

**Bachelor's Degree Thesis**

**A New Completely Decentralized  
Communication System Over IP**

Nicholas Helke

August 4, 2013

Prof. Stephan Robert  
HEIG-VD



## **Abstract**

The Internet Protocol (IP) was designed to be resilient to topology changes. It should therefore be possible to build a resilient communication system on top of IP, in fact consumer communication systems deployed today all rely to some extent on central servers. This makes these communication systems vulnerable should the central servers become unreachable for any reason.

This memoir proposes a strategy for new communication system over IP that is completely decentralized, i.e. that doesn't rely on central servers at all. This strategy uses the BitTorrent network's Distributed Hash Table (DHT) (which is based on Kademlia) to map users to IP addresses in a completely decentralized manner. An application network protocol and a proof of concept software application which implement the proposed strategy are also discussed in this memoir.



## Terms of reference

The purpose of the bachelor's degree project of which this memoir is the report is to produce a new completely decentralized communication system over Internet Protocol (IP).

The key words in [3] are to be interpreted in the context of this chapter as they are defined in said document.

The development of this new system is divided into two parts:

1. A theoretical strategy for establishing a communication over IP in a completely decentralized manner. This strategy must meet the following requirements:
  - a) The strategy must enable a conforming implementation to establish a point-to-point connection over IP between two nodes with public Internet Protocol Version 4 (IPv4) addresses where the user need only know the public key of participating node he wishes to contact.
  - b) The strategy can not rely on a central server to establish a mapping from public keys to IPv4 address, i.e. the strategy must be "pure" Peer-to-Peer (P2P) according to Definition 2 of [29]
  - c) A central seeding server or list of servers may be used to seed fresh installations of a conforming implementation with the necessary information to join the system. Once a fresh node has been seeded however, it must be able to continue to function without ever relying on the seeding server or servers again.
  - d) The strategy must be resilient in the face of network splits, i.e. communication must remain possible within the each subnets after the split.
  - e) The strategy must enable an established connection to transport arbitrary data (e.g. text, voice, video or files) in both directions between the nodes.
  - f) The strategy may employ a variety of other strategies in order to extend the applicability of this strategy to nodes behind firewalls and/or Network Address Translation (NAT).
2. A practical implementation of the above strategy in the form of a proof of concept software application and associated application network protocol. This proof of concept application and network application protocol must meet the following requirements:
  - a) The network application protocol must be architecture, platform and language agnostic.

- b) The network application protocol must allow for text communication and must be extendible to support arbitrary forms of communication (e.g. voice, video or files). This may be achieved by building feature negotiation into the protocol.
- c) The proof of concept application must support Linux and may use any Graphical User Interface (GUI) available on the platforms it supports.

# Contents

<b>1. Introduction</b>	<b>1</b>
<b>2. Existing works</b>	<b>3</b>
2.1. Bitmessage . . . . .	3
2.2. Completely decentralised DNS . . . . .	4
2.2.1. Dot-P2P . . . . .	4
2.2.2. Namecoin . . . . .	5
<b>3. Challenges of Completely Decentralised Communication</b>	<b>7</b>
3.1. Finding Peers in a Completely Decentralised Manner . . . . .	7
3.1.1. Decentralised data storage . . . . .	7
3.1.2. BitTorrent’s Mainline DHT . . . . .	9
3.1.3. Kademia . . . . .	10
3.2. Establishing the connection . . . . .	14
3.3. Security . . . . .	16
3.3.1. Obfuscation . . . . .	16
3.3.2. Handshake . . . . .	18
3.3.3. Encryption . . . . .	19
3.3.4. Web of Trust . . . . .	19
3.4. Ensuring resilience to topology changes . . . . .	20
<b>4. Implementation</b>	<b>23</b>
4.1. Designing the program’s architecture and choosing a programming language . . . . .	23
4.2. Reference implementation’s application architecture . . . . .	26
4.3. Finding peers . . . . .	27
4.4. Establishing a connection . . . . .	31
4.4.1. Using SSDP and UPnP to create port mappings in compatible routers	31
4.4.2. On not implementing PCP <i>née</i> NAT-PMP . . . . .	39
4.4.3. Confronting our hypotheses with reality . . . . .	39
4.5. Security . . . . .	41
4.5.1. Public Key Infrastructure . . . . .	41
4.5.2. Adapting the handshake to symmetric NAT initiators . . . . .	41
4.6. Communication protocol . . . . .	42
4.6.1. Basic protocol . . . . .	42
4.6.2. Extending the protocol . . . . .	43

<b>5. Conclusion</b>	<b>47</b>
<b>Bibliography</b>	<b>49</b>
<b>Acronyms</b>	<b>53</b>
<b>Appendices</b>	
<b>A. Dictator Breaker User Manual</b>	<b>A1</b>
A.1. How to run . . . . .	A1
A.2. Command line flags . . . . .	A1
<b>B. Dictator Breaker Code</b>	<b>B1</b>
B.1. <a href="https://go.nhelke.com/dictator-breaker/main.go">go.nhelke.com/dictator-breaker/main.go</a> . . . . .	B1
B.2. <a href="https://go.nhelke.com/dictator-breaker/ws.go">go.nhelke.com/dictator-breaker/ws.go</a> . . . . .	B7
B.3. <a href="https://go.nhelke.com/dictator-breaker/messaging/message.go">go.nhelke.com/dictator-breaker/messaging/message.go</a> . . . . .	B9
B.4. <a href="https://go.nhelke.com/dictator-breaker/messaging/message_test.go">go.nhelke.com/dictator-breaker/messaging/message_test.go</a> . . . . .	B10
B.5. <a href="https://go.nhelke.com/dictator-breaker/security/openpgp.go">go.nhelke.com/dictator-breaker/security/openpgp.go</a> . . . . .	B10
B.6. <a href="https://go.nhelke.com/dictator-breaker/security/openpgp_test.go">go.nhelke.com/dictator-breaker/security/openpgp_test.go</a> . . . . .	B14
B.7. <a href="https://go.nhelke.com/dictator-breaker/views/chat.html">go.nhelke.com/dictator-breaker/views/chat.html</a> . . . . .	B15
B.8. Known issues . . . . .	B17
<b>C. GoUPNPc</b>	<b>C1</b>
C.1. <a href="https://github.com/nhelke/goupnpc/README.mdown">github.com/nhelke/goupnpc/README.mdown</a> . . . . .	C1
C.2. <a href="https://github.com/nhelke/goupnpc/cmd.go">github.com/nhelke/goupnpc/cmd.go</a> . . . . .	C2
C.3. <a href="https://github.com/nhelke/goupnpc/goupnp/goupnp.go">github.com/nhelke/goupnpc/goupnp/goupnp.go</a> . . . . .	C3
C.4. <a href="https://github.com/nhelke/goupnpc/goupnp/goupnp_test.go">github.com/nhelke/goupnpc/goupnp/goupnp_test.go</a> . . . . .	C9
C.5. <a href="https://github.com/nhelke/goupnpc/goupnp/backend.go">github.com/nhelke/goupnpc/goupnp/backend.go</a> . . . . .	C10
C.6. <a href="https://github.com/nhelke/goupnpc/goupnp/backend_test.go">github.com/nhelke/goupnpc/goupnp/backend_test.go</a> . . . . .	C14
C.7. <a href="https://github.com/nhelke/goupnpc/goupnp/ssdp.go">github.com/nhelke/goupnpc/goupnp/ssdp.go</a> . . . . .	C15
C.8. <a href="https://github.com/nhelke/goupnpc/goupnp/ssdp_test.go">github.com/nhelke/goupnpc/goupnp/ssdp_test.go</a> . . . . .	C19
C.9. Known issues . . . . .	C21
<b>D. Patches submitted to open source projects</b>	<b>D1</b>
D.1. Patch submitted to <a href="https://github.com/nictuku/dht">github.com/nictuku/dht</a> . . . . .	D1
D.2. Patch submitted to <a href="https://github.com/monnand/dhcx">github.com/monnand/dhcx</a> . . . . .	D5
<b>E. GNU General Public License Version 2</b>	<b>E1</b>



# 1. Introduction

It was the Arab Spring that first alerted me to the need for a completely decentralized communication system<sup>1</sup>. The falling dictatorships desperately attempted to thwart the revolutions by cutting and disrupting internet communication[8]. The internet was designed to be and is resistant to topology changes, however the countries of the Arab Spring had a limited number of backbone connections to the outside world, which were all under government control. The failing governments cut themselves off from the rest of the world by destroying the routing tables of the edge routers. This was surprisingly effective at disrupting communications.

The number of people relying on American and European based services had been underestimated, or at least there was no widely known data about the penetration of such services. It turns out Facebook, Gmail, Skype and Twitter were the principal tools of communication for the revolutionaries and so when they were cut off from the rest of the world they were no longer able to communicate internally because they were relying on server-based services based outside their isolated subnet. Despite still being physically connected to one another, they were unable to communicate as they no longer had access to the central server.

A worrying trend amongst all countries developing and developed alike are increasing violations of civil liberties such as [12]. A decentralised communication therefore would also provide some protection against unlawful surveillance although this is more debatable given the governments increased regulation and control of Internet Service Providers (ISPs).

This brought to light the necessity for a completely decentralized communication system over Internet Protocol (IP), one that would be as resilient as IP is itself. It should also be mentioned that such a system would also be beneficial to users on accidental topology changes, such as the recent failure of SEA-ME-WE 4 due to criminal sabotage[1, 2] or cases of accidental failure.

Some effort has gone into structuring this memoir for linear reading. It has not however been possible to eliminate forward-references completely. In order to fully grasp some of the shortcomings of the solutions in chapter 2 an understanding of some notions of the theory from chapter 3 can be helpful. This forward-reference can be avoided by skipping chapter 2 entirely. It is of no consequence to the comprehension of the solution proposed by this memoir, it serves principally to show an alternative decentralized communication system.

Other P2P solutions of all sorts are explored in chapter 2. In particular, Bitmessage[31], a project with similar intent, is described, analysed and its shortcomings, which prevent

---

<sup>1</sup>Completely decentralized is to be understood as a "pure" Peer-to-Peer (P2P) system as in Definition 2 of [29].

## *1. Introduction*

it from meeting the terms of reference of this memoir (cf. page iii), are exposed.

Chapter 3 covers the challenges of such a communication system. Lessons from past attempts are drawn and the various solutions envisaged in this memoir are described and evaluated.

Chapter 4 describes the process through which a proof of concept application, which is part of the author's bachelors degree, was developed and the protocol which was developed along side it and which others are welcome to use.

## 2. Existing works

Before embarking on this project I scoured the internet for similar and not-so-much-similar-but-related projects. After sifting through the countless communication systems that call themselves Peer-to-Peer (P2P), but that still rely on some central infrastructure (“hybrid” P2P), I found the following couple of interesting related projects. Although there are most certainly some that I am still unaware of.

### 2.1. Bitmessage

Bitmessage[31] is a proposal for an anonymous decentralized communication system, first described in a white paper by Jonathan Warren in October 2012. Its abstract reads:

We propose a system that allows users to securely send and receive messages, and subscribe to broadcast messages, using a trustless decentralized peer-to-peer protocol. Users need not exchange any data beyond a relatively short (around 36 character) address to ensure security and they need not have any concept of public or private keys to use the system. It is also designed to mask non-content data, like the sender and receiver of messages, from those not involved in the communication.

With the exception of the last sentence, the above abstract would aptly describe the proof of concept application described in chapter 4. The two solutions are, however, very different.

Bitmessage is limited by design to non-real-time communication, so is more akin to e-mail—and perhaps a slight regression over push e-mail. This alone makes it incompatible with the terms of reference of this project.

Bitmessage1 is very heavily inspired by Bitcoin[23]. Messages are sent to all the users of the system on a best effort basis, this is necessary for the solution to provide anonymity to both the sender and receiver of messages. In order to protect the system from flooding, which would be catastrophic as all messages must go to all users, proof-of-work must be appended to messages.

Sending messages to all users also serves as the strategy whereby peers find one another. Essentially peers do not find one another, instead they are relatively confident their message will reach its destination as all messages go to everyone. This is clearly not an efficient strategy and one which is more akin to the Gnutella generation (cf. section 3.1 for an explanation of Gnutella’s part in P2P systems) of linear complexity flooding protocols of P2P systems than the more recent logarithmic P2P protocols.

This strategy of flooding everyone does however provide plausible deniability. As all messages are encrypted using public key cryptography it cannot be surmised from the

## 2. Existing works

message contents alone who the indented recipient is. Given that each message is sent to everyone, determining the recipient by tracking the Internet Protocol (IP) address the message goes to is not possible.

The proof-of-work, which protects the network from flooding, is to take about four minutes on standard hardware. If this proof-of-work also is an effective way of combating spam. Each message requires this four-minute proof-of-work, for the average user this has only a marginal cost in the form of energy costs and slight delay of four minutes in message delivery. Spammers would have to spend those four minutes on each message as the proof-of-work is applied to messages after they have been encrypted with the public key of their respective recipients.

Bitmessage is not a competitor to this paper. It presents an interesting alternative approach, which trades real-time communication for anonymity, plausible deniability and effective elimination of spam.

## 2.2. Completely decentralised DNS

The major challenge to completely decentralized communication is the mapping of users via some identifier to IP addresses. The Domain Name System (DNS) is certainly the most widely deployed and used system for mapping identifiers, in this case fully qualified domain names, to IP addresses. DNS is also a distributed system, it is not however completely decentralized. In fact it is organized in a tree, managed by Internet Corporation for Assigned Names and Numbers (ICANN). The DNS system is therefore subject to the whims of ICANN. Fortunately ICANN takes its responsibility seriously and has not yet censored a top-level domain.

Network Information Centres (NICs) on the other hand have been known to seize domains. In 2010 the United States Department of Justice and Homeland Security started to seize certain dot-com domain names without warning and without possible the possibility of appeal. This has prompted several attempts at creating truly decentralized domain name systems.

### 2.2.1. Dot-P2P

It appears that the first serious attempt in this field, insofar as it was picked up by the Electronic Frontier Foundation (EFF)<sup>1</sup> was Dot-P2P.

Unfortunately aside from the EFF's mention of it, there is not much I can say about it, as in July 2013, when I was attempting to gather information about it, its website<sup>2</sup> was down, seemingly abandoned.

---

<sup>1</sup><https://www.eff.org/deeplinks/2010/12/constructive-direct-action-against-censorship>

<sup>2</sup><http://dot-p2p.org/>

### 2.2.2. Namecoin

A more recent attempt, and one which has the merit of still being active is Namecoin<sup>3</sup>. Like Bitmessage (cf. section 2.1), Namecoin seems to be following a trend of reusing concepts from Bitcoin[23]. In fact the only way to register a Namecoin domain is to exchange Bitcoins for Namecoins which serve as the currency within the Namecoin ecosystem. Unfortunately I fear its close tie to Bitcoin will ultimately prevent it from ever going mainstream given how controversial Bitcoin is outside its niche. Although, who knows, that may change.

---

<sup>3</sup>[http://dot-bit.org/Main\\_Page](http://dot-bit.org/Main_Page)

## 2. Existing works

## 3. Challenges of Completely Decentralised Communication

The use of central servers is inherently vulnerable as a central server's role is to be a known location where information can be fetched or stored. That makes it also a known location which can be blocked, attacked or otherwise disabled. In completely decentralised communication Alice must be able to contact Bob without resorting to the use of a central directory server or list of servers.

In section 3.1 the challenge of finding other peers without using a central directory server is explored.

It is not sufficient to find the peer you wish to talk to, one must also be able to establish a connection with that peer. Unfortunately most internet devices are behind firewalls and/or Network Address Translation (NAT) which are designed to enable devices to reach servers but often completely block connection attempts initiated from devices on the WAN. Section 3.2 discusses the challenges and possible work-around solutions for to the problem of actually establishing a connection to Bob.

Security concerns and possible approaches to address them are the subject of section 3.3.

Any solution to the prior challenges, is not worth retaining as a solution to completely decentralised communication if it is as vulnerable to topology changes as server-based communication solutions. The whole point of completely decentralised communication is to approach the resilience of plain old Internet Protocol (IP) as much as possible. Section 3.4 discusses the resilience of the discussed solutions.

### 3.1. Finding Peers in a Completely Decentralised Manner

Alice wishes to know how to reach Bob. She cannot use a phone book. She is only allowed to contact random people who may or may not know Bob but who can help her by suggesting other people more likely to know Bob. This in essence is the challenge to finding peers in a completely decentralised manner or a so-called "pure" Peer-to-Peer (P2P) system.

#### 3.1.1. Decentralised data storage

Distributing data over many hosts, a practice known as sharding in the database world, is quite common for large databases. The practice has also been around for some time and is quite well understood. It is however very different and becomes a much harder problem when one does not control the hosts that are to collaborate.

### 3. Challenges of Completely Decentralised Communication

The problem of finding information in a decentralized manner over a collaborating but not centrally controlled set of hosts has been the result of quite some interest ever since the original Napster's collapse. Napster served to illustrate the great power and scalability of P2P systems. It relied however on central servers for indexing the content available within the network. Rüdiger Schollmeier defines such systems as "hybrid" P2P networks in [29]. Napster's collapse revealed a weakness in "hybrid" P2P networks, which rely on central servers to some degree, and led to a series of breakthroughs in P2P data storage.

The first reaction wasn't a breakthrough by any means and it is mentioned here only for historic accuracy. Following Napster's demise the first networks to be born from the ashes did away with central servers by using *request flooding*. Request flooding can hardly be called a strategy, it consists of contacting every peer you have ever known, on the off chance that they are still online and that they have the information you are looking for. Gnutella was perhaps the most popular such system and arguably its only contribution to the field of P2P data storage and retrieval was a study by Saroiu et al of probable future uptime as a function of past observed uptime [28]. The results of this study informed the replication parameters used later in Kademia[21], the solution ultimately settled on for this project.

A far more interesting development and one that can be called a breakthrough was the development of a range of new Distributed Hash Tables (DHTs). CAN[25], Chord[30], Pastry[27] and Tapestry[34], to name just four. Kademia also belongs to this set of DHTs, however chronologically it came a couple of years after the aforementioned DHTs and included new features which presented some significant advantages for use cases such as BitTorrent or Emule's DHTs. In particular:

1. It requires less network traffic than the previously mentioned DHTs in order to converge on the requested data.
2. It also dynamically adjusts to load, caching hotspots upstream on the various convergence paths, thereby distributing the load and protecting the random authoritative host(s) from suffering a Distributed Denial of Service (DDoS).

All these solutions have in common the following basic strategy which is still in force today, when developing DHTs:

1. Participating nodes are uniformly spread over some address space. This address space is usually the 160-bit address space and forms what is known as the *overlay network*.
2. Information is spread over the same address space as the nodes usually by using a hashing function and in particular SHA1, hence the 160-bit address space.
3. A function is defined that maps data in the address space to the nearest node or  $k$  nodes in the address space, where  $k$  is known as the replication factor. This is commonly known as the *distance metric*.



4. An algorithm is supplied that given the previous points converges as quickly as possible on the requested data.

The perennity of data in the DHT is dependent on the storage strategy, in particular how it replicates data, and on a sufficient number of nodes remaining online at all times in order for the replication to have a chance succeed. Additionally in a “pure” P2P DHT the number of online hosts must be large enough to ensure that it is highly probable that one of the hosts last known is still online and connected to the rest of the overlay network, as this is the only way that we can rejoin the network while staying true to “pure” P2P. We speak of a critical mass. This is of paramount importance for a system relying solely on a “pure” P2P DHT to be itself perennial.

### 3.1.2. BitTorrent’s Mainline DHT

By far the most widely deployed and used DHT is BitTorrent’s so-called Mainline DHT. Because of this, it is the most perennial DHT currently in use. It is the Mainline DHT’s popularity and therefore perennity which led the communication system described in this memoir to choose this as a backing DHT thereby ensuring immediate critical-mass. Section 4.3 explains how piggy-backing on the actual Mainline DHT designed for and used by BitTorrent clients for tracking other BitTorrent peers is possible and achieved.

It should be noted that choosing based on popularity for the sake of perennity does not necessarily come at the cost of technical inferiority. It is likely that technical criteria dictated the popularity of the given DHT in the first place. It may come at the cost of the latest and greatest developments in the field, as a popular solution has taken time to become popular and once popular cannot as easily be changed.

In fact, in the specific case of the BitTorrent network’s Mainline DHT which chose Kademia, there are few subsequent DHTs that are technically superior. The principal drawback of Kademia is its vulnerability to Sybil-attacks, a vulnerability that subsequent DHT designs have attempted to address.

Sybil-attacks are attacks where a large number of malicious peers join a network in order to have a sufficient mass to cause damage or otherwise disrupt service. It used not to be economically viable to perform Sybil attacks, however illegal botnets and legal virtual private servers for hire by the hour make this a far more feasible endeavour. In fact it is rumoured that the big movie studios and record labels hire companies to infest the Mainline DHT with peers that return bad results and/or that log users trying to obtain certain Torrents, data which they later use to report users to the authorities.

A little anecdote about subsequent DHTs and why the final choice for this project remained Kademia. Petar Maymoukov, one of Kademia’s authors, has since developed Tonika<sup>1</sup> based on the theory in [16], however it is no longer maintained and is not known to be used in any widely distributed solution, so was discarded (despite the reference implementation being in Go, the language chosen in section 4.1).

Another prominent subsequent Sybil-proof DHT is Whānau, developed by Chris Lesniewski-Laas and M. Frans Kaashoek in [18]. It boasts strong protection to Sybil-attacks, up to

<sup>1</sup><https://pdos.csail.mit.edu/~petar/5ttt.org/>

### 3. Challenges of Completely Decentralised Communication

$O(\frac{n}{\log n})$  attack edges.<sup>2</sup> Despite having been tested Whānau on PlanetLab<sup>3</sup>, which a chat application no-less, it remains a theoretical exercise for the time being as no widely distributed solution has picked up on it. A shame given that Whānau even boasts constant-time  $O(1)$  lookups!

Neither of these promising projects therefore has reached the critical mass needed for it to be of interest to this project.

This is a good opportunity to remind ourselves that ultimately we must produce a usable completely distributed messaging system. This is probably the area where this goal has had the strongest influence in essentially determining the chosen technology, i.e. Kademia. It is more important that from the get go the solution work and that it work with the inevitable large fluctuations of users in its beginnings. There would be little point in choosing a solution which addresses Kademia's shortcomings if that solution would then not be sustainable with the inevitably few early-adopters, even if theoretically down the line, with more users, it would work.

In spite of its vulnerability to some distributed attacks Kademia is quite ingenious. It has already been mentioned that Kademia has some features not supported by the other mentioned DHTs, and it is easy, in light of these features and the unavailability of the mentioned subsequent works, to understand Andrew Loewenstern's choice for the Mainline DHT in [19].

#### 3.1.3. Kademia

In order to understand how Kademia works and how it succeeds in being so efficient, we must first go over the strategic points that all DHTs define:

1. Kademia uses random 160-bit identifiers which ideally should be completely random in order to ensure uniform distribution, in practice they are pseudo-randomly generated based on the current time, in order to obtain reasonable uniform distribution, and on the machine's MAC address, in order to ensure uniqueness.
2. Data in Kademia is traditionally spread over the overlay network address space using the SHA1 of the data itself. In the specific case of BitTorrent's Mainline DHT, it turns out that all torrents have a so called *infohash* which represents the data in the torrent and is in actual fact the SHA1 of the torrent's data, leading to a trivial association from any given torrent to the overlay network address space.
3. The bit-wise exclusive-or (XOR) of 160-bit addresses is used as the distance metric for Kademia. It is this breakthrough choice that enables the additional features that Kademia provides, as we will soon see.
4. Finally, the convergence algorithm Kademia uses has logarithmic time-complexity  $O(\log n)$  and converges along a constant path, a fact which comes from the prop-

---

<sup>2</sup>Attack edges are those edges in the overlay network graph, which connect bad and/or malicious nodes, with the sub-graph of good nodes.

