

The I²C-bus and how to use it (including specifications)

1995 update

1.0 THE I²C-BUS BENEFITS DESIGNERS AND MANUFACTURERS

In consumer electronics, telecommunications and industrial electronics, there are often many similarities between seemingly unrelated designs. For example, nearly every system includes:

- Some intelligent control, usually a single-chip microcontroller
- General-purpose circuits like LCD drivers, remote I/O ports, RAM, EEPROM, or data converters
- Application-oriented circuits such as digital tuning and signal processing circuits for radio and video systems, or DTMF generators for telephones with tone dialling.

To exploit these similarities to the benefit of both systems designers and equipment manufacturers, as well as to maximize hardware efficiency and circuit simplicity, Philips developed a simple bidirectional 2-wire bus for efficient inter-IC control. This bus is called the Inter IC or I²C-bus. At present, Philips' IC range includes more than 150 CMOS and bipolar I²C-bus compatible types for performing functions in all three of the previously mentioned categories. All I²C-bus compatible devices incorporate an on-chip interface which allows them to communicate directly with each other via the I²C-bus. This design concept solves the many interfacing problems encountered when designing digital control circuits.

Here are some of the features of the I²C-bus:

- Only two bus lines are required; a serial data line (SDA) and a serial clock line (SCL)
- Each device connected to the bus is software addressable by a unique address and simple master/ slave relationships exist at all times; masters can operate as master-transmitters or as master-receivers
- It's a true multi-master bus including collision detection and arbitration to prevent data corruption if two or more masters simultaneously initiate data transfer
- Serial, 8-bit oriented, bidirectional data transfers can be made at up to 100 kbit/s in the standard mode or up to 400 kbit/s in the fast mode
- On-chip filtering rejects spikes on the bus data line to preserve data integrity
- The number of ICs that can be connected to the same bus is limited only by a maximum bus capacitance of 400 pF.

Figure 1 shows two examples of I²C-bus applications.

1.1 Designer benefits

I²C-bus compatible ICs allow a system design to rapidly progress directly from a functional block diagram to a prototype. Moreover, since they 'clip' directly onto the I²C-bus without any additional external interfacing, they allow a prototype system to be modified or upgraded simply by 'clipping' or 'unclipping' ICs to or from the bus.

Here are some of the features of I²C-bus compatible ICs which are particularly attractive to designers:

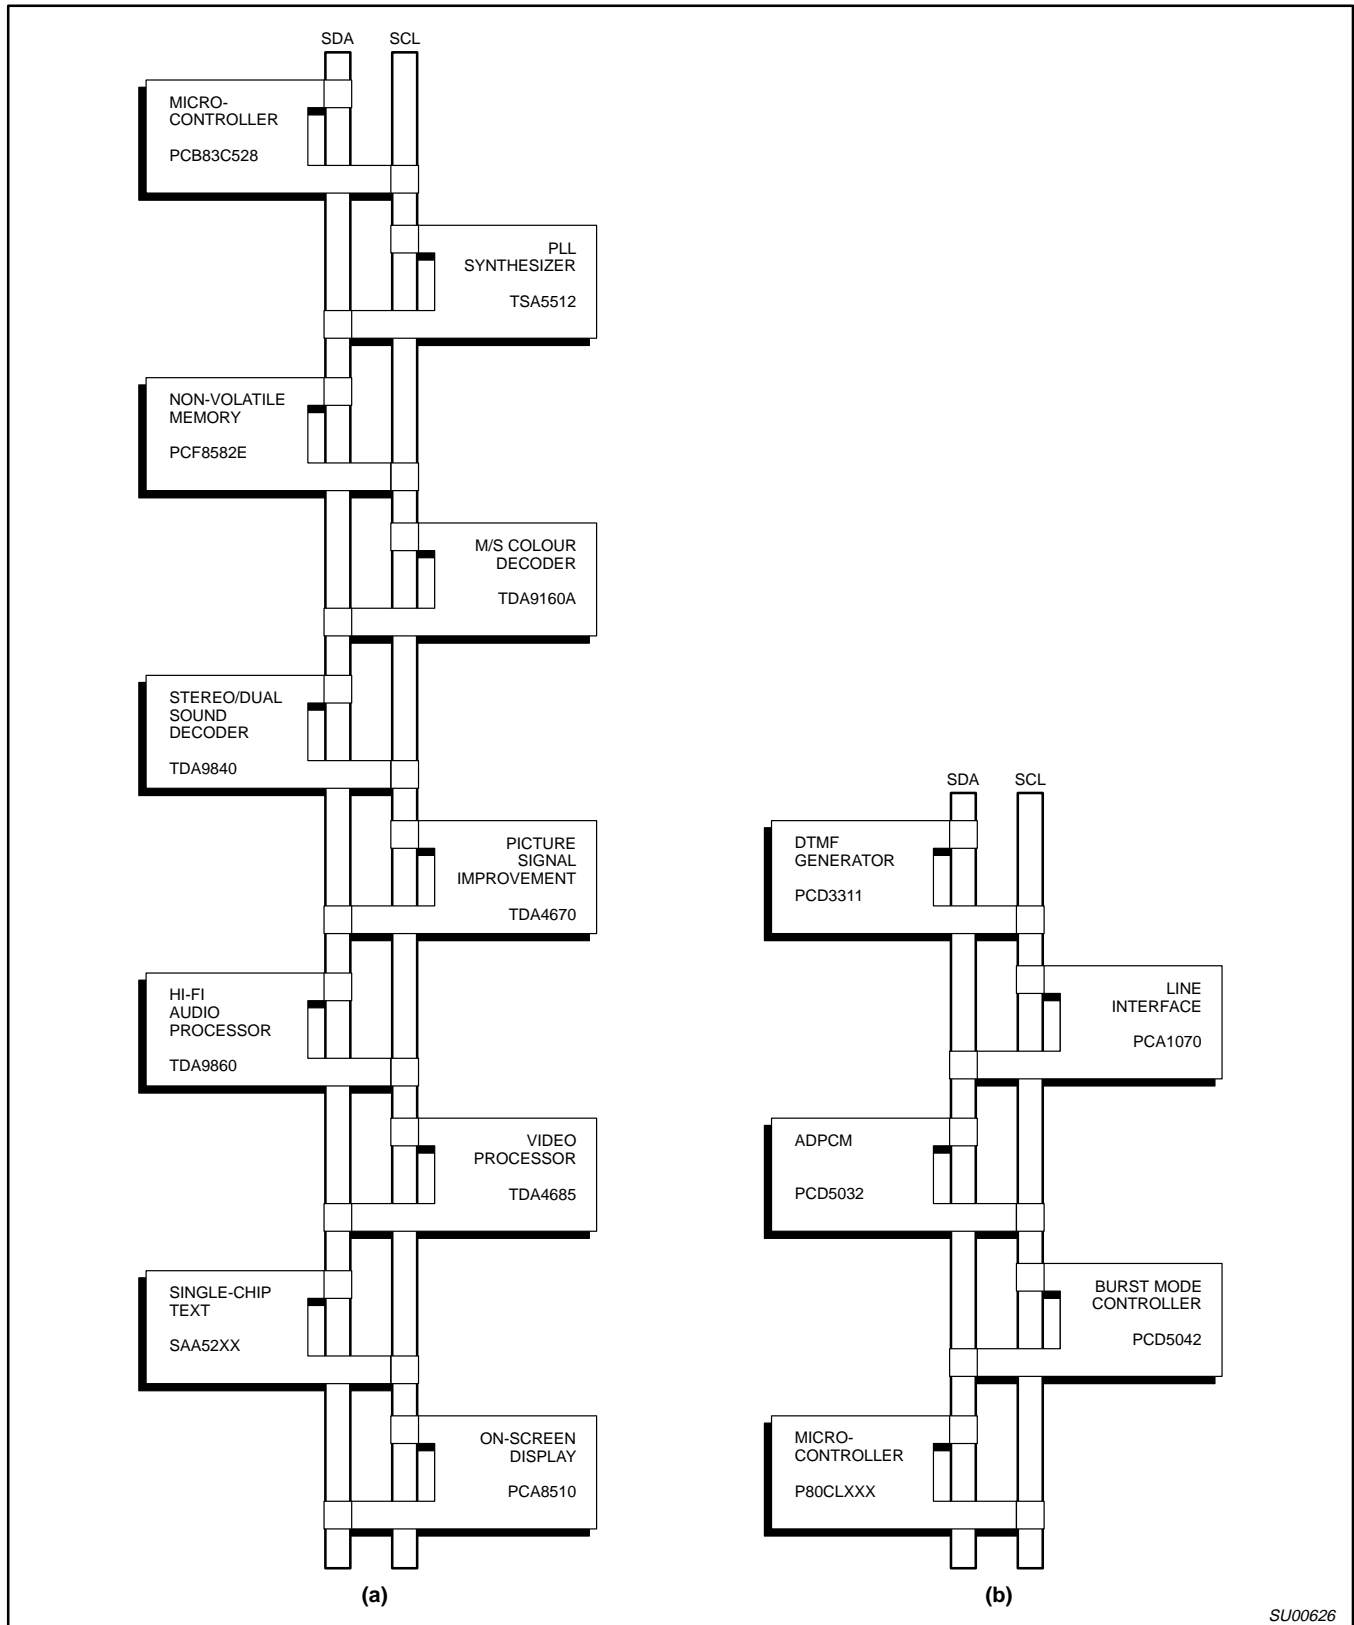
- Functional blocks on the block diagram correspond with the actual ICs; designs proceed rapidly from block diagram to final schematic
- No need to design bus interfaces because the I²C-bus interface is already integrated on-chip
- Integrated addressing and data-transfer protocol allow systems to be completely software-defined
- The same IC types can often be used in many different applications
- Design-time reduces as designers quickly become familiar with the frequently used functional blocks represented by I²C-bus compatible ICs
- ICs can be added to or removed from a system without affecting any other circuits on the bus
- Fault diagnosis and debugging are simple; malfunctions can be immediately traced
- Software development time can be reduced by assembling a library of reusable software modules.

