

# Bluetooth for Programmers

**Albert Huang**

`albert@csail.mit.edu`

**Larry Rudolph**

`rudolph@csail.mit.edu`

## **Bluetooth for Programmers**

by Albert Huang and Larry Rudolph

Copyright © 2005 Albert Huang, Larry Rudolph

TODO

# Chapter 1. Introduction

In a single phrase, Bluetooth is *a way for devices to communicate with each other wirelessly over short distances*. A comprehensive set of documents, called the Bluetooth Specifications, describes in gory detail exactly how they accomplish this, but the basic idea is about wireless, short-range communication.

TODO

## 1.1. Understanding Bluetooth as a software developer

Developing applications that make use of Bluetooth communication is straightforward and easy although it may seem difficult due to its unusually wide scope. Technologies names or specifications, often refer to something very specific and with a narrow scope. Ethernet, for example, describes how to connect a bunch of machines together to form a simple network, but that's about it. TCP/IP describe two specific communication protocols that form the basis of the Internet, but they're just two protocols. Similarly, HTTP is the basis behind the World-Wide-Web, but also boils down to a simple protocol. But if someone asked you to describe the Internet, where would you start? What would you explain? You might describe Ethernet, TCP/IP, email, or the World-Wide-Web, or all of them at once. The hard part is knowing where to start because there is so much to describe at so many different levels. On the other hand, if a software developer approached you and wanted to know about Internet programming - how to connect one computer on the Internet to the other and send data back and forth, you probably wouldn't bother describing the details of Ethernet or email, precisely because they are both technologies aren't central to answering that question. Sure, you might mention email as an example of what Internet programming can accomplish, or describe Ethernet to give context on how the connections are implemented, but what you'd really want to describe is TCP/IP programming.

In many ways, the word Bluetooth is like the word Internet because it encompasses a wide range of subjects. Similar to Ethernet or USB, Bluetooth defines a lot of physical on-the-wire stuff like on which radio frequencies to transmit and how to modulate and demodulate signals. Similar to Voice-over-IP protocols used in many Internet applications, Bluetooth also describes how to transmit audio between devices. But Bluetooth also specifies everything in between! It's no wonder that the Bluetooth specifications are thousands upon thousands of pages.

Despite all that Bluetooth encompasses, a programmer only needs to know a small fraction of what's laid out in the specifications. When a software developer approaches to ask about how to get started with Bluetooth programming, you really only need to describe how to connect one Bluetooth device to another, and how to transfer data between the two. Sure, it helps to know a bit about the rest of Bluetooth, but there's no need to go into the specifics of the algorithms that Bluetooth devices use to choose on their radio frequencies. The bad news is that Bluetooth is more than just a replacement for a USB or ethernet cable. Most network application do not need to if their machine is connected to the network via a physical ethernet cable or a wireless 802.11 connection, they do need to know if the connection is Bluetooth. The good news, is that they do not need to know very much.

## 1.2. Bluetooth Programming Concepts

The previous section gave a general overview of Bluetooth as a communications technology, and information that's useful to know about Bluetooth but isn't absolutely necessary to create functional programs. This section focuses specifically on explaining the parts of Bluetooth that concern a software developer. Throughout the rest of this chapter, we'll often present Bluetooth concepts side by side with concepts from Internet programming. This is in part because the vast majority of network programmers are already familiar with TCP/IP to some degree. It is also because Bluetooth programming shares so much in common with Internet programming, and it makes sense to explain a new idea in terms of an old idea when they're not all that different.

Although Bluetooth was designed from the ground up, independently of the Ethernet and TCP/IP protocols, it is quite reasonable to think of Bluetooth programming in the same way as Internet programming. Fundamentally, they have the same principles of one device communicating and exchanging data with another device.

The different parts of network programming can be separated into several components

- Choosing a device with which to communicate
- Figuring out how to communicate with it
- Making an outgoing connection
- Accepting an incoming connection
- Sending and receiving data

Some of these components do not apply to all models of network programming. In a connectionless model, for example, there is no notion of establishing a connection. Some parts can be trivial in certain scenarios and quite complex in another. If the numerical IP address of a server is hard-coded into a client program, for example, then choosing a device is no choice at all. In other cases, the program may need to consult numerous lookup tables and perform several queries before it knows its final communication endpoint.

### 1.2.1. Choosing a communication partner

Every Bluetooth chip ever manufactured is imprinted with a globally unique 48-bit address, which we will refer to as the *Bluetooth address* or *device address*. This is identical in nature to the MAC addresses of Ethernet<sup>1</sup>, and both address spaces are actually managed by the same organization - the IEEE Registration Authority. These addresses are assigned at manufacture time and are intended to be unique and remain static for the lifetime of the chip. It conveniently serves as the basic addressing unit in all of Bluetooth programming.

For one Bluetooth device to communicate with another, it must have some way of determining the other device's Bluetooth address. This address is used at all layers of the Bluetooth communication process,

from the low-level radio protocols to the higher-level application protocols. In contrast, TCP/IP network devices that use Ethernet as their data link layer discard the 48-bit MAC address at higher layers of the communication process and switch to using IP addresses. The principle remains the same, however, in that the unique identifying address of the target device must be known to communicate with it.

In both cases, the client program will often not have advance knowledge of these target addresses. In Internet programming, the user will typically supply a host name, such as `www.kernel.org`, which the client must translate to a physical IP address using the Domain Name System (DNS). In Bluetooth, the user will typically supply some user-friendly name, such as "My Phone", and the client translates this to a numerical address by searching nearby Bluetooth devices and checking the name of each device.

### 1.2.1.1. Device Name

Since humans do not deal well with 48-bit numbers like `0x000EED3D1829` (in much the same way we do not deal well with numerical IP addresses like `64.233.161.104`), Bluetooth devices will almost always have a user-friendly name. This name is usually shown to the user in lieu of the Bluetooth address to identify a device, but ultimately it is the Bluetooth address that is used in actual communication. For many machines, such as cell phones and desktop computers, this name is configurable and the user can choose an arbitrary word or phrase. There is no requirement for the user to choose a unique name, which can sometimes cause confusion when many nearby devices have the same name. When sending a file to someone's phone, for example, the user may be faced with the task of choosing from 5 different phones, each of which is named "My Phone".

Although names in Bluetooth differ from Internet names in that there is no central naming authority and names can sometimes be the same, the client program still has to translate from the user-friendly names presented by the user to the underlying numerical addresses. In TCP/IP, this involves contacting a local nameserver, issuing a query, and waiting for a result. In Bluetooth, where there are no nameservers, a client will instead broadcast inquiries to see what other devices are nearby and query each detected device for its user-friendly name. The client then chooses whichever device has a name that matches the one supplied by the user.

### 1.2.1.2. Searching for nearby devices

THIS SHOULD REALLY BE A SIDE NOTE

Device discovery, the process of searching for and detecting nearby Bluetooth devices is often a confusing and irritating subject for Bluetooth developers and users. Why's that, you might ask? Well, the source of this aggravation stems from the fact that it can take a long time to detect nearby Bluetooth devices. To be specific, if you have a Bluetooth cell phone and a Bluetooth laptop sitting next to each other on your desk and you want your phone to make a connection to your laptop, it will usually take an average of 5 seconds before it detects your laptop, and sometimes as long as 10-15 seconds. This might not seem like that much time, but if you put it in context and realize that while it's performing its search, the phone is changing frequencies more than a thousand times a second and there are only 79 possible frequencies<sup>2</sup> that it can transmit on, then you'd start to wonder why they don't find each other in the

blink of an eye. The technical reasons for this aren't very interesting, but it's mostly due to the result of a strangely designed search algorithm. Suffice to say, device discovery may often take much longer than you'd like it to.

## 1.2.2. Choosing a transport protocol

Once our client application has determined the address of the host machine it wants to connect to, it must determine which transport protocol to use. This section describes the Bluetooth transport protocols closest in nature to the most commonly used Internet protocols. Consideration is also given to how the programmer might choose which protocol to use based on the application requirements.

Both Bluetooth and Internet programming involve using numerous different transport protocols, some of which are stacked on top of others. In TCP/IP, many applications use either TCP or UDP, both of which rely on IP as an underlying transport. TCP provides a connection-oriented method of reliably sending data in streams, and UDP provides a thin wrapper around IP that unreliably sends individual datagrams of fixed maximum length. There are also protocols like RTP for applications such as voice and video communications that have strict timing and latency requirements.

While Bluetooth does not have exactly equivalent protocols, it does provide protocols which can often be used in the same contexts as some of the Internet protocols.

### 1.2.2.1. RFCOMM

The RFCOMM protocol provides roughly the same service and reliability guarantees as TCP. Although the specification explicitly states that it was designed to emulate RS-232 serial ports (to make it easier for manufacturers to add Bluetooth capabilities to their existing serial port devices), it is quite simple to use it in many of the same scenarios as TCP.

In general, applications that use TCP are concerned with having a point-to-point connection over which they can reliably exchange streams of data. If a portion of that data cannot be delivered within a fixed time limit, then the connection is terminated and an error is delivered. Along with its various serial port emulation properties that, for the most part, do not concern network programmers, RFCOMM provides the same major attributes of TCP.

The biggest difference between TCP and RFCOMM from a network programmer's perspective is the choice of port number. Whereas TCP supports up to 65535 open ports on a single machine, RFCOMM only allows for 30. This has a significant impact on how to choose port numbers for server applications, and is discussed shortly.

### 1.2.2.2. L2CAP

UDP is often used in situations where reliable delivery of every packet is not crucial, and sometimes to avoid the additional overhead incurred by TCP. Specifically, UDP is chosen for its best-effort, simple datagram semantics. These are the same criteria that L2CAP satisfies as a communications protocol.

