March 20, 2019

Managed adaptive spectrum sharing for BLE

Brian Cornell

Managed adaptive spectrum sharing for BLE

ABSTRACT

Many Bluetooth Low Energy (BLE) advertising devices operate primarily in a transmit-only mode. Prior to start of transmissions, such devices do not test for the presence of other BLE devices that might currently be occupying the spectrum. To minimize collisions, such BLE transmitters add randomness to transmit timing. This technique fails when too many advertisers are in a confined area, as the chance of collisions goes up with the number of advertisers. This disclosure describes managed adaptive spectrum sharing for BLE (MASS BLE), a technique that uses time division multiple access to enable co-located BLE advertisers to share spectrum with graceful degradation of functionality.

KEYWORDS

- Bluetooth Low Energy (BLE)
- Time division multiple access (TDMA)
- BLE advertising
- Spectrum sharing
- Adaptive re-assignment
- Window randomization

BACKGROUND

Many Bluetooth Low Energy (BLE) advertising devices operate primarily in a transmit-only mode. Prior to start of transmissions, they do not test for the presence of other BLE devices that might currently be occupying the spectrum. To minimize collisions, such BLE transmitters send little data and randomly jitter their transmit timing. This technique fails when too many advertisers are in a confined area, as the chance of collisions goes up with the number of advertisers. As collisions go up, the visibility of advertisers decreases. For example, with completely random advertisements within every 200 milliseconds window, simulations show that only about 36 devices can be sharing a channel before there is a 1% chance of missing at least one device each second.

DESCRIPTION

This disclosure describes managed adaptive spectrum sharing for BLE (MASS BLE), a technique that uses time division multiple access to enable many co-located BLE advertisers to share spectrum with graceful degradation of functionality. Some of the features of the techniques are as follows.

- Time division multiplexing is used to share a channel.
- Multiple channels are used to spread load.
- Variable length windows are used to gracefully scale the number of spectrum users.
- A precisely-timed management signal is used to maintain low power characteristics.
- Adaptive reassignment enables period downsizing.
- Window randomization prevents consistent conflict with non-participating devices.

Steady State Operation

This section describes how the system works in steady state with devices already participating in the scheme.

Basic Operation

The MASS BLE system involves a single management system and any number of advertising members. The devices dedicate one or more BLE channels to the MASS BLE system. On each of these channels, time is divided into windows which are further divided into slices:

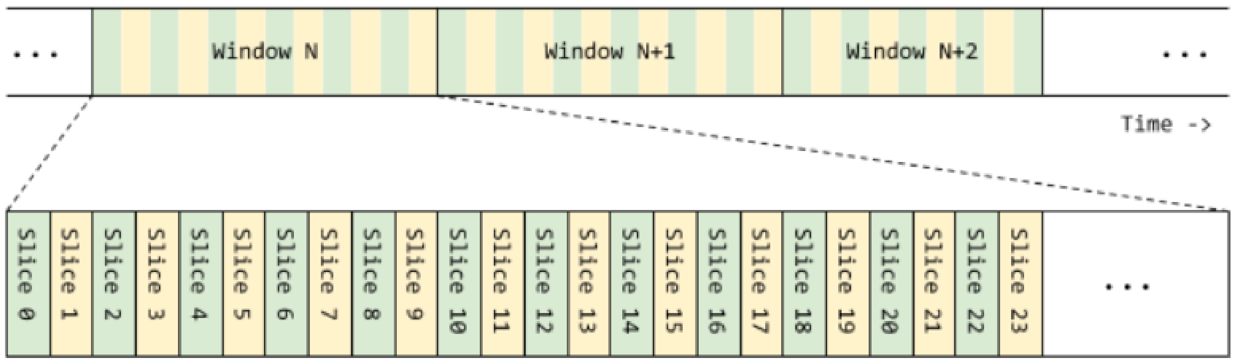

Each slice is long enough to allow a single BLE advertising packet with margin on both sides. Each window is long enough for every member using that channel to get a slice with some extra slices allocated to the management system.

At the beginning of each window, a small number of slices are allocated to the management system. These slices are used to broadcast information to every member including

- the current slice index;
- the number of slices before the start of the next window;
- the number of slices in each of a small number of following windows;
- a small set of data for allocation or reallocation (discussed more below); etc.

A few slices after the management slices are allocated to device onboarding. The rest of the slices are either free or allocated to a single member for an advertising packet.

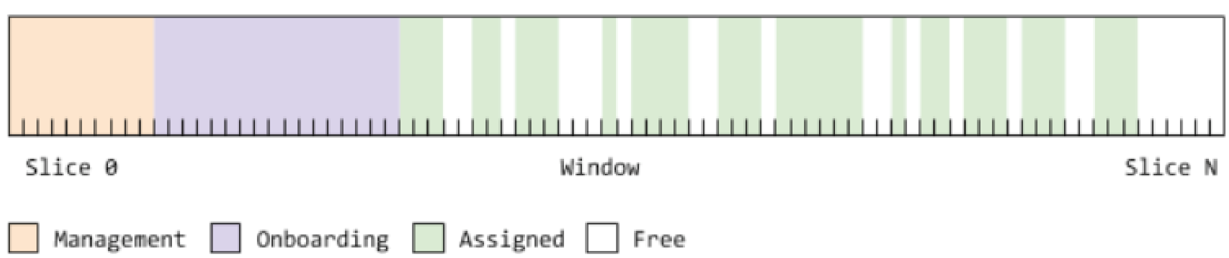

Each member can use the management slices in order to establish the timing of each slice within the window as well as when the next management slices will be sent. This allows the member to then sleep until its assigned slice when it can wake up and transmit a single packet.

The member can then sleep again until the next management slices will be transmitted. As a result, each member is able to sleep the majority of the time and conserve battery power.

