

10 Deadlock Revisited: Circular Structures

Some networks cannot be analysed using a client-server labelling

- a ring of processes invariably deadlocks
- plausible solutions are presented and then discounted
- deadlock avoidance strategies are described
- an argument is developed that shows that the final system will not deadlock

In previous chapters the concept of a client-server design pattern has been introduced and it has then been applied to a number of simple examples. The primary requirement of the pattern is that any resulting network should not contain any circuits of client and server labels. Needless to say, that if we have a ring of processes then a circuit is inevitable. Hence, we shall investigate a ring of processes to explore how, even though the client-server pattern cannot be applied, we can construct a system that is deadlock free.

The aim of the application is to construct a message passing structure from one node to another by providing a set of message passing elements that connect each node to the next. The simplest way of doing this is to create a ring of message passing nodes to which message sender and receiver processes are attached. Figure 10-1 shows the basic structure with a client-server labelling that demonstrates immediately that deadlock will occur, even ignoring the effect of the Sender and Receiver processes. It is obvious that the set of channels that connect the Ring Element processes has to be broken in some way. Deadlock will occur trivially when every Ring Element attempts to either input or output a message at the same time. Thus we have to find a way of breaking the ring

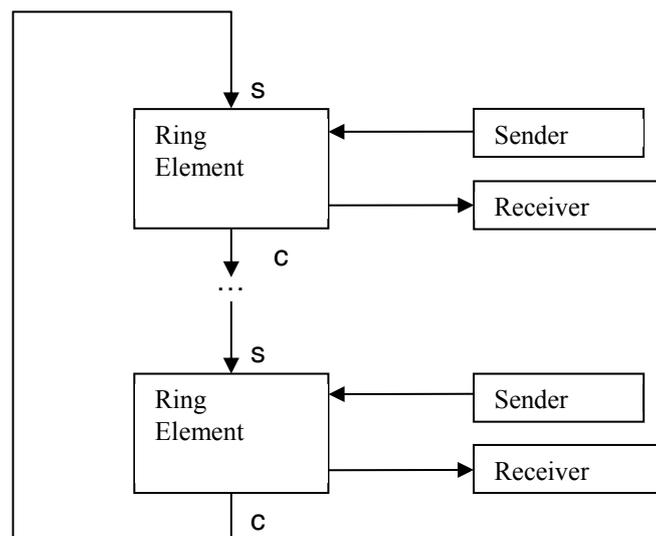


Figure 10-1 The Basic Message passing System

10.1 A First Sensible Attempt

The simplest way of breaking the ring of channels connecting the Ring Element processes is to add another element to the ring which does the input and output operations in a different order to that undertaken by the Ring Elements. This will mean that there is at least one element on the ring that is always able to undertake an input operation if all the other Ring Elements are trying to output to the ring. This is shown in Figure 10-2.

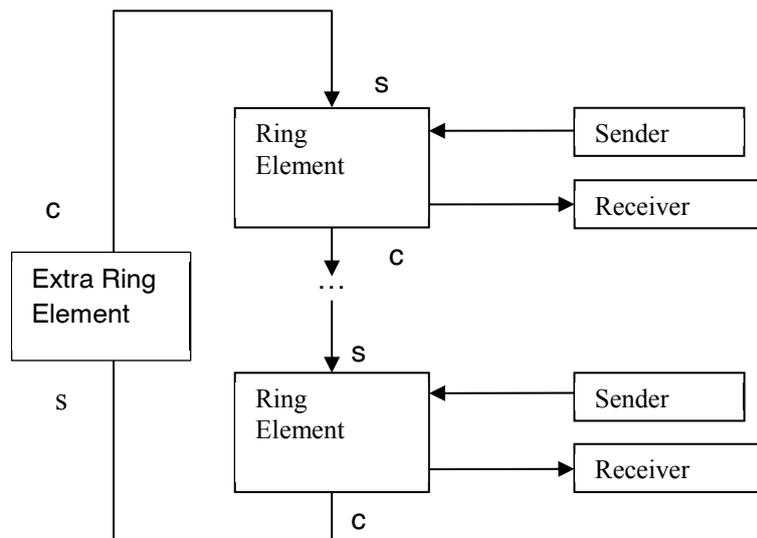


Figure 10-2 Adding the Extra Ring Element

The client-server labelling has not altered and still indicates a problem but we now know that the Extra Ring Element undertakes its input – output operations in a different order to the Ring Elements. The behaviour of a Ring Element is shown in Listing 10-1.

```

10 class RingElementv0 implements CSProcess {
11
12     def ChannelInput fromRing
13     def ChannelOutput toRing
14     def ChannelInput fromLocal
15     def ChannelOutput toLocal
16     def int element
17
18     void run () {
19         def RING = 0
20         def LOCAL= 1
21         def ringAlt = new ALT ( [ fromRing, fromLocal ] )
22         while (true) {
23             def index = ringAlt.priSelect()
24             switch (index) {

