

FEATURES

- Complete supervisory and sequencing solution for up to 10 supplies**
- 10 supply fault detectors enable supervision of supplies to better than 1% accuracy**
- 5 selectable input attenuators allow supervision:**
 - Supplies up to 14.4 V on VH**
 - Supplies up to 6 V on VP1–4**
- 5 dual-function inputs, VX1–5:**
 - High impedance input to supply fault detector with thresholds between 0.573 V and 1.375 V**
 - General-purpose logic input**
- 10 programmable output drivers (PDO1–10):**
 - Open collector with external pull-up**
 - Push/pull output, driven to VDDCAP or VPn**
 - Open collector with weak pull-up to VDDCAP or VPn**
 - Internally charge-pumped high drive for use with external N-FET (PDO1–6 only)**
- Sequencing engine (SE) implements state machine control of PDO outputs:**
 - State changes conditional on input events**
 - Enables complex control of boards**
 - Power-up and power-down sequence control**
 - Fault event handling**
 - Interrupt generation on warnings**
 - Watchdog function can be integrated in SE**
 - Program software control of sequencing through SMBus**
- Open-loop margining solution for 6 voltage rails**
- 6 voltage output 8-bit DACs (0.300 V to 1.551 V) allow voltage adjustment via dc/dc converter trim/feedback node**
- Device powered by the highest of VP1–4, VH for improved redundancy**
- User EEPROM: 256 bytes**
- Industry-standard 2-wire bus interface (SMBus)**
- Guaranteed PDO low with VH, VPn = 1.2 V**
- 40-lead 6 mm × 6 mm LFCSP and**
- 48-lead 7 mm × 7 mm TQFP packages**

FUNCTIONAL BLOCK DIAGRAM




Figure 1.

APPLICATIONS

- Central office systems**
- Servers/routers**
- Multivoltage system line cards**
- DSP/FPGA supply sequencing**
- In-circuit testing of margined supplies**

GENERAL DESCRIPTION

The ADM1067 is a configurable supervisory/sequencing device that offers a single-chip solution for supply monitoring and sequencing in multiple supply systems. In addition to these functions, the ADM1067 integrates six 8-bit voltage output DACs. These circuits can be used to implement an open-loop margining system, which enables supply adjustment by altering either the feedback node or reference of a dc/dc converter using the DAC outputs.

(continued on Page 3)

Rev. 0

Information furnished by Analog Devices is believed to be accurate and reliable. However, no responsibility is assumed by Analog Devices for its use, nor for any infringements of patents or other rights of third parties that may result from its use. Specifications subject to change without notice. No license is granted by implication or otherwise under any patent or patent rights of Analog Devices. Trademarks and registered trademarks are the property of their respective owners.

TABLE OF CONTENTS

General Description	3	Sequence Detector.....	19
Specifications.....	4	Monitoring Fault Detector	19
Pin Configurations and Function Descriptions	7	Timeout Detector	19
Absolute Maximum Ratings.....	8	Fault Reporting.....	19
Thermal Characteristics	8	Supply Margining	20
ESD Caution.....	8	Overview	20
Typical Performance Characteristics	9	Open-Loop Margining	20
Powering the ADM1067	12	Writing to the DACs	20
Inputs.....	13	Choosing the Size of the Attenuation Resistor	21
Supply Supervision.....	13	DAC Limiting/Other Safety Features	21
Programming the Supply Fault Detectors.....	13	Applications Diagram	22
Input Comparator Hysteresis.....	14	Communicating with the ADM1067	23
Input Glitch Filtering	14	Configuration Download at Power-Up	23
Supply Supervision with VXn Inputs.....	14	Updating the Configuration	23
VXn Pins as Digital Inputs.....	15	Updating the Sequencing Engine.....	24
Outputs	16	Internal Registers.....	24
Supply Sequencing through Configurable Output Drivers..	16	EEPROM	24
Sequencing Engine	17	Serial Bus Interface.....	24
Overview.....	17	Write Operations	26
Warnings.....	17	Read Operations.....	28
SMBus Jump/Unconditional Jump	17	Outline Dimensions	30
Sequencing Engine Application Example	18	Ordering Guide	30

REVISION HISTORY

10/04—Revision 0: Initial Version

SPECIFICATIONS

VH = 3.0 V to 14.4 V¹, VPn = 3.0 V to 6.0 V¹, TA = -40°C to +85°C, unless otherwise noted.

Table 1.

Parameter	Min	Typ	Max	Unit	Test Conditions/Comments
POWER SUPPLY ARBITRATION					
VH, VPn	3.0			V	Minimum supply required on one of VPn, VH
VP			6.0	V	Maximum VDDCAP = 5.1 V, typical
VH			14.4	V	VDDCAP = 4.75 V
VDDCAP	2.7	4.75	5.4	V	Regulated LDO output
CVDDCAP	10			μF	Minimum recommended decoupling capacitance
POWER SUPPLY					
Supply Current, I _{VH} , I _{VPn}		4.2	6	mA	VDDCAP = 4.75 V, PDO1–10 off, DACs off, ADC off
Additional Currents					
All PDO FET Drivers On		1		mA	VDDCAP = 4.75 V, PDO1-6 loaded with 1 μA each, PDO7–10 off
Current Available from VDDCAP			2	mA	Maximum additional load that can be drawn from all PDO pull-ups to VDDCAP
DACs Supply Current		2.2		mA	6 DACs on with 100 μA maximum load on each
ADC Supply Current		1		mA	Running round-robin loop
EEPROM Erase Current		10		mA	1 ms duration only, VDDCAP = 3 V
SUPPLY FAULT DETECTORS					
VH Pin					
Input Attenuator Error		±0.05		%	Midrange and high range
Detection Ranges					
High Range	6		14.4	V	
Midrange	2.5		6	V	
VPn Pins					
Input Attenuator Error		±0.05		%	Low range and midrange
Detection Ranges					
Midrange	2.5		6	V	
Low Range	1.25		3	V	
Ultralow Range	0.573		1.375	V	No input attenuation error
VX Pins					
Input Impedance	1			MΩ	
Detection Ranges					
Ultralow Range	0.573		1.375	V	No input attenuation error
Absolute Accuracy			±1	%	VREF error + DAC nonlinearity + comparator offset error + input attenuation error
Threshold Resolution		8		Bits	
Digital Glitch Filter		0		μs	Minimum programmable filter length
		100		μs	Maximum programmable filter length
BUFFERED VOLTAGE OUTPUT DACs					
Resolution		8		Bits	
Code 0x80 Output Voltage					6 DACs are individually selectable for centering on one of four output voltage ranges
Range 1	0.592	0.6	0.603	V	
Range 2	0.796	0.8	0.803	V	
Range 3	0.997	1	1.003	V	
Range 4	1.247	1.25	1.253	V	
Output Voltage Range		601.25		mV	Same range, independent of center point
LSB Step Size		2.36		mV	
INL			±0.75	LSB	Endpoint corrected

