

# Contents

<b>1</b>	<b>Encoding</b>	<b>1</b>
1.1	The <b>Message</b> type . . . . .	2
1.1.1	Message types . . . . .	2
<b>2</b>	<b>Links and Receivers</b>	<b>3</b>
2.1	The <b>Receiver</b> . . . . .	3
2.2	The <b>Link</b> . . . . .	3
2.2.1	Concrete methods . . . . .	4
2.2.2	Abstract methods . . . . .	5
2.3	The <b>LinkManager</b> . . . . .	5
2.3.1	Methods . . . . .	5
<b>3</b>	<b>Address resolution protocol</b>	<b>7</b>
3.1	Why do we need this? . . . . .	7
3.2	The mapping function . . . . .	7
3.3	Implementation . . . . .	8
3.3.1	The entry type . . . . .	8
3.3.2	Making an ARP request . . . . .	8
3.3.3	Caching . . . . .	13
3.4	The API . . . . .	14
3.4.1	Mock links . . . . .	14
3.4.2	Resolution . . . . .	14
3.4.3	Shutting it down . . . . .	15
<b>4</b>	<b>Routing and forwarding</b>	<b>15</b>
4.1	A route . . . . .	16
4.1.1	Methods . . . . .	17
4.1.2	Route equality . . . . .	17
4.2	The router . . . . .	18
4.2.1	The routing table . . . . .	18
4.2.2	Handling of ingress traffic . . . . .	22

## 1 Encoding

The encoding and decoding for the twine protocol messages is accomplished via the **MessagePack** format. This is a format of which allows one to encode data structures into a byte stream and send them over the wire. It is a *format* because it is standardized - meaning all languages which have a message pack library can decode twine messages if they re-implement the simple routines for the various messages - *all the hard work is accomplished by the underlying message pack library used.*

## 1.1 The Message type

TODO: Add this

### 1.1.1 Message types

The type of a message is the first field which one will consider. We store this as an enum value called `MType`, it is defined below:

```
1  /**
2   * Message type
3   */
4  public enum MType
5  {
6     /**
7      * An unknown type
8      *
9      * Used for developer
10     * safety as this would
11     * be the value for
12     * `MType.init` and hence
13     * implies you haven't
14     * set the `Message`'s
15     * type field
16     */
17     UNKNOWN,
18
19     /**
20     * A route advertisement
21     * message
22     */
23     ADV,
24
25     /**
26     * Unicast data
27     * packet
28     */
29     DATA,
30
31     /**
32     * An ARP request
33     * or reply
34     */
35     ARP
36 }
```

It is this type which will aid us in decoding the `byte[] payload` field with the

intended interpretation.

## 2 Links and Receivers

So-called *links* are *receivers* are terms referred to throughout the documentation and understanding what they are and how they relate to each other is important for what is to follow.

### 2.1 The Receiver

A *receiver* is a relatively simple type, the interface is defined as follows:

```
1  /**
2  * A subscriber could be a router that wants
3  * to subscribe to data coming in from this
4  * interface
5  */
6  public interface Receiver
7  {
8      /**
9       * On reception of the provided data from
10      * the given link-layer address over
11      * the given `Link`
12      *
13      * Params:
14      *   source = the source `Link`
15      *   recv = the received data
16      *   srcAddr = the source link-layer address
17      */
18      public void onReceive(Link source, byte[] recv, String srcAddr);
19 }
```

As you can probably understand from the just of it, it is basically a handler for *ingress* traffic whereby the first argument is the data itself and the second must be the link-layer address the traffic is sourced from. Any class which implements the `Receiver` interface may be (as you will see later) attached to a `Link` such that it can have data passed to it.

### 2.2 The Link

A *Link* is provides us with a method to send data to a destination link-layer address and be notified when we receive packets from link-layer addressed hosts over said link.

