




HALOGEN

ACOUSTIC ECHO CANCELLATION GUIDE



Table of Contents

CHAPTER 1: Introduction to Conferencing	1
Conferencing Basics	1
AEC Introduction	2
CHAPTER 2: HAL System Conferencing Solution	4
Basic Halogen Conferencing Configuration	4
Basic Halogen Conferencing with Voice Lift	5
Introduction to Halogen Conference Room Combine	7
Halogen Conference Switchboard	11
CHAPTER 3: Conclusion	13


CHAPTER 1: Introduction to Conferencing

This document provides an introduction to HAL System conferencing solutions. If you are not familiar with Halogen software, a quick review is advised before jumping into conferencing.

We begin with a brief review of the most basic conference application, communication between two rooms. We discuss half-duplex communication where audio only travels in one direction at a time and duplex communication, where audio travels in both directions at the same time. A need for Acoustic Echo Cancellation (AEC) when using duplex communication is explained, followed by a brief explanation of how AEC works. After a review of these basics, we present a number of basic Halogen conference configurations that demonstrate the use of available Halogen conferencing blocks to solve common conferencing needs.

Conferencing Basics

The following diagram shows a very basic conference application with communication between two rooms:



If you implement the above example with half-duplex there is no echo because audio only goes one way at a time and there is no echo return path. Half-duplex is problematic as it does not allow a normal conversation to take place. Most conferencing applications prefer duplex operation and often require it since this provides communication in both directions at all times, allowing far more natural communication.

The problem with duplex communication without AEC is that it can result in echo at one or both ends of the conference. This occurs when a near-end microphone picks up far-end audio from the overhead loudspeaker and returns a delayed signal to the far-end, which is then sent back to the near-end. This delayed feedback causes an echo that can impede intelligibility and/or make a system unstable. The type of connection between the far-end and near-end determines the length of the delay, which may be from several milliseconds up to a couple hundred milliseconds.

It is possible to reduce echo by carefully selecting room acoustics, microphone type and placement, loudspeaker type and location and volume levels. Even when these elements of a design are optimal, it is not possible to fully eliminate echo without the use of active acoustic echo cancellation. However, paying close attention to these elements of design improves the performance of any system with or without AEC. We refer to echo reduction


achieved through acoustic, microphone and loudspeaker choices as Echo Return Loss (ERL), which indicates the reduction in echo before AEC is applied. The higher this reduction, the better the system will perform with AEC.

AEC Introduction

As stated earlier, in addition to the desired near-end speech, the local microphone picks up a portion of far-end speech from the near-end loudspeaker. AEC reduces echo by subtracting this unwanted portion of the microphone signal before sending it to the far-end. We refer to this additional active echo reduction as Echo Return Loss Enhancement (ERLE). Total Echo Reduction (TER) is the combined echo reduction achieved with ERL and ERLE. The remaining audio consists of echo-free speech and some miscellaneous acoustic sounds, such as room noise. For this reason, AEC often includes ambient noise reduction to further improve the signal sent to the far-end. The following diagram shows our original two room conference application with the addition of AEC. Note that to remove echo from both rooms, each room requires AEC processing.



NOTE: When echo is heard, it is the far-end audio system that is the culprit and needs work, not the local room.



How does AEC accomplish this task? First, an AEC algorithm models the room. It does this by comparing the reference signal (typically the same signal sent to the near-end loudspeaker) to what the microphone picks up. The model includes everything in the signal path between the AEC Reference and the AEC Local Mic In. Once the AEC algorithm has modeled a room, it is able to look at the AEC Reference, estimate what will appear at the AEC Local Mic Input and subtract the estimated signal from what the microphone picks up.



NOTE: An AEC algorithm's ability to remove echo is dependent on the accuracy of its room model.

Because the ability to remove echo is dependent on the accuracy of the room model, it is important to avoid unnecessary changes during operation that vary the relationship between AEC Reference and AEC Local Mic In. We'll talk more about this later when we discuss where to locate signal processing in a system and the necessity of applying equal processing to room and reference signals. (See "Basic Halogen Conferencing with Voice Lift" on page 5)

CHAPTER 1: Introduction to Conferencing


One misconception is that AEC simply subtracts the reference from what is received from a microphone input. This is not true. AEC only subtracts the portion of the reference signal that the room model estimates will be picked up by the microphone.

CHAPTER 2: HAL System Conferencing Solution

The following examples provide an introduction to Halogen conferencing solutions by demonstrating basic configurations and introducing specialized Halogen processing blocks that support conferencing. More detailed information about individual conferencing blocks may be found in the reference section of Halogen Help.