Parameter	Min	Typ	Max	Unit	Test Conditions/Comments
DNL			±0.4	LSB	
Gain Error			1	%	
Load Regulation		-4		mV	Sourcing Current, $I_{REFOUTMA X} = -200\mu A$
		2		mV	Sinking Current, $I_{REFOUTMA X} = 100\mu A$
Maximum Load Capacitance			50	pF	
Settling Time into 50 pF Load			2	μs	
Load Regulation		2.5		mV	Per mA
PSRR		60		dB	DC
		40		dB	100 mV step in 20 ns with 50 pF load
REFERENCE OUTPUT					
Reference Output Voltage	2.043	2.048	2.053	V	No load
Load Regulation		-0.25		mV	Sourcing current, $I_{DACnMAX} = -100\mu A$
		0.25		mV	Sinking current, $I_{DACnMAX} = 100\mu A$
Minimum Load Capacitance	1			μF	Capacitor required for decoupling, stability
Load Regulation		2		mV	Per 100 μA
PSRR		60		dB	DC
PROGRAMMABLE DRIVER OUTPUTS					
High Voltage (Charge Pump) Mode (PDO1–6)					
Output Impedance		500		kΩ	
V_{OH}	11	12.5	14	V	$I_{OH} = 0$
	10.5	12	13.5	V	$I_{OH} = 1\mu A$
I_{OUTAVG}		20		μA	$2 V < V_{OH} < 7 V$
Standard (Digital Output) Mode (PDO1–10)					
V_{OH}	2.4			V	V_{PU} (pull-up to V_{DDCAP} or V_{PN}) = 2.7 V, $I_{OH} = 0.5 mA$
			4.5	V	V_{PU} to $V_{PN} = 6.0 V$, $I_{OH} = 0 mA$
	$V_{PU} - 0.3$			V	$V_{PU} \leq 2.7 V$, $I_{OH} = 0.5 mA$
V_{OL}	0		0.50	V	$I_{OL} = 20 mA$
I_{OL}^2			20	mA	Maximum sink current per PDO pin
I_{SINK}^2			60	mA	Maximum total sink for all PDOs
$R_{PULL-UP}$		20		kΩ	Internal pull-up
$I_{SOURCE} (VPn)^2$			2	mA	Current load on any VPn pull-ups, that is, total source current available through any number of PDO pull-up switches configured onto any one
Three-State Output Leakage Current			10	μA	$V_{PDO} = 14.4 V$
Oscillator Frequency	90	100	110	kHz	All on-chip time delays derived from this clock
DIGITAL INPUTS (VXn, A0, A1, MUP, MDN)					
Input High Voltage, V_{IH}	2.0			V	Maximum $V_{IN} = 5.5 V$
Input Low Voltage, V_{IL}			0.8	V	Maximum $V_{IN} = 5.5 V$
Input High Current, I_{IH}	-1			μA	$V_{IN} = 5.5 V$
Input Low Current, I_{IL}			1	μA	$V_{IN} = 0$
Input Capacitance		5		pF	
Programmable Pull-Down Current, $I_{PULL-DOWN}$		20		μA	$V_{DDCAP} = 4.75$, $T_A = 25^\circ C$, if known logic state is required
SERIAL BUS DIGITAL INPUTS (SDA, SCL)					
Input High Voltage, V_{IH}	2.0			V	
Input Low Voltage, V_{IL}			0.8	V	
Output Low Voltage, V_{OL}^2			0.4	V	$I_{OUT} = -3.0 mA$
SERIAL BUS TIMING					
Clock Frequency, f_{SCLK}			400	kHz	
Bus Free Time, t_{BUF}	4.7			μs	
Start Setup Time, $t_{SU;STA}$	4.7			μs	
Start Hold Time, $t_{HD;STA}$	4			μs	

ADM1067

Parameter	Min	Typ	Max	Unit	Test Conditions/Comments
SCL Low Time, t_{LOW}	4.7			μs	$V_{\text{IN}} = 0$
SCL High Time, t_{HIGH}	4			μs	
SCL, SDA Rise Time, t_r			1000	μs	
SCL, SDA Fall Time, t_f			300	μs	
Data Setup Time, $t_{\text{SU;DAT}}$	250			ns	
Data Hold Time, $t_{\text{HD;DAT}}$	5			ns	
Input Low Current, I_{IL}			1	μA	
SEQUENCING ENGINE TIMING					
State Change Time		10		μs	

¹ At least one of the VH, VP1-4 pins must be $\geq 3.0\text{ V}$ to maintain the device supply on VDDCAP.

² Specification is not production tested, but is supported by characterization data at initial product release.

PIN CONFIGURATIONS AND FUNCTION DESCRIPTIONS




Figure 3. LFCSP Pin Configuration




Figure 4. TQFP Pin Configuration

Table 2. Pin Function Descriptions

Pin No.		Mnemonic	Description
LFCSP	TQFP		
13	1, 12–13, 16, 24–25, 36–37, 48	NC	No connection.
1–5	2–6	VX1–5	High Impedance Inputs to Supply Fault Detectors. Fault thresholds can be set from 0.573 V to 1.375 V. Alternatively, these pins can be used as general-purpose digital inputs.
6–9	7–10	VP1–4	Low Voltage Inputs to Supply Fault Detectors. Three input ranges can be set by altering the input attenuation on a potential divider connected to these pins, the output of which connects to a supply fault detector. These pins allow thresholds from 2.5 V to 6.0 V, 1.25 V to 3.00 V, and 0.573 V to 1.375 V.
10	11	VH	High Voltage Input to Supply Fault Detectors. Three input ranges can be set by altering the input attenuation on a potential divider connected to this pin, the output of which connects to a supply fault detector. This pin allows thresholds from 6.0 V to 14.4 V and 2.5 V to 6.0 V.
11	14	AGND	Ground Return for Input Attenuators.
12	15	REFGND	Ground Return for On-Chip Reference Circuits.
14	17	REFOUT	2.048 V Reference Output.
15–20	18–23	DAC1–6	Voltage Output DACs. These pins default to high impedance at power-up.
21–30	26–35	PDO10–1	Programmable Output Drivers.
31	38	PDOGND	Ground Return for Output Drivers.
32	39	VCCP	Central Charge-Pump Voltage of 5.25 V. A reservoir capacitor must be connected between this pin and GND.
33	40	A0	Logic Input. This pin sets the seventh bit of the SMBus interface address.
34	41	A1	Logic Input. This pin sets the sixth bit of the SMBus interface address.
35	42	SCL	SMBus Clock Pin. Open-drain output requires external resistive pull-up.
36	43	SDA	SMBus Data I/O Pin. Open-drain output requires external resistive pull-up.
37	44	MUP	Digital Input. Forces DACs to their lowest value, causing the voltage at the feedback node to drop, which is compensated for by the output voltage of the supply increasing, thus margining up.
38	45	MDN	Digital Input. Forces DACs to their highest value, causing the voltage at the feedback node to rise, which is compensated for by the output voltage of the supply decreasing, thus margining down.
39	46	VDDCAP	Device Supply Voltage. Linearly regulated from the highest of the VP1–4, VH pins to a typical of 4.75 V.
40	47	GND	Supply Ground.

ABSOLUTE MAXIMUM RATINGS

Table 3.

Parameter	Rating
Voltage on VH Pin	16 V
Voltage on VP Pins	7 V
Voltage on VX Pins	−0.3 V to +6.5 V
Voltage on MUP, MDN Pins	−0.3 V to +5 V
Input Current at Any Pin	±5 mA
Package Input Current	±20 mA
Maximum Junction Temperature (T _J max)	150°C
Storage Temperature Range	−65°C to +150°C
Lead Temperature, Soldering Vapor Phase, 60 s	215°C

Stresses above those listed under Absolute Maximum Ratings may cause permanent damage to the device. This is a stress rating only; functional operation of the device at these or any other conditions above those indicated in the operational section of this specification is not implied. Exposure to absolute maximum rating conditions for extended periods may affect device reliability.

THERMAL CHARACTERISTICS

40-lead LFCSP package: $\theta_{JA} = 25^{\circ}\text{C/W}$

48-lead TQFP package: $\theta_{JA} = 14.8^{\circ}\text{C/W}$

ESD CAUTION

ESD (electrostatic discharge) sensitive device. Electrostatic charges as high as 4000 V readily accumulate on the human body and test equipment and can discharge without detection. Although this product features proprietary ESD protection circuitry, permanent damage may occur on devices subjected to high energy electrostatic discharges. Therefore, proper ESD precautions are recommended to avoid performance degradation or loss of functionality.