L2CAP, by default, provides a connection-oriented<sup>3</sup> protocol that sends individual datagrams of fixed maximum length. The default maximum packet size is 672 bytes, but this can be negotiated up to 65535 bytes. Being fairly customizable, L2CAP can be configured for varying levels of reliability. To provide this service, the transport protocol that L2CAP is built on<sup>4</sup> employs a transmit/acknowledgement scheme, where unacknowledged packets are retransmitted. There are three policies an application can use:

- never retransmit
- retransmit until success or total connection failure (the default)
- drop a packet and move on to queued data if a packet hasn't been acknowledged after a specified time limit (0-1279 milliseconds). This is useful when data must be transmitted in a timely manner.

Never retransmitting and dropping packets after a timeout are often referred to as *best-effort* communications. Trying to deliver a packet until it has been acknowledged or the entire connection fails is known as *reliable* communications. Although Bluetooth does allow the application to use best-effort instead of reliable communication, several caveats are in order. The reason for this is that adjusting the delivery semantics for a single L2CAP connection to another device affects *all* L2CAP connections to that device. If a program adjusts the delivery semantics for an L2CAP connection to another device, it should take care to ensure that there are no other L2CAP connections to that device. Additionally, since RFCOMM uses L2CAP as a transport, all RFCOMM connections to that device are also affected. While this is not a problem if only one Bluetooth connection to that device is expected, it is possible to adversely affect other Bluetooth applications that also have open connections.

The limitations on relaxing the delivery semantics for L2CAP aside, it serves as a suitable transport protocol when the application doesn't need the overhead and streams-based nature of RFCOMM, and can be used in many of the same situations that UDP is used in.

Given this suite of protocols and different ways of having one device communicate with another, an application developer is faced with the choice of choosing which one to use. In doing so, we will typically consider the delivery reliability required and the manner in which the data is to be sent. As shown above and illustrated in Table 1-1, we will usually choose RFCOMM in situations where we would choose TCP, and L2CAP when we would choose UDP.

**Table 1-1. A comparison of the requirements that would lead us to choose certain protocols. Best-effort streams communication is not shown because it reduces to best-effort datagram communication.**

Requirement	Internet	Bluetooth
Reliable, streams-based	TCP	RFCOMM

Requirement	Internet	Bluetooth
Reliable, datagram	TCP	RFCOMM or L2CAP with infinite retransmit
Best-effort, datagram	UDP	L2CAP (0-1279 ms retransmit)

### 1.2.3. Port numbers and the Service Discovery Protocol

The second part of figuring out how to communicate with a remote machine, once a numerical address and transport protocol are known, is to choose the port number. Almost all Internet transport protocols in common usage are designed with the notion of port numbers, so that multiple applications on the same host may simultaneously utilize a transport protocol. Bluetooth is no exception, but uses slightly different terminology. In L2CAP, ports are called *Protocol Service Multiplexers*, and can take on odd-numbered values between 1 and 32767. Don't ask why they have to be odd-numbered values, because you won't get a convincing answer. In RFCOMM, *channels* 1-30 are available for use. These differences aside, both protocol service multiplexers and channels serve the exact same purpose that ports do in TCP/IP. L2CAP, unlike RFCOMM, has a range of reserved port numbers (1-1023) that are not to be used for custom applications and protocols. This information is summarized in Table 1-2. Throughout the rest of this book, we'll often use the word *port* instead of protocol service multiplexer and channel, mostly for clarity.

**Table 1-2. Port numbers and their terminology for various protocols**

protocol	terminology	reserved/well-known ports	dynamically assigned ports
TCP	port	1-1024	1025-65535
UDP	port	1-1024	1025-65535
RFCOMM	channel	none	1-30
L2CAP	PSM	odd numbered 1-4095	odd numbered 4097 - 32765

In Internet programming, server applications traditionally make use of well known port numbers that are chosen and agreed upon at design time. Client applications will use the same well known port number to connect to a server. The main disadvantage to this approach is that it is not possible to run two server applications which both use the same port number. Due to the relative youth of TCP/IP and the large number of available port numbers to choose from, this has not yet become a serious issue.

The Bluetooth transport protocols, however, were designed with many fewer available port numbers, which means we cannot choose an arbitrary port number at design time. Although this problem is not as significant for L2CAP, which has around 15,000 unreserved port numbers, RFCOMM has only 30 different port numbers. A consequence of this is that there is a greater than 50% chance of port number collision with just 7 server applications. In this case, the application designer clearly should not arbitrarily choose port numbers. The Bluetooth answer to this problem is the Service Discovery Protocol (SDP).



Instead of agreeing upon a port to use at application design time, the Bluetooth approach is to assign ports at runtime and follow a publish-subscribe model. The host machine operates a server application, called the SDP server, that uses one of the few L2CAP reserved port numbers. Other server applications are dynamically assigned port numbers at runtime and register a description of themselves and the services they provide (along with the port numbers they are assigned) with the SDP server. Client applications will then query the SDP server (using the well defined port number) on a particular machine to obtain the information they need.

### 1.2.3.1. Service ID

This raises the question of how do clients know which service description is the one they are looking for. The easy answer would be to just assign every single service a unique identifier and be done with it. This approach has been done before, and the Internet Engineering Task Force has a standard method for developers to independently come up with their own 128-bit Universally Unique Identifiers (UUID). This is the basic idea around which SDP revolves, and this identifier is called the service's *Service ID*. Specifically, a developer chooses this UUID at design time and when the program is run, it registers its Service ID with the SDP server for that device. A client application trying to find a specific service would query the SDP server on each device it finds to see if the device offers any services with that same UUID.

UUIDs are typically referred to as a hyphen-separated series of digits of the form "XXXXXXXX-XXXX-XXXX-XXXX-XXXXXXXXXXXX", where each 'X' is a hexadecimal digit. The first segment of 8 digits corresponds to bits 1-32 of the UUID, the next segment of 4 digits is bits 33-36, and so on.

### 1.2.3.2. Service Class ID list

Although a Service ID by itself can take us a pretty long way in terms of identifying services and finding the one we want, it's really meant for custom applications built by a single development team. The Bluetooth designers wanted to distinguish between these custom applications and classes of applications that all do the same thing. For example, two different companies might both release Bluetooth software that provides audio services over Bluetooth. Even though they're completely different programs written by different people, they both do the same thing. To handle this, Bluetooth introduces a second UUID, called the *Service Class ID*. Now, the two different programs can just advertise the same Service Class ID, and all will be well in Bluetooth Land. Of course, this is only useful if the two companies agree on which Service Class ID to use.

Another thought to consider is this: what if I have a single application that can provide multiple services? For example, many Bluetooth headsets can function as a simple headphone and speaker, and advertise that service class; but they also are capable of controlling a phone call - answering an incoming call, muting the microphone, hanging up, and so on. Although it's possible to just register two separate services in this case, each with a specific service class, the Bluetooth designers chose to allow every service to have a list of service classes that the service provides. So while a single service can only have *one* Service ID, it can have many *Service Class IDs*.

NOTE: Technically, the Bluetooth specification demands that every SDP service record have an Service Class ID list with at least one entry. I think that's stupid. The Linux Bluetooth implementation does not enforce this. Should we mention this?

### Bluetooth Reserved UUIDs

Similar to the way L2CAP and TCP have reserved port numbers for special purposes, Bluetooth also has reserved UUIDs. These are mostly used for identifying predefined service classes, but also for transport protocols and profiles (Bluetooth profiles are described in Section 1.3.5). Usually, you'll see them referred to as 16-bit or 32-bit values, but they do correspond to full 128-bit UUIDs.

To get the full 128-bit UUID from a 16-bit or 32-bit number, take the *Bluetooth Base UUID* (00000000-0000-1000-8000-00805F9B34FB) and replace the leftmost segment with the 16-bit or 32-bit value. Mathematically, this is the same as:

$$128\_bit\_UUID = 16\_or\_32\_bit\_number * 2^{96} + * Bluetooth\_Base\_UUID$$

### 1.2.3.3. SDP attributes

So far, we've described SDP as a way to figure out what port and protocol a particular application service is running on, using a Service ID or a Service Class ID as a lookup key. A more general way to think of SDP is as an information database. Every record advertised by SDP is actually a list of *attributes*, where each attribute is in turn an *[ ID, value ]* pair. The attribute *ID* is a 16-bit unsigned integer that specifies the type of attribute, and the actual attribute data is described in the *value* field. A client application looking for a service can search on any of these attributes, although most will usually search on the two already mentioned.

The data in the *value* field is not restricted to only UUIDs, and can also be an integer, a boolean value, a text string, a list of any of those types, or even a list of lists. Attributes values can be of variable length - up to 4 GB long, although you'd have to be a little crazy to actually try that. All of this makes SDP a powerful way of describing services, but also makes it a bit complicated and sometimes tedious to work with.

Bluetooth defines several reserved attribute IDs which always have a special meaning, and the rest can be used any way an application designer wishes to. Some of the more common reserved attributes are:

#### Service class ID list

A list of service class UUIDs that the service provides.

#### Service ID

A single UUID identifying the specific service.

Service Name

A text string containing the name of the service.

Service Description

A text string describing the service provided.

Protocol descriptor list

A list of protocols and port numbers used by the service.

Profile descriptor list

A list of Bluetooth profile descriptor that the service complies with. Bluetooth Profiles are described in Section 1.3.5. Each descriptor consists of a UUID and a version number.

DIAGRAM!!!