In addition to these advantages, the CMOS ICs in the I²C-bus compatible range offer designers special features which are particularly attractive for portable equipment and battery-backed systems.

They all have:

- Extremely low current consumption
- High noise immunity
- Wide supply voltage range
- Wide operating temperature range.

The I²C-bus and how to use it (including specifications)



SU00626

Figure 1. Two examples of I²C-bus applications
(a) a high performance highly-integrated TV set; (b) DECT cordless phone base-station

The I²C-bus and how to use it (including specifications)

1.2 Manufacturer benefits

I²C-bus compatible ICs don't only assist designers, they also give a wide range of benefits to equipment manufacturers because:


- The simple 2-wire serial I²C-bus minimizes interconnections so ICs have fewer pins and there are not so many PCB tracks; result — smaller and less expensive PCBs
- The completely integrated I²C-bus protocol eliminates the need for address decoders and other 'glue logic'
- The multi-master capability of the I²C-bus allows rapid testing and alignment of end-user equipment via external connections to an assembly-line computer
- The availability of I²C-bus compatible ICs in SO (small outline), VSO (very small outline) as well as DIL packages reduces space requirements even more.

These are just some of the benefits. In addition, I²C-bus compatible ICs increase system design flexibility by allowing simple

construction of equipment variants and easy upgrading to keep designs up-to-date. In this way, an entire family of equipment can be developed around a basic model. Upgrades for new equipment, or enhanced-feature models (i.e. extended memory, remote control, etc.) can then be produced simply by clipping the appropriate ICs onto the bus. If a larger ROM is needed, it's simply a matter of selecting a microcontroller with a larger ROM from our comprehensive range. As new ICs supersede older ones, it's easy to add new features to equipment or to increase its performance by simply unclipping the outdated IC from the bus and clipping on its successor.

1.3 The ACCESS.bus

Another attractive feature of the I²C-bus for designers and manufacturers is that its simple 2-wire nature and capability of software addressing make it an ideal platform for the ACCESS.bus (Figure 2). This is a lower-cost alternative for an RS-232C interface for connecting peripherals to a host computer via a simple 4-pin connector (see Section 19.0).



SU00311A

Figure 2. The ACCESS.bus — a low-cost alternative to an RS-232C interface

The I²C-bus and how to use it (including specifications)

2.0 INTRODUCTION TO THE I²C-BUS SPECIFICATION

For 8-bit digital control applications, such as those requiring microcontrollers, certain design criteria can be established:

- A complete system usually consists of at least one microcontroller and other peripheral devices such as memories and I/O expanders
- The cost of connecting the various devices within the system must be minimized
- A system that performs a control function doesn't require high-speed data transfer
- Overall efficiency depends on the devices chosen and the nature of the interconnecting bus structure.

In order to produce a system to satisfy these criteria, a serial bus structure is needed. Although serial buses don't have the throughput capability of parallel buses, they do require less wiring and fewer IC connecting pins. However, a bus is not merely an interconnecting wire, it embodies all the formats and procedures for communication within the system.

Devices communicating with each other on a serial bus must have some form of protocol which avoids all possibilities of confusion,

data loss and blockage of information. Fast devices must be able to communicate with slow devices. The system must not be dependent on the devices connected to it, otherwise modifications or improvements would be impossible. A procedure has also to be devised to decide which device will be in control of the bus and when. And, if different devices with different clock speeds are connected to the bus, the bus clock source must be defined. All these criteria are involved in the specification of the I²C-bus.

3.0 THE I²C-BUS CONCEPT

The I²C-bus supports any IC fabrication process (NMOS, CMOS, bipolar). Two wires, serial data (SDA) and serial clock (SCL), carry information between the devices connected to the bus. Each device is recognised by a unique address — whether it's a microcontroller, LCD driver, memory or keyboard interface — and can operate as either a transmitter or receiver, depending on the function of the device. Obviously an LCD driver is only a receiver, whereas a memory can both receive and transmit data. In addition to transmitters and receivers, devices can also be considered as masters or slaves when performing data transfers (see Table 1). A master is the device which initiates a data transfer on the bus and generates the clock signals to permit that transfer. At that time, any device addressed is considered a slave.