```

```
25     case RING:
26         def packet = (RingPacket) fromRing.read()
27         if ( packet.destination == element )
28             toLocal.write(packet)
29         else
30             toRing.write (packet)
31         break
32     case LOCAL:
33         def packet = (RingPacket) fromLocal.read()
34         toRing.write (packet)
35         break
36     }
37 }
38 }
39 }
```

Listing 10-1 The Ring Element Process Behaviour (Print Statements Omitted)

A Ring Element alternates over inputs from the ring and from its local sender process {21}. In a loop {22} it determines the enabled alternative, giving priority to inputs from the ring {23}. An enabled input from the ring is read {26} as a `RingPacket`, and if the message is for this element, it is written to the local receiver {28}, otherwise it is written to the ring {30} for onward transmission. If the enabled alternative is an input from the local sender then it is read {33} and written to the ring {34}.



"I studied English for 16 years but...
...I finally learned to speak it in just six lessons"
Jane, Chinese architect

ENGLISH OUT THERE

Click to hear me talking before and after my unique course download



The behaviour of the Extra Ring Element is shown in Listing 10-2.

```
10 class ExtraElement implements CSProcess {
11
12     def ChannelInput fromRing
13     def ChannelOutput toRing
14
15     void run () {
16         def packet = new RingPacket (source:-1, destination:-1,
17                                     value:-1, full: false )
18         while (true) {
19             toRing.write( packet )
20             packet = (RingPacket) fromRing.read()
21         }
22     }
23 }
```

Listing 10-2 The Behaviour of the Extra Ring Element (print Statements Omitted)

Given that the Ring Elements initially input from the ring, or local Sender process then the Extra Ring Element has to output a packet, so that a Ring Element has a packet to read. A `RingPacket` is defined {16, 17} which is then written to the ring {19}. Thereafter the process simply reads a `RingPacket` from the ring {20} and then outputs it to the ring {19}. The empty packet will continue to circulate forever.

10.1.1 Evaluation

The accompanying examples package contains a version of this first attempt `c10.examples.Runv0.groovy` that has print statements inserted within it to show the effect of this solution formulation. The user is able to indicate the number of nodes in the network when the system is executed. The messages received by each receiver process are displayed using a `GConsole` process. A network with 4 nodes will additionally have the extra node numbered as node 0. The output changes with each execution of the network but seldom terminates, though on occasion it has terminated. In a 4 node system each node should receive 3 messages from each of the other nodes, that is, each should receive 9 messages. Typically, no node receives all its messages and some nodes receive no messages. Inspection of the system console print messages indicates that the extra node does indeed output its empty packet and that this is read by the next node in the ring. This means that the other nodes have no input on their ring input channels and so they read a message from their local sender process. The sender processes attempt to send their messages as quickly as possible. This then has the effect of sending many messages on to the ring, which at some stage may deadlock when every node, including the extra node attempt to undertake an input or an output operation. Just when this occurs depends on the particular execution sequence. It is obvious that we have to find a way of managing the number of messages in the ring.

10.2 An Improvement

A simple improvement can be seen quite easily. If a node sends a message to another node on behalf of its local node then the receiving node undertakes to send a message back to the original source that the message has been read. This means that each node can only ever have one packet on the ring at any one time. On the first part of its journey it contains the desired message and then once it has been processed by the destination node it is returned with an empty flag. The definition of the `RingPacket` used to send messages around the system is shown in Listing 10-3. The property `source` {12} gives the number of the node that sent the message and `destination` {13} is the node to which it is to be sent. The actual message is contained in the property `value` {14} and the Boolean `full` {15} indicates whether the packet contains a message or is just an empty packet. A `toString` method is provided to enable printing of the packet on the console window and also on the `GConsole` processes.

```
10 class RingPacket implements Serializable, JCSPCopy {
11
12     def int source
13     def int destination
14     def int value
15     def boolean full
16
17     def copy () {
18         def p = new RingPacket ( source: this.source,
19                                 destination: this.destination,
20                                 value: this.value,
21                                 full: this.full)
22         return p
23     }
24
25     def String toString () {
26         def s = "Packet [ s: ${source}, d: ${destination}, v: ${value}, f: ${full} ] "
27         return s
28     }
29 }
```