#### 1.2.3.4. Is SDP even necessary?

In this section, we've seen how to avoid the pitfalls of fixed port numbers and how a client program can use SDP to find the specific Bluetooth service it's looking for. Knowing this, we should also keep in mind that SDP is not even required to create a Bluetooth application. It is perfectly possible to revert to the TCP/IP way of assigning port numbers at design time and hope to avoid port conflicts, and this might often be done to save some time. In controlled settings such as the computer laboratory or an in-house project, this is quite reasonable. Ultimately, however, to create a portable application that will run in the greatest number of scenarios, the application should use dynamically assigned ports and SDP.

#### 1.2.4. Communicating using sockets

It turns out that choosing which machine to connect to and how to connect are the most difficult parts of Bluetooth programming. Once the transport protocol and port number to communicate on are chosen, the rest of Bluetooth communications is essentially the same type of programming most network programmers are already accustomed to: sockets! A server application waiting for an incoming Bluetooth connection is conceptually the same as a server application waiting for an incoming Internet connection, and a client application attempting to establish an outbound connection behaves the same whether it is using RFCOMM, L2CAP, TCP, or UDP. For this reason, extending the socket programming framework to encompass Bluetooth is a natural approach. In this section, we'll give a brief introduction to the concepts behind socket programming. Like the rest of this chapter, we won't distract you with any code yet, just give an overview of what's involved. If you're already a seasoned veteran with socket programming, then you can skip this section, but if you're new to sockets, then read on!

### 1.2.4.1. Introducing the Socket

DIAGRAM!!! A *socket* in network programming represents the endpoint of a communication link. The idea is that from a software application's point of view, all data being passing through the link must go into or come out of the socket. First used in the 4.2BSD operating system, sockets have since become the de-facto standard for network programming.

To establish a Bluetooth connection, a program must first `create` a socket that will serve as the endpoint of the connection. Sockets are used for all types of network programming, so the first thing to do is specify what kind of socket it's going to be. In Bluetooth programming, we'll almost always be creating either L2CAP or RFCOMM sockets, so that all data sent over the sockets will be sent using the correct protocol.

When first created, the socket is not yet connected and can't be used yet for communication. To connect it, however, the application must decide if the socket will be used as a server socket to listen for incoming connections, or as a client socket to establish an outgoing connection. The process of connecting the socket depends on this choice, so we'll look at each case separately.

### 1.2.4.2. Client sockets

Client sockets are easy to understand and straightforward to use. Once the socket has been created, the client program only needs to issue the `connect` command, specifying which device to connect to, and on which port. The operating system then takes care of all the lower level details, reserving resources on the local Bluetooth adapter, searching for the remote device, forming a piconet, and establishing a connection. Once the socket is connected, it can be used for data transfer.

### 1.2.4.3. Server / Listening sockets

To get a useful data connection out of a server socket (also called listening sockets), there are three steps an application must take. First, it must `bind` the socket to local Bluetooth resources, specifying which Bluetooth adapter and which port number to use<sup>5</sup>. Second, the socket must be placed into `listening` mode. This indicates to the operating system that it should listen for connection requests on the adapter and port number chosen during the bind step. Finally, the application uses the bound and listening socket to `accept` incoming connections.

One of the major differences between a server socket and a client socket is that the server socket first created by the application can never be used for actual communication. Instead, what happens is each time the server socket accepts a new incoming connection, it spawns a brand-new socket that represents the newly established connection. The server socket then goes back to listening for more connection requests, and the application should use the newly created socket to communicate with the client.

DIAGRAM!!!

#### 1.2.4.4. Communicating using a connected socket

Once a Bluetooth application has a connected socket, using it to communicate is simple. The `send` and `receive` commands are used to... well, send and receive data. When the application is finished, it simply invokes the `close` command to disconnect the socket. Closing a listening server socket unbinds the port and stops accepting incoming connections.

#### 1.2.4.5. Nonblocking sockets with `select`

TODO

#### 1.2.4.6. Socket summary

To briefly summarize, socket programming is a multi-step process that involves 8 main operations.

Create

Allocates an unconnected socket.

Connect (client)

Establishes an outgoing connection. Implicitly forms a piconet if necessary.

Bind (server)

Reserves a port number on a local Bluetooth adapter.

Listen (server)

Instructs the operating system to begin accepting incoming connections.

Accept (server)

Waits for incoming connections.

Receive

Receive incoming data on a Bluetooth connection.

Send

Send data to the remote device of a Bluetooth connection.

Close

Disconnects a connected socket, or shuts down a listening socket.

## 1.3. Useful things to know about Bluetooth

One does not need to know very much about section

### 1.3.1. Communications range

Bluetooth devices are divided into three power classes, the only difference between them is the transmission power levels used. Table 1-3 summarizes their differences. Almost all Bluetooth-enabled cell phones, headsets, laptops, and other consumer-level Bluetooth devices are class 2 devices. There are many class 1 USB devices for sale to consumers. It is the higher class that determines the properties. If a class 1 USB device communicates with a class 2 Bluetooth cell phone, the range of the Bluetooth radio is limited by the cell phone. Class 3 Bluetooth device are rare, as their limited range heavily restricts their usefulness.

**Table 1-3. The three Bluetooth power classes**

Power class	Transmission power level	Advertised range
1	100 mW	100 meters
2	2.5 mW	10 meters
3	< 1 mW	< 1 meter

The ranges listed here are only rough estimates used for advertising purposes. In practice, one can see a much larger range when there aren't many obstructions between two devices, and a smaller range when there's a lot of radio interference or objects in between. People are actually quite good at blocking Bluetooth signals, mostly because water (which constitutes around 60% of the human body) does a great job absorbing radio waves at the frequencies used by Bluetooth. Distance is only related to the transmission power. Further distances may have higher error rates and a device might be seen outside its low-error operating range.

### 1.3.2. Communications Speed

It is also difficult to give a reliable number on the bandwidth of a Bluetooth communications channel, but ballpark figures do help. Theoretically, two Bluetooth devices have a maximum asymmetric data rate of 723.2 kilobits per second (kb/s) and a maximum symmetric data rate of 433.9 kb/s. Here, asymmetric means that only one Bluetooth device is transmitting, and symmetric means that both are transmitting to each other. In practice, the transfer rates you're likely to see will be a bit less since there's always going to be a bit of noise on wireless communications channels as well as some transport protocol overhead on each packet transmitted.

Like all wireless communications methods, the strength of a Bluetooth signal deteriorates quadratically with the distance from the source. Since weaker signals are much more likely to be corrupted by noise, the maximum communication speed between two Bluetooth devices is strongly limited by how far apart

they are. Unless you can closely control the distance and obstructions between two Bluetooth devices, it's a good idea to design a protocol that can tolerate lower communication speeds or dropped packets.

Bluetooth devices that conform to the Bluetooth 2.0 specification, which was released in late 2004, have a theoretical limit triple that of older devices (2178.1 kb/s asymmetric, 1306.9 kb/s symmetric), but at the time of this writing (October, 2005) there aren't very many Bluetooth 2.0 devices available on the market, and the vast majority of existing devices are limited by the older data rates.

### 1.3.3. Radio Frequencies and Channel Hopping

Bluetooth devices all operate in the 2.4 GHz frequency band. This means that it uses the same radio frequencies as microwaves, 802.11, and some cordless phones (the kind that attach to land lines, not cell phones). What makes Bluetooth different from the other technologies is that it divides the 2.4 GHz band into 79 channels and employs channel hopping techniques so that Bluetooth devices are always changing which frequencies they're transmitting and receiving on.

DIAGRAM!! For comparison, take a look at the way 802.11b and 802.11g work. Both of these wireless networking technologies divide the 2.4 GHz band into 14 channels that are 5 MHz wide. When a wireless network is setup, the network administrator chooses one of these channels and all 802.11 devices on that wireless network will always transmit on the radio frequency for that channel (sometimes this is done automatically by the wireless access point). If there are many wireless networks in the same area, like in an apartment building where every apartment has its own wireless router, then chances are that some of these networks will collide with each other and their overall performance will suffer.

Bluetooth, like 802.11, divides the 2.4 GHz band into channels, but that's where the similarity ends. For starters, Bluetooth has 79 channels instead of 14, and the channels are narrower (1 MHz wide instead of 5 MHz). The big difference, though, is that Bluetooth devices never stay on the same channel. An actively communicating Bluetooth device changes channels every 625 microseconds (1600 times per second). It tries to do this in a fairly random order so that no one channel is used much more than any other channel. Of course, two Bluetooth devices that are communicating with each other must hop channels together so that they're always transmitting and receiving on the same frequencies.

Supposedly, all this hopping around makes Bluetooth more robust to interference from nearby sources of evil radio waves, and allows for many Bluetooth networks to co-exist in the same place. Newer versions of Bluetooth (1.2 and greater) go even further and use *adaptive frequency hopping*, where devices will specifically avoid channels that are noisy and have high interference, (e.g. a channel that coincides with a nearby 802.11 network). How much it actually helps is debatable, but it certainly makes Bluetooth a lot more complicated than the other wireless networking technologies.

### 1.3.4. Bluetooth networks - piconets, scatternets, masters,

## and slaves

To support the intricacies of a pseudorandom channel hopping scheme, the Bluetooth designers came up with some even more confusing terminology that you might hear a lot, but doesn't matter all that much when developing Bluetooth software. Since it's mentioned in a lot of Bluetooth literature, we'll describe it here, but don't put too much effort into remembering it.