TYPICAL PERFORMANCE CHARACTERISTICS




Figure 5. V_{DDCAP} vs. V_{VP1}




Figure 8. I_{VP1} vs. V_{VP1} (VP1 Not as Supply)




Figure 6. V_{DDCAP} vs. V_{VH}




Figure 9. I_{VH} vs. V_{VH} (VH as Supply)




Figure 7. I_{VP1} vs. V_{VP1} (VP1 as Supply)




Figure 10. I_{VH} vs. V_{VH} (VH Not as Supply)




Figure 11. V_{PDO1} (FET Drive Mode) vs. I_{LOAD}




Figure 14. Transient Response of DAC Code Change into Typical Load




Figure 12. V_{PDO1} (Strong Pull-Up V_P) vs. I_{LOAD}




Figure 15. Transient Response of DAC to Turn-On from HI-Z State




Figure 13. V_{PDO1} (Weak Pull-Up to V_P) vs. I_{LOAD}




Figure 16. DAC Output vs. Temperature




Figure 17. REFOUT vs. Temperature

POWERING THE ADM1067

The ADM1067 is powered from the highest voltage input on either the positive-only supply inputs (VPn) or the high voltage supply input (VH). This technique offers improved redundancy as the device is not dependent on any particular voltage rail to keep it operational. The same pins are used for supply fault detection (discussed later in the next section). A VDD arbitrator on the device chooses which supply to use. The arbitrator can be considered an OR'ing of five LDOs together. A supply comparator chooses which of the inputs is highest and selects this one to provide the on-chip supply. There is minimal switching loss with this architecture (~ 0.2 V), resulting in the ability to power the ADM1067 from a supply as low as 3.0 V. Note that the supply on the VXn pins cannot be used to power the device.

An external capacitor to GND is required to decouple the on-chip supply from noise. This capacitor should be connected to the VDDCAP pin, as shown in Figure 18. The capacitor has another use during brownouts (momentary loss of power). Under these conditions, when the input supply (VPn or VH) dips transiently below V_{DD} , the synchronous rectifier switch immediately turns off so that it does not pull V_{DD} down. The V_{DD} cap can then act as a reservoir to keep the device active until the next highest supply takes over the powering of the device. 10 μ F is recommended for this reservoir/decoupling function.

Note that when two or more supplies are within 100 mV of each other, the supply that takes control of V_{DD} first keeps control.

For example, if VP1 is connected to a 3.3 V supply, then V_{DD} powers up to approximately 3.1 V through VP1. If VP2 is then connected to another 3.3 V supply, VP1 still powers the device, unless VP2 goes 100 mV higher than VP1.




Figure 18. VDD Arbitrator Operation

04635-022

INPUTS

SUPPLY SUPERVISION

The ADM1067 has ten programmable inputs. Five of these are dedicated supply fault detectors (SFDs). These dedicated inputs are called VH and VP1–4 by default. The other five inputs are labeled VX1–VX5 and have dual functionality. They can be used as either supply fault detectors, with similar functionality to VH and VP1–4, or CMOS/TTL-compatible logic inputs to the devices. Therefore, the ADM1067 can have up to ten analog inputs, a minimum of five analog inputs and five digital inputs, or a combination. If an input is used as an analog input, it cannot be used as a digital input. Therefore, a configuration requiring ten analog inputs has no digital inputs available. Table 5 shows the details of each of the inputs.




Figure 19. Supply Fault Detector Block

PROGRAMMING THE SUPPLY FAULT DETECTORS

The ADM1067 has up to ten supply fault detectors (SFDs) on its ten input channels. These highly programmable reset generators enable the supervision of up to ten supply voltages. The supplies can be as low as 0.573 V and as high as 14.4 V. The inputs can be configured to detect an undervoltage fault (the input voltage droops below a preprogrammed value), an overvoltage fault (the input voltage rises above a preprogrammed value) or an out-of-window fault (undervoltage or overvoltage). The thresholds can be programmed to an 8-bit resolution in registers provided in the ADM1067. This translates to a voltage resolution that is dependent on the range selected.

The resolution is given by

$$\text{Step Size} = \text{Threshold Range}/255$$

Therefore, if the high range is selected on VH, the step size can be calculated as follows:

$$(14.4 \text{ V} - 4.8 \text{ V})/255 = 37.6 \text{ mV}$$

Table 4 lists the upper and lower limit of each available range, the bottom of each range (V_B), and the range itself (V_R).

Table 4. Voltage Range Limits

Voltage Range (V)	V_B (V)	V_R (V)
0.573 to 1.375	0.573	0.802
1.25 to 3.00	1.25	1.75
2.5 to 6.0	2.5	3.5
4.8 to 14.4	4.8	9.6

The threshold value required is given by

$$V_T = (V_R \times N)/255 + V_B$$

where:

V_T is the desired threshold voltage (UV or OV).

V_R is the voltage range.

N is the decimal value of the 8-bit code.

V_B is the bottom of the range.

Reversing the equation, the code for a desired threshold is given by

$$N = 255 \times (V_T - V_B)/V_R$$

For example, if the user wants to set a 5 V OV threshold on VP1, the code to be programmed in the PS1OVTH register (discussed in the AN-698 application note) is given by

$$N = 255 \times (5 - 2.5)/3.5$$

Therefore, $N = 182$ (1011 0110 or 0xB6).

Table 5. Input Functions, Thresholds, and Ranges

Input	Function	Voltage Range (V)	Maximum Hysteresis	Voltage Resolution (mV)	Glitch Filter (μ s)
VH	High V Analog Input	2.5–6.0	425 mV	13.7	0–100
		4.8–14.4	1.16 V	37.6	0–100
VPn	Positive Analog Input	0.573–1.375	97.5 mV	3.14	0–100
		1.25–3.00	212 mV	6.8	0–100
		2.5–6.0	425 mV	13.7	0–100
		0.573–1.375	97.5 mV	3.14	0–100
VXn	High Z Analog Input	0.573–1.375	97.5 mV	3.14	0–100
	Digital Input	0–5	N/A	N/A	0–100

INPUT COMPARATOR HYSTERESIS

The UV and OV comparators shown in Figure 19 are always looking at VPn. To avoid chattering (multiple transitions when the input is very close to the set threshold level), these comparators have digitally programmable hysteresis. The hysteresis can be programmed up to the values shown in Table 5.

The hysteresis is added after a supply voltage goes out of tolerance. Therefore, the user can program how much above the UV threshold the input must rise again before a UV fault is deasserted. Similarly, the user can program how much below the OV threshold an input must fall again before an OV fault is deasserted.

The hysteresis figure is given by

$$V_{HYST} = V_R \times N_{THRESH}/255$$

where:

V_{HYST} is the desired hysteresis voltage.

N_{THRESH} is the decimal value of the 5-bit hysteresis code.

Note that N_{THRESH} has a maximum value of 31. The maximum hysteresis for the ranges is listed in Table 5.

INPUT GLITCH FILTERING

The final stage of the SFDs is a glitch filter. This block provides time-domain filtering on the output of the SFD comparators. This allows the user to remove any spurious transitions such as supply bounce at turn-on. The glitch filter function is additional to the digitally programmable hysteresis of the SFD comparators. The glitch filter timeout is programmable up to 100 μ s.

For example, when the glitch filter timeout is 100 μ s, any pulses appearing on the input of the glitch filter block that are less than 100 μ s in duration are prevented from appearing on the output of the glitch filter block. Any input pulse that is longer than 100 μ s does appear on the output of the glitch filter block. The output is delayed with respect to the input by 100 μ s. The filtering process is shown in Figure 20.