Table 1. Definition of I²C-bus terminology

TERM	DESCRIPTION
Transmitter	The device which sends the data to the bus
Receiver	The device which receives the data from the bus
Master	The device which initiates a transfer, generates clock signals and terminates a transfer
Slave	The device addressed by a master
Multi-master	More than one master can attempt to control the bus at the same time without corrupting the message
Arbitration	Procedure to ensure that, if more than one master simultaneously tries to control the bus, only one is allowed to do so and the message is not corrupted
Synchronization	Procedure to synchronize the clock signals of two or more devices

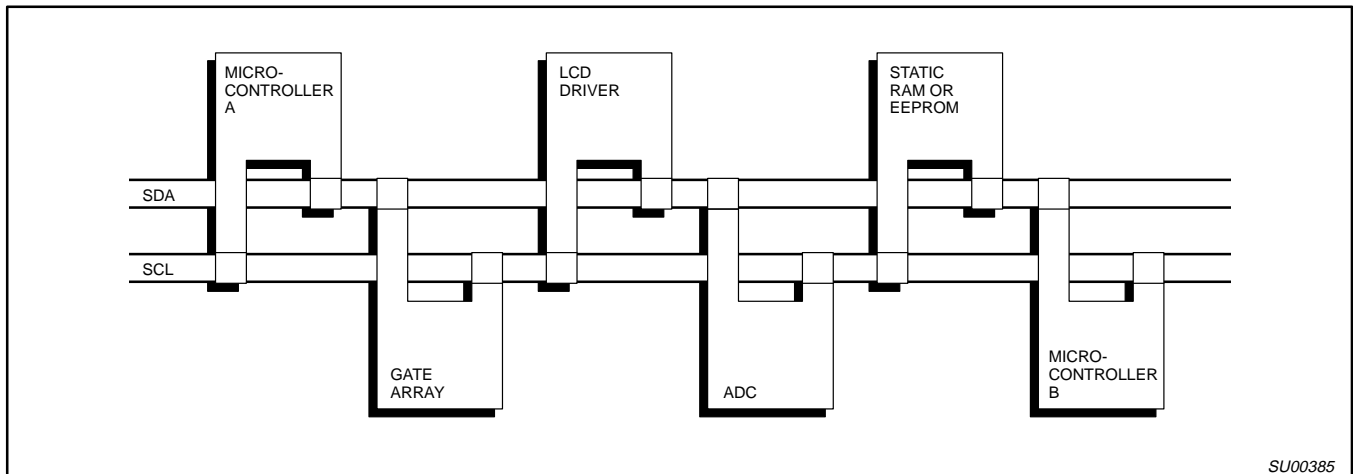


Figure 3. Example of an I²C-bus configuration using two microcontrollers

The I²C-bus and how to use it (including specifications)

The I²C-bus is a multi-master bus. This means that more than one device capable of controlling the bus can be connected to it. As masters are usually micro-controllers, let's consider the case of a data transfer between two micro-controllers connected to the I²C-bus (Figure 3). This highlights the master-slave and receiver-transmitter relationships to be found on the I²C-bus. It should be noted that these relationships are not permanent, but only depend on the direction of data transfer at that time. The transfer of data would proceed as follows:

1. Suppose microcontroller A wants to send information to microcontroller B:
 - microcontroller A (master), addresses microcontroller B (slave)
 - microcontroller A (master-transmitter), sends data to microcontroller B (slave-receiver)
 - microcontroller A terminates the transfer.
2. If microcontroller A wants to receive information from microcontroller B:
 - microcontroller A (master) addresses microcontroller B (slave)
 - microcontroller A (master-receiver) receives data from microcontroller B (slave-transmitter)
 - microcontroller A terminates the transfer.

Even in this case, the master (microcontroller A) generates the timing and terminates the transfer.

The possibility of connecting more than one microcontroller to the I²C-bus means that more than one master could try to initiate a data

transfer at the same time. To avoid the chaos that might ensue from such an event — an arbitration procedure has been developed. This procedure relies on the wired-AND connection of all I²C interfaces to the I²C-bus.

If two or more masters try to put information onto the bus, the first to produce a 'one' when the other produces a 'zero' will lose the arbitration. The clock signals during arbitration are a synchronized combination of the clocks generated by the masters using the wired-AND connection to the SCL line (for more detailed information concerning arbitration see Section 7.0).

Generation of clock signals on the I²C-bus is always the responsibility of master devices; each master generates its own clock signals when transferring data on the bus. Bus clock signals from a master can only be altered when they are stretched by a slow-slave device holding-down the clock line, or by another master when arbitration occurs.

4.0 GENERAL CHARACTERISTICS

Both SDA and SCL are bidirectional lines, connected to a positive supply voltage via a pull-up resistor (see Figure 4). When the bus is free, both lines are HIGH. The output stages of devices connected to the bus must have an open-drain or open-collector in order to perform the wired-AND function. Data on the I²C-bus can be transferred at a rate up to 100 kbit/s in the standard-mode, or up to 400 kbit/s in the fast-mode. The number of interfaces connected to the bus is solely dependent on the bus capacitance limit of 400pF.

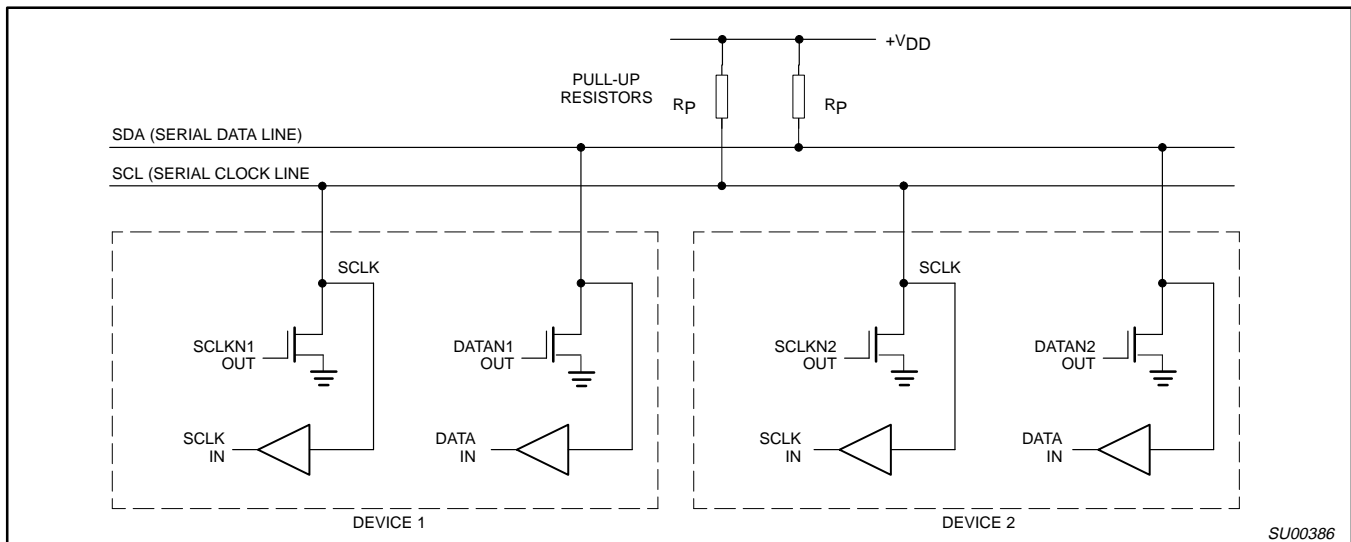


Figure 4. Connection of I²C-bus devices to the I²C-bus

SU00386

The I²C-bus and how to use it (including specifications)




Figure 5. Bit transfer on the I²C-bus




Figure 6. START and STOP conditions

5.0 BIT TRANSFER

Due to the variety of different technology devices (CMOS, NMOS, bipolar) which can be connected to the I²C-bus, the levels of the logical '0' (LOW) and '1' (HIGH) are not fixed and depend on the associated level of V_{DD} (see Section 15.0 for Electrical Specifications). One clock pulse is generated for each data bit transferred.

