

Enhanced Full-duplex Speakerphone IC

Features

- Single-chip, full-duplex, hands-free operation
- Optional Tx Noise Guard
- Programmable attenuation during double-talk
- Optional 34 dB microphone preamplifier
- Dual channel AGC'ed volume controls with mute
- Dual integrated 80 dB IDR codecs
- Speech-trained Network and Acoustic Echo Cancellers
- Rx and Tx supplementary echo suppression
- Configurable half-duplex training mode
- Powerdown mode
- Microcontroller Interface

General Description

Most modern speakerphones use half-duplex operation, which alternates transmission between the far-end talker and the speakerphone user. This is done to ensure stability because the acoustic coupling between the speaker and microphone is much higher in speakerphones than in handsets where the coupling is mechanically suppressed.

The CS6422 enables full-duplex conversation using echo cancellation and suppression in a single-chip solution. The CS6422 can easily replace existing half-duplex speakerphone ICs with a huge increase in conversation quality.

The CS6422 consists of telephone & audio interfaces, two codecs and an echo-cancelling DSP.

ORDERING INFORMATION

See [page 48](#).

CDB6422 Evaluation Board

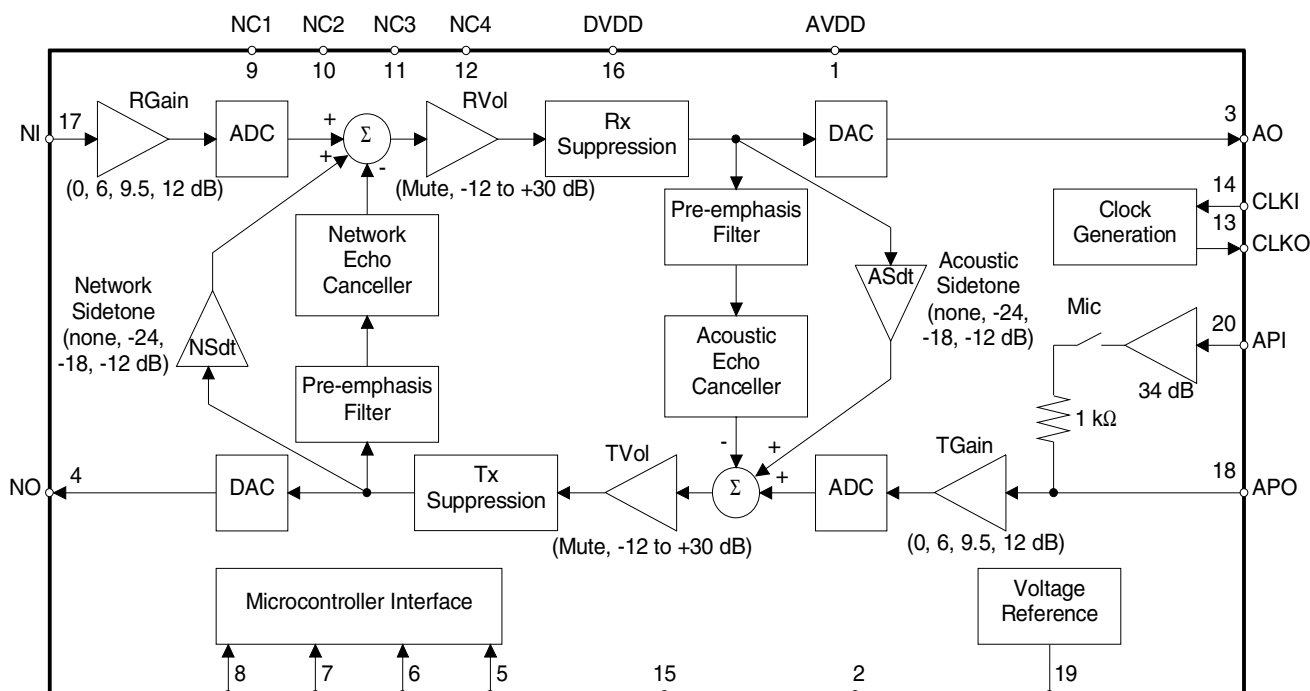


TABLE OF CONTENTS

1. CHARACTERISTICS AND SPECIFICATIONS	5
ABSOLUTE MAXIMUM RATINGS	5
RECOMMENDED OPERATING CONDITIONS	5
POWER CONSUMPTION	5
ANALOG CHARACTERISTICS	5
ANALOG TRANSMISSION CHARACTERISTICS.....	6
MICROPHONE AMPLIFIER	6
DIGITAL CHARACTERISTICS	6
SWITCHING CHARACTERISTICS	7
2. OVERVIEW	9
3. FUNCTIONAL DESCRIPTION	9
3.1 Analog Interface	9
3.1.1 Acoustic Interface	10
3.1.2 Network Interface	11
3.2 Microcontroller Interface	11
3.2.1 Description	11
3.2.2 Register Definitions	12
3.3 Register 0	13
3.3.1 Mic - Microphone Preamplifier Enable	14
3.3.2 HDD - Half-Duplex Disable.....	14
3.3.3 GB - Graded Beta.....	14
3.3.4 RVol - Receive Volume Control.....	14
3.3.5 TSD - Transmit Suppression Disable	14
3.3.6 ACC - Acoustic Coefficient Control	15
3.3.7 TSMde - Transmit Suppression Mode.....	15
3.4 Register 1	16
3.4.1 THDet - Transmit Half-Duplex Detection Threshold.....	17
3.4.2 Taps - AEC/NEC Tap Allocation	17
3.4.3 TVol - Transmit Volume Control.....	17
3.4.4 RSD - Receive Suppression Disable.....	17
3.4.5 NCC - Network Coefficient Control.....	17
3.4.6 AuNECD - Auto re-engage NEC Disable	17
3.5 Register 2	18
3.5.1 RHDet - Receive Half-Duplex Detection Threshold	19
3.5.2 RSThd - Receive Suppression Threshold	19
3.5.3 NseRmp - Noise estimator Ramp rate	19

3.5.4 HDly - Half-Duplex Holdover Delay	19
3.5.5 HHold - Hold in Half-Duplex on Howl	19
3.5.6 TDSRmp - Tx Double-talk Suppression Ramp rate	19
3.5.7 RDSRmp - Rx Double-talk Suppression Ramp rate	20
3.5.8 IdITx - half-duplex Idle return-to-Transmit	20
3.6 Register 3	21
3.6.1 TSAtt - Transmit Suppression Attenuation	22
3.6.2 PCSen- Path Change Sensitivity	22
3.6.3 TDbtS - Tx Double-talk Suppression attenuation	22
3.6.4 RDbtS - Rx Double-talk Suppression attenuation	22
3.6.5 TSThd - Transmit Suppression Threshold	22
3.6.6 TSBias - Transmit Suppression Bias	22
3.7 Register 4	23
3.7.1 AErle - AEC Erle threshold	24
3.7.2 AFNse - AEC Full-duplex Noise threshold	24
3.7.3 NErle - NEC Erle threshold	24
3.7.4 NFNse - NEC Full-duplex Noise threshold	24
3.7.5 RGain - Receive Analog Gain	24
3.7.6 TGain - Transmit Analog Gain	24
3.8 Register 5	25
3.8.1 HwID - Howl detector Disable	26
3.8.2 TD - Tone detector Disable	26
3.8.3 APCD - Acoustic Path Change detector Disable	26
3.8.4 NPCD - Network Path Change detector Disable	26
3.8.5 APFD/NPFD - Acoustic Pre-emphasis Filter Disable/Network Pre-emphasis Filter Disable	26
3.8.6 AECD - Acoustic Echo Canceller Disable	27
3.8.7 NECD - Network Echo Canceller Disable	27
3.8.8 ASdt - Acoustic Sidetone level	27
3.8.9 NSdt - Network Sidetone level	27
3.9 Reset	28
3.9.1 Cold Reset	28
3.9.2 Warm Reset	28
3.9.3 Reset Timer	28
3.10 Clocking	28
3.11 Power Supply	29
3.11.1 Power Down Mode	29
3.11.2 Noise and Grounding	29
4. DESIGN CONSIDERATIONS	31
4.1 Algorithmic Considerations	31
4.1.1 Full-Duplex Mode	31
4.1.1.1 Theory of Operation	31
4.1.1.2 Adaptive Filter	32
4.1.1.2.1 Pre-Emphasis	32
4.1.1.2.2 Graded Beta	33
4.1.1.3 Update Control	33
4.1.1.4 Speech Detection	33
4.1.2 Half-Duplex Mode	34
4.1.2.1 Idle Return to Transmit	34
4.1.3 AGC	34
4.1.4 Suppression	35
4.1.4.1 Transmit Suppression	36
4.1.4.2 Receive Suppression	36

4.1.4.3 Double-talk Attenuation	36
4.1.4.4 Noise Guard	37
4.2 Circuit Design	37
4.2.1 Interface Considerations	37
4.2.1.1 Analog Interface	37
4.2.1.2 Microcontroller Interface	37
4.2.2 Grounding Considerations	38
4.2.3 Layout Considerations	38
4.3 System Design	38
4.3.1 Gain Structure	38
4.3.2 Testing Issues	39
4.3.2.1 ERLE	39
4.3.2.2 Convergence Time	40
4.3.2.3 Half-Duplex Switching	40
5. PIN DESCRIPTIONS	41
6. GLOSSARY	44
7. PACKAGE DIMENSIONS	46

LIST OF FIGURES

Figure 1. CLKI Timing	7
Figure 2. Reset Timing	7
Figure 3. Microcontroller Interface Timing	7
Figure 4. Typical Connection Diagram (Microphone Preamplifier Enabled)	8
Figure 5. Typical Connection Diagram (Microphone Preamplifier Disabled)	8
Figure 6. Analog Interface	10
Figure 7. Microcontroller Interface	12
Figure 8. Suggested Layout	29
Figure 9. Ground Planes	30
Figure 10. Simplified Acoustic Echo Canceller Block Diagram	31
Figure 11. How the AGC works (TVol = +30 dB)	35

LIST OF TABLES

Table 1. Full scale voltages for each gain stage	11
Table 2. MCR Control Register Mapping	12
Table 3. Register 0 Bit Definitions	13
Table 4. Register 1 Bit Definitions	16
Table 5. Register 2 Bit Definitions	18
Table 6. Register 3 Bit Definitions	21
Table 7. Register 4 Bit Definitions	23
Table 8. Register 5 Bit Definitions	25

1. CHARACTERISTICS AND SPECIFICATIONS

ABSOLUTE MAXIMUM RATINGS

Parameter	Symbol	Min	Max	Units
DC Supply (AVDD, DVDD)		-0.3	6.0	V
Input Current (Except supply pins)	I_{in}	-10	+10	mA
Input Voltage	Analog V_{ina} Digital V_{ind}	-0.3 -0.3	AVDD+0.3 DVDD+0.3	V
Ambient Operating Temperature	T_A	-40	85	°C
Storage Temperature	T_{stg}	-65	150	°C

WARNING: Operation beyond these limits may result in permanent damage to the device.
Normal operation is not guaranteed at these extremes.

RECOMMENDED OPERATING CONDITIONS

Parameter	Symbol	Min	Typ	Max	Units
DC Supply (AVDD, DVDD)		4.5	5.0	5.5	V
Ambient Operating Temperature	Commercial T_{AOp} Industrial	0 -40	25 25	70 85	°C

POWER CONSUMPTION ($T_A = 25^\circ\text{C}$, DVDD = AVDD = 5 V, $f_{XTAL} = 20.480\text{ MHz}$) (Note 1)

Parameter	Symbol	Min	Typ	Max	Units
Power Supply Current, Analog ($\overline{RST}=0$)	P_{DA0}			1	mA
Power Supply Current, Analog ($\overline{RST}=1$)	P_{DA}		10	20	mA
Power Supply Current, Digital ($\overline{RST}=0$)	P_{DD0}			1	mA
Power Supply Current, Digital ($\overline{RST}=1$)	P_{DD}		50	80	mA

Notes: 1. AO and NO outputs are not loaded.

ANALOG CHARACTERISTICS ($T_A = 25^\circ\text{C}$, DVDD = AVDD = 5 V, $f_{XTAL} = 20.480\text{ MHz}$)

Parameter	Symbol	Min	Typ	Max	Units
Input Offset Voltage (APO, NI)			2.12		V
Output Offset Voltage (AO, NO)			2.12		V
Transmit Group Delay (Note 2)				6	ms
Receive Group Delay (Note 2)				6	ms
Input Impedance (APO, NI)	Z_{in}		1.5		MΩ
Load Impedance (AO, NO)	Z_{load}	10			kΩ
Power Supply Rejection (1 kHz)			40		dB

Notes: 2. These parameters are guaranteed by design or by characterization.