Listing 10-3 The `RingPacket` Class definition

A first running, `c10.examples.Runv1.groovy`, of this modification typically results in even worse performance than the initial version. On reflection this is obvious. The Extra Ring Element process still outputs an empty packet onto the ring and thus there will be no space for the messages to rotate around the ring. The solution is to modify the Extra Ring Element process so that it provides an empty space on the ring of nodes so that a communication can take place. This behaviour is shown in Listing 10-4. This does mean that the Extra Ring Element has to read {17} and then write {18} a packet, the same as all the other nodes.

```
10 class ExtraElementv1 implements CSProcess {
11
12     def ChannelInput fromRing
13     def ChannelOutput toRing
14
15     void run () {
16         while (true) {
17             def packet = (RingPacket) fromRing.read()
18             toRing.write( packet )
19         }
20     }
21 }
```

Listing 10-4 The Modified Behaviour of The Extra Ring Element (Print Statements Omitted)

The execution of `c10.examples.Runv1a.groovy` now results in the proper operation of the network with all Receivers getting and outputting the expected messages. The solution does however have some limitations in that only one packet is ever in circulation for each node as shown in the behaviour given in Listing 10-5.

The solution uses an alternative with pre-conditions to control the input of messages either from the ring or from the local sender {19–24}. Initially, messages can be input either from the ring or from the local sender {23, 24}. The `index` of the enabled alternative is determined using a `select` method call {26}.

Excellent Economics and Business programmes at:



university of
 groningen



“The perfect start
of a successful,
international career.”

CLICK HERE
to discover why both socially
and academically the University
of Groningen is one of the best
places for a student to be

www.rug.nl/feb/education



For messages read from the ring {29}, it is first determined whether the message has its destination at this element {30}. It is then necessary to determine whether or not the packet is full {31}.

```
10  class RingElementv1 implements CProcess {
11
12  def ChannelInput fromRing
13  def ChannelOutput toRing
14  def ChannelInput fromLocal
15  def ChannelOutput toLocal
16  def int element
17
18  void run () {
19      def RING = 0
20      def LOCAL= 1
21      def ringAlt = new ALT ( [ fromRing, fromLocal ] )
22      def preCon = new boolean[2]
23      preCon[RING] = true
24      preCon[LOCAL] = true
25      while (true) {
26          def index = ringAlt.select(preCon)
27          switch (index) {
28              case RING:
29              def packet = (RingPacket) fromRing.read()
30              if ( packet.destination == element ) {
31                  if ( packet.full ) {
32                      toLocal.write(packet.copy())
33                      packet.destination = packet.source
34                      packet.source = element
35                      packet.full = false
36                      toRing.write(packet)
37                  }
38                  else
39                      preCon[LOCAL] = true
40              }
41              else
42                  toRing.write (packet)
43              break
44              case LOCAL:
45              def packet = (RingPacket) fromLocal.read()
46              toRing.write (packet)
47              preCon[LOCAL] = false
48              break
49          }
50      }
51  }
52 }
```

Listing 10-5 The Ring Element That Expects A Returned Empty Packet (Print Statements Omitted)

If the packet is full then we can write a `copy` of the packet to the local receiver process {32}. After which we can update the content of the packet for its return journey to its originating node, because a copy was written to the local receiver process. The `destination` and `source` properties of the packet are updated accordingly {33, 34}, the `packet.full` indication is set `false` {35} and the revised packet written to the ring {36}. If the received packet is not `full` {38} then this is a returned packet and the ring element process can now input a message from its local sender, requiring an update to the associated pre-condition {39}.

If the initial packet was not destined for this node element {41} then it is simply written to the ring {42}.

Messages read from the local sender process {45} are immediately written to the ring {46} and the pre-condition controlling input from the local sender is set `false` {47}. As described above, this pre-condition will only be set `true`, when the returned empty packet has been received.

10.2.1 Evaluation

This solution, though functional, does still have some performance limitations in that an element has to wait for a sent packet to be returned before the next message can be sent. This means that on average half the network is filled with empty packets. The next solution removes this restriction by allowing the reuse of an empty packet if a node is ready to send a message from its local sender process.