The main operation of the management system follows this process:

1. Transmit management packet with current configuration.
2. Repeat the management packet with increasing slice index N_M times in order to increase the chance of each member seeing it.
3. Listen for any new members to onboard.
4. Receive packets from members.
5. Perform bookkeeping to compute the next windows and adapt as necessary.

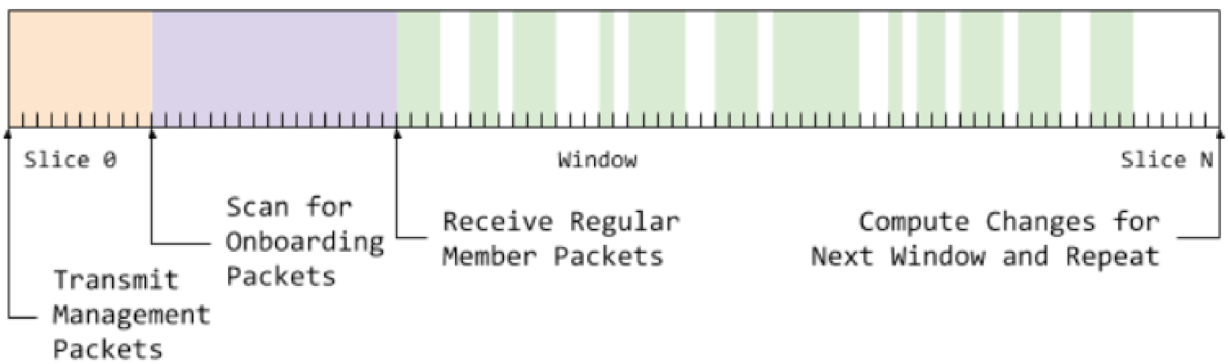

The main operation of each member follows this process:

1. Listen for up to N_M slice lengths to receive the management packet.
2. Perform any necessary updates based on the management packet.
3. Sleep until assigned slice S.
4. Wake up and transmit a single packet.
5. Sleep until beginning of next window.

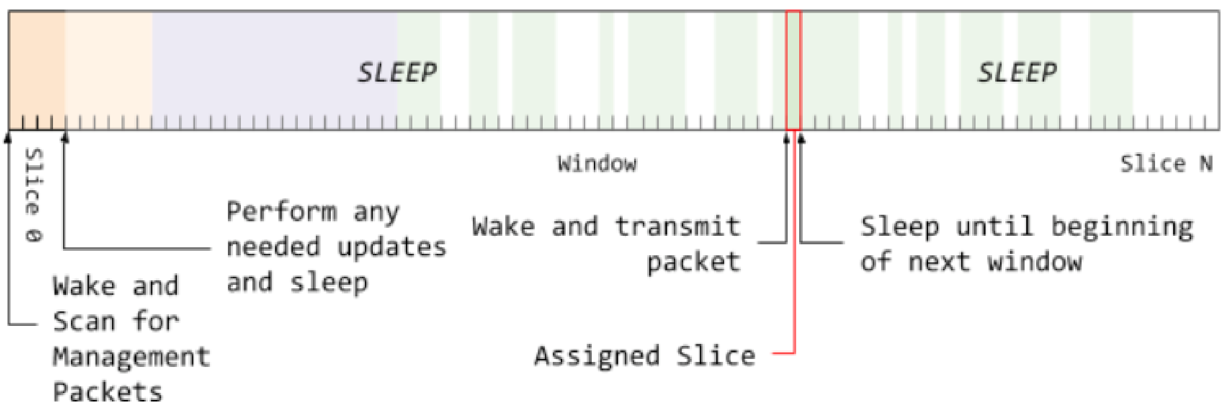

Graceful Degradation

With the MASS BLE system, advertising frequency of each member is controlled by the window length assigned by the management system. If the windows average 200 milliseconds in length, the members will advertise at an average of 5 Hz.

Because of this property, the management system can gracefully degrade the system for all members by increasing the window size in order to fit new members. This degradation allows members to fully utilize the dedicated channels at the maximum possible transmit rate that would support every member.

As an example, if the system needed to fit 450 members on a single channel it may increase the window size to 500 slices with 10 allocated to management, 20 to onboarding, and 20 free for new members. If each slice were 0.5 milliseconds in length this would decrease each member to advertising at 4 Hz, but the channel would remain over 90% utilized.

Just as the system can gracefully degrade, it can also recover as members leave. This may require dynamic reallocation as described below in order to move members to earlier slices in the window. Once the end of the window is unused, it can be removed and the transmit rate increases.

The graceful degradation is shown in the following diagram.