ANALOG TRANSMISSION CHARACTERISTICS ($T_A = 25^\circ\text{C}$, $DVDD = AVDD = 5\text{ V}$, $f_{XTAL} = 20.480\text{ MHz}$, $RVol=TVol=RGain=TGain= 0\text{ dB}$, $HDD=TSD=RSD=1$, analog inputs and outputs loaded with resistors and capacitors as shown in the typical connection diagram, Figure 4)

Parameter	Symbol	Min	Typ	Max	Units
Idle Channel Noise C-Message weighted (0-4 kHz) (Inputs grounded through a capacitor) C-Message weighted (0-4 kHz) Psophometrically weighted (0-4 kHz)			-80 11 -78	-73	dBV dBnC0 dBm0p
Signal-to-Noise Ratio C-Message weighted (0-4 kHz) (1.0 V_{rms} , 1 kHz sine wave input)	SNR	73	80		dB
Total Harmonic Distortion C-Message weighted (0-4 kHz) (1.0 V_{rms} , 1 kHz sine wave input)	THD		0.030	0.1	%
Programmable Analog Gain RGain/TGain = 00 RGain/TGain = 01 RGain/TGain = 10 RGain/TGain = 11			0 6 9.5 12		dB
Volume Control Stepsize (TVol/RVol)			3		dB
ADC Full-scale Voltage Input		0.9	1.0		V_{rms}
DAC Full-scale Voltage Output			1.0	1.2	V_{rms}
ADC Noise Floor C-Message weighted (0-4 kHz)			-83		dBV
DAC Noise Floor, DAC muted C-Message weighted (0-4 kHz)			-83		dBV

MICROPHONE AMPLIFIER ($T_A = 25^\circ\text{C}$, $DVDD = AVDD = 5\text{ V}$, $f_{XTAL} = 20.480\text{ MHz}$)

Parameter	Symbol	Min	Typ	Max	Units
Gain ($Z_{source} = 50\Omega$)	A_{mic}		34		dB
Signal-to-Noise Ratio C-Message weighted (0-4 kHz)	SNRm		70		dB
Input Impedance	Z_{inm}		8		$k\Omega$
Input Offset Voltage	V_{offm}		2.12		V

DIGITAL CHARACTERISTICS ($T_A = 25^\circ\text{C}$, $DVDD = AVDD = 5\text{ V}$, $f_{XTAL} = 20.480\text{ MHz}$)

Parameter	Symbol	Min	Typ	Max	Units
High-Level Input Voltage	V_{IH}	$DVDD-1.0$			V
Low-Level Input Voltage	V_{IL}			1.0	V
Input Leakage Current	I_{leak}			10	μA
Input Capacitance	C_{IN}		5		pF

ANALOG TRANSMISSION CHARACTERISTICS ($T_A = -40^{\circ}\text{C}$ to 85°C , $DVDD = AVDD = 5\text{ V}$, $f_{XTAL} = 20.480\text{ MHz}$, $RVol=TVol=RGain=TGain= 0\text{ dB}$, $HDD=TSD=RSD=1$, analog inputs and outputs loaded with resistors and capacitors as shown in the typical connection diagram, Figure 4)

Parameter	Symbol	Min	Typ	Max	Units
Idle Channel Noise C-Message weighted (0-4 kHz) (Inputs grounded through a capacitor) C-Message weighted (0-4 kHz) Psophometrically weighted (0-4 kHz)			-80 11 -78	-72	dBV dBmC0 dBm0p
Signal-to-Noise Ratio C-Message weighted (0-4 kHz) (1.0 V_{rms} , 1 kHz sine wave input)	SNR	72	80		dB
Total Harmonic Distortion C-Message weighted (0-4 kHz) (1.0 V_{rms} , 1 kHz sine wave input)	THD		0.030	0.1	%
Programmable Analog Gain RGain/TGain = 00 RGain/TGain = 01 RGain/TGain = 10 RGain/TGain = 11			0 6 9.5 12		dB
Volume Control Step size (TVol/RVol)			3		dB
ADC Full-scale Voltage Input		0.9	1.0		V_{rms}
DAC Full-scale Voltage Output			1.0	1.2	V_{rms}
ADC Noise Floor C-Message weighted (0-4 kHz)			-83		dBV
DAC Noise Floor, DAC muted C-Message weighted (0-4 kHz)			-83		dBV

MICROPHONE AMPLIFIER ($T_A = 25^{\circ}\text{C}$, $DVDD = AVDD = 5\text{ V}$, $f_{XTAL} = 20.480\text{ MHz}$)

Parameter	Symbol	Min	Typ	Max	Units
Gain ($Z_{source} = 50\Omega$)	A_{mic}		34		dB
Signal-to-Noise Ratio C-Message weighted (0-4 kHz)	SNRm		70		dB
Input Impedance	Z_{inm}		8		$k\Omega$
Input Offset Voltage	V_{offm}		2.12		V

DIGITAL CHARACTERISTICS ($T_A = 25^{\circ}\text{C}$, $DVDD = AVDD = 5\text{ V}$, $f_{XTAL} = 20.480\text{ MHz}$)

Parameter	Symbol	Min	Typ	Max	Units
High-Level Input Voltage	V_{IH}	$DVDD-1.0$			V
Low-Level Input Voltage	V_{IL}			1.0	V
Input Leakage Current	I_{leak}			10	μA
Input Capacitance	C_{IN}		5		pF

SWITCHING CHARACTERISTICS

Parameter	Symbol	Min	Typ	Max	Units
Digital input rise time	t_{rise}			1.0	μs
\overline{RST} low time	t_{RSTL}	1.0			μs
CLKI frequency	f_{CLKI}		20.480		MHz
CLKI duty cycle	t_{LCLKI}	40	50	60	%
CLKI high or low time	t_{HLCLKI}	19.5			ns
Min \overline{DRDY} falling to \overline{DRDY} falling (CLKI = 20.480 MHz)	t_{DRDY}		125		μs
STROBE high or low time	$t_{HLSTROBE}$	55			ns
\overline{DRDY} falling to STROBE rising setup time	t_{sDRDY}	30			ns
DATA valid to STROBE rising setup time	t_{sDATA}	30			ns
STROBE rising to DATA valid hold time	t_{hDATA}	30			ns
STROBE rising to \overline{DRDY} rising hold time	t_{hDRDY}	30			ns
Min \overline{RST} rising to 4th extra STROBE pulse (cold reset)	t_{cRST}		110		ms
Max \overline{RST} rising to 4th extra STROBE pulse (warm reset)	t_{wRST}		100		ms




Figure 1. CLKI Timing

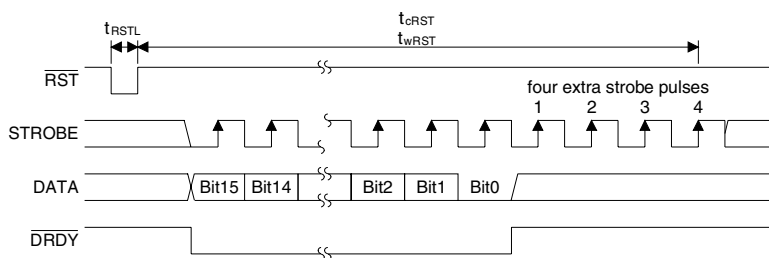


Figure 2. Reset Timing




Figure 3. Microcontroller Interface Timing

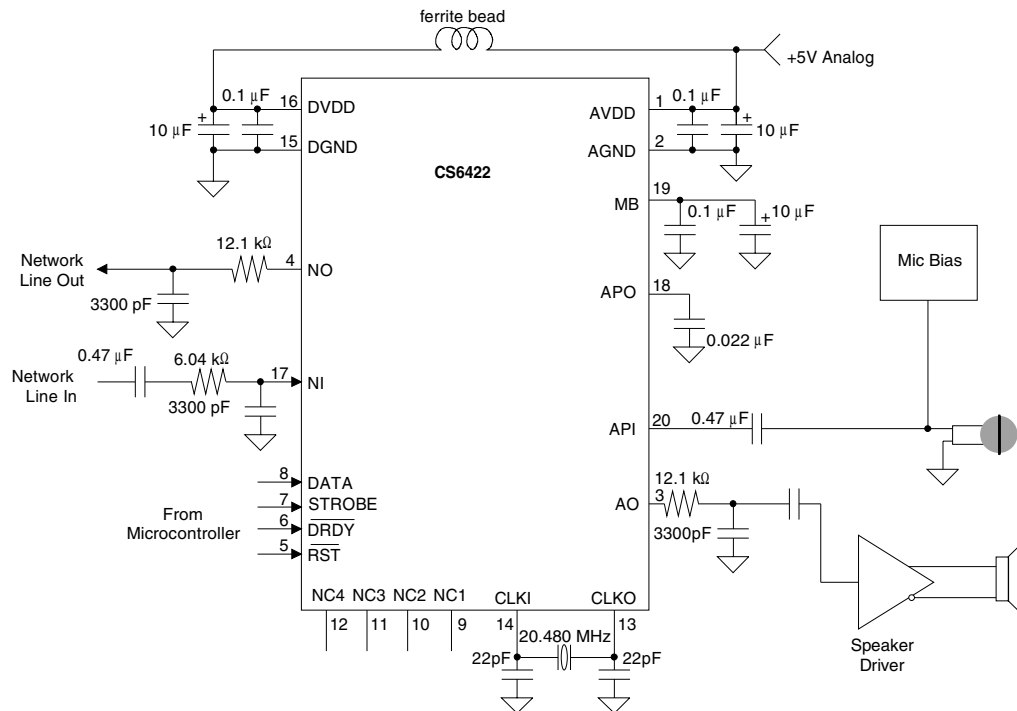


Figure 4. Typical Connection Diagram (Microphone Preamplifier Enabled)

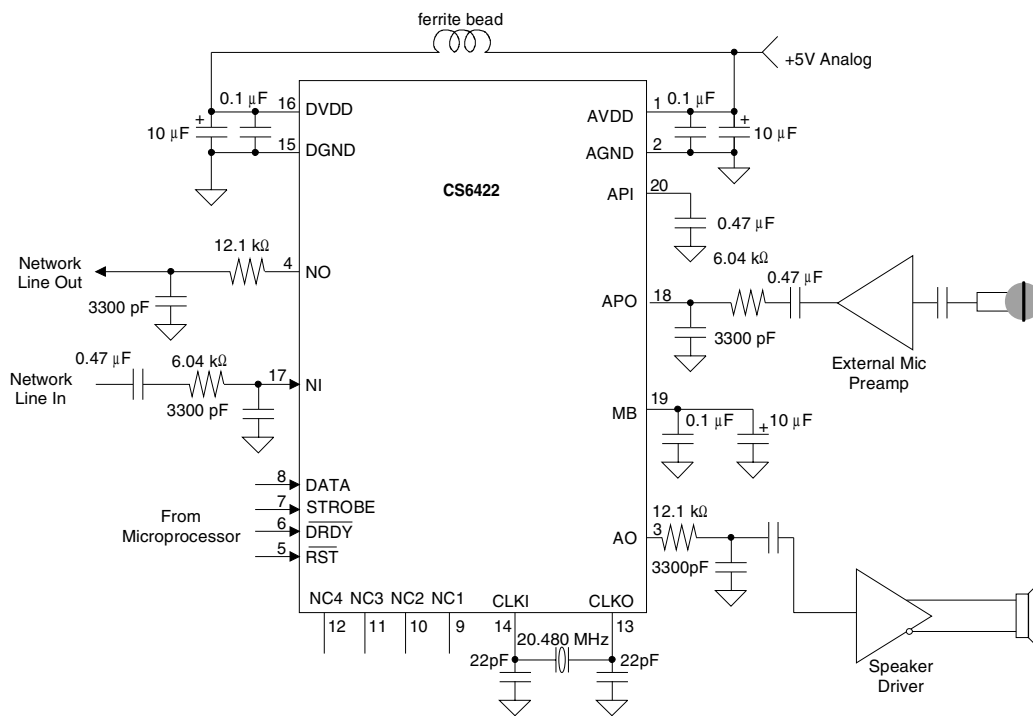


Figure 5. Typical Connection Diagram (Microphone Preamplifier Disabled)

2. OVERVIEW

The CS6422 is a full-duplex speakerphone chip for use in hands-free communications with telephony quality audio. Common applications include speakerphones, inexpensive video-conferencing, and hands-free cellular phone car kits. The CS6422 requires very few external components and allows system control through a microcontroller interface.

Hands-free communication through a microphone and speaker typically results in acoustic feedback or howling because the loop gain of the system exceeds unity by the time audio amplitudes are adjusted to a reasonable level. The solution to the howling problem has typically been half-duplex, where either the transmit or the receive channel is active, never both at the same time. This prevents instability, but diminishes the overall communication quality by clipping words and forcing each talker to speak in turn.

Full-duplex conversation, where both transmit and receive channels are active simultaneously, is the conversation quality we enjoy when using handsets. Full-duplex for hands-free communications is achieved in the CS6422 using a digital signal processing technique called “Echo Cancellation.” The end result is a more natural conversation than half-duplex, with no awkward breaks and pauses, allowing both parties to speak simultaneously.

Echo Cancellation reduces overall loop gain and the acoustic coupling between speaker and microphone. This coupling reduction prevents the annoying effect of hearing one’s own delayed speech, which is worsened when there is delay in the system, such as vocoder delay in digital cellular phones.

The CS6422 is a complete system implementation of a Digital Signal Processor with RAM and program ROM, running Echo Cancellation algorithms developed at Crystal Semiconductor using custom-er input, integrated with two delta-sigma codecs. The CS6422 is intended to provide a full-duplex

speakerphone solution with a minimum of design effort while displacing existing half-duplex speakerphone chips.

3. FUNCTIONAL DESCRIPTION

The CS6422 is divided into four external interface blocks. The analog interfaces connect the device to the transmit and receive paths. Control functions are accessible through the microcontroller interface. Two pins accommodate either a crystal or an externally applied digital clock signal. Analog and digital power and ground are provided through four pins.

3.1 Analog Interface

In a speakerphone application, one input of the CS6422 connects to the signal from the microphone, called the near-end or transmit input, and one output connects to the speaker. The output that leads to the speaker is called the near-end or receive output. Together, the input and output that connect to the microphone and speaker form the Acoustic Interface.

The signal received at the near-end input is passed to the far-end or transmit output after acoustic echo cancellation. This signal is sent to the telephone line. The signal from the telephone line is received at the far-end input, also called the receive input, and this signal is passed to the receive output after network echo cancellation. The far-end input and output form the Network Interface.