<sup>3</sup><https://www.planet-lab.org/>

erties of the XOR distance metric and which enables the caching of data upstream from the authoritative nodes, thereby distributing the load and preventing DDoS

Let us first note that bit-wise XOR is a valid distance metric:

$$d(x, y) = x \oplus y \quad (3.1)$$

$$d(x, x) = 0 \quad (3.2)$$

$$d(x, y) > 0 \Leftrightarrow x \neq y \quad (3.3)$$

$$d(x, y) = d(y, x) \quad \forall x, y \quad (3.4)$$

which also respects the triangle inequality:

$$d(x, y) + d(y, z) \geq d(x, y) \oplus d(y, z) \quad (3.5)$$

$$\wedge d(x, y) \oplus d(y, z) = d(x, z) \quad (3.6)$$

$$\Rightarrow d(x, y) + d(y, z) \geq d(x, z) \quad (3.7)$$

Of more interest to us however is the following property, which makes Kademia scale very well. Given  $x$ , a node, and  $\delta$ , a distance, there exists only one node,  $y$ , such that:

$$d(x, y) = \delta \quad (3.8)$$

Chord and the seminal Kademia paper call this property being *unidirectional*. This property combined with Kademia's routing algorithm means that any given node tends to converge on requested nodes along different but constant paths. This means that nodes closer to the querying node along the convergence path can cache the data thereby protecting the node closest to the requested information and therefore the authoritative node from DDoS. In order to better see how this works, it is time to review how Kademia converges on peers.

Kademia's overlay network can be seen as a binary tree where each successive level represents successive bits of the nodes' IDs. In keeping with reality not all IDs are used and therefore we represent the tree stopping at unique prefixes. This also aids us in fitting the tree on the page.

This tree is depicted in Figure 3.1. We note that in this binary tree the position of highest order bit of the xor-distance which represents the length of the common prefix between two nodes corresponds with how far up the tree one must go to reach the other node.

Next we introduce an extremely important concept for Kademia's node-finding algorithm. Each and every node in the Kademia overlay network views and sub-divides the complete tree differently. For the purposes of this example, we will consider a node with unique prefix  $0011 \dots$ . In figure 3.1 this node is marked in black and the sub-trees  $0011 \dots$  divides the tree into are surrounded by grey ovals. We note that these sub-trees correspond to nodes with prefixes differing in successive bits. Note also in descending the tree along the chosen node's path, how each sub-tree is an order of magnitude smaller than the previous one.

### 3. Challenges of Completely Decentralised Communication




Figure 3.1.: Kademlia binary tree. The black dot shows the location of node  $0011 \dots$  in the tree. Grey ovals show subtrees in which node  $0011 \dots$  must have a contact. (Source: [21])

We now note how if each node only knows a single node for each of its individual sub-trees that is sufficient to be able to find any other given node.<sup>4</sup> As every node has divided his immediate neighbourhood into smaller sub-trees, if the given node we know in a given sub-tree is not the node we are looking for, we are certain that the node we do know will know a closer node to the node we are looking for (if it doesn't know it directly) as it has divided the tree according to its location.

In actual fact, in order to work around so-called node *churn*, each node tracks  $k$  nodes in each sub-tree. These lists are known as *k-buckets*.  $k$  is known as the replication factor and is chosen to ensure with high probability that at least one of the nodes will remain online at any given time. Additionally the list is curated periodically to eliminate unreachable nodes and resort the list by uptime of each node, as [28] has shown that the longer a node stays online, the longer it is likely to remain online.

An example of the path taken to converge on a given address is illustrated in figure 3.2. We note how our starting node,  $0011 \dots$ , is looking for a node,  $1110 \dots$ , but only knows  $101 \dots$  in the  $k$ -bucket where the highest order bit differs.  $0011 \dots$  contacts  $101 \dots$  and asks him if he is the closest active node to  $1110 \dots$  or who he knows who is closer than him to that node-ID.  $101 \dots$  only knows  $1101 \dots$  in his corresponding  $k$ -bucket and he returns this to  $0011 \dots$ . In turn  $1101 \dots$  returns the only node known to him,  $11110 \dots$ . Finally upon contacting  $11110 \dots$ ,  $0011 \dots$  gets the address of node  $1110 \dots$ .

The only networking requirement of Kademlia is P2P User Datagram Protocol (UDP) connections.

<sup>4</sup>We will not reproduce the proof here, it can be found in [21].




Figure 3.2.: Locating a node by its ID. Here the node with prefix 0011 finds the node with prefix 1110 by successively learning of and querying closer and closer nodes. The line segment on top represents the space of 160-bit IDs, and shows how the lookups converge to the target node. Below we illustrate RPC messages made by 0011. The first RPC is to node 101, already known to 0011. Subsequent RPCs are to nodes returned by the previous RPC. (Source: [21])

### 3.2. Establishing the connection

Having found Bob in the proverbial haystack that is the internet, Alice is going to want to establish a connection with him. The choice of the term communication is deliberate as once a connection is established (and depending on the bandwidth), Alice and Bob can exchange anyway they want. Obvious choices are text chat, Voice over IP (VoIP) or video. This memoir will only provide text as the expertise and codecs needed for VoIP or video fall outside the scope of this memoir. Regardless of the chosen medium of communication or the IP version used, a majority (to not say all) users are located behind firewalls, designed to refuse uninitiated incoming connections. This makes establishing “pure” P2P connections a real challenge, a challenge which is the subject of this section.

There is or was some hope that this would no longer be an issue when Internet Protocol Version 6 (IPv6) finally gets widely deployed, as each device will in theory be directly addressable from the internet and NAT will no longer be an impedance to establishing connections. Sadly early indications are that firewalls built into IPv6-compatible consumer routers will by default refuse uninitiated connections, meaning that the same contrived strategies developed to circumvent Internet Protocol Version 4 (IPv4) NAT, will unfortunately continue to be relevant after the switch to IPv6, in order to allow incoming connections.

The terms of reference of this memoir call only for IPv4 connectivity, so we will focus on that in this section, we should however keep in mind that some of this work will be relevant should this project ever be extended for IPv6 connectivity.

The cleanest means to enable P2P connections to be established through NAT and arguably the only means that respects the definition of a “pure” P2P system, is the use of port mapping protocols to negotiate the opening of incoming ports on the router or Internet Gateway Device (IGD) and to discover the publicly addressable IP address that should be announced to the world. There are two main port mapping protocols, Universal Plug and Play (UPnP) and Port Control Protocol (PCP), both are discussed in more detail in section 4.4.

Not all IGDs support port mapping protocols. These protocols are considered by many to represent security risks and so are generally disabled (even if the hardware supports them) on most (to not say all) enterprise IGDs. As this limitation concerns all P2P networks, both “pure” and “hybrid”, wishing to establish P2P connections between peers, other solutions have already been developed, however there is no panacea.

Session Traversal Utilities for NAT (STUN), defined in [26], is a protocol enabling nodes behind NAT to discover not only their publicly addressable IP address but also their external port number. This information is only valid however for non-symmetric NATs as defined in [20]. Additionally STUN’s reliance on servers discovered using Domain Name System (DNS) service records, means that using this strategy immediately downgrades any P2P network to a “hybrid” P2P network.

Some of the principals of STUN can potentially be used in a “pure” P2P system as long as each peer acts as a STUN server. There are however security concerns with such a strategy. An unauthenticated STUN server could potentially perform a DDoS attack by providing all nodes with the IP address of the victim as their own, the nodes

would in turn publish this to the DHT thereby getting other nodes to DDoS the targeted IP address. If we only perform STUN procedure after we have authenticated the peer as trustworthy (a concept we will discuss in section 3.3), we can then only realistically join the network as long as one of our contacts remains not only online, but accessible in case of network split, meaning he must be in the same subnet after the split. These make for unrealistic expectations and reliance on such a strategy would weaken the system's resilience as a whole.

It should also be noted as long as one of the two peers which are attempting to establish a connection is publicly accessible (perhaps thanks to port mapping) a connection can be established between the two. It is only if both users are not publicly accessible that additional methods must be employed.

Commercial communication solutions use Traversal Using Relays around NAT (TURN) as defined in [20], or in plain English, a *relay server*. As with STUN while the standard is designed around the use of a server, the principals can be reused in a decentralised solution, although it may no longer strictly be considered a "pure" P2P solution, according to [29], as the peers are no longer all equal. This is by no means a novel idea; Skype uses them. In fact relay servers in a decentralised system are known as *super-peers*.

As noted above as long as one of the nodes is publicly accessible, a communication can be established. The addition of super-nodes therefore enables two non-publicly accessible nodes to nevertheless communicate by using a third-party super-node which is publicly accessible as an intermediary for the connection.

Super-peers are far from ideal. Skype was infamous for automatically using clients running their software as super-peers or *supernodes*, without warning their users or asking for their consent beyond an article in their terms of service. A more conscientious way, would be to ask potential super-peers whether they are willing to act as such, thereby allowing them to reach this decision consciously and to factor in considerations like their bandwidth costs. This of course will mean that only a small minority of the minority of potential super-peers will actually play that role, as many users unfamiliar with the concept will, in doubt, simply turn down the option, if given the chance. This probably factored into Skype's decision to simply require potential super-peers to act as such in their terms of service.

The observant reader will have noticed the use of the past tense in the last paragraph. The reason is that Skype has since abandoned its super-peers, as reported in [11], in favour of centrally controlled relay servers that act like the old super-peers.

Another downside to super-peers is that they potentially open up the protocol to eaves-dropping. An obvious solution is to use encrypted connections, thereby theoretically preventing super-peers from knowing the contents of what you exchange with other peers. However in an age where governments are building super-computers with the potential to brute-force encrypted connections in a matter of minutes, governments might be tempted to set themselves up as super-peers and thereby eaves-drop on many communications within the network. Additionally in case of political tension, the governments could easily turn off their super-peer network thereby disrupting communications on the network, if the network grew to rely on super-peers. Latency is another issue with the use of relays, particularly for VoIP or video. The additional encapsulation

### 3. Challenges of Completely Decentralised Communication

required to route the packets via a relay would add to the latency induced by the lower network layers.

Interestingly ever since Skype has centralized its super-peers, it has started sharing data with the National Security Agency (NSA) as part of PRISM, as reported in [22].

These considerations mean that using super-peers actually has the potential to weaken the resilience of the network. While it may enable more users to connect to the network in the short-term it also increases greatly the number of users who would suffer service disruptions if the network were to be disrupted. Section 3.4 goes into more detail of the resilience of the chosen strategy to topology changes.

For the sake of exhaustiveness, I must point out that in light of recent revelations about the practices of western governments' security services, it would seem that even western governments have control over the Internet Service Providers (ISPs) that serve most if not all end-users. This means that resilience to topology changes might be of little relevance, if the government can simply tell ISPs to shut down. Concerns over privacy and resilience with super-peers pale in comparison to these concerns. However the solutions to these concerns of government over-surveillance are an issue for a political science paper and not a computer science issue.

## 3.3. Security

We have already discussed in section 3.2 the challenges and potential issues with passing your data around third parties and how encryption could be used to help protect against such attacks. Encryption is also just generally a good idea. There is no reason to offer would-be eavesdroppers the opportunity to do so. In this section we will discuss possible encryption techniques and the possibility of obfuscating the protocol.

This is not however the only security concern. As with any P2P communication system that maps users to IP addresses caution must be taken to prevent third parties from tracking the location of users by tracking the IP addresses they connect from. In this section we will therefore also discuss what steps we propose to mask to a certain extent this information.

### 3.3.1. Obfuscation

The use of Mainline DHT, as discussed in section 3.1, means that we announce ourselves to BitTorrent clients as a BitTorrent client participating in the swarm with the infohash (the BitTorrent terminology for a Torrent's SHA1 hash) of our public key. This means that a non-standard BitTorrent client which has been programmed to attempt to download any Torrent it hears about through the DHT might attempt to connect to our client. This fact actually gave me the idea that I might down the line add an actual BitTorrent client to the program and offer up the actual public key file if a BitTorrent request was received, easily identifiable by the protocol header, cf. figure 3.3.

This simple and plain text header coupled with the enormous amount of traffic gen-



```

00000000 19 42 69 74 54 6f 72 72 65 6e 74 20 70 72 6f 74 |.BitTorrent prot|
00000010 6f 63 6f 6c                                     |ocol|
00000014

```

Figure 3.3.: Hex dump of BitTorrent protocol header as defined in [6]

erated by BitTorrent prompted ISPs to throttle BitTorrent traffic<sup>5</sup>.

In response some BitTorrent clients (but notably not the reference implementation) introduced Message Stream Encryption (MSE)<sup>6</sup>. MSE uses Diffie-Hellman key exchange, the infohash of the torrent the connection's initiator is attempting to download as a shared secret and random noise to obfuscate traffic between BitTorrent clients including importantly the headers. The solution uses the exchanged key to encrypt the rest of traffic using RC4. The solution is not cryptographically secure. It is designed to be a cheap (resource-wise) obfuscation solution and nothing more.

I toyed with the idea of using the same method so as to provide users of the system with plausible deniability, if for some reason the use of this system was to become illegal someday, somewhere. Statistical detection, it transpires in [15], is capable of detecting BitTorrent traffic even with MSE.

Consider voice communication. Talking is made up of moments of sound separated by pauses, particularly in conversation when the speaker is considering what to say on the spot. An additional tell-tale sign of voice communication is that generally only one person speaks at a time. Tracking the volume of the packets sent between hosts enables one to recognise these patterns of silence, talk and only one user at a time. In fact Charles V. Wright et al show in [33] that it is not only possible to recognise the use of VoIP even over encrypted channels, but also to detect certain spoken phrases with characteristic patterns.

This statistical detection means it would very likely also be possible to differentiate the traffic of the proposed communication system from that of BitTorrent traffic rendering null and void the possibility of plausible deniability. Unfortunately I came by this information only after I had already written and submitted a patch (cf. appendix D.2) to a Go library for Diffie-Hellman key exchange<sup>7</sup> which was missing support for arbitrary parameters not defined in the seminal Requests for Comments (RFCs) [14, 17].

As mentioned above, MSE is not cryptographically secure, so its use would have added an additional layer of encryption to the layer necessary to ensure private communication. Given the statistical detection and the fact that the traffic would in any case have to be encrypted by other means as well, I decided to abandon the idea of obfuscating the traffic. Steganography is therefore left as an exercise for the reader.

<sup>5</sup><https://torrentfreak.com/new-data-exposes-bittorrent-throttling-isps-120809/>

<sup>6</sup>[http://wiki.vuze.com/w/Message\\_Stream\\_Encryption](http://wiki.vuze.com/w/Message_Stream_Encryption)

<sup>7</sup><https://github.com/monnand/dhcx>

### 3. Challenges of Completely Decentralised Communication

#### 3.3.2. Handshake

Anyone can announce themselves as you on the Mainline DHT, by design. One of the first things I tried when I had the DHT code working was looking up the hash of my public key, which I am fairly confident does not correspond with a real Torrent. I discovered that countless peers announced themselves as participating in such a torrent, even though the chance of such a Torrent existing is less than your chances of being mauled by an escaped gorilla<sup>8</sup>.

It turns out that there are numerous bad nodes out there, some sponsored by record companies, others by movie studios, and who knows who else. The roles of these bad nodes can be to pollute the DHT with incorrect information, track usage, a combination of the former and the latter, or something else yet.

This means that when we search for a given person, we cannot be certain which of the many hosts that announce themselves are genuine. Obviously the only way to be sure is to attempt a handshake with each and every peer. Some will simply drop the connection, other will mimic the BitTorrent protocol, and there may well be other, as yet, unobserved behaviours. Regardless we must design the handshake in advance to protect against malicious peers that might be set up to track users. The handshake should also be hardened against *replay attacks* which would-be spammers might be tempted to exploit in order to send spam on the network, should it ever garner a sufficient following.

A good handshake should therefore ensure the following:

1. The handshake must protect the initiator from identifying himself to an unauthorized user who might otherwise impersonate one of your friends in order to track you by logging your IP address every time your node attempts to connect.
2. The handshake receiver must be reasonably protected against replay attacks, where an unauthorised host who has somehow capture an initiator's packet sends it to the receiver in the hopes of discovering if the receiver is present and whether the receiver responds to the captured initiation packet.
3. The handshake initiator must be reasonably protected against replayed responses, where a malicious host attempts to fool the initiator into believing it is an authorised peer by replying with a previously captured good response from a real authorised peer.

Meeting these requirements can be done using public-key cryptography, by following the following procedure:

1. The handshake initiator signs his/her public IP address, the current date and time in Coordinated Universal Time (UTC) and a random authentication token with his/her private key and then encrypts message with the public key of the receiver.
2. The receiver decrypts the message, checks the signature. To guard against replay attacks, the receiver checks the IP address and the timestamp.

---

<sup>8</sup><http://stackoverflow.com/questions/4676828#4681221>

3. The receiver replies to the initiator with his IP address, timestamp and the random token received from the initiator, signed with the receiver's private key and then encrypted with the initiator's public key. At this point the receiver is reasonably certain of the receiver's authenticity and starts awaiting messages from the initiator.
4. The initiator checks the IP address and timestamp of the returned packet as well as the random token sent to the receiver. If the tests pass, the initiator can be reasonably certain of the authenticity of the receiver and can start exchanging messages with him/her.

The use of the tuple (IP address, timestamp, random token) means that we can safely relax the timestamp resolution to one hour thereby removing the need for synchronized time between hosts.<sup>9</sup> This tuple forms the main protection against replay attacks, it is unlikely in any given hour that an attacker will be able to take over a given IP address and see the same random token be generated.

As all messages in the handshake are encrypted with the public keys of their intended recipients, the best an attacker could hope to achieve with a replay attack is user tracking. That is an attacker could confirm the location of a given user by replaying the packets of a contact of that user. The proposed system cannot be expected to detect cases where the private key has been compromised.

### 3.3.3. Encryption

Public-key cryptography is computationally expensive. Like many before, we propose using public-key cryptography only for the handshake. Thereafter Advanced Encryption Standard (AES) should be used. The attentive reader will recall our handshake procedure includes a random token. This random token corresponds to the AES key used for all communication after the handshake.

You will also recall that at the end of our handshake procedure the receiver is in the waiting state, waiting for messages from the initiator. The initiator must therefore send a ping over the AES channel if the handshake completes to tell the receiver that the AES channel is established and symmetric, meaning that either host may send messages on the channel asynchronously. The actual network protocol is discussed in section 4.6.

### 3.3.4. Web of Trust

In order to prevent spam, Web of Trust (WoT) is used in lieu of expensive and environmentally unfriendly proof of work as used by Bitmessage (cf. section 2.1). This does however mean that new people may not contact you before exchanging their public keys with you. This strikes me as an acceptable trade off. Additionally I have no plan

---

<sup>9</sup>It dawns on me that time synchronisation between coordinating but not centrally controlled hosts would make for an interesting paper. The problem of time synchronisation between centrally controlled hosts is a well known textbook example.

### 3. Challenges of Completely Decentralised Communication

to replace the existing Pretty Good Privacy (PGP) key-servers. In case of degraded connection, this may prevent temporarily the creation of new contacts. A possible work around for this would be manually exchange keys, using physical media like a memory card or stick. Another option, and one which could conceivably be build into the application at a later date, would be to use the BitTorrent protocol, seeing as nodes announce themselves as participating in the torrent corresponding with the public key. We have already seen in section 3.3, that BitTorrent has a distinctive header. It would be relatively easy add a minimal BitTorrent client to the program that would serve up the public key, upon detecting a BitTorrent connection. This option however would be reserved as a stop-gap solution for those cases when the standard PGP key-servers are unavailable.

#### 3.4. Ensuring resilience to topology changes

Alice is connected to many different peers participating in the communication system. All of a sudden a backbone connection is lost. Alice now wishes to contact Bob. A physical link exists between Alice and Bob, however who is to say whether Alice knows random people she can contact between her and Bob who are in the same subnet formed by the lost uplink connection.

It turns out unfortunately that this is impossible to predict and will depend on many different factors outside of the control of this project. One worrying fact even without topology changes is that [7] found that due to variations in implementations of routing code not all the supposedly compatible Mainline DHT clients conform sufficiently to Kademia in order to ensure deterministic convergence on requested nodes.

We can attempt to anticipate what might happen, however without a full scale test this will remain purely theoretical behaviour. What should happen when the network splits, is that users on either side of the split will all of a sudden see a number of the nodes they have in their  $k$ -buckets disappear. It will and should be able to compensate for this provided that at least one node per  $k$ -bucket is still accessible. Since all nodes will behave in this manner, the result will be two completely independent Kademia DHTs. One on each side of the split.

For instance without tests it is difficult to estimate how probable it is for more than two DHTs to exist after the physical split. Intuitively it seems unlikely, however theoretically it could still happen.

Consider three subsets of nodes  $A, B, C$ . Before a topology change, they are all physically connected to one another, however it just so happens that nodes in  $A$  only have nodes from  $A \cup B$  in their routing tables, while  $C$  only has nodes from  $C \cup B$  in its routing tables. If we now split the network between  $A$  and  $B$  and  $C$  and  $B$  such that  $A$  and  $C$  are still physically but  $A$  is not, we have created a case where there are three overlay-subnets for only two physical subnets.

Additionally these overlay networks may not be able to join back together automatically. The most effective way to merge split overlay networks is to reseed routing tables using a common centralised DHT node. Under normal circumstances this seeding

server is used by freshly installed clients to enable them to seed their routing tables. It is however reasonable to assume that if a network split occurs these are not normal circumstances. These seeding routers, like all central infrastructure are a weak link, and could relatively easily be attacked in coordination with a network split so as to prevent the networks from merging without a software update changing the seeding server address.

While I was at San Jose State University (SJSU) finishing this thesis, I came across a platform, called PlanetLab<sup>10</sup>, for testing P2P programs like this one. I tried to take advantage of it while I was in San Jose. Jon Pearce, dean of computer science at SJSU, kindly put me in touch with the engineer responsible for SJSU's participation in PlanetLab, Xiao Su. Unfortunately she was away while I was in San Jose and was unable to grant me access to the platform while on the road. I was therefore most unfortunately unable to confront the above theory to reality.

In summary the theory is that a network should split into just two overlay-subnets and the two networks should be able to merge again provided enough time and common the seed of nodes. Unfortunately without PlanetLab it was impossible to test.

---

<sup>10</sup><https://www.planet-lab.org/>

### *3. Challenges of Completely Decentralised Communication*

In theory, theory and practice are  
the same. In practice, they are not.

---

(Albert Einstein)

## 4. Implementation

This chapter discusses the implementation of a proof of concept software application implementing the theoretical solutions discussed in Chapter 3 and an associated network application protocol.

The implementation details are discussed through concrete examples in the Go programming language<sup>1</sup>, which was used in the development of the proof of concept application. The application network protocol is however, as required by the terms of reference (cf. page iii), architecture, platform and language agnostic. It should therefore nonetheless be possible to implement a conforming application in another language by following this chapter.