5.1 Data validity

The data on the SDA line must be stable during the HIGH period of the clock. The HIGH or LOW state of the data line can only change when the clock signal on the SCL line is LOW (see Figure 5).

5.2 START and STOP conditions

Within the procedure of the I²C-bus, unique situations arise which are defined as START and STOP conditions (see Figure 6).

A HIGH to LOW transition on the SDA line while SCL is HIGH is one such unique case. This situation indicates a START condition.

A LOW to HIGH transition on the SDA line while SCL is HIGH defines a STOP condition.

START and STOP conditions are always generated by the master. The bus is considered to be busy after the START condition. The bus is considered to be free again a certain time after the STOP condition. This bus free situation is specified in Section 15.0.

Detection of START and STOP conditions by devices connected to the bus is easy if they incorporate the necessary interfacing hardware. However, microcontrollers with no such interface have to sample the SDA line at least twice per clock period in order to sense the transition.

The I²C-bus and how to use it (including specifications)

6.0 TRANSFERRING DATA

6.1 Byte format

Every byte put on the SDA line must be 8-bits long. The number of bytes that can be transmitted per transfer is unrestricted. Each byte has to be followed by an acknowledge bit. Data is transferred with the most significant bit (MSB) first (Figure 7). If a receiver can't receive another complete byte of data until it has performed some other function, for example servicing an internal interrupt, it can hold the clock line SCL LOW to force the transmitter into a wait state. Data transfer then continues when the receiver is ready for another byte of data and releases clock line SCL.

In some cases, it's permitted to use a different format from the I²C-bus format (for CBUS compatible devices for example). A message which starts with such an address can be terminated by generation of a STOP condition, even during the transmission of a byte. In this case, no acknowledge is generated (see Section 9.1.3).

6.2 Acknowledge

Data transfer with acknowledge is obligatory. The acknowledge-related clock pulse is generated by the master. The transmitter releases the SDA line (HIGH) during the acknowledge clock pulse.

The receiver must pull down the SDA line during the acknowledge clock pulse so that it remains stable LOW during the HIGH period of this clock pulse (Figure 8). Of course, set-up and hold times (specified in Section 15.0) must also be taken into account.

Usually, a receiver which has been addressed is obliged to generate an acknowledge after each byte has been received, except when the message starts with a CBUS address (see Section 9.1.3).

When a slave-receiver doesn't acknowledge the slave address (for example, it's unable to receive because it's performing some real-time function), the data line must be left HIGH by the slave. The master can then generate a STOP condition to abort the transfer.

If a slave-receiver does acknowledge the slave address but, some time later in the transfer cannot receive any more data bytes, the master must again abort the transfer. This is indicated by the slave generating the not acknowledge on the first byte to follow. The slave leaves the data line HIGH and the master generates the STOP condition.

If a master-receiver is involved in a transfer, it must signal the end of data to the slave-transmitter by not generating an acknowledge on the last byte that was clocked out of the slave. The slave-transmitter must release the data line to allow the master to generate a STOP or repeated START condition.




Figure 7. Data transfer on the I²C-bus

SU00363




Figure 8. Acknowledge on the I²C-bus

SU00387

The I²C-bus and how to use it (including specifications)

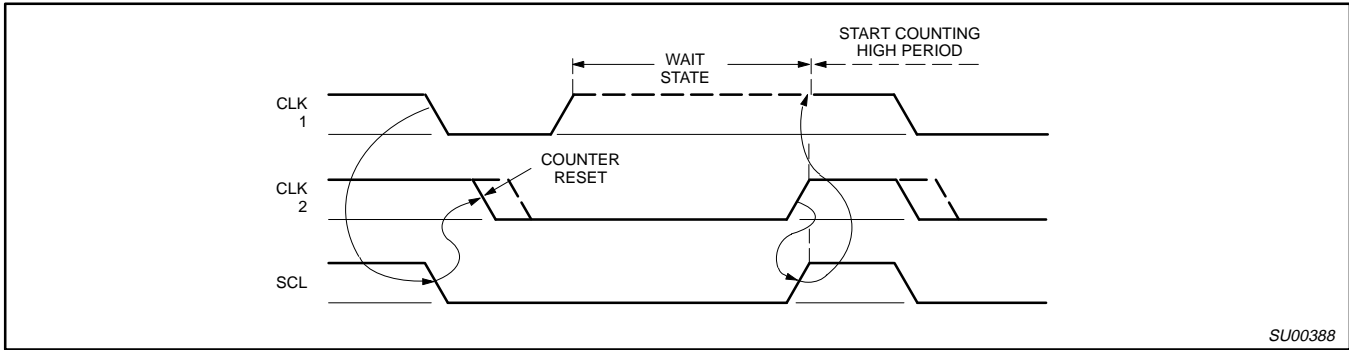


Figure 9. Clock synchronization during the arbitration procedure

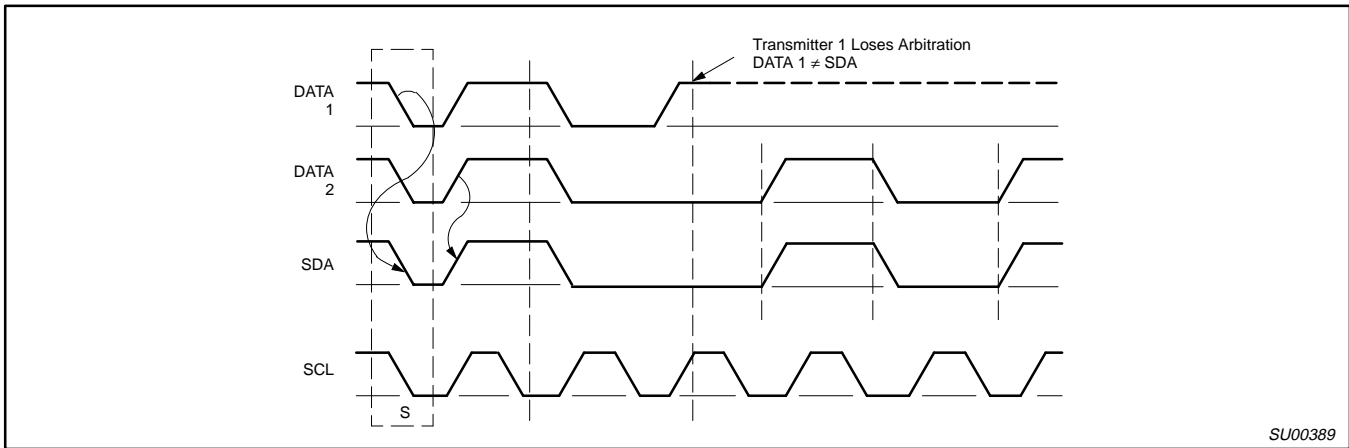


Figure 10. Arbitration procedure of two masters

7.0 ARBITRATION AND CLOCK GENERATION

7.1 Synchronization

All masters generate their own clock on the SCL line to transfer messages on the I²C-bus. Data is only valid during the HIGH period of the clock. A defined clock is therefore needed for the bit-by-bit arbitration procedure to take place.