The analog interfaces are physically implemented using delta sigma converters running at an output word rate of 8 kHz, resulting in a passband from DC to 4 kHz. Because the inputs are analog to digital converters (ADCs), anti-aliasing and full-scale input voltage must be kept in mind. The ADCs expect a single-pole RC filter with a corner at 8 kHz, and they are post-compensated internally to prevent any resulting passband droop. The ADCs also expect a maximum of $0.9 V_{\text{rms}}$ ($2.5 V_{\text{pp}}$) at their inputs (which are biased around 2.12 VDC). A signal

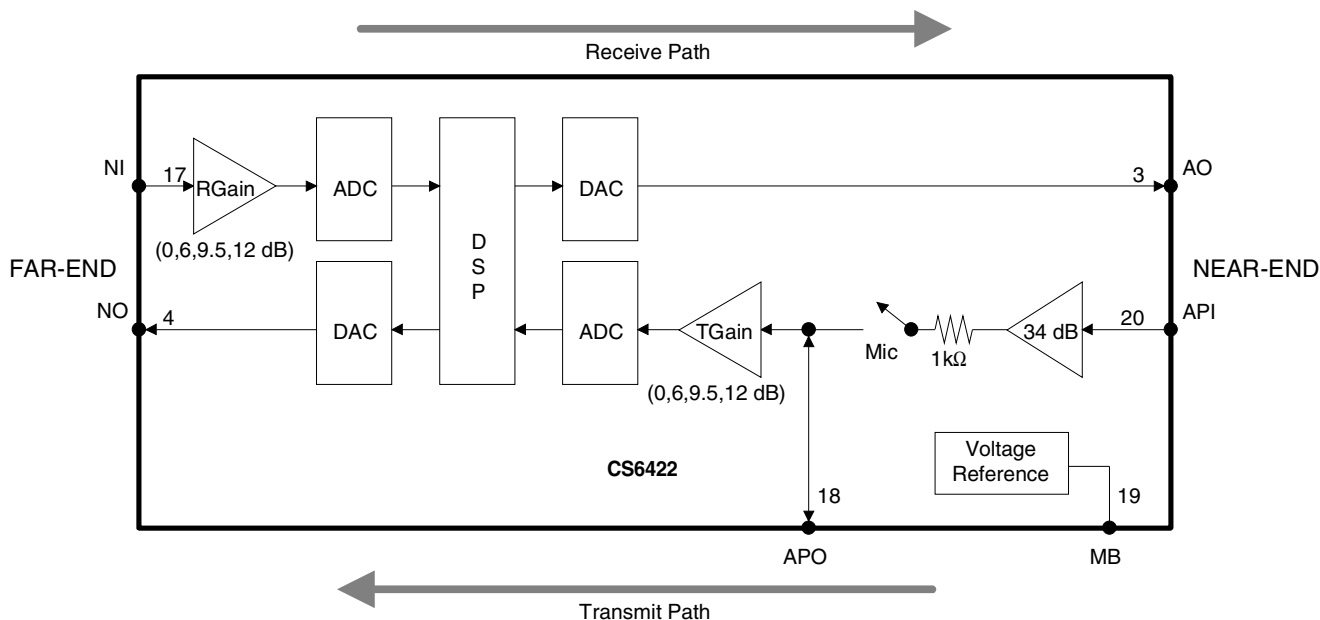


Figure 6. Analog Interface

of higher amplitude will clip the ADC input and will result in poor echo canceller performance. See Section 4., “Design Considerations” for more details.

The outputs are delta-sigma digital to analog converters (DACs) and have similar requirements to the ADCs. The DACs are pre-compensated to expect a single-pole RC filter with a corner frequency at 4 kHz. The full scale voltage output from a DAC is $1.1 V_{\text{rms}}$ ($3.1 V_{\text{pp}}$) maximum, $1 V_{\text{rms}}$ typical, biased around 2.12 VDC.

3.1.1 Acoustic Interface

The pins API (pin 20), APO (pin 18), AO (pin 3), and MB (pin 19) form the Acoustic Interface. A block diagram of the Acoustic Interface is shown in Figure 6.

API and APO are, respectively, the input and output of the built-in microphone pre-amplifier. The pre-amplifier is an inverting amplifier with a fixed

gain of 34 dB biased around an input offset voltage of 2.12 V. APO is the output of the pre-amplifier after a 1 kΩ (typical) resistor. The circuitry connected to the amplifier input must present low source impedance ($<100 \Omega$) to the API pin or the gain will be reduced. When using the internal mic preamp, a $0.022 \mu\text{F}$ capacitor should be placed between APO and ground to provide the anti-aliasing filter required by the ADC, as shown in Figure 4. The pre-amplifier may be bypassed by clearing the ‘Mic’ bit (Register 0, bit 15) using the Microcontroller Interface (see Section 3.2, “Microcontroller Interface”). If the internal mic preamp is not used, a $0.022 \mu\text{F}$ capacitor should be tied between API and ground, and APO should be driven directly. In this case, the signal into APO must be low-pass filtered by a single-pole RC filter with a corner frequency at 8 kHz (see Figure 5).

Following the pre-amplifier is a programmable analog gain stage, called TGain, which is controlled

through the Microcontroller Interface. This gain stage allows gains of 0 dB, 6 dB, 9.5 dB, and 12 dB to be added prior to the ADC input. The default gain stage setting is 0 dB.

The signal at APO should not exceed $2.5 V_{pp}$ at the 0 dB gain stage setting. If a different gain setting is used, then the full-scale signal at APO must also change. Table 1 shows full-scale voltages as measured at APO for the given programmable gain:

Gain Setting	Full-scale Voltage
0 dB	$2.5 V_{pp}$
6 dB	$1.25 V_{pp}$
9.5 dB	$0.84 V_{pp}$
12 dB	$0.63 V_{pp}$

Table 1. Full scale voltages for each gain stage

MB serves to provide decoupling for the internal voltage reference, and must have a $0.1 \mu F$ and a $10 \mu F$ capacitor to ground for bypass. *Noise on MB will strongly influence the overall analog performance of the CS6422.*

The acoustic output, AO, should connect to a single-pole low-pass RC network with a corner frequency of 4 kHz, which will filter out-of-band components. The full-scale voltage swing at AO is $3.1 V_{pp}$ maximum, $1 V_{rms}$ typical. AO is capable of driving a load of 10 k Ω or more.

3.1.2 Network Interface

The pins NI (pin 17) and NO (pin 4) form the Network Interface. The details of the Network Interface are shown in Figure 6.

NI is the input from the telephone network into the CS6422. The signal into NI must be low pass filtered by a single-pole RC filter with a corner frequency of 8 kHz.

RGain, a programmable analog gain stage accessible through the Microcontroller Interface, amplifies signals received at NI. This gain stage allows a gain of 0 dB, 6 dB, 9.5 dB, or 12 dB to be added

prior to the ADC input. The default gain stage setting for the network side is 0 dB.

The signal at NI should not exceed $2.5 V_{pp}$ at the 0 dB gain stage setting. If another gain setting is selected, then the full-scale signal at NI will change. Table 1 shows full-scale voltages as measured at NI for the given programmable gain.

The output to the telephone network side, NO, should connect to a single pole RC network with a corner frequency at 4 kHz, which will filter out-of-band components. The maximum swing NO is capable of producing is $3.1 V_{pp}$ maximum, $1 V_{rms}$ typical. NO is capable of driving a load of 10 k Ω or more.

3.2 Microcontroller Interface

The registers and control functions of the CS6422 are accessible through the Microcontroller Interface, which consists of three pins: DATA (pin 8), STROBE (pin 7), and \overline{DRDY} (pin 6). These inputs can connect to the outputs of a microcontroller to allow write-only access to the 16-bit Microcontroller Control Register (MCR).

3.2.1 Description

The Microcontroller Interface is implemented by a serial shift register that is clocked by STROBE and gated by \overline{DRDY} . The microcontroller begins the transaction by setting \overline{DRDY} low while STROBE is low. The most significant bit (MSB), Bit 15, of the 16-bit data word should be presented to the DATA pin and then STROBE should be brought high to shift the data bit into the CS6422. STROBE should be brought low again so it is ready to shift the next bit into the shift register. The next data bit should then be presented to the DATA pin ready to be latched by the rising edge of STROBE. This procedure repeats for all sixteen bits as shown in Figure 7. After the last bit (Bit 0) has been shifted in, \overline{DRDY} should be brought high to indicate the conclusion of the transfer, and *four or more extra*

STROBE pulses must be applied to latch the data into the CS6422.

Since the MCR is a shift register, the *STROBE* can be run arbitrarily slowly with a duty cycle limited only by the minimum high and low time specified in “Switching Characteristics”. *The Microcontroller Interface is polled at 125 μs intervals, so register writes must be spaced at least 125 μs apart or the register contents may be overwritten.*

3.2.2 Register Definitions

The six control registers accessible through the MCR are described in detail in the following tables. These registers are addressed by bits b3-0 of the MCR. *Bit ‘b0’ must always be ‘0’*. Table 2 shows the register map with the default settings. Tables 3 through 8 show the control registers in more detail.

The Register Map at the top of each register description shows the names of all the bits, with their reset values below the bitfield name. The reset value can also be found in the Word column of the bit-field summary as indicated by an ‘*’.




Figure 7. Microcontroller Interface

#	b15	b14	b13	b12	b11	b10	b9	b8	b7	b6	b5	b4	b3-0
0	Mic 1	HDD 0	GB 10			RVol 0100			TSD 0	ACC 00		TSMde 0	0000
1	THDet 00		Taps 10			TVol 1010			RSD 0	NCC 00		AuNECD 0	0010
2	RHDet 00		RSThd 00		NseRmp 00		HDly 00		HHold 0	TDSRmp 0	RDSRmp 0	IdITx 0	0100
3	TSAtt 00		PCSen 0		TDbtS 000		RDbtS 00		TSThd 00		TSBias 00		0110
4	AErle 00		AFNse 00		NErle 00		NFNse 00		RGain 00		TGain 00		1000
5	HwID 0	TD 0	APCD 0	NPCD 0	APFD 0	NPFD 0	AECD 0	NECD 0	ASdt 00		NSdt 00		1010

Table 2. MCR Control Register Mapping

3.3 Register 0

b15	b14	b13	b12	b11	b10	b9	b8	b7	b6	b5	b4	b3	b2	b1	b0
Mic	HDD	GB		RVol				TSD	ACC		TSMde	0	0	0	0
1	0	10		0100				0	00		0				
A				4				0				0			

Bits	Name	Function	Word	Operation
15	Mic	Microphone preamplifier enable	0 1*	disable mic preamp enable mic preamp
14	HDD	Half-Duplex Disable	0* 1	enable half-duplex disable half-duplex
13-12	GB	Graded Beta	00 01 10* 11	0.00 dB/ms 0.75 dB/ms 0.38 dB/ms 0.19 dB/ms
11-8	RVol	Rx Volume control	0000 0001 --- 0100* --- 1010 1011 --- 1101 1110 1111	+30 dB +27 dB --- +18 dB --- +0 dB -3 dB --- -9 dB -12 dB mute
7	TSD	Tx Suppression Disable	0* 1	enable Tx suppression disable Tx suppression
6-5	ACC	AEC Coefficient Control	00* 01 10 11	Normal Clear Freeze reserved
4	TSMde	Tx Suppression Mode	0* 1	enable noise guard disable noise guard

* Denotes reset value

Table 3. Register 0 Bit Definitions

3.3.1 MIC - MICROPHONE PREAMPLIFIER ENABLE

The microphone preamplifier described in Section 3.1.1, “Acoustic Interface” is enabled by default, but may be disabled by setting Mic to ‘0’. Refer to Section 3.1.1, “Acoustic Interface” for more details on using the Microphone Preamplifier.

3.3.2 HDD - HALF-DUPLEX DISABLE

In normal operation, the CS6422 will be in a half-duplex mode if the echo canceller is not providing enough loop gain reduction to prevent howling. This half-duplex mode is active at power-up while the adaptive filter begins to train. Half-duplex mode prevents howling and also masks the convergence process.

In some cases, such as when measuring convergence speed (see Section 4.3.2, “Testing Issues”), the half-duplex mode is undesirable. By default, the half-duplex mode is enabled.

3.3.3 GB - GRADED BETA

The room-size adjustment scheme called “graded beta,” provided for the acoustic echo canceller in the CS6422, is controlled by GB. The network echo canceller does not support graded beta.

Graded beta is an architectural enhancement to the CS6422 which takes advantage of the fact that acoustic echoes tend to decay exponentially with time. The CS6422 can increase the beta, or update gain, for the coefficients of the adaptive filter which occur earlier in time and decrease it for those that occur later in time, which increases convergence speed while maintaining stability. In order to make this improvement, there is an implicit assumption that the decay rate of the echo is known. The graded beta control allows the system designer to adjust this. For very acoustically live rooms, use either no decay (00) or slight decay (11). Cars and acoustically dead rooms can benefit from the most rapid decay (01).

3.3.4 RVOL - RECEIVE VOLUME CONTROL

Volume in the receive path is set by RVol. The volume control in the receive direction is implemented by a peak-limiting automatic gain control (AGC) and digital attenuation at the near-end output DAC.

The AGC is discussed in detail in Section 4., “Design Considerations”. See Section 4.1.3, “AGC” for a full explanation of how it functions.

When the reference level is set to +0 dB, the AGC is disabled. Volume control is implemented by digital attenuation in 3 dB steps from this point on down. The maximum gain is +30 dB and the minimum is -12 dB. The lowest gain setting (1111) mutes the receive path.

The default setting for RVol is +18 dB.

3.3.5 TSD - TRANSMIT SUPPRESSION DISABLE

The Transmit Supplementary Echo Suppression function is a non-linear echo control mechanism. Transmit Suppression introduces TSAtt (see Register 3) of attenuation into the transmit path when it is engaged. When TSMde = ‘1’, the transmit suppressor engages when there is speech detected in the receive path and no near-end speech is present. When TSMde = ‘0’, the default case, the transmit suppressor engages when there is no near-end speech present. When near-end speech is present, the suppression attenuation is removed. By default, the transmit suppression function is enabled.

3.3.6 ACC - ACOUSTIC COEFFICIENT CONTROL

The coefficients of the AEC adaptive filters in the CS6422 are controlled by ACC. The default position (00) yields normal operation, which means the coefficients are free to adjust themselves to the echo path in order to cancel echo. When set to the clear position (01), the adaptive filter coefficients are all held at zero, so the echo canceller is effectively disabled. Note that unless the half-duplex mode is disabled, this will force the CS6422 into half-duplex mode. The freeze position (10) causes the coefficients to retain their current values and not change.

