Four Cambridge Center, Cambridge, MA 02142 • Tel: 866-833-4069 • www.thingmagic.com 4/12

Basic GEN2 Protocol Settings Demystified

Dan Ratner

The EPCglobal Gen2 protocol was designed to accommodate a wide variety of applications, so nearly every aspect of its operation can be adjusted to optimize its behavior for a particular set of requirements. Out of the box, ThingMagic readers are configured to perform well in average circumstances, and many users never have to touch the Gen2 settings. For those that must, the information provided here will make optimization of the reader's performance a simple step-by-step process.

Introduction

It is common for new users of EPCglobal Gen2 technology to be surprised that they cannot immediately achieve the performance promised in reader and tag vendors' literature. Although some discrepancy can be due to the vendor having measured their performance under carefully controlled conditions, a much more likely cause is non-optimal reader settings, particularly those relating to the EPCglobal Gen2 ("Gen2") protocol. A vendor's results can nearly always be duplicated by users who are aware of the Gen2 setting options and can tune them to optimize overall system performance. This application note will explain the purpose and impact of each of the relevant Gen2 settings, followed by a step-by-step process for optimizing them for best results.

To understand the relevance of the Gen2 settings, it is necessary to have a basic understanding of how UHF RFID works. At its most basic level, the reader powers up the tag at the same time it is communicating with it. A good analogy for this kind of system would be if two people are trying to send messages from one mountaintop to another at night, where one has only a flashlight (representing the reader) and the other has a mirror (representing the tag). The person with the flashlight can send messages by turning the flashlight on and off, but when he wants a response, he must keep the flashlight on so his partner can signal back with the mirror. Layer on top of that the possibility that there may be many people on other mountaintops, also with mirrors, ready to send messages back to the "reader", and you have an idea of the communication challenge the Gen2 protocol was designed to handle.

The settings which govern the tag's behavior are all controlled by the reader via messages it sends when it communicates with one or more tags. These settings do not remain in effect long. Some last for a single exchange with the reader, some until all tags present have responded, some until the tag powers down, and some for a fixed period of time. After a few minutes, however, the Gen2 protocol expects all tags to fall back into their default state so a second reader would automatically know what initial state the tags are in when it encounters them.

There are quite a few Gen2 settings, but they can be grouped into 3 categories:

1. Settings that control how the reader communicates to the tags
2. Settings that control how the tag communicates back to the reader
3. Settings that control when tags respond relative to each other to avoid communication collisions

Also, as you learn more about the Gen2 options, keep in mind which aspects of RFID performance you are looking to optimize for your application. Typical choices are:

1. Maximum read distance
2. Minimum time to read every tag in the field at least once
3. Minimum time to read a single tag traveling through the field
4. Maximum number of responses from any tag in the shortest amount of time

The last case, 4, is rarely required in practical applications, but it is mentioned along with the others because it is often the first thing users try to do when they test a reader in a lab environment. Your goals may also include minimizing negative influences on performance, such as reducing reader-to-reader interference when many readers are located in close proximity. Gen2 settings can help there, too.

Reader-to-Tag Settings

Although not directly related to Gen2 settings, the role of transmit power is so central to successful RFID applications, a few words must be said about it. In order to have successful communications between a reader and a tag, the tag must be able to be powered-up by the reader's signal. No communication at all occurs if the tag remains dormant and does not respond to the reader. To achieve maximum performance, always configure the reader to transmit at the highest permitted power into an antenna of highest permitted gain, allowed by local regulations. (In regions governed by the FCC regulations, this would be a +30 dBm power level into an antenna with a maximum gain of 9 dBiC, if circularly polarized, and 6dBiL if linearly polarized.)

There are two primary Gen2 settings that control reader-to-tag communications, "Tari" and "link rate". In our flashlight and mirror analogy, if the communication "protocol" was Morse code, the "Tari" would control the length of a dot (or dash) and the link rate would control how quickly the dots and dashes are sent. Just as with Morse code, the maximum speed at which words can be sent is limited by the length of the dots and dashes . ThingMagic adjusts both of these Gen2 settings together, so as the user selects a smaller Tari, the link rate is automatically increased. Tari values offered are 6.25 usec, 12.5 usec, and

25 usec, which are automatically paired with link rates of 40 kbps, 80 kbps and 160 kbps by the ThingMagic reader.

