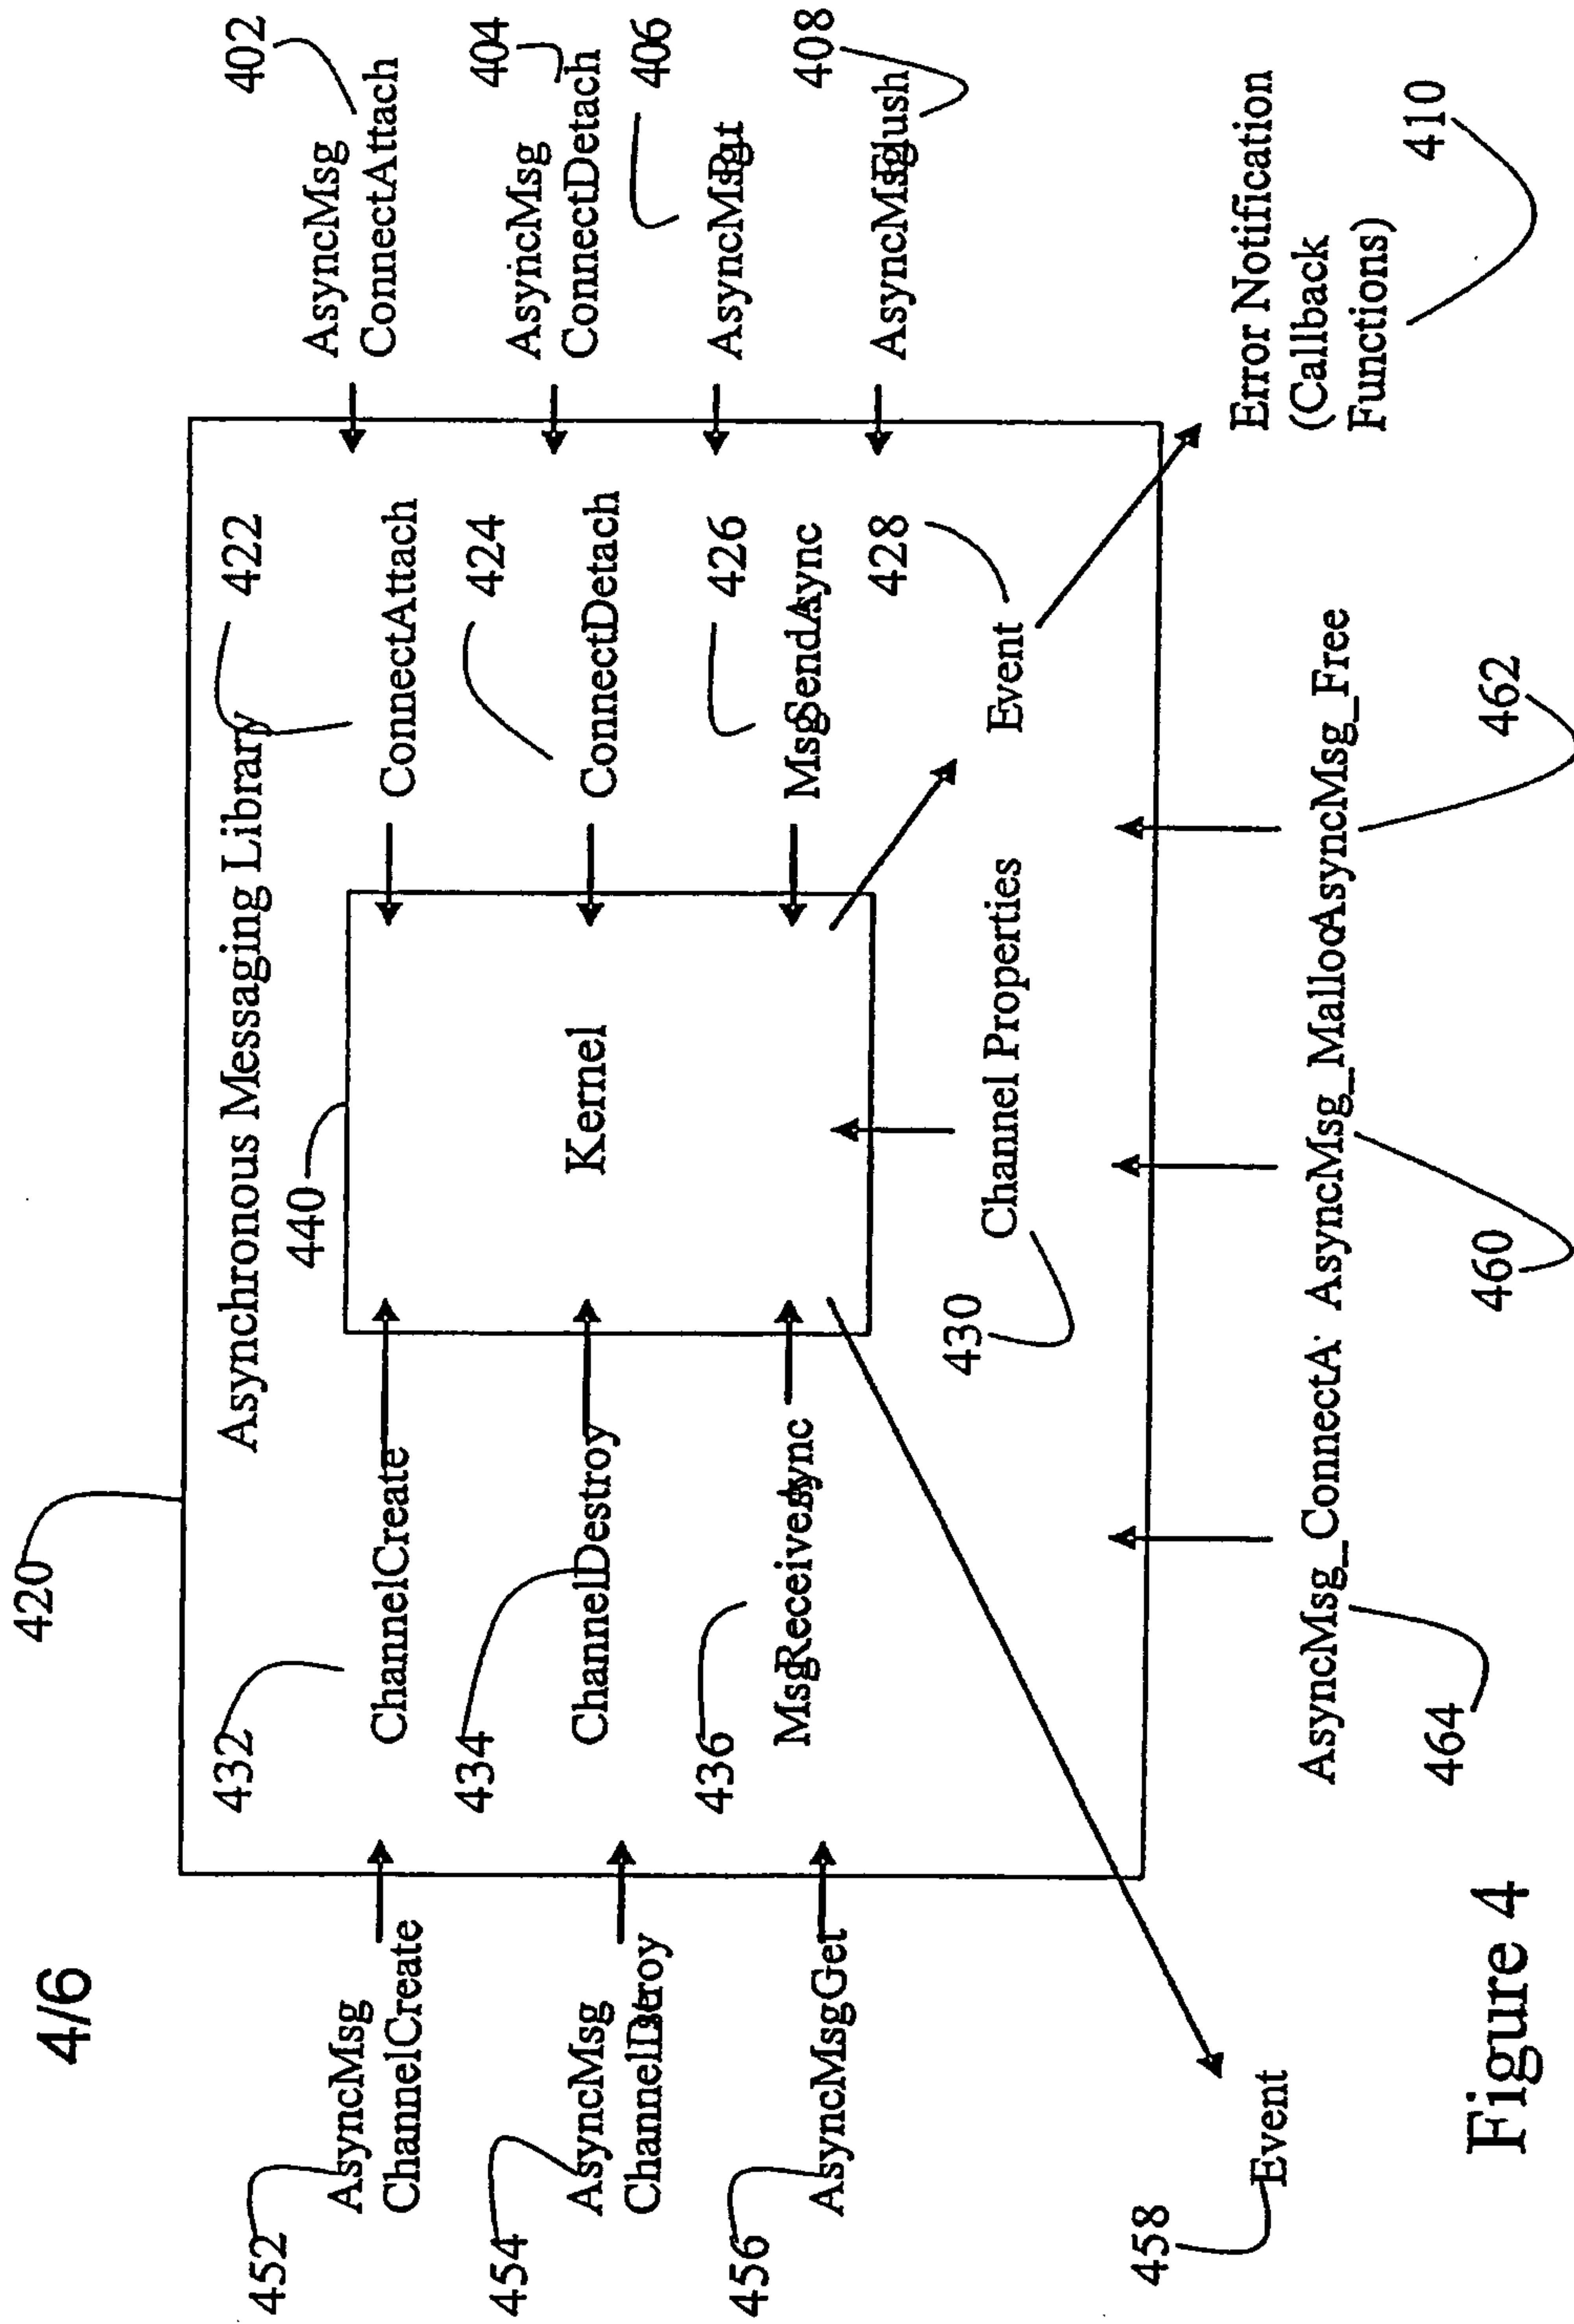