A *link* is composed out of a few things:

1. A list of *receivers*

- These are the currently attached receivers which is to be called serially (one after the other) whenever a data packet arrives over this link.
  - Given a link with two **Receiver**(s) attached, then in an example whereby the bytes [66, 65, 65, 66] arrive over the link then that that byte array would be copied to the attached
2. A *source address*
    - We must have a way to determine the source address of the link such that it can be used for various procedures such as ARP
  3. A *transmission* and *broadcast* mechanism
    - We need a way to send unicast (traffic directed to a singular *given* host) and also to broadcast to all those attached to the link

### 2.2.1 Concrete methods

There are a few methods which relate to the **Receiver**(s). These are shown below and essentially are for adding, removing and enumerating receivers for this link:

Method name	Description
<code>attachReceiver(Receiver receiver)</code>	This attaches the given receiver to this <b>Link</b> , meaning packets will be copied to it
<code>removeReceiver(Receiver receiver)</code>	Removes the given receiver from this <b>Link</b> meaning that packets will no longer be copied to it
<code>auto getRecvCnt()</code>	Returns the number of receivers attached to this <b>Link</b>

**2.2.1.1 Implementing your driver** As part of implementing your driver, i.e. by method of extending/sub-classing the **Link** class, you will implement the mechanism (however you go about it) by which will extract data from your link-layer and extract the network-layer part (the twine data payload of your link-layer packet)

and then what do you do with it?

Well, you will want to make this data available to any of the **Receiver**(s) which are attached currently. you want to *pass it up* to the handlers. This can be safely done by calling the `receive(...)` method as shown below:

Method name	Description
<code>receive(byte[] recv, string srcAddr)</code>	This is to be called when the <b>Link</b> sub-class (implementation) has network-layer traffic to provide

Calling this method iterates over every attached **Receiver** and calls their respective `onReceive(...)` methods.

Note: that the `srcAddr` must contain the link-layer source address.

## 2.2.2 Abstract methods

There are a few more methods to take note of, these are not available as an already-implemented set of methods in the `Link` class, and hence must be overridden.

**2.2.2.1 Implementing your driver... *again*** Whilst the usage of the aforementioned `receive(byte[], string)` method had to do with processing *ingress* traffic, these methods require an implementation for handling *egress* traffic.

Method name	Description
<code>void transmit(byte[] xmit, string addr)</code>	Link-implementation specific for driver to send data to a specific destination address
<code>void broadcast(byte[] xmit)</code>	Link-implementation specific for driver to broadcast to all hosts on its broadcast domain
<code>string getAddress()</code>	Link-implementation specific for driver to report its address

Note: The last method, `getAddress()`, must return the `Link`'s link-layer address.

## 2.3 The LinkManager

The `LinkManager` is a module which consumes a single `Receiver` is present to manage the complexity of `Link` management. It is rather simple however. Whenever one requests that a `Link` is to be added to the manager then we shall take the single `Receiver` and attach it to the given link. This helps contain this attachment process in a separate module such that the `Router` need only pass itself in (as it is a `Receiver`) and call the `addLink(Link)` method whenever it wants to attach a link and ensure that packets incoming from them make its way to our `Router`.

### 2.3.1 Methods

The methods that are made available are shown below:

Method name	Description
<code>addLink(Link link)</code>	Adds this link such that we will receive data packets from it onto our <code>Receiver</code>
<code>removeLink(Link link)</code>	Removes this link and ensures we no longer receive data packets from it to our <code>Receiver</code>
<code>Link[] getLinks()</code>	Get a list of all attached links

We look at the implementation for the `addLink(Link)` method below to illustrate the ideas mentioned earlier:

```
1  /**
2   * Manages links in a manner that
3   * a single `Receiver` can be responsible
4   * for all such links attached or
5   * to be attached in the future
6   */
7  public class LinkManager
8  {
9      ...
10
11     /**
12      * Constructs a new `LinkManager`
13      * with the given receiver
14      *
15      * Params:
16      *   receiver = the receiver
17      *   to use
18      */
19     this(Receiver receiver)
20     {
21         this.receiver = receiver;
22         this.linksLock = new Mutex();
23     }
24
25     /**
26      * Adds this link such that we will
27      * receive data packets from it onto
28      * our `Receiver`
29      *
30      * Params:
31      *   link = the link to add
32      */
33     public final void addLink(Link link)
34     {
35         this.linksLock.lock();
36
37         scope(exit)
38         {
39             this.linksLock.unlock();
40         }
41
42         // Add link
43         this.links.insertAfter(this.links[], link);
```

```

44
45         // Receive data from this link
46         link.attachReceiver(this.receiver);
47     }
48
49     ...

```

As you can see we add the provided `Link` to a list of links and then also attach the `Receiver`, which we consumed during construction of the `LinkManager`, to this link such that data packets from the link can be sent up to the `Receiver`.