The process by which the application's architecture was determined and the reasons behind the choice of Go as a programming language are discussed in section 4.1.

The shaping of the reference implementation's internals are the subject of section 4.2.

Section 4.3 discusses the use of Kademila[21], a Distributed Hash Table (DHT), to find the Internet Protocol (IP) addresses of node in the network. More specifically the section discusses the use of the BitTorrent network's instance of Kademlia, known as the Mainline DHT [19].

Section 4.4 discusses the connection process, a Go library I wrote for the purpose of mapping ports on Universal Plug and Play (UPnP) Internet Gateway Devices (IGDs) and confronts the theory of traversing Network Address Translation (NAT) with reality.

Section 4.5 discusses the implementation of the security features discussed in section 3.3 and in particular the reworking of the handshake procedure to work around NAT.

Section 4.6 discusses the communication protocol and its features to support arbitrary communication mediums including but not limited to text, voice, video and files.

### 4.1. Designing the program's architecture and choosing a programming language

The strategy described in Chapter 3 has a significant drawback in that it is slow. Median lookup times were found in [7] to be over a minute. Although in my experiments I found it to be much faster, often less than 30 seconds. This still makes for an unacceptable launch time for an interactive user oriented application in this day and age. This led me to consider decoupling the application into two parts, a daemon which starts with the user session and a front-end application which merely presents the communications and availability information collected by the daemon.

---

<sup>1</sup><http://golang.org/ref/spec>

#### 4. Implementation

A considerable advantage to decoupling the backend and the frontend is that it forces the creation of a well defined Application Programming Interface (API) between the two. A well defined API greatly simplifies the creation of a multitude of different frontends for a wide variety of device and platforms. This in turn encourages the development of native frontends on each platform, which greatly improves the user experience on each platform, thereby increasing the good-will towards the application. An important consideration for increasing an open-source application's audience, where there is no money for marketing.

Splitting the program into separate daemons and front-ends is also necessary in order to efficiently support smartphones. While far more efficient than BitMessage (cf. section 2.1), the strategy of this paper still is not appropriate for mobile applications where power efficiency is the primary concern. The number of connections required and the fact that the application must keep a socket open for incoming connections makes it a poor match for mobile devices.

Client-server communication is far more efficient than Peer-to-Peer (P2P) communication. Decoupling the back-end would allow a mobile client to be developed which would let the back-end run on a server. Obviously this defeats the whole purpose of this communication system, however it can be seen as just one possible tier in a tiered solution.

The tiers can be organised from most energy efficient and least resilient to least efficient and most resilient. Users not concerned with delegating their backend process to a third party could opt for an even more efficient tier which could use vendor provided push notifications services to further increase energy efficiency on the mobile device. The three tiers are described below and illustrated in figure 4.1.


1. The bottom tier from a reliability perspective and the top tier from a power efficiency perspective is to delegate the running of the daemon to a central server controlled by the mobile application developer. The central server can then use vendor controlled push services (e.g. Apple Push Notifications<sup>2</sup> or Google Cloud Messaging<sup>3</sup>) to push events to the client as needed, leaving the application free to suspend completely when not in use. Figure 4.1(a)
2. The second tier would be for the mobile application to connect to an instance of the daemon running on a machine controlled by the user (e.g. his desktop computer). This would then only require the mobile application to maintain one connection, with the daemon. This would require the application to stay in the background, thereby draining the battery more than in the first tier. Figure 4.1(b)
3. The third and most resilient tier would be used as a fallback should either or both of the first two fail and would consist of running the daemon on the mobile terminal. Figure 4.1(c)

---


<sup>2</sup><https://developer.apple.com/notifications/>

<sup>3</sup><https://developer.android.com/google/gcm/>






(a) Most efficient tier, where communication with the mobile device is delegated to a third party



(b) Efficient tier, where the mobile device keeps a single connection open with the user's own PC, which runs the backend



(c) Most resilient tier, where the phone runs the backend locally and connects directly to the other clients of the communication system

Figure 4.1.: The various possible tiers which can be used by a mobile device to participate in the communication system from most efficient and least resilient to least efficient and most resilient

## 4. Implementation

Ideally a mobile implementation of this communication system would provide all three tiers and would automatically switch between them, or at least let the user switch between them as they become viable or not.

Initially I thought of using ubiquitous C. C has the advantage of being very portable. C even works on Android, using the Native Development Kit (NDK). C also has a vast array of well supported (and sometimes well written) libraries for a wide range of use cases.

C also has considerable disadvantages and is a ruthlessly low-level by today's standards. It is neither type-safe nor memory-safe, almost inevitably leading to segmentation faults during development (or worse during deployment) which are a real bane to debug.

I also considered two Java Virtual Machine (JVM) languages, Java or Scala<sup>4</sup>. Both Java and Scala are memory-safe and Scala is type-safe to boot. I ultimately ruled them out as too bloated for a daemon and because Apple is no longer shipping the Java Runtime with new Macs thereby severely diminishing its attractiveness as a cross platform solution.

I ultimately settled on Go.

Go is a relatively young open-source system programming language developed by Robert Griesemer, Rob Pike and Ken Thompson at Google starting in 2007. It is both memory-safe and type-safe, it has garbage-collection and sports a C-like syntax. It is a fast compiled language which produces statically linked binaries only slightly larger than C binaries that link to the local *libC*, but are still an order of magnitude smaller than what would be possible with the JVM. It is simple (thanks to [4]) to cross compile Go programs for all the platforms it supports and, importantly, it provides a unified network interface regardless of the kernel. The only relevant platform not yet supported by Go is Darwin/ARM, however given that there is no facility to run daemons on stock versions of iOS. It can therefore be argued that this is a non issue for this project as only the first two tiers are possible in any case on that platform.

My inner autodidact also wanted to try this language that it had learn the previous summer but had not yet put to good use.

### 4.2. Reference implementation's application architecture

The terms of reference call for a reference implementation. We have already discussed in section 4.1 generally our plans for the reference implementation. In this section we will discuss the plan for the application's inner structure and then in the subsequent sections we will actually see the code.

Now is also the time to announce that the codename of the reference implementation is **Dictator Breaker**. You will see this mainly in the code in appendix B.

One of the reasons I chose Go is its notoriety for good, compiled, cross-platform support, as covered in section 4.1. There are however some platform specific issues that no abstraction layer can efface. In particular there are no so-called UNIX sockets on

---

<sup>4</sup><http://www.scala-lang.org/>

Windows. This is an issue for communication between the backend and the frontend; a UNIX socket would have been ideal for this.

In order to work around this issue I decided to go a different route and develop a WebSocket server bound to the loopback interface, that would serve the Graphical User Interface (GUI) to users in their browsers. Serving a GUI as Hyper-Text Markup Language (HTML)/JavaScript (JS) over Hyper-Text Transport Protocol (HTTP) is not new or indeed scarce. Probably the most widely deployed such program is the Common Unix Printing System (CUPS), which is installed on all Macs and is the principal printer subsystem in use in Linux-based distributions.

Another platform specific note is for tiered solutions for smartphones. For the purposes of simplicity the reference implementation includes the encryption in the backend to avoid having to implement it for each different platform in each different backend. This works well on desktop machines where the communication between the backend and the frontend happen over the loopback interface and therefore are safe from snooping (except from root). In the tiered solutions proposed for smartphone communication in Figure 4.1 the channel between the backend and the front end is not secure except for 4.1(c) which is equivalent to the desktop case. In order to reduce the need to substantially modify the backend for the smartphone use-case, I structured divided the program into the components seen in figure 4.2.

The plan was that only the Protocol Manager would have had to be adapted for the mobile use-case and that the interfaces, be they desktop or mobile, would sit just above it.

If you look closely you will note a circular dependency in the graph of figure 4.2. This circular dependency and the fact that the interface turned out to need to influence happenings as far down as the connection manager led me to the realisation that the structure didn't properly match what I needed for the data and events that I was processing.

In stead I opted for a monolithic event or run loop. This loop, in the main function would take care of all incoming events, be they from the OS or from the GUI and would dispatch them as necessary.

The structure and packages in the final architecture are illustrated in figure 4.3.

### 4.3. Finding peers

Something I found amusing when it came to implementing the finding of peers using BitTorrent's Mainline DHT is that, for the first time I can remember, I was unable to find a maintained C library. This fact played a role in the choice of Go. The best I could find for C was BitDHT by Robert "DrBob" Fernie, however it was last updated in 2002. I contacted DrBob who informed me that he had moved on to other projects.

Finding a well maintained self-contained library to interface with the Mainline DHT of BitTorrent for Go, turned out to be a leisurely walk in the park. I searched for "dht" on GoDoc<sup>5</sup> and immediately came across a "dht" package maintained by Yves Junqueira

<sup>5</sup><http://godoc.org/> – GoDoc is a central website which keeps up-to-date documentation on all open-

#### 4. Implementation

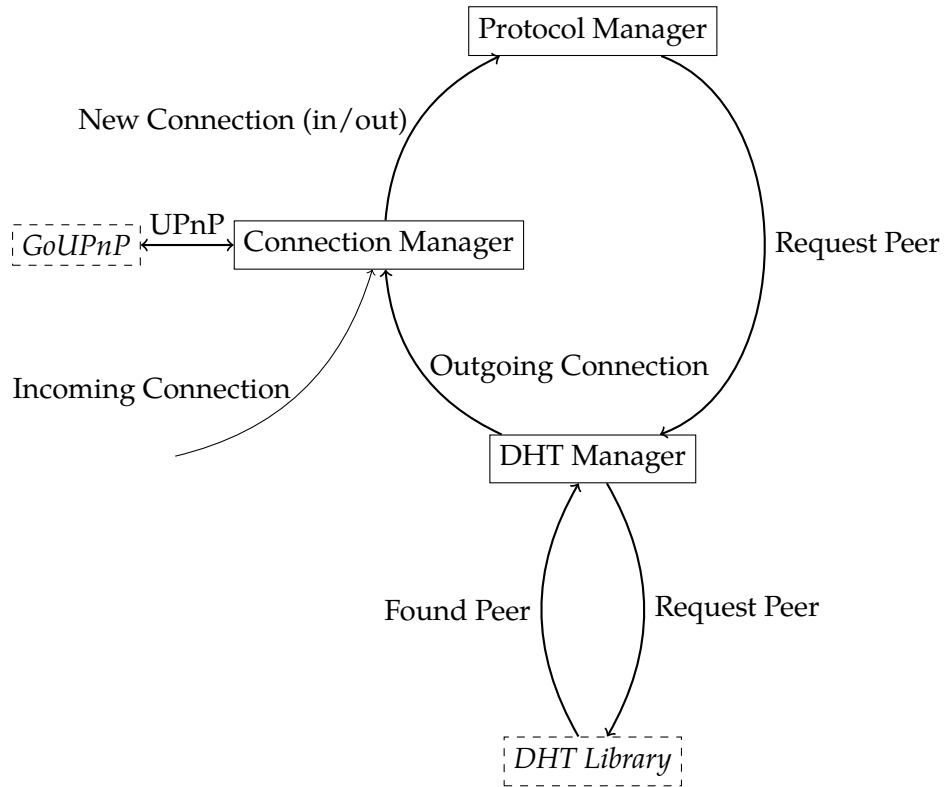


Figure 4.2.: The various components of the original plan for the reference implementation and the messages that they exchange. **Solid rectangles** represent the various components of the originally planned reference implementation. **Arrows** represent OO messages passed between components. **Arrows** not originating in a component, are external events, passed to the program by the OS. **Italic dashed rectangles** represent external libraries.

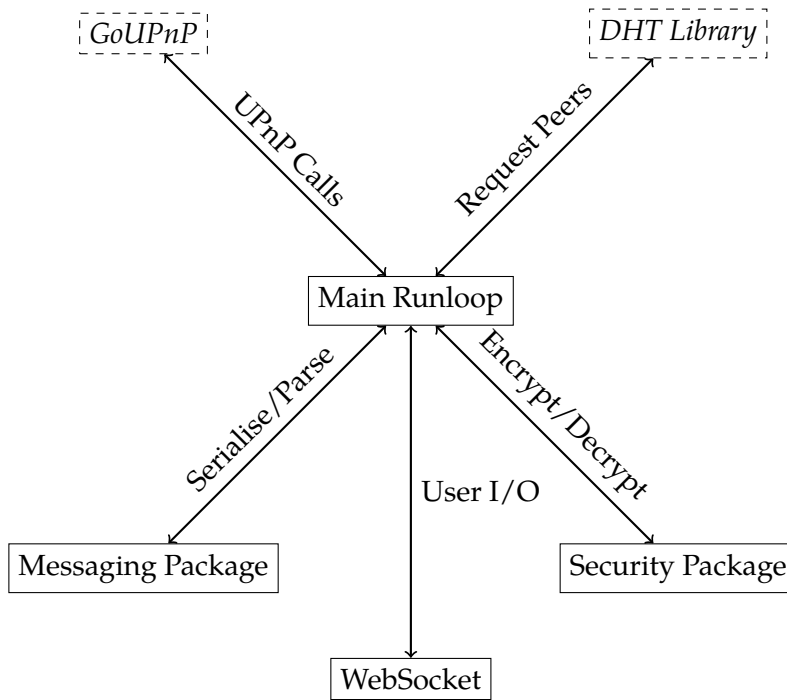


Figure 4.3.: The ultimately settled upon architecture. **Solid rectangles** represent the internal packages of the application. *Italic dashed rectangles* represent external libraries. **Arrows** represent what the **Main Runloop** exchanges with the other package or it uses them for. All OS events are handled by the **Main Runloop**

#### 4. Implementation

on GitHub<sup>6</sup>. At the time there was only one other result which was completely inactive and incompatible with BitTorrent's Mainline DHT.

The package's interface is remarkable simple. A minimal example to get up and running might Go (pun intended) something like this:

```
1 // Create a new DHT instance
2 // Arguments are:
3 //   - local port number
4 //   - target number of peers in routing table
5 //   - a boolean indicating whether to store the routing table on disk
6 d, err := NewDHTNode(6881, 100, false)
7 if err != nil {
8     fmt.Println(err)
9     return
10 }
11
12 // After checking above that there was no error, we can safely fire up the
13 // DHT's thread
14 go d.DoDHT()
15
16 // Now we can search for a particular key or infohash in BitTorrent parlance
17 // Arguments are:
18 //   - the binary infohash stored as a '\0'-terminated string
19 //   - a boolean indicating whether we are downloading the specified
20 //     infohash or just acting as a passive node
21 d.PeersRequest(string(infoHash), false)
22
23 // We can get the responding peer(s) by listening (or in this case iterating)
24 // over all the results returned by the DHT's PeersRequestResults channel
25 for infohash, peers := d.PeersRequestResults {
26     // Do something useful with the received peers
27 }
```

During my testing of the library I wished at one point to run several instances of the DHT on single machine. This prompted me to go ahead and set the port number to 0, which by convention means to let the OS auto-assign a free port. This actually revealed a bug in the library.

As well documented as the simple public API of the library is, the internal working of the library were not particularly well documented or indeed commented. This made tracking down the bug challenging until it was suggested somewhere that I use a profiling tool's call graph to gain a better understanding of the relationships between the private functions. I ran the Go profiling tool against the `find_infohash_and_wait` example that is bundled with the library's code and which is essentially the above example

---

source Go libraries. It does this by checking their code out of their respective Version Control System (VCS) when ever a user checks the project's documentation page. This is made possible by the way the Go programming language constructs package names to include the Universal Resource Locator (URL) of the VCS system that contains the code.

<sup>6</sup><https://github.com/nictuku/dht>

code wrapped in a main function. This produced the call graph in figure 4.4.

The bug came from the fact that the library did not check what port the OS returned and simply took at face value the port number the library user supplied, this meant that the library was reporting to other peers that it was reachable at port 0 which of course makes no sense. I wrote a patch, cf. appendix D.1, for the library and submitted it to Junqueira who was very encouraging and after asking me to review one change I made merged my patch into the upstream project.<sup>7</sup>

The specification of the Mainline DHT is that it announces the Transport Control Protocol (TCP) port at which the participating peer can be reached. The specification does not define what to do if a single IP address publishes more than one port number for a given IP address. Some libraries actually replace the port number in the existing entry. This is something I only realised during the implementation phase. This does unfortunately potentially introduce non-deterministic behaviour if there are multiple hosts sharing a single address.

## 4.4. Establishing a connection

As established in section 3.2, many if not most end-user's hosts are behind NAT, which must somehow be circumvented in order to enable P2P communication.

We discussed not using super-peers and therefore must map ports on routers as our only way in. Fortunately many routers support it. Unfortunately not all. This may therefore remain a theoretical exercise.

### 4.4.1. Using SSDP and UPnP to create port mappings in compatible routers

For the purposes of this project I developed a small Go library for controlling a UPnP enabled IGD as specified in [9].

When I was still thinking about using the C programming language, I happened upon the MiniUPnP library<sup>8</sup> (and MiniUPnPc UPnP port mapping client). The project includes a UPnP client to change mappings on UPnP enabled IGDs which is really quite simple. When I determined that I would write my own Go library I decided to base its API on it.

The library is called GoUPnP. It is GNU Public License Version 2 (GPLv2) licensed and available from GitHub<sup>9</sup>. In this section we will development process behind the library and associated binary, GoUPnPc (the *c* stands for client).

It is important first, to know the process by which a client creates a port forwarding rule on a UPnP enabled IGD:

1. The client uses the Simple Service Discovery Protocol (SSDP) [10] to discover the location and Simple Object Access Protocol (SOAP) endpoint of a UPnP enable

<sup>7</sup>This was very satisfying process, getting one's code merged into an upstream open-source project, I am encouraged to submit more patches to projects now.

<sup>8</sup><http://miniupnp.free.fr/> or <http://miniupnp.tuxfamily.org/>

<sup>9</sup><https://github.com/nhelke/goupnp>

## 4. Implementation

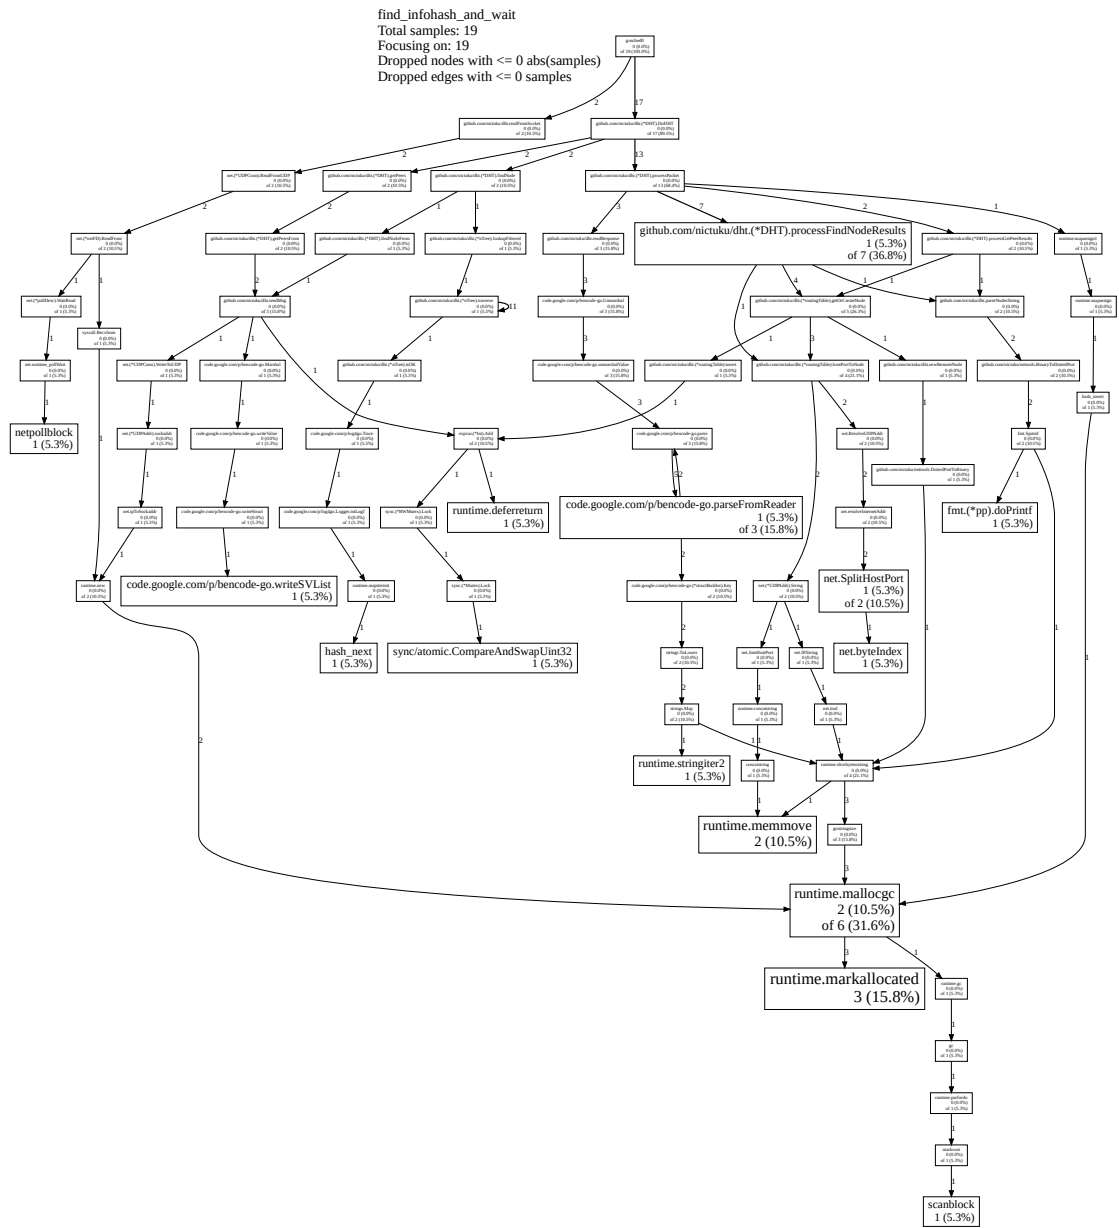


Figure 4.4.: The call graph from profiling a run of the `find_infohash_and_wait` was useful to understand the relationships between the various undocumented private functions of the library. This is hard to see on paper, but seeing as it is a vectorial graph, those with a digital copy of this document can blow it up without losing any quality.



## IGD on the Local Area Network (LAN)

- a) First the client sends a multicast HTTP Request with method M-SEARCH and a ST header indicating the type of device the client wishes to find.
  - b) The client then listens on the User Datagram Protocol (UDP) socket it used to send the search for HTTP replies from available devices matching the requested type and in particular the SOAP endpoint for further UPnP interaction.
2. Once the device has been discovered, the client can control it through SOAP calls. For port mapping a single request suffices in most cases.
- a) The client request a particular port mapping from the IGD.
  - b) The IGD replies with the mapping or an error
  - c) If an error was returned the client may try again from 2a

Before writing my own UPnP port mapping library I did search for SSDP libraries written in Go. This search produced two results, a completely undocumented package and a reasonably well documented subpackage of *go-sonos*<sup>10</sup> which looked quite promising. After failing to be able to get the example code to produce any results, closer inspection of the library's code revealed:

```
func (this *ssdpDefaultManager) ssdpIncludeNotification(msg *ssdpNotifyMessage) {
    /*TODO*/
}
```

So much for documentation then! While I was looking into patching the missing parts of the library I uncovered three other issues. Its API required, for some unknown reason, that the API-user choose which port number to use for its otherwise completely internal socket. The problem was that if the port was already occupied the API didn't simply return an error, it panicked and exited the whole program. This struck me as a rather crippling race-time condition, how can I ever be sure that the port I think is free will still be free when the API gets around to creating the socket. This issue coupled with concerns I had about the maintainability of the code which was not in its own repository, led me to implement just the parts of SSDP I needed and integrate them directly into the code of GoUPnPc.