DIAGRAM!! Two or more Bluetooth devices that are communicating with each other and using the same channel hopping configuration (so that they're always using the same frequencies) form a Bluetooth *piconet*. A piconet can have up to 8 devices total. That's pretty straightforward. But how do they all agree on which frequencies to use and when to use them? That's where the *master* comes in. One device on every piconet is designated the master, and has two roles. The first is to tell the other devices (the *slaves*) which frequencies to use - the slaves all agree on the frequencies dictated by the master. The second is to make sure that the devices communicate in an orderly fashion by taking turns.

DIAGRAM!! To better understand the master device's second role, we'll compare it again with how 802.11 works. In 802.11, there is no such thing as an orderly way of transmitting. If a device has a data packet to send to another, it waits until no other device is transmitting, then waits a little more, and then transmits. If the recipient got the message, then it replies with an acknowledgment. If the sender doesn't get the acknowledgment, then it tries again. You can see how this can get messy when a lot of 802.11 devices are trying to transmit at the same time. Bluetooth, on the other hand, uses a turn-based transmission scheme, where the master of a piconet essentially informs every device when to transmit, and when to keep quiet. The big advantage here is that the data transfer rates on a Bluetooth piconet will be somewhat predictable, since every device will always have its turn to transmit. It's like the difference between a raucous town meeting where everyone is shouting to get their voice heard, and a moderated discussion where the moderator gives everyone who raises their hands a chance to speak.

The last bit of Bluetooth networking terminology here is the *scatternet*. It's theoretically possible for a single Bluetooth device to participate in more than one piconet. In practice, a lot of devices don't support this ability, but it is possible. When this happens, the two different piconets are collectively called a scatternet. Despite the impressive name, don't get too excited because scatternets don't really do a whole lot. In fact, they don't do anything at all. In order for two devices to communicate, they must be a part of the same piconet. Being part of the same scatternet doesn't help, and the device that joins the two piconets (by participating in both of them) doesn't have any special routing capabilities. Scatternet is just a name, and nothing more.

To be clear, the reason all this talk about piconets, scatternets, masters, and slaves doesn't matter is that for the most part, all of this network formation and master-slave role selection is handled automatically by Bluetooth hardware and low-level device drivers. As software developers, all we need to care about is setting up a connection between two Bluetooth devices, and the piconet issue is taken care of for us. But it does help to know what the terms mean.



### 1.3.5. Bluetooth Profiles + RFCs

Along with the simple TCP, IP, and UDP transport protocols used in Internet programming, there are a host of other protocols to specify, in great detail, methods to route data packets, exchange electronic mail, transfer files, load web pages, and more. Once standardized by the Internet Engineering Task Force in the form of Request For Comments (RFCs) <sup>6</sup>, these protocols are generally adopted by the wider Internet community. Similarly, Bluetooth also has a method for proposing, ratifying, and standardizing protocols and specifications that are eventually adopted by the Bluetooth community. The Bluetooth equivalent of an RFC is a Bluetooth Profile.

Due to the short-range nature of Bluetooth, the Bluetooth Profiles tend to be complementary to the Internet RFCs, with emphasis on tasks that can assume physical proximity. For example, there is a profile for exchanging physical location information <sup>7</sup>, a profile for printing to nearby printers <sup>8</sup>, and a profile for using nearby modems <sup>9</sup> to make phone calls. There is even a specification for encapsulating TCP/IP traffic in a Bluetooth connection, which really does reduce Bluetooth programming to Internet programming.

If you find yourself needing to implement one of the Bluetooth Profiles, you can find the specification and all the details for that particular profile on the Bluetooth website <http://www.bluetooth.org/spec>, where they are freely distributed.

## Notes

1. <http://www.ietf.org/rfc/rfc0826.txt>
2. The device discovery process actually only uses 24 of the 79 channels, which makes it even sillier
3. The L2CAP specification actually allows for both connectionless and connection-based channels, but connectionless channels are rarely used in practice. Since sending "connectionless" data to a device requires joining its piconet, a time consuming process that is merely establishing a connection at a lower level, connectionless L2CAP channels afford no advantages over connection-oriented channels.
4. Asynchronous Connection-Less logical transport
5. Most computers only have one Bluetooth adapter, so choosing a Bluetooth adapter isn't much of a choice at all
6. <http://www.ietf.org/rfc.html>
7. Local Positioning Profile
8. Basic Printing Profile
9. Dial Up Networking Profile

# Chapter 2. Bluetooth programming with Python

## - PyBluez

Now that we have an understanding of the concepts needed for Bluetooth programming, it's time to get our hands dirty and learn how to implement each of those different parts. To do this, we're going to use Python as a learning tool. Why Python, you might ask? Why not Java, or C, or (insert your favorite language here)? There are two answers to that question. The short answer is that it's just plain easy, as we'll soon find out. The long answer is that Python is a versatile and powerful dynamically typed object oriented language, providing syntactic clarity along with built-in memory management so that the programmer can focus on the algorithm at hand without worrying about memory leaks or matching braces. Additionally, there's no need to worry about compiling object files or linking against libraries or setting the correct classpaths because, for our purposes, Python "Just Works".

The only tricky part we have to deal with before getting started is making sure that we add Bluetooth support to Python. Although Python has a large and comprehensive standard library, Bluetooth is not yet part of the standard distribution. Enter PyBluez, a Python extension that provides Python programmers with access to system Bluetooth resources on GNU/Linux computers. Once we have this installed, as described in *TODO*, we're ready to get up and running.

**Note:** If you're not very comfortable with Python, don't worry! The examples used in this chapter use only the simplest parts of Python possible, and you should think of reading through the examples as if you're reading pseudocode.

## 2.1. Choosing a communication partner

Following the steps outlined in Chapter 1, the first action a Bluetooth program should take is to choose a communication partner. Example 2-1 shows a Python program that looks for a nearby device with the user-friendly name "My Phone". An explanation of the program follows.

### Example 2-1. findmyphone.py

```
from bluetooth import *

target_name = "My Phone"
target_address = None

nearby_devices = discover_devices()

for address in nearby_devices:
    if target_name == lookup_name( address ):
        target_address = address
        break
```

```

if target_address is not None:
    print "found target bluetooth device with address ", target_address
else:
    print "could not find target bluetooth device nearby"

```

A Bluetooth device is uniquely identified by its address, so choosing a communication partner amounts to picking a Bluetooth address. If only the user-friendly name of the target device is known, then two steps must be taken to find the correct address. First, the program must scan for nearby Bluetooth devices. The function `discover_devices` does this and returns a list of addresses of detected devices. Next, the program uses `lookup_name` to connect to each detected device, request its user-friendly name, and compare the result to the desired name. In this example, we just assumed that the user is always looking for the Bluetooth device named "My Phone", but we could also display the names of all the Bluetooth devices and prompt the user to choose one. Pretty easy, right?

PyBluez represents a Bluetooth address as a string of the form "xx:xx:xx:xx:xx", where each x is a hexadecimal character representing one byte of the 48-bit address, with the most significant byte listed first. Bluetooth devices in PyBluez will always be identified using an address string of this form. In the previous example, if the target device had address "01:23:45:67:89:AB", we might see the following output:

```

# python findmyphone.py
found target bluetooth device with address 01:23:45:67:89:AB

```

`discover_devices` is used in this example without any arguments, which should be sufficient for most situations, but there are a couple ways we can tweak it. When a Bluetooth device is detected during a scan, its address is cached for up to a few minutes. By default, `discover_devices` will return addresses from this cache in addition to devices that were actually detected in the current scan. To avoid these cached results, set the `flush_cache` parameter to `True`. We can also control the amount of time that `discover_devices` spends scanning with the `duration` parameter, which is specified in integer units of 1.28 seconds. This somewhat strange number is a consequence of the Bluetooth specification - device scans always last a multiple of *exactly* 1.28 seconds. It's usually not a good idea to decrease this below the default value of 8 (10.24 seconds).

`lookup_name` also takes a parameter that controls how long it spends searching. If `lookup_name` is not able to determine the user-friendly name of the specified Bluetooth device within a default value of 10 seconds, then it gives up and returns `None`. Setting the `timeout` parameter, a floating point number specified in seconds, adjusts this timeout.

An important property of Bluetooth to keep in mind is that wireless communication is never perfect, so `discover_devices()` will sometimes fail to detect devices that are in range, and `lookup_name()` will sometimes return `None` when it shouldn't. Unfortunately, it's impossible for the program to know whether these failures were a result of a bad signal or if the remote devices really aren't there any more. In these cases, it may be a good idea to try a few times, or to adjust the search durations.

## 2.2. Communicating with RFCOMM

Example 2-2 and Example 2-3 show the basics of how to establish a connection using an RFCOMM socket, transfer some data, and disconnect. In the first example, a server application waits for and accepts a single connection on RFCOMM port 1, receives a bit of data and prints it on the screen. The second example, the client program, connects to the server, sends a short message, and then disconnects.

### Example 2-2. rfcomm-server.py

```
from bluetooth import *

port = 1

server_sock=BluetoothSocket( RFCOMM )
server_sock.bind(("",port))
server_sock.listen(1)

client_sock, client_info = server_sock.accept()
print "Accepted connection from ", client_info

data = client_sock.recv(1024)
print "received [%s]" % data

client_sock.close()
server_sock.close()
```

### Example 2-3. rfcomm-client.py

```
from bluetooth import *

server_address = "01:23:45:67:89:AB"
port = 1

sock=BluetoothSocket( RFCOMM )
sock.connect((server_address, port))

sock.send("hello!!")

sock.close()
```

In the socket programming model, a socket represents an endpoint of a communication channel. Sockets are not connected when they are first created, and are useless until a call to either `connect` (client application) or `accept` (server application) completes successfully. Once a socket is connected, it can be used to send and receive data until the connection fails due to link error or user termination.