### 3 Address resolution protocol

The *address resolution protocol* or **ARP** is a standalone module which performs the mapping of a given layer-3 address  $addr_{NL}$  to another address, the link-layer address, which we will call  $addr_{LL}$ .

#### 3.1 Why do we need this?

The reason that we require this  $addr_{LL}$  is because when we need to send data to a host we do so over a link which is indicated in the routing table for said packet.

However, links don't speak the same network-layer protocol as twine - they speak whatever protocol they implement - i.e. Ethernet via `LIInterface` or the in-memory `PipedLink`. Needless to say there is also always a requirement of such a mapping mechanism because several links may be backed by a different link-layer protocols in their `Link` implementation and therefore we cannot marry ourselves to only one link-layer protocol - *we need something dynamic*.

#### 3.2 The mapping function

We now head over to the technical side of things. Before we jump directly into an analysis of the source code it is worth considering what this procedure means in a mathematical sense because at a high-level this is what the code is modeled on.

If we have a router  $r_1$  which has a set of links  $L = \{l_1, l_2\}$  and we wanted to send a packet to a host addresses  $h_1$  and  $h_2$  which are accessible over  $l_1$  and  $l_2$  respectively then the mapping function would appear as such:

$$\begin{aligned}
 (h_1, l_1) &\rightarrow addr_{LL_1} \\
 (h_2, l_2) &\rightarrow addr_{LL_2}
 \end{aligned}$$

On the right hand side the  $addr_{LL_1}$  and  $addr_{LL_2}$  are the resolved link-layer addresses.

$$(h_i, l_i) \rightarrow addr_{LL_i}$$

Therefore we discover that we have the above mapping function which requires the network-layer  $h_i$  address you wish to resolve and the link  $l_i$  over which the resolution must be done, this then mapping to a single scalar - the link-layer address,  $addr_{LL_i}$ .

### 3.3 Implementation

We will begin the examination of the code at the deepest level which models this mathematical function earlier, first, after which we will then consider the code which calls it and how that works.

#### 3.3.1 The entry type

Firstly let us begin with the definition of the in-memory data type which holds the mapping details. this is known as the `ArpEntry` struct and it is shown in part below:

```
1 public struct ArpEntry
2 {
3     private string l3Addr;
4     private string l2Addr;
5
6     ...
7
8     public bool isEmpty()
9     {
10        return this.l3Addr == "" && this.l2Addr == "";
11    }
12
13    public static ArpEntry empty()
14    {
15        return ArpEntry("", "");
16    }
```

Please note the methods `isEmpty()`. An entry is considered empty if both its network-layer and link-layer fields have an empty string in them, this is normally accomplished by calling the `empty()` static method in order to construct such an `ArpEntry`.

#### 3.3.2 Making an ARP request

The code to make an ARP request is in the `regen(Target target)` method and we will now go through it line by line.

**3.3.2.1 Setting up the request and link** Firstly we are provided with a `Target`, this encapsulates the network-layer address and the `Link` instance we want to request over. We now extract both of those items into their own variables:

```
1 // use this link
2 Link link = target.getLink();
3
4 // address we want to resolve
5 String addr = target.getAddr();
```

Before we make the request we will need a way to receive the response, therefore we attach ourselves, the `ArpManager`, as a `Receiver` to the link:

```
1 // attach as a receiver to this link
2 link.attachReceiver(this);
3
4 logger.dbg("attach done");
```

This provides us with a callback method which will be called by the `Link` whenever it receives *any* traffic. It is worth noting that such a method will not run on the thread concerning the code we are looking at now but rather on the thread of the `Link`'s - we will discuss later how will filter it and deliver the result to us, *but for now - back to the code*.

**3.3.2.2 Encoding and sending the request** Now that we know what we want to request and over which link we can go ahead and encode the ARP request message and broadcast it over the link:

```
1 // generate the message and send request
2 Arp arpReq = Arp.newRequest(addr);
3 Message msg;
4 if(toMessage(arpReq, msg))
5 {
6     link.broadcast(msg.encode());
7     logger.dbg("arp req sent");
8 }
9 else
10 {
11     logger.error("Arp failed but oh boy, at the encoding level");
12 }
```

As you can see we make use of the `broadcast(byte[])` method, this is handled by the link's implementation according to its link-layer protocol.