Basic Halogen Conferencing Configuration

The following Halogen processing map demonstrates a simple conference configuration. Remember that a complete conference system employs a similar configuration for near-end and far-end rooms:



This configuration is more advanced than the basic two room conference diagrammed in the AEC Introduction, but the same basic principles apply. This configuration may be expanded to support any number of microphones and AV sources as well as a variety of room processing. Eight channels of drag and drop AEC are available in an EXP7x Halogen expansion device. All other processing blocks in these examples are available in a HAL1x device.

In this example, From VoIP 1 is a signal from the far-end and To VoIP 1 is the signal returned to the far-end. These signals are used to establish a duplex connection between the two conference rooms. An AEC block is used for each near-end microphone. Each AEC block includes AEC with ambient noise reduction, Automatic Gain Control (AGC) and Parametric Equalizer (PEQ). As shown in the following image, AGC and PEQ functions follow AEC processing, allowing adjustment without interfering with room modeling.



Following the AEC blocks is an automatic Gain-sharing Mixer. This mixer distributes gain between any number of microphones based on received signal amplitude and priority. Following the Gain-sharing Mixer is a standard Mixer that adds local AV sources such as CD, DVD or Blu-ray players and laptops to what is sent to the far-end. In this way, echo-free audio from any number of room microphones and AV sources is available for the far-end.


The configuration above does not require local speech reinforcement so inputs to the room are simply a signal from the far-end and local AV sources. Output to the far-end consists of echo free speech from the two near-end microphones and local AV sources.

Additional processing for the room includes Compression, PEQ and Level control. Earlier we mentioned the importance of a reference signal accurately representing the signal output to a room in order to maintain an accurate room model for the AEC algorithm. The configuration satisfies this requirement by locating the AEC Reference after this additional processing. When providing a linked level control for the room, it is important to include the level control before the reference so that the room output signal and the reference track each other.

Basic Halogen Conferencing with Voice Lift

Now that we're more comfortable with the basic principles of conferencing, let's add Voice Lift to a configuration. Voice Lift provides near-end microphone sound reinforcement for audio within a room so that everyone in the near-end room can hear clearly those who are speaking. The configuration shown below supports local speech reinforcement. What is sent to the far-end is unchanged from our previous example. The difference is in the Room Processing portion, which now supports local speech reinforcement. While this adds a bit of complexity, Halogen provides the blocks required to keep it relatively simple.

CHAPTER 2: HAL System Conferencing Solution



Microphone inputs for the room are taken before the AEC blocks and mixed with a Gain-sharing Mixer as shown. This signal is then mixed with the signal from the far-end and local AV sources. At the Tracking Compressor input, we now have two available mixes for the room and reference: one with from-far-end, local AV and local microphones and one with just from-far-end and local AV. There are a couple of different ways to use these two mixes.

First, let's consider an example with a head table where the microphones are located and a larger audience area that requires Voice Lift. Voice Lift is not required at the head table so loudspeakers in this zone use the mix without Voice Lift. And because the mix without Voice Lift is most closely coupled with the microphones, it is also used for the AEC reference. The mix with the microphone signal serves the audience area. This is a very basic mix-minus application. More complex mix-minus applications are outside the scope of this document, but the principles are the same.

Second, let's consider an example where microphones are less segregated and more tightly coupled with local sound reinforcement. This application is a bit less straightforward. Which signal should be used for the reference? Well, it depends. If ERL is poor, a better choice is to use the mix without near-end microphones for the AEC reference as doing so avoids adding AEC artifacts to speech sent to the far-end. If ERL is very good, using the mix with near-end microphones may improve performance. It is a misconception that including the local microphone mix in the reference subtracts all speech in the AEC block. AEC only subtracts the portion of speech it estimates will arrive from the local loudspeaker and in some cases, this can be a good thing. When in doubt, the reference should use the mix without the microphone signal.

In both scenarios, we need to ensure that both mixes receive identical signal processing as shown above. Halogen provides multichannel and tracking DSP blocks that help satisfy this requirement. In this example, a Tracking Compressor applies compression based exclusively on the Room signal and applies the required gain adjustment to both Room Out and Reference Out signals. The compressor is followed by a Multichannel PEQ and a multichannel Level control that apply identical equalization and level adjustment to both signals. It is important to Link to a level control that affects both signal paths and is located before the Reference output. Without these

special blocks, a system designer would need to set up independent processing blocks for each signal path and somehow ensure that they stay in sync.

Halogen also offers Multichannel Shelf and Tracking ANC blocks.