Using a network dump of the SSDP packets generated by MiniUPnPc, its code and the standard [10] I implemented just what was necessary of SSDP in order to find the IGD and nothing more.

SSDP uses HTTP over UDP for communication. One of Go's strengths is its rich standard library which includes an HTTP package. Initially when I was writing my solution I used the facilities of the standard HTTP package to modify the used network transport. However I encountered some issues with URL handling. Most notably the standard library URL package would not accept `*` as a valid URL, rightly so, however SSDP uses `*` as a path with its M-SEARCH verb.

<sup>10</sup><https://github.com/ianr0bkny/go-sonos>

#### 4. Implementation

My initial reaction was to just work around this issue by overloading the URL package as well. Soon, however, my work-around code had surpassed the Source Lines of Code (SLOC) that hand crafting the UDP packet would require, so I switched techniques and in the process learned a valuable lesson.

We are told time and time again, all through are studies, that we should be lazy and reuse monolith libraries, which exists for almost anything you can imagine, and which, we are told, have been developed and tested by better men than we are.

I would tend to agree that such libraries often provide unparalleled quality and reliability as they are developed by specialists in that field who have a very good understanding of all the aspects and can therefore not only develop good backend code but a usually very well thought out specific APIs. Most of the time these specific APIs save the API-users from making mistakes, however they can get in your way when you are working on or with an edge-case.

These unforeseen edge-cases do not fit the standard. It would be wrong to modify the HTTP or any standard library to allow such a bizarre and wrong URL or other edge-case. So the correct solution, and the lesson learned, is to find a work-around which impacts the standard library in the least possible way and, importantly, does not depend on the standard library's internal behaviour as my initial work around accessing the encapsulated RawString field inside the standard library's URL object did.

My second attempt involved manually crafting the request, using plain old string, and sending it on a raw UDP socket, receiving the response on the raw socket and then passing that to the HTTP response parsing code from the standard library. It is a testament to the quality of the standard libraries that this was not only possible but relatively easy, in particular the accessibility of the HTTP response parsing function.

We now take a look at some actual code, corresponding to what we have just discussed. It is not the actual code used in the published version of the GoUPnP library, that can be found in function `discoverIGDDescriptionURL` in `ssdp.go` in appendix C. Instead we review some snippets based on the actual code, but cleared of some of the error checking that is done in the real code for the sake of conciseness.

First we quickly take a look at the constants used in said function and in particular we note the format string, which is used to hand craft the HTTP over UDP request.

```
const (  
    ssdpIPv4Addr = "239.255.255.250"  
    ssdpPort     = 1900  
    format       = "M-SEARCH * HTTP/1.1\r\n" +  
                  "HOST: %s:%d\r\n" +  
                  "ST: %s\r\n" +  
                  "MAN: \\"ssdp:discover"\r\n" +  
                  "MX: %d\r\n" +  
                  "\r\n"  
)
```

We recognise in the format string the headers of an HTTP request and we note the use of `*` as a URL.

This is an opportunity for us to introduce a Go language construct. The `const` parenthetical is just a shortcut to avoid repeating `const` on each line. It also works with `var` and `import`. For instance, instead of writing:

```
import "fmt"
import "net"
```

one could write:

```
import (
    "fmt"
    "net"
)
```

SSDP can only search for one device type at a time or all devices at once. Searching for all devices at once would potentially yield unreasonably many results as more and more household appliances connect to our LAN and are UPnP-enabled such as network-attached storage or smart televisions.

IGDs come in many different flavours. We have already excluded using the catch-all mechanism with SSDP. Unfortunately therefore, in order to find each and every available IGD, we must search for the following devices on the network, one at a time:

```
var deviceTypes = []string{
    "urn:schemas-upnp-org:device:InternetGatewayDevice:1",
    "urn:schemas-upnp-org:service:WANIPConnection:1",
    "urn:schemas-upnp-org:service:WANPPConnection:1",
    "upnp:rootdevice",
}
```

The last one is a catch all, the standard shouldn't require it but as MiniUPnP uses it as a last resort I decided to add it as well.

For each successive type we try to get a response:

```
1   for i := 0; i < len(deviceTypes); i++ {
2       conn, err := net.ListenUDP("udp4", allIf)
3       if err == nil {
4           // We want to timeout and move on to the next type after a couple of
5           // seconds
6           conn.SetDeadline(time.Now().Add(timeout))
7           // Send multicast request
8           conn.WriteToUDP([]byte(fmt.Sprintf(format, ssdpIPv4Addr, ssdpPort,
9               deviceTypes[i], timeout/time.Second)), broadcast)
10          // Allocate a buffer for the response
11          buf := make([]byte, 1500)
12          for {
13              // Get a response; the above timeout is still in effect as it
14              // should be, so that we never wait indefinitely on any given
15              // iteration
16              n, addr, err := conn.ReadFromUDP(buf)
17              if err != nil {
```

#### 4. Implementation

```
18         l4g.Info(err)
19         break
20     }
21     // As the http package's response parsing
22     // function requires an associated request
23     // object we let the http package parse our
24     // request as well
25     req, err := http.ReadRequest(bufio.NewReader(bytes.NewReader(
26         requestString)))
27     if err != nil {
28         // Failure to parse the request represents an assertion
29         // failure as we crafted the request ourselves and have
30         // ensured its validity
31         panic(err)
32     }
33     // This is the magic line where we use the http
34     // package to read the response
35     resp, err := http.ReadResponse(bufio.NewReader(bytes.NewReader(
36         buf[:n])), req)
37     // Interpret the response and break out of loop if successful
38     ...
39     }
40     } else {
41         l4g.Warn(err)
42     }
43 }
44 // If we get here we could not find any UPnP devices
```

As every single method in the library requires at least one round-trip to the IGD, and in order to offer API-users more flexibility in the way they use the API I decided to make all the methods asynchronous. One great thing about Go is that concurrency is build directly into the language. Go uses *channels* for communication between *green-threads* which it calls *goroutines*. These channels are very similar to Ada task entry-points, both of which are based on bounded-buffers.

We will use the `GetConnectionStatus` method as an example:

```
1 // This method fetches the status of the IGD.
2 //
3 // Errors are indicated by the channel closing before a ConnectionStatus is
4 // returned. Listeners should therefore check at the very least for nil, better
5 // still for channel closure.
6 //
7 // NOTA BENE the channel closes after a successive ConnectionStatus has been
8 // send on it, in order to not leak resources.
9 func (self *IGD) GetConnectionStatus() (ret chan *ConnectionStatus) {
10     // We initialise the channel
11     ret = make(chan *ConnectionStatus)
12
13     // We go do the work in a separate goroutine, the closure has access to the
```

```

14 // channel we just instantiated so we will be able to manipulate it.
15 go func() {
16     resp, ok := self.soapRequest("GetStatusInfo",
17         statusRequestStringReader(self.upnptype))
18     // In actual fact this goroutine is far more complicated because
19     // obtaining the IGD's status requires several soapRequests
20     // which all have to be successful
21     if ok {
22         // This again is a simplification, just know that ip is
23         // defined from traversing resp
24         ...
25         ret <- &ConnectionStatus{true, ip}
26         return
27     } else {
28         // Error handling
29         ...
30     }
31     // Note how if there is an error, the channel is closed without
32     // ever being send anything, upon success, the channel is only
33     // closed after the ConnectionStatus is consumed
34     close(ret)
35 }()
36
37 // We immediately return the channel to the caller
38 return
39 }

```

There are quite a few things that need to be noted about this code. First let us quickly review the method signature on line 9. Right after the keyword `func` and before the function name we see a parenthetical naming and specifying the object type that this method will act on. In `go` the presence of this parameter before the function name is the difference between a function and a method that can only be called on instances. After the empty parentheses indicating that this method takes no arguments we see another parenthetical, similar to the first. This is the return declaration. We could just have indicated the type without parentheses, however naming our return value is advantageous here.

It lets us assign directly to it on line 11, where we initialise the channel, and then, as it counts as a normal variable within the scope of the function, and therefore is captured by the closure, so we can freely use it without passing it in the goroutine's parameters.

Now let us see how easy it is to start a goroutine This can be seen on line 15. This is only made easier by Go's support for anonymous functions and closures. The combination of these features is used to avoid polluting the package namespace with non-reusable functions (anonymous function) and saves us the trouble of passing the channel to the goroutine as it automatically capture the calling stack frame (closure).

On line 25 we note how objects are send down a channel with the `<-` operator. Note that this is blocking meaning that we will not reach line 36, where the channel is closed, before the object is consumed by a listener on the channel.

#### 4. Implementation

On the same line the `&` has the same meaning as in C and returns a reference *in lieu* of the actual struct. What follows the ampersand is merely Go's compound literal syntax. In C this would usually result in non-deterministic behaviour as you would be returning a reference to an object on a stack frame that by its very definition is about to be popped off the stack and will no longer exist after the function returns. Go's compiler has *escape analysis* which detects what objects are leaving the scope and automatically generates the code at compile time to allocate them on the heap at runtime.<sup>11</sup>

The documentation (lines 1–8) and the comment on lines 31–33 assume some understanding of the way channels work. It has been said before that channels are bounded-buffers. Each object send on the channel must be consumed, however if there are several listeners (or consumers) one is chosen pseudo-randomly to receive the object and all others are left in the waiting state. All waiting listeners on a channel are unconditionally released upon channel closure, thereby eliminating one possible cause of deadlocking.

We now turn to the use of this function. While there are no limits to what one could do around calling this function, there are two cases I would like to review, the synchronous and the asynchronous. The use of channels in Go makes it trivial to do either.

In order to make the function blocking or synchronous, one simply needs to instantly consume the channel:

```
status := <-igd.GetConnectionStatus()
```

The `:=` is a nifty way to declare a variable and let the compiler figure out the type. The above is therefore equivalent to:

```
var status *ConnectionStatus = <-igd.GetConnectionStatus()
```

More importantly though, we must note how consuming the channel uses the same `<-` operator used to send to it, only this time the channel goes on the right of the arrow. This is actually quite intuitive.

Recall that the method returns a channel. We can also just store the reference to the channel, do something else for a bit, and consume it later, at our leisure:

```
statusChannel := igd.GetConnectionStatus()
// Do something else for a bit, the gorouting we started in GetConnectionStatus
// is working away at the same time
...
status := <-statusChannel
```

We note how using this pattern of returning a channel for long-running functions gives the API-user the choice of two very different paradigms without the overhead of having separate asynchronous and synchronous APIs.

For this reason the entire GoUPnP library uses this pattern, as you can see in the documentation available from GoDoc<sup>12</sup>.

---

<sup>11</sup>This feature is a very good example of something that Go does much better than C. This is precisely the sort of improved language feature that I did not wish to deprive myself of by using C.

<sup>12</sup><http://godoc.org/github.com/nheike/goupnpc/goupnp>

#### 4.4.2. On not implementing PCP *née* NAT-PMP

NAT Port Mapping Protocol (NAT-PMP) [5] or Port Control Protocol (PCP) [32] as it should henceforth be known turns out to be more difficult to implement—in Go at least. It is actually the far superior protocol, and is probably easier to implement than UPnP using C, as it doesn't use tedious and expensive-to-parse eXtensible Markup Language (XML).

In contrast to UPnP's two-step discovery-then-control process, PCP goes straight to step two and relies on the OS's routing table to determine the location of the IGD. This is far more sensible than using SSDP, and eliminates potential abuses or honest mistakes. The UPnP protocol does not check that the IGD it is controlling is actually an IGD being used by the current host, as it never checks the local routing table.

So while PCP may well be technically superior and a more efficient solution than UPnP, the problem is that there are no cross-platform system calls to inspect a kernel's routing table. A possible solution would be to let a platform dependent front-end communicate the discovered IGD to the platform independent backend, however this does not work with the premise of starting the backend as a daemon. Fortunately this problem can be tabled for the time being as UPnP is more prevalent. The only big name to support only PCP in their routers is Apple (they in fact authored the original NAT-PMP), however they are hardly a big player when it comes to the home-router market.


#### 4.4.3. Confronting our hypotheses with reality

After implementing the GoUPnP library we were able to start testing hypothesis from section 3.2, i.e. that communication is possible between two peers with each other in their contact lists provided one of them is publicly addressable from the internet, either directly (rare) or through port mapped NAT (more common).


Fortunately during our testing our hypotheses were largely confirmed. We did however encounter two scenarios which we had not anticipated:

1. The so-called double-NAT. To our surprise this scenario turned up more than once, we had initially discarded it, assuming it to be of marginal importance. In a bizarre twist often the NAT closest to the client supported UPnP but only mappings to the second LAN, no control over the Wide Area Network (WAN). This case is illustrated in figure 4.5(d).
2. The second case which we had not anticipated was certain corporate and higher educational institutions have firewalls that actively detect Mainline DHT traffic (and indeed other BitTorrent traffic) and completely prevent it. This case completely surprised us. We did not expect the Mainline DHT which uses a separate protocol to BitTorrent's transfer protocol to be blocked in this manner. This case is illustrated in figure 4.5(e)


#### 4. Implementation




(a) Base case where Alice and Bob are both publicly addressable. Communication is possible.




(b) Alice is behind a NAT and Bob behind a UPnP-enabled router, enabling him to open the necessary ports for Alice to contact him. Communication is possible.



(c) Alice and Bob are behind NATs. Neither NAT is UPnP-enabled. Communication is **not** possible.



(d) Alice is directly connected to the internet while and Bob is behind a double NAT where the second one is UPnP-enabled. Bob can control the UPnP-enabled IGD but this only opens ports to the second LAN. Communication is **not** possible.



(e) Alice is behind a firewall that blocks outgoing Mainline DHT packets. Communication is **not** possible. Alice doesn't know where any other peers are without access to the DHT.

Figure 4.5.: Various connection scenarios where Alice is trying to initiate a connection with Bob



## 4.5. Security

### 4.5.1. Public Key Infrastructure

As already discussed in section 3.3, Pretty Good Privacy (PGP) will be used for our Public Key Infrastructure (PKI). Fortunately there is a Go library for PGP called *openpgp*<sup>13</sup>.

One fiasco I experienced while first trying the library was that the canonical example<sup>14</sup> would panic at runtime:

```
panic: crypto: requested hash function is unavailable
```

It turns out that the *openpgp* package has a bug<sup>15</sup>. It fails to link in a required library. In Go linking is done automatically based on the packages that are imported into the application. This failure can be fixed by manually importing the missing package:

```
import _ "code.google.com/p/go.crypto/ripemd160"
```

The underscore indicates that we do not use the package ourselves, but merely wish to link it into the program. Since Go binaries are statically linked, to avoid bloating binaries unnecessarily it is a compile error to import a package and not use it. The underscore override this on per package basis.

PGP was originally supposed to be solely the handshake procedure. Ultimately it was decided to use PGP for all peer exchanges. If you recall the theoretical plan was to use Advanced Encryption Standard (AES) once the handshake was completed. AES unfortunately was too low level. It would have required rolling our own padding and boundary detecting code. PGP packets conveniently take care of all this for us already. Very convenient for passing each message on to the protocol decoder, discussed in section 4.6. A potential draw back of PGP over AES is resource requirements. During development we did not however encounter any throughput issues with using PGP instead of AES. It is possible that this might become an issue with other higher throughput low-latency media such as voice and/or video, but these cases can be handled on an *ad hoc* basis as the protocol is designed to be extendible.

### 4.5.2. Adapting the handshake to symmetric NAT initiators

As discussed in section 3.2, sometimes the initiator will be behind a NAT that cannot automatically be controlled using UPnP. In this case the initiator has no reliable way of ascertaining in a decentralized manner his public IP address before the handshake. In order to handle this case while still meeting the self-imposed requirements for the handshake, the procedure must be adapted.

We note first of all that while the initiator does not always know his/her own IP address, the receiver always knows his/her IP address either because the receiver is directly connected to the internet and has a public IP or is at least behind an automatically controllable IGD and has been able to establish this information using UPnP.

<sup>13</sup><https://code.google.com/p/go.crypto/openpgp>

<sup>14</sup><https://www.imperialviolet.org/2011/06/12/goopenpgp.html>

<sup>15</sup>[https://groups.google.com/d/topic/golang-nuts/BLkPi\\_JvMtI/discussion](https://groups.google.com/d/topic/golang-nuts/BLkPi_JvMtI/discussion)

## 4. Implementation

```
{ "name": "Nicholas Helke", "email": "nhelke@gmail.com", "key": 138 }  
&  
d4:name14:Nicholas Helke5:email16:nhelke@gmail.com3:keyi138ee
```

Figure 4.6.: A comparison of a JSON document (above) and a *bencoded* document (below) representing the same data.

Instead of supplying the receiver with our IP address we encrypt his/her address in the packet. This prevents the packet from being replayed to other hosts belonging to the same user. The handshake is now vulnerable to replay attacks within the hour targeting the same user at the same location. In order to remove this possibility, receivers must keep track of random tokens they have received and not allow the same token to be used twice within the timestamp tolerance period, i.e. one hour.

This should prevent all replay attacks, except when a malicious host intercepts the initial handshake packet and the receiver never receives it. In this case the malicious host can replay that packet from any IP address up to one hour after capture.

Another important change compared with the theoretical solution discussed in section 3.3 is not using AES with the token as a key, but instead, continuing to use PGP. This may be changed down the road, but in the mean time we are not constrained by computing resources and PGP conveniently detects individual messages and buffers them into a single readable byte slice for us.

### 4.6. Communication protocol

The terms of reference on page iii call for a platform-independent network protocol to be defined. This section discusses the design of a conforming network protocol. Initially as the only requirement is text communication, a basic protocol meeting those requirements is presented. The terms of reference also require that the protocol be extendible. In a later subsection the means by which the protocol might be extended are discussed and in particular a strategy to enable the use of UDP (necessary for voice and video for example) while the rest of the protocol is based on TCP.

#### 4.6.1. Basic protocol

In order to meet the requirement of having a platform-independent protocol, and given that the DHT protocol already uses *bencoding* for serialisation, I decided to do the same.

*Bencoding*[6, § bencoding] is the serialisation format invented by Bram Cohen for Torrent files and since used throughout the BitTorrent eco-system, including the Mainline DHT. It is far more compact than XML and in some ways easier to parse than JavaScript Object Notation (JSON), as strings are prepended with their size, enabling the parser to dynamically allocate the correct size items and not have to worry about resizing them. It is however less human readable as we can see in figure 4.6

The theoretical advantage of *bencoding* over JSON may not actually hold up in the real world. There are far more JSON libraries than there are *bencode* ones and I would not be surprised if all that competition has produced faster JSON libraries (relative to *bencode*). Fortunately should we ever decide to make such a change it can be done in a backwards compatible way as all the protocol's messages are dictionaries and the starting character of a dictionary in *bencode* is "d" whereas JSON uses "{".

We have just revealed in passing that all the messages of the protocol are *bencoded* dictionaries. There are five different message types. The fields contained in the various dictionaries of the different message types are documented below:

**SYN** This is the initiating message of the handshake procedure. Its fields are:

- token** A cryptographically secure random string
- dest** The destination IP address in decimal-dotted format
- timestamp** The current epoch time as an integer

**SYNACK** This is the response to a handshake challenge.

- token** The token received from the initiator in its *SYN*-message
- source** The source IP address in decimal-dotted format
- timestamp** The current epoch time as an integer

**Ping** This is a ping request. It can double as an *ACK* for the handshake procedure. Its fields are:

- token** A random token (to track the response)

**Pong** This is a ping response.

- token** The token received in the *ping* request

**Text Message** This is a text message. Its only field is:

- txtMsg** The message contents

#### 4.6.2. Extending the protocol

The use of dictionaries makes the protocol easily extendible simply by ensuring that new message types have dictionary keys that differ from the above described ones. This lets anyone develop extensions to the protocol.

In order to enable organic enhancements to the protocol by third parties, we specify that clients must ignore message types they do not recognize and need not reply to them. It is the responsibility of protocol extenders to develop the message-types and message-passing diagram necessary in order to establish that a client supports a given extension. Additionally, in order to better guard ourselves against accidental message-type conflicts, protocol extenders should prepend their dictionary keys with something unique to them, such as a domain name they control.

#### 4. Implementation

The terms of reference (page iii) are very clear, the protocol must support arbitrary communication including but not limited to voice and video. Unfortunately the above does not work for real time voice and video as it uses TCP, whereas real time multimedia needs UDP.

Another potential advantage of UDP is that it works with full cone NATs. Full cone NATs are defined in [20] to be NATs that accept incoming UDP packets from any host to a translated address of a host behind the NAT once at least one packet has been send out, thereby establishing the translation. Additionally Saikat Guha and Paul Francis estimate in [13] 70.1%' of the world's NATs to be full cone. This means that if we could use UDP instead of TCP even for the basic protocol and text, we would be able to achieve better penetration than the penetration discussed in section 4.4, as even routers not supporting UPnP could now serve as endpoints.

Our control protocol however is based on TCP and therefore makes the assumption that its transport is reliable. In order to use this protocol over UDP either the protocol would have to be changed to build handle retransmits or—

It turns out that Torrent clients, which suffer from similar issues when running in fully decentralised mode, have already solved this question by developing Micro Transport Protocol ( $\mu$ TP)<sup>16</sup>.  $\mu$ TP is a reliable transport layer build on top of UDP which provides higher level protocols with a means to send traffic reliably over UDP.

As the Mainline DHT only reports a single port number for a given IP address,  $\mu$ TP connections are normally accepted on the same port as DHT communication. This has the added bonus of reducing the number of ports that need to be mapped to the outside world. It does however mean that the two distinct protocols, KRPC and  $\mu$ TP, must be multiplexed.

The  $\mu$ TP protocol header (version 1) is illustrated in figure 4.7. KRPC is defined as a bencoded dictionary which means that every KRPC packet starts with a single ASCII-coded 'd'. Figure 4.8 shows how irrespective of the type of  $\mu$ TP packet, the last four bits of the first byte do not match between the two protocols and will not until a version 4 is released, at which time one could easily choose to skip that value to avoid ambiguity. This means that is is indeed possible to multiplex the two protocols.

This nonetheless would require substantive modifications to the DHT library to enable multiplexing and demultiplexing of the two protocols. Two possibilities were examined:

1. The first would be to patch the DHT library to delegate network communication to the calling application. This it turns out is pretty unrealistic as the entire library was based on the idea of it handling its own socket and in particular throttling it as necessary to defend against bad nodes.
2. The second solution is to add an optional handler that the library can call if it doesn't recognize the incoming packet as starting with a 'd'. This solution is probably the way to go as it would not only be easier to implement but also, unlike the former, it does not require that the API be changed and therefore stands a better chance of being merged into the upstream project.

---

<sup>16</sup>Originally known as  $\mu$ Torrent Transport Protocol in the seminal paper[24]




Figure 4.7.:  $\mu$ TP version 1 header defined in [24]




Figure 4.8.: Comparison of the first 8 bits of  $\mu$ TP version 1 and KRPC

Extending the application to also support incoming connections when behind a full cone NAT would require either the handshake to be modified one more, as a host behind a full cone NAT does not know its public IP address and therefore cannot protect itself adequately against replay attacks. Alternatively the handshake can be left unchanged, however a means must be found for this host to establish reliably its public IP address. As we have already discussed in section 3.2 this is far from trivial to achieve in a fully decentralised manner.