**3.3.2.3 Waiting for a response** We now have to wait for a response and not just any response. It has to be an ARP reply for the particular network-layer address we requested earlier.

This is done with the following code:

```
1 // wait for reply
2 string llAddr;
3 bool status = waitForLLAddr(addr, llAddr);
4
5 ...
```

As you can see we have this call to a method called `waitForLLAddr(addr, llAddr)`. This method will block for us and can wake up if it is signaled to by the callback method running on the Link's thread (as mentioned previously).

---

```
1 Stopwatch timer = Stopwatch(AutoStart.yes);
2
3 // todo, make timeout-able (todo, make configurable)
4 while(timer.peek() < this.timeout)
5 {
6     this.waitLock.lock();
7
8     scope(exit)
9     {
10        this.waitLock.unlock();
11    }
12
13    this.waitSig.wait(dur!("msecs")(500)); // todo, duty cycle if missed notify but also hel
14
15    // scan if we have it
16    string* llAddr = l3Addr in this.addrIncome;
17    if(llAddr != null)
18    {
19        string llAddrRet = *llAddr;
20        this.addrIncome.remove(l3Addr);
21        llAddrOut = llAddrRet; // set result
22        return true; // did not timeout
23    }
24 }
25
26 return false; // timed out
```

Because it is implemented using a condition variable, it could potentially miss a signal from the calling `notify()` if we only call `wait()` on our thread *after* the link's thread has called `notify()`. Therefore, we make our `wait()` wake up every now and then by using a timed-wait, to check if the data has been filled in by the other thread.

Second of all, what we do after retrying from `wait(Duration)` is check if the

*requested network-layer address* has been resolved or not - this is that filtering I was mentioning earlier. This is important as we don't want to wake up for *any* ARP response, but only the one which matches our `addr` requested.

Thirdly, this also gives us a chance to check the while-loop's condition so that we can see if we have timed out (waited long enough) for an ARP response.

---

After all is done, the resulting entry is placed in a globally accessible `string[string] addrIncome` which is protected by the `waitLock` for both threads contending it. We then continue:

```
1  ...
2
3  // if timed out
4  if(!status)
5  {
6      logger.warn("Arp failed for target: ", target);
7      return ArpEntry.empty();
8  }
9  // on success
10 else
11 {
12     ArpEntry arpEntry = ArpEntry(addr, llAddr);
13     logger.info("Arp request completed: ", arpEntry);
14     return arpEntry;
15 }
```

We now check, as I said, if the entry is valid or not. If we timed-out then we would have returned `false`. Now, as we shall see later, we will still have to return *some* `ArpEntry` because that is the signature of our method, `regen(Target target)`. Thus, if we failed to get an `ArpEntry` we then return one generated by `ArpEntry.empty()`, else we return the actual entry that we received.

**3.3.2.4 Catching responses** I have mentioned that the thread which waits for a matching ARP response to come in (the one which calls the `wait(Duration)` above. So then, the question is - which thread is the one calling `notify()` on the condition variable and under which scenarios?

---

Recall that we attached the `ArpManager` as a `Receiver` to the `Link` object which was passed into the `regen(Target)` method:

```
// use this link
Link link = target.getLink();

// address we want to resolve
```

```

string addr = target.getAddr();

// attach as a receiver to this link
link.attachReceiver(this);

logger.dbg("attach done");

```

---

Now the reason for this is that whenever traffic is received on a `Link` it will copy the `byte[]` containing the payload to each attached `Receiver`.

This means that the `ArpManager` will receive all packets from a given link, the question is - which ones to we react to? Well that's easy. Below I show you the `onReceive(Link src, byte[] data, string srcAddr)` method which the arp manager overrides. This is called every time a given link receives data:

```

1  /**
2   * Called by the `Link` which received a packet which
3   * may be of interest to us
4   *
5   * Params:
6   *   src = the `Link` from where the packet came from
7   *   data = the packet's data
8   *   srcAddr = the link-layer source address
9   */
10 public override void onReceive(Link src, byte[] data, string srcAddr)
11 {
12     Message recvMesg;
13     if(Message.decode(data, recvMesg))
14     {
15         // payload type
16         if(recvMesg.getType() == MType.ARP)
17         {
18             Arp arpMesg;
19             if(recvMesg.decodeAs(arpMesg))
20             {
21                 logger.dbg("arpMesg, received: ", arpMesg, "from: ", srcAddr);
22                 ArpReply reply;
23                 if(arpMesg.getReply(reply))
24                 {
25                     logger.info("ArpReply: ", reply);
26
27                     // place and wakeup waiters
28                     placeLLAddr(reply.networkAddr(), reply.llAddr());
29                 }