A few things worth noting:

1. It is important to place all dynamics processing and any other blocks that a user may adjust during operation before the reference location. Failing to do so will adversely affect AEC operation.
2. It is equally important to not locate dynamics processing or any other blocks that a user may adjust during operation between a microphone and an AEC block as this adversely affects AEC operation.
3. It does not adversely affect AEC performance to place fixed equalization on a microphone input before it is split between the AEC and local Voice Lift mixes. This is often highly desired as low-cut and PEQ requirements are more dependent on the microphone and you avoid duplicating processing for AEC and the local mix.

Introduction to Halogen Conference Room Combine

Now that we are getting comfortable with the basics, it's time to dive into something that many fear and few have attempted: what would it take to include conferencing in a room combine application? It's not that hard to conceptualize, but often difficult to implement using typical tools. All you need to do is properly combine microphones, AV and *from far end* inputs, acquire an appropriate reference and output signals to the appropriate rooms and *to far end* destinations.

Halogen provides a highly integrated Conference Room Combine block that greatly simplifies this process. The image below shows a basic conference room combine configuration. For this example the number of microphones and AV sources is limited to one per room and the number of rooms is limited to two in order to simplify the configuration and allow us to focus how the system works. The same principles apply as you add more rooms, microphones and AV sources.




CHAPTER 2: HAL System Conferencing Solution


Those familiar with a standard Halogen Room Combine block will notice some similarities. The Conferencing Room Combine block includes the same Distributed Program Bus, Local Automixer (Auto Mixer), Local AV Mixer (Mixer), Paging Zone, Room Out and Record Out and base rooms combine in the same manner. In the Conference Room Combine, each base room also includes a Far End Auto Mixer, a From Far End input, a To Far End output and a Reference Out output.

- **Far End Auto Mixer** is where you connect any microphone that needs to be heard at a *far end*. To reduce echo, this configuration includes an AEC block before each Far End Auto Mixer input. You may connect a microphone to a Far End Auto Mixer, Local Auto Mixer or both. Connecting microphones to the Local Auto Mixer allows them to be heard in the local room.
- **From Far End** is where a From Far End conference connection into a room is made. This may be from VoIP, Video Codec, POTs or any other *far end* conference source. When combining rooms, not all rooms need to have a conference input, but more than one room may have this input as in this example. When rooms combine, the room combine block mixes together all of the From Far End signals as we'll discuss later.
- **Reference Out** provides a reference signal that feeds an AEC reference input. The room processor derives this signal by mixing local AV and *from far end* sources as in earlier examples. An identical reference is sent to each room when rooms are combined. If unique signal processing is applied to each base room as shown on page 11, each base room reference needs to have this same processing applied before being used.
- **To Far End** provides a signal for a *far end*. As with the reference output, the block sends an identical signal to all combined rooms. In the above example there is a VoIP line in each room and either may require a To Far End signal.

A look inside a Room Processor provides a clearer view of its structure. It shows a typical room combine assortment of I/O plus the additional I/O required for conferencing. As with the standard Room Combine block, a similar Room Processor exists for Room A, Room B and Room A+B.



The following figure shows the Room A+B processor. Note that signals from the two rooms combine in the **Far End Auto Mixer**, the **Local Auto Mixer**, the **Local AV Mixer** and the **From Far End Mixer** blocks.



The **Far End Auto Mixer** combines all connected microphones from both rooms needed for the *far end*. The **To Far End Mixer** allows a custom mix of *far end* microphones, **Local AV** sources and **Distributed Program Bus** (DPB) selection. The block splits this composite To Far End signal out to each room as shown. How the **Record Out**, **Reference Out** and **Room Out** are mixed is better understood by taking a closer look at the Room Matrix Mixer which is central to all conferencing room processors.

The following figure shows the properties dialog for this mixer:

CHAPTER 2: HAL System Conferencing Solution



The Record Out signal consists of a unique mix of the four available sources. You can use this signal for recording an event or supplying audio to an overflow room. You could also use this signal in conjunction with the Room Out to provide one zone with Voice Lift and another without.

The Room Out signal also consists of a unique mix of the four available sources.

Remember we discussed a need for the AEC reference to match what is sent to a room? In the case of the **Room Matrix Mixer**, this means that if you add a source to a room mix, the same needs to happen in the reference mix. In a similar way, if you delete a source from the room mix, you also need to delete it from the reference mix. In addition, you also need to make corresponding changes to the reference signal when adjusting a crosspoint level or room output level. For this reason, the **Room Matrix Mixer**'s crosspoint mix levels and output level control follow the settings for the **Room Out**. The **Room Matrix Mixer** dialog shows **Reference Out** controls in half-tone to indicate that you may not adjust them independently from the **Room Out** mix. Displaying the reference controls in this way also clearly shows that the block applies identical settings to both the **Room Out** and the **Reference Out**. When linked level controls are required for a room, you should link to Level controls in the **Room Matrix Mixer** to ensure that **Room Out** and **Reference Out** track.