3.3.7 TSMDE - TRANSMIT SUPPRESSION MODE

TSMde enables the Noise Guard feature of the CS6422. Noise Guard is a noise squelch feature that operates in the transmit path (from the near-end microphone to the far-end speaker). In traditional hands-free systems where the near-end talker is located in a noisy environment, the near-end system will remain in transmit mode and send that noise to the far-end listener. This creates a real problem if the listener is using a traditional half-duplex speakerphone because the far-end phone will stay in receive mode, thus preventing the far-end talker from being heard. Noise Guard eliminates this problem by squelching the transmit channel at the near-end unless near-end speech is detected, permitting the far-end speakerphone to switch normally during the conversation.

Noise Guard is also useful in cellular hands-free car applications because it prevents car noise from reaching the far-end while the near-end talker is silent.

Noise Guard is usually disabled when “half-duplex Idle return-to-Transmit” is enabled. See the Register 2 description for more information. Noise Guard is enabled by default.

3.4 Register 1

b15	b14	b13	b12	b11	b10	b9	b8	b7	b6	b5	b4	b3	b2	b1	b0
THDet		Taps		TVol				RSD	NCC		AuNECD	0	0	1	0
00		10		1010				0	00		0				
2				A				0				2			

Bits	Name	Function	Word	Operation
15-14	THDet	Tx Half-duplex Detection threshold	00* 01 10 11	6 dB 9 dB 12 dB reserved
13-12	Taps	AEC/NEC Tap allocation	00 01 10* 11	444/64 (55.5 ms/8 ms) 380/128 (47.5 ms/16 ms) 316/192 (39.5 ms/24 ms) 252/256 (31.5 ms/32 ms)
11-8	TVol	Tx Volume control	0000 0001 --- 0100 --- 1010* 1011 --- 1101 1110 1111	+30 dB +27 dB --- +18 dB --- +0 dB -3 dB --- -9 dB -12 dB mute
7	RSD	Rx Suppression Disable	0* 1	enable Rx suppression disable Rx suppression
6-5	NCC	NEC Coefficient Control	00* 01 10 11	Normal Clear Freeze reserved
4	AuNECD	Auto re-engage NEC Disable	0* 1	enable Auto NEC disable Auto NEC

* Denotes reset value

Table 4. Register 1 Bit Definitions

3.4.1 THDET - TRANSMIT HALF-DUPLEX DETECTION THRESHOLD

The sensitivity of the speech detector controls channel switching and ownership in half-duplex mode. The transmit speech detector registers speech if the transmit channel signal power is THDet above the noise floor of the transmit channel.

3.4.2 TAPS - AEC/NEC TAP ALLOCATION

The CS6422 has a total of 63.5 ms of echo canceller taps that it can partition for use by the network and acoustic echo cancellers. By default, the CS6422 allocates 39.5 ms for the AEC and 24 ms for the NEC. See NERle, NFNse, AERle, and AFNse in Register 4, and AECD and NECD in Register 5 for more options when an echo path is nonexistent.

3.4.3 TVOL - TRANSMIT VOLUME CONTROL

Volume in the transmit path is controlled by TVol. Like receive volume, the transmit volume is controlled by an AGC. See RVol in Register 0 for more details. The default setting for TVol is +0 dB.

3.4.4 RSD - RECEIVE SUPPRESSION DISABLE

The Receive Supplementary Echo Suppression function is a non-linear echo control mechanism. Supplementary Echo Suppression attenuates signals in the receive direction by 24 dB when far-end speech is absent in the receive path. The attenuation is released only when the receive channel is active. By default, the receive suppression function is enabled.

3.4.5 NCC - NETWORK COEFFICIENT CONTROL

The NEC adaptive filter's coefficients are controlled by NCC. See ACC in Register 0 for more details. The default setting for NCC is Normal mode.

3.4.6 AUNECD - AUTO RE-ENGAGE NEC DISABLE

AuNECD works in conjunction with NFNse in the determination of whether the Network Echo Cancellor should be enabled or disabled. If the CS6422 determines that a network coupling path does not exist and disables the NEC (which can occur only if NFNse is set to a non-zero value), then AuNECD allows the DSP to re-enable the NEC if at some point during the call a network path appears.

An example occurs in a digital PBX environment. Initially, a 4-wire 'intercom' call is placed between two stations. The CS6422 at the near-end determines that a network path is not present and disables the NEC. During the call, one of the stations conferences in a call from an external analog line. A network coupling path is introduced by the addition of the analog line due to the impedance mismatch at the 2-to-4 wire converter. If AuNECD is enabled, the CS6422 at the near-end will detect the presence of the network coupling path and re-enable the NEC automatically, drop to half-duplex, and re-train.

3.5 Register 2

b15	b14	b13	b12	b11	b10	b9	b8	b7	b6	b5	b4	b3	b2	b1	b0
RHDet		RSThd		NseRmp		HDly		HHold	TDSRmp	RDSRmp	IdITx	0	1	0	0
00		00		00		00		0	0	0	0				
0				0				0				4			

Bits	Name	Function	Word	Operation
15-14	RHDet	Rx Half-duplex Detection threshold	00* 01 10 11	6 dB 9 dB 12 dB reserved
13-12	RSThd	Rx Suppression Threshold	00* 01 10 11	6 dB 9 dB 12 dB reserved
11-10	NseRmp	Noise estimator Ramp rate	00* 01 10 11	3 dB/s 6 dB/s 12 dB/s reserved
9-8	HDly	half-duplex Holdover Delay	00* 01 10 11	200 ms 100 ms 150 ms reserved
7	HHold	Hold in half-duplex on Howl	0* 1	disable HHold enable HHold
6	TDSRmp	Tx Double-talk Suppression Ramp rate	0* 1	slow normal
5	RDSRmp	Rx Double-talk Suppression Ramp rate	0* 1	slow normal
4	IdITx	half-duplex Idle return-to-Transmit	0* 1	disable IdITx enable IdITx

* Denotes reset value

Table 5. Register 2 Bit Definitions

3.5.1 RHDET - RECEIVE HALF-DUPLEX DETECTION THRESHOLD

The sensitivity of the speech detector controls channel switching and ownership in half-duplex mode. The receive speech detector registers speech if the receive channel signal power is RHDet above the noise floor for the receive channel.

3.5.2 RSTHD - RECEIVE SUPPRESSION THRESHOLD

This parameter sets the threshold for far-end speech detection for disengaging receive suppression. The speech detector that disengages the receive suppression has its sensitivity controlled by RSThd. The suppression is inserted into the receive path unless signal from the far-end exceeds the receive channel noise power by RSThd, in which case speech is assumed to be detected and the suppression is defeated until speech is no longer detected. Decreasing RSThd to make the speech detector more sensitive could result in false detections due to spurious noise events which may cause an unpleasant noise modulation at the near-end. Increasing RSThd makes it robust to spurious noise, but may suppress weak far-end talkers. RSThd does not affect the ability of the receive suppressor to attenuate residual network echo.

3.5.3 NSERMP - NOISE ESTIMATOR RAMP RATE

The background noise power estimators increase at a programmable rate until the background noise power estimate equals the current input power estimate. The background noise power estimators quickly track drops in the current input power estimate. Choose large values of NseRmp if the environment is expected to have rapidly varying noise levels. Choose small values of NseRmp if the environment is expected to have relatively constant noise power.

3.5.4 HDLY - HALF-DUPLEX HOLDOVER DELAY

After a channel goes idle in the half-duplex mode of operation, a change of channel ownership is inhibited for HDly in order to prevent false switching due to echoes. The half-duplexor will be more immune to false switching if this delay is longer, but it will also prevent a fast response to legitimate channel changes. Short values of HDly mimic a more full-duplex like behavior, but may be susceptible to false switching due to echo.

3.5.5 HHOLD - HOLD IN HALF-DUPLEX ON HOWL

This is a control flag which, if enabled, holds the system in half-duplex when a howl event is detected. The system may transition to full-duplex if the flag is subsequently cleared. The default state of HHold is 'disabled', thus when a howl is detected, the CS6422 will temporarily drop into half-duplex, retrain, and transition back into full-duplex on its own.

3.5.6 TDSRMP - TX DOUBLE-TALK SUPPRESSION RAMP RATE

When "Tx Double-talk Suppression attenuation" (TDbtS, Register 3) is set to a non-zero value, the CS6422 will introduce a programmable amount of attenuation into the transmit path during a double-talk event, that is, when the near-end talker and far-end talker are speaking simultaneously. TDSRmp controls the decay rate of the transmit double-talk attenuation (the attack rate is ~40 ms).

The 'slow' setting of TDSRmp results in an attenuation decay rate of about 1 second. The 'normal' setting of TDSRmp results in an attenuation decay rate of about 100 ms.

3.5.7 RDSRMP - RX DOUBLE-TALK SUPPRESSION RAMP RATE

When “Rx Double-talk Suppression attenuation” (RDbtS, Register 3) is set to a non-zero value, the CS6422 will introduce a programmable amount of attenuation into the receive path during a double-talk event. RDSRmp controls the decay rate of the receive double-talk attenuation (the attack rate is ~40 ms).

The ‘slow’ setting of RDSRmp results in an attenuation decay rate of about 1 second. The ‘normal’ setting of RDSRmp results in an attenuation decay rate of about 100 ms.

3.5.8 IDLTX - HALF-DUPLEX IDLE RETURN-TO-TRANSMIT

When IdlTx is enabled, the CS6422’s half-duplex engine will automatically switch into <Transmit> mode from the <Idle> state. The <Idle> state is entered when the previously active channel has been silent for the time period set by HDly (half-duplex Holdover Delay) in Register 2.

The use of IdlTx permits a full-duplex-like behavior when operating in half-duplex at the beginning of a call. This benefit is most noticeable when the listener at the far end is using a handset.

When TSMde is set to ‘0’ (Noise Guard enabled), IdlTx is usually disabled. IdlTx is disabled by default.

3.6 Register 3

b15	b14	b13	b12	b11	b10	b9	b8	b7	b6	b5	b4	b3	b2	b1	b0
TSAtt		PCSen	TDbtS			RDbtS		TSThd		TSBias		0	1	1	0
00		0	000			00		00		00					
0				0				0				6			

Bits	Name	Function	Word	Operation
15-14	TSAtt	Tx Suppression Attenuation	00*	18 dB
			01	12 dB
			10	24 dB
			11	reserved
13	PCSen	Path Change Sensitivity	0*	high sensitivity
			1	low sensitivity
12-10	TDbtS	Tx Double-talk Suppression attenuation	000*	0 dB
			001	3 dB
			010	6 dB
		
			110	18 dB
9-8	RDbtS	Rx Double-talk Suppression attenuation	111	21 dB
			00*	0 dB
			01	3 dB
			10	6 dB
7-6	TSThd	Tx Suppression Threshold	11	9 dB
			00*	15 dB
			01	12 dB
			10	9 dB
5-4	TSBias	Tx Suppression Bias	11	18 dB
			00*	15 dB
			01	21 dB
			10	reserved

* Denotes reset value

Table 6. Register 3 Bit Definitions

3.6.1 TSATT - TRANSMIT SUPPRESSION ATTENUATION

This parameter sets the amount of attenuation inserted into the transmit path when transmit suppression is engaged.

3.6.2 PCSEN- PATH CHANGE SENSITIVITY

The Acoustic Interface is likely to have many path changes as people move about in the room where the full-duplex speakerphone is being used. The sensitivity of the path change detector can be changed with the PCSen bit. Set PCSen to '0' for high sensitivity and '1' for low sensitivity.

In any adaptive echo cancelling system, there is a trade-off between hearing echo and remaining in full-duplex when the acoustic path changes. When PCSen is set to '0' for high sensitivity, the CS6422 will tend to drop to half-duplex in the event of a path change, preventing the far-end listener from hearing echo as the adaptive filter adjusts to the new path.

When PCSen is set to '1' for low sensitivity, the CS6422 will tend to remain in full-duplex during the path change, and the far-end listener may hear some residual echo as the adaptive filter adjusts to the new path.

3.6.3 TDBTS - TX DOUBLE-TALK SUPPRESSION ATTENUATION

This parameter controls the amount of attenuation that is added to the transmit channel during double-talk, that is, when parties at both ends of the link are speaking simultaneously.

3.6.4 RDBTS - RX DOUBLE-TALK SUPPRESSION ATTENUATION

This parameter controls the amount of attenuation that is added to the receive path during double-talk.

3.6.5 TSTHD - TRANSMIT SUPPRESSION THRESHOLD

This parameter sets the ERLE requirement for discrimination between echo and near-end speech by the transmit suppressor. See Section 4.1.4.1, "Transmit Suppression" for full details.

3.6.6 TSBIAS - TRANSMIT SUPPRESSION BIAS

This bias level affects the ease with which near-end speech may break-in or be attenuated by far-end echo which causes the transmit suppressor to engage. See Section 4.1.4.1, "Transmit Suppression" for full details.

3.7 Register 4

b15	b14	b13	b12	b11	b10	b9	b8	b7	b6	b5	b4	b3	b2	b1	b0
AErle		AFNse		NErle		NFNse		RGain		TGain		1	0	0	0
00		00		00		00		00		00					
0				0				0				8			

Bits	Name	Function	Word	Operation
15-14	AErle	AEC Erle threshold	00* 01 10 11	24 dB 18 dB 30 dB reserved
13-12	AFNse	AEC Full-duplex Noise threshold	00* 01 10 11	zero -42 dB -54 dB reserved
11-10	NErle	NEC Erle threshold	00* 01 10 11	24 dB 18 dB 30 dB reserved
9-8	NFNse	NEC Full-duplex Noise threshold	00* 01 10 11	zero -42 dB -54 dB reserved
7-6	RGain	Rx analog Gain	00* 01 10 11	0 dB 6 dB 9.5 dB 12 dB
5-4	TGain	Tx analog Gain	00* 01 10 11	0 dB 6 dB 9.5 dB 12 dB

* Denotes reset value

Table 7. Register 4 Bit Definitions

3.7.1 AERLE - AEC ERLE THRESHOLD

The CS6422 will allow full-duplex operation when the ERLE provided by the AEC exceeds the value programmed at AERle. See also AFNse. See Section 6., “Glossary” for a definition of ERLE.