10.3 A Final Resolution

The behaviour shown in Listing 10-6 shows the behaviour modification required to use an empty packet, as it passes through a node that is ready to output a local message.

```
10 class RingElementv2 implements CSProcess {
11
12     def ChannelInput fromRing
13     def ChannelOutput toRing
14     def ChannelInput fromLocal
15     def ChannelOutput toLocal
16     def int element
17
18     void run () {
19         def RING = 0
20         def LOCAL= 1
21         def ringAlt = new ALT ( [ fromRing, fromLocal ] )
22         def preCon = new boolean[2]
23         preCon[RING] = true
24         preCon[LOCAL] = true
25         def emptyPacket = new RingPacket ( source: -1, destination: -1 ,
26                                           value: -1 , full: false)
27         def localBuffer = new RingPacket()
```

```
28     def localBufferFull = false
29     toRing.write ( emptyPacket )
30     while (true) {
31         def index = ringAlt.select(preCon)
32         switch (index) {
33             case RING:
34                 def ringBuffer = (RingPacket) fromRing.read()
35                 if ( ringBuffer.destination == element ) {
36                     toLocal.write(ringBuffer)
37                     if ( localBufferFull ) {
38                         toRing.write ( localBuffer )
39                         preCon[LOCAL] = true
40                         localBufferFull = false
41                     }
42                     else {
43                         toRing.write ( emptyPacket )
44                     }
45                 }
46                 else {
47                     if ( ringBuffer.full ) {
48                         toRing.write ( ringBuffer )
49                     }
50                     else {
51                         if ( localBufferFull ) {
52                             toRing.write ( localBuffer )
53                             preCon[LOCAL] = true
54                             localBufferFull = false
55                         }
56                         else {
57                             toRing.write ( emptyPacket )
58                         }
59                     }
60                 }
61             break
62         case LOCAL:
63             localBuffer = fromLocal.read()
64             preCon[LOCAL] = false
65             localBufferFull = true
66             break
67         } // end switch
68     }
69 }
70 }
```

Listing 10-6 The Final Ring Element Process (Print Statements Omitted)

The setup of the preconditions and the alternative are the same as the previous version {19–24}. An `emptyPacket` is defined {25, 26} as is a buffer {27} to hold messages from the local sender process. A Boolean flag, `localBufferFull` {28} is used to signify whether or not the `localBuffer` is full. The first action each `RingElement` node undertakes is to output an `emptyPacket` {29}, which has the effect of initialising the system. In general, this initial empty packet will only pass as far as the next node, which by then will have input a message from its local sender and will thus be able to use this empty packet. The extra element has the revised behaviour given in Listing 10-4. As before, the `index` of the enabled alternative is determined {31} and the appropriate `case` selected {32}.

If the `selected` alternative is to read an input packet from the local sender process {62} this is read into the `localBuffer` {63}, the pre-condition flag for this alternative is set `false` {64} and the `localBufferFull` flag set `true`{65}. This does not cause the packet to be written to the ring, merely to get it ready to be written.

If the `selected` alternative relates to an input from the ring then the message packet is read into a `ringBuffer` {34} and the subsequent processing is determined by the state of that message. If the destination of the message is for this node {35} then the message is written to the local receiver process {36}. This means that the node can output the `localBuffer` to the ring if it is full and update the flags associated with the buffer {37–40}; otherwise an `emptyPacket` is written to the ring {43}. If the `ringBuffer` does not have this node as its destination {46} then if the `ringBuffer` is full it is simply written to the ring {47–48}, otherwise the `localBuffer` is processed in the way described previously {51–58}.



LIGS University
based in Hawaii, USA

is currently enrolling in the
Interactive Online **BBA, MBA, MSc,**
DBA and PhD programs:

- ▶ enroll **by October 31st, 2014** and
- ▶ **save up to 11%** on the tuition!
- ▶ pay in 10 installments / 2 years
- ▶ Interactive Online education
- ▶ visit www.ligsuniversity.com to find out more!

Note: LIGS University is not accredited by any nationally recognized accrediting agency listed by the US Secretary of Education. More info [here](#).



10.3.1 Evaluation

This final version has resulted in a solution that routes messages around a circular network, which is inherently prone to deadlock. This version does not suffer from the drawbacks of the previous solution in that an empty packet only travels around the network until it comes to a node that needs to send a message from its `localBuffer` to another node. An argument has been presented that explains why deadlock will not occur because client-server labelling does not provide a categorical solution and furthermore indicates that deadlock will occur.

10.4 Summary

This chapter has analysed a set of processes that inherently tend to deadlock. Two algorithms have been developed that overcome the problems. The benefit of one solution over the other has been explained, though this is difficult to measure unless the system is run over a real network, where each Ring Element process can be placed on a specific processor of that network.

.....Alcatel-Lucent 

www.alcatel-lucent.com/careers

What if you could build your future and create the future?

One generation's transformation is the next's status quo. In the near future, people may soon think it's strange that devices ever had to be "plugged in." To obtain that status, there needs to be "The Shift".