As in our previous examples, we do not include *near end* microphones in the reference. If your application is suitable for near end microphones to be included in the reference, you can use the room output for both the room and the reference. As shown earlier, the mix without local microphones may be appropriate for use in loudspeakers co-located with a head table.

Once **Room Out** and **Reference Out** signals leave the **Room Matrix Mixer**, we still need to ensure that they track each other. In the **Room Processor**, a **Paging Zone** block follows the **Room Matrix Mixer**. As with a standard Halogen Room Combine block, paging combines as rooms combine. Paging into a room ducks all other signals by

a set amount and adds paging audio to the room. When this happens, the reference needs to follow. The Conferencing Room Combine block provides a special “tracking” **Paging Zone** block able to duck and add page audio to the reference as well as the room. Note that this behavior is required whether the **Reference Out** is used for an AEC reference, mix-minus zone or both.

Now is a good time to discuss why the block sends an identical reference signal to each room. Consider the configuration shown below. For demonstration purposes, this configuration includes different processing for each room output (often a real world requirement). In particular, Room-A includes a compressor while Room-B includes a **PEQ**. As we’ve learned all along, the AEC reference signal must always match the room output signal, so the configuration uses a **Tracking Compressor** in Room-A to ensure that the reference matches the room output. In a similar manner, the configuration uses a **Multichannel PEQ** block in Room-B.

When you combine the two rooms, the Conferencing Room Combine block’s **Reference Out A** and **Reference Out B** are the same signal. The additional processing that creates Reference A and Reference B is different. Deciding which reference to use in each of the AEC blocks depends on where the microphones for the *far end* are located. Use Reference A if all microphones are located at a head table in Room A. Use Reference B if all microphones are located at a head table in Room B. If some microphones are in Room A and some in Room B, you may use both references -- with Reference A sent to the AEC blocks for microphones located in Room A and Reference B sent to the AEC blocks located in Room B as shown below.




Halogen Conference Switchboard

Halogen provides another useful block: the **Conference Switchboard**, demonstrated in the figure below. This example configuration has three rooms co-located in a building that from time to time need to conference with each other. Further, any combination of rooms may want to participate in a conference at any given time. This example does not necessarily use VoIP or other Telephony between these local rooms, as these add unnecessary delay and can degrade sound quality compared to a direct connection.

You could accomplish the required routing with a matrix mixer, but doing so requires ensuring paired From-Far-End and To-Far-End routing for each room. This is difficult to do without setting up complex presets to adjust multiple matrix mixer levels and/or cross-points simultaneously. In addition, it is highly desirable to create intuitive system controls for the end-user, which is difficult with the matrix mixer approach.

CHAPTER 2: HAL System Conferencing Solution

The **Conference Switchboard** solves these problems in an elegant way. Each Room has a From-Switchboard and To-Switchboard connection with the **Conference Switchboard** block. If a user is in the Green Room she simply selects conference participants from a list. A DR remote, a Web Control, and/or a 3rd-party control system can provide intuitive control for her to make this selection. In the example shown below, the Green Room user has selected Blue and Amber rooms for a conference. The **Conference Switchboard** supports up to 32 channels in any combination of room and telephony participants.



CHAPTER 3: Conclusion

While there are endless possibilities for conference configurations, Halogen provides the tools necessary to keep tasks manageable while maintaining a high degree of flexibility. Drag and drop AEC via an EXP7x, Conference Room Combine, Conference Switchboard, Tracking Compressor, Tracking ANC and multichannel Level, PEQ and Shelf blocks comprise a powerful set of tools that help manage these complex tasks. Examples in this document cover the basic principals required to design both basic systems and highly complex systems.

AMX® and the AMX logo are registered trademarks of AMX.

Stardraw Control is a trademark of Stardraw.com Ltd.

Crestron® is a registered trademark of Crestron Electronics, Inc.

Lexan® is a registered trademark of General Electric.

Decora® is a registered trademark of Leviton.

Windows® is a registered trademark of Microsoft Corporation in the United States and other countries.

Apple, Mac, Macintosh, iTunes, Safari, QuickTime, GarageBand, and OS X are registered trademarks of Apple Inc., registered in the U.S. and other countries.

Halogen is a trademark of Rane Corporation.