In any event I think the case for the extensibility of this project’s protocol has been successfully made.

#### *4. Implementation*

## 5. Conclusion

My work on this thesis was divided into two parts. A first part, in parallel with classes, during which I explored almost every and any avenue remotely relevant to my subject. As is to be expected, I met with dead-ends along the way. This is the general idea of discovery phase and means that during the second part of the project, the part that is spent full time on the project, the relevant paths are known, and can be explored to greater depths while simultaneously finishing the accompanying project.

In finishing up a project, the conclusion gives one a valuable opportunity to look back and draw lessons from things done right and things done wrong. We can ponder the “what ifs” and what we might do differently, were we ever to redo the project.

The theoretical aspects of this project are quite solid I would argue. It is clearly possible in theory to develop a completely distributed communication system over Internet Protocol (IP). Additionally the solution discussed in this paper would, in theory, have helped in at least Egypt and Libya. In developing the theory however it became clear to me that this is merely one round in the constant game of cat and mouse and that it is very likely that means could be found to disrupt this theoretical solution. Perhaps the most alarming realisation was the extent to which even western governments control and censor the internet. This leads one to wonder whether the internet can be relied on at all.

As I was concluding this work, in the last week, it dawned on me how much of my work is based theory or standards or builds on elements from the BitTorrent ecosystem. This nonetheless leads me to my only theoretical “what if”. What if I had known of or discovered during the course of this project other, different ecosystems that I might have built on. Bitcoin seems to be inspiring other people to come up with new decentralised systems such as Bitmessage or Namecoin, both mentioned in chapter 2. I still maintain that those systems do not lend themselves to real time communication, however they do seem to be better at traversing Network Address Translation (NAT).

On the implementation front I completed every requirement of the terms of reference. That being said my hope was always to produce widely usable solution, one which would essentially run itself and provide better overall reliability than any existing server based solutions. The discovery along the way that NAT is not nearly as easily traversed as I had hoped, was very disappointing. This means that the overall usability of the solution is severely limited.

While it is true that the solution proposed in this paper is, under certain circumstances (such as those in Egypt and Libya during the Arab Spring), more resilient than server based solutions, users will not judge it on this basis.

In photographers’ circles it is oft said that “the best camera is the one you have with you.” This adage perfectly summarises the issue with the solution proposed in this

## 5. Conclusion

paper. Most users are not concerned with elegant technical solutions to rare edge-cases. Even if a solution provides unparalleled access in such cases, users are more concerned with whether a solution works day to day for them.

It does not help that the main case where my solution cannot receive connections is from behind corporate firewalls. It is in these circumstances most people spend the majority time on their computers. Telling them that this solution will work between homes if their country or region's internet was ever, all of a sudden, to become detached from the rest of the internet, is not sufficient to garner their interest.

I truly valued this opportunity to work on my own on a project. It was far from easy working on alone, particularly after my experiences working in teams on various projects all through my studies. A lesson I will have learned from this project is that it is easier to work in teams. In a team, if in doubt about a theoretical or implementation decision, you can always bounce ideas off another team member. Working alone, I had to justify every decision I made, using other people's works, to make sure that I was not going adrift. Although I also feel a certain sense of satisfaction and pride at having successfully carried this project to fruition alone.



## Bibliography

- [1] Egypt arrests as undersea internet cable cut off Alexandria. <http://www.bbc.co.uk/news/world-middle-east-21963100>, March 2013.
- [2] Charles Arthur. Undersea internet cables off Egypt disrupted as navy arrests three. <http://www.guardian.co.uk/technology/2013/mar/28/egypt-undersea-cable-arrests>, March 2013.
- [3] Scott Bradner. Key words for use in RFCs to Indicate Requirement Levels. RFC 2119 (Best Current Practice), March 1997.
- [4] Dave Cheney. An introduction to cross compilation with Go. <http://dave.cheney.net/2012/09/08/an-introduction-to-cross-compilation-with-go>, September 2012.
- [5] S. Cheshire and M. Krochmal. NAT Port Mapping Protocol (NAT-PMP). RFC 6886 (Informational), April 2013.
- [6] Bram Cohen. The BitTorrent protocol specification. [http://www.bittorrent.org/beps/bep\\_0003.html](http://www.bittorrent.org/beps/bep_0003.html), January 2008.
- [7] Scott A. Crosby and Dan S. Wallach. An analysis of BitTorrent's two Kademlia-based DHTs, 2007.
- [8] Alberto Dainotti, Claudio Squarcella, Emile Aben, Kimberly C. Claffy, Marco Chiesa, Michele Russo, and Antonio Pescapé. Analysis of country-wide internet outages caused by censorship. In *Proceedings of the 2011 ACM SIGCOMM Conference on Internet Measurement Conference, IMC '11*, pages 1–18, New York, NY, USA, 2011. ACM.
- [9] UPnP Forum. WANIPConnection:1. Service template version 1.01, November 2001. Standardized DCP.
- [10] UPnP Forum. UPnP. Device architecture version 1.1, October 2008.
- [11] Dan Goodin. Skype replaces P2P supernodes with Linux boxes hosted by Microsoft. <http://arstechnica.com/business/2012/05/skype-replaces-p2p-supernodes-with-linux-boxes-hosted-by-microsoft/>, May 2012.
- [12] Glenn Greenwald. NSA collecting phone records of millions of Verizon customers daily. <http://www.guardian.co.uk/world/2013/jun/06/nsa-phone-records-verizon-court-order>, June 2013.

## Bibliography

- [13] Saikat Guha and Paul Francis. Characterization and measurement of TCP traversal through NATs and firewalls. In *Proceedings of the 5th ACM SIGCOMM conference on Internet Measurement, IMC '05*, pages 18–18, Berkeley, CA, USA, 2005. USENIX Association.
- [14] D. Harkins and D. Carrel. The Internet Key Exchange (IKE). RFC 2409 (Proposed Standard), November 1998. Obsoleted by RFC 4306, updated by RFC 4109.
- [15] Erik Hjelmvik and Wolfgang John. Breaking and improving protocol obfuscation. Technical Report 2010-05, Department of Computer Science and Engineering, Chalmers University of Technology, 2010. ISSN 1652-926X.
- [16] Jonathan A. Kelner and Petar Maymounkov. Electric routing and concurrent flow cutting. *CoRR*, abs/0909.2859, 2009.
- [17] T. Kivinen and M. Kojo. More Modular Exponential (MODP) Diffie-Hellman groups for Internet Key Exchange (IKE). RFC 3526 (Proposed Standard), May 2003.
- [18] Chris Lesniewski-Laas and M. Frans Kaashoek. Whānau: a sybil-proof distributed hash table. In *Proceedings of the 7th USENIX conference on Networked systems design and implementation, NSDI'10*, pages 8–8, Berkeley, CA, USA, 2010. USENIX Association.
- [19] Andrew Loewenstern. BitTorrent enhancement proposal #5: DHT protocol. [http://www.bittorrent.org/beps/bep\\_0005.html](http://www.bittorrent.org/beps/bep_0005.html), January 2008.
- [20] R. Mahy, P. Matthews, and J. Rosenberg. Traversal Using Relays around NAT (TURN): Relay Extensions to Session Traversal Utilities for NAT (STUN). RFC 5766 (Proposed Standard), April 2010.
- [21] Petar Maymounkov and David Mazières. Kademia: A peer-to-peer information system based on the xor metric. In *Revised Papers from the First International Workshop on Peer-to-Peer Systems, IPTPS '01*, pages 53–65, London, UK, UK, 2002. Springer-Verlag.
- [22] Declan McCullagh. NSA docs boast: Now we can wiretap Skype video calls. [http://news.cnet.com/8301-13578\\_3-57593339-38/](http://news.cnet.com/8301-13578_3-57593339-38/), July 2013.
- [23] Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system, May 2009.
- [24] Arvid Norberg. BitTorrent enhancement proposal #29: uTorrent transport protocol. [http://www.bittorrent.org/beps/bep\\_0029.html](http://www.bittorrent.org/beps/bep_0029.html), June 2009.
- [25] Sylvia Ratnasamy, Paul Francis, Mark Handley, Richard Karp, and Scott Shenker. A scalable content-addressable network. In *Proceedings of the 2001 conference on Applications, technologies, architectures, and protocols for computer communications, SIGCOMM '01*, pages 161–172, New York, NY, USA, 2001. ACM.

- [26] J. Rosenberg, R. Mahy, P. Matthews, and D. Wing. Session Traversal Utilities for NAT (STUN). RFC 5389 (Proposed Standard), October 2008.
- [27] Antony I. T. Rowstron and Peter Druschel. Pastry: Scalable, decentralized object location, and routing for large-scale peer-to-peer systems. In *Proceedings of the IFIP/ACM International Conference on Distributed Systems Platforms Heidelberg, Middleware '01*, pages 329–350, London, UK, UK, 2001. Springer-Verlag.
- [28] Stefan Saroiu, P. Krishna Gummadi, and Steven D. Gribble. A measurement study of peer-to-peer file sharing systems. Technical Report UW-CSE-01-06-02, University of Washington, July 2001.
- [29] R. Schollmeier. A definition of peer-to-peer networking for the classification of peer-to-peer architectures and applications. In *Peer-to-Peer Computing, 2001. Proceedings. First International Conference on*, pages 101–102, 2001.
- [30] Ion Stoica, Robert Morris, David Karger, M. Frans Kaashoek, and Hari Balakrishnan. Chord: A scalable peer-to-peer lookup service for internet applications. In *Proceedings of the 2001 conference on Applications, technologies, architectures, and protocols for computer communications, SIGCOMM '01*, pages 149–160, New York, NY, USA, 2001. ACM.
- [31] Jonathan Warren. Bitmessage: A peer-to-peer message authentication and delivery system. <http://www.bitmessage.org>, November 2012.
- [32] D. Wing, S. Cheshire, M. Boucadair, R. Penno, and P. Selkirk. Port Control Protocol (PCP). RFC 6887 (Proposed Standard), April 2013.
- [33] Charles V. Wright, Lucas Ballard, Scott E. Coull, Fabian Monrose, and Gerald M. Masson. Uncovering spoken phrases in encrypted voice over IP conversations. *ACM Trans. Inf. Syst. Secur.*, 13(4):35:1–35:30, December 2010.
- [34] Ben Y. Zhao, Ling Huang, Jeremy Stribling, Sean C. Rhea, Anthony D. Joseph, and John D. Kubiatowicz. Tapestry: A resilient global-scale overlay for service deployment. *IEEE Journal on Selected Areas in Communications*, 22:41–53, 2004.

## *Bibliography*

## Acronyms

- μTP** Micro Transport Protocol. 44
- AES** Advanced Encryption Standard. 19, 41, 42
- API** Application Programming Interface. 24, 30, 31, 33, 34, 36, 38, 44
- CUPS** Common Unix Printing System. 27
- DDoS** Distributed Denial of Service. 8, 11, 14, 15
- DHT** Distributed Hash Table. i, 8–10, 15, 16, 18, 20, 23, 27, 30, 31, 39, 42, 44
- DNS** Domain Name System. 4, 14
- EFF** Electronic Frontier Foundation. 4
- GPLv2** GNU Public License Version 2. 31, 51
- GUI** Graphical User Interface. iv, 27
- HTML** Hyper-Text Markup Language. 27
- HTTP** Hyper-Text Transport Protocol. 27, 33, 34
- ICANN** Internet Corporation for Assigned Names and Numbers. 4
- IGD** Internet Gateway Device. 14, 23, 31, 33, 35, 36, 39, 41
- IP** Internet Protocol. i, iii, 1, 4, 7, 14–16, 18, 19, 23, 31, 41–45, 47
- IPv4** Internet Protocol Version 4. iii, 14
- IPv6** Internet Protocol Version 6. 14
- ISP** Internet Service Provider. 1, 16, 17
- JS** JavaScript. 27
- JSON** JavaScript Object Notation. 42, 43
- JVM** Java Virtual Machine. 26

## *Acronyms*

- LAN** Local Area Network. 33, 35, 39
- MSE** Message Stream Encryption. 17
- NAT** Network Address Translation. iii, 7, 14, 23, 31, 39, 41, 44, 45, 47
- NAT-PMP** NAT Port Mapping Protocol. 39
- NDK** Native Development Kit. 26
- NIC** Network Information Centre. 4
- NSA** National Security Agency. 16
- OO** Object Oriented. 28
- OS** Operating System. 27–31, 39
- P2P** Peer-to-Peer. iii, 1, 3, 7–9, 12, 14–16, 21, 24, 31
- PCP** Port Control Protocol. 14, 39
- PGP** Pretty Good Privacy. 20, 41, 42
- PKI** Public Key Infrastructure. 41
- RFC** Request for Comments. 17
- RPC** Remote Procedure Call. 13
- SLOC** Source Lines of Code. 34
- SOAP** Simple Object Access Protocol. 31, 33
- SSDP** Simple Service Discovery Protocol. 31, 33, 35, 39
- STUN** Session Traversal Utilities for NAT. 14, 15
- TCP** Transport Control Protocol. 31, 42, 44
- TURN** Traversal Using Relays around NAT. 15
- UDP** User Datagram Protocol. 12, 33, 34, 42, 44
- UPnP** Universal Plug and Play. 14, 23, 31, 33, 35, 39, 41, 44
- URL** Universal Resource Locator. 30, 33, 34
- UTC** Coordinated Universal Time. 18

**VCS** Version Control System. 30

**VoIP** Voice over IP. 14, 15, 17

**WAN** Wide Area Network. 39

**WoT** Web of Trust. 19

**XML** eXtensible Markup Language. 39, 42

## *Acronyms*



## Appendix A.

# Dictator Breaker User Manual

### A.1. How to run

You run Dictator Breaker by calling its binary from the terminal (on some platforms it may also be possible to double-click the binary and a terminal will launch automatically).

The program will spit out some output among which a URL where the user interface can be found and which must be viewed in order to proceed.

Once the interface has been loaded in the browser, you will be prompted to enter the email address and passphrase of the private key from your local GPG keyring. Once you have done this, you should be able to follow in the terminal as the backend attempts to connect to your GPG contacts who are also using Dictator Breaker.

New connections and messages should appear in your browser window.

In order to send a message to an active connection prefix it with the connection identifier reported on the webpage when the connection was initialised.

Close the program either by closing the browser window or by closing the terminal.

### A.2. Command line flags

`-cleanupPeriod=15m0s`: How often to ping nodes in the network to see if they are reachable.

`-http=8080`: port from which to serve the HTML/JS GUI on the loopback interface

`-maxNodes=500`: Maximum number of nodes to store in the routing table, in memory. This is the primary configuration for how noisy or aggressive this node should be. When the node starts, it will try to reach `maxNodes/2` as quick as possible, to form a healthy routing table.

`-port=0`: use this port for the DHT and communication (UDP and TCP), 0 means let the operating system choose a random available port

`-rateLimit=100`: Maximum packets per second to be processed. Beyond this limit they are silently dropped. Set to -1 to disable rate limiting.

`-routers="1.a.magnets.im:6881,router.utorrent.com:6881"`: Comma separated IP:Port address of the DHT routeirs used to bootstrap the DHT network.

`-savePeriod=5m0s`: How often to save the routing table to disk.

`-storeDHT=true`: whether to store and load the DHT routing table from/to disk



## Appendix B.

### Dictator Breaker Code

For archival and review purposes, the source of Dictator Breaker is included below as it stood at submission time:

#### B.1. [go.nhelke.com/dictator-breaker/main.go](https://github.com/nhelke/dictator-breaker)

```
1 package main
2
3 import (
4     "code.google.com/p/go.crypto/openpgp"
5     l4g "code.google.com/p/log4go"
6     "flag"
7     "fmt"
8     "github.com/nhelke/dht"
9     "github.com/nhelke/goupnpc/goupnp"
10    "go.nhelke.com/dictator-breaker/messaging"
11    "go.nhelke.com/dictator-breaker/security"
12    "math/rand"
13    "net"
14    "net/http"
15    "os"
16    "strings"
17    "time"
18 )
19
20 const (
21     targetNumPeers = 60
22 )
23
24 var (
25     port = flag.Int("port", 0,
26         "use this port for the DHT and communication (UDP and TCP), "+
27         "0 means let the operating system choose a random available port")
28     storeDHTRoutingTable = flag.Bool("storeDHT", true,
29         "whether to store and load the DHT routing table from/to disk")
30     interfacePort = flag.Int("http", 8080,
31         "port from which to serve the HTML/JS GUI on the loopback interface")
32 )
```

## Appendix B. Dictator Breaker Code

```
33     connections = make([]*conn, 0, 200)
34     newConnection = make(chan *conn)
35     expConnection = make(chan *conn)
36     newIncMessage = make(chan messaging.Message)
37     newOutMessage = make(chan messaging.Message)
38     me            string
39     ws            *wsConn
40 )
41
42 type conn struct {
43     net      net.Conn
44     contact  *openpgp.Entity
45     in       []string
46     out      []string
47     send     chan map[string]interface{}
48     closed   chan struct{}
49 }
50
51 const (
52     txtMessageKey = "txtMsg"
53     pingKey       = "ping"
54     pongKey       = "pong"
55 )
56
57 // This function should be called the thread that handled the new connection,
58 // i.e. either accept's or dial's thread, and it only returns after the
59 // connection has dropped. It manages its its addition and removal from the
60 // connection slice itself.
61 func (c conn) Run() {
62     c.send = make(chan map[string]interface{})
63     c.closed = make(chan struct{})
64
65     newConnection <- &c
66     defer func() { expConnection <- &c }()
67
68     go func() {
69         for {
70             select {
71                 case pack := <-c.send:
72                     err := security.EncryptBencode(c.net, c.contact, pack)
73                     if err != nil {
74                         l4g.Warn(err)
75                         close(c.closed)
76                         return
77                     }
78
79                 case _, open := <-c.closed:
80                     if !open {
81                         c.net.Close()
82                     }
83             }
84         }
85     }
86 }
```

```

82         return
83     }
84 }
85 }
86 }()
87
88 for {
89     pack, signer, err := security.DecryptBencode(c.net)
90     if err != nil {
91         l4g.Warn(err)
92         close(c.closed)
93         return
94     } else if signer != c.contact {
95         l4g.Warn("Unexpected signer (%v) on connection %v@%v", signer,
96             c.contact, c.net.RemoteAddr())
97         close(c.closed)
98         return
99     }
100
101     if msg, ok := pack[txtMessageKey]; ok {
102         c.in = append(c.in, msg.(string))
103         ws.out <- fmt.Sprintf("%s says > %s", signer, msg)
104     }
105     if ping, ok := pack[pingKey]; ok {
106         pack := map[string]interface{}{
107             pongKey: ping,
108         }
109         c.send <- pack
110     }
111     if _, ok := pack[pongKey]; ok {
112         l4g.Warn("Sucessfully pinged %v@%v", c.contact, c.net.RemoteAddr())
113     }
114 }
115 }
116
117 func (c conn) SendMessage(msg string) {
118     pack := map[string]interface{}{
119         txtMessageKey: msg,
120     }
121     c.send <- pack
122 }
123
124 func (c conn) Ping() {
125     pack := map[string]interface{}{
126         txtMessageKey: security.RandomToken(),
127     }
128     c.send <- pack
129 }
130

```

## Appendix B. Dictator Breaker Code

```
131 func main() {
132     // Remove this line in production
133     l4g.NewDefaultLogger(l4g.DEBUG)
134
135     flag.Parse()
136     if len(flag.Args()) > 0 {
137         flag.PrintDefaults()
138         os.Exit(1)
139     }
140
141     igd := <-goupnp.DiscoverIGD()
142     if igd == nil {
143         l4g.Warn("No UPnP IGD found")
144     } else {
145         if *port == 0 {
146             *port = rand.Intn(65536-1024) + 1024
147         }
148         l4g.Warn("Using port %v", *port)
149         tcpPortMap := <-igd.AddLocalPortRedirection(uint16(*port), goupnp.TCP)
150         udpPortMap := <-igd.AddLocalPortRedirection(uint16(*port), goupnp.UDP)
151         if tcpPortMap == nil || udpPortMap == nil {
152             l4g.Error("Unable to map ports")
153         }
154         status := <-igd.GetConnectionStatus()
155         if status != nil {
156             me = status.IP.String()
157         }
158     }
159
160     httpPort := fmt.Sprintf(":%d", *interfacePort)
161     go http.ListenAndServe(httpPort, nil)
162     l4g.Warn("Listening on http://localhost%s", httpPort)
163
164     for {
165         connChan := make(chan *wsConn)
166         connCC <- connChan
167
168         ws = <-connChan
169
170         ws.out <- "GPG Email:"
171         user := <-ws.in
172         ws.out <- "GPG Password:"
173         pass := <-ws.in
174
175         err := security.ReadKeyring(user, pass)
176         if err != nil {
177             l4g.Warn(err)
178             ws.ws.Close()
179             continue

```

```

180     } else {
181         break
182     }
183 }
184
185 d, err := dht.NewDHTNode(*port, targetNumPeers, *storeDHTRoutingTable)
186 if err != nil {
187     l4g.Critical(err)
188     os.Exit(1)
189 }
190
191 tick := time.Tick(1 * time.Minute)
192 for {
193     select {
194     case infoHashPeers := <-d.PeersRequestResults:
195         for ih, peers := range infoHashPeers {
196             if len(peers) > 0 {
197                 for _, peer := range peers {
198                     addr := dht.DecodePeerAddress(peer)
199                     l4g.Debug("peer found for infohash %x@%v", addr, ih)
200                     dial(string(ih), addr)
201                 }
202             }
203         }
204     case newConnection := <-newConnection:
205         connections = append(connections, newConnection)
206     case expConnection := <-expConnection:
207         for i := 0; i < len(connections); i++ {
208             if connections[i] == expConnection {
209                 connections[i] = connections[len(connections)-1]
210                 connections = connections[0 : len(connections)-1]
211                 break
212             }
213         }
214     case <-tick:
215         for _, c := range security.Contacts() {
216             ih := security.FingerprintToString(c.PrimaryKey.Fingerprint)
217             d.PeersRequest(ih, me != "" && ih == security.MyFingerprint())
218         }
219     }
220 }
221 }
222
223 // This should be called on a separate thread
224 func listen(port int) {
225     ln, err := net.Listen("tcp", fmt.Sprintf(":%d", port))
226     if err != nil {
227         l4g.Critical("Failed to open TCP socket on same port as UDP", err)
228         os.Exit(5)