Factors to consider when selecting the Tari value (and therefore the link rate) are:

1. Will I be writing data to the tag often? The Gen2 protocol is designed to minimize communication from the reader to the tag (during an inventory round, the messages sent are very short.) If obtaining the EPC of the tag is essentially all that is required, then the Tari/link-rate can be optimized for other factors – they won't affect performance that much. If tags are being written to, then decreasing the Tari (to increase the link rate) could have noticeable benefits.

2. Are there many other readers in the area? Without going into too much explanation, faster communication on a channel can cause interference with adjacent channels if the channels are closely spaced. If there are other readers in the area and performance of the first reader seems to degrade as more readers are turned on, then the system might be experiencing adjacent channel interference between readers. This is rarely seen when readers are operated in the North American (FCC) region as it contains 50 well-spaced channels, but in the EU region where there are only 4 channels, this could become an issue. In this case, you would want to try increasing the Tari to decrease the link rate and see if the performance of the first reader improves.

Tag-to-Reader Settings

The tag uses a slightly different signaling method to communicate back to the reader than the reader uses to communicate with the tag. There are two settings that you can use to control the tag-to-reader communication method, the link frequency (called "Backscatter Link Frequency" to clearly distinguish it from the reader's link rate) and "M" value. Both the Backscatter Link Frequency (often abbreviated "BLF") and "M" value modify the way the tag communicates back to the reader. The BLF is the raw signaling rate. Data rates supported by ThingMagic readers are 250 kHz and 640 kHz. The M value essentially controls how many times a symbol is repeated. An M of 2 means that each symbol is repeated twice. M values of 1 (FM0), 2, 4, and 8 are supported by ThingMagic readers. When there is no repetition (M=1) this mode is referred to as "FM0" because of other slight distinctions not relevant to this discussion. The Gen2 protocol provides this option to repeat symbols to maximize the chances that the reader can decode a very weak signal from the tag. Just as it is easier to understand someone who is whispering in a noisy room if they repeat everything they say several times, so too, the RFID reader will decode a weak signal more reliably if "M" is 2 or greater.

ThingMagic readers currently support a BLF of 640 kHz together with an "M" value of FM0 to achieve the highest tag-to-reader data rate, around 400 tags per second. Alternatively, a BLF of 250 kHz can be combined with "M" options of FM0, 2, 4, or 8. At BLF=250 kHz and M=8 the tag read rate drops to around 100 tags per second, but the increased sensitivity could nearly double the read distance compared to settings of 640 kHz/FM0 for sensitive battery-assisted tags. There is one additional tradeoff to consider when selecting theses values. Just as higher data rates from reader-to-tag communication increased the likelihood of reader-to-reader interference when many readers are present, so too can higher rates from tag-to-reader cause unwanted adjacent channel interference. How to balance these concerns will be addressed in the section on optimization.

Tag Contention Settings

The remaining Gen2 settings control how soon, and how often, a tag responds to a reader. If all tags in the field responded immediately to the reader, their responses would collide and the reader would rarely be able to decode responses from any tags. To avoid this, the Gen2 protocol provides a variable number of time slots in which tags can respond. The reader announces to the tag population the number of response opportunities it will give to them (calculated from a value it sends, called "Q"). Each tag randomly picks a number within this range and the reader starts announcing response slots. Essentially the reader announces the start of a round and then repeatedly shouts "next". Each tag keeps count of how many "nexts" there has been and responds when their selected slot comes around. The "Q" value that is announced is a number from 0 to 15. The number of slots is 2^Q power. The number of slots increases rapidly with the "Q" value: 1, 2, 4, 8, 16, etc.

If "Q" is too large, many response opportunities will go unanswered and the inventory round will take longer than it should. If "Q" is too small, one or more tags will be more likely to respond in the same time slot, resulting in several inventory rounds having to be run before all tags respond. In the extreme case, where there are many more tags in the field than slots available, response collisions will always occur and the inventory round will never complete successfully. The ideal number of slots has been determined to be around 1.5 times the number of tags in the field. If the tag population will be variable, the user can set the ThingMagic reader to automatically determine the Q based on the results of successive rounds (it increases the "Q" if there are many collisions, and decreases the "Q" if there are too many unanswered response slots.)