Clock synchronization is performed using the wired-AND connection of I²C interfaces to the SCL line. This means that a HIGH to LOW transition on the SCL line will cause the devices concerned to start counting off their LOW period and, once a device clock has gone LOW, it will hold the SCL line in that state until the clock HIGH state is reached (Figure 9). However, the LOW to HIGH transition of this clock may not change the state of the SCL line if another clock is still within its LOW period. The SCL line will therefore be held LOW by the device with the longest LOW period. Devices with shorter LOW periods enter a HIGH wait-state during this time.

When all devices concerned have counted off their LOW period, the clock line will be released and go HIGH. There will then be no difference between the device clocks and the state of the SCL line, and all the devices will start counting their HIGH periods. The first device to complete its HIGH period will again pull the SCL line LOW.

In this way, a synchronized SCL clock is generated with its LOW period determined by the device with the longest clock LOW period, and its HIGH period determined by the one with the shortest clock HIGH period.

7.2 Arbitration

A master may start a transfer only if the bus is free. Two or more masters may generate a START condition within the minimum hold time ($t_{HD;STA}$) of the START condition which results in a defined START condition to the bus.

Arbitration takes place on the SDA line, while the SCL line is at the HIGH level, in such a way that the master which transmits a HIGH level, while another master is transmitting a LOW level will switch off its DATA output stage because the level on the bus doesn't correspond to its own level.

Arbitration can continue for many bits. Its first stage is comparison of the address bits (addressing information is in Sections 9.0 and 13.0). If the masters are each trying to address the same device, arbitration continues with comparison of the data. Because address and data information on the I²C-bus is used for arbitration, no information is lost during this process.

A master which loses the arbitration can generate clock pulses until the end of the byte in which it loses the arbitration.

If a master also incorporates a slave function and it loses arbitration during the addressing stage, it's possible that the winning master is trying to address it. The losing master must therefore switch over immediately to its slave-receiver mode.

Figure 10 shows the arbitration procedure for two masters. Of course, more may be involved (depending on how many masters are connected to the bus). The moment there is a difference

The I²C-bus and how to use it (including specifications)

between the internal data level of the master generating DATA 1 and the actual level on the SDA line, its data output is switched off, which means that a HIGH output level is then connected to the bus. This will not affect the data transfer initiated by the winning master.

Since control of the I²C-bus is decided solely on the address and data sent by competing masters, there is no central master, nor any order of priority on the bus.

Special attention must be paid if, during a serial transfer, the arbitration procedure is still in progress at the moment when a repeated START condition or a STOP condition is transmitted to the I²C-bus. If it's possible for such a situation to occur, the masters involved must send this repeated START condition or STOP condition at the same position in the format frame. In other words, arbitration isn't allowed between:

- A repeated START condition and a data bit
- A STOP condition and a data bit
- A repeated START condition and a STOP condition.

7.3 Use of the clock synchronizing mechanism as a handshake

In addition to being used during the arbitration procedure, the clock synchronization mechanism can be used to enable receivers to cope with fast data transfers, on either a byte level or a bit level.

On the byte level, a device may be able to receive bytes of data at a fast rate, but needs more time to store a received byte or prepare another byte to be transmitted. Slaves can then hold the SCL line LOW after reception and acknowledgement of a byte to force the master into a wait state until the slave is ready for the next byte transfer in a type of handshake procedure.

On the bit level, a device such as a microcontroller without, or with only a limited hardware I²C interface on-chip can slow down the bus clock by extending each clock LOW period. The speed of any master is thereby adapted to the internal operating rate of this device.

8.0 FORMATS WITH 7-BIT ADDRESSES

Data transfers follow the format shown in Figure 11. After the START condition (S), a slave address is sent. This address is 7 bits long followed by an eighth bit which is a data direction bit (R/W) — a 'zero' indicates a transmission (WRITE), a 'one' indicates a request for data (READ). A data transfer is always terminated by a STOP condition (P) generated by the master. However, if a master still wishes to communicate on the bus, it can generate a repeated START condition (Sr) and address another slave without first generating a STOP condition. Various combinations of read/write formats are then possible within such a transfer.

Possible data transfer formats are:

- **Master-transmitter transmits to slave-receiver. The transfer direction is not changed (Figure 12)**
- **Master reads slave immediately after first byte (Figure 13).** At the moment of the first acknowledge, the master-transmitter becomes a master-receiver and the slave-receiver becomes a slave-transmitter. This acknowledge is still generated by the slave. The STOP condition is generated by the master
- **Combined format (Figure 14).** During a change of direction within a transfer, the START condition and the slave address are both repeated, but with the R/W bit reversed. If a master receiver sends a repeated START condition, it has previously sent a not acknowledge (A).

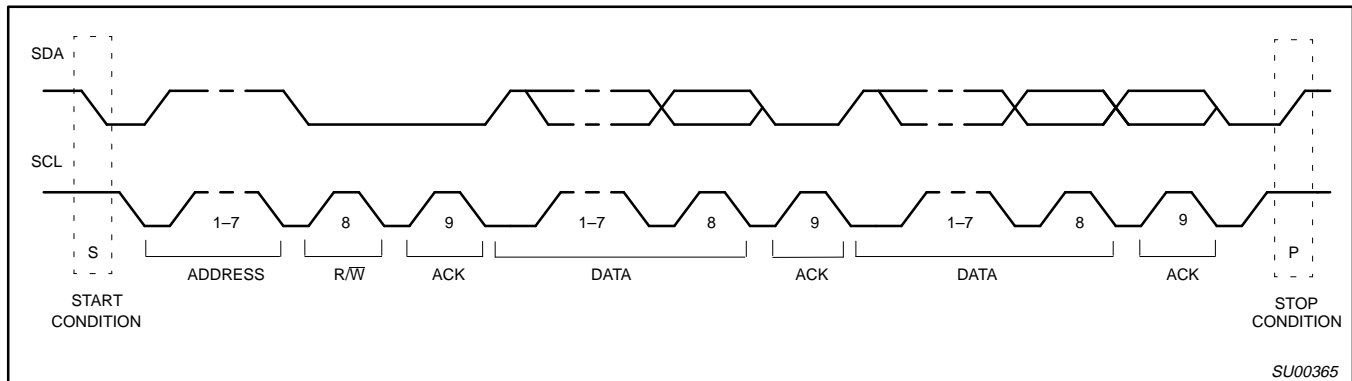


Figure 11. A complete data transfer

The I²C-bus and how to use it (including specifications)




Figure 12. A master-transmitter addresses a slave receiver with a 7-bit address. The transfer direction is not changed




Figure 13. A master reads a slave immediately after the first byte




Figure 14. Combined format

NOTES:

1. Combined formats can be used, for example, to control a serial memory. During the first data byte, the internal memory location has to be written. After the START condition and slave address is repeated, data can be transferred.
2. All decisions on auto-increment or decrement of previously accessed memory locations etc. are taken by the designer of the device.
3. Each byte is followed by an acknowledgement bit as indicated by the A or A-bar blocks in the sequence.
4. I²C-bus compatible devices must reset their bus logic on receipt of a START or repeated START condition such that they all anticipate the sending of a slave address.

The I²C-bus and how to use it (including specifications)

9.0 7-BIT ADDRESSING

(SEE SECTION 13.0 FOR 10-BIT ADDRESSING)

The addressing procedure for the I²C-bus is such that the first byte after the START condition usually determines which slave will be selected by the master. The exception is the 'general call' address which can address all devices. When this address is used, all devices should, in theory, respond with an acknowledge. However, devices can be made to ignore this address. The second byte of the general call address then defines the action to be taken. This procedure is explained in more detail in Section 9.1.1.