Figure 20. Input Glitch Filter Function

SUPPLY SUPERVISION WITH VXn INPUTS

The VXn inputs have two functions. They can be used as either supply fault detectors or digital logic inputs. When selected as an analog (SFD) input, the VXn pins have very similar functionality to the VH and VPn pins. The major difference is that the VXn pins have only one input range: 0.573 V to 1.375 V. Therefore, these inputs can directly supervise only the very low supplies. However, the input impedance of the VXn pins is high, allowing an external resistor divide network to be connected to the pin. Thus, any supply can be potentially divided down into the input range of the VXn pin and supervised. This enables the ADM1067 to monitor other supplies such as +24 V, +48 V, and -5 V.

An additional supply supervision function is available when the VXn pins are selected as digital inputs. In this case, the analog function is available as a second detector on each of the dedicated analog inputs, VP1-4 and VH. The analog function of VX1 is mapped to VP1, VX2 is mapped to VP2, and so on. VX5 is mapped to VH. In this case, these SFDs can be viewed as a secondary or warning SFD.

The secondary SFDs are fixed to the same input range as the primary SFD. They are used to indicate warning levels rather than failure levels. This allows faults and warnings to be generated on a single supply using only one pin. For example, if VP1 is set to output a fault if a 3.3 V supply droops to 3.0 V, VX1 can be set to output a warning at 3.1 V. Warning outputs are available for readback from the status registers. They are also OR'ed together and fed into the sequencing engine (SE), allowing warnings to generate interrupts on the PDOs. Therefore, in the example above, if the supply droops to 3.1 V, a warning is generated, and remedial action can be taken before the supply drops out of tolerance.

VXn PINS AS DIGITAL INPUTS

As mentioned previously, the VXn input pins on the ADM1067 have dual functionality. The second function is as a digital input to the device. Therefore, the ADM1067 can be configured for up to five digital inputs. These inputs are TTL/CMOS-compatible. Standard logic signals can be applied to the pins: RESET from reset generators, PWRGOOD signals, fault flags, manual resets, and so on. These signals are available as inputs to the SE, and can be used to control the status of the PDOs. The inputs can be configured to detect either a change in level or an edge.

When configured for level detection, the output of the digital block is a buffered version of the input. When configured for edge detection, once the logic transition is detected, a pulse of programmable width is output from the digital block. The width is programmable from 0 μ s to 100 μ s.

The digital blocks feature the same glitch filter function that is available on the SFDs. This enables the user to ignore spurious transitions on the inputs. For example, the filter can be used to debounce a manual reset switch.

When configured as digital inputs, each of the VXn pins has a weak (10 μ A) pull-down current source available for placing the input in a known condition, even if left floating. The current source, if selected, weakly pulls the input to GND.




Figure 21. VXn Digital Input Function

04635-027

OUTPUTS

SUPPLY SEQUENCING THROUGH CONFIGURABLE OUTPUT DRIVERS

Supply sequencing is achieved with the ADM1067 using the programmable driver outputs (PDOs) on the device as control signals for supplies. The output drivers can be used as logic enables or as FET drivers.

The sequence in which the PDOs are asserted (and, therefore, the supplies are turned on) is controlled by the sequencing engine (SE). The SE determines what action is to be taken with the PDOs based on the condition of the inputs of the ADM1067. Therefore, the PDOs can be set up to assert when the SFDs are in tolerance, the correct input signals are received on the VXn digital pins, no warnings are received from any of the inputs of the device, and so on. The PDOs can be used for a variety of functions. The primary function is to provide enable signals for LDOs or dc/dc converters, which generate supplies locally on a board. The PDOs can also be used to provide a POWER_GOOD signal when all the SFDs are in tolerance, or a RESET output if one of the SFDs goes out of specification (this can be used as a status signal for a DSP, FPGA, or other microcontroller).

The PDOs can be programmed to pull up to a number of different options. The outputs can be programmed as follows:

- Open-drain (allowing the user to connect an external pull-up resistor)
- Open-drain with weak pull-up to V_{DD}
- Push/pull to V_{DD}
- Open-drain with weak pull-up to VP_n
- Push/pull to VP_n
- Strong pull-down to GND
- Internally charge-pumped high drive (12 V, PDO1–6 only)

The last option (available only on PDO1–6) allows the user to directly drive a voltage high enough to fully enhance an external N-FET, which is used to isolate, for example, a card-side voltage

from a backplane supply (a PDO can sustain greater than 10.5 V into a 1 μ A load). The pull-down switches can also be used to drive status LEDs directly.

The data driving each of the PDOs can come from one of three sources. The source can be enabled in the PnPDOCFG configuration register (see the AN-698 application note for details).

The data sources are

- Output from the SE.
- Directly from the SMBus. A PDO can be configured so that the SMBus has direct control over it. This enables software control of the PDOs. Therefore, a microcontroller can be used to initiate a software power-up/power-down sequence.
- On-Chip Clock. A 100 kHz clock is generated on the device. This clock can be made available on any of the PDOs. It can be used, for example, to clock an external device such as an LED.

By default, the PDOs are pulled to GND by a weak (20 k Ω) on-chip pull-down resistor. This is also the condition of the PDOs on power-up, until the configuration is downloaded from EEPROM and the programmed setup is latched. The outputs are actively pulled low once a supply of 1 V or greater is on VP_n or VH . The outputs remain high impedance prior to 1 V appearing on VP_n or VH . This provides a known condition for the PDOs during power-up. The internal pull-down can be overdriven with an external pull-up of suitable value tied from the PDO pin to the required pull-up voltage. The 20 k Ω resistor must be accounted for in calculating a suitable value. For example, if PDO_n must be pulled up to 3.3 V, and 5 V is available as an external supply, the pull-up resistor value is given by

$$3.3 \text{ V} = 5 \text{ V} \times 20 \text{ k}\Omega / (R_{UP} + 20 \text{ k}\Omega)$$

Therefore,

$$R_{UP} = (100 \text{ k}\Omega - 66 \text{ k}\Omega) / 3.3 = 10 \text{ k}\Omega$$




Figure 22. Programmable Driver Output

SEQUENCING ENGINE

OVERVIEW

The ADM1067's sequencing engine (SE) provides the user with powerful and flexible control of sequencing. The SE implements a state machine control of the PDO outputs, with state changes conditional on input events. SE programs can enable complex control of boards such as power-up and power-down sequence control, fault event handling, interrupt generation on warnings, and so on. A watchdog function that verifies the continued operation of a processor clock can be integrated into the SE program. The SE can also be controlled via the SMBus, giving software or firmware control of the board sequencing.

The SE state machine comprises 63 state cells. Each state has the following attributes:

- Monitors signals indicating the status of the 10 input pins, VP1 to VP4, VH, and VX1 to VX5.
- Can be entered from any other state.
- Three exit routes move the state machine on to a next state: sequence detection, fault monitoring, and timeout.
- Delay timers for the sequence and timeout blocks can be programmed independently, and change with each state change. The range of timeouts is from 0 ms to 400 ms.
- Output condition of the 10 PDO pins is defined and fixed within a state.

- Transition from one state to the next is made in less than 20 μ s, which is the time needed to download a state definition from EEPROM to the SE.




Figure 23. State Cell

The ADM1067 offers up to 63 state definitions. The signals monitored to indicate the status of the input pins are the outputs of the SFDs.

WARNINGS

The SE also monitors warnings. These warnings can be generated when the ADC readings violate their limit register value or when the secondary voltage monitors on VP1–4 and VH. The warnings are all OR'ed together and are available as a single warning input to each of the three blocks that enable exiting from a state.

SMBus JUMP/UNCONDITIONAL JUMP

The SE can be forced to advance to the next state unconditionally. This enables the user to force the SE to advance. Examples of where this might be used include moving to a margining state or debugging a sequence. The SMBus jump or go-to command can be seen as another input to sequence and timeout blocks, which provide an exit from each state.