A Bluetooth socket in PyBluez is represented as an instance of the `BluetoothSocket` class, and almost all communications will use methods of this class. The constructor takes in only one parameter specifying the type of socket. This can be either `RFCOMM`, as used in these examples, or `L2CAP`, which is

described in the next section. The construction of the socket is the same for both client and server sockets.

An RFCOMM `BluetoothSocket` used to accept incoming connections must be attached to operating system resources with the `bind` method. `bind` takes in a single parameter - a tuple specifying the address of the local Bluetooth adapter to use and a port number to listen on. Usually, there is only one local Bluetooth adapter or it doesn't matter which one to use, so the empty string indicates that any local Bluetooth adapter is acceptable. Once a socket is bound, a call to `listen` puts the socket into listening mode and it is then ready to accept incoming connections with the `accept` method.

`accept` returns two values - a brand new `BluetoothSocket` object connected to the client, and the connection information as a `address, port` pair - `address` corresponds to the Bluetooth address of the connected client and `port` is the port number on the client's side of the connection.

Client programs do not need to call `bind` or the other two server-specific functions, but instead use the `connect` method to establish an outgoing connection. Like `bind`, `connect` also takes a tuple specifying an address and port number, but in this case the address can't be empty and must be a valid Bluetooth address. In Example 2-3, the client tries to connect to the Bluetooth device with address "01:23:45:67:89:AB" on port 1. This example, and Example 2-2, assumes that all communication happens on RFCOMM port 1. Section 2.4 shows how to dynamically choose ports and use SDP to search for which port a server is operating on.

Once a socket is connected, the `send` and `recv` methods can be used to, well... send and receive data. `recv` takes a parameter specifying the maximum amount of data to receive, specified in bytes, and returns the next data packet on the connection. To send a packet of data over a connection, simply pass it to `send`, which queues it up for delivery.

Once an application is finished with its Bluetooth communications, it can disconnect by calling the `close` method on a connected socket. So how does one side detect when the other has disconnected? The `recv` method will return an empty string. This is the only case where `recv` does that, which makes it a reliable way of knowing when the connection has been terminated.

We've left out error handling code in these examples for clarity, but the process is fairly straightforward. If any of the Bluetooth operations fail for some reason (e.g. connection timeout, no local bluetooth resources are available, etc.) then a `BluetoothError` is raised with an error message indicating the reason for failure.

## 2.3. Communicating with L2CAP

Example 2-4 and Example 2-5 demonstrate the basics of using L2CAP as a transport protocol. You'll notice that using L2CAP sockets is almost identical to using RFCOMM sockets.

**Example 2-4. l2cap-server.py**

```

from bluetooth import *

port = 0x1001

server_sock=BluetoothSocket( L2CAP )
server_sock.bind(("",port))
server_sock.listen(1)

client_sock,address = server_sock.accept()
print "Accepted connection from ",address

data = client_sock.recv(1024)
print "received [%s]" % data

client_sock.close()
server_sock.close()

```

**Example 2-5. l2cap-client.py**

```

from bluetooth import *

sock=BluetoothSocket(L2CAP)

bd_addr = "01:23:45:67:89:AB"
port = 0x1001

sock.connect((bd_addr, port))

sock.send("hello!!")

sock.close()

```

Aside from passing in L2CAP as a parameter to the BluetoothSocket constructor instead of RFCOMM, the only major difference between these examples and the RFCOMM examples from the previous section is the choice of port number. Remember that in L2CAP, we're strictly limited to odd-valued port numbers between 4097 and 32765. Usually, we'll use hexadecimal notation when referring to L2CAP port numbers, just because they tend to look a little cleaner.

**2.3.1. Maximum Transmission Unit**

As a datagram-based protocol, packets sent on L2CAP connections have an upper size limit. Although this has a small default value of 672 bytes, it can be adjusted. Each device at the endpoint of a connection maintains an *incoming maximum transmission unit (MTU)*, which specifies the maximum size packet it can receive. If both devices adjust their incoming MTU settings, then it is possible to change the MTU of

the entire connection beyond the 672 byte default up to 65535 bytes and as low as 48 bytes. In PyBluez, the `set_l2cap_mtu` function is used to adjust this value.

```
set_l2cap_mtu( l2cap_sock, new_mtu )
```

This method is fairly straightforward, and takes two parameters. `l2cap_sock` should be a connected L2CAP BluetoothSocket, and `new_mtu` is an integer specifying the incoming MTU for the local computer. Calling this function affects only the specified socket, and does not change the MTU for any other socket. Here's an example of how we might use it to raise the MTU:

```
l2cap_sock = BluetoothSocket( L2CAP )
.
. # connect the socket. This must be done before setting the MTU!
.
set_l2cap_mtu( l2cap_sock, 65535 )
```

If you do find yourself using this function, don't forget that both devices involved in a connection should raise their MTU settings. It is possible for each side to have a different MTU, but that just gets confusing.

### 2.3.2. Best-effort transmission

Although we expressed reservations about using best-effort L2CAP channels in Section 1.2.2.2, there are some cases where we might prefer best-effort semantics over reliable semantics. For example, if we're sending time-critical data such as audio or video data, it may be more important to forget about a few bad packets and keep sending at a constant data rate so that the connection doesn't "skip". Adjusting the reliability semantics of a connection in PyBluez is also a simple task, and can be done with the `set_packet_timeout` function.

```
set_packet_timeout( address, timeout )
```

`set_packet_timeout` takes a Bluetooth address and a timeout, specified in milliseconds, as input and tries to adjust the packet timeout for all L2CAP and RFCOMM connections to that device. The process must have superuser privileges, and there must be an active connection to the specified address. The effects of adjusting this parameter will last as long as any active connections are open, including those which outlive the Python program. If all connections to the specified Bluetooth device are closed and new ones are re-established, then the connection reverts to the default of never timing out.

## 2.4. Service Discovery Protocol

So far we've seen how to detect nearby Bluetooth device and establish the two main types of data transport connections, all using fixed Bluetooth address and port numbers that were determined at design time. To build a truly robust Bluetooth application service, we should use dynamically allocated port numbers. In doing so, we also need to give client applications a way to determine which port the service is running on. After all, what's the point of having a server running on a random port if the clients can't

find it? Here, we'll see how to use the Service Discovery Protocol (SDP) for this purpose. To get started, Example 2-6 and Example 2-7 show the RFCOMM client and server from Section 2.2 modified to use dynamic port numbers and SDP. An explanation follows the examples.

**Example 2-6. rfcomm-server-sdp.py**

```
from bluetooth import *

port = get_available_port( RFCOMM )

server_sock=BluetoothSocket( RFCOMM )
server_sock.bind(("",port))
server_sock.listen(1)

advertise_service( server_sock, "Bluetooth Serial Port",
                   service_classes = [ SERIAL_PORT_CLASS ],
                   profiles = [ SERIAL_PORT_PROFILE ] )

client_sock, client_info = server_sock.accept()
print "Accepted connection from ", client_info

data = client_sock.recv(1024)
print "received [%s]" % data

client_sock.close()
server_sock.close()
```

**Example 2-7. rfcomm-client-sdp.py**

```
import sys
from bluetooth import *

service_matches = find_service( name = "Bluetooth Serial Port",
                                uuid = SERIAL_PORT_CLASS )

if len(service_matches) == 0:
    print "couldn't find the service!"
    sys.exit(0)

first_match = service_matches[0]
port = first_match["port"]
name = first_match["name"]
host = first_match["host"]

print "connecting to ", host

sock=BluetoothSocket( RFCOMM )
sock.connect((host, port))
sock.send("hello!!")
sock.close()
```



You'll notice right away that these examples aren't much different from the ones we saw in Section 2.2. Instead of hard-coding a port number, the server dynamically allocates a port number. After creating a bound and listening socket, the server then advertises an SDP service and continues on in the same manner as the previous examples. The client, instead of hardcoding a Bluetooth address and port number, searches for a service record and uses that information to establish a connection. In the next few pages, we'll see some more details on how all this happens.

### 2.4.1. Dynamically allocating port numbers

Instead of using a predetermined port number, a Bluetooth server application can use the `get_available_port` function to find an unused port number.

```
free_port = get_available_port( protocol )
```

This function takes a single parameter, `protocol`, which can be either `L2CAP` or `RFCOMM` and specifies which protocol the application will use. It checks each port starting from the lowest number and returns the first one that isn't being used. The server application can then use `free_port` in a call to `bind`. If no ports are available at all, then it returns `None`.

`get_available_port` only identifies free ports, and doesn't reserve them, so your application should make a call to `bind` immediately afterwards. It is possible that, in the few milliseconds of time between identifying the free port and binding it, another application could sneak by and "steal" the port number. If this happens, `bind` will raise a `BluetoothError`, so you can just repeat the process. This should almost never happen, but if you want to have a completely bug-free program that guards against this problem, you could do the following:

```
from bluetooth import *
socket = BluetoothSocket( RFCOMM )
while True:
    free_port = get_available_port( RFCOMM )
    try:
        socket.bind( ( "", free_port ) )
        break
    except BluetoothError:
        print "couldn't bind to ", free_port

# listen, accept, and the rest of the program...
```

## 2.4.2. Advertising a service

Once an application has a bound and listening socket, it can advertise a service with the local SDP server. This is done with the `advertise_service` function.

```
advertise_service( sock, name, service_id="", service_classes=[],
                  profiles=[], provider="", description="" )
```

Only the first two parameters to this function, `sock` and `name` are required, and the rest have empty defaults.

*sock*

A `BluetoothSocket` object that must already be bound and listening.