9.1 Definition of bits in the first byte

The first seven bits of the first byte make up the slave address (Figure 15). The eighth bit is the LSB (least significant bit). It determines the direction of the message. A 'zero' in the least significant position of the first byte means that the master will write information to a selected slave. A 'one' in this position means that the master will read information from the slave.




Figure 15. The first byte after the START procedure

When an address is sent, each device in a system compares the first seven bits after the START condition with its address. If they match, the device considers itself addressed by the master as a slave-receiver or slave-transmitter, depending on the R/W bit.

A slave address can be made-up of a fixed and a programmable part. Since it's likely that there will be several identical devices in a system, the programmable part of the slave address enables the maximum possible number of such devices to be connected to the I²C-bus. The number of programmable address bits of a device depends on the number of pins available. For example, if a device has 4 fixed and 3 programmable address bits, a total of 8 identical devices can be connected to the same bus.

The I²C-bus committee coordinates allocation of I²C addresses. Further information can be obtained from the Philips representatives listed on the back cover. Two groups of eight addresses (0000XXX and 1111XXX) are reserved for the purposes shown in Table 2. The bit combination 11110XX of the slave address is reserved for 10-bit addressing (see Section 13.0).

Table 2. Definition of bits in the first byte

SLAVE ADDRESS	R/ bit	DESCRIPTION
0000 000	0	General call address
0000 000	1	START byte
0000 001	X	CBUS address
0000 010	X	Address reserved for different bus format
0000 011	X	Reserved for future purposes
0000 1XX	X	
1111 1XX	X	
1111 0XX	X	10-bit slave addressing

NOTES:

1. No device is allowed to acknowledge at the reception of the START byte.
2. The CBUS address has been reserved to enable the inter-mixing of CBUS compatible and I²C-bus compatible devices in the same system. I²C-bus compatible devices are not allowed to respond on reception of this address.
3. The address reserved for a different bus format is included to enable I²C and other protocols to be mixed. Only I²C-bus compatible devices that can work with such formats and protocols are allowed to respond to this address.

9.1.1 General call address

The general call address is for addressing every device connected to the I²C-bus. However, if a device doesn't need any of the data supplied within the general call structure, it can ignore this address by not issuing an acknowledgement. If a device does require data from a general call address, it will acknowledge this address and behave as a slave-receiver. The second and following bytes will be acknowledged by every slave-receiver capable of handling this data. A slave which cannot process one of these bytes must ignore it by not acknowledging. The meaning of the general call address is always specified in the second byte (Figure 16).

There are two cases to consider:

- When the least significant bit B is a 'zero'
- When the least significant bit B is a 'one'.

When bit B is a 'zero'; the second byte has the following definition:

- 00000110 (H'06'). Reset and write programmable part of slave address by hardware. On receiving this 2-byte sequence, all devices designed to respond to the general call address will reset and take in the programmable part of their address. Precautions have to be taken to ensure that a device is not pulling down the SDA or SCL line after applying the supply voltage, since these low levels would block the bus
- 00000100 (H'04'). Write programmable part of slave address by hardware. All devices which define the programmable part of their address by hardware (and which respond to the general call address) will latch this programmable part at the reception of this two byte sequence. The device will not reset.
- 00000000 (H'00'). This code is not allowed to be used as the second byte.

Sequences of programming procedure are published in the appropriate device data sheets.

The remaining codes have not been fixed and devices must ignore them.

The I²C-bus and how to use it (including specifications)

When bit B is a 'one'; the 2-byte sequence is a 'hardware general call'. This means that the sequence is transmitted by a hardware master device, such as a keyboard scanner, which cannot be programmed to transmit a desired slave address. Since a hardware master doesn't know in advance to which device the message has to be transferred, it can only generate this hardware general call and its own address — identifying itself to the system (Figure 17).

The seven bits remaining in the second byte contain the address of the hardware master. This address is recognised by an intelligent device (e.g. a microcontroller) connected to the bus which will then

direct the information from the hardware master. If the hardware master can also act as a slave, the slave address is identical to the master address.

In some systems, an alternative could be that the hardware master transmitter is set in the slave-receiver mode after the system reset. In this way, a system configuring master can tell the hardware master-transmitter (which is now in slave-receiver mode) to which address data must be sent (Figure 18). After this programming procedure, the hardware master remains in the master-transmitter mode.




Figure 16. General call address format




Figure 17. Data transfer from a hardware master-transmitter




Figure 18. Data transfer by a hardware-transmitter capable of dumping data directly to slave devices
(a) Configuring master sends dump address to hardware master
(b) Hardware master dumps data to selected slave

The I²C-bus and how to use it (including specifications)

9.1.2 START byte

Microcontrollers can be connected to the I²C-bus in two ways. A microcontroller with an on-chip hardware I²C-bus interface can be programmed to be only interrupted by requests from the bus. When the device doesn't have such an interface, it must constantly monitor the bus via software. Obviously, the more times the microcontroller monitors, or polls the bus, the less time it can spend carrying out its intended function. There is therefore a speed difference between fast hardware devices and a relatively slow microcontroller which relies on software polling.

In this case, data transfer can be preceded by a start procedure which is much longer than normal (Figure 19). The start procedure consists of:

- A START condition (S)
- A START byte (00000001)
- An acknowledge clock pulse (ACK)
- A repeated START condition (Sr).

After the START condition S has been transmitted by a master which requires bus access, the START byte (00000001) is transmitted. Another microcontroller can therefore sample the SDA line at a low sampling rate until one of the seven zeros in the START byte is detected. After detection of this LOW level on the SDA line, the microcontroller can switch to a higher sampling rate to find the repeated START condition Sr which is then used for synchronization.

A hardware receiver will reset on receipt of the repeated START condition Sr and will therefore ignore the START byte.

An acknowledge-related clock pulse is generated after the START byte. This is present only to conform with the byte handling format used on the bus. No device is allowed to acknowledge the START byte.

9.1.3 CBUS compatibility

CBUS receivers can be connected to the I²C-bus. However, a third bus line called DLEN must then be connected and the acknowledge bit omitted. Normally, I²C transmissions are sequences of 8-bit bytes; CBUS compatible devices have different formats.

In a mixed bus structure, I²C-bus devices must not respond to the CBUS message. For this reason, a special CBUS address (0000001X) to which no I²C-bus compatible device will respond, has been reserved. After transmission of the CBUS address, the DLEN line can be made active and a CBUS-format transmission (Figure 20) sent. After the STOP condition, all devices are again ready to accept data.

Master-transmitters can send CBUS formats after sending the CBUS address. The transmission is ended by a STOP condition, recognised by all devices.

NOTE: If the CBUS configuration is known, and expansion with CBUS compatible devices isn't foreseen, the designer is allowed to adapt the hold time to the specific requirements of the device(s) used.