Table 6. Sample Sequence State Entries

State	Sequence	Timeout	Monitor
IDLE1	If VX1 is low, go to state IDLE2.		
IDLE2	If VP1 is okay, go to state EN3V3.		
EN3V3	If VP2 is okay, go to state EN2V5.	If VP2 is not okay after 10 ms, go to state DIS3V3.	If VP1 is not okay, go to state IDLE1.
DIS3V3	If VX1 is high, go to state IDLE1.		
EN2V5	If VP3 is okay, go to state PWRGD.	If VP3 is not okay after 20 ms, go to state DIS2V5.	If VP1 or VP2 is not okay, go to state FSEL2.
DIS2V5	If VX1 is high, go to state IDLE1.		
FSEL1	If VP3 is not okay, go to state DIS2V5.		If VP1 or VP2 is not okay, go to state FSEL2.
FSEL2	If VP2 is not okay, go to state DIS3V3.		If VP1 is not okay, go to state IDLE1.
PWRGD	If VX1 is high, go to state DIS2V5.		If VP1, VP2, or VP3 is not okay, go to state FSEL1.

SEQUENCING ENGINE APPLICATION EXAMPLE

The application in this section demonstrates the operation of the SE. Figure 24 shows how the simple building block of a single SE state can be used to build a power-up sequence for a 3-supply system.

Table 7 lists the PDO outputs for each state in the same SE implementation. In this system, the presence of a good 5 V supply on VP1 and the VX1 pin held low are the triggers required for a power-up sequence to start. The sequence intends to turn on the 3.3 V supply next, then the 2.5 V supply (assuming successful turn-on of the 3.3 V supply). Once all three supplies are good, the PWRGD state is entered, where the SE remains until a fault occurs on one of the three supplies, or it is instructed to go through a power-down sequence by VX1 going high.

Faults are dealt with throughout the power-up sequence on a case-by-case basis. The following sections, which describe the individual blocks, use this sample application to demonstrate the state machine's actions.




Figure 24. Sample Application Flow Diagram

Table 7. PDO Outputs for Each State

PDO Outputs	IDLE1	IDLE2	EN3V3	EN2V5	DIS3V3	DIS2V5	PWRGD	FSEL1	FSEL2
PDO1 = 3V3ON	0	0	1	1	0	1	1	1	1
PDO2 = 2V5ON	0	0	0	1	1	0	1	1	1
PDO3 = FAULT	0	0	0	0	1	1	0	1	1

SEQUENCE DETECTOR

The sequence detector block is used to detect when a step in a sequence has been completed. It looks for one of the inputs to the SE to change state, and is most often used as the gate on successful progress through a power-up or power-down sequence. A timer block is included in this detector, which can insert delays into a power-up or power-down sequence, if required. Timer delays can be set from 10 μ s to 400 ms. Figure 25 is a block diagram of the sequence detector.




Figure 25. Sequence Detector Block Diagram

The sequence detector can also help to identify monitoring faults. In the sample application shown in Figure 24, the FSEL1 and FSEL2 states first identify which of the VP1, VP2, or VP3 pins has faulted, and then they take the appropriate action.

MONITORING FAULT DETECTOR

The monitoring fault detector block is used to detect a failure on an input. The logical function implementing this is a wide OR gate, which can detect when an input deviates from its expected condition. The clearest demonstration of the use of this block is in the PWRGD state, where the monitor block indicates that a failure on one or more of the VP1, VP2, or VP3 inputs has occurred.

No programmable delay is available in this block, because the triggering of a fault condition is likely to be caused when a supply falls out of tolerance. In this situation, the user would want to react as quickly as possible. Some latency occurs when moving out of this state, however, because it takes a finite amount of time ($\sim 20 \mu$ s) for the state configuration to download from EEPROM into the SE. Figure 26 is a block diagram of the monitoring fault detector.




Figure 26. Monitoring Fault Detector Block Diagram

TIMEOUT DETECTOR

The timeout detector allows the user to trap a failure to make proper progress through a power-up or power-down sequence.

In the sample application shown in Figure 24, the timeout next-state transition is from the EN3V3 and EN2V5 states. For the EN3V3 state, the signal 3V3ON is asserted upon entry to this state (on the PDO1 output pin) to turn on a 3.3 V supply. This supply rail is connected to the VP2 pin, and the sequence detector looks for the VP2 pin to go above its UV threshold, which is set in the supply fault detector (SFD) attached to that pin.

The power-up sequence progresses when this change is detected. If, however, the supply fails (perhaps due to a short circuit overloading this supply), then the timeout block traps the problem. In this example, if the 3.3 V supply fails within 10 ms, then the SE moves to the DIS3V3 state and turns off this supply by bringing PDO1 low. It also indicates that a fault has occurred by taking PDO3 high. Timeout delays of from 100 μ s to 400 ms can be programmed.

FAULT REPORTING

The ADM1067 has a fault latch for recording faults. Two registers are set aside for this purpose. A single bit is assigned to each input of the device, and a fault on that input sets the relevant bit. The contents of the fault register can be read out over the SMBus to determine which input(s) faulted. The fault register can be enabled/disabled in each state. This ensures that only real faults are captured and not, for example, undervoltage trips when the SE is executing a power-down sequence.

SUPPLY MARGINING

OVERVIEW

It is often necessary for the system designer to adjust supplies, either to optimize their level or force them away from nominal values to characterize the system performance under these conditions. This is a function typically performed during an in-circuit test (ICT), such as when the manufacturer, for example, wants to guarantee that the product under test functions correctly at nominal supplies minus 10%.

OPEN-LOOP MARGINING

The simplest method of margining a supply is to implement an open-loop technique. A popular method for this is to switch extra resistors into the feedback node of a power module, such as a dc/dc converter or low dropout regulator (LDO). The extra resistor alters the voltage at the feedback or trim node and forces the output voltage to margin up or down by a certain amount.

The ADM1067 can perform open-loop margining for up to six supplies. The six on-board voltage DACs (DAC1–6) can drive into the feedback pins of the power modules to be margined. The simplest circuit to implement this function is an attenuation resistor, which connects the DACn pin to the feedback node of a dc/dc converter. When the DACn output voltage is set equal to the feedback voltage, no current flows in the attenuation resistor, and the dc/dc output voltage does not change. Taking DACn above the feedback voltage forces current into the feedback node, and the output of the dc/dc converter is forced to fall to compensate for this. The dc/dc output can be forced high by setting the DACn output voltage lower than the feedback node voltage. The series resistor can be split in two, and the node between them decoupled with a capacitor to ground. This can help to decouple any noise picked up from the board. Decoupling to a ground local to the dc/dc converter is recommended.




Figure 27. Open-Loop Margining System Using the ADM1067

The ADM1067 can be commanded to margin a supply up or down over the SMBus by updating the values on the relevant DAC output.

To implement open-loop margining:

1. Disable the six DACn outputs.
2. Set the DAC output voltage equal to the voltage on the feedback node.
3. Enable the DAC.
4. Assert MUP (drive logic high). The DAC voltage moves down to the value set in the DNLIM register (see Application Note AN-698). The output of the dc/dc converter rises to compensate for this, that is, margin up.
5. Assert MDN (drive logic high). The DAC voltage moves down to the value set in the DPLIM register (see Application Note AN-698). The output of the dc/dc converter drops to compensate for this, that is, margin down.

Steps 1 to 3 ensure that when the DACn output buffer is turned on it has little effect on the dc/dc output. The DAC output buffer is designed to power up without glitching. It does this by first powering up the buffer to follow the pin voltage and does not drive out onto the pin at this time. Once the output buffer is properly enabled, the buffer input is switched over to the DAC, and the output stage of the buffer is turned on. Output glitching is negligible.

