

PART I

Location Services

- ▶ **CHAPTER 1:** Introducing the Android Location Service
- ▶ **CHAPTER 2:** Determining a Device's Current Location
- ▶ **CHAPTER 3:** Tracking Device Movement
- ▶ **CHAPTER 4:** Proximity Alerts

1

Introducing the Android Location Service

WHAT'S IN THIS CHAPTER?

- Providing overview of how location information is provided in Android
- Presenting an overview of GPS
- Discussing why A-GPS is used in Android
- Providing an overview of the network location provider

Location information is becoming increasingly important in the world of mobile development. Apps that were once location agnostic now make use of location information to provide a richer user experience. Being able to combine a simple web search engine with up-to-the-minute location information allows Android devices to provide a level of functionality that was previously not possible. The capability to easily retrieve and provide location data to apps is becoming a major feature of today's mobile platforms. Android provides this functionality with its location service.

Android's location service provides access to facilities that can be used to determine a device's current location. This information can be used for a wide variety of functions and can allow a device and the software that runs on it to have a better understanding of its surroundings.

METHODS USED TO DETERMINE LOCATION

Android makes use of different methods to provide location information to an app. In Android, these facilities are called *location providers*, and each has its own unique set of strengths and weaknesses. In addition, because location providers have such unique characteristics, they each lend themselves to be used differently in different situations.


The following sections give some high-level explanations as to how the different location acquisition methods work. Although an app has little control over how the providers work, it can decide which location provider to use. Understanding how each provider works goes a long way in understanding its limitations and characteristics.

GPS Provider

The Global Positioning System (GPS) uses a system of satellites orbiting the planet to help a receiver (an Android handset in this case) determine its current location. The term *GPS* refers to the entire GPS system, which consists of satellites, receivers, and the control stations that monitor and adjust it. The receiver that is located in the phone is useless without the rest of the system.

How It Works

In general, a GPS receiver uses information from the GPS satellites orbiting the earth to calculate its current location. The GPS system contains 27 satellites that continually orbit the earth, transmitting information to would-be receivers. Each satellite follows a defined path, ensuring that at least four satellites are “visible” from any point on earth at any given time. Being able to have a “line of sight” to at least four satellites is necessary to determine location using GPS. Figure 1-1 shows a depiction of the GPS satellite constellation.



SOURCE: [HTTP://GPS.GOV/MULTIMEDIA/IMAGES](http://gps.gov/multimedia/images)

FIGURE 1-1: GPS satellite constellation

Each GPS satellite in the constellation continuously transmits its current position (ephemeris data) and almanac data. The almanac data includes data about each satellite in the constellation, including orbiting data as well as information about the overall state of the system as a whole. To say it another way, ephemeris data is information about a single satellite, and almanac data is information about every satellite. Every satellite transmits both. Though both the ephemeris data and almanac data provide location data for a given satellite, the ephemeris data provides accuracy for location calculation.

To calculate its location, a GPS receiver must be able to determine its distance from multiple satellites. It does this using the ephemeris data. Included in the data that is transmitted from the satellite, along with the position data, is the time at which the transmission started. Each GPS satellite contains a highly accurate timekeeping mechanism that allows the satellite to keep its time in sync with the rest of the satellites. To produce an accurate location calculation, the GPS satellites and GPS receivers must have their clocks highly synchronized. Even the slightest difference in time can cause large errors when computing location.

Using the transmission start time, the GPS receiver can calculate the time it took for the transmission to be received (the receiver knows when the transmission ended). This calculation is made with the assumption that the radio waves that transmit the data travel at the speed of light in a vacuum (which is not always the case). Using the start time, end time, and a constant for the speed of light, a GPS receiver can calculate the distance of the satellite from the receiver.

Using the distance from multiple satellites, the GPS receiver can triangulate its current location. Essentially, the point at which all the spheres intersect is the location of the receiver. A minimum of three satellites is needed to determine a two-dimensional location (latitude and longitude). Communications from additional satellites allow a GPS receiver to determine additional positional information such as altitude. A GPS receiver will not limit itself to only four satellites. In general as the number of satellites from which the receiver can receive data increases, so does the accuracy of the location (there is an upper limit, however).

GPS is useful for determining current location, but it does have some drawbacks (especially for mobile platforms), one of which is the time it can take to calculate the current position. Before the location can be calculated, multiple satellites must be found. Many satellites are orbiting the earth, but only a handful can be “seen” at any given time because most will be below the horizon and blocked by the earth (remember, a line of sight is needed). The almanac used by the GPS system can provide assistance in determining which satellites should be used for a given location at a given time. However, if the GPS does not have a relatively current almanac, it will need to have the almanac data transmitted by a GPS satellite. This can be a slow process.

GPS Improvements

Although standard GPS can provide accurate location data, the limitations it imposes make it difficult for mobile devices to use it. To help circumvent some the limitations of standard GPS, modern mobile devices make use of assisted GPS (A-GPS) and possibly simultaneous GPS (S-GPS).