```

## Appendix B. Dictator Breaker Code

```
229     }
230     for {
231         conn, err := ln.Accept()
232         if err != nil {
233             l4g.Error(err)
234             continue
235         }
236         go accept(conn)
237     }
238 }
239
240 // This should be called on a separate thread
241 func accept(c net.Conn) {
242     l4g.Debug("Accepted connection from %v", c.RemoteAddr())
243
244     reply, signer, err := security.DecryptBencode(c)
245     if err != nil {
246         l4g.Warn(err)
247         c.Close()
248         return
249     }
250
251     token := reply["token"]
252
253     if me != reply["dest"] ||
254        !security.WithinOneHour(reply["timestamp"].(int)) {
255         l4g.Warn(err)
256         c.Close()
257         return
258     }
259
260     synack := map[string]interface{}{
261         "token":    token,
262         "timestamp": time.Now().Unix(),
263         "source":   me,
264     }
265
266     err = security.EncryptBencode(c, signer, synack)
267     if err != nil {
268         l4g.Warn(err)
269         c.Close()
270         return
271     }
272
273     conn{
274         net:    c,
275         contact: signer,
276     }.Run()
277 }
```



```

278
279 // This should be called on a separate thread
280 func dial(ih, addr string) {
281     conn, err := net.Dial("tcp", addr)
282     if err != nil {
283         l4g.Debug(err)
284         conn.Close()
285         return
286     }
287
288     token := security.RandomToken()
289     remote := strings.Split(addr, ":")[0]
290     syn := map[string]interface{}{
291         "token":    token,
292         "dest":     remote,
293         "timestamp": time.Now().Unix(),
294     }
295
296     err = security.EncryptBencode(conn, security.GetKeyByFingerprint(ih), syn)
297     if err != nil {
298         l4g.Warn(err)
299         conn.Close()
300         return
301     }
302
303     synack, signer, err := security.DecryptBencode(conn)
304     if err != nil {
305         l4g.Warn(err)
306         conn.Close()
307         return
308     }
309
310     if ih != security.FingerprintToString(signer.PrimaryKey.Fingerprint) ||
311        synack["token"] != token ||
312        synack["source"] != remote ||
313        !security.WithinOneHour(synack["timestamp"].(int)) {
314        l4g.Warn("Bad SYNACK: %v", synack)
315    }
316
317 }

```

## B.2. go.nhelke.com/dictator-breaker/ws.go

```

1 package main
2
3 import (
4     "fmt"
5     "net/http"

```

## Appendix B. Dictator Breaker Code

```
6
7     "code.google.com/p/go.net/websocket"
8     l4g "code.google.com/p/log4go"
9 )
10
11 var (
12     connCC      = make(chan chan *wsConn)
13     incomingWs = make(chan string)
14 )
15
16 type userPass struct {
17     user, pass string
18 }
19
20 type wsConn struct {
21     ws *websocket.Conn
22     in  <-chan string
23     out chan<- string
24 }
25
26 func APISocket(ws *websocket.Conn) {
27     select {
28     case connChan := <-connCC:
29         l4g.Debug("Accepting WebSocket connection from %v", ws.RemoteAddr())
30
31         var (
32             in  = make(chan string)
33             out = make(chan string)
34         )
35
36         go func() {
37             for {
38                 var line string
39                 _, err := fmt.Fscanln(ws, &line)
40                 if err != nil {
41                     l4g.Error(err)
42                     break
43                 }
44                 in <- line
45             }
46         }()
47
48         connChan <- &wsConn{
49             ws: ws,
50             in: in,
51             out: out,
52         }
53
54         for {
```

```

55         select {
56             case str := <-out:
57                 fmt.Fprintln(ws, str)
58         }
59     }
60     default:
61         fmt.Fprintln(ws, "There is already an active session connected to this backend")
62         ws.Close()
63     }
64 }
65
66 func init() {
67     http.Handle("/api", websocket.Handler(APISocket))
68     http.Handle("/", http.FileServer(http.Dir("views")))
69 }

```

### B.3. go.nhelke.com/dictator-breaker/messaging/message.go

```

1  package messaging
2
3  import (
4      "encoding/json"
5      "go.nhelke.com/dictator-breaker/security"
6      "io"
7      "net"
8  )
9
10 type Message struct {
11     conn *net.Conn `json:"-"`
12
13     Type string
14     Token string
15
16     Body map[string]interface{} `json:",omitempty"`
17 }
18
19 func (m *Message) WriteTo(w io.Writer) (err error) {
20     b, err := json.Marshal(m)
21     if err == nil {
22         _, err = w.Write(b)
23     }
24     return
25 }
26
27 func Ping() (m Message) {
28     m.Type = "ping"
29     m.Token = security.RandomToken()
30     return

```

## Appendix B. Dictator Breaker Code

```
31 }
32
33 func TextMessage(text string) (m Message) {
34     m.Type = "textMessage"
35     m.Token = security.RandomToken()
36     m.Body = make(map[string]interface{})
37     m.Body["txt"] = text
38     return
39 }
40
41 func (m *Message) Ack() (ack Message) {
42     ack = *m
43     ack.Type = "ack"
44     ack.Body = nil
45     return
46 }
```

### B.4. [go.nhelke.com/dictator-breaker/messaging/message\\_test.go](https://go.nhelke.com/dictator-breaker/messaging/message_test.go)

```
1 package messaging
2
3 import (
4     "bytes"
5     "testing"
6 )
7
8 func TestMessageMarshalling(t *testing.T) {
9     msg := Ping()
10    b := new(bytes.Buffer)
11    msg.WriteTo(b)
12    str := string(b.Bytes())
13    t.Log(str)
14    if str != `{"Type":"ping","Token":"4"}` {
15        t.Fail()
16    }
17 }
```

### B.5. [go.nhelke.com/dictator-breaker/security/openpgp.go](https://go.nhelke.com/dictator-breaker/security/openpgp.go)

```
1 package security
2
3 import (
4     bencode "code.google.com/p/bencode-go"
5     "code.google.com/p/go.crypto/openpgp"
6     l4g "code.google.com/p/log4go"
7     "crypto/rand"
8     "errors"
```

```

9     "io"
10    "os"
11    "path"
12    "strings"
13    "time"
14    // This is needed by openpgp, failure to link it into the binary
15    // results in cryptic panics at runtime
16    _ "code.google.com/p/go.crypto/ripemd160"
17 )
18
19 var (
20     alice    *openpgp.Entity
21     pubring  openpgp.EntityList
22 )
23
24 func ReadKeyring(email, password string) error {
25     secringFile, err := os.Open(GPGFolderPath("secring.gpg"))
26     if err != nil {
27         return err
28     }
29
30     secring, err := openpgp.ReadKeyRing(secringFile)
31     if err != nil {
32         return err
33     }
34
35     alice = getKeyByEmail(secring, email)
36     if alice == nil {
37         return errors.New("Requested private key not found")
38     }
39
40     err = alice.PrivateKey.Decrypt([]byte(password))
41     if err != nil {
42         return err
43     }
44
45     pubringFile, err := os.Open(GPGFolderPath("pubring.gpg"))
46     if err != nil {
47         return err
48     }
49
50     pubring, err = openpgp.ReadKeyRing(pubringFile)
51     if err != nil {
52         return err
53     }
54
55     pubring = append(pubring, alice)
56
57     return nil

```

## Appendix B. Dictator Breaker Code

```
58 }
59
60 func GPGFolderPath(filename string) (dir string) {
61     env := os.Environ()
62     for _, e := range env {
63         if strings.HasPrefix(e, "HOME=") {
64             dir = strings.SplitN(e, "=", 2)[1]
65             dir = path.Join(dir, ".gnupg")
66             path.Join(dir, filename)
67             break
68         }
69     }
70     return
71 }
72
73 // Shamelessly taken from the seemingly canonical OpenPGP example for Go at
74 // https://www.imperialviolet.org/2011/06/12/goopenpgp.html
75 func getKeyByEmail(keyring openpgp.EntityList, email string) *openpgp.Entity {
76     for _, entity := range keyring {
77         for _, ident := range entity.Identities {
78             if ident.UserId.Email == email {
79                 return entity
80             }
81         }
82     }
83
84     return nil
85 }
86
87 func GetKeyByFingerprint(fingerprint string) *openpgp.Entity {
88     return getKeyByFingerprint(pubring, fingerprint)
89 }
90
91 func getKeyByFingerprint(keyring openpgp.EntityList, fingerprint string) *openpgp.Entity {
92     for _, entity := range keyring {
93         if FingerprintToString(entity.PrimaryKey.Fingerprint) == fingerprint {
94             return entity
95         }
96     }
97
98     return nil
99 }
100
101 func Contacts() openpgp.EntityList {
102     return pubring
103 }
104
105 func MyFingerprint() string {
106     return FingerprintToString(alice.PrimaryKey.Fingerprint)
```

```

107 }
108
109 func FingerprintToString(fingerprint [20]byte) string {
110     return string(fingerprint[:])
111 }
112
113 func EncryptBencode(conn io.Writer, bob *openpgp.Entity, msg map[string]interface{}) (err error) {
114     plain, err := openpgp.Encrypt(conn, openpgp.EntityList{bob}, alice, nil, nil)
115     if err != nil {
116         l4g.Warn(err)
117     }
118     err = bencode.Marshal(plain, msg)
119     if err != nil {
120         l4g.Warn(err)
121     }
122     plain.Close()
123
124     return
125 }
126
127 func DecryptBencode(conn io.Reader) (res map[string]interface{},
128     signer *openpgp.Entity, err error) {
129     msg, err := openpgp.ReadMessage(conn, pubring, nil, nil)
130     if err != nil {
131         l4g.Warn(err)
132         return
133     }
134
135     if !msg.IsEncrypted || !msg.IsSigned {
136         err = errors.New("Message was not encrypted or not signed")
137         return
138     }
139
140     signer = msg.SignedBy.Entity
141
142     err = bencode.Unmarshal(msg.UnverifiedBody, res)
143     if err != nil {
144         l4g.Warn(err)
145         return
146     }
147
148     if msg.SignatureError != nil {
149         err = errors.New("MESSAGE SIGNATURE WRONG! MESSAGE TAMPERING DETECTED!")
150     }
151     return
152 }
153
154 func RandomToken() string {
155     b := make([]byte, 32)

```

## Appendix B. Dictator Breaker Code

```
156     n, err := io.ReadFull(rand.Reader, b)
157     if n != len(b) || err != nil {
158         panic("Programming error, there is no reason for this to fail")
159     }
160     return string(b)
161 }
162
163 func WithinOneHour(epoch int) bool {
164     now := time.Now()
165     oneHourFromNow := now.Add(1 * time.Hour)
166     oneHourAgo := now.Add(-1 * time.Hour)
167     x := time.Unix(int64(epoch), 0)
168     return oneHourFromNow.After(x) && oneHourAgo.Before(x)
169 }
```

### B.6. go.nhelke.com/dictator-breaker/security/openpgp\_test.go

```
1 package security
2
3 import (
4     "bytes"
5     "fmt"
6     "io/ioutil"
7     "testing"
8
9     "code.google.com/p/go.crypto/openpgp"
10    _ "code.google.com/p/go.crypto/ripemd160"
11 )
12
13 func TestHandshake(t *testing.T) {
14     alice, err := openpgp.NewEntity("Alice", "", "alice@example.com", nil)
15     if err != nil {
16         t.Error(err)
17     }
18     bob, err := openpgp.NewEntity("Bob", "", "bob@example.com", nil)
19     if err != nil {
20         t.Error(err)
21     }
22
23     buf := new(bytes.Buffer)
24
25     plain, err := openpgp.Encrypt(buf, openpgp.EntityList{bob}, alice, nil, nil)
26     if err != nil {
27         t.Error(err)
28     }
29
30     handshake := "Whatever"
31     fmt.Fprintf(plain, handshake)
```



```

32
33     plain.Close()
34
35     msg, err := openpgp.ReadMessage(buf, openpgp.EntityList{alice, bob}, nil, nil)
36     if err != nil {
37         t.Error(err)
38     }
39
40     if !msg.IsEncrypted || !msg.IsSigned {
41         t.Fail()
42     }
43
44     res, err := ioutil.ReadAll(msg.UnverifiedBody)
45     if err != nil {
46         t.Error(err)
47     }
48
49     if msg.SignatureError != nil || string(res) != handshake {
50         t.Fail()
51     }
52
53 }

```

## B.7. go.nhelke.com/dictator-breaker/views/chat.html

```

1 <!DOCTYPE html>
2
3 <head>
4 <meta charset="utf-8" />
5
6 <title>WebSocket Test</title>
7
8 <script src="https://ajax.googleapis.com/ajax/libs/jquery/2.0.0/jquery.min.js"></script>
9 <script src="/jquery-2.0.0.js"></script>
10
11 <style>
12 @import url(http://fonts.googleapis.com/css?family=Inconsolata:400,700);
13 * {
14 font-family: 'Inconsolata', monospace;
15 }
16 #input {
17     border: thin solid black;
18 }
19 </style>
20
21 </head>
22 <body>
23 <h2>Dictator Breaker</h2>

```

## Appendix B. Dictator Breaker Code

```
24
25 <div id="output"></div>
26 <div id="input" contenteditable="true"></div>
27
28 <script>
29
30   var wsUri = "ws://localhost:8080/api";
31   var output;
32
33   function onOpen(evt)
34   {
35       var msg = $('<div>WebSocket connection opened.</div>');
36       msg.hide()
37       msg.appendTo($("#output"));
38       msg.slideDown();
39   }
40
41   function onClose(evt)
42   {
43       alert("DISCONNECTED");
44   }
45
46   function onMessage(evt)
47   {
48       var msg = $('<div class="incoming">' + evt.data + '</div>');
49       msg.hide()
50       msg.appendTo($("#output"));
51       msg.slideDown();
52   }
53
54   function onError(evt)
55   {
56       alert(evt.data);
57   }
58
59   $(function () {
60       websocket = new WebSocket(wsUri);
61       websocket.onopen = function(evt) { onOpen(evt) };
62       websocket.onclose = function(evt) { onClose(evt) };
63       websocket.onmessage = function(evt) { onMessage(evt) };
64       websocket.onerror = function(evt) { onError(evt) };
65
66       var input = jQuery("#input");
67       var output = $("#output");
68       input.focus();
69       input.keypress(function(e) {
70           if(e.which == 13) {
71               var msg = input.clone().removeAttr("id").removeAttr("contenteditable");
72               websocket.send(msg[0].innerHTML + "\n");
```

```
73         msg.hide()
74         msg.addClass("outgoing");
75         msg.appendTo(output);
76         msg.slideDown();
77         input.empty();
78     return false;
79     }
80     });
81     });
82 </script>
83
84 </body>
85
86 </html>
87
88
```

## B.8. Known issues

1. The resources for the user interface must be placed in a subdirectory of the current working directory, called "views".

*Appendix B. Dictator Breaker Code*

## Appendix C.

### GoUPnPc

GoUPnPc is a GNU Public License Version 2 (GPLv2) licensed<sup>1</sup> library for controlling port mappings of Universal Plug and Play (UPnP)-enabled Internet Gateway Devices (IGDs). Its canonical source is GitHub<sup>2</sup>. Its documentation is available from GoDoc<sup>3</sup>.

For archival and review purposes, the source of GoUPnPc is included below as it stood at submission time:

#### C.1. [github.com/nhelke/goupnpc/README.mdown](https://github.com/nhelke/goupnpc/README.mdown)

```
GoUPnPc
=====
```

```
GoUPnPc is a MiniUPnPc inspired library I am in the process of developing as
part of my Bachelor's degree thesis. For the time being only IGDv1 is supported.
It offers an API very close to the options provided by the UPnPc command line
client.
```

```
Documentation is available at <http://godoc.org/github.com/nhelke/goupnpc>
```

```
License
-----
```

```
Copyright (C) 2013 Nicholas Helke
```

```
This program is free software; you can redistribute it and/or
modify it under the terms of the GNU General Public License
as published by the Free Software Foundation; either version 2
of the License, or (at your option) any later version.
```

```
This program is distributed in the hope that it will be useful,
but WITHOUT ANY WARRANTY; without even the implied warranty of
MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the
GNU General Public License for more details.
```

---

<sup>1</sup>As required by the license, a copy is available as appendix E

<sup>2</sup><https://github.com/nhelke/goupnpc>

<sup>3</sup><http://godoc.org/github.com/nhelke/goupnpc>

You should have received a copy of the GNU General Public License along with this program; if not, write to the Free Software Foundation, Inc., 51 Franklin Street, Fifth Floor, Boston, MA 02110-1301, USA.

## C.2. [github.com/nhelke/goupnpc/cmd.go](https://github.com/nhelke/goupnpc/cmd.go)

```
1 // This command is useful both to test the associated goupnp library and
2 // its source serves as an example of how to use said library.
3 //
4 // Usage instructions can be obtained by running it without any arguments.
5 package main
6
7 import (
8     "fmt"
9     "os"
10    "strconv"
11    "time"
12
13    l4g "code.google.com/p/log4go"
14    "github.com/nhelke/goupnpc/goupnp"
15 )
16
17 func main() {
18     l4g.AddFilter("stdout", l4g.WARNING, l4g.NewConsoleLogWriter())
19
20     if len(os.Args) < 2 {
21         printUsage()
22     } else {
23         discover := goupnp.DiscoverIGD()
24         if os.Args[1] == "s" {
25             igd := <-discover
26             status := <-igd.GetConnectionStatus()
27             fmt.Printf("%+v\n", status)
28         } else if os.Args[1] == "l" {
29             igd := <-discover
30             for portMapping := range igd.ListRedirections() {
31                 fmt.Println(portMapping)
32             }
33         } else if os.Args[1] == "a" {
34             igd := <-discover
35             port, _ := strconv.Atoi(os.Args[2])
36             proto := goupnp.ParseProtocol(os.Args[3])
37             myMapping := <-igd.AddLocalPortRedirection(uint16(port), proto)
38             fmt.Printf("%+v\n", myMapping)
39         } else {
40             printUsage()
41         }
42     }
43 }
```

```

43
44     time.Sleep(1 * time.Second)
45 }
46
47 func printUsage() {
48     fmt.Println(
49         `Usage: goupnpc s
50             Print IGD Status
51             goupnpc a port protocol
52             Add local port mapping with internal and external ports equal to
53             port and protocol equal to, well I will let you guess
54             goupnpc l
55             Lists all port mappings on the IGD
56     NOTABENE No error checking is performed, if anything goes wrong, it will
57     probably panic on nil or something`)
58 }

```

### C.3. [github.com/nhelke/goupnpc/goupnp/goupnp.go](https://github.com/nhelke/goupnpc/goupnp/goupnp.go)

```

1 // A small library for using the port mapping controls of UPnP-enabled IGDs
2 package goupnp
3
4 import (
5     "errors"
6     "fmt"
7     "io/ioutil"
8     "net"
9     "net/http"
10    "net/url"
11    "strings"
12
13    l4g "code.google.com/p/log4go"
14 )
15
16 const (
17     TCP protocol = 1 << iota
18     UDP
19 )
20
21 // This type provides all the information about port mappings.
22 // It also serves as a handle returned by AddLocalPortRedirection() for use with
23 // DeletePortRedirection().
24 type PortMapping struct {
25     InternalPort uint16
26     ExternalPort uint16
27     Protocol     protocol
28     InternalHost net.IP
29     Description  string

```

## Appendix C. GoUPNPc

```
30     Enabled      bool
31     Lease        uint
32 }
33
34 func (self *PortMapping) String() string {
35     return fmt.Sprintf(self.InternalHost, ":", self.InternalPort, "<=",
36         self.ExternalPort, self.Protocol, ` "` , self.Description, ` "` (`,
37         self.Enabled, ", ", self.Lease, ")")
38 }
39
40 // This opaque type provides a handle to a discovered IGD
41 // Use DiscoverIGD() to obtain such a handle.
42 //
43 // NOT A BENE Using instances of this struct not returned by the appropriate
44 // function call has undefined behaviour
45 type IGD struct {
46     controlURL *url.URL
47     upnptype   string
48     iface      net.IP
49 }
50
51 func (self *IGD) String() string {
52     return self.controlURL.String()
53 }
54
55 // This function returns a channel which will be sent the first IGD it finds in
56 // traversing `net.InterfaceAddrs()` with IP addresses in the private network
57 // range.
58 //
59 // The channel this function returns should be listened on to avoid leaking
60 // goroutines. Additionally the listener must check whether the channel the
61 // value returned by the channel against nil, to ensure that an IGD was indeed
62 // found.
63 func DiscoverIGD() (ret chan *IGD) {
64     // Create the channel we will return
65     ret = make(chan *IGD)
66
67     // Do the work asynchronously
68     go func() {
69         // For each and every local address in the private network range
70         bindLocalAddrs := localPrivateAddrs()
71         l4g.Debug("Found %d private network interfaces", len(bindLocalAddrs))
72         for i := 0; i < len(bindLocalAddrs); i++ {
73             // Use SSDP to search for a UPnP-enabled IGD
74             descURL, ok := discoverIGDDescriptionURL(bindLocalAddrs[i])
75
76             if ok {
77                 // If we found one, we go fetch its description XML
78                 resp, err := http.Get(descURL.String())
```



```

79     if err == nil {
80         // We got something back, lets not leak it
81         defer resp.Body.Close()
82         // We read in the whole description into memory We might
83         // envisage at a later date putting an upperbound on the
84         // buffer, however there is no risk of buffer overflow, so
85         // it is a low priority
86         body, err := ioutil.ReadAll(resp.Body)
87         if err == nil {
88             l4g.Debug("Description XML:\n%s", string(body))
89             // Parse the XML and extract relevant information
90             upnptype, controlURL, err := getConnectionControlURL(body)
91             if err == nil {
92                 var igd IGD
93                 // It worked, lets now try and wrap it in an igd struct
94                 igd.controlURL, err = url.Parse(controlURL)
95                 if err != nil {
96                     l4g.Warn("Failed to parse URL %v", controlURL)
97                 } else {
98                     // The URL was good, lets track the type as
99                     // well, in order to make the correct calls down
100                    // the line
101                    igd.upnptype = upnptype
102                    // We now add the local binding address to
103                    // enable the simple AddLocalPortRedirection
104                    // method
105                    igd.iface = bindLocalAddrs[i].IP
106
107                    ret <- &igd
108                }
109            } else {
110                l4g.Warn("Bad XML: %v", err)
111            }
112        } else {
113            l4g.Warn("Error reading response")
114        }
115    }
116 }
117 }
118
119 // If we get here we did not find an IGD or have already passed the
120 // information to the channel and it has been read, so we close the
121 // channel This will have the effect of returning nil and will indicate
122 // the closure to listeners.
123 close(ret)
124 }()
125 return
126 }
127