*name*

A short text string describing the name of the service.

*service\_id*

Optional. The service ID of the service, specified as a string of the form "XXXXXXXX-XXXX-XXXX-XXXX-XXXXXXXXXXXX", where each 'X' is a hexadecimal digit.

*service\_classes*

Optional. A list of service class IDs, each of which can be specified as a full 128-bit UUID in the form "XXXXXXXX-XXXX-XXXX-XXXX-XXXXXXXXXXXX", or as a reserved 16-bit UUID in the form "XXXX". A number of predefined UUIDs can be used here, such as `SERIAL_PORT_CLASS`, or `BASIC_PRINTING_CLASS`. See the PyBluez documentation for a full list of predefined service class IDs.

*profiles*

Optional. A list of profiles. Each item of the list should be a `(uuid, version)` tuple. A number of predefined profiles can be used here, such as `SERIAL_PORT_PROFILE`, or `LAN_ACCESS_PROFILE`. See the PyBluez documentation for a full list of predefined profiles.

*provider*

Optional. A short text string describing the provider of the service.

*description*

Optional. A short text string describing the actual service.

Calling `advertise_service` will register a service record with the local SDP server. To unregister the service, use the function `stop_advertising`.

```
stop_advertising( sock )
```

This function takes a single parameter, *sock*, which is the socket originally used to advertise the service. Another way to unregister a service is to simply close the socket, which will automatically call `stop_advertising`.

### 2.4.3. Searching for and browsing services

To find a single service, or get a listing of services on one or multiple nearby Bluetooth devices, we use the function `find_service`.

```
results = find_service( name = None, uuid = None, address = None )
```

Without any arguments at all, `find_service` returns a listing of all services offered by all nearby Bluetooth devices. If there are a lot of Bluetooth devices in range, this could take a long time! Three optional parameters to this function can be used to filter the search results:

*name*

Optional. Restricts search results to services with this name. In the special case that this is "localhost", then the local SDP server is searched.

*uuid*

Optional. Restricts search results to services with any attribute value matching this *uuid*. Note that the matching UUID could be either the service ID, or an entry in the service class ID list, or an entry in the profiles list.

*address*

Optional. Only searches the Bluetooth device with this *address*.

The results of this search is a list of dictionary objects. Each dictionary has eight keys, which describe the corresponding service. The value for a key may be `None`, which indicates that it wasn't specified in the service record. The keys and their values are:

"host "

The bluetooth address of the device advertising the service

"name "

The name of the service being advertised.

"description"

A description of the service.

"provider"

The provider of the service.

"protocol"

A text string indicating which transport protocol the service is using. This can take on one of three values: "RFCOMM", "L2CAP", or "UNKNOWN".

"port"

If "protocol" is either "RFCOMM" or "L2CAP", then this is an integer indicating which port number the service is running on.

"service-classes"

A list of service class IDs, in the same format as used for `advertise_service`

"profiles"

A list of profiles, in the same format as used for `advertise_service`

## 2.5. Advanced usage

Although the techniques described in this chapter so far should be sufficient for most Bluetooth applications with simple and straightforward requirements, some applications may require more advanced functionality or finer control over the Bluetooth system resources. This section describes asynchronous Bluetooth communications and the `_bluetooth` module.

### 2.5.1. Asynchronous socket programming with `select`

In the communications routines described so far, there is usually some sort of waiting involved. During this time, the controlling thread blocks and can't do anything else, such as respond to user input or display progress information. To avoid these pitfalls of *synchronous* programming, it is possible to use multiple threads of control, with one thread dedicated to each task that requires some waiting. That can get quite hairy and complicated, though, so instead we'll turn to using *asynchronous* techniques as a solution.

The first step in asynchronous programming is to switch the sockets to *non-blocking* mode, so that all the operations that would block (wait) beforehand return immediately instead. The idea is "Don't wait for something to happen. Just get it started and we'll figure it out later". To switch a socket into non-blocking mode, use the `setblocking` method and pass it `False`. Conversely, to switch back into blocking mode, pass it `True`. For example:

```
from bluetooth import *
sock = BluetoothSocket( RFCOMM )
sock.setblocking( False )
s.bind(("", get_available_port( RFCOMM )))
# ...
```

The `setblocking` method must be called on every socket that you want to switch to nonblocking mode. This includes sockets that are returned by the `accept` method.

The next step in asynchronous programming is the "Figure it out" step, where the program determines if anything happened. The idea here is to consolidate all of the things a program can wait on into one place. Then, when anything happens, some data is received or the user types something or a timer fires, the program can deal with it immediately. To do this, we can use the `select` module, which comes as part of the standard Python distribution. Within the `select` module is the `select` function, which is what we'll be using extensively.

```
from select import *

can_rd, can_wr, has_exc = select( to_read, to_write, to_exc, [timeout] )
```

`select` can wait for three different types of events - read events, write events, and exceptions. The first three parameters are lists of objects - which list an object is in determines which type of event `select` will detect for that object. An object can be in multiple lists. As soon as `select` detects an event, it returns three more lists, each of which contains objects from the original lists where event activity was detected. The fourth parameter to `select` is optional and specifies a timeout as a floating point number in seconds. If no events are detected before the timeout elapses, then `select` returns three empty lists.

So what exactly are the different types of events? Some of these should be pretty obvious, but others have been shoehorned in. Table 2-1 summarizes which list to put a socket in for detecting specific events.

**Table 2-1. `select` events**

event	list
outgoing connection established (client)	write
data received on socket	read
incoming connection accepted (server)	read
can send data (i.e. send buffer not full)	write
disconnected	read

You'll notice a couple things here. First, the third list for exceptions isn't used at all. `select` is meant to be used for all different types of objects, and the third list is used elsewhere, just not in Bluetooth. Second, we didn't mention searching for nearby devices or SDP. We'll talk about the device discovery process next, but unfortunately there aren't yet any asynchronous techniques for SDP. In this case, you'll have to rely on threads for non-blocking operations, but hopefully that will change in the future.

## 2.5.2. Asynchronous device discovery

Asynchronously searching for nearby devices and determining their user-friendly names can also be done with `select`, but is a bit more complicated and involves the use of a new class, the

```

// build a command packet to send to the bluetooth microcontroller
cmd_param.handle = cr->conn_info->handle;
cmd_param.flush_timeout = htobs(timeout);
rq.ogf = OGF_HOST_CTL;
rq.ocf = 0x28;
rq.cparam = &cmd_param;
rq.clen = sizeof(cmd_param);
rq.rparam = &cmd_response;
rq.rlen = sizeof(cmd_response);
rq.event = EVT_CMD_COMPLETE;

// send the command and wait for the response
status = hci_send_req( dd, &rq, 0 );
if( status != 0 ) goto cleanup;

if( cmd_response.status ) {
    status = -1;
    errno = bt_error(cmd_response.status);
}

cleanup:
    free(cr);
    if( dd >= 0 ) close(dd);
    return status;
}

```

On success, the packet timeout for the low level connection to the specified device is set to `timeout * 0.625` milliseconds. A timeout of 0 is used to indicate infinity, and is how to revert back to a reliable connection. The bulk of this function is comprised of code to construct the command packets and response packets used in communicating with the Bluetooth controller. The Bluetooth Specification defines the structure of these packets and the magic number 0x28. In most cases, BlueZ provides convenience functions to construct the packets, send them, and wait for the response. Setting the packet timeout, however, seems to be so rarely used that no convenience function for it currently exists.

## 3.6. Chapter Summary

This chapter has provided an introduction to Bluetooth programming with BlueZ. The concepts covered in chapter 2 were presented here in greater detail with examples on how to implement them in BlueZ. Many other useful aspects of BlueZ were left out for brevity. Specifically, the command line tools and utilities that are distributed with BlueZ, such as `hciconfig`, `hcitool`, `sdptool`, and `hcidump`, are not described here. These utilities, which are invaluable to a serious Bluetooth developer, are already well documented. Only the simplest aspects of using the Service Discovery Protocol were covered - just enough to search for and advertise services. Additionally, other socket types such as `BTPROTO_SCO` and `BTPROTO_BNEP` were left out, as they are not crucial to forming a working knowledge of programming with BlueZ. Unfortunately, as of now there is no official API reference to refer to, so more curious readers are advised to download and examine the BlueZ source code <sup>4</sup>.

## Notes

1. <http://www.bluez.org/lists.html> (<http://www.bluez.org/lists.html>)
2. <https://www.bluetooth.org/foundry/assignnumb/document/baseband>
3. Bluetooth terminology refers to this as the ACL connection
4. available at <http://www.bluez.org>

# Chapter 4. Bluetooth development tools

**Note:** need to re-word this introduction now that the chapter is after 2 and 3

There are three major parts of the Bluetooth subsystem in Linux - the kernel level routines, the `libbluetooth` development library, and the user level tools and daemons. Roughly speaking, the kernel part is responsible for managing the Bluetooth hardware resources that are attached to a machine, wrestling with all the different types of bluetooth adapters that are out there, and presenting a unified interface to the rest of the system that allows any Bluetooth application to work with any Bluetooth hardware.

The `libbluetooth` development library takes the interface exposed by the Linux kernel and provides a set of convenient data structures and functions that can be used by Bluetooth programmers. It abstracts some of the most commonly performed operations (such as detecting nearby Bluetooth devices) and provides simple functions that can be invoked to perform common tasks.

The user-level tools are the programs that a typical end-user or programmer might use to leverage the computer's Bluetooth capabilities, while the daemons are constantly running programs that use the Bluetooth development library to manage the system's Bluetooth resources in the ways configured by the user. The BlueZ developers strive to make these tools and daemons as straightforward to use as possible, while also providing enough flexibility to meet every user's needs. As a software developer, you'll be interacting with the user-level tools the most, so we'll focus on introducing them in this chapter.