A-GPS

A-GPS uses the mobile network to transmit the GPS almanac along with other pieces of information to a mobile device. This use of the mobile network allows for faster transmission of the almanac, which may lead to faster determination of the device’s current location. In addition, because the almanac contains information about all of the GPS satellites, the device will know the approximate location of the GPS satellites in its line of sight. This will also improve the time it takes to acquire a GPS location.

Examination of the GPS configuration file provides some insight into where the A-GPS data comes from. Listing 1-1 shows an example of a GPS configuration file that is used in Android for a device located in North America.

LISTING 1-1: An example of a GPS configuration file located in /system/etc/gps.conf

```
NTP_SERVER=north-america.pool.ntp.org
XTRA_SERVER_1=http://xtra1.gpsonextra.net/xtra.bin
XTRA_SERVER_2=http://xtra2.gpsonextra.net/xtra.bin
XTRA_SERVER_3=http://xtra3.gpsonextra.net/xtra.bin
```

Listing 1-1 shows that the GPS configuration file can specify the location of the A-GPS data to download (`XTRA_SERVER_1`, `XTRA_SERVER_2`, and `XTRA_SERVER_3`) as well as a Network Time Protocol (NTP) server that can be used to coordinate time (`NTP_SERVER`). NTP can be used to force coordination of time. This is important because GPS relies heavily on the clocks of a GPS receiver and the GPS satellites being in sync. Although the use of NTP does not guarantee true time synchronization down to the millisecond, it does help to prevent large time differences. Because of the numbers used in calculating times, like the speed of light, a small difference in time can lead to large inaccuracies in location calculations.

Though most users can read `/system/etc/gps.conf`, increased permissions are required to write to the file. Generally, users should not need to edit this file.

S-GPS

Devices that use standard GPS may use the same hardware to communicate with GPS satellites and make mobile phone calls. This means that only one of these actions can take place at a time. S-GPS addresses this issue by adding additional hardware that allows the GPS radio and the cellular network radio to be operational simultaneously. The ability to have two radios active can speed up GPS data transmission because it allows the data to be received while the cellular network radio is active.

Limitations

Although GPS can provide the most accurate location data, it does have limitations that may be difficult to work around. First is the fact that a GPS receiver needs a clear path to a GPS satellite. This means that GPS receivers are unlikely to work indoors, and may even have problems outside in areas where the sky is not visible (such as dense forests). Additionally, because multiple GPS satellites are needed to produce location information, it may take a substantial amount of time to acquire a location. This is exacerbated by that fact that devices may contain low-powered GPS radios. For these reasons, other sources of location information are sometimes needed.

Objects that obstruct a GPS signal may cause the signal to be reflected before it reaches the GPS receiver. As stated earlier, the time it takes a signal to reach the GPS receiver is used to calculate the distance between the GPS satellite and the GPS receiver. GPS signals that are reflected off of objects have a different path from the GPS satellite to the GPS receiver and cause the distance calculation to be erroneous. These types of errors are called *multipath errors* and can cause the location to appear to jump from one place to another. This is often seen in urban areas where GPS signals frequently bounce off of tall buildings.

Controlling GPS

For most cases, GPS should “just work” as far as app developers are concerned. Typically, there will be no reason to interfere with the source of the A-GPS data, or when the A-GPS data should be purged and reinitialized.

However, Android does provide an API for controlling certain aspects of GPS data. The `LocationManager` class (which is introduced in detail in the next chapter) contains a `sendExtraCommand()` method that can be used to manipulate device GPS state. The `LocationManager.sendExtraCommand()` method takes three parameters: a string specifying the location provider, the extra command, and a `Bundle` that provides additional information for performing the command.

At the time of this writing, the GPS location provider supports only three extra commands:

- `delete_aiding_data`
- `force_time_injection`
- `force_extra_injection`

The `delete_aiding_data` command is used to remove the A-GPS data that has been previously downloaded. It is the only extra command that makes use of the `Bundle` parameter, which is used to control what A-GPS data should be removed. The `Bundle` can contain `Boolean` values with keys to indicate which data to remove. The keys can be any of the following strings:

- `ephemeris`
- `almanac`
- `position`
- `time`
- `iono`
- `utc`
- `health`
- `svdir`
- `scsteer`
- `sadata`
- `rti`
- `cellldb-info`
- `all`

Passing a null for the `Bundle` causes all the A-GPS data to be removed.

The `force_time_injection` command causes the current time to be retrieved from the configured NTP server and updated for the purposes of GPS calculations.

The `force_extra_injection` command causes the A-GPS data to be downloaded from one of the configured servers and used by the GPS location provider.

Network Provider

In Android, network-based location can use different methods for determining the location of a device. As of this writing, the network location provider can provide location information using cell towers, or based on wireless network information.

Using Wireless Network Access Points

Providing location information based on wireless network access points is one of the ways that Android supports location resolution with the network provider. Although it does require that the Wi-Fi radio is active, the Wi-Fi radio often consumes less battery power than the GPS hardware.