The margining method above assumes that margin up and margin down DAC limits have been preloaded into the ADM106x and that only one margin up and margin down level are required. Alternatively, a DACn output level can be dynamically altered by an SMBus write to that DACn output register.

WRITING TO THE DACs

Four DAC ranges are offered. They can be placed with midcode (Code 0x7F) at 0.6 V, 0.8 V, 1.0 V, and 1.25 V. These voltages are placed to correspond to the most common feedback voltages. Centering the DAC outputs in this way provides the best use of the DAC resolution. For most supplies, it is possible to place the DAC midcode at the point where the dc/dc output is not modified, thereby giving half of the DAC range to margin up and the other half to margin down.

The DAC output voltage is set by the code written to the DACn register. The voltage is linear with the unsigned binary number in this register. Code 0x7F is placed at the midcode voltage, as described previously.

The output voltage is given by the following equation:

$$DAC\ Output = (DACn - 0x7F)/255 \times 0.6015 + V_{OFF}$$

where V_{OFF} is one of the four offset voltages.

There are 256 DAC settings available. The midcode value is located at DAC Code 0x7F as close as possible to the middle of the 256 code range. The full output swing of the DACs is +302 mV (+128 codes) and -300 mV (-127 codes) around the selected midcode voltage. The voltage range for each midcode voltage is shown in Table 8.

Table 8. Ranges for Midcode Voltages

Midcode Voltage (V)	Minimum Voltage Output (V)	Maximum Voltage Output (V)
0.6	0.300	0.902
0.8	0.500	1.102
1.0	0.700	1.302
1.25	0.950	1.552

CHOOSING THE SIZE OF THE ATTENUATION RESISTOR

How much this DAC voltage swing affects the output voltage of the dc/dc converter that is being margined is determined by the size of the attenuation resistor, R3.

Because the voltage at the feedback pin remains constant, the current flowing from the feedback node to GND via R2 is a constant. Also, the feedback node itself is high impedance. This means that the current flowing through R1 is the same as the current flowing through R3. Therefore, direct relationship exists between the extra voltage drop across R1 during margining and the voltage drop across R3.

This relationship is given by the equation

$$\partial V_{OUT} = \frac{R1}{R3} (V_{FB} - V_{DACOUT})$$

where:

∂V_{OUT} is the change in V_{OUT} .

V_{FB} is the voltage at the feedback node of the dc/dc converter.

V_{DACOUT} is the voltage output of the margining DAC.

This equation demonstrates that, if the user wants the output voltage to change by ± 300 mV, then $R1 = R3$. If the user wants the output voltage to change by ± 600 mV, then $R1 = 2 \times R3$, and so on.

It is best to use the full DAC output range to margin a supply. Choosing the attenuation resistor in this way provides the most resolution from the DAC. In other words, with one DAC code change, the smallest effect on the dc/dc output voltage is induced. If the resistor is sized up to use a code such as 27(dec) to 227(dec) to move the dc/dc output by $\pm 5\%$, then it takes 100 codes to move 5% (each code moves the output by 0.05%). This is beyond the readback accuracy of the ADC, but should not prevent the user from building a circuit to use the most resolution.

DAC LIMITING/OTHER SAFETY FEATURES

Limit registers (called DPLIMn and DNLIMn) on the device offer the user some protection from firmware bugs, which can cause catastrophic board problems by forcing supplies beyond their allowable output ranges. Essentially, the DAC code written into the DACn register is clipped such that the code used to set the DAC voltage is actually given by

$$\begin{aligned} \text{DAC Code} &= \text{DACn}, \text{ DACn} \geq \text{DNLIMn and DACn} \leq \text{DPLIMn} \\ &= \text{DNLIMn}, \quad \text{DACn} < \text{DNLIMn} \\ &= \text{DPLIMn}, \quad \text{DACn} > \text{DPLIMn} \end{aligned}$$

In addition, the DAC output buffer is three-stated, if $\text{DNLIMn} > \text{DPLIMn}$. By programming the limit registers in this way, the user can make it very difficult for the DAC output buffers to be turned on at all during normal system operation (these are among the registers downloaded from EEPROM at startup).

APPLICATIONS DIAGRAM




Figure 28. Applications Diagram

COMMUNICATING WITH THE ADM1067

CONFIGURATION DOWNLOAD AT POWER-UP

The configuration of the ADM1067 (UV/OV thresholds, glitch filter timeouts, PDO configurations, and so on) is dictated by the contents of RAM. The RAM is comprised of digital latches that are local to each of the functions on the device. The latches are double-buffered and have two identical latches, Latch A and Latch B. Therefore, when an update to a function occurs, the contents of Latch A are updated first, and then the contents of Latch B are updated with identical data. The advantages of this architecture are explained in detail in this section.

The latches are volatile memory and lose their contents at power-down. Therefore, the configuration in the RAM must be restored at power-up by downloading the contents of the EEPROM (nonvolatile memory) to the local latches. This download occurs in steps, as follows:

1. With no power applied to the device, the PDOs are all high impedance.
2. When 1 V appears on any of the inputs connected to the VDD arbitrator (VH or VPn), the PDOs are all weakly pulled to GND with a 20 kΩ impedance.
3. When the supply rises above the undervoltage lockout of the device (UVLO is 2.5 V), the EEPROM starts to download to the RAM.
4. The EEPROM downloads its contents to all Latch As.
5. Once the contents of the EEPROM are completely downloaded to the Latch As, the device controller signals all Latch As to download to all Latch Bs simultaneously, completing the configuration download.
6. At 0.5 ms after the configuration download completes, the first state definition is downloaded from EEPROM into the SE.

Note that any attempt to communicate with the device prior to the completion of the download causes the ADM1067 to issue a no acknowledge (NACK).

UPDATING THE CONFIGURATION

After power-up, with all the configuration settings loaded from EEPROM into the RAM registers, the user might need to alter the configuration of functions on the ADM1067, such as changing the UV or OV limit of an SFD, changing the fault output of an SFD, or adjusting the rise time delay of one of the PDOs.

The ADM1067 provides several options that allow the user to update the configuration over the SMBus interface. The following options are controlled in the UPDCFG register:

1. Update the configuration in real time. The user writes to RAM across the SMBus and the configuration is updated immediately.
2. Update the Latch As without updating the Latch Bs. With this method, the configuration of the ADM1067 remains unchanged and continues to operate in the original setup until the instruction is given to update the Latch Bs.
3. Change EEPROM register contents without changing the RAM contents, and then download the revised EEPROM contents to the RAM registers. Again, with this method, the configuration of the ADM1067 remains unchanged and continues to operate in the original setup until the instruction is given to update the RAM.

The instruction to download from the EEPROM in Option 3 is also a useful way to restore the original EEPROM contents, if revisions to the configuration are unsatisfactory. For example, if the user needs to alter an OV threshold, this can be done by updating the RAM register as described in Option 1. However, if the user is not satisfied with the change and wants to revert to the original programmed value, then the device controller can issue a command to download the EEPROM contents to the RAM again, as described in Option 3, restoring the ADM1067 to its original configuration.

The topology of the ADM1067 makes this type of operation possible. The local, volatile registers (RAM) are all double-buffered latches. Setting Bit 0 of the UPDCFG register to 1 leaves the double-buffered latches open at all times. If Bit 0 is set to 0, then, when a RAM write occurs across the SMBus, only the first side of the double-buffered latch is written to. The user must then write a 1 to Bit 1 of the UPDCFG register. This generates a pulse to update all the second latches at once. EEPROM writes occur in a similar way.

The final bit in this register can enable or disable EEPROM page erasure. If this bit is set high, the contents of an EEPROM page can all be set to 1. If low, then the contents of a page cannot be erased, even if the command code for page erasure is programmed across the SMBus. The bitmap for the UPDCFG register is shown in the AN-698 application note. A flow chart for download at power-up and subsequent configuration updates is shown in Figure 29.