```

```

30
31         ...
32     ...
33     ...
34     ...
35 }

```

What we do here is we attempt to decode each incoming packet into our `Message` type, then further check if it is an ARP-typed message. If this is the case then we check if it is an ARP request (because as we have seen, ARP requests are **not** handled here).

```

1  /**
2   * Called by the thread which has an ARP response
3   * it would like to pass off to the thread waiting
4   * on the condition variable
5   *
6   * Params:
7   *   l3Addr = the network layer address
8   *   llAddr = the link-layer address
9   */
10 private void placeLLAddr(string l3Addr, string llAddr)
11 {
12     this.waitLock.lock();
13
14     scope(exit)
15     {
16         this.waitLock.unlock();
17     }
18
19     this.waitSig.notify(); // todo, more than one or never?
20
21     this.addrIncome[l3Addr] = llAddr;
22 }

```

If this is the case then we will place the link-layer address into a key-value map where the key is the network-layer address and the value is the link-layer address. After this we wake up the sleeping thread by calling `notify()`.

### 3.3.3 Caching

I mentioned that there is caching involved. The involvement is that all `ArpEntry`'s are stored in a `CacheMap!(ArpEntry)` which means that they will exist in there for some period of time and then be evicted.

If an entry has not yet been cached-in then it is created on demand when you do `map.get(Target)`. Now remember the `regen(Target)` method? Well, thats

the regeneration method that we supply this cache map upon instantiation - therefore it works as expected.

### 3.4 The API

We have now discussed the gritty internals which aid us in creating requests, awaiting replies and then returning the matched entry. We now must move over to the publicly facing API of the `ArpManager`. This really just contains a single method:

```
Optional!(ArpEntry) resolve(string networkAddr, Link onLink)
```

The way this method works is that it will return an `Optional!(ArpEntry)`, meaning that you can test to see if the arp resolution process succeeded or failed (i.e. timed-out for example) using code that looks akin to what shall follow.

---

I have prepared an example which can illustrate the usage of the `ArpManager`. In fact this example is part of a unittest which tests the various scenarios that can occur with the manager itself.

#### 3.4.1 Mock links

Firstly we setup a pseudo-link. This is a sub-class of the `Link` class which is specifically configured to respond **only** to ARP requests and only to those which a mapping exists for.

In this example I configure two mappings of network-layer addresses to link-layer addresses:

$$\begin{aligned}(host_{A_{13}}, dummyLink) &\rightarrow host_{A_{12}} \\ (host_{B_{13}}, dummyLink) &\rightarrow host_{B_{12}}\end{aligned}$$

The code to do this is as follows:

```
1 // Map some layer 3 -> layer 2 addresses
2 string[string] mappings;
3 mappings["hostA:13"] = "hostA:12";
4 mappings["hostB:13"] = "hostB:12";
5
6 // create a dummy link that responds with those mappings
7 ArpRespondingLink dummyLink = new ArpRespondingLink(mappings);
```

#### 3.4.2 Resolution

We then must create an `ArpManager` we can use for the resolution process:

```
ArpManager man = new ArpManager();
```

Now we are ready to attempt resolution. I first try to resolve the link-layer address of the network-layer address `hostA:13` by specifying it along with the mock link, `dummyLink`, which we created earlier:

```
1 // try resolve address `hostA:13` over the `dummyLink` link (should PASS)
2 Optional!(ArpEntry) entry = man.resolve("hostA:13", dummyLink);
3 assert(entry.isPresent());
4 assert(entry.get().llAddr() == mappings["hostA:13"]);
```

In the above case the mapping succeeds and we get an `ArpEntry` returned from `entry.get()`, upon which I extract the link-layer address by calling `llAddr()` on it and comparing it to what I expected, `mappings["hostA:13"]` - which maps to `hostA:12`.