How It Works

Wi-Fi-based location detection works by having a device track what Wi-Fi access points it can detect and the current signal strength of those access points. The device then makes a query to the Google location service (which is different from the Android location service), which provides location data based on the Wi-Fi information. The Wi-Fi information collected by the device includes the mandatory access control (MAC) addresses of the Wi-Fi access points that are in range and the strength of the signal being received from those access points.

To provide location information based on visible Wi-Fi access points, the Google location service must obtain information about Wi-Fi access points and their locations. This information is collected by Android devices when a user enables use of Google's location service in the Location Settings screen. Figure 1-2 shows the confirmation screen that is presented to a user when enabling the Google location service as a source of location data.

Pressing Agree on this screen allows the device to record Wi-Fi information as well as current location information (possibly provided by GPS) and transmit this information to Google. This essentially allows Google use each and every Android device as a way to update the Wi-Fi location information and constantly maintain up-to-date data.

One of the main benefits of the Wi-Fi location source is that it allows devices to acquire location information in areas where GPS cannot provide location data. As stated in the previous section, GPS is problematic when used indoors or even in an urban environment where tall buildings can cause signal problems. In contrast, an urban environment may increase the accuracy of Wi-Fi-based location because of the abundance of Wi-Fi networks available to determine a device's current location.

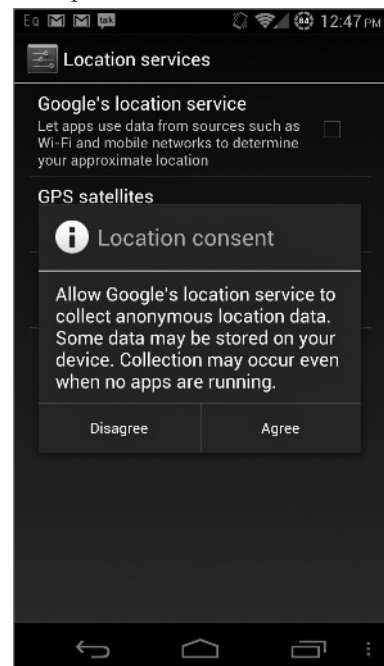


FIGURE 1-2: Confirmation screen displayed when enabling Google's location service

Limitations

As with GPS, using Wi-Fi networks as a source of location information does have its limitations. First, to determine the location, Wi-Fi networks must be in range. Additionally, the networks must have a publicly broadcasted service set identifier (SSID) that has not been configured to be ignored by Android. Access points that have an SSID that ends in `_nomap` will not have their information sent to the Google location service.

Additionally, changes to the location of Wi-Fi access points can cause inaccuracies in the location data that is produced. For example, many people now have wireless networks in their homes for daily use. Assuming an Android device has been configured to use the Google location service, Android would have sent the access point MAC address and location to the Google location service. If the user were to change the location of that access point (take it to a vacation home, for example), the location service might determine the device to be in the wrong location when the Wi-Fi location source is used.

Although the location service does allow for access point location to be updated via Android devices, Google does not allow users to explicitly set the location of an access point. An Android device will push the information to the location service, which may wait until other devices can confirm the change before the location service is updated.

Using Cell IDs

In addition to using Wi-Fi information to determine device location, Android can also use the cellular network. The cellular network is used in a similar way as Wi-Fi access points to determine device location.

How It Works

To function properly, a cellular device must be in contact with a cell tower. As a device moves, it may connect to a different cell tower as the signal strength of an approaching cell tower becomes stronger. Knowing the unique ID of the tower that a device is currently connected to, and possibly the towers that a device was previously connected to, can provide insight to where the device is located assuming the location of a given cell tower is known.

Android and the Google location service work together to map cell tower IDs to location data in a way that is similar to Wi-Fi data. Once a device has been configured to use the network provider, it collects data on the current cell tower ID in addition to the visible wireless networks. For cell towers, this data includes the cell tower the device is currently connected to and the device's current GPS location. With this information, the Google location service can develop a “map” of cell towers that includes their locations.

By again allowing Android devices to update cell ID information, the Google location service can maintain a constantly updated store of information that increases in accuracy as the number of entries increases.

When a device needs to find its current location, it sends the ID of the cell tower it is currently connected to, as well as historic information about past cell towers it has used, to the Google location service. With this information, the Google location service can provide information about the

device's current location based on the data it has about the cell tower network. If the IDs of multiple cell towers are sent to the Google location service, it can use triangulation to provide increased location accuracy. The Google location service cannot do this if the device submits only a single cell tower ID.

Limitations

The limitations for using cell tower IDs are similar to the limitations that exist when using Wi-Fi networks to determine location. However, because the location of cell towers is less likely to change than the location of wireless access points, some of the complications that may exist when using Wi-Fi access points are removed.

However, just like Wi-Fi access point data, the Google location service must have data on the cell tower IDs that are sent by a device in order to provide location data.

SUMMARY

This chapter presented an overview of how the location providers available in Android work and discussed their limitations. The decision of which location provider to use in which situation can be a complex topic, and is discussed at length in the following chapters.