There are two additional Gen2 settings which control whether, and how often, a tag participates in an inventory round. When initially setting up the response queue, the reader also sends two other pieces of information to all tags: (1) how long they should delay until they re-respond (called the "session" value) and (2) what state ("A" or "B") they should be in to participate in the inventory round. The "A" state indicates that the tag has not yet responded to an inventory round. The "B" state indicates that it has. The session setting determines how long the tags wait in the "B" state before returning to the "A" state so they can participate again.

The "Session" setting, which controls how often tags respond to inventory rounds, has 4 options, but two behave identically.

Session "0": Prepare to respond again as soon as RF power drops
Session "1": Prepare to respond again between 0.5 and 5 seconds after first response
Session "2" or "3": Do not respond again for at least 2 seconds

Selection of the session setting is nearly always based on the expected number of tags in the field. You usually want all tags in the field to have had a chance to respond once before the first tags start responding again. For ThingMagic's default settings, the tag read rate is around 200 tags per second. With these settings, you would want to use session 0 up to around 100 tags in the field, and session 1 up to around 400 tags, and session "2" (or "3") for more than 400 tags.

The ThingMagic "target" setting is only important if you want to force tags to re-respond more often than they otherwise would. There are two relevant choices, "A" and "A then B". "A" means that the reader always looks for tags that are in the "A" state. Tags held in their "B" state by their session timer are ignored. "A then B" tells the reader to read all the tags in the "A" state, then read all the tags in the "B" state, and keep repeating this process. Tags in their "B" state will respond to a "B" query and return immediately to their "A" state, regardless if there was time left on their session timer. This accelerated read rate is only useful in two cases:

(1) The application needs to read the same tag multiple times in order to determine more than the tag's identity, for example, if an attempt is being made to determine whether the tag is moving toward or away from the antenna by monitoring the tag's returning signal level.

(2) The user is attempting to estimate the reader's performance in the presence of many tags by reading fewer tags over and over.

Optimizing Gen2 settings

If the previous information on Gen2 settings is new to you, the task of optimizing all these settings together will likely appear daunting to you at this point. Fortunately, we have created a step-by-step procedure for achieving the desired results.

(1) Set the reader up to achieve the maximum read distance possible (Tari=6.25 usec, BLF=250 kHz, M=8) at the anticipated transmit level (usually the highest allowed by regulations unless the read-field size must be limited). Place tags at the maximum read distance required for your application and ensure that all tags are read reliably.

(2) Set up your "Q" and session value for the expected tag population (2^Q should be greater than 1.5 times the number of tags). If the number of tags will vary widely, use the ThingMagic "automatic" setting.

(3) Set the session value. The maximum-distance BLF and M settings will read around 100 tags per second, so the session should be set to "0" if you anticipate up to 50 tags, "1" for up to 200 tags, and "2" (or "3") if you anticipate there being more tags in the field.

(4) Now determine if these settings meet your requirements for tag read rate (or the ability to read a tag moving through the field). If the read rate (or single-tag read time) is too slow, start decreasing the "M" value (keeping the BLF at 250 kHz) until the desired read rate is achieved. (If FM0 and 250 kHz is insufficient, try FM0 and 640 kHz.) Now check your system performance at the required distance again to make sure your tags can still be read reliably. If it is still OK, keep these settings. If not, you will have to decide on the best compromise values for M and BLF that balance read-distance with read-rate.

(5) If the application calls for it, add more readers with the same settings. Check the performance of the first reader while the other readers are operating. If it has not changed, leave its settings alone. If its performance has decreased, then reader-to-reader interference is probably occurring. First increase the Tari on all readers to see if that helps (this change will impact the other performance factors least.) If that does not work, try decreasing the BLF and increasing the M value to see if that eliminates the problem. If so, you must select the BLF and M values which give you an acceptable level of reduced reader-to-reader interference, read rate, and read distance.

(6) If your application requires a few tags to respond repeatedly, use session="0" and leave the target="A" setting alone. If there are enough tags in the field to require using session "1" or "2", but you need the tags to re-respond more rapidly than the session timers allow, use the target="A then B" setting.

At this point, your Gen2 settings will be optimized and your system will be operating at peak performance. If you are still experiencing difficulties, additional resources are available from ThingMagic. For more information, visit www.thingmagic.com.