3.7.2 AFNSE - AEC FULL-DUPLEX NOISE THRESHOLD

AFNse works in conjunction with AERle to determine when the CS6422 should transition into full-duplex operation. AFNse specifies a noise level. If the current noise level at the near-end input is greater than AFNse, then AERle is used to determine if full-duplex is allowed, that is, the AEC must provide at least AERle of cancellation in order for the CS6422 to transition to full-duplex.

If the noise level is below AFNse, the CS6422 uses an internal estimate of asymptotic performance to determine whether or not to transition to full-duplex. If AFNse is zero, AERle is used as the exclusive full-duplex criterion.

3.7.3 NERLE - NEC ERLE THRESHOLD

The CS6422 will allow full-duplex operation only when the ERLE provided by the NEC exceeds the threshold set by NERle. See also NFNse. See Section 6., “Glossary” for a definition of ERLE.

3.7.4 NFNSE - NEC FULL-DUPLEX NOISE THRESHOLD

NFNse works in conjunction with NERle to determine when the CS6422 should transition into full-duplex operation. If the noise level at the far-end input is greater than NFNse, then NERle is used to determine if full-duplex is allowed. If the noise level is below the level of NFNse, the CS6422 uses an internal estimate of asymptotic performance to determine whether or not to transition to full-duplex. If NFNse is zero, NERle is always used as the exclusive full-duplex criterion.

If NFNse is non-zero, then the CS6422 will automatically disable the NEC if a network coupling path is not detected. Thus in systems in which the presence of a network path is not known, NFNse should be set to a non-zero value. See also AuNECD.

3.7.5 RGAIN - RECEIVE ANALOG GAIN

RGain selects the amount of additional on-chip analog gain to be supplied to the network input of the CS6422. The output of this amplifier stage feeds the receive path ADC, and can supply 0 dB, 6 dB, 9.5 dB, or 12 dB of gain to the signal path. The gain setting defaults to 0 dB.

Note: Changing the analog gain will change the full-scale voltage as applied to the input pin. Make sure that the ADC input does not clip with the gain stage on.3.

3.7.6 TGAIN - TRANSMIT ANALOG GAIN

3.8 Register 5

b15	b14	b13	b12	b11	b10	b9	b8	b7	b6	b5	b4	b3	b2	b1	b0
HwID	TD	APCD	NPCD	APFD	NPFD	AECD	NECD	ASdt		NSdt		1	0	1	0
0	0	0	0	0	0	0	0	0	0	0	0				
0				0				0				A			

Bits	Name	Function	Word	Operation
15	HwID	Howl detector Disable	0* 1	enable howl detector disable howl detector
14	TD	Tone detector Disable	0* 1	enable tone detector disable tone detector
13	APCD	Acoustic Path Change detector Disable	0* 1	enable PC detector disable PC detector
12	NPCD	Network Path Change detector Disable	0* 1	enable PC detector disable PC detector
11	APFD	Acoustic Pre-emphasis Filter Disable	0* 1	enable filter disable filter
10	NPFD	Network Pre-emphasis Filter Disable	0* 1	enable filter disable filter
9	AECD	Acoustic Echo Canceller Disable	0* 1	enable AEC disable AEC
8	NECD	Network Echo Canceller Disable	0* 1	enable NEC disable NEC
7-6	ASdt	Acoustic Sidetone level	00* 01 10 11	none -24 dB -18 dB -12 dB
5-4	NSdt	Network Sidetone level	00* 01 10 11	none -24 dB -18 dB -12 dB

* Denotes reset value

Table 8. Register 5 Bit Definitions

3.8.1 HWLD - HOWL DETECTOR DISABLE

This is a diagnostic parameter that is normally set to '0'.

In normal operation, the CS6422 will clear both the AEC and NEC coefficients, dropping the device into half-duplex operation, whenever an instability event is detected. Such an event can be caused by excessive loop gain, a major path change, or mistraining of the echo cancellers.

Setting HwLD to '1' prevents the instability detector from clearing the echo cancellers' coefficients.

3.8.2 TD - TONE DETECTOR DISABLE

This is a diagnostic parameter that is normally set to '0'.

In normal operation, the tone detector responds to the detection of tones in the receive path. If the CS6422 is in half-duplex mode, the tone detector will clear the AEC coefficients and force the half-duplex engine into <Receive> mode to allow the tone to pass through, independent of the presence of signals at the near-end microphone.

If the CS6422 is in full-duplex mode when a tone is detected, the tone detector will momentarily freeze the AEC coefficients to prevent false training.

3.8.3 APCD - ACOUSTIC PATH CHANGE DETECTOR DISABLE

This diagnostic bit is normally set to '0'.

The purpose of the acoustic path change detector is to respond to drastic path changes by clearing the AEC coefficients to facilitate rapid and accurate convergence to the new path.

Disabling the acoustic path change detector prevents it from clearing the AEC coefficients, thus forcing the filter to 'adapt out' of the path change, which typically takes longer and is less accurate than adapting from a cleared state.

3.8.4 NPCD - NETWORK PATH CHANGE DETECTOR DISABLE

This diagnostic bit is normally set to '0'.

The purpose of the network path change detector is to respond to drastic path changes by clearing the NEC coefficients to facilitate rapid and accurate convergence to the new path.

Disabling the network path change detector prevents it from clearing the NEC coefficients, thus forcing the filter to 'adapt out' of the path change, which typically takes longer and is less accurate than adapting from a cleared state.

3.8.5 APFD/NPFD - ACOUSTIC PRE-EMPHASIS FILTER DISABLE/NETWORK PRE-EMPHASIS FILTER DISABLE

These diagnostic bits are normally set to '0'.

The pre-emphasis filter helps the adaptive filter correctly model the coupling path by attenuating lower frequency information. This is done because high-frequency information more accurately describes the echo path, that is, low frequency information is more spatially ambiguous.

Sometimes it is useful to disable the pre-emphasis filter when performing ERLE tests using white noise, since the filter will tend to prevent the adaptive filter from cancelling the low frequency components of the signal, resulting in artificially low ERLE measurements.

3.8.6 AECD - ACOUSTIC ECHO CANCELLER DISABLE

Setting this bit to a '1' disables the Acoustic Echo Canceller. The AEC is removed from the signal path and is not considered in the half/full-duplex decision making process.

3.8.7 NECD - NETWORK ECHO CANCELLER DISABLE

Setting this bit to a '1' disables the Network Echo Canceller. The NEC is removed from the signal path and is not considered in the half/full-duplex decision making process.

3.8.8 ASDT - ACOUSTIC SIDETONE LEVEL

This control allows the introduction of a linear coupling path for the AEC to train on. The real acoustic path is superimposed on this path and both are cancelled by the AEC.

The use of an acoustic sidetone is beneficial in environments where the real acoustic path may be highly variable, faint, or distorted, such as in hands-free automotive applications. This control is usually set to 'none'.

3.8.9 NSDT - NETWORK SIDETONE LEVEL

This control allows the introduction of a linear coupling path for the NEC to train on. The real network path is superimposed on this path and both are cancelled by the NEC.

The use of a network sidetone is beneficial in environments where the real network path is faint or distorted. This control is usually set to 'none'.

3.9 Reset

A hardware reset, initiated by bringing \overline{RST} low for at least t_{RSTL} and then high again, must be applied after initial power-on.

When \overline{RST} is held low, the various internal blocks of the CS6422 are powered down. When \overline{RST} is brought high, the oscillator is enabled and approximately 4 ms later, all digital clocks begin operating. The ADCs and DACs are calibrated and all internal digital initializations occur.

The CS6422 supports two reset modes, cold reset and warm reset. The reset mode is selected by completing a write of a specified value to the MCR within T_{wRST} of the rising edge of \overline{RST} . If no writes to the MCR occur within T_{cRST} , then a cold reset is initiated by default at the end of the T_{cRST} time period.

The value written to the MCR determines the behavior of the CS6422:

- 1) a value of '0x0000' will initiate a cold reset when the reset timer expires. This is the default behavior of the device.
- 2) a value of '0x0006' will initiate a warm reset when the reset timer expires.
- 3) a value of '0x8000' will initiate a cold reset immediately, bypassing the reset timer.
- 4) a value of '0x8006' will initiate a warm reset immediately, bypassing the reset timer.

Values (#2) through (#4) above are interpreted as legitimate register writes (to register 0 for (#3) and to register 3 for (#2) and (#4)) of the CS6422. Therefore, it is important to follow the first register write with another write containing the proper settings for register 0 or register 3.

3.9.1 Cold Reset

Cold reset initializes all the components of the CS6422. The ADCs and DACs are reset, the echo canceller memories and registers are cleared, and the default settings of the MCR are restored.

3.9.2 Warm Reset

Warm reset is like cold reset except that the echo canceller coefficients and certain key variables are not cleared, but instead keep their pre-reset value. This gives the CS6422 a headstart in adapting to its environment if the echo environment is relatively stable, assuming a cold reset occurred at least once since power up.

3.9.3 Reset Timer

Another special reset option is to exit the T_{wRST} reset timer before the T_{wRST} has elapsed. This timer halts device operation until the analog bias voltages have had time to settle. The early-exit option should be used only in applications in which the T_{wRST} start-up delay is unacceptable.

3.10 Clocking

The clock for the converters and DSP is provided via the clocking pins, CLKI (pin 14) and CLKO (pin 13). A 20.480 MHz parallel resonant crystal placed between these two pins and loaded with 22 pF capacitors will allow the on-chip oscillator to provide this system clock. Alternatively, the CLKI pin may be driven by a CMOS level clock signal. The clock may vary from 20.480 MHz by up to 10%, however, this will change the sampling rate of the converters and echo canceller, which will affect the bandwidth of the analog signals and the duration of echo that the echo canceller can accommodate. CLKO is not connected when CLKI is driven by the CMOS signal.

3.11 Power Supply

The pins AVDD (pin 1) and AGND (pin 2) power the analog sections of the CS6422, and DVDD (pin 16) and DGND (pin 15) power the digital sections. This distinction is important because internal to the part, the digital power supply is likely to contain high-frequency energy. The analog power supply is kept clean internally by drawing current from a different pin, thereby achieving high performance in the codecs and the microphone preamplifier.

The digital supply of the CS6422 should not be connected to the system digital supply, if there is one, as the CS6422 has internal timing mechanisms designed to minimize the detrimental effects of its own digital noise, but cannot use these to compensate for externally introduced digital noise. The CS6422 digital power supply should be derived from its analog power supply through a ferrite bead with low ($< 1 \Omega$) DC impedance.

3.11.1 Power Down Mode

Typical power consumption of the CS6422 is 60 mA, assuming normal operating conditions. This current consumption can be further reduced by invoking the powerdown mode, which is entered by holding $\overline{\text{RST}}$ low. Holding $\overline{\text{RST}}$ low will power down all the internal blocks of the CS6422 and stop the oscillator. In powerdown mode, current consumption drops to less than 2 mA.

3.11.2 Noise and Grounding

Since the CS6422 is a mixed-signal integrated circuit, the system designer must pay special attention to layout and decoupling to minimize noise coupling. The three best methods to reduce noise when using the CS6422 are to properly decouple the power supplies, to separate the system analog and digital power and ground (all power and ground pins of the CS6422 should tie to the analog power supply), and to route signals on the board carefully.




Figure 8. Suggested Layout

Figure 8 shows the suggested placement of decoupling capacitors for the power supplies. Note that the trace length from the power pin to the capacitors is minimized. Also note that the smaller valued capacitor is placed closer to the pin than the larger valued capacitor. The smaller capacitor decouples high frequency noise and the larger capacitor attenuates lower frequencies.

The separation of analog and digital power and ground is done in two ways. The power is separated by deriving the digital power for the CS6422 from the analog through a ferrite bead to isolate analog

from digital, as shown in Figure 9. The ferrite bead serves as a low-pass filter to remove CS6422 digital switching noise from the analog power supply. The ground is separated by isolating all the digital components of the system board on one ground plane and all the analog and linear components on a different ground plane. The CS6422 should be placed over the analog ground plane. This prevents digital switching noise from the digital components of the board from coupling into the converters and aliasing into the passband.

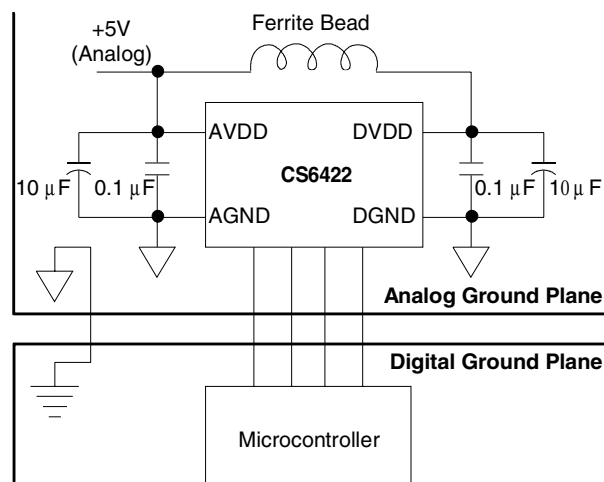


Figure 9. Ground Planes

4. DESIGN CONSIDERATIONS

When designing the CS6422 into a system, it is important to keep several considerations in mind. These concerns can be loosely grouped into three categories: algorithmic considerations, circuit design considerations, and system design considerations.

4.1 Algorithmic Considerations

The CS6422 facilitates full-duplex hands-free communication via many algorithms running on the Digital Signal Processor that is the core of the CS6422. Among these are the algorithms that perform the adaptive filtering, the half-duplex switching, digital volume control, and supplementary echo suppression.

4.1.1 Full-Duplex Mode

Full-duplex hands-free communication is achieved through a technique called adaptive filtering. The basic principle behind adaptive filtering is that the acoustic path between speaker and microphone can be modeled by a transfer function which can be dynamically determined by an adaptive digital filter. This principle assumes good update control and speech/tone detection algorithms to prevent the filter from mistraining.