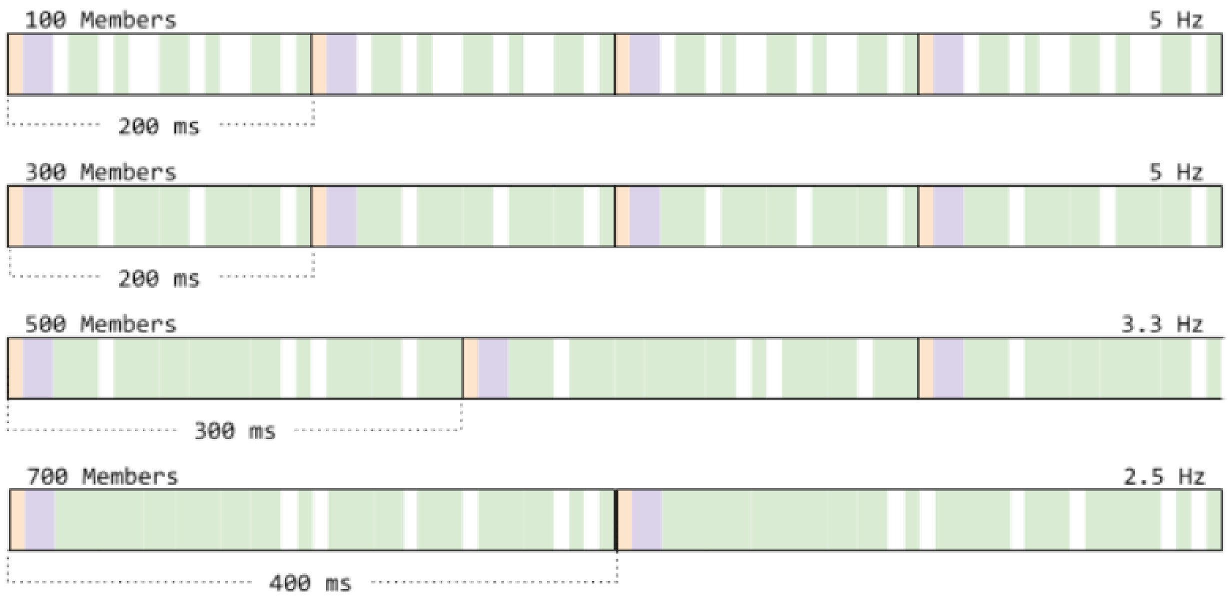

Randomized Windows

While devices using MASS BLE may share the spectrum well, other BLE devices can use the same channels. This introduces the chance of collision with an external device. Since devices often transmit on a fairly regular cadence, regularly sized windows would have the potential to lopsidedly disadvantage certain members who are more likely to collide with an external device. Because of this, randomness is added to the window length to position members at different places relative to a regular interval.

This randomization involves adding a random number of unused slices at the end of the window within a certain percentage. This has the effect of moving every member by that many slices, decreasing the chances of repetitive collisions. For example, if a window is 500 slices, the management system may randomly add between 0 and 50 slices before the next window starts.

Because of this randomness, the management system must communicate the slice count for multiple windows in case a member misses the management packets for one.

Management Collision Avoidance

While missing an advertisement from a member due to collision may not be a problem, missing the management signal could impact the entire window or cause a member to drop. As such, it is important to attempt to avoid collision for the management system. The main mechanism for this is through repetition of the management signal through multiple slices.

The management system uses a small number of slices (N_M) to transmit the management signal repeatedly. Each slice transmits the same data except that the slice index increases to let members know when the actual window started. N_M can be tuned, but expected values are around 5-10 packets.

Device Onboarding

While the above section discusses what happens in steady state, members must join MASS BLE in order to participate in that operation. This section discusses the process of joining the system, also known as onboarding a member.

Discovery

The first part of device onboarding is that the device needs to determine that the MASS BLE system is present and operational. In addition, this must be done in a power efficient way since devices need discovery at all times, at least when they are in motion.

System Detection

Discovery is done in two phases for power optimization. The first phase is system detection and the second is management synchronization.

The system detection phase is the lowest power and is simply designed to determine whether MASS BLE is in use in the surrounding area. In order to make this determination, the device listens briefly on one of the dedicated MASS BLE channels. The device only needs to listen for a small number of slices such as 10 or 20. One of three things can happen then

1. Device detects a management signal during at least one slice:
    a. Immediately proceed to management synchronization and skip management scan.
2. Device detects that a significant percentage of the slices are used, such as over 50%:
    a. System is likely in use on this channel.
    b. Proceed to management synchronization with a management scan.
3. Device detects few or no advertisements:
    a. System is not active, go to sleep and check again later.

If the MASS BLE system is not heavily loaded and there are many free slices, the management system should insert extra management packets in free slices such that the detection duration will always see at least one management packet or heavy usage.

In order to perform this process in a power efficient manner, the detection should only happen when the device is in motion and even then infrequently, such as every 10 seconds. For this reason, the management signal must be present in a slightly wider area than where MASS BLE is needed or risk devices entering the area and not detecting MASS BLE for up to the detection interval.

Management Synchronization

Once the device determines that MASS BLE is likely in use in the area, it must synchronize with the management system in order to establish timing and participate properly in the system.

If a management signal was not detected during system detection, synchronization begins by listening for a longer period on the MASS BLE channel, such as for half a second. The management system can ensure that such a duration would include management signals even with larger windows by reserving some slices in the middle of the window for additional management signals. If no management signal is detected during this scan, the device can assume that MASS BLE is not actually active and return to system detection.

Once a management signal has been identified either during system detection or the above scan, the device can use the information in that packet to identify the timing of the following windows. The device can then proceed to issue a slice request as below.

The following diagram shows the detection and synchronization process:

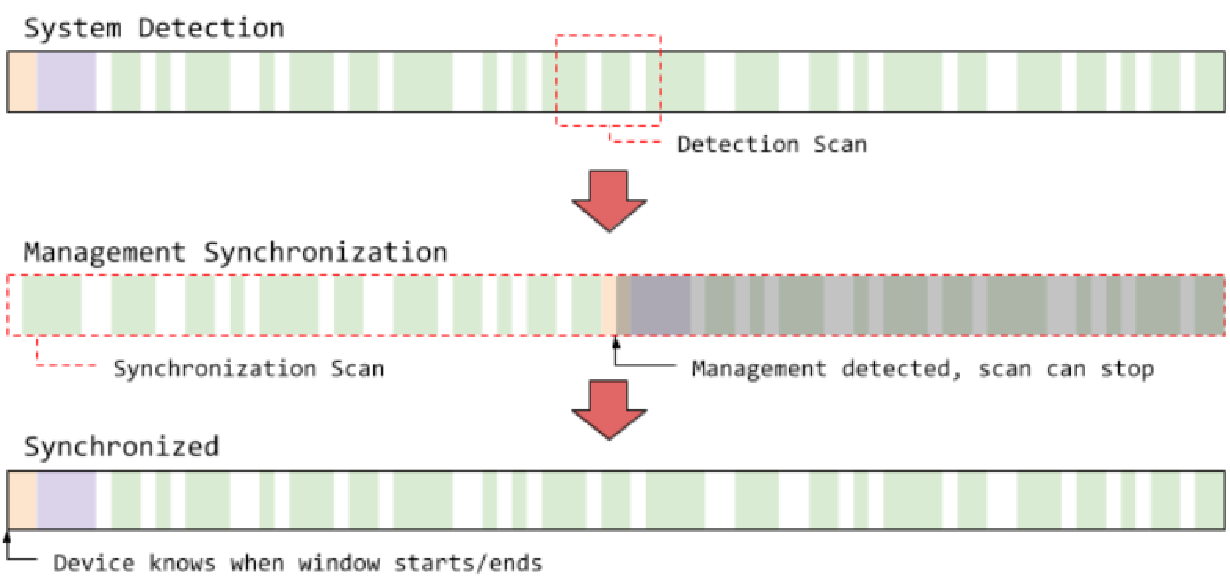

<u>*Slice Request*</u>

The slice request process involves a request from the member, allocation from the system, activation of the member, and reclamation of missed allocations.

Request from Member

As mentioned in the steady state operation section, a number of slices after the management packets are dedicated to member onboarding. The new member selects one of these slices at random and uses it to transmit a packet to request allocation from the system. This allocation is a special packet including a unique ID for the member. The ID can be completely ephemeral and randomly generated as it is only used for the device to recognize a response or reallocation.

The slice used is randomized to reduce the chance of two potential members selecting the same slice and colliding. In the case of a collision, the member will not receive a response and will select a new random onboarding slice in the next window for its request.

Allocation of a Slice

Once the management system detects the onboarding request for a member, it allocates a slice to the member in its internal accounting. This may necessitate increasing the window size as mentioned in the graceful degradation section of steady state operation. This slice will belong to that member until reallocation or device removal as described in later sections.

Once the system has assigned a slice, it advertises that assignment in all management packets for the next window that has space. This data includes the ID given by the device, the channel to use in case of multiple MASS BLE channels, and the slice allocated to the device. This data will be transmitted until the member is detected using its assigned slice. If there are more new member allocations to advertise than fit in the management packet, advertisements may alternate in subsequent packets or windows to accommodate the load.

There is a rare special case where the ID sent by the new member is already in use in the system. When this happens, the management packet sends a special allocation response that tells the new member to select a new ID and try again.

Activation of the Member

Once the member detects the allocation response in the management packet, one of three things may happen:

1. The response is an allocation for the same channel that the device is currently synchronized to:
    a. The device considers itself active and transmits in that slice every window until further notice.
2. The response is an allocation for a different channel:
    a. The device performs a management synchronization scan as described in discovery on the new channel.
    b. Once synchronized to the management signal on that channel, the device considers itself active and transmits in its assigned slot every window until further notice.
3. The response indicates an ID collision:
    a. The device selects a new ID and transmits a new request.

If the member does not detect a response, it continues to send onboarding requests until it does. Some windows may be skipped potentially with exponential backoff to handle the potential load of many new members at the same time.

Reclamation of Missed Allocations

Even if a device misses an allocation response, it will generally be received in one of the following windows when it is repeated. There is however a case where a device sends the onboarding request but then leaves the area before it can become active. In this case, the management system would be left advertising the new allocation indefinitely. The system would slowly degrade as more and more management packet space was given to orphaned allocations.

Because of this, the management system must give up on an allocation and reclaim the slice if a device is not detected as active after a timeout. This process involves first ceasing the advertisement of the allocation and waiting for another timeout to ensure the device isn't in the process of activating. If the device is still not detected, then the slice can be deallocated and returned to the free pool. The management system may also advertise the deallocation in the management packets for a short duration.

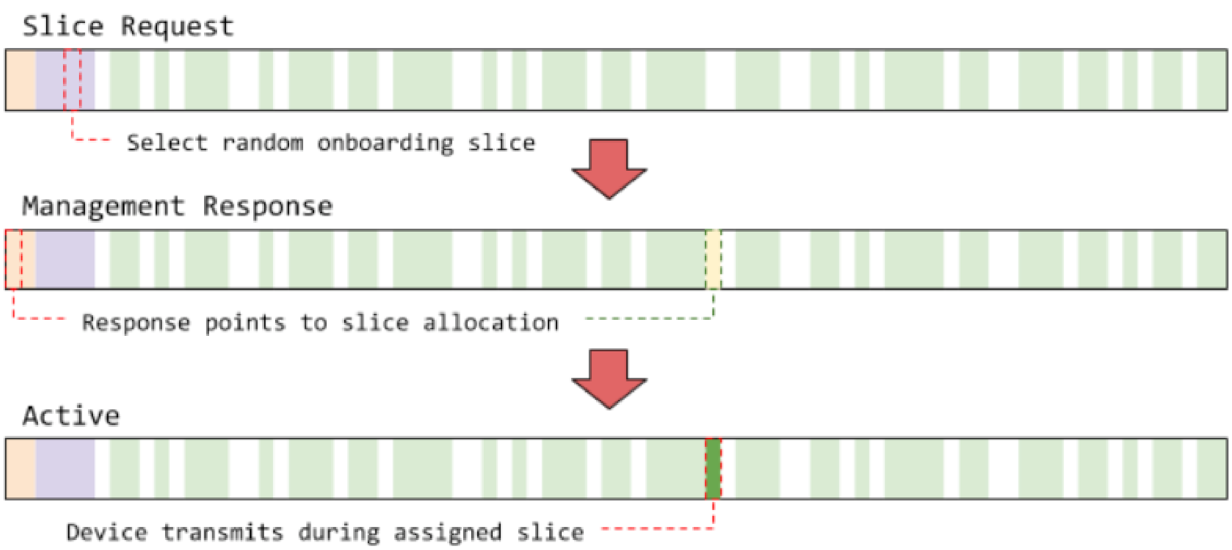

Device Removal

There are three components to device removal: the management system freeing slices no longer in use, members stopping operation, and removing devices using slices not currently assigned to them.