Figure 29. Configuration Update Flow Diagram

UPDATING THE SEQUENCING ENGINE

Sequencing engine (SE) functions are not updated in the same way as regular configuration latches. The SE has its own dedicated 512-byte EEPROM for storing state definitions, providing 63 individual states with a 64-bit word each (one state is reserved). At power-up, the first state is loaded from the SE EEPROM into the engine itself. When the conditions of this state are met, the next state is loaded from EEPROM into the engine, and so on. The loading of each new state takes approximately 10 μ s.

To alter a state, the required changes must be made directly to EEPROM. RAM for each state does not exist. The relevant alterations must be made to the 64-bit word, which is then uploaded directly to EEPROM.

INTERNAL REGISTERS

The ADM1067 contains a large number of data registers. The principal registers are the address pointer register and the configuration registers.

Address Pointer Register

This register contains the address that selects one of the other internal registers. When writing to the ADM1067, the first byte of data is always a register address, which is written to the address pointer register.

Configuration Registers

These registers provide control and configuration for various operating parameters of the ADM1067.

EEPROM

The ADM1067 has two 512-byte cells of nonvolatile, electrically erasable, programmable read-only memory (EEPROM), from Register Addresses 0xF800 to 0xFBFF. The EEPROM is used for permanent storage of data that is not lost when the ADM1067 is powered down. One EEPROM cell contains the configuration data of the device; the other contains the state definitions for the SE. Although referred to as read-only memory, the EEPROM can be written to as well as read from via the serial bus in exactly the same way as the other registers.

The major differences between the EEPROM and other registers are

- An EEPROM location must be blank before it can be written to. If it contains data, it must first be erased.
- Writing to EEPROM is slower than writing to RAM.
- Writing to the EEPROM should be restricted, because it has a limited write/cycle life of typically 10,000 write operations due to the usual EEPROM wear-out mechanisms.

The first EEPROM is split into 16 (0 to 15) pages of 32 bytes each. Pages 0 to 6, starting at Address 0xF800, hold the configuration data for the applications on the ADM1067 (the SFDs, PDOs, and so on). These EEPROM addresses are the same as the RAM register addresses, prefixed by F8. Page 7 is reserved. Pages 8 to 15 are for customer use.

Data can be downloaded from EEPROM to RAM in one of the following ways:

- At power-up, when Pages 0 to 6 are downloaded.
- By setting Bit 0 of the UDOWNLD register (0xD8), which performs a user download of Pages 0 to 6.

SERIAL BUS INTERFACE

The ADM1067 is controlled via the serial system management bus (SMBus). The ADM1067 is connected to this bus as a slave device, under the control of a master device. It takes approximately 1 ms after power-up for the ADM1067 to download from its EEPROM. Therefore, access to the ADM1067 is restricted until the download is completed.

Identifying the ADM1067 on the SMBus

The ADM1060 has a 7-bit serial bus slave address. The device is powered up with a default serial bus address. The five MSBs of the address are set to 01101; the two LSBs are determined by the logical states of Pins A1 and A0. This allows the connection of four ADM1067s to one SMBus.

The device also has several identification registers (read-only), which can be read across the SMBus. Table 9 lists these registers with their values and functions.

Table 9. Identification Register Values and Functions

Name	Address	Value	Function
MANID	0xF4	0x41	Manufacturer ID for Analog Devices
REVID	0xF5	0x00	Silicon revision
MARK1	0xF6	0x00	S/w brand
MARK2	0xF7	0x00	S/w brand

General SMBus Timing

Figure 30, Figure 31, and Figure 32 are timing diagrams for general read and write operations using the SMBus. The SMBus specification defines specific conditions for different types of read and write operations, which are discussed in the Write Operations and Read Operations sections.

The general SMBus protocol operates as follows:

1. The master initiates data transfer by establishing a start condition, defined as a high-to-low transition on the serial data-line SDA, while the serial clock-line SCL remains high. This indicates that a data stream follows. All slave peripherals connected to the serial bus respond to the start condition and shift in the next 8 bits, consisting of a 7-bit slave address (MSB first) plus a $\overline{R/W}$ bit. This bit determines the direction of the data transfer, that is, whether data is written to or read from the slave device (0 = write, 1 = read).

The peripheral whose address corresponds to the transmitted address responds by pulling the data line low during the low period before the ninth clock pulse, known as the acknowledge bit, and holding it low during the high period of this clock pulse.

All other devices on the bus remain idle while the selected device waits for data to be read from or written to it. If the $\overline{R/W}$ bit is a 0, the master writes to the slave device. If the $\overline{R/W}$ bit is a 1, the master reads from the slave device.

2. Data is sent over the serial bus in sequences of nine clock pulses, eight bits of data followed by an acknowledge bit from the slave device. Data transitions on the data line must occur during the low period of the clock signal and remain stable during the high period, because a low-to-high transition when the clock is high might be interpreted as a stop signal. If the operation is a write operation, the first data byte after the slave address is a command byte. This tells the slave device what to expect next. It might be an instruction telling the slave device to expect a block write, or it might simply be a register address that tells the slave where subsequent data is to be written. Because data can flow in only one direction, as defined by the $R/\overline{W}$ bit, sending a command to a slave device during a read operation is not possible. Before a read operation, it might be necessary to perform a write operation to tell the slave what sort of read operation to expect and/or the address from which data is to be read.
3. When all data bytes have been read or written, stop conditions are established. In write mode, the master pulls the data line high during the 10th clock pulse to assert a stop condition. In read mode, the master device releases the SDA line during the low period before the ninth clock pulse, but the slave device does not pull it low. This is known as no acknowledge. The master then takes the data line low during the low period before the tenth clock pulse, then high during the tenth clock pulse to assert a stop condition.




Figure 30. General SMBus Write Timing Diagram




Figure 31. General SMBus Read Timing Diagram




Figure 32. Serial Bus Timing Diagram

SMBus Protocols for RAM and EEPROM

The ADM1067 contains volatile registers (RAM) and nonvolatile registers (EEPROM). User RAM occupies address locations from 0x00 to 0xDF; EEPROM occupies addresses from 0xF800 to 0xFBFF.

Data can be written to and read from both RAM and EEPROM as single data bytes. Data can be written only to unprogrammed EEPROM locations. To write new data to a programmed location, it must first be erased. EEPROM erasure cannot be done at the byte level. The EEPROM is arranged as 32 pages of 32 bytes each, and an entire page must be erased.

Page erasure is enabled by setting Bit 2 in the UPDCFG register (Address 0x90) to 1. If this bit is not set, page erasure cannot occur, even if the command byte (0xFE) is programmed across the SMBus.

WRITE OPERATIONS

The SMBus specification defines several protocols for different types of read and write operations. The following abbreviations are used in the diagrams:

- S Start
- P Stop
- R Read
- W Write
- A Acknowledge
- $\overline{A}$ No acknowledge

The ADM1067 uses the following SMBus write protocols.

Send Byte

In a send byte operation, the master device sends a single command byte to a slave device, as follows:

1. The master device asserts a start condition on SDA.
2. The master sends the 7-bit slave address followed by the write bit (low).
3. The addressed slave device asserts ACK on SDA.
4. The master sends a command code.
5. The slave asserts ACK on SDA.
6. The master asserts a stop condition on SDA and the transaction ends.

In the ADM1067, the send byte protocol is used for two purposes:

- To write a register address to RAM for a subsequent single byte read from the same address, or a block read or write starting at that address, as shown in Figure 33.




Figure 33. Setting a RAM Address for Subsequent Read

- To erase a page of EEPROM memory. EEPROM memory can be written to only if it is unprogrammed. Before writing to one or more EEPROM memory locations that are already programmed, the page or pages containing those locations must first be erased. EEPROM memory is erased by writing a command byte.