4.1.1.1 Theory of Operation

Figure 10 illustrates how the adaptive filter can cancel echo and reduce loop gain. The echo path of the system is between points B and C: the speaker to microphone coupling. A signal injected at A (sometimes called a “training signal”) is sent both to B, the input of the echo path, and to F, the input of the adaptive filter. The signal at B is modified by the acoustic transducers (speaker and microphone) and the environment, and received at point C (as an “Echo”). Meanwhile, assume that the adaptive filter has exactly the right transfer function to match the echo path BC, and so the signal at point D is approximately equal to the signal at point C. After these are subtracted by the summing element, all that is left is the error signal at point E, which should be very small.

If a person were to speak into the microphone at point C, that signal would pass through the summing element unchanged because the adaptive filter had no comparable input to subtract out. In this manner, the person at A and the person at C may simultaneously speak and A will not hear his own echo.

In the real world, the echo path is not static. It will change, for example, when people move in the

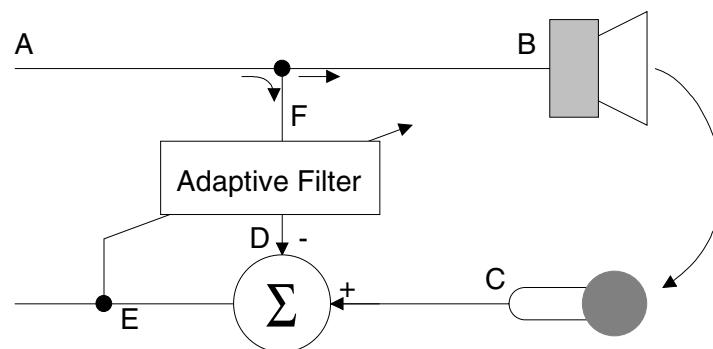


Figure 10. Simplified Acoustic Echo Canceller Block Diagram

room, when someone moves the speaker or the microphone, or when someone drops a piece of paper on top of the speaker. So, the filter needs to adapt to modify its transfer function to match that of the environment. It does so by measuring the error signal at point E and trying to minimize it. This signal is fed back to the adaptive filter to measure performance and how best to adapt, or train.

The trouble arises when the person at the near-end (C) speaks: the error signal will be non-zero, but the adaptive filter should not change. If it tries to train to the near-end signal, the adaptive filter has no way to reduce the error signal, because there is no input to the filter, and therefore no output from it. The adaptive filter would mistrain.

To prevent this mistraining, the echo canceller uses double-talk detection algorithms to determine when to update. These update control algorithms are the heart of most echo canceller implementations.

The worst case situation for the CS6422 is when parties at both ends are speaking and the person at the near-end is moving. In this case, the echo canceller will cease to adapt because of the double-talk, but the echo will not be optimally reduced because of the change in path.

4.1.1.2 Adaptive Filter

The adaptive filter in the CS6422 uses an algorithm called the “Normalized Least-Mean-Square (NLMS)” update algorithm to learn the echo path transfer function. This Finite Impulse Response (FIR) filter has 508 taps, which can model up to 63.5 ms of total path response at a sampling rate of 8kHz. The coverage time is calculated by the following formula:

$$\left(\frac{1}{8\text{kHz}}\right) \times 508 = 63.5 \text{ ms.}$$

The CS6422’s adaptive filter, like all FIR filters, only models Linear and Time Invariant (LTI) sys-

tems. So, any non-linearity in the echo path can not be modeled by the adaptive filter and the resulting signals will not be cancelled. Signal clipping and poor-quality speakers are very common sources of non-linearity and distortion.

A common integration problem for echo cancellers is signal clipping in the echo path. For example, if a speaker driver is driven to its rails, the distortion of the speech may be hard to perceive, but it is very bad for the echo canceller. The technique of over-driving the speaker has been used in half-duplex phones to provide good low-level signal gain at the expense of distortion with high amplitude signals. Since this does not work for the CS6422, an AGC mechanism has been introduced to provide equivalent behavior without clipping. See Section 4.1.3, “AGC” for more details.

Another common problem is speaker quality. A poor quality speaker which is perfectly acceptable for a half-duplex speakerphone, may limit the echo canceller’s performance in a full-duplex speakerphone. The distortion elements are not modeled by the adaptive filter and so limit its effectiveness. Speakers should have better than 2% THD performance to not impede the adaptive filter.

Volume control should be implemented using the CS6422 Microcontroller Interface. A real-time external change in the gain of the speaker driver results in a change in the transfer function of the echo path, and will force the adaptive filter to readapt. If the volume control is done before the input to the adaptive filter, the echo path does not change, and retraining is not necessary. Another side benefit of the CS6422 volume control is that it transparently provides dynamic range compression through the AGC function.

4.1.1.2.1 Pre-Emphasis

The typical training signal for the adaptive filter is speech, but most adaptive filters train optimally with white noise. Speech has very different spectral

characteristics than white noise because of its quasi-periodic nature.

Research at Crystal has shown that quasi-periodic signals cause the formation of spurious non-zero coefficients within the adaptive filter at tap intervals determined by the periodicity of the signal. This results in small changes in period being very destructive to the adaptive filter's performance.

One mechanism the CS6422 uses to prevent this filter corruption with speech is to pre-emphasize the signal sent to the adaptive filter so that much of the low frequency content is removed.

The CS6422 works very well with a speech training signal because of the pre-emphasis filter. White noise training signals, however, result in sub-optimal performance, so when testing with white noise, it is recommended that the pre-emphasis filters be disabled.

4.1.1.2.2 Graded Beta

The update gain of an adaptive filter, sometimes called the “beta”, is the rate at which the filter coefficients can change. If beta is too low, the adaptive filter will be slow to adapt. Conversely, if it is too high, the filter will be unstable and will create unwanted noise in the system.

In most echo canceller implementations, the beta is a fixed value for all the filter coefficients. In some situations, though, through knowledge of the characteristics of echo path response, the beta can be varied for groups of coefficients. This preserves stability by allowing the beta to be higher for some coefficients and compensating by reducing beta below nominal for others.

For example, acoustic echo tends to decay exponentially, so the first taps need to be larger than the later taps. Having a beta profile that matches the expected response path enhances the echo canceller's ability to correctly and accurately model the acoustic path. Furthermore, this has an added ben-

efit of suppressing the spurious taps mentioned in Section 4.1.1.2.1, “Pre-Emphasis”.

The Microcontroller Interface allows four settings for graded beta: none, 0.19 dB/ms, 0.38 dB/ms, and 0.75 dB/ms. Use 0.75 dB/ms for acoustically dead rooms or cars, and 0.19 dB/ms or no grading of beta for large, or acoustically live rooms.

4.1.1.3 Update Control

As mentioned in Section 4.1.1.1, “Theory of Operation”, the update control algorithms are the heart of any useful echo canceller implementation. Aside from telling the adaptive filter when to adapt, they are responsible for correcting performance when the path changes more quickly than the filter can respond. For example, if the adaptive filter is actually adding signal power instead of cancelling it, the update control algorithms will reset the adaptive filter to cleared coefficients, forcing it to restart.

4.1.1.4 Speech Detection

The CS6422 detects speech by using power estimators to track deviations from a background noise power level. The power estimators filter and average the raw incoming samples from the ADC.

A background noise level is established by a register that increases 3 dB at intervals determined by NseRmp (Register 2, bits 11 and 10). When the power estimator level rises, the background noise level will slowly increase to try to match it. When the power estimator level is below the background noise level, the background noise level adjusts quickly to match the power estimator level. This method allows significant flexibility in tracking the background noise level.

Speech is detected when the power estimator level rises above the background noise level by a given threshold. The half-duplex receive speech detector threshold is set by RHDet (Register 2, bits 15 and 14), the half-duplex transmit speech detector threshold is set by THDet (Register 1, bits 15 and

14), and the receive suppression speech detector threshold is set by RSThd (Register 2, bits 13 and 12). The transmit speech detectors for both half-duplex and suppression default to 6 dB.

Note that constant power signals which persist for long durations, such as tones or white noise from a signal generator, will be detected as speech only as long as the background noise level has not risen to within the speech detection threshold of the signal power. When a tone has persisted for long enough, the background noise level will be equal to the power estimator level, and so the tone will no longer be considered speech. This duration is dependent upon the power difference between the signal and the ambient noise power, as well as NseRmp. It should be noted that the CS6422 has a tone detector to prevent updates when tones are present and allow tones to persist regardless of the speech detectors.

4.1.2 Half-Duplex Mode

In cases where the system relies on the echo canceller for stability, a fail-safe mechanism must be in place for instances when the echo canceller is not performing adequately. The CS6422 implements a half-duplex mode to guarantee communication even when the echo canceller is disabled.

When the CS6422 is first powered on, or emerges from a reset, the echo canceller coefficients are cleared, and the echo cancellers provide no benefit at this point. The half-duplex mode is on to prevent howling and echo from interfering with communication. Once the CS6422's adaptive filters have adapted sufficiently, the half-duplex mode is automatically disabled, and full-duplex communication can occur.

The half-duplex mode allows three states: <Transmit>, <Receive>, and <Idle>. In the <Transmit> state, the transmit channel is open and the receive channel is muted. The <Receive> state mutes the transmit channel. The <Idle> state is an internal state which is used to enhance switching decision

making. The CS6422 must be <Idle> before it will allow a state change between <Transmit> and <Receive>.

The half-duplex controller can be susceptible to echo, so a holdover timer is provided to help prevent false switching. Holdover will force the channel to remain in its current state for a fixed duration after speech has stopped. HDly (Register 2, bits 9 and 8) sets the duration of the holdover. Longer holdover will tend to make interrupting more difficult, but will be more robust to spurious switching caused by echo.

4.1.2.1 Idle Return to Transmit

When enabled, this feature causes the CS6422 to return to <Transmit> mode from an <Idle> state when operating in half-duplex. This simulates full-duplex-like behavior during the periods of half-duplex operation at the beginning of a call.

4.1.3 AGC

The CS6422 implements a peak-limiting AGC in both the transmit and receive directions in order to boost low-level signals without compromising performance when high amplitude signals are present. The technique effectively results in dynamic range compression.

The AGC works by setting a reference level based on the value represented by TVol (Register 1, bits 11-8) for the transmit direction and RVol (Register 0, bits 11-8) for the receive direction. If the signal from the input is above this reference, it is attenuated to the reference level with an attack time of 125 μ s. This attenuation level decays with a time constant of 30 ms unless another signal greater than the reference level is detected. After the attenuation, a post-scaler scales the reference level to full-scale (the maximum digital code), which amplifies all signals by the difference between the reference level and full-scale.

For example, Figure 11 shows how the AGC works with a reference level of +30 dB (Word = 0000).

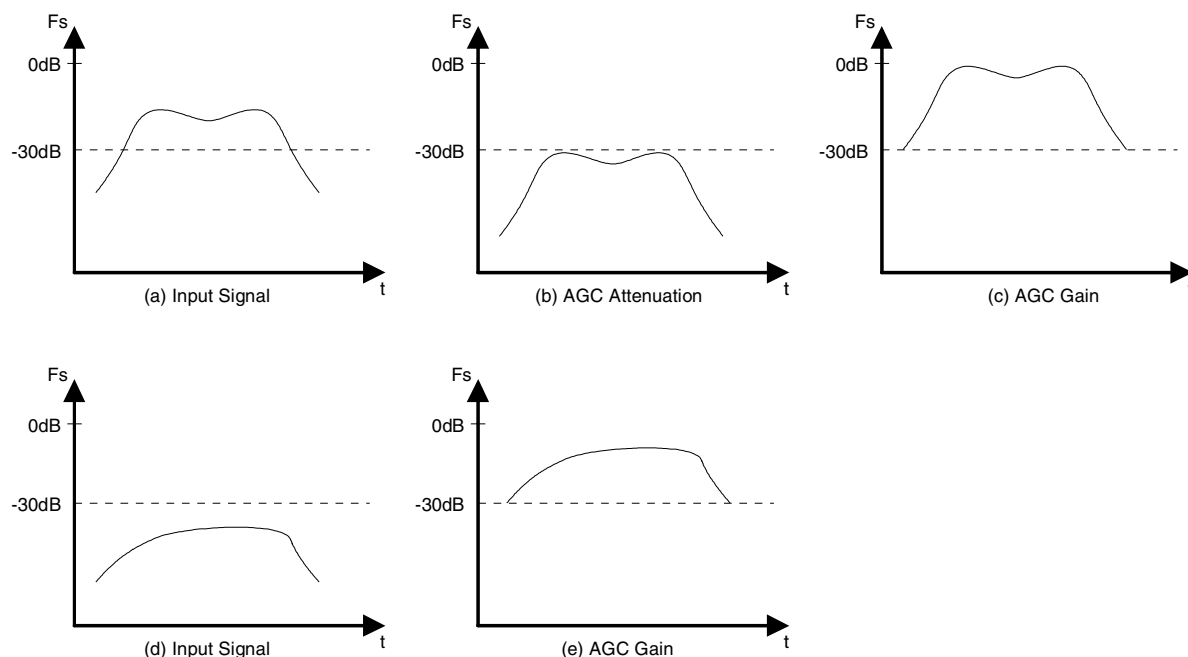


Figure 11. How the AGC works (TVol = +30 dB)

Any signal greater than 30 dB below full-scale (a), is scaled down to 30 dB (b). This signal is then scaled up +30 dB (the reference level) to provide the final output (c). Note that the combination of attenuation and gain results in less than +30 dB total gain being applied. If the input signal is below 30 dB below full-scale (d), no attenuation is added and the full +30 dB of gain is applied to the signal (e).

When the reference level is set to +0 dB, the AGC is effectively disabled. Volume control is implemented by digital attenuation in 3 dB steps from this point on down. The maximum gain is +30 dB, and the minimum is -12 dB in 3 dB steps. The lowest gain setting (1111) mutes the signal path. The signal scaling takes place in between the two cancellers, and so does not disturb the echo canceller as changing gain in the echo path, for example at the speaker driver, would (see Section 4.1.1.2, “Adaptive Filter” for more details).