Management Slice Deallocations

In order to free slices no longer used by members, the management system should monitor slice usage. If a slice goes unused for many windows in a row, the management system must assume that the member has either stopped participating or left the area of MASS BLE. The number of windows can be tuned, but somewhere around 10-20 is likely to be reasonable.

After the requisite number of windows have had no transmission during an assigned slice, the management system should deallocate the slice. The slice should first be put into a pending state where the deallocation is advertised in management packets for a few windows. This notifies a member that is actually in the area or has returned to the area that it has lost its slice and must onboard again to participate.

Once the deallocation has been advertised for a number of slices, the management system can consider the removal complete. At that point, the slice can be freed for assignment to a different member.

Member Cessation

A member in the MASS BLE system stops participating if it does not see the management packets for a series of contiguous windows. After a few windows, the member will no longer know the window size and will lose synchronization. At this point, the member should scan for a management signal as in the management synchronization step of device onboarding.

After attempting resynchronization, one of two things may happen:

1. The member does not see any management signal and is unable to synchronize.
    a. The member should consider itself no longer in a MASS BLE area and should return to occasional system detection scans.
2. The member is able to synchronize to the management signal again.
    a. The member should check the management signal to see if it has been deallocated:

i. If the member has been deallocated, it should select a new ID and request a new slice.

ii. If the member has not been deallocated, it should resume transmitting in its previously assigned slice.

Orphan Removal

There may be cases where a member thinks it is part of the MASS BLE system, but the management system is not tracking it. This may happen for example if the management system removed a device but the device did not realize. This will result in the device either transmitting in a free slice or colliding with another member.

When the management system detects either consistent collisions in a particular slice or consistent use of an unassigned slice, it must attempt to resolve the issue. Since it is not an assigned usage, the system cannot use an ID to remove the member. Instead it must use a slice based removal. The management signal should support a removal indication that provides a slice instead of an ID. If an assigned member is also using this slice, the management system should consider that member removed.

When any device sees the slice it is using advertised for slice based removal, it should immediately stop using the slice and proceed to request a new slice with a new ID using the process in device onboarding.

If the management system does not see the slice freed up after the slice based removal, there may be a non-compliant device using the slice. The management system should consider the slice unavailable and not use or assign it. The non-compliant device and transmissions should be logged for further investigation.

Dynamic Reallocation

As the usage of MASS BLE in an area decreases from peak, the slices in use will become fragmented with many gaps like in the below diagram. The system cannot decrease the window size and thereby increase advertising frequency without consolidating the remaining members closer to the start of the window. Two alternatives for how the management system can perform this consolidation are provided below.

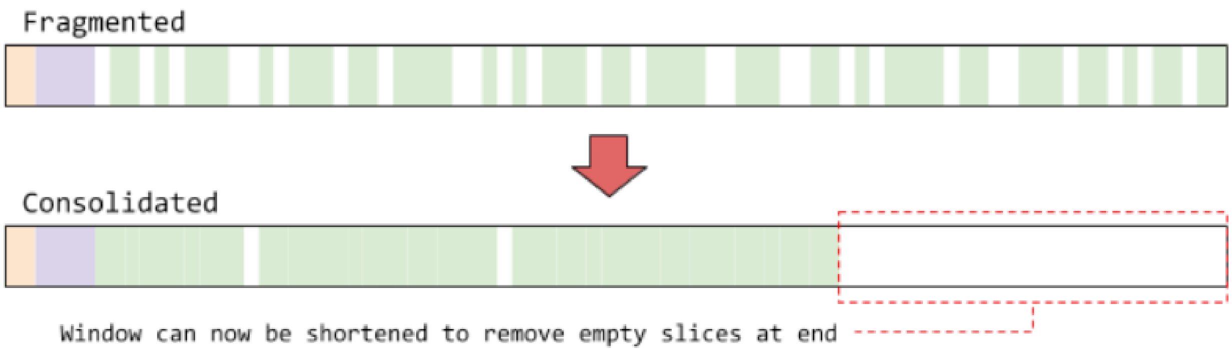

Member Defragmentation

Member defragmentation is a process whereby the management system notifies a member that it has been assigned a new slice and then monitors for the member's transition between slices. The following process is used:

1. Management system selects a new slice S for member M.
2. Management system advertises new allocation M→S in management packets.
3. Management system begins monitoring slice S usage.
4. Member sees M→S reallocation in management packets.
5. Member immediately switches to slice S.
6. Management system sees usage of S, removes previous allocation for M.

Using this process, the member can be moved from one slice to another with a brief interval where two slices are allocated to it. This process can be repeated for each member that needs to be moved until the window is compacted. This is shown in the below diagram.