Figure 4

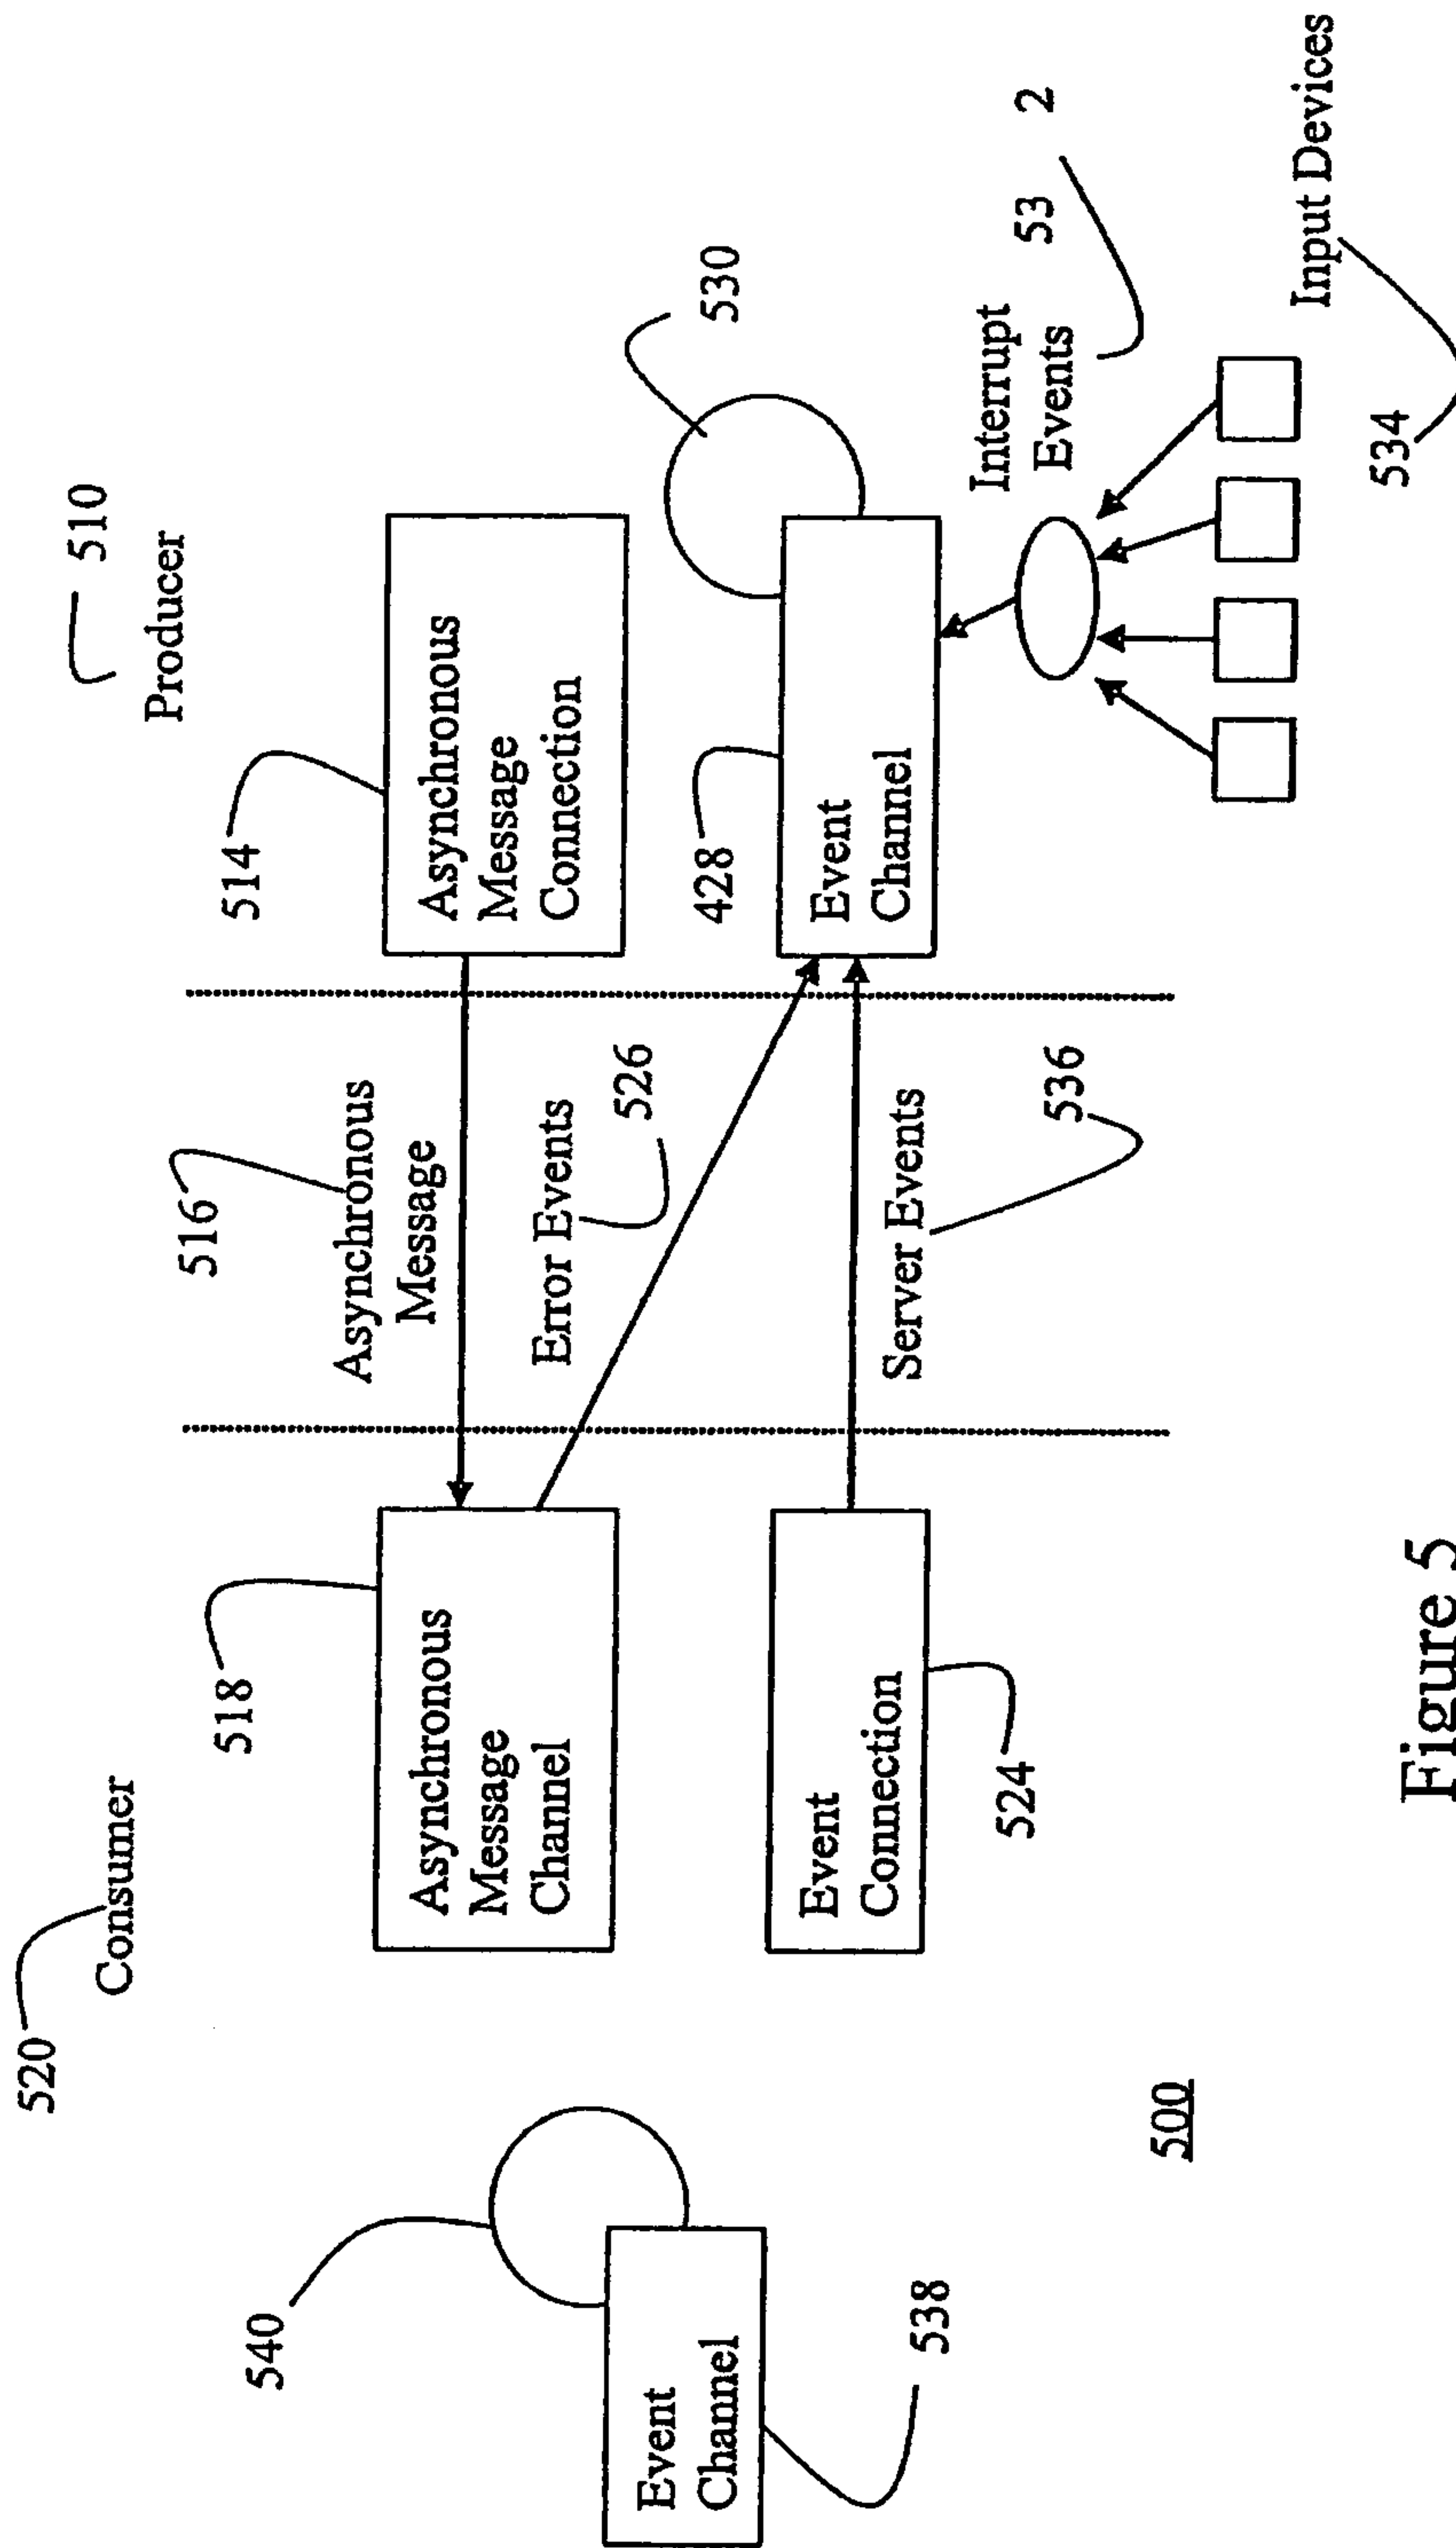