4.1.4 Suppression

Echo cancellation is somewhat of a misnomer in that echo is merely attenuated, not entirely cancelled. Some residual echo still exists after the summing node. This residual echo, though low in amplitude, may be audible when the near-end talker is not speaking. Suppression further attenuates the echoed signal, preventing the far-end listener from hearing echo.

The CS6422 employs supplementary echo suppression which adds attenuation on top of the cancellation to remove the residual echo. For example, the CS6422 will engage extra attenuation in the transmit path whenever only the far-end talker is speaking. However, if the near-end talker starts speaking, this attenuation is removed and the system relies on the near-end talker’s speech to mask the residual echo.

Suppression may cause some modulation of the perceived background noise which may be distracting to some users. As a result, it may be desirable to limit the suppression attenuation to the minimum necessary. The CS6422 provides TSAtt (Register 3, bits 15 and 14) to control the amount of attenuation introduced by suppression in the transmit channel. Receive suppression attenuates by 24 dB.

4.1.4.1 *Transmit Suppression*

When TSMde = '1' (Noise Guard 'off'), the transmit suppressor attenuates the transmit path when only far-end speech is present. When TSMde = '0' (Noise Guard 'on'), the suppressor attenuates when the transmit channel is idle, that is, when no near-end speech is present.

The purpose of Transmit Suppression is to mask residual echo by inserting additional loss/attenuation in the transmit path in the scenario when far-end speech is present; the residual echo, if any, in double-talk is masked by near-end speech, assuming reasonable levels of ERLE.

There are four controls that govern the behavior of Transmit Suppression. These are TSThd (Register 3, bits 7 and 6), TSAtt (Register 3, bits 15 and 14), TSBias (Register 3, bits 5 and 4), and TSMde (Register 0, bit 4). TSThd is the primary control and should be adjusted before changing the value of TSBias from its default setting. TSThd sets the ERLE expectation to be used in discriminating between near-end speech and far-end echo. This control setting will by far predominate in affecting the manner in which Transmit Suppression behaves.

TSAtt controls the amount of attenuation added to the transmit path when the transmit suppressor engages.

TSBias is a secondary control. This is to be adjusted after the system designer is more or less satisfied with the behavior of Transmit Suppression with the TSThd set. It affects the ease with which a near-end talker may disengage Transmit Suppression and

keep it disengaged. We recommend using larger values of TSBias relative to TSThd settings in order to facilitate ease of near-end speech transmission. For example, the default setting for TSThd is 15 dB and 18 dB for TSBias.

In some scenarios, especially when the dynamic range of volume control is significantly large, we also recommend the use of different combinations of TSThd and TSBias setting relative to output volume of the acoustic interface. Specifically, higher volume levels may call for larger values of TSThd.

TSMde controls the Noise Guard feature. When TSMde = '0' (Noise Guard enabled), the transmit suppressor is engaged when no near-end speech is present. When TSMde = '1' (Noise Guard disabled), the transmit suppressor is engaged only when far-end speech is present in the absence of near-end speech.

4.1.4.2 *Receive Suppression*

The receive suppressor is nominally attenuating unless far-end speech is present. This behavior is more consistent with behavior observed in modern speakerphones, and helps keep noise levels low.

One side effect of this scheme is that a constant power signal, such as noise from a noise generator or a tone, will eventually be attenuated when the background noise level estimate turns off the receive suppression speech detector. See Section 4.1.1.4, "Speech Detection" for more details.

RSThd (Register 2, bits 13 and 12) sets the speech detection threshold of the suppressor's speech detector. This control is normally set to the same value as RHDet. See Section 4.1.1.4, "Speech Detection" for more details.

4.1.4.3 *Double-talk Attenuation*

In full-duplex hands-free to full-duplex hands-free scenarios (where a call exists between two full-duplex speakerphones), stability problems can arise at higher volume levels due to the acoustic coupling

loop (near-end speaker/mic to far-end speaker/mic) during a double-talk scenario (where both near-end and far-end parties are talking at the same time).

The CS6422 implements an optional attenuation feature that introduces a programmable amount of loss in the transmit and/or receive directions during double-talk to alleviate stability concerns without sacrificing speaker volume. This allows for higher speaker volume levels at both ends of the call without compromising stability.

4.1.4.4 Noise Guard

Noise Guard is an optional noise squelch feature that operates in the transmit path (near-end microphone to far-end speaker). In traditional systems, if the near-end talker is located in a noisy environment, the near-end system will remain in transmit mode and transmit that noise to the far-end listener. While this may be bothersome to the far-end listener using a standard handset, this creates a real problem if the listener is using a traditional half-duplex speakerphone because the far-end phone may stay in receive mode and not allow the far-end talker to be heard. Noise guard eliminates this problem by squelching the transmit channel at the near-end unless near-end speech is detected, permitting the far-end speakerphone to switch normally during the conversation.

4.2 Circuit Design

The design of the CS6422 interface circuitry plays an important role in achieving optimum performance. The actual circuit design is important, especially the analog interface. Proper grounding and layout will help minimize the noise that might get coupled into the CS6422.

4.2.1 Interface Considerations

Of the CS6422 interfaces, the analog interface and the microcontroller interface are the most important to pay special attention to during circuit design. The analog interface especially will

determine how well the echo canceller can perform.

4.2.1.1 Analog Interface

The Analog Interface feeds information about the echo path to the adaptive filter, so it is critical that this interface be well designed. Using high-quality transducers and circuits that guarantee low-distortion and minimal clipping are essential to the success of any echo canceller based design.

As mentioned in Section 4.1.1.2, “Adaptive Filter”, the adaptive filter assumes that the echo path is linear and time-invariant. As such, poor quality speakers are a common cause of poor echo canceller performance due to their high distortion. Speakers must be selected with their linearity in mind. In general, the speaker should have less than 2% Total Harmonic Distortion (THD). This will result in distortion terms 34 dB below the desired signal, enough headroom for the echo canceller to function adequately.

The other major consideration in the design of the analog interface is that the circuitry that processes the transducer signals not clip or distort it. For example, a common problem is the use of a speaker amplifier with a fixed gain, which clips when driving the speaker. Although the distortion may not be objectionable to the human ear, it will prevent the adaptive filter from modeling the path correctly. Speakers and microphones which worked for half-duplex speakerphones will not necessarily work for full-duplex speakerphones. Microphone amplifier circuitry is also suspect when looking for sources of clipping and distortion.

4.2.1.2 Microcontroller Interface

The Microcontroller Interface is the only asynchronous digital connection to the CS6422, so it is the most likely place for digital noise coupling to be a problem. The interface itself is fairly straightforward and requires only three pins from a microcontroller.

The three pins that comprise the Microcontroller Interface are STROBE, DATA, and $\overline{\text{DRDY}}$. Also, four extra clocks are required after $\overline{\text{DRDY}}$ is brought high in order to latch the data into the CS6422, as is shown in Figure 7.

4.2.2 Grounding Considerations

Proper grounding of the CS6422 is necessary for optimal performance from this mixed-signal device. The CS6422 should be considered an analog device for grounding purposes.

The digital sections of the CS6422 are synchronized with its ADCs and DACs to minimize the effects of digital noise coupling. However, for external digital devices that are asynchronous with respect to the CS6422, precautions should be taken to minimize the chances of digital noise coupling into the CS6422.

A design with the CS6422 should have a separate ground plane for any digital devices. For example, a system microcontroller should be on a digital ground plane with its control lines leading to the CS6422 in the shortest reasonable distance. The CS6422 itself should lie completely on the analog ground plane.

4.2.3 Layout Considerations

The physical layout of the traces and components around the CS6422 will also strongly affect the performance of the device. Special attention must be paid to decoupling capacitors, the crystal oscillator, and the input anti-aliasing filters.

The decoupling capacitors for the power supplies of the CS6422 should be placed as close as possible to the power pins for best performance. There are two capacitors per pin: the 0.1 μF capacitor needs to be closest to the pin to decouple the high frequency components, and the larger cap can be farther away. The MB pin is the most critical as it connects directly to the on-chip voltage reference. AVDD and DVDD are secondary to MB with respect to priority.

The crystal oscillator should be placed as close as possible to reduce the distance that the high frequency signals must travel. If the crystal is placed too far away, the trace inductance may cause problems with oscillator startup.

The next concern with placement is the input anti-aliasing filters for the ADC inputs. NI has an RC low-pass network with a corner frequency of 8 kHz. The capacitor of this low-pass network should be placed very close to the pin so that there is very little exposed trace to pick up noise. If the on-chip microphone amplifier is used, the 0.022 μF capacitor on APO will provide the appropriate cut-off frequency, and so should be placed close to the APO pin. If the on-board preamplifier is not used, APO will have the same RC network as NI, and should be treated similarly.

The connections from the controller to the Microcontroller Interface should be short straight traces, if possible. The traces should not run very close to any digital clocks to avoid cross coupling.

4.3 System Design

The CS6422 is ultimately only one part of a bigger full-duplex hands-free system. In order for that system to work well, it needs to be properly balanced. The distribution of the system gains will make or break the echo canceller. In order to judge performance, however, the system integrator must be armed with the means to test the product.

4.3.1 Gain Structure

The distribution of the system gains is an important design consideration to keep in mind. Gain distribution is an intricate balancing act where the system integrator tries to maximize dynamic range while minimizing noise, and at the same time, getting excellent echo canceller performance.

The basic constraint on getting good echo canceller performance is that the maximum output should not clip when coupled to the input. For example, if in a speakerphone, AO provides 1.1 V_{rms} to a

speaker, the reflections reaching the microphone should present no more than $0.9 V_{\text{rms}}$ to the Acoustic ADC. In fact, it is advisable to allow 6 dB or even 12 dB of margin, such that in the above example, the signal present at the Acoustic ADC is $250 \text{ mV}_{\text{rms}}$.

After this coupling level is established, the desired signal gain must be established. To continue from the previous example, the transmit gain must be adjusted to make sure the near-end talker is easy to hear at the far-end. If the signal from the near-end talker clips at the ADC, it is not significant to the echo path because the AEC should not be updating anyway.

In general, to minimize noise system gain should be concentrated before the ADC. However, this is not practical in all cases, mostly because of the coupling constraint. The CS6422 offers the AGC'd gains provided by TVol and RVol to help provide the desired transmit and receive gains.

The CS6422 offers two different programmable gain sources: TGain/RGain and TVol/RVol. TGain and RGain provide analog gain at the input to the ADC of 0 dB, 6 dB, 9.5 dB, or 12 dB. TVol and RVol introduce digital gain and attenuation in 3 dB steps. The difference is significant in that the digital gain will gain up the noise of the ADC as well as the desired signal, whereas the analog gain will not. Furthermore, gains introduced by TVol and RVol will not result in clipping, since both gains are AGC'd, unlike the gains at TGain and RGain which are not.

4.3.2 Testing Issues

The following tests are suggestions for measuring echo canceller and half-duplex performance.

4.3.2.1 ERLE

Echo Return-Loss Enhancement (ERLE) is a measure of the attenuation that an echo canceller provides. The number is an expression of the ratio of

the level of signal without the echo canceller compared to the level of signal with the echo canceller.

When measuring ERLE, it is important that any potential signal loops be broken; so to measure the ERLE of the Acoustic Canceller, the NO output should be disconnected from the rest of the network. This will prevent feedback which could occur when all of the CS6422's failsafes are disabled.

The following example outlines the steps necessary to measure the ERLE of the acoustic echo canceller.

It is important to choose a good test signal for the tests to be valid. As mentioned in Section 4.1.1.2, "Adaptive Filter", the CS6422 does not work optimally with white noise. The best signal to use would be a repeatable speech signal, like a recording of someone counting or saying "ah."

Use the Microcontroller Interface to disable transmit and receive suppression, half-duplex, and the Network Echo Canceller. The gains should be set appropriate for good system performance.

The first measurement is a baseline figure of performance with no echo canceller. Use the Microcontroller Interface to clear the acoustic canceller coefficients. Inject the test signal at NI and measure the rms voltage at NO. This measurement gives the baseline coupling level (denominator).

Use the Microcontroller Interface to set the acoustic canceller coefficients to normal which will allow the adaptive filter to adapt. Inject the test signal at NI and allow a few seconds for the filter to adapt. Again, measure the rms voltage at NO. This measurement gives the cancelled echo level (numerator).

Convert both voltages to decibels and subtract the echo cancelled level from the baseline level to calculate the ERLE. At the factory, with known good components, we typically see 30 dB of ERLE with speech.

4.3.2.2 Convergence Time

Convergence time is a measure of how quickly the adaptive filter can model the echo path. From cleared coefficients, the training signal is injected into the echo canceller and the time for the ERLE to reach a given threshold value is the convergence time. Different customers will have different threshold levels, so Crystal does not specify convergence time.

The following example will measure convergence time for the acoustic echo canceller:

Set up the system as for the ERLE test. Clear the acoustic canceller coefficients through the Microcontroller Interface. Apply the training signal to NI, set the coefficients to normal, and simultaneously start a timer. Once the measured ERLE reaches the threshold the system designer desires, stop the timer. The elapsed time is the convergence time. A good value for the threshold would be the AErle value from Register 3, since this would be the time for the CS6422 to go from half-duplex mode to full-duplex mode.

A good tool for this measurement is a digital storage oscilloscope set to a slow sweep so that about five seconds of signal is shown on the screen. One channel of the oscilloscope should monitor the ADC input (for an uncanceled reference), and another channel should monitor the echo cancelled output. This technique is especially effective when speech is the training signal.

We see about 2-5 seconds of training time using known good equipment. This time assumes continuous speech as the training signal. Pauses will extend the convergence time.