---

We do a similar example for the other host:

```
1 // try resolve address `hostB:13` over the `dummyLink` link (should PASS)
2 entry = man.resolve("hostB:13", dummyLink);
3 assert(entry.isPresent());
4 assert(entry.get().llAddr() == mappings["hostB:13"]);
```

---

Lastly, I wanted to show what a failure would look like. With this we expect that `entry.isPresent()` would return `false` and therefore stop right there:

```
1 // try top resolve `hostC:13` over the `dummyLink` link (should FAIL)
2 entry = man.resolve("hostC:13", dummyLink);
3 assert(entry.isPresent() == false);
```

This resolution fails because our `ArpRespondingLink`, our *dummy link*, doesn't respond to mapping requests of the kind  $(host_{B13}, dummyLink)$ .

### 3.4.3 Shutting it down

We need to shut down the `ArpManager` when we shut down the whole system, this is then accomplished by running its destructor:

```
// shut down the arp manager
destroy(man);
```

## 4 Routing and forwarding

Routing is the process by which one announces their routes to others, whilst forwarding is the usage of those learnt routes in order to facilitate the transfer

of packets from one endpoint to another through a network of inter-connected routers.

## 4.1 A route

Before we can get into the process of routing we must first have a conception of a *route* itself.

A route consists of the following items:

1. A *destination*
  - Describes to whom this route is for, i.e. a route to *who*
2. A *link*
  - The `Link` object over which we can reach this host
3. A *gateway*
  - This is who we would need to forward the packet *via* in order to get the packet either to the final destination (in such a case the *gateway* = *destination*) or the next-hop gateway that we must forward via (*gateway*  $\neq$  *destination*)
4. A *distance*
  - This is metric which doesn't affect *how* packets are forwarded but rather how routes that have the same matching *destination* are tie-broken.
  - Given routes  $r = \{r_1, r_2\}$  and a function  $d(r_i)$  which returns the distance we shall install the route  $r_i$  which has the lowest distance, hence  $r_{installed} = r_i$  where  $d(r_i) = \min(d(r))$  (TODO: fix this maths)
5. A *timer* and *lifetime*
  - We have *timer* which ticks upwards and a *lifetime* which allows us to check when the *timer* > *lifetime* which signifies that the route has expired, indicating that we should remove it from the routing table.

And in code this can be found as the `Route` struct shown below:

```
1  /**
2   * Represents a route
3   */
4  public struct Route
5  {
6     private string dstKey; // destination
7     private Link ll; // link to use
8     private string viaKey; // gateway (can be empty)
9     private ubyte dst; // distance
10
11     private Stopwatch lifetime; // timer
12     private Duration expirationTime; // maximum lifetime
13
14     ...
15 }
```

### 4.1.1 Methods

Some important methods that we have are the following (there are more but these are ones that hold under certain conditions that are not so obvious, therefore I would like to explicitly mention them):

Method	Description
<code>isDirect()</code>	Returns <code>true</code> when <i>gateway = destination</i> , otherwise <code>false</code>
<code>isSelfRoute()</code>	Returns <code>true</code> if the <code>Link</code> is <code>null</code> , otherwise <code>false</code>

### 4.1.2 Route equality

Lastly, route equality is something that is checked as part of the router's code, so we should probably show how we have overrode the `opEquals(Route)` method. This is the method that is called when two `Route` structs are compared for equality using the `==` operator.

Our implementation goes as follows:

```
1 public struct Route
2 {
3     ...
4
5     /**
6      * Compares two routes with one
7      * another
8      *
9      * Params:
10     *   r1 = first route
11     *   r2 = second route
12     * Returns: `true` if the routes
13     * match exactly, otherwise `false`
14     */
15     public static bool isSameRoute(Route r1, Route r2)
16     {
17
18         return r1.destination() == r2.destination() &&
19             r1.gateway() == r2.gateway() &&
20             r1.distance() == r2.distance() &&
21             r1.link() == r2.link();
22     }
23
24     /**
25     * Compares this `Route` with
26     * another
```

```

27     *
28     * Params:
29     *   rhs = the other route
30     * Returns: `true` if the routes
31     * are identical, `false` otherwise
32     */
33     public bool opEquals(Route rhs)
34     {
35         return isSameRoute(this, rhs);
36     }
37 }

```

## 4.2 The router

The `Router` class is the main component of the twine system. Everything such as `Link` objects and so forth make a part of the router's way or working. The router performs several core tasks which include:

1. Maintaining the routing table
  - This means we advertise all routes present in the routing table to other routers over the available links
  - It also means checking the routing table every *now and then* for routes which ought to be expired
  - Receiving advertised routes from other nodes and checking if they should be installed into the table
2. Traffic management
  - Support for installing a message handler which will run whenever traffic destined to you arrives
  - Forwarding traffic on behalf of others; to its final destination
  - Allowing the sending of traffic to other nodes

### 4.2.1 The routing table

The routing table is at the heart of handling egress and forward-intended traffic. It is relatively simple as well, in fact this is the routing table itself:

```

// routing tables
private Route[string] routes;
private Mutex routesLock;

```

There are then several methods which manipulate this routing table by locking it, performing some action and then releasing said lock:

Method	Description
<code>Optional!(Route) findRoute(string)</code>	Given the destination network-layer address this returns an <code>Optional</code> potentially containing the found <code>Route</code>
<code>installRoute(Route route)</code>	Checks if the given route should be installed and, if so, installs it.
<code>dumpRoutes()</code>	This is a debugging method which prints out the routing table in ASCII form
<code>installSelfRoute()</code>	Installs a route to yourself (destination is the result of <code>getPublicKey()</code> )
<code>Route[] getRoutes()</code>	Returns a list of all the currently installed routes
<code>routeSweep()</code>	Checks all routes and evicts those which have expired
<code>advertiseLoop()</code>	Sends out modified routes from the routing table (with us as the <code>via</code> ) on an interval whilst we are running

**4.2.1.1 The self-route** You would have seen the `installSelfRoute()` but are probably wondering what that is. Well, it is actually called in the constructor (`this()`) and is there such that you will have a route in your routing table with a distance of 0 (meaning it will never be replaced) and with a destination to your public key. What this means is the the route advertising mechanism will be able to advertise your presence to other routers - that's it.

```

1  /**
2   * Installs a route to ourselves
3   * which has a distance of `0`,
4   * a destination of our public
5   * key and no link
6   */
7  private void installSelfRoute()
8  {
9      Route selfR = Route(getPublicKey(), null, 0);
10     installRoute(selfR);
11 }

```

As you can see it doesn't have much difference to it than any other route being installed, besides, perhaps - the fact that its `Link` is null. This is such that when you call `isSelfRoute()` (on the `Route` struct) that it will report itself as such.

**4.2.1.2 Installing of routes** Let's take a closer look at `installRoute(Route route)` because I would like to explain the logic that is used to determine whether or not a given route, received from an advertisement, is installed into the routing table or not.

```

1 private void installRoute(Route route)
2 {
3     this.routesLock.lock();
4
5     scope(exit)
6     {
7         this.routesLock.unlock();
8     }
9
10    Route* cr = route.destination() in this.routes;
11
12    ...

```

Firstly as you have seen we lock the routing table mutex to make sure we don't get any inconsistent changes to the routing table during usage (remember that we will be modifying it and others could be doing so as well). We then also set a `scope(exit)` statement which means that upon any exiting of this level of scope we will unlock the mutex. Lastly we then get a pointer to the `Route` in the table at the given key. Remember the routing table was a `Route[string]` which means the `string`, *the key*, is the destination address of the incoming route in this case. The *value* would be the found `Route*` if any.

```

1     ...
2
3     // if no such route installs, go ahead and install it
4     if(cr is null)
5     {
6         this.routes[route.destination()] = route;
7     }

```

As you can see above we first check if the pointer was null, which indicates no route to said destination existed. Therefore we will then install the incoming route at that destination.

```

1     ...
2
3     // if such a route exists, then only install it if it
4     // has a smaller distance than the current route
5     else
6     {
7         if(route.distance() < (*cr).distance())
8         {
9             this.routes[route.destination()] = route;
10        }
11        else
12        {
13            // if matched route is the same as incoming route

```

```

14         // then simply refresh the current one
15         if(*cr == route)
16         {
17             cr.refresh();
18         }
19     }
20 }
21 }

```

However, if a route did exist then we need to check some things before we install it. Namely, we only install the route if the predicate of  $d(r_{incoming}) < d(r_{current})$  where  $d(r_i)$  is the distance metric of a given route  $r_i$ . If this is *not* the case then we do not install the route. However, we do do a check to see if the incoming route is identical (must have been the same router advertising a route we received from it earlier) then we simply refresh it (reset its timer) instead of storing it again, if that is not the case we don't change anything.