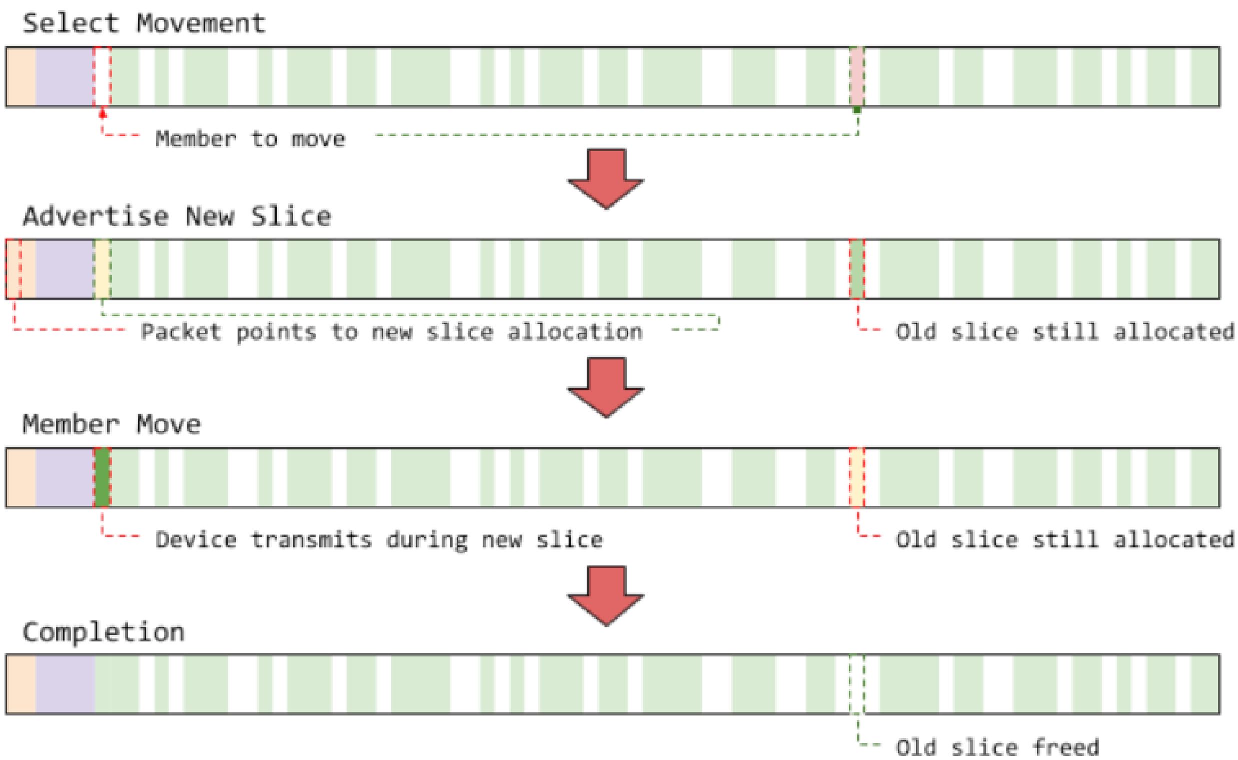

Active Pruning

Active pruning is a process whereby the management system intentionally shrinks the window size such that a slice belonging to an active member is no longer present. The member must detect that it has been pruned and request a new slice using the slice request process of onboarding.

If this process is used, the member must monitor the window size advertised in the management packets. If the window size does not fit the member's allocated slice, the member should consider itself pruned. At that point, the member should send out a slice request using the same ID as it previously used. The management system will see the slice request and assign a new slice to the member.

When the management system sees a new slice request for the pruned ID, it can consider the old slice vacated. If it does not see a new slice request, it should follow the device removal process to clear the slice. Otherwise, increasing the window again could lead to the member continuing its usage of the slice if the member did not notice the pruning.

The advantage of this approach is that the window compacting can happen sooner. The disadvantage is that members that are pruned may be unable to transmit for a few windows while they await a new slice and there may be additional load to the onboarding slices for these requests.

This diagram illustrates the active pruning process:

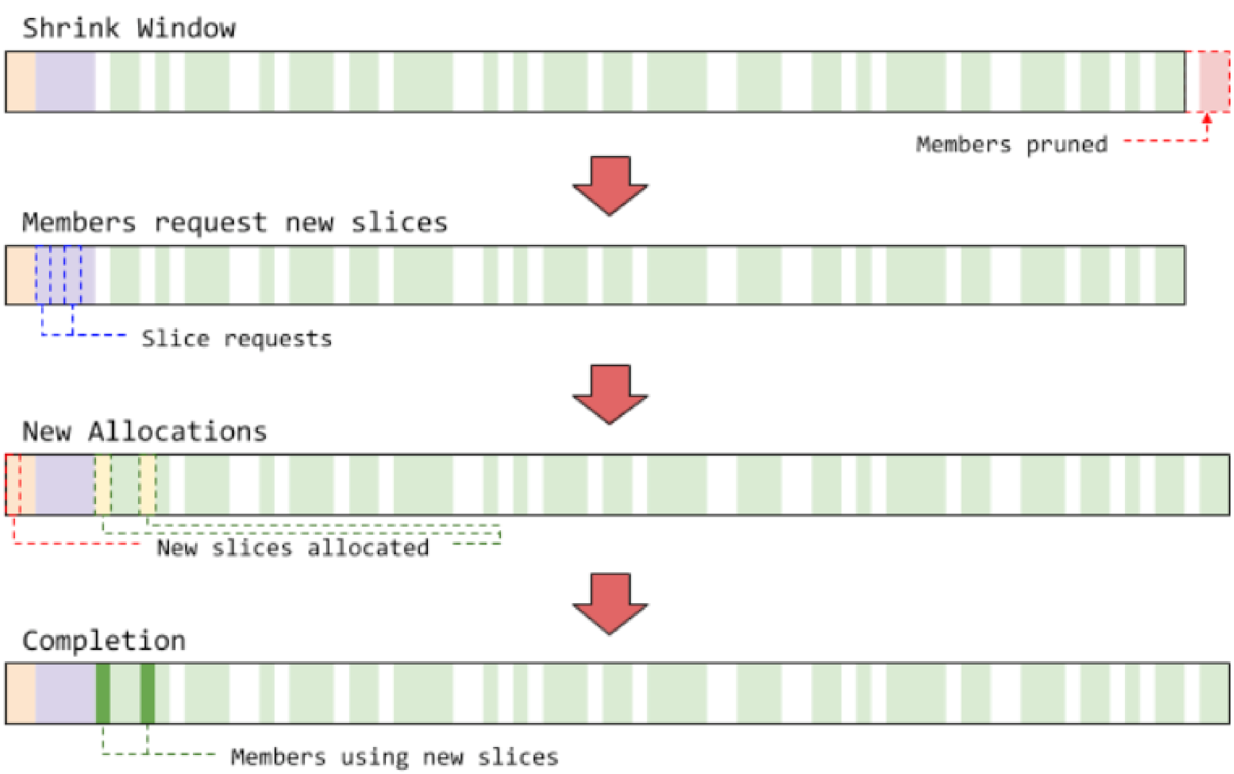

Other Considerations

While the above sections discuss the overall implementation of a MASS BLE system, there are some other considerations and potential alternatives discussed here for practical implementations.