Figure 19. START byte procedure

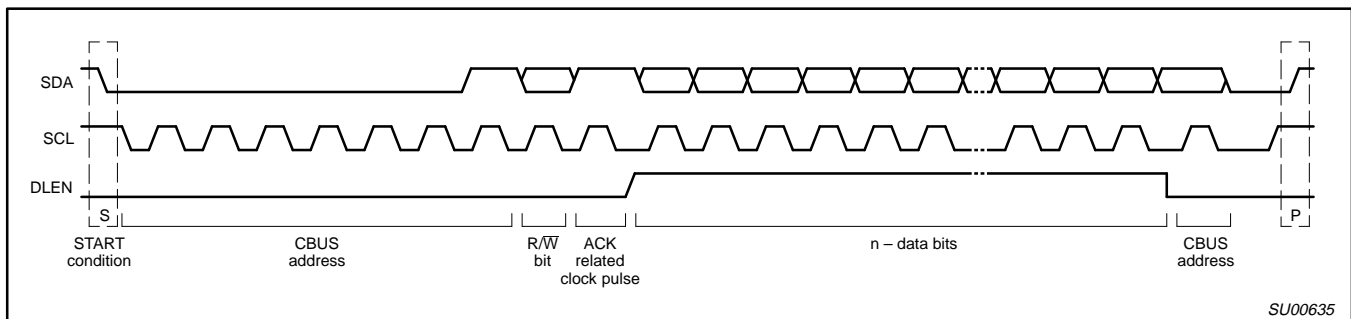


Figure 20. Data format of transmissions with CBUS transmitter/receiver

The I²C-bus and how to use it (including specifications)

10.0 ELECTRICAL CHARACTERISTICS FOR I²C-BUS DEVICES

The electrical specifications for the I/Os of I²C-bus devices and the characteristics of the bus lines connected to them are given in Tables 3 and 4 in Section 15.0.

I²C-bus devices with fixed input levels of 1.5 V and 3 V can each have their own appropriate supply voltage. Pull-up resistors must be connected to a 5V ± 10% supply (Figure 21). I²C-bus devices with input levels related to V_{DD} must have one common supply line to which the pull-up resistor is also connected (Figure 22).

When devices with fixed input levels are mixed with devices with input levels related to V_{DD}, the latter devices must be connected to one common supply line of 5 V ± 10% and must have pull-up resistors connected to their SDA and SCL pins as shown in Figure 23.

Input levels are defined in such a way that:

- The noise margin on the LOW level is 0.1 V_{DD}
- The noise margin on the HIGH level is 0.2 V_{DD}
- As shown in Figure 24, series resistors (R_S) of, e.g., 300Ω can be used for protection against high-voltage spikes on the SDA and SCL lines (due to flash-over of a TV picture tube, for example).




Figure 21. Fixed input level devices connected to the I²C-bus




Figure 22. Devices with wide supply voltage range connected to the I²C-bus

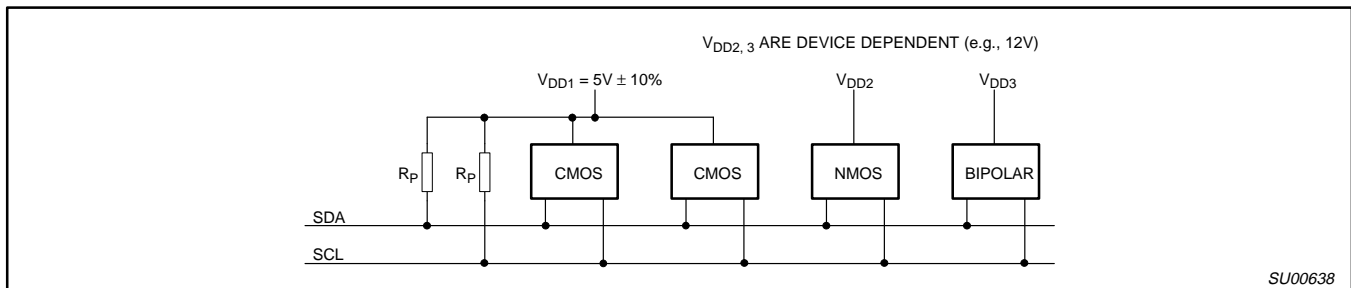


Figure 23. Devices with input levels related to V_{DD} (supply V_{DD1}) mixed with fixed input level devices (supply V_{DD2, 3}) on the I²C-bus

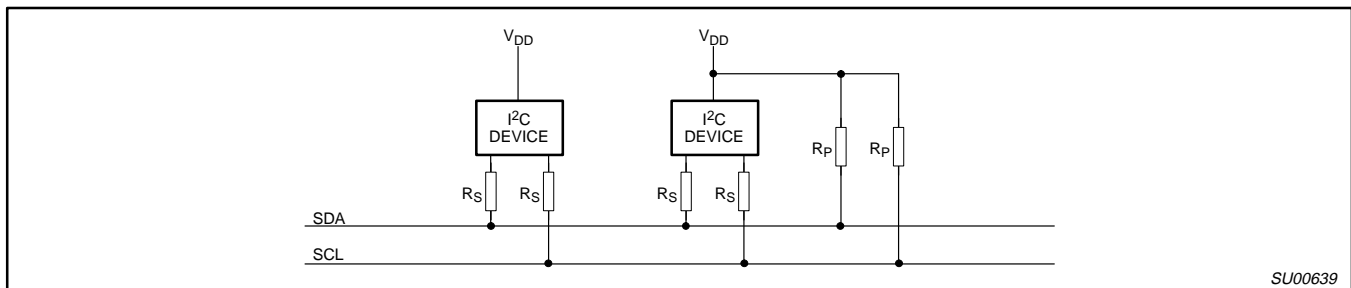


Figure 24. Series resistors (R_S) for protection against high-voltage spikes

The I²C-bus and how to use it (including specifications)

10.1 Maximum and minimum values of resistors R_p and R_s

For standard-mode I²C-bus devices, the values of resistors R_p and R_s in Figure 24 depend on the following parameters:

- Supply voltage
- Bus capacitance
- Number of connected devices (input current + leakage current).

The supply voltage limits the minimum value of resistor R_p due to the specified minimum sink current of 3 mA at V_{OLmax} = 0.4 V for the output stages. V_{DD} as a function of R_{p min} is shown in Figure 25. The desired noise margin of 0.1V_{DD} for the LOW level,

limits the maximum value of R_s. R_{s max} as a function of R_p is shown in Figure 26.

The bus capacitance is the total capacitance of wire, connections and pins. This capacitance limits the maximum value of R_p due to the specified rise time. Figure 27 shows R_{p max} as a function of bus capacitance.

The maximum HIGH level input current of each input/output connection has a specified maximum value of 10 μA. Due to the desired noise margin of 0.2V_{DD} for the HIGH level, this input current limits the maximum value of R_p. This limit depends on V_{DD}. The total HIGH level input current is shown as a function of R_{p max} in Figure 28.




Figure 25. Minimum value of R_p as a function of supply voltage with the value of R_s as a parameter




Figure 26. Maximum value of R_s as a function of the value of R_p with supply voltage as a parameter




Figure 27. Maximum value of R_p as a function of bus capacitance for a standard-mode I²C-bus




Figure 28. Total HIGH level input current as a function of the maximum value of R_p with supply voltage as a parameter

The I²C-bus and how to use it (including specifications)

11.0 EXTENSIONS TO THE I²C-BUS SPECIFICATION

The I²C-bus with a data transfer rate of up to 100 kbit/s and 7-bit addressing has now been in existence for more than ten years with an unchanged specification. The concept is accepted world-wide as a de facto standard and hundreds of different types of I²C-bus compatible ICs are available from Philips and other suppliers. The I²C-bus specification is now extended with the following two features:

- A **fast-mode** which allows a fourfold increase of the bit rate to 0 to 400 kbit/s
- **10-bit addressing** which allows the use of up to 1024 additional addresses.