There are six command-line tools provided with BlueZ that are indispensable when configuring Bluetooth on a machine and debugging applications. We'll give some short descriptions here on how they're useful, and show some examples on how to use them. For full information on how to use them, you should consult the `man` pages that are distributed with the tools, or invoke each tool with the `-h` flag. This section serves mainly to give you an idea of what the tools are and which one to use for what scenario.

## 4.1. `hciconfig`

`hciconfig` is used to configure the basic properties of Bluetooth adapters. When invoked without any arguments, it will display the status of the adapters attached to the local machine. In all other cases, the usage follows the form:

```
# hciconfig <device> <command> <arguments...>
```

where `<device>` is usually `hci0` (`hci1` specifies the second Bluetooth adapter if you have two, `hci2` is the third, and so on). Most of the commands require superuser privileges. Some of the most useful ways to use this tool are:



Display the status of recognized Bluetooth adapters

```
# hciconfig
hci0:  Type: USB
      BD Address: 00:0F:3D:05:75:26 ACL MTU: 192:8 SCO MTU: 64:8
      UP RUNNING PSCAN ISCAN
      RX bytes:505075 acl:31 sco:0 events:5991 errors:0
      TX bytes:25758 acl:24 sco:0 commands:1998 errors:0
```

Each Bluetooth adapter recognized by BlueZ is displayed here. In this case, there is only one adapter, `hci0`, and it has Bluetooth Address `00:0F:3D:05:75:26`. The "UP RUNNING" part on the second line indicates that the adapter is enabled. "PSCAN" and "ISCAN" refer to Inquiry Scan and Page Scan, which are described a few paragraphs down. The rest of the output is mostly statistics and a few device properties.

Enable / Disable an adapter

The `up` and `down` commands can be used to enable and disable a Bluetooth adapter. To check whether or not a device is enabled, use `hciconfig` without any arguments.

```
# hciconfig hci0 down
# hciconfig
hci0:  Type: USB
      BD Address: 00:0F:3D:05:75:26 ACL MTU: 192:8 SCO MTU: 64:8
      DOWN
      RX bytes:505335 acl:31 sco:0 events:5993 errors:0
      TX bytes:25764 acl:24 sco:0 commands:2000 errors:0
# hciconfig hci0 up
# hciconfig
hci0:  Type: USB
      BD Address: 00:0F:3D:05:75:26 ACL MTU: 192:8 SCO MTU: 64:8
      UP RUNNING PSCAN ISCAN
      RX bytes:505075 acl:31 sco:0 events:5991 errors:0
      TX bytes:25758 acl:24 sco:0 commands:1998 errors:0
```

Display and change the user-friendly name of an adapter.

The `name` command is fairly straightforward, and can be used to display and change the user-friendly name of the Bluetooth adapter.

```
# hciconfig hci0 name
hci0:  Type: USB
      BD Address: 00:0F:3D:05:75:26 ACL MTU: 192:8 SCO MTU: 64:8
      Name: 'Trogdor'
# hciconfig hci0 name 'StrongBad'
# hciconfig hci0 name
hci0:  Type: USB
      BD Address: 00:0F:3D:05:75:26 ACL MTU: 192:8 SCO MTU: 64:8
      Name: 'StrongBad'
```

"Hide" an adapter, or show it to the world.

The Inquiry Scan and Page Scan settings for a Bluetooth adapter determine whether it is detectable by nearby Bluetooth devices, and whether it will accept incoming connection requests, respectively. Don't be confused by the names! These control whether the adapter *responds* to inquiries and to pages (connection requests), not whether it makes them.<sup>1</sup>

**Table 4-1. Inquiry Scan and Page Scan**

Inquiry Scan	Page Scan	Interpretation	command
On	On	This is the default. The adapter is detectable by other Bluetooth devices, and will accept incoming connection requests	<code>piscan</code>
Off	On	Although not detectable by other Bluetooth devices, the adapter still accepts incoming connection requests by devices that already know the Bluetooth address of the adapter.	<code>piscan</code>
On	Off	The adapter is detectable by other Bluetooth devices, but it will not accept any incoming connections. This is mostly useless.	<code>iscan</code>
Off	Off	The adapter is not detectable by other Bluetooth devices, and will not accept any incoming connections.	<code>noscan</code>

For example, the following invocation disables both Inquiry Scan and Page Scan for the first Bluetooth adapter.

```
# hciconfig hci0 noscan
```

There are many more ways to use `hciconfig`, all of which are described in the help text (`hciconfig -h`) and the man pages (`man hciconfig`). The key thing to remember is that `hciconfig` is the tool to

use for any non-connection related settings for a Bluetooth adapter.

NOTE: Changes made by `hciconfig` are only temporary, and the effects are erased after a reboot or when the device is disabled and enabled again. `hcid.conf` should be used To make a change permanent (e.g. to permanently change the user-friendly name).

NOTE: The name `hciconfig` comes from the term Host Controller Interface (HCI). It refers to the protocol that a computer uses to communicate with the Bluetooth microcontroller that resides on the computer's Bluetooth adapter. HCI is used to do all the dirty work of configuring the adapter and setting up connections. The commands `hciconfig` and `hcitool` are so named to emphasize that they are used for the low-level Bluetooth operations that, while important, can't actually be used for communicating with other Bluetooth devices.

## 4.2. hcitool

`hcitool` has two main uses. The first is to search for and detect nearby Bluetooth devices, and the second is to test and show information about low-level Bluetooth connections. In a sense, `hcitool` picks up where `hciconfig` ends - once the Bluetooth adapter starts communicating with other Bluetooth devices.

### Detecting Nearby Bluetooth devices

`hcitool scan` searches for nearby Bluetooth devices and displays their addresses and user-friendly names.

```
# hcitool scan
Scanning ...
    00:11:22:33:44:55      Cell Phone
    AA:BB:CC:DD:EE:FF      Computer-0
    01:23:45:67:89:AB      Laptop
    00:12:62:B0:7B:27      Nokia 6600
```

In this invocation, four Bluetooth devices were found. Detecting the addresses of nearby Bluetooth devices and looking up their user-friendly names are actually two separate processes, and conducting the name lookup can often take quite a long time. If you don't need the user-friendly names, then `hcitool inq` is useful for only performing the first part of the search - finding the addresses of nearby devices.

### Testing low-level Bluetooth connections

`hcitool` can be used to create piconets of Bluetooth devices and show information about locally connected piconets. Remember that piconets are just an ugly consequence of Bluetooth's fancy frequency hopping techniques. When we're writing Bluetooth software, we won't have to worry

about these low level details, just like we won't have to worry about instructing the Bluetooth adapter on which radio frequencies to use. So for application programming, this part of `hcitool` is strictly of educational use, because BlueZ automatically takes care of piconet formation and configuration in the process of establishing higher-level RFCOMM and L2CAP connections.

If you're curious about using `hcitool` for basic piconet configuration, then the `hcitool cc` and `hcitool con` commands are the first places to start. `hcitool cc` forms a piconet with another device, and is fairly straightforward to use. For example, to join a piconet with the device `00:11:22:33:44:55`

```
# hcitool cc 00:11:22:33:44:55:66
```

`hcitool con` can then be used to show information about existing piconets.

```
# hcitool con
Connections:
  < ACL 00:11:22:33:44:55 handle 47 state 1 lm MASTER
```

Here, the output of `hcitool con` tells us that the local Bluetooth adapter is the master of one piconet, and the device `00:11:22:33:44:55` is a part of that piconet. For details on the rest of the output, see the `hcitool` documentation.

NOTE: A fairly common mistake is to try to use `hcitool` to create data transport connections between two Bluetooth devices. It's important to know that even if two devices are part of the same piconet, a higher-level connection needs to be established before any application-level data can be exchanged. Creating the piconet is only the first step in the communications process.

## 4.3. sdptool

`sdptool` has two uses. The first is for searching and browsing the Service Discovery Protocol (SDP) services advertised by nearby devices. This is useful for seeing what Bluetooth profiles are implemented by another Bluetooth device such as a cellular phone or a headset. The second is for basic configuration of the SDP services offered by the local machine.

Browsing and searching for services

`sdptool browse [addr]` retrieves a list of services offered by the Bluetooth device with address `addr`. Leaving `addr` out causes `sdptool` to check all nearby devices. If `local` is used for the address, then the local SDP server is checked instead. Each service record found is then briefly described. A typical service record might look like this:

```
# sdptool browse 00:11:22:33:44:55
```

```

Browsing 00:11:22:33:44:55
Service Name: Bluetooth Serial Port
Service RecHandle: 0x10000
Service Class ID List:
  "Serial Port" (0x1101)
Protocol Descriptor List:
  "L2CAP" (0x0100)
  "RFCOMM" (0x0003)
    Channel: 1
Language Base Attr List:
  code_ISO639: 0x656e
  encoding:    0x6a
  base_offset: 0x100
Profile Descriptor List:
  "Serial Port" (0x1101)
    Version: 0x0100

```

Here, the device 00:11:22:33:44:55 is advertising a single service called "Bluetooth Serial Port" that's operating on RFCOMM channel 1. The service has the UUID 0x1101, and also adheres to the Bluetooth Serial Port Profile, as indicated by the profile descriptor list at the bottom. In general, this information should be sufficient for an application to determine whether or not this is the service that it's looking for (has UUID 0x1101), and how to connect to it (use RFCOMM channel 1).