The master sends a command code that tells the slave device to erase the page. The ADM1067 command code for a page erasure is 0xFE (1111 1110). Note that, for a page erasure to take place, the page address has to be given in the previous write word transaction (see the Write Byte/Word section). Also, Bit 2 in the UPDCFG register (Address 0x90) must be set to 1.




Figure 34. EEPROM Page Erasure

As soon as the ADM1067 receives the command byte, page erasure begins. The master device can send a stop command as soon as it sends the command byte. Page erasure takes approximately 20 ms. If the ADM1067 is accessed before erasure is complete, it responds with a no acknowledge (NACK).

Write Byte/Word

In a write byte/word operation, the master device sends a command byte and one or two data bytes to the slave device, as follows:

- The master device asserts a start condition on SDA.
- The master sends the 7-bit slave address followed by the write bit (low).
- The addressed slave device asserts ACK on SDA.
- The master sends a command code.
- The slave asserts ACK on SDA.
- The master sends a data byte.
- The slave asserts ACK on SDA.
- The master sends a data byte (or asserts a stop condition at this point).
- The slave asserts ACK on SDA.
- The master asserts a stop condition on SDA to end the transaction.

In the ADM1067, the write byte/word protocol is used for three purposes:

- To write a single byte of data to RAM. In this case, the command byte is the RAM address from 0x00 to 0xDF and the only data byte is the actual data, as shown in Figure 35.




Figure 35. Single Byte Write to RAM

- To set up a 2-byte EEPROM address for a subsequent read, write, block read, block write, or page erase. In this case, the command byte is the high byte of the EEPROM address from 0xF8 to 0xFB. The only data byte is the low byte of the EEPROM address, as shown in Figure 36.




Figure 36. Setting an EEPROM Address

Note, for page erasure, that because a page consists of 32 bytes, only the three MSBs of the address low byte are important. The lower five bits of the EEPROM address low byte specify the addresses within a page and are ignored during an erase operation.

- To write a single byte of data to EEPROM. In this case, the command byte is the high byte of the EEPROM address from 0xF8 to 0xFB. The first data byte is the low byte of the EEPROM address, and the second data byte is the actual data, as shown in Figure 37.




Figure 37. Single Byte Write to EEPROM

Block Write

In a block write operation, the master device writes a block of data to a slave device. The start address for a block write must have been set previously. In the ADM1067, a send byte operation sets a RAM address, and a write byte/word operation sets an EEPROM address, as follows:

- The master device asserts a start condition on SDA.
- The master sends the 7-bit slave address followed by the write bit (low).
- The addressed slave device asserts ACK on SDA.
- The master sends a command code that tells the slave device to expect a block write. The ADM1067 command code for a block write is 0xFC (1111 1100).

- The slave asserts ACK on SDA.
- The master sends a data byte that tells the slave device how many data bytes are being sent. The SMBus specification allows a maximum of 32 data bytes in a block write.
- The slave asserts ACK on SDA.
- The master sends N data bytes.
- The slave asserts ACK on SDA after each data byte.
- The master asserts a stop condition on SDA to end the transaction.




Figure 38. Block Write to EEPROM or RAM

Unlike some EEPROM devices that limit block writes to within a page boundary, there is no limitation on the start address when performing a block write to EEPROM, except

- There must be at least N locations from the start address to the highest EEPROM address (0xFBFF), to avoid writing to invalid addresses.
- If the addresses cross a page boundary, both pages must be erased before programming.

Note that the ADM1067 features a clock extend function for writes to EEPROM. Programming an EEPROM byte takes approximately 250 μ s, which would limit the SMBus clock for repeated or block write operations. The ADM1067 pulls SCL low and extends the clock pulse when it cannot accept any more data.

READ OPERATIONS

The ADM1067 uses the following SMBus read protocols.

Receive Byte

In a receive byte operation, the master device receives a single byte from a slave device, as follows:

- The master device asserts a start condition on SDA.
- The master sends the 7-bit slave address followed by the read bit (high).
- The addressed slave device asserts ACK on SDA.
- The master receives a data byte.
- The master asserts no acknowledge on SDA.
- The master asserts a stop condition on SDA, and the transaction ends.

In the ADM1067, the receive byte protocol is used to read a single byte of data from a RAM or EEPROM location whose address has previously been set by a send byte or write byte/word operation, as shown in Figure 39.




Figure 39. Single Byte Read from EEPROM or RAM

Block Read

In a block read operation, the master device reads a block of data from a slave device. The start address for a block read must have been set previously. In the ADM1067, this is done by a send byte operation to set a RAM address, or a write byte/word operation to set an EEPROM address. The block read operation itself consists of a send byte operation that sends a block read command to the slave, immediately followed by a repeated start and a read operation that reads out multiple data bytes, as follows:

- The master device asserts a start condition on SDA.
- The master sends the 7-bit slave address followed by the write bit (low).
- The addressed slave device asserts ACK on SDA.
- The master sends a command code that tells the slave device to expect a block read. The ADM1067 command code for a block read is 0xFD (1111 1101).
- The slave asserts ACK on SDA.
- The master asserts a repeat start condition on SDA.
- The master sends the 7-bit slave address followed by the read bit (high).
- The slave asserts ACK on SDA.
- The ADM1067 sends a byte-count data byte that tells the master how many data bytes to expect. The ADM1067 always returns 32 data bytes (0x20), which is the maximum allowed by the SMBus 1.1 specification.
- The master asserts ACK on SDA.
- The master receives 32 data bytes.
- The master asserts ACK on SDA after each data byte.
- The master asserts a stop condition on SDA to end the transaction.




Figure 40. Block Read from EEPROM or RAM

Error Correction

The ADM1067 provides the option of issuing a PEC (packet error correction) byte after a write to RAM, a write to EEPROM, a block write to RAM/EEPROM, or a block read from RAM/EEPROM. This enables the user to verify that the data received by or sent from the ADM1067 is correct. The PEC byte is an optional byte sent after that last data byte has been written to or read from the ADM1067. The protocol is as follows:

- 1. The ADM1067 issues a PEC byte to the master. The master checks the PEC byte and issues another block read, if the PEC byte is incorrect.

- 2. A no acknowledge (NACK) is generated after the PEC byte to signal the end of the read.

Note that the PEC byte is calculated using CRC-8. The frame check sequence (FCS) conforms to CRC-8 by the polynomial

$C(x) = x^8 + x^2 + x^1 + 1$

See the SMBus 1.1 specification for details.

An example of a block read with the optional PEC byte is shown in Figure 41.




Figure 41. Block Read from EEPROM or RAM with PEC

OUTLINE DIMENSIONS




Figure 42. 40-Lead Lead Frame Chip Scale Package [LFCSP]
(CP-40)
Dimensions shown in millimeters




Figure 43. 48-Lead Thin Plastic Quad Flat Package [TQFP]
(SU-48)
Dimensions shown in millimeters

ORDERING GUIDE

Model	Temperature Range	Package Description	Package Option
ADM1067ACP	−40°C to +85°C	40-Lead LFCSP	CP-40
ADM1067ACP-REEL	−40°C to +85°C	40-Lead LFCSP	CP-40
ADM1067ACP-REEL7	−40°C to +85°C	40-Lead LFCSP	CP-40
ADM1067ASU	−40°C to +85°C	48-Lead TQFP	SU-48
ADM1067ASU-REEL	−40°C to +85°C	48-Lead TQFP	SU-48
ADM1067ASU-REEL7	−40°C to +85°C	48-Lead TQFP	SU-48
EVAL-ADM1067LFEB		ADM1067 Evaluation Kit (LFSCP Version)	
EVAL-ADM1067TQEB		ADM1067 Evaluation Kit (TQFP Version)	

NOTES

ADM1067

NOTES