Figure 5

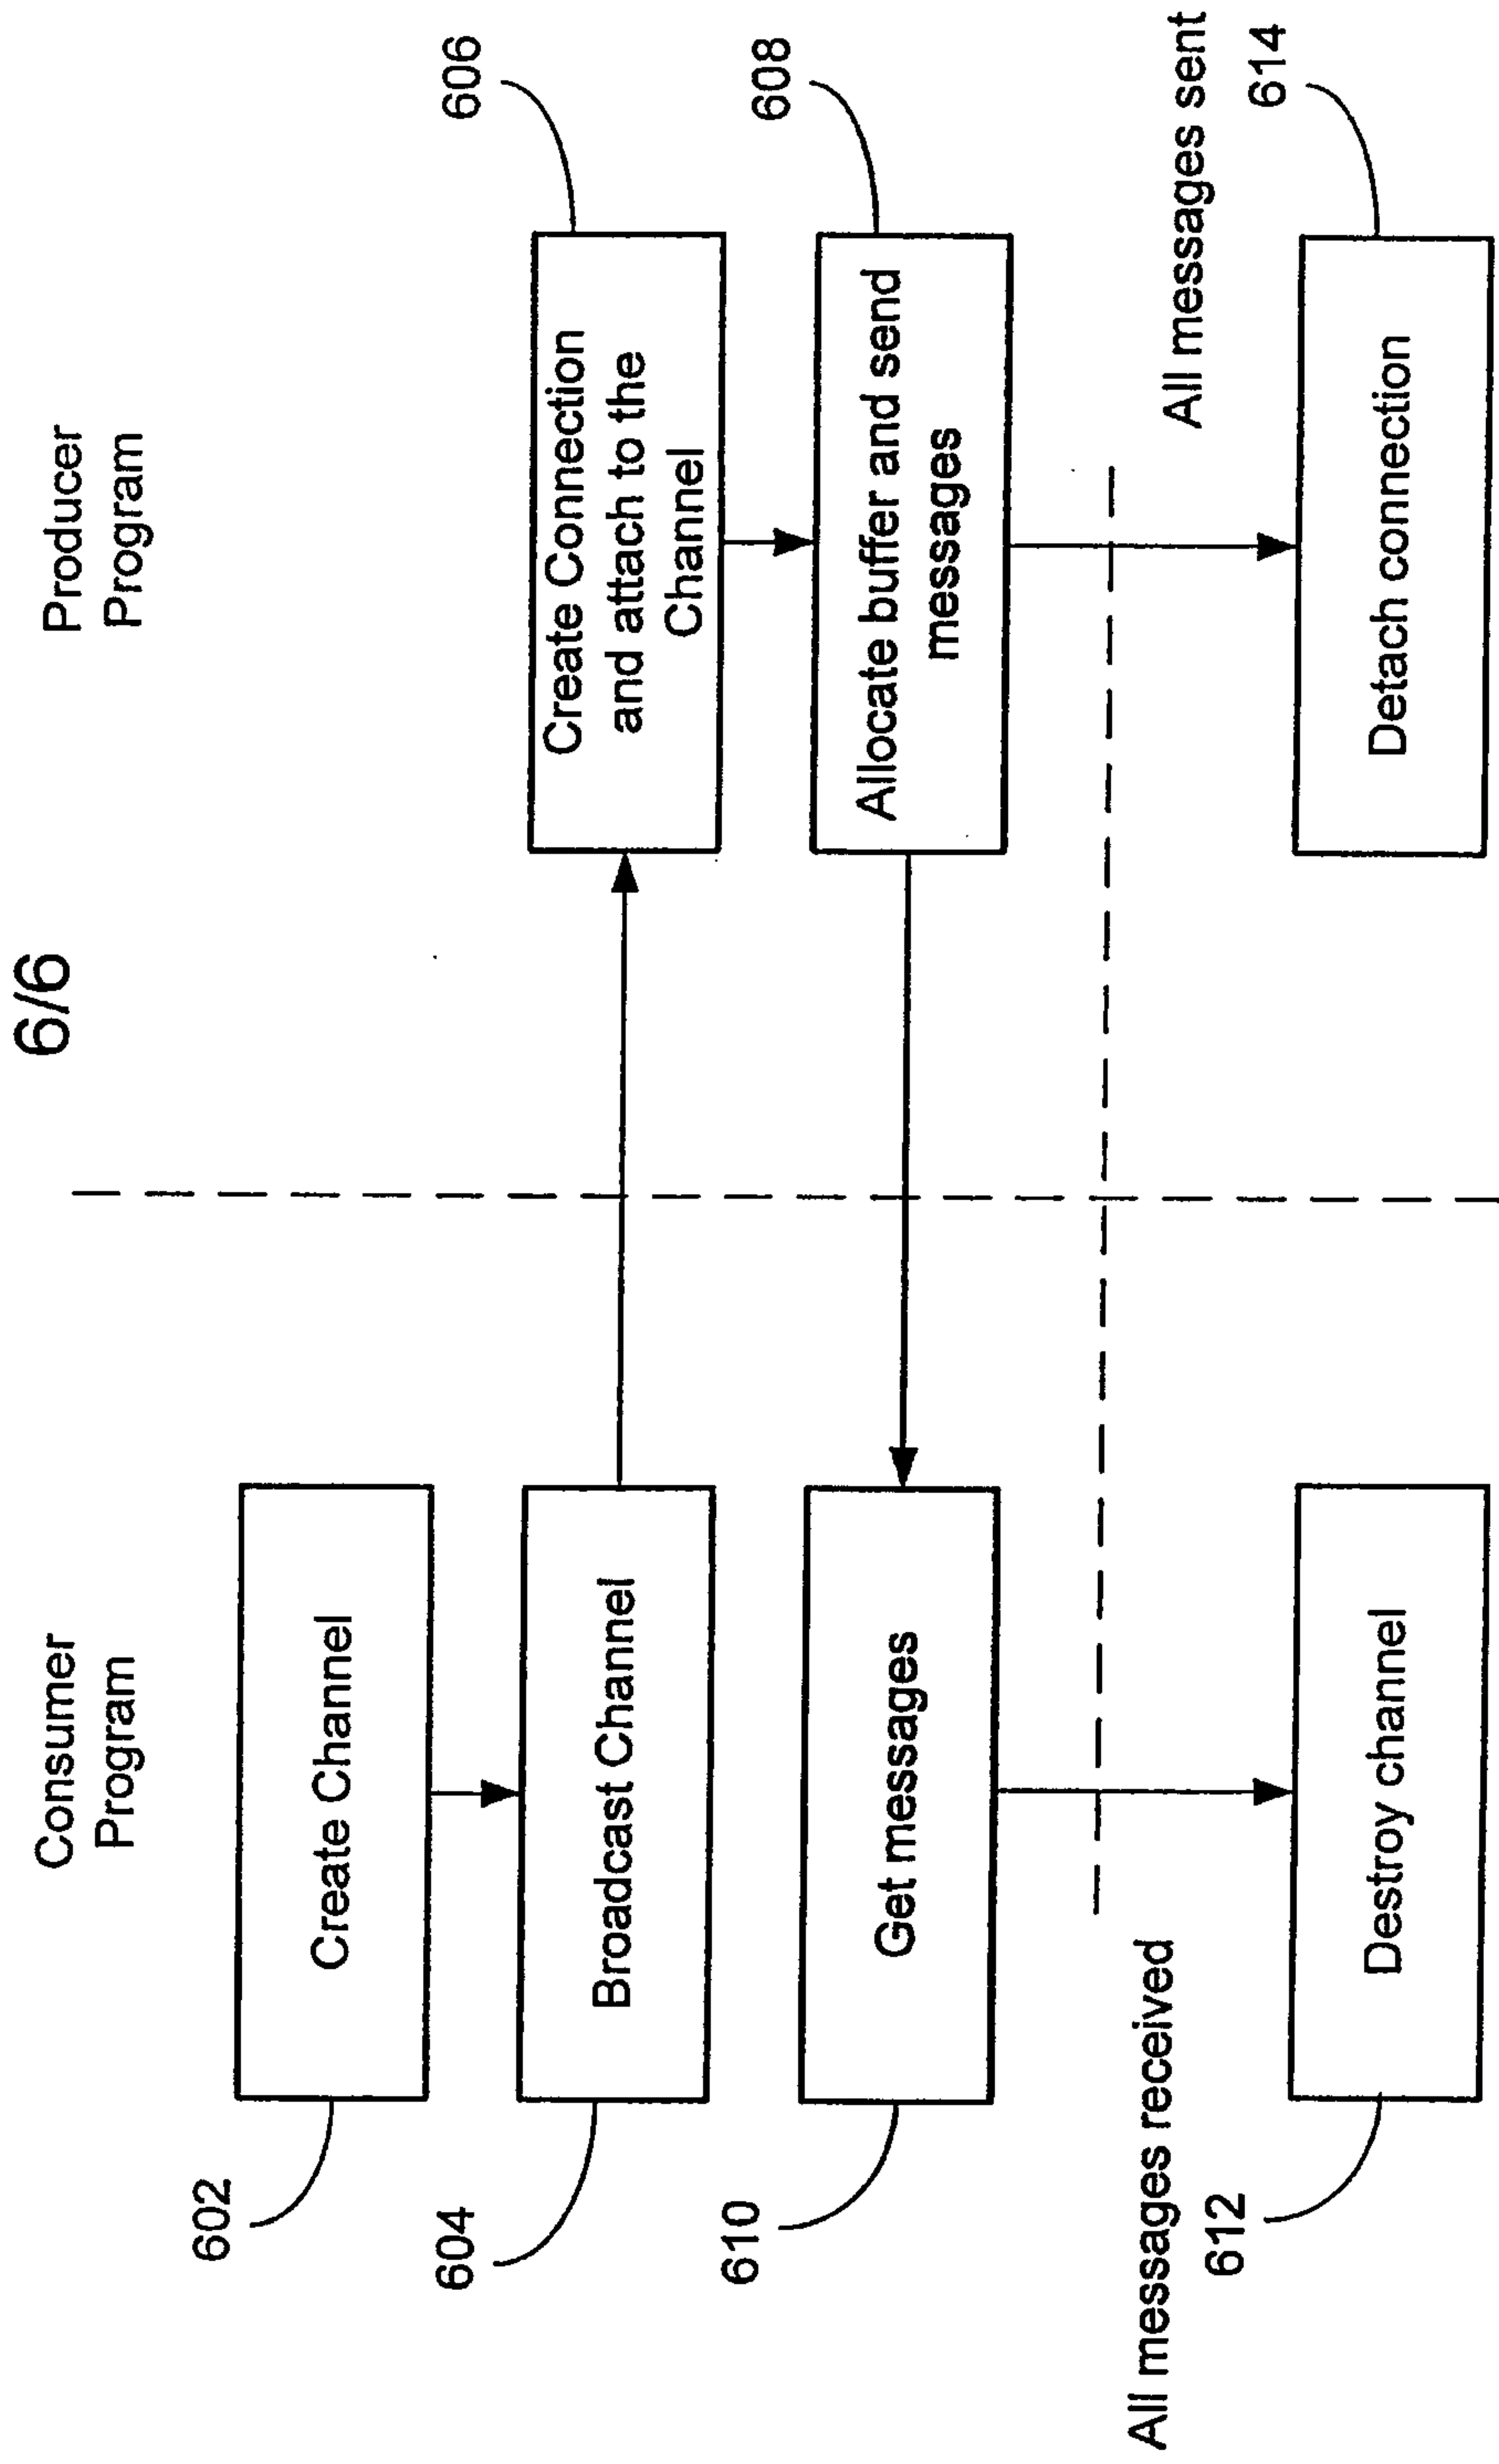


Figure 6