```

## Appendix C. GoUPNPc

```
128 type ConnectionStatus struct {
129     Connected bool
130     IP         net.IP
131 }
132
133 // This method fetches the status of the IGD.
134 //
135 // Errors are indicated by the channel closing before a ConnectionStatus is
136 // returned. Listeners should therefore check at the very least for nil, better
137 // still for channel closure.
138 //
139 // NOTA BENE the channel closes after a successive ConnectionStatus has been
140 // send on it, in order to not leak resources.
141 func (self *IGD) GetConnectionStatus() (ret chan *ConnectionStatus) {
142     // We initialise the channel
143     ret = make(chan *ConnectionStatus)
144
145     // We go do the work in a separate goroutine, the closure has access to the
146     // channel we just instanciated so we will be able to manipulate it.
147     go func() {
148         x, ok := self.soapRequest("GetStatusInfo", statusRequestStringReader(self.upnptype))
149         if ok && strings.EqualFold(x.Body.Status.NewConnectionStatus, "Connected") {
150             y, ok := self.soapRequest("GetExternalIPAddress", externalIPRequestStringReader(s
151
152             if ok {
153                 ipString := y.Body.IP.NewExternalIPAddress
154                 ip := net.ParseIP(ipString)
155                 if ip != nil {
156                     ret <- &ConnectionStatus{true, ip}
157                     return
158                 } else {
159                     l4g.Warn("Failed to parse IP string %v", ipString)
160                 }
161             } else {
162                 l4g.Warn("Failed to get IP address after estabilishing the connection was ok"
163             }
164             } else if ok && strings.EqualFold(x.Body.Status.NewConnectionStatus, "Disconnected")
165                 ret <- &ConnectionStatus{false, nil}
166             }
167             close(ret)
168         }()
169
170     // We immediately return the channel to the caller
171     return
172 }
173
174 // This method creates a port mapping on the IGD with internal, external ports
175 // and protocol respectively equal to the passed port argument (bis) and
176 // protocol
```

```

177 //
178 // Errors are indicated by the channel closing before a PortMapping is returned.
179 // Listeners should therefore check at the very least for nil, better still
180 // for channel closure.
181 //
182 // NOTA BENE the channel closes after a successive PortMapping has been send on
183 // it, in order to not leak resources.
184 func (self *IGD) AddLocalPortRedirection(port uint16, proto protocol) (ret chan *PortMapping) {
185     ret = make(chan *PortMapping)
186
187     go func() {
188         description := fmt.Sprintf("goupnp %s %d %s", self.iface, port, proto)
189         _, ok := self.soapRequest("AddPortMapping",
190             createPortMappingStringReader(self.upnptype, port,
191                 proto, self.iface, description))
192         if ok {
193             portMapping := PortMapping{
194                 InternalPort: port,
195                 ExternalPort: port,
196                 Enabled:        true,
197                 Description:  description,
198                 InternalHost: self.iface,
199                 Protocol:    proto,
200             }
201
202             ret <- &portMapping
203         }
204         close(ret)
205     }()
206
207     return
208 }
209
210 // BUG(nhelke): NOT IMPLEMENTED ALWAYS RETURNS ERROR
211 //
212 // Please feel free to submit a pull request to
213 // https://github.com/nhelke/goupnpc and I will be sure to merge it.
214 func (self *IGD) DeletePortRedirection(portMappings ...*PortMapping) (ret chan error) {
215     ret = make(chan error)
216     go func() {
217         ret <- errors.New("Sorry, I haven't implemented this yet. " +
218             "Feel free to submit a pull request to github.com/nhelke/goupnpc " +
219             "and I will be sure to merge it.")
220         close(ret)
221     }()
222     return ret
223 }
224
225 // This method returns a buffered channel which should be iterated over. The

```

## Appendix C. GoUPNPc

```
226 // channel is closed on after the last port mapping, so iterating over the
227 // channel will not loop forever.
228 func (self *IGD) ListRedirections() (ret chan *PortMapping) {
229     ret = make(chan *PortMapping, 10)
230
231     go func() {
232         var (
233             ok bool = true
234             i uint = 0
235             x *soapEnvelope
236         )
237         for ; ; i++ {
238             x, ok = self.soapRequest("GetGenericPortMappingEntry",
239                 portMappingRequestStringReader(self.upnptype, i))
240             if ok {
241                 portMapping := PortMapping{
242                     InternalPort: x.Body.PortMapping.InternalPort,
243                     ExternalPort: x.Body.PortMapping.ExternalPort,
244                     Enabled: x.Body.PortMapping.Enabled != 0,
245                     Description: x.Body.PortMapping.Description,
246                     InternalHost: net.ParseIP(x.Body.PortMapping.InternalClient),
247                 }
248                 portMapping.Protocol = ParseProtocol(x.Body.PortMapping.Protocol)
249                 ret <- &portMapping
250             } else {
251                 close(ret)
252                 break
253             }
254         }
255     }()
256
257     return
258 }
259
260 // Function for parsing a protocol in string form to protocol type for use with
261 // this library's methods. Only TCP and UDP are supported.
262 func ParseProtocol(proto string) (ret protocol) {
263     switch {
264     case strings.EqualFold("tcp", proto):
265         ret = TCP
266     case strings.EqualFold("udp", proto):
267         ret = UDP
268     }
269     return
270 }
271
272 // This function returns true if and only if the passed IP address belongs to
273 // one of the ranges reserved in RFC 1918 for use in private networks
274 //
```

```

275 // This function is only part of this package as it is used internally and is
276 // public as it is deemed useful for developers to ascertain whether or not a
277 // given external IP address such as one returned by GetConnectionStatus is
278 // public or not and as creating a standalone package just for this one function
279 // seemed excessive.
280 func IsPrivateIPAddress(addr net.IP) bool {
281     ip4 := addr.To4()
282     if ip4 == nil || !ip4.IsGlobalUnicast() {
283         return false
284     }
285     var (
286         aAddr = net.IPv4(10, 0, 0, 0)
287         aMask = net.IPv4Mask(255, 0, 0, 0)
288
289         bAddr = net.IPv4(172, 16, 0, 0)
290         bMask = net.IPv4Mask(255, 240, 0, 0)
291
292         cAddr = net.IPv4(192, 168, 0, 0)
293         cMask = net.IPv4Mask(255, 255, 0, 0)
294     )
295
296     return ip4.Mask(aMask).Equal(aAddr) ||
297         ip4.Mask(bMask).Equal(bAddr) ||
298         ip4.Mask(cMask).Equal(cAddr)
299 }

```

#### C.4. [github.com/nhelke/goupnpc/goupnp/goupnp\\_test.go](https://github.com/nhelke/goupnpc/goupnp/goupnp_test.go)

```

1 package goupnp
2
3 import (
4     "net"
5     "testing"
6 )
7
8 func TestIsPrivateIPAddress(t *testing.T) {
9     privateAddrs := []net.IP{
10         net.IPv4(192, 168, 2, 32),
11         net.IPv4(10, 230, 46, 52),
12         net.IPv4(172, 22, 8, 61),
13     }
14
15     for i := 0; i < len(privateAddrs); i++ {
16         if !IsPrivateIPAddress(privateAddrs[i]) {
17             t.Errorf("Incorrectly did not identify %v as a private IPv4 Address", privateAddrs[i])
18         }
19     }
20 }

```

## Appendix C. GoUPNPc

```
21     publicAddr := []net.IP{
22         net.IPv4(184, 85, 61, 15),
23         net.IPv4(137, 164, 29, 67),
24         net.IPv4(8, 8, 8, 8),
25     }
26
27     for i := 0; i < len(publicAddr); i++ {
28         if IsPrivateIPAddress(publicAddr[i]) {
29             t.Errorf("Incorrectly identified %v as a private IPv4 Address", publicAddr[i])
30         }
31     }
32 }
```

### C.5. [github.com/nhelke/goupnpc/goupnp/backend.go](https://github.com/nhelke/goupnpc/goupnp/backend.go)

```
1 package goupnp
2
3 import (
4     "bytes"
5     "encoding/xml"
6     "fmt"
7     "io"
8     "io/ioutil"
9     "net"
10    "net/http"
11
12    l4g "code.google.com/p/log4go"
13 )
14
15 type protocol int
16
17 func (self protocol) String() string {
18     switch self {
19     case TCP:
20         return "TCP"
21     case UDP:
22         return "UDP"
23     default:
24         return "#(Bad Protocol Value)"
25     }
26 }
27
28 type deviceElement struct {
29     FriendlyName string `xml:"friendlyName"`
30     Manufacturer string `xml:"manufacturer"`
31
32     Services []struct {
33         ServiceType string `xml:"serviceType"`
```

```

34     ControlURL string `xml:"controlURL"`
35 } `xml:"serviceList>service"`
36
37 Devices []deviceElement `xml:"deviceList>device",omitempty`
38 }
39
40 type deviceDescription struct {
41     XMLName xml.Name `xml:"urn:schemas-upnp-org:device-1-0 root"`
42
43     SpecVersion struct {
44         Major int `xml:"major"`
45         Minor int `xml:"minor"`
46     } `xml:"specVersion"`
47
48     URLBase string
49
50     Device deviceElement `xml:"device"`
51 }
52
53 const (
54     connectionTypeStringWANIP = "urn:schemas-upnp-org:service:WANIPConnection:1"
55     connectionTypeStringWANPPP = "urn:schemas-upnp-org:service:WANPPPConnection:1"
56 )
57
58 func statusRequestStringReader(upnptype string) io.Reader {
59     return bytes.NewReader([]byte(fmt.Sprintf(statusRequestString, upnptype)))
60 }
61
62 const statusRequestString = `<?xml version="1.0"?>
63 <s:Envelope xmlns:s="http://schemas.xmlsoap.org/soap/envelope/"
64 s:encodingStyle="http://schemas.xmlsoap.org/soap/encoding/">
65 <s:Body>
66 <u:GetStatusInfo xmlns:u="%s">
67 </u:GetStatusInfo>
68 </s:Body>
69 </s:Envelope>
70 `
71
72 const externalIPRequestString = `<?xml version="1.0"?>
73 <s:Envelope xmlns:s="http://schemas.xmlsoap.org/soap/envelope/"
74 s:encodingStyle="http://schemas.xmlsoap.org/soap/encoding/">
75 <s:Body>
76 <u:GetExternalIPAddress xmlns:u="%s"></u:GetExternalIPAddress>
77 </s:Body>
78 </s:Envelope>
79 `
80
81 func externalIPRequestStringReader(upnptype string) io.Reader {
82     return bytes.NewReader([]byte(fmt.Sprintf(externalIPRequestString, upnptype)))

```

## Appendix C. GoUPNPc

```
83 }
84
85 const portMappingRequestString = `<?xml version="1.0"?>
86 <s:Envelope xmlns:s="http://schemas.xmlsoap.org/soap/envelope/" ` +
87     `s:encodingStyle="http://schemas.xmlsoap.org/soap/encoding/"><s:Body>` +
88     `<u:GetGenericPortMappingEntry xmlns:u="%s"><NewPortMappingIndex>%d` +
89     `</NewPortMappingIndex></u:GetGenericPortMappingEntry></s:Body></s:Envelope>
90 `
91
92 func portMappingRequestStringReader(upnptype string, index uint) io.Reader {
93     str := fmt.Sprintf(portMappingRequestString, upnptype, index)
94     return bytes.NewReader([]byte(str))
95 }
96
97 const createPortMappingString = `<?xml version="1.0"?>
98 <s:Envelope xmlns:s="http://schemas.xmlsoap.org/soap/envelope/" ` +
99     `s:encodingStyle="http://schemas.xmlsoap.org/soap/encoding/"><s:Body>` +
100    `<u:AddPortMapping xmlns:u="%s"><NewRemoteHost></NewRemoteHost>` +
101    `<NewExternalPort>%d</NewExternalPort><NewProtocol>%s</NewProtocol>` +
102    `<NewInternalPort>%d</NewInternalPort>` +
103    `<NewInternalClient>%s</NewInternalClient><NewEnabled>1</NewEnabled>` +
104    `<NewPortMappingDescription>%s</NewPortMappingDescription>` +
105    `<NewLeaseDuration>0</NewLeaseDuration>` +
106    `</u:AddPortMapping></s:Body></s:Envelope>
107 `
108
109 func createPortMappingStringReader(upnptype string, port uint16,
110     proto protocol, localAddr net.IP, description string) io.Reader {
111     str := fmt.Sprintf(createPortMappingString, upnptype, port, proto, port,
112         localAddr, description)
113     return bytes.NewReader([]byte(str))
114 }
115
116 type soapEnvelope struct {
117     XMLName xml.Name `xml:"http://schemas.xmlsoap.org/soap/envelope/ Envelope"`
118
119     Body struct {
120         IP struct {
121             XMLName xml.Name `xml:"GetExternalIPAddressResponse"`
122
123             NewExternalIPAddress string
124         }
125         Status struct {
126             XMLName xml.Name `xml:"GetStatusInfoResponse"`
127
128             NewConnectionStatus string
129         }
130         PortMapping soapPortMapping `xml:"GetGenericPortMappingEntryResponse"`
131     }
132 }
```



```

132 }
133
134 type soapPortMapping struct {
135     Protocol      string `xml:"NewProtocol"`
136     ExternalPort  uint16 `xml:"NewExternalPort"`
137     InternalPort  uint16 `xml:"NewInternalPort"`
138     InternalClient string `xml:"NewInternalClient"`
139     Enabled       int    `xml:"NewEnabled"`
140     Description   string `xml:"NewPortMappingDescription"`
141     Lease         uint   `xml:"NewLeaseDuration"`
142 }
143
144 func (self *IGD) soapRequest(requestType string,
145     requestXML io.Reader) (x *soapEnvelope, ok bool) {
146     req, err := http.NewRequest("POST", self.controlURL.String(), requestXML)
147     if err != nil {
148         panic("Programming Error: This hand crafted http.Request should not be bad")
149     }
150     req.Header.Add("Content-Type", "text/xml")
151     req.Header.Add("SOAPAction",
152         `"+self.upnptype+"#"+requestType+"`)
153     req.Header.Add("Connection", "Close")
154     req.Header.Add("Cache-Control", "no-cache")
155     req.Header.Add("Pragma", "no-cache")
156
157     resp, err := http.DefaultClient.Do(req)
158     if err == nil {
159         // We got something back, lets not leak it
160         defer resp.Body.Close()
161
162         if resp.StatusCode != http.StatusOK {
163             return
164         }
165
166         body, err := ioutil.ReadAll(resp.Body)
167         if err == nil {
168             l4g.Debug("SOAP Response:\n%s", string(body))
169             err := xml.Unmarshal(body, &x)
170             if err == nil {
171                 ok = true
172             } else {
173                 l4g.Warn(err)
174             }
175         } else {
176             l4g.Warn(err)
177         }
178     } else {
179         l4g.Warn(err)
180     }

```

```
181
182     return
183 }
```

## C.6. [github.com/nhelke/goupnpc/goupnp/backend\\_test.go](https://github.com/nhelke/goupnpc/goupnp/backend_test.go)

```
1 package goupnp
2
3 import (
4     "encoding/xml"
5     "testing"
6 )
7
8 const exampleBelkinSOAP = `<?xml version="1.0"?>
9 <SOAP-ENV:Envelope xmlns:SOAP-ENV="http://schemas.xmlsoap.org/soap/envelope/"
10 SOAP-ENV:encodingStyle="http://schemas.xmlsoap.org/soap/encoding/"
11 <SOAP-ENV:Body>
12 <m:GetStatusInfoResponse
13 xmlns:m="urn:schemas-upnp-org:service:WANIPConnection:1">
14 <NewConnectionStatus>Connected</NewConnectionStatus>
15 <NewLastConnectionError>ERROR_NONE</NewLastConnectionError>
16 <NewUptime>194979</NewUptime></m:GetStatusInfoResponse></SOAP-ENV:Body>
17 </SOAP-ENV:Envelope>
18 `
19
20 func TestSOAPParsing(t *testing.T) {
21     var x soapEnvelope
22     err := xml.Unmarshal([]byte(exampleBelkinSOAP), &x)
23     if err != nil {
24         t.Errorf("%v", err)
25     } else if status := x.Body.Status.NewConnectionStatus; status != "Connected" {
26         t.Errorf("Status incorrectly parsed as %v", status)
27     }
28 }
29
30 const exampleBelkinPortMappingResponse = `<?xml version="1.0"?>
31 <SOAP-ENV:Envelope xmlns:SOAP-ENV="http://schemas.xmlsoap.org/soap/envelope/"
32 SOAP-ENV:encodingStyle="http://schemas.xmlsoap.org/soap/encoding/"
33 <SOAP-ENV:Body><m:GetGenericPortMappingEntryResponse
34 xmlns:m="urn:schemas-upnp-org:service:WANIPConnection:1">
35 <NewRemoteHost></NewRemoteHost>
36 <NewExternalPort>5900</NewExternalPort>
37 <NewProtocol>TCP</NewProtocol>
38 <NewInternalPort>5901</NewInternalPort>
39 <NewInternalClient>192.168.2.5</NewInternalClient>
40 <NewEnabled>1</NewEnabled>
41 <NewPortMappingDescription>cPM.Port.Map.ee97f96de8c1647a</NewPortMappingDescription>
42 <NewLeaseDuration>0</NewLeaseDuration>
```

```

43 </m:GetGenericPortMappingEntryResponse>
44 </SOAP-ENV:Body>
45 </SOAP-ENV:Envelope>
46 `
47
48 func TestPortMappingResponseParsing(t *testing.T) {
49     var x soapEnvelope
50     err := xml.Unmarshal([]byte(exampleBelkinPortMappingResponse), &x)
51     if err != nil {
52         t.Errorf("%v", err)
53     } else {
54         referenceMapping := soapPortMapping{Protocol: "TCP", ExternalPort: 5900,
55             InternalPort: 5901, InternalClient: "192.168.2.5", Enabled: 1,
56             Description: "cPM.Port.Map.ee97f96de8c1647a"}
57         if portMapping := x.Body.PortMapping; portMapping != referenceMapping {
58             t.Errorf("Port mapping incorrectly parsed as %#v", portMapping)
59         }
60     }
61 }

```

## C.7. [github.com/nhelke/goupnpc/goupnp/ssdp.go](https://github.com/nhelke/goupnpc/goupnp/ssdp.go)

```

1 package goupnp
2
3 import (
4     "bufio"
5     "bytes"
6     "encoding/xml"
7     "errors"
8     "fmt"
9     "net"
10    "net/http"
11    "net/url"
12    "time"
13
14    l4g "code.google.com/p/log4go"
15 )
16
17 // Returns all local interface IP addresses in the private network range
18 // They are traversed in the order returned by `net.InterfaceAddrs()`
19 func localPrivateAddrs() (ret []*net.UDPAddr) {
20     addrs, err := net.InterfaceAddrs()
21     if err == nil {
22         for i := 0; i < len(addrs); i++ {
23             if ip, ok := addrs[i].(*net.IPNet); ok {
24                 if IsPrivateIPAddress(ip.IP) {
25                     l4g.Debug("Found private addr %v", ip.IP)
26                     ret = append(ret, &net.UDPAddr{ip.IP, 0, ""})

```

## Appendix C. GoUPNPc

```
27         }
28     }
29 }
30 } else {
31     l4g.Warn(err)
32 }
33 return
34 }
35
36 // This function implements the strict minimum of SSDP in order to discover the
37 // an IGD on the passed localBindAddr. The function blocks until a UPnP enabled
38 // IGD is found or timeout of four seconds expires. Timeouts smaller than 3
39 // seconds are unreasonable This function's behavior is not defined if the
40 // passed localBindAddr is not an IP address in the private network range. You
41 // may wish to use goupnp.localPrivateAddrs() to obtain a list of valid such
42 // addresses for the localhost.
43 func discoverIGDDescriptionURL(localBindAddr *net.UDPAddr) (u *url.URL, ok bool) {
44     const (
45         ssdpIPv4Addr = "239.255.255.250"
46         ssdpPort      = 1900
47         format        = "M-SEARCH * HTTP/1.1\r\n" +
48             "HOST: %s:%d\r\n" +
49             "ST: %s\r\n" +
50             "MAN: \"ssdp:discover\"\r\n" +
51             "MX: %d\r\n" +
52             "\r\n"
53     )
54
55     // These are the various device types we need to M-SEARCH the local subnet
56     // for. The last one is a fallback copied from MiniUPnPc's behavior and is
57     // unlikely to yield usable results
58     //
59     // This slice is sorted from most specific device type to the most general.
60     // Be advised that the below loop relies on this ordering.
61     var deviceTypes = []string{
62         "urn:schemas-upnp-org:device:InternetGatewayDevice:1",
63         "urn:schemas-upnp-org:service:WANIPConnection:1",
64         "urn:schemas-upnp-org:service:WANPPPPConnection:1",
65         "upnp:rootdevice",
66     }
67
68     multicastAddr, err := net.ResolveUDPAddr("udp4", fmt.Sprintf("%s:%d",
69         ssdpIPv4Addr, ssdpPort))
70     if err != nil {
71         panic("Programming error: Our UDPAddr is incorrect")
72     }
73
74     conn, err := net.ListenUDP("udp4", localBindAddr)
75     var timeout time.Duration = 4 * time.Second
```

```

76 if err == nil {
77     // For each device type, M-SEARCH for it, return the first one found
78     // As deviceTypes is sorted from most specific to least specific type
79     // returning the first should work fine.
80     for i := 0; i < len(deviceTypes); i++ {
81         // We write our own request *à la main* as trying to use Go's
82         // standard library's HTTP package turns out to be require more
83         // code than writing the request by hand, because of the non-
84         // standard URL
85         requestString := []byte(fmt.Sprintf(format, ssdpIPv4Addr, ssdpPort,
86             deviceTypes[i], timeout/time.Second))
87         // Allocate a buffer for the response
88         buf := make([]byte, 1500)
89         // We want to timeout and move on to the next type after a couple of
90         // seconds
91         conn.SetDeadline(time.Now().Add(timeout))
92         // Send multicast request
93         conn.WriteToUDP(requestString, multicastAddr)
94         // Get a response; the above timeout is still in effect as it
95         // should be
96         n, addr, err := conn.ReadFromUDP(buf)
97         if err == nil {
98             // Parse and interpret the response and break if successful
99             l4g.Debug("Received %d bytes from %v", n, addr)
100            req, err := http.ReadRequest(bufio.NewReader(bytes.NewReader(
101                requestString)))
102            if err != nil {
103                // Failure to parse the request represents an assertion
104                // failure as we crafted the request ourselves and have
105                // ensured its validity
106                panic(err)
107            }
108            resp, err := http.ReadResponse(bufio.NewReader(bytes.NewReader(
109                buf[:n])), req)
110            if err == nil {
111                // We got something back, lets not leak it
112                defer resp.Body.Close()
113                l4g.Debug("Discovered device returned:\n%v", resp.Header)
114                // We extract the description URL returned in the Location
115                // header. The UPnP standard ensure
116                urls := resp.Header["Location"]
117                // We must check that the Location header exists as required
118                // by the standard to avoid panicking if we get a bad
119                // response missing a Location header.
120                if len(urls) > 0 {
121                    // We have the location, bundle it up into a url.URL
122                    // object and return it
123                    u, err = url.Parse(urls[0])
124                    ok = err == nil

```

## Appendix C. GoUPNPc

```
125         return
126     } else {
127         l4g.Warn("Response did not contain Location header:\n%v",
128             resp.Header)
129     }
130     } else {
131         l4g.Warn(err)
132     }
133     } else {
134         l4g.Warn(err)
135     }
136 }
137 } else {
138     l4g.Warn(err)
139 }
140 // If we get here we could not find any UPnP devices
141 return // ok is false by default, signaling this failure
142 }
143
144 func extractConnectionControlURL(d deviceElement) (upnptype, url string, ok bool) {
145     for i := 0; i < len(d.Services); i++ {
146         if serviceType := d.Services[i].ServiceType; serviceType == connectionTypeStringWANIP
147             serviceType == connectionTypeStringWANPPP {
148             return serviceType, d.Services[i].ControlURL, true
149         }
150     }
151     for i := 0; i < len(d.Devices); i++ {
152         if upnptype, url, ok = extractConnectionControlURL(d.Devices[i]); ok {
153             return
154         }
155     }
156     return
157 }
158
159 func getConnectionControlURL(body []byte) (upnptype, url string, err error) {
160     var x deviceDescription
161     err = xml.Unmarshal(body, &x)
162     if err == nil {
163         var ok bool
164         upnptype, url, ok = extractConnectionControlURL(x.Device)
165         if !ok {
166             err = errors.New("Control URL not found")
167         } else {
168             // The URLs in the DeviceDescription elements are relative
169             url = x.URLBase + url
170         }
171     }
172     return
173 }
```