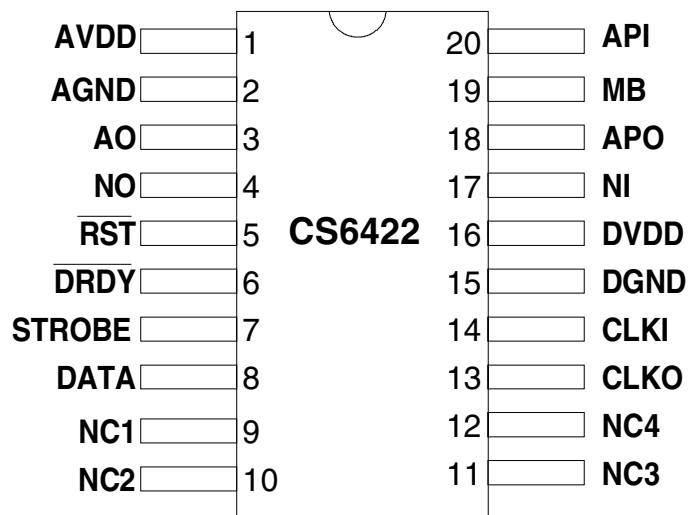
4.3.2.3 Half-Duplex Switching

Although the CS6422 transitions from half-duplex to full-duplex operation from reset after only a few utterances are passed through the system, the performance of half-duplex is critical to the end-user in cases where the echo canceller is not adequate. The half-duplex switching characteristics can be subjectively tested with the following procedure:

Set the CS6422 Microcontroller Interface to the nominal register values for the system, but clear the acoustic and network echo canceller coefficients. This will force the CS6422 to remain in half-duplex mode.

The most useful test of practical performance found at Crystal has been the “alternating counting test.” In this test the person at the near-end counts all the odd numbers and the person at the far-end counts all the even numbers. This tests the interruptibility of the half-duplexer. During testing, system parameters for the half-duplex may need to be changed to accommodate the level of performance expected for the product. See Section 4.1.2, “Half-Duplex Mode” and Section 3.2.2, “Register Definitions” for more details.

5. PIN DESCRIPTIONS



Analog Interface

AO - Acoustic Interface Output, Pin 3

Analog voltage output for the acoustic side (near-end output/receive output). Maximum output signal is $1.1 V_{rms}$ ($3.1 V_{pp}$). This output can drive down to $10 k\Omega$ and is usually followed by a speaker driver. The output is pre-compensated to expect a single-pole RC low pass filter with a corner frequency of 4 kHz.

NO - Network Interface Output, Pin 4

Analog voltage output for the network side (far-end output/transmit output). Maximum output signal is $1.1 V_{rms}$ ($3.1 V_{pp}$). This output can drive down to $10 k\Omega$. The output is pre-compensated to expect a single-pole RC low pass filter with a corner frequency of 4 kHz.

API - Acoustic Interface Preamplifier Input, Pin 20

Input to the acoustic side microphone preamplifier. Signal source resistance at this pin will reduce the 34 dB gain inherent in the preamplifier. The maximum input signal level to avoid clipping is $20 mV_{rms}$ ($57 mV_{pp}$), assuming default settings.

APO - Acoustic Interface Preamplifier Output, Pin 18

Output of the acoustic side microphone preamplifier and input to the acoustic side analog-to-digital converter (near-end input/transmit input). This input expects a single-pole RC anti-aliasing filter with a corner frequency of 8 kHz. Maximum signal level before clipping at this point is $0.9 V_{rms}$ ($2.5 V_{pp}$), assuming default settings for TGain.

MB - Microphone Bias Voltage Output, Pin 19

Output of 3.5 VDC provides the internal voltage reference for the CS6422. MB must be decoupled with a $10 \mu F$ and $0.1 \mu F$ capacitor to prevent noise from affecting the on-chip voltage reference. MB must not be connected to any load.

NI - Network Interface Input, Pin 17

Input to the network side analog-to-digital converter (far-end input/receive input). This input expects a single-pole RC anti-aliasing filter with a corner frequency of 8 kHz. Maximum signal level before clipping at this point is $0.9 V_{\text{rms}}$ ($2.5 V_{\text{pp}}$), assuming default settings for RGain.

Microcontroller Interface **$\overline{\text{RST}}$ - Active Low Reset Input, Pin 5**

When $\overline{\text{RST}}$ is held low, the CS6422 is put into a low power mode with all functional blocks idle. When $\overline{\text{RST}}$ goes high, the CS6422 is started in a known state.

 $\overline{\text{DRDY}}$ - Active Low Microcontroller Interface Data Ready Input, Pin 6

$\overline{\text{DRDY}}$ is a low pulse used to gate valid input data into the Microcontroller Interface.

STROBE - Microcontroller Interface Clock Input, Pin 7

The rising edge of STROBE latches DATA into the Microcontroller Interface while $\overline{\text{DRDY}}$ is low.

DATA - Microcontroller Interface Data Input, Pin 8

DATA is latched into the Microcontroller Interface on the rising edge of STROBE.

Clock**CLKI - Clock Oscillator Input, Pin 14**

A 20.480 MHz parallel-resonant crystal should be connected between CLKI and CLKO. Alternatively, CLKI may be driven directly with an 20.480 MHz CMOS level clock.

CLKO - Clock Oscillator Output, Pin 13

A 20.480 MHz parallel-resonant crystal should be connected between CLKI and CLKO. CLKO must be left floating if CLKI is driven directly with a CMOS level clock.

Power Supply**AVDD - Analog Supply, Pin 1**

+5 Volt analog power supply.

AGND - Analog Ground, Pin 2

Analog ground reference.

DVDD - Digital Supply, Pin 16

+5 Volt digital power supply.

DGND - Digital Ground, Pin 15

Digital ground reference.

*Miscellaneous***NC1 - No Connect, Pin 9**

Must be floating for normal operation.

NC2 - No Connect, Pin 10

Must be floating for normal operation.

NC3 - No Connect, Pin 11

Must be floating for normal operation.

NC4 - No Connect, Pin 12

Must be floating for normal operation.

6. GLOSSARY

Echo

A signal that returns to its source after some delay.

Network Echo

Echo resulting from signal reflection due to an impedance mismatch in a 2-to-4 wire converter (hybrid).

Acoustic Echo

Echo created by signal propagation in a room from a speaker to a microphone.

Reverberation

Local information that bounces around the room before it reaches the microphone. An example of reverberation is when your back is to the speakerphone, and your voice bounces off the wall before it reaches the microphone.

Near-End

The location with the acoustic interface (speaker and microphone).

Far-End

The location connected to the network interface.

Transmit Path

The signal path from Near-End input to Far-End output.

Receive Path

The signal path from Far-End input to Near-End output.

Full-Duplex

The state when both Transmit and Receive paths are simultaneously active.

Half-Duplex

The state when either Transmit or Receive path is active.

Supplementary Echo Suppression

Dynamic attenuation placed in the opposite path of the active path to mask residual echo. For example, if the receive path is active, the transmit path is attenuated. When both paths are simultaneously active, the suppression attenuation is removed. See Section 4.1.4, "Suppression" for more details.

Howling

In full-duplex operation, both the microphone and speaker are active at the same time, which, in conjunction with the reflection off the hybrid, creates a closed loop. The signal coupling between the speaker and the microphone can cause feedback oscillation or howling. This happens when the coupling between the speaker and microphone is strong enough to increase the system's closed loop gain above unity.

Acoustic Coupling

The strength of the output signal from the speaker that is received at the microphone input.

Adaptive Filter

A digital FIR filter that adjusts its coefficients to match a transfer function, such as the echo path between the speaker and microphone. The adaptive filter is able to compensate for different and changing conditions, such as someone moving in the room.

Echo Path

The acoustic echo path describes the acoustic coupling between the speaker and the microphone. It describes both the magnitude and delay characteristics of the echoed signal. It is affected by the speaker, microphone, phone housing, room, objects in the room, movement, and the talker. The network echo path is comprised of the transfer function between NO and NI.

Path Change

A change in the transfer function that describes the Echo Path. Changes in the acoustic echo path are most commonly due to motion in the room or gain changes at an external speaker. Network echo path is most easily changed by picking up an extension or hanging up the phone.

AGC

The CS6422 implements a peak-limiting Automatic Gain Control to allow a greater dynamic range without clipping the signal. See Section 4.1.3, “AGC” for details on how it works.

Doubletalk

The condition occurring when both Near End and Far End talkers are speaking simultaneously.

ERLE

Echo Return-Loss Enhancement is the amount of attenuation of echo signal an echo canceller provides (not counting Suppression) as measured in dB. ERLE is a measure of the echo canceller's performance. The larger the value for ERLE, the better the echo cancellation.

Coverage Time


The CS6422 echo canceller has 508 taps and it can sample an analog signal at an 8 kHz rate. $512 \times 1/8 \text{ kHz} = 63.5 \text{ ms}$. Sound travels through air at a rate of around 1 ft/ms. Thus the echo canceller can be used in a room with walls 32 feet away, discounting multiple reflections. But remember that at this distance, most of the echo has been attenuated due to the physical separation. The majority of the acoustic coupling comes from the first arrival, or directly from the speaker to the microphone. The first signal is by far the strongest.

Convergence Time

A high quality echo canceller is continuously modifying its internal model of the echo path characteristics (See Section 4.1.1.2, “Adaptive Filter”). When the model is complete, the echo canceller will be able to cancel echo to the extent of its rated capabilities. Convergence time is the duration it takes the echo canceller to train itself, from cleared coefficients, and switch to full-duplex operation, in the presence of speech.

7. PACKAGE DIMENSIONS

20L SOIC (300 MIL BODY) PACKAGE DRAWING



DIM	INCHES			MILLIMETERS		
	MIN	NOM	MAX	MIN	NOM	MAX
A	0.093	0.098	0.104	2.35	2.50	2.65
A1	0.004	0.008	0.012	0.10	0.20	0.30
b	0.013	0.017	0.020	0.33	0.43	0.51
C	0.009	0.011	0.013	0.23	0.28	0.32
D	0.496	0.504	0.512	12.60	12.80	13.00
E	0.291	0.295	0.299	7.40	7.50	7.60
e	0.040	0.050	0.060	1.02	1.27	1.52
H	0.394	0.407	0.419	10.00	10.34	10.65
L	0.016	0.025	0.050	0.40	0.64	1.27
∞	0°	4°	8°	0°	4°	8°

JEDEC #: MS-013

Controlling Dimension is Inches/Chip Pac

Controlling Dimension is Millimeters/Jedec

ORDERING INFORMATION

Model	Temperature	Package
CS6422-CS	0 to +70 °C	20-pin SOIC
CS6422-CSZ (Lead Free)		
CS6422-IS	-40 to +85 °C	
CS6422-ISZ (Lead Free)		

ENVIRONMENTAL, MANUFACTURING, & HANDLING INFORMATION

Model Number	Peak Reflow Temp	MSL Rating*	Max Floor Life
CS6422-CS	240 °C	2	365 Days
CS6422-CSZ (Lead Free)	260 °C	3	7 Days
CS6422-IS	240 °C	2	365 Days
CS6422-ISZ (Lead Free)	260 °C	3	7 Days

* MSL (Moisture Sensitivity Level) as specified by IPC/JEDEC J-STD-020.

REVISION HISTORY

Revision	Date	Changes
PP4	JUL 2001	Preliminary Release
F1	SEP 2005	Updated device ordering info. Updated legal notice. Added MSL data..

Contacting Cirrus Logic Support

For all product questions and inquiries contact a Cirrus Logic Sales Representative.

To find the one nearest to you go to www.cirrus.com

IMPORTANT NOTICE

Cirrus Logic, Inc. and its subsidiaries ("Cirrus") believe that the information contained in this document is accurate and reliable. However, the information is subject to change without notice and is provided "AS IS" without warranty of any kind (express or implied). Customers are advised to obtain the latest version of relevant information to verify, before placing orders, that information being relied on is current and complete. All products are sold subject to the terms and conditions of sale supplied at the time of order acknowledgment, including those pertaining to warranty, indemnification, and limitation of liability. No responsibility is assumed by Cirrus for the use of this information, including use of this information as the basis for manufacture or sale of any items, or for infringement of patents or other rights of third parties. This document is the property of Cirrus and by furnishing this information, Cirrus grants no license, express or implied under any patents, mask work rights, copyrights, trademarks, trade secrets or other intellectual property rights. Cirrus owns the copyrights associated with the information contained herein and gives consent for copies to be made of the information only for use within your organization with respect to Cirrus integrated circuits or other products of Cirrus. This consent does not extend to other copying such as copying for general distribution, advertising or promotional purposes, or for creating any work for resale.

CERTAIN APPLICATIONS USING SEMICONDUCTOR PRODUCTS MAY INVOLVE POTENTIAL RISKS OF DEATH, PERSONAL INJURY, OR SEVERE PROPERTY OR ENVIRONMENTAL DAMAGE ("CRITICAL APPLICATIONS"). CIRRUS PRODUCTS ARE NOT DESIGNED, AUTHORIZED OR WARRANTED FOR USE IN AIRCRAFT SYSTEMS, MILITARY APPLICATIONS, PRODUCTS SURGICALLY IMPLANTED INTO THE BODY, AUTOMOTIVE SAFETY OR SECURITY DEVICES, LIFE SUPPORT PRODUCTS OR OTHER CRITICAL APPLICATIONS. INCLUSION OF CIRRUS PRODUCTS IN SUCH APPLICATIONS IS UNDERSTOOD TO BE FULLY AT THE CUSTOMER'S RISK AND CIRRUS DISCLAIMS AND MAKES NO WARRANTY, EXPRESS, STATUTORY OR IMPLIED, INCLUDING THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR PARTICULAR PURPOSE, WITH REGARD TO ANY CIRRUS PRODUCT THAT IS USED IN SUCH A MANNER. IF THE CUSTOMER OR CUSTOMER'S CUSTOMER USES OR PERMITS THE USE OF CIRRUS PRODUCTS IN CRITICAL APPLICATIONS, CUSTOMER AGREES, BY SUCH USE, TO FULLY INDEMNIFY CIRRUS, ITS OFFICERS, DIRECTORS, EMPLOYEES, DISTRIBUTORS AND OTHER AGENTS FROM ANY AND ALL LIABILITY, INCLUDING ATTORNEYS' FEES AND COSTS, THAT MAY RESULT FROM OR ARISE IN CONNECTION WITH THESE USES.

Cirrus Logic, Cirrus, and the Cirrus Logic logo designs are trademarks of Cirrus Logic, Inc. All other brand and product names in this document may be trademarks or service marks of their respective owners.