Multiple Management Transmitters

While the discussion above considers the management system as a single unit, in practice there may be multiple transmitters required to cover the desired area. This section discusses options for implementing multiple management transmitters.

Time Synchronization

The naive implementation of multiple transmitters is to precisely synchronize their clocks and transmit management packets simultaneously from each transmitter. For small installations this may be sufficient, but larger installations will suffer degradation due to the finite speed of light. One bit in BLE is transmitted for 1 microsecond. For two transmitters to send the same signal without problems, those 1 microsecond windows must mostly overlap. While accurate synchronization of transmit time may be feasible, the distance between transmitters and the receiver may impact that synchronization.

As an example, let’s say a receiver is seeing the signal from two transmitters, one of which is 150m closer as in the below diagram. Even if the signal started out at the exact same time, the 150m extra traveled by one of the signals introduces a 0.5 microsecond deviation. It may be possible for a receiver to still discern signal in these conditions, but it may be degraded.

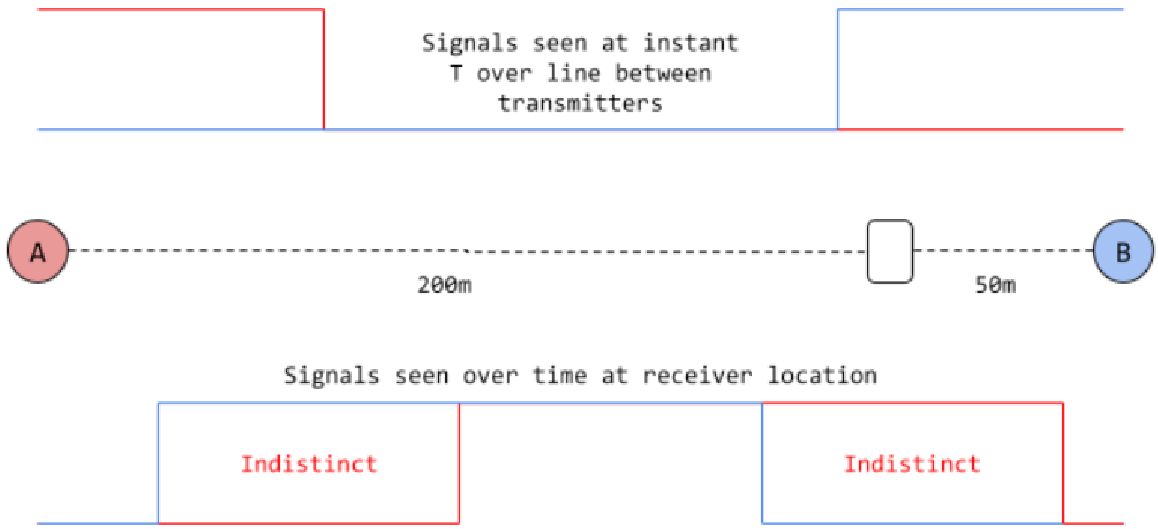

Grid Time Division

In order to minimize interference from nearby management transmitters, the transmitters can be laid out in a grid where each transmitter only transmits a subset of the management packets. The transmission pattern can be crafted such that only one transmitter near any given point transmits each packet. This should limit the chance of a member seeing packets from multiple transmitters with conflicting timing.

Triangular Grid: The ideal grid for maximizing the number of packets each transmitter can send consists of transmitters put at the vertices of equilateral triangles as in the below diagram. Each vertex of each triangle transmits ⅓ of the packets. In the below diagram, the packets are divided into subsets A, B, and C and each vertex is labeled with the subset that it transmits.

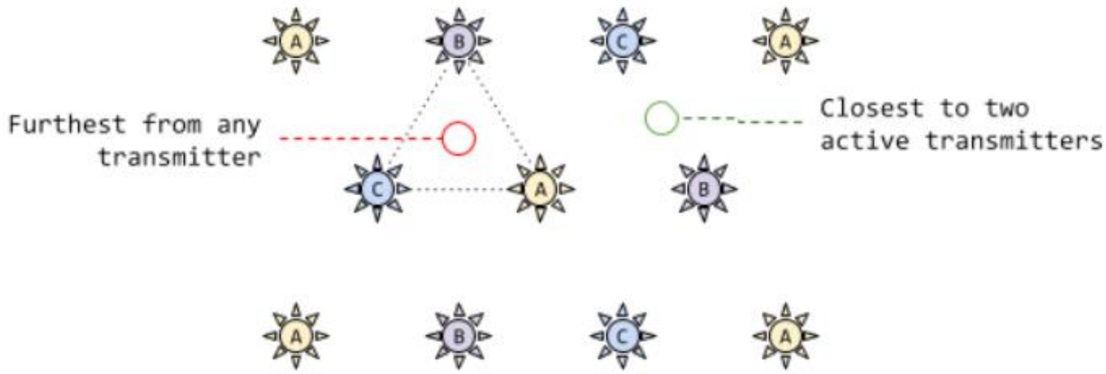

With this grid, when any transmitter is transmitting, none of its six neighbors is transmitting. In order for every point within the triangle to be covered by at least one transmitter, the center of each triangle must be within range. This means the distance between transmitters should be at most $\sqrt{3}$ times the range of one transmitter. Ideally, a point equidistant from two transmitters using the same subset should not see the signal from either. This means the distance between transmitters is at least $\frac{2}{\sqrt{3}}$times the range of one transmitter. Splitting the difference gives us an ideal separation of $\frac{5}{6}\sqrt{3}$ times the range of one transmitter.