## C.8. github.com/nhelke/goupnpc/goupnp/ssdp\_test.go

```

1 package goupnp
2
3 import (
4     "testing"
5 )
6
7 func TestDescriptionParsing(t *testing.T) {
8     const belkinDescription string = `<?xml version="1.0"?>
9 <root xmlns="urn:schemas-upnp-org:device-1-0">
10     <specVersion>
11         <major>1</major>
12         <minor>0</minor>
13     </specVersion>
14     <URLBase>http://192.168.2.1:80</URLBase>
15     <device>
16         <deviceType>urn:schemas-upnp-org:device:InternetGatewayDevice:1</deviceType>
17         <friendlyName>Belkin N150 Wireless Router</friendlyName>
18         <manufacturer>Belkin International</manufacturer>
19         <manufacturerURL>http://www.Belkin.com</manufacturerURL>
20         <modelDescription>Wireless Router with Ethernet Switch</modelDescription>
21         <modelName>N150 Wireless Router</modelName>
22         <modelNumber>F9K1001</modelNumber>
23         <modelURL>http://www.Belkin.com</modelURL>
24         <serialNumber>201223GB303099</serialNumber>
25         <UDN>uuid:upnp-InternetGatewayDevice-1_0-08863bf24378</UDN>
26         <UPC>00000-00001</UPC>
27         <serviceList>
28             <service>
29                 <serviceType>urn:schemas-upnp-org:service:Layer3Forwarding:1</serviceType>
30                 <serviceId>urn:upnp-org:serviceId:L3Forwarding1</serviceId>
31                 <controlURL>/upnp/service/Layer3Forwarding</controlURL>
32                 <eventSubURL>/upnp/service/Layer3Forwarding</eventSubURL>
33                 <SCPDURL>/upnp/service/L3Frwd.xml</SCPDURL>
34             </service>
35         </serviceList>
36     </deviceList>
37     <device>
38         <deviceType>urn:schemas-upnp-org:device:WANDevice:1</deviceType>
39         <friendlyName>Belkin N150 Wireless Router</friendlyName>
40         <manufacturer>Belkin International</manufacturer>
41         <manufacturerURL>http://www.Belkin.com</manufacturerURL>
42         <modelDescription>Wireless Router with Ethernet Switch</modelDescription>
43         <modelName>N150 Wireless Router</modelName>
44         <modelNumber>F9K1001</modelNumber>
45         <modelURL>http://www.Belkin.com</modelURL>
46         <serialNumber>201223GB303099</serialNumber>
47         <UDN>uuid:upnp-WANDevice-1_0-08863bf24378</UDN>

```

Appendix C. GoUPNPc

```
48         <UPC>00000-00001</UPC>
49         <serviceList>
50             <service>
51                 <serviceType>urn:schemas-upnp-org:service:WANCommonInterfaceConfig:1</
52                 <serviceId>urn:upnp-org:serviceId:WANCommonInterfaceConfig</serviceId>
53                 <controlURL>/upnp/service/WANCommonInterfaceConfig</controlURL>
54                 <eventSubURL>/upnp/service/WANCommonInterfaceConfig</eventSubURL>
55                 <SCPDURL>/upnp/service/WANICf.g.xml</SCPDURL>
56             </service>
57         </serviceList>
58         <deviceList>
59             <device>
60                 <deviceType>urn:schemas-upnp-org:device:WANConnectionDevice:1</device
61                 <friendlyName>Belkin N150 Wireless Router</friendlyName>
62                 <manufacturer>Belkin International</manufacturer>
63                 <manufacturerURL>http://www.Belkin.com</manufacturerURL>
64                 <modelDescription>Wireless Router with Ethernet Switch</modelDescript
65                 <modelName>N150 Wireless Router</modelName>
66                 <modelName>N150 Wireless Router</modelName>
67                 <modelNumber>F9K1001</modelNumber>
68                 <modelURL>http://www.Belkin.com</modelURL>
69                 <serialNumber>201223GB303099</serialNumber>
70                 <UDN>uuid:upnp-WANConnectionDevice-1_0-08863bf24378</UDN>
71                 <UPC>00000-00001</UPC>
72                 <serviceList>
73                     <service>
74                         <serviceType>urn:schemas-upnp-org:service:WANIPConnection:1</
75                         <serviceId>urn:upnp-org:serviceId:WANIPConnection</serviceId>
76                         <controlURL>/upnp/service/WANIPConnection</controlURL>
77                         <eventSubURL>/upnp/service/WANIPConnection</eventSubURL>
78                         <SCPDURL>/upnp/service/WANIPcn.xml</SCPDURL>
79                     </service>
80                 </serviceList>
81             </device>
82         </deviceList>
83     </device>
84     <presentationURL>/index.html</presentationURL>
85 </device>
86 </root>
87 `
88
89     upnptype, url, err := getConnectionControlURL([]byte(belkinDescription))
90     if upnptype != connectionTypeStringWANIP || url != "http://192.168.2.1:80/upnp/service/WA
91         t.Errorf("Type: %v, URL: %v, Error: %v", upnptype, url, err)
92     }
93 }
```



## **C.9. Known issues**

1. The GoUPnP package must be compiled with Go version 1.1 or later, because of a bug in the standard "http" package in prior versions of Go.



## Appendix D.

### Patches submitted to open source projects

During this project, a couple of issues with some of the libraries used came to my attention and I submitted these two patches for them:

#### D.1. Patch submitted to [github.com/nicktuku/dht](https://github.com/nicktuku/dht)

This patch is discussed in section 4.3

```

1 From 5eadf6d5a3c57e83e17ce9ceba87b658c46dc553 Mon Sep 17 00:00:00 2001
2 From: Nicholas Helke <nhelke@gmail.com>
3 Date: Thu, 30 May 2013 17:56:40 +0200
4 Subject: [PATCH] Added flag to enable CPU profiling and adapted API to support
5 port number 0 (auto-assign)
6
7 Profiling
8 -----
9 In order to support profiling it was necessary to modify the main loop and
10 draining function to ensure that the main function exits cleanly.
11 It turns out the profile is not written to disk if the program does not exit
12 cleanly.
13
14 Auto-assigned port number
15 -----
16 I was a little surprised that this was not supported, I guess most torrent
17 client determine their port number and only then set up the DHT.
18 Nonetheless less I think this is an interesting feature to have.
19 ---
20 dht.go | 16 ++++++-
21 examples/find_infohash_and_wait/main.go | 66 ++++++-----
22 krpc.go | 2 +-
23 3 files changed, 67 insertions(+), 17 deletions(-)
24
25 diff --git a/dht.go b/dht.go
26 index d8c2e8f..6eeddca 100644
27 --- a/dht.go
28 +++ b/dht.go
29 @@ -164,6 +164,13 @@ func (d *DHT) PeersRequest(ih string, announce bool) {
30     l4g.Info("DHT: torrent client asking more peers for %x.", ih)
31 }
```

## Appendix D. Patches submitted to open source projects

```
32
33 +// Port returns the port number assigned to the DHT. This is useful when
34 +// when initialising the DHT with port 0, i.e. automatic port assignment,
35 +// in order to retrieve the actual port number used.
36 +func (d *DHT) Port() int {
37 +   return d.port;
38 +}
39 +
40 // AddNode informs the DHT of a new node it should add to its routing table.
41 // addr is a string containing the target node's "host:port" UDP address.
42 func (d *DHT) AddNode(addr string) {
43 @@ -200,6 +207,11 @@ func (d *DHT) DoDHT() {
44     return
45     }
46     d.conn = socket
47 +
48 + // Update the stored port number in case it was set 0, meaning it was
49 + // set automatically by the system
50 + d.port = socket.LocalAddr().(*net.UDPAddr).Port
51 +
52     bytesArena := newArena(maxUDPPacketSize, 500)
53     go readFromSocket(socket, socketChan, bytesArena)
54
55 @@ -228,7 +240,7 @@ func (d *DHT) DoDHT() {
56     rateLimit = 10
57     }
58     }
59 -   14g.Info("DHT: Starting DHT node %x.", d.nodeId)
60 +   14g.Info("DHT: Starting DHT node %x on port %d.", d.nodeId, d.port)
61
62     for {
63         select {
64 @@ -285,7 +297,7 @@ func (d *DHT) DoDHT() {
65             d.processPacket(p)
66             tokenBucket -= 1
67         } else {
68 -             // In the future it might be better to avoid dropping things like ping r
69 +             // TODO In the future it might be better to avoid dropping things like p
70             totalDroppedPackets.Add(1)
71         }
72     } else {
73 diff --git a/examples/find_infohash_and_wait/main.go b/examples/find_infohash_and_wait/main.g
74 index 3582216..16afece 100644
75 --- a/examples/find_infohash_and_wait/main.go
76 +++ b/examples/find_infohash_and_wait/main.go
77 @@ -1,5 +1,5 @@
78 -// Runs a node on UDP port 11221 that attempts to collect 100 peers for an
79 -// infohash, then keeps running as a passive DHT node.
80 +// Runs a node on system selected UDP port that attempts to collect 100 peers for
```

```

81  +// an infohash, then keeps running as a passive DHT node.
82  //
83  // IMPORTANT: if the UDP port is not reachable from the public internet, you
84  // may see very few results.
85  @@ -17,22 +17,41 @@ import (
86      "flag"
87      "fmt"
88      "os"
89  +   "os/signal"
90  +   "runtime/pprof"
91      "time"
92
93      l4g "code.google.com/p/log4go"
94  -   "github.com/nictuku/dht"
95  +   "github.com/nhelke/dht"
96      "net/http"
97  )
98
99  const (
100     httpPortTCP = 8711
101  -   dhtPortUDP  = 11221
102  +   dhtPortUDP  = 0 // 0 to let operating system automatically assign a free port
103  +)
104  +
105  +var (
106  +   quit          = make(chan bool)
107  +   interrupt     = make(chan os.Signal)
108  +   cpuprofile   = flag.String("cpuprofile", "", "write CPU profile to file")
109  +)
110
111  func main() {
112     flag.Parse()
113  +   if *cpuprofile != "" {
114  +       f, err := os.Create(*cpuprofile)
115  +       if err != nil {
116  +           l4g.Critical("Unable to create CPU profile file: %v", err)
117  +       } else {
118  +           pprof.StartCPUProfile(f)
119  +           defer pprof.StopCPUProfile()
120  +       }
121  +   }
122  +
123  +   flag.Parse()
124     // Change to l4g.DEBUG to see *lots* of debugging information.
125  -   l4g.AddFilter("stdout", l4g.WARNING, l4g.NewConsoleLogWriter())
126  +   l4g.AddFilter("stdout", l4g.INFO, l4g.NewConsoleLogWriter())
127     if len(flag.Args()) != 1 {
128         fmt.Fprintf(os.Stderr, "Usage: %v <infohash>\n\n", os.Args[0])
129         fmt.Fprintf(os.Stderr, "Example infohash: d1c5676ae7ac98e8b19f63565905105e3c4c37a2\n")

```

*Appendix D. Patches submitted to open source projects*

```
130 @@ -62,24 +81,43 @@ func main() {
131     go d.DoDHT()
132     go drainresults(d)
133
134 -   for {
135 -       // Give the DHT some time to "warm-up" its routing table.
136 -       time.Sleep(5 * time.Second)
137 +   // Signal handling is necessary so that we exit in a clean state capable
138 +   // of producing an optional CPU profile
139 +   signal.Notify(interrupt, os.Interrupt)
140
141 -       d.PeersRequest(string(ih), false)
142 +F:
143 +   for {
144 +       select {
145 +       case <-time.After(5 * time.Second):
146 +           // Give the DHT some time to "warm-up" its routing table.
147 +           // TODO Possibly create a channel to let the DHT advise us when it is
148 +           // nice and hot
149 +           d.PeersRequest(string(ih), false)
150 +       case <-interrupt:
151 +           break F
152 +       }
153     }
154 +
155 +   quit <- true
156 +   <-quit
157 }
158
159 // drainresults loops, printing the address of nodes it has found.
160 func drainresults(n *dht.DHT) {
161 -   fmt.Println("===== DHT")
162     l4g.Warn("Note that there are many bad nodes that reply to anything you ask.")
163     l4g.Warn("Peers found:")
164 -   for r := range n.PeersRequestResults {
165 -       for _, peers := range r {
166 -           for _, x := range peers {
167 -               l4g.Warn("%v", dht.DecodePeerAddress(x))
168 +F:
169 +   for {
170 +       select {
171 +       case r := <-n.PeersRequestResults:
172 +           for _, peers := range r {
173 +               for _, x := range peers {
174 +                   l4g.Warn("%v", dht.DecodePeerAddress(x))
175 +               }
176 +           }
177 +       case <-quit:
178 +           break F
```

```

179     }
180   }
181 +   quit <- true
182 }
183 diff --git a/krpc.go b/krpc.go
184 index e3ba095..6aa1ac6 100644
185 --- a/krpc.go
186 +++ b/krpc.go
187 @@ -196,7 +196,7 @@ func listen(listenPort int) (socket *net.UDPConn, err error) {
188     // debug.Printf("DHT: Listening for peers on port: %d\n", listenPort)
189     listener, err := net.ListenPacket("udp4", ":"+strconv.Itoa(listenPort))
190     if err != nil {
191 -         // debug.Println("DHT: Listen failed:", err)
192 +         l4g.Critical("DHT: Listen failed:", err)
193     }
194     if listener != nil {
195         socket = listener.(*net.UDPConn)
196 --
197 1.8.3.1
198

```

## D.2. Patch submitted to [github.com/monnand/dhcx](https://github.com/monnand/dhcx)

This patch is discussed in section 3.3

```

1 From 0c763f0ff37752f3a48aca7abb363fb9e255388b Mon Sep 17 00:00:00 2001
2 From: Nicholas Helke <nhelke@gmail.com>
3 Date: Mon, 1 Jul 2013 14:46:04 -0700
4 Subject: [PATCH] Added a CreateGroup method to dhgroup.go
5
6 Also added an associated test to dhcx_test.go and fleshed out the documentation
7 in dhgroup.go
8
9 Fixes monnand#1
10 ---
11 dhgroup.go | 19 ++++++
12 dhcx_test.go | 17 ++++++
13 2 files changed, 34 insertions(+), 2 deletions(-)
14
15 diff --git a/dhgroup.go b/dhgroup.go
16 index ea3bec4..e4cfd93 100644
17 --- a/dhgroup.go
18 +++ b/dhgroup.go
19 @@ -47,6 +47,10 @@ func (self *DHGroup) GeneratePrivateKey(randReader io.Reader) (key *DHKey, err e
20     return
21 }
22
23 +// This function fetches a DHGroup by its ID as defined in either RFC 2409 or
24 +// RFC 3526.

```

*Appendix D. Patches submitted to open source projects*

```
25  +//
26  +// If you are unsure what to use use group ID 0 for a sensible default value
27  func GetGroup(groupID int) (group *DHGroup, err error) {
28      if groupID <= 0 {
29          groupID = 14
30 @@ -77,6 +81,21 @@ func GetGroup(groupID int) (group *DHGroup, err error) {
31      return
32  }
33
34  +// This function enables users to create their own custom DHGroup.
35  +// Most users will not however want to use this function, and should prefer
36  +// the use of GetGroup which supplies DHGroups defined in RFCs 2409 and 3526
37  +//
38  +// WARNING! You should only use this if you know what you are doing. The
39  +// behavior of the group returned by this function is not defined if prime is
40  +// not in fact prime.
41  +func CreateGroup(prime, generator *big.Int) (group *DHGroup) {
42  +    group = &DHGroup{
43  +        g: generator,
44  +        p: prime,
45  +    }
46  +    return
47  +}
48  +
49  func (self *DHGroup) ComputeKey(pubkey *DHKey, privkey *DHKey) (key *DHKey, err error) {
50      if self.p == nil {
51          err = errors.New("DH: invalid group")
52 diff --git a/dhcx_test.go b/dhcx_test.go
53 index c2a7597..aad31da 100644
54 --- a/dhcx_test.go
55 +++ b/dhcx_test.go
56 @@ -19,13 +19,14 @@ package dhcx
57
58  import (
59      "fmt"
60  +    "math/big"
61      "testing"
62  )
63
64  type peer struct {
65  -    priv *DHKey
66  +    priv *DHKey
67      group *DHGroup
68  -    pub *DHKey
69  +    pub *DHKey
70  }
71
72  func newPeer(g *DHGroup) *peer {
73 @@ -88,3 +89,15 @@ func TestKeyExchange(t *testing.T) {
```



```
74     }
75 }
76
77 +func TestCustomGroupKeyExchange(t *testing.T) {
78 +   p, _ := new(big.Int).SetString("FFFFFFFFFFFFFFFFFC90FDAA22168C234C4C6628B80DC1CD129024E088A67CC74020B
79 +   g := new(big.Int).SetInt64(2)
80 +   group := CreateGroup(p, g)
81 +   p1 := newPeer(group)
82 +   p2 := newPeer(group)
83 +
84 +   err := exchangeKey(p1, p2)
85 +   if err != nil {
86 +       t.Errorf("%v", err)
87 +   }
88 +}
89 --
90 1.8.3.1
91
```

*Appendix D. Patches submitted to open source projects*

## Appendix E.

# GNU General Public License Version 2

As some of the code of this project is licensed under the GNU Public License Version 2 (GPLv2), I am legally required to distribute the license herewith:

GNU GENERAL PUBLIC LICENSE

Version 2, June 1991

Copyright (C) 1989, 1991 Free Software Foundation, Inc.,  
51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA  
Everyone is permitted to copy and distribute verbatim copies  
of this license document, but changing it is not allowed.

### Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change free software--to make sure the software is free for all its users. This General Public License applies to most of the Free Software Foundation's software and to any other program whose authors commit to using it. (Some other Free Software Foundation software is covered by the GNU Lesser General Public License instead.) You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs; and that you know you can do these things.

To protect your rights, we need to make restrictions that forbid anyone to deny you these rights or to ask you to surrender the rights. These restrictions translate to certain responsibilities for you if you distribute copies of the software, or if you modify it.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their

rights.

We protect your rights with two steps: (1) copyright the software, and (2) offer you this license which gives you legal permission to copy, distribute and/or modify the software.

Also, for each author's protection and ours, we want to make certain that everyone understands that there is no warranty for this free software. If the software is modified by someone else and passed on, we want its recipients to know that what they have is not the original, so that any problems introduced by others will not reflect on the original authors' reputations.

Finally, any free program is threatened constantly by software patents. We wish to avoid the danger that redistributors of a free program will individually obtain patent licenses, in effect making the program proprietary. To prevent this, we have made it clear that any patent must be licensed for everyone's free use or not licensed at all.

The precise terms and conditions for copying, distribution and modification follow.

GNU GENERAL PUBLIC LICENSE  
TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

0. This License applies to any program or other work which contains a notice placed by the copyright holder saying it may be distributed under the terms of this General Public License. The "Program", below, refers to any such program or work, and a "work based on the Program" means either the Program or any derivative work under copyright law: that is to say, a work containing the Program or a portion of it, either verbatim or with modifications and/or translated into another language. (Hereinafter, translation is included without limitation in the term "modification".) Each licensee is addressed as "you".

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running the Program is not restricted, and the output from the Program is covered only if its contents constitute a work based on the Program (independent of having been made by running the Program). Whether that is true depends on what the Program does.

1. You may copy and distribute verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and give any other recipients of the Program a copy of this License along with the Program.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Program or any portion of it, thus forming a work based on the Program, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

a) You must cause the modified files to carry prominent notices stating that you changed the files and the date of any change.

b) You must cause any work that you distribute or publish, that in whole or in part contains or is derived from the Program or any part thereof, to be licensed as a whole at no charge to all third parties under the terms of this License.

c) If the modified program normally reads commands interactively when run, you must cause it, when started running for such interactive use in the most ordinary way, to print or display an announcement including an appropriate copyright notice and a notice that there is no warranty (or else, saying that you provide a warranty) and that users may redistribute the program under these conditions, and telling the user how to view a copy of this License. (Exception: if the Program itself is interactive but does not normally print such an announcement, your work based on the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Program, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Program, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Program.

In addition, mere aggregation of another work not based on the Program with the Program (or with a work based on the Program) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may copy and distribute the Program (or a work based on it,

*Appendix E. GNU General Public License Version 2*

under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you also do one of the following:

- a) Accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
- b) Accompany it with a written offer, valid for at least three years, to give any third party, for a charge no more than your cost of physically performing source distribution, a complete machine-readable copy of the corresponding source code, to be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
- c) Accompany it with the information you received as to the offer to distribute corresponding source code. (This alternative is allowed only for noncommercial distribution and only if you received the program in object code or executable form with such an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for making modifications to it. For an executable work, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the executable. However, as a special exception, the source code distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

If distribution of executable or object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place counts as distribution of the source code, even though third parties are not compelled to copy the source along with the object code.

4. You may not copy, modify, sublicense, or distribute the Program except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense or distribute the Program is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

5. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Program or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by

modifying or distributing the Program (or any work based on the Program), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Program or works based on it.

6. Each time you redistribute the Program (or any work based on the Program), the recipient automatically receives a license from the original licensor to copy, distribute or modify the Program subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties to this License.

7. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Program at all. For example, if a patent license would not permit royalty-free redistribution of the Program by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system, which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

8. If the distribution and/or use of the Program is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Program under this License may add an explicit geographical distribution limitation excluding

*Appendix E. GNU General Public License Version 2*

those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.

9. The Free Software Foundation may publish revised and/or new versions of the General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of this License, you may choose any version ever published by the Free Software Foundation.

10. If you wish to incorporate parts of the Program into other free programs whose distribution conditions are different, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

END OF TERMS AND CONDITIONS



## How to Apply These Terms to Your New Programs

If you develop a new program, and you want it to be of the greatest possible use to the public, the best way to achieve this is to make it free software which everyone can redistribute and change under these terms.

To do so, attach the following notices to the program. It is safest to attach them to the start of each source file to most effectively convey the exclusion of warranty; and each file should have at least the "copyright" line and a pointer to where the full notice is found.

```
{{description}}  
Copyright (C) {{year}} {{fullname}}
```

```
This program is free software; you can redistribute it and/or modify  
it under the terms of the GNU General Public License as published by  
the Free Software Foundation; either version 2 of the License, or  
(at your option) any later version.
```

```
This program is distributed in the hope that it will be useful,  
but WITHOUT ANY WARRANTY; without even the implied warranty of  
MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the  
GNU General Public License for more details.
```

```
You should have received a copy of the GNU General Public License along  
with this program; if not, write to the Free Software Foundation, Inc.,  
51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA.
```

Also add information on how to contact you by electronic and paper mail.

If the program is interactive, make it output a short notice like this when it starts in an interactive mode:

```
Gnomovision version 69, Copyright (C) year name of author  
Gnomovision comes with ABSOLUTELY NO WARRANTY; for details type `show w'.  
This is free software, and you are welcome to redistribute it  
under certain conditions; type `show c' for details.
```

The hypothetical commands `show w' and `show c' should show the appropriate parts of the General Public License. Of course, the commands you use may be called something other than `show w' and `show c'; they could even be mouse-clicks or menu items--whatever suits your program.

You should also get your employer (if you work as a programmer) or your school, if any, to sign a "copyright disclaimer" for the program, if necessary. Here is a sample; alter the names:

```
Yoyodyne, Inc., hereby disclaims all copyright interest in the program  
'Gnomovision' (which makes passes at compilers) written by James Hacker.
```

*Appendix E. GNU General Public License Version 2*

{signature of Ty Coon}, 1 April 1989  
Ty Coon, President of Vice

This General Public License does not permit incorporating your program into proprietary programs. If your program is a subroutine library, you may consider it more useful to permit linking proprietary applications with the library. If this is what you want to do, use the GNU Lesser General Public License instead of this License.