`sdptool search` can be used to search nearby devices for a specific service, but it can only look for a handful of predefined services. It is not able to search for a service with an arbitrary UUID, this must be done programmatically. Because of this, `sdptool browse` will generally be more useful for testing and debugging applications that use SDP (e.g. to check that a service is being advertised correctly).

### Basic service configuration

`sdptool add <name>` can be used to advertise a set of predefined services, all of which are standardized Bluetooth Profiles. It cannot be used to advertise an arbitrary service with a user-defined UUID, this must be done programmatically. This means it won't be very useful for advertising a custom service.

`sdptool del <handle>` can be used to un-advertise a local service. The SDP server maintains a handle for each service that identifies it to the server - essentially a pointer to the service record. To find the handle, just look at the description of the service using `sdptool browse` and look for the line that says "Service RecHandle: ". Using the example above, the Serial Port service has the handle 0x10000, so if we were using that machine, we could issue the following command to stop advertising the service:

```
# sdptool del 0x10000
```

`sdptool` also provides commands for modifying service records (e.g. to change a UUID), that you could actually use, but probably don't want to. These, along with the `add` and `del` commands exist

more so that programmers can look at the source code of `sdptool` for examples on how to do the same in their own applications. Advertising and configuring services with C and Python are described in later chapters of this book, but you can always download the BlueZ source code at <http://www.bluez.org> and see how it's done with `sdptool`.

## 4.4. hcidump

For low-level debugging of connection setup and data transfer, `hcidump` can be used to intercept and display all Bluetooth packets sent and received by the local machine. This can be very useful for determining how and why a connection fails, and lets us examine at exactly what stage in the connection process did communications fail. `hcidump` requires superuser privileges.

When run without any arguments, `hcidump` displays summaries of Bluetooth packets exchanged between the local computer and the Bluetooth adapter as they appear. This includes packets on device configuration, device inquiries, connection establishment, and raw data. Incoming packets are preceded with the ">" greater-than symbol, and outgoing packets are preceded with the "<" less-than symbol. The length of each packet (`plen`) is also shown. For example, if we started `hcidump` in one command shell and issued the command `hcitool inq` in another, the output of `hcidump` might look like this:

```
# hcidump
HCI sniffer - Bluetooth packet analyzer ver 1.23
device: hci0 snap_len: 1028 filter: 0xffffffff
< HCI Command: Inquiry (0x01|0x0001) plen 5
> HCI Event: Command Status (0x0f) plen 4
> HCI Event: Inquiry Result (0x02) plen 15
> HCI Event: Inquiry Complete (0x01) plen 1
```

Here, we can see that one command (Inquiry) was sent out instructing the Bluetooth adapter to search for nearby devices, and three packets of size 5, 4, and 15 bytes were received: information on the status of the command, an inquiry result indicating that a nearby device was detected, and another status packet once the inquiry completed. You'll notice that used this way, `hcidump` only provides basic summaries of the packets, which is not always enough for debugging. One option is to use the `-x` flag, which causes `hcidump` to display the raw contents of every packet in hexadecimal format along with their ASCII decodings. Used in the above example, we might see the following:

```
# hcidump -X
HCI sniffer - Bluetooth packet analyzer ver 1.23
device: hci0 snap_len: 1028 filter: 0xffffffff
< HCI Command: Inquiry (0x01|0x0001) plen 5
  0000: 33 8b 9e 08 00                                     3....
> HCI Event: Command Status (0x0f) plen 4
  0000: 00 01 01 04                                         ....
```

```
> HCI Event: Inquiry Result (0x02) plen 15
0000: 01 26 75 05 3d 0f 00 01 02 00 00 01 3e d6 1f .&u.=.....>...
> HCI Event: Inquiry Complete (0x01) plen 1
0000: 00 .
```

Okay, so unless you've memorized the Bluetooth specification and can decode the raw binary packets in your head, maybe that's not as useful as we'd like. While `hcidump -x` is great for very low-level debugging of raw packets, the `-v` option gives us a nice compromise. `hcidump -v` will display as much information as it can gather from each packet, and summarize the ones it can't interpret. If used together with `-x`, it will still provide all the information for packets that it can decode, but will also show the raw hexadecimal data for all the other packets (these tend to be application-level data packets). Repeating our example once again, we might see this:

```
# hcidump -X -V
HCI sniffer - Bluetooth packet analyzer ver 1.23
device: hci0 snap_len: 1028 filter: 0xffffffff
< HCI Command: Inquiry (0x01|0x0001) plen 5
  lap 0x9e8b33 len 8 num 0
> HCI Event: Command Status (0x0f) plen 4
  Inquiry (0x01|0x0001) status 0x00 ncmd 1
> HCI Event: Inquiry Result (0x02) plen 15
  bdaddr 00:0F:3D:05:75:26 mode 1 clkoffset 0x1fd5 class 0x3e0100
> HCI Event: Inquiry Complete (0x01) plen 1
  status 0x00
```

Now, we see the packets decoded according to the Bluetooth specification, which are probably mostly meaningless to you right now, but would make sense if you found the need to read the parts of the Bluetooth specification on device inquiry. Since this is a simple example, `hcidump` is able to fully decode each packet, so we don't see any raw hexadecimal data.

As with the other utilities, there are many more ways to use `hcidump` for debugging and low-level display of Bluetooth packet communication that you can find out by reading the help text included with BlueZ.

## 4.5. `l2ping`

`l2ping` sends echo packets to another Bluetooth device and waits for a response. An echo packet is a special type of L2CAP packet that contains no meaningful data - when a Bluetooth device receives an echo packet, it should just send (echo) the packet back to the originator. This is useful for testing and analyzing L2CAP communications with another Bluetooth device. If two devices are communicating, but seem a little sluggish, then `l2ping` can provide timing information on how long it takes to send and receive packets of a certain size. The only required parameter is the address of the Bluetooth device to "ping". For example, to send echo packets to the device `01:23:45:67:89:AB`:

```
# l2ping -c 5 01:23:45:67:89:AB
Ping: 01:23:45:67:89:AB from 00:D0:F5:00:0E:B5 (data size 44) ...
44 bytes from 01:23:45:67:89:AB id 0 time 60.87ms
44 bytes from 01:23:45:67:89:AB id 1 time 55.97ms
44 bytes from 01:23:45:67:89:AB id 2 time 50.96ms
44 bytes from 01:23:45:67:89:AB id 3 time 51.94ms
44 bytes from 01:23:45:67:89:AB id 4 time 48.93ms
```

`l2ping` continues sending packets until stopped by pressing `Ctrl-C`. Other command line arguments let us control the size of the packets sent, the delay between packets, how many to send, and so on. For details on how to use these capabilities, invoke `l2ping -h`.

## 4.6. rfcmm

The `rfcomm` tool lets us establish arbitrary RFCOMM connections and treat them like serial ports. Although the RFCOMM protocol was described in the previous chapter as a general purpose transport protocol, one of its original purposes was to emulate a serial port connection between two devices. The idea was that device manufacturers who had serial-port capable devices would only need to add a Bluetooth chip to the end of the serial port controller, which requires much less modification to the original device than replacing the serial port controller. In fact, Bluetooth was even marketed as a "wireless serial cable". To utilize the serial-port emulation capabilities of Bluetooth in Linux, we use the `rfcomm` tool.

`rfcomm` can be used to connect to another device or to listen for incoming connections. A special device file is created for each connection, which user-level programs can read and write to like regular files. Data written to the device file is transmitted over Bluetooth, and reading from the device file retrieves the data received over the connection. When the device file is closed, the Bluetooth connection is terminated.

To listen for an incoming connection, we first choose which device file to bind it to. Typically, we'll use `/dev/rfcommX`, where `X` ranges from 0 - 9. Next, we choose an RFCOMM port number to listen on. To listen on RFCOMM port 20 and connect it to `/dev/rfcomm0`, we'd use the `rfcomm listen` command like this:

```
# rfcmm listen /dev/rfcomm0 20
```

Similarly, to establish an outgoing connection and serial port, we'd use the `rfcomm connect` command, but we would also specify the address of the Bluetooth device to connect to:

```
# rfcmm connect /dev/rfcomm0 01:23:45:67:89:AB 20
```

Keep in mind that in both these examples, the special device file `/dev/rfcomm0` is not a valid file until the `rfcomm` commands successfully complete. The other way of using `rfcomm` to establish outgoing



connections is to use the `rfcomm bind` command to create the device file, and only establish the Bluetooth connection when a separate program tries to access the device file. For example:

```
# rfcomm bind /dev/rfcomm0 01:23:45:67:89:AB 20
```

Using `rfcomm` in this way is sort of saying "When a program opens `/dev/rfcomm0`, make a connection to the Bluetooth device `01:23:45:67:89:AB` and send all data through that file. But if no program ever access that file, don't bother making the connection"

## 4.7. `uuidgen`

TODO

## 4.8. Obtaining BlueZ and PyBluez

**Note:** this should be an appendix

Instructions for installing the BlueZ development libraries can be found at the BlueZ website: <http://www.bluez.org> (<http://www.bluez.org>). Most modern Linux distributions should have this packaged somehow. For example, on Debian-based systems:

```
apt-get install libbluetooth1-dev bluez-utils
```

On Fedora:

```
yum install bluez-devel
```

Similarly, instructions for installing PyBluez can be found at the PyBluez website: <http://org.csail.mit.edu/pybluez>. PyBluez is included with a few Linux distributions, but TODO

## Notes

1. The idea is that Inquiry Scan and Page Scan control whether the adapter *scans* for inquiries and pages, in the same way that you might use your eyes to scan around to see if anyone is talking to you. Confusing!