Rectangular Grid: Another approach to the grid is a rectangular grid where transmitters are placed at the corners of squares. With this approach, each transmitter could transmit either ¼ or ½ of the packets. The below diagram labels each intersection with A, B, C, or D, but an alternative would be to have A and D transmit the same packets and B and C transmit the same packets.

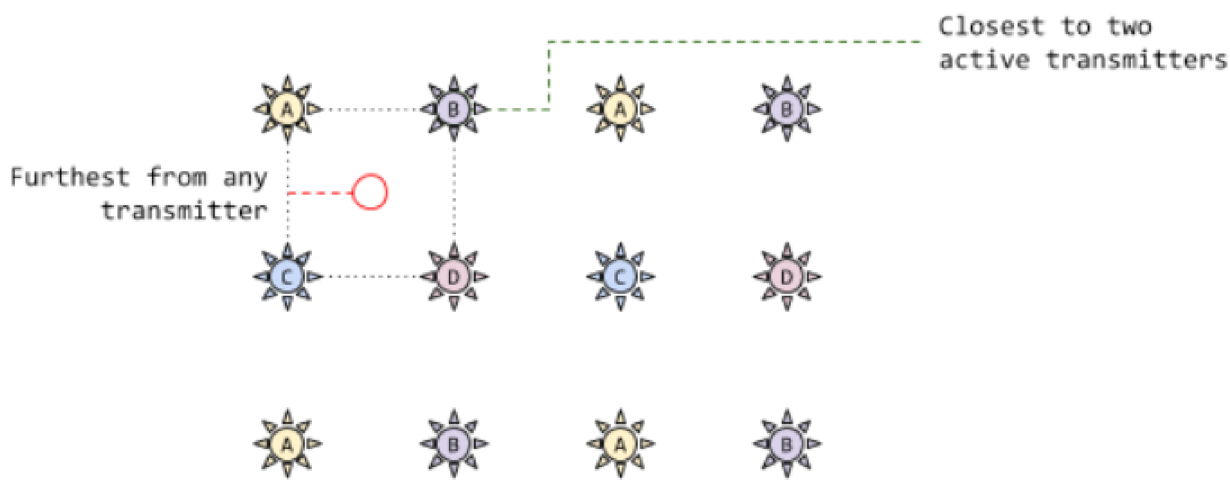

With this grid the center of each square is the furthest from each transmitter, meaning the separation between transmitters can be at most $\sqrt{2}$ times the range of one transmitter. With transmitters transmitting every fourth packet, the midpoint between active transmitters is exactly where another transmitter is placed, meaning the distance should be at least the range of one transmitter. For this arrangement the difference can be split such that the transmitters are placed at $\frac{1+\sqrt{2}}{2}$ times the range of one transmitter.

If A+D and B+C are combined to have each transmitter transmit half of the points, the center of the square becomes the midpoint between active transmitters. This means that the separation should be at least $\sqrt{2}$ times the range of the transmitter, which is the same as the maximum. This makes the center of the square degenerate in that in order to be covered it must see the same strength signal from two active transmitters. For this to work correctly, these transmitters must be synchronized well so that the signals from those two transmitters appear as one.

Interleaved Management Packets

An alternative to having the management system send a series of management packets at the beginning of each window would be to evenly space them throughout the window between other member slices.

The advantage of evenly spacing the management packets is that during system detection a new device is much more likely to see a management packet and not require an additional scan. In addition, if a management scan is required, it can be shorter.

The disadvantage of this approach is that ongoing members may miss the packet at the beginning of the window due to collision or interference and would be required to scan for longer in order to find another packet. This could lead to higher power usage by the members.

Multiple Member Packets

Some members may want to transmit more than one different packet. There are two ways they can achieve this easily. The first is for the member to alternate which packet it sends during its slice in subsequent windows. As an example, the member may want to transmit packet B at 1/10 the rate of packet A. This could be done by transmitting packet A for 10 slices in a row and then transmitting packet B in the next slice. This would interrupt the pattern for packet A but would achieve the desired ratio.

The second alternative is for the member to request a second slice with a different ID. The member can use one slice for packet A and one for packet B. The disadvantage of this approach is that it creates more load on the system and doesn't allow for the packets to transmit at different rates. If the member only transmits packet B during 1/10th of its slices, it may end up getting removed by the management system.

Member Slice Sharing

It may be desirable for members to be able to transmit at a lower frequency than the overall MASS BLE system. A straightforward solution to this would be to only transmit on a subset of the windows, but this leads to less than ideal spectrum usage and could lead to the management system thinking the member is not present.

One could implement a system where during onboarding the member could request a subset of a slice rather than a full slice. The member could request ⅛ of the frequency for example, and the management system would tell the member how to identify the windows during which it should transmit. This would require that the management signal include a window index in order to identify the subset.

With this approach, the management system could assign the same slice to multiple members that don't want a full scale transmission. This would lead to higher spectrum usage overall. As a caveat, the management system would have to take the member's scale into account when determining things like device removal and re-allocation.

Alternatively, scale can also be increased by using additional channels or changing the transmit rate of advertisers. Transmit power may also be tuned in order to decrease the physical area in which devices' transmissions can collisions.

CONCLUSION

This disclosure describes managed adaptive spectrum sharing for BLE (MASS BLE), a technique that uses time division multiple access to enable co-located BLE advertisers to share spectrum with graceful degradation of functionality.