**4.2.1.3 Advertising of routes** The advertising of routes is implemented in the `advertiseLoop()` which runs on its own thread and will wake up at a fixed interval in order to perform two operations:

1. Checking for evicted routes
  - By calling `routeSweep()`
2. Sending out advertisements
  - This is explained below

---

We now analyze this loop below:

```

1 // Check for and evict expired routes
2 routeSweep();
3
4 // advertise to all links
5 Link[] selected = getLinkMan().getLinks();
6 logger.info("Advertising to ", selected.length, " many links");

```

---

As we can see above we sweep the routing table firstly by a call to `routeSweep()`, this is implemented as follows:

```

1 this.routesLock.lock();
2
3 scope(exit)
4 {
5     this.routesLock.unlock();
6 }

```

```

7
8 foreach(string destination; this.routes.keys())
9 {
10     Route cro = this.routes[destination];
11
12     // if it has expired and is not a self-route (never expire it)
13     if(cro.hasExpired() && !cro.isSelfRoute())
14     {
15         this.routes.remove(destination);
16         logger.warn("Expired route '", cro, "'");
17     }
18 }

```

It is relatively simple, lock the table check all which have expired, and if they have then remove them from the table. Finalizing by unlocking the table's lock.

**Note:** We check `!cro.isSelfRoute()` because we don't want to expire our self-route, else if we do it we will cease to advertise our prescencce to neighboring routers after the initial sweep after starting the router.

---

We also see how we are enumerating all `Link(s)` which are attached to the router (via its `LinkManager` (returned by `getLinkMan()`)). We would like to advertise all the routes in our table over all of these links.

```

1 // advertise each route in table
2 foreach(Route route; getRoutes())
3 {
4     logger.info("Advertising route '", route, "'");
5     string dst = route.destination();
6     string via = this.getPublicKey(); // routes must be advertised as if they're from me now
7     ubyte distance = route.distance();
8
9     Advertisement advMesg = Advertisement.newAdvertisement(dst, via, distance);
10    Message message;
11    if(toMessage(advMesg, message))
12    {
13        logger.dbg("Sending advertisement on '", link, "'...");
14        link.broadcast(message.encode()); // should have a return value for success or failure
15        logger.info("Sent advertisement");
16    }
17    else
18    {
19        // todo, handle failure to encode
20        logger.error("Failure to encode, developer error");
21    }
22 }

```

23  
24 ...

The advertising of routes works as follows. Given a route  $r_i$  in our routing table, we construct a new route,  $r_{iout}$  of which has all

#### 4.2.2 Handling of ingress traffic

We now move over to the way in which the `Router` receives data packets from its attached links. Recall earlier we described how the `LinkManager` takes in a `Receiver` and kept tracks of all the `Link(s)` requested to be added *and*, if added, would attach the `Receiver` to it.

Well, now we have a single method in the router, `onReceive(Link link, byte[] data, string srcAddr)` which is responsible for handling data packets coming from all of these attached links. It actually calls a method called `process(Link link, byte[] data, string srcAddr)`, hence copying its arguments in. We therefore will look at this method:

```
1 logger.dbg("Received data from link '", link, "' with ", data.length, " many bytes (l1Src: '
2
3 Message recvMesg;
4 if(Message.decode(data, recvMesg))
5 {
6     logger.dbg("Received from link '", link, "' message: ", recvMesg);
7
8     ...
9 }
10 else
11 {
12     logger.warn("Received message from '", link, "' but failed to decode");
13 }
```

What we do here is attempt to decode the incoming bytes into a `Message`, the outermost message encapsulation.

```
1 // Process message
2 MType mType = recvMesg.getType();
3 switch(mType)
4 {
5     // Handle ADV messages
6     case MType.ADV:
7         handle_ADV(link, recvMesg);
8         break;
9     // Handle ARP requests
10    case MType.ARP:
11        handle_ARP(link, srcAddr, recvMesg);
12        break;
```

```
13     // Handle DATA messages
14     case MType.DATA:
15         handle_DATA(link, srcAddr, recvMesg);
16         break;
17     default:
18         logger.warn("Unsupported message type: '", mType, "'");
19 }
```

If the decoding succeeds we then move ahead to determine the type of message and calling the correct method depending on the type.