There are two reasons for these extensions to the I²C-bus specification:

- New applications will need to transfer a larger amount of serial data and will therefore demand a higher bit rate than 100 kbit/s. Improved IC manufacturing technology now allows a fourfold speed increase without increasing the manufacturing cost of the interface circuitry
- Most of the 112 addresses available with the 7-bit addressing scheme have been issued more than once. To prevent problems with the allocation of slave addresses for new devices, it is desirable to have more address combinations. About a tenfold increase of the number of available addresses is obtained with the new 10-bit addressing.

All new devices with an I²C-bus interface are provided with the fast-mode. Preferably, they should be able to receive and/or transmit at 400 kbit/s. The minimum requirement is that they can synchronize with a 400 kbit/s transfer; they can then prolong the LOW period of the SCL signal to slow down the transfer. Fast-mode devices must be downward-compatible which means that they must still be able to communicate with 0 to 100 kbit/s devices in a 0 to 100 kbit/s I²C-bus system.

Obviously, devices with a 0 to 100 kbit/s I²C-bus interface cannot be incorporated in a fast-mode I²C-bus system because, since they cannot follow the higher transfer rate, unpredictable states of these devices would occur.

Slave devices with a fast-mode I²C-bus interface can have a 7-bit or a 10-bit slave address. However, a 7-bit address is preferred because it is the cheapest solution in hardware and it results in the shortest message length. Devices with 7-bit and 10-bit addresses can be mixed in the same I²C-bus system regardless of whether it is a 0 to 100 kbit/s standard-mode system or a 0 to 400 kbit/s fast-mode system. Both existing and future masters can generate either 7-bit or 10-bit addresses.

12.0 FAST-MODE

In the fast-mode of the I²C-bus, the protocol, format, logic levels and maximum capacitive load for the SDA and SCL lines quoted in the

previous I²C-bus specification are unchanged. Changes to the previous I²C-bus specification are:

- The maximum bit rate is increased to 400 kbit/s
- Timing of the serial data (SDA) and serial clock (SCL) signals has been adapted. There is no need for compatibility with other bus systems such as CBUS because they cannot operate at the increased bit rate
- The inputs of fast-mode devices must incorporate spike suppression and a Schmitt trigger at the SDA and SCL inputs
- The output buffers of fast-mode devices must incorporate slope control of the falling edges of the SDA and SCL signals
- If the power supply to a fast-mode device is switched off, the SDA and SCL I/O pins must be floating so that they don't obstruct the bus lines
- The external pull-up devices connected to the bus lines must be adapted to accommodate the shorter maximum permissible rise time for the fast-mode I²C-bus. For bus loads up to 200pF, the pull-up device for each bus line can be a resistor; for bus loads between 200pF and 400pF, the pull-up device can be a current source (3mA max.) or a switched resistor circuit as shown in Figure 37.

13.0 10-BIT ADDRESSING

The 10-bit addressing does not change the format in the I²C-bus specification. Using 10 bits for addressing exploits the reserved combination 1111XXX for the first seven bits of the first byte following a START (S) or repeated START (Sr) condition as explained in Section 9.1. The 10-bit addressing does not affect the existing 7-bit addressing. Devices with 7-bit and 10-bit addresses can be connected to the same I²C-bus, and both 7-bit and 10-bit addressing can be used in a standard-mode system (up to 100 kbit/s) or a fast-mode system (up to 400 kbit/s).

Although there are eight possible combinations of the reserved address bits 1111XXX, only the four combinations 11110XX are used for 10-bit addressing. The remaining four combinations 11111XX are reserved for future I²C-bus enhancements.

13.1 Definition of bits in the first two bytes

The 10-bit slave address is formed from the first two bytes following a START condition (S) or a repeated START condition (Sr).

The first seven bits of the first byte are the combination 11110XX of which the last two bits (XX) are the two most-significant bits (MSBs) of the 10-bit address; the eighth bit of the first byte is the R/W bit that determines the direction of the message. A 'zero' in the least significant position of the first byte means that the master will write information to a selected slave. A 'one' in this position means that the master will read information from the slave.

If the R/W bit is 'zero', then the second byte contains the remaining 8 bits (XXXXXXXX) of the 10-bit address. If the R/W bit is 'one', then the next byte contains data transmitted from a slave to a master.

The I²C-bus and how to use it (including specifications)

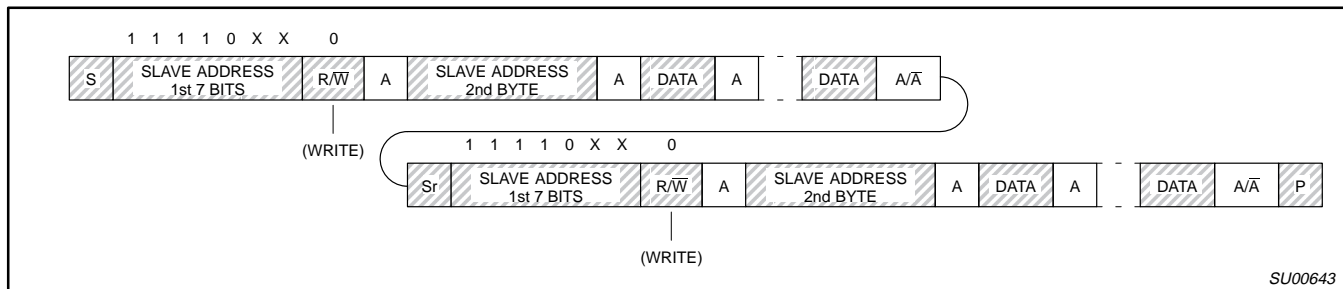


Figure 32. Combined format. A master transmits data to two slaves, both with 10-bit addresses

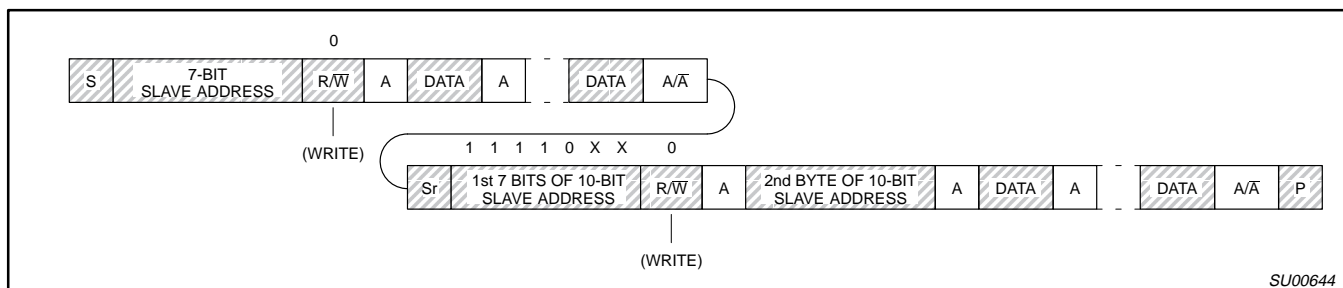


Figure 33. Combined format. A master transmits data to two slaves, one with a 7-bit address, and one with a 10-bit address

NOTES:

1. Combined formats can be used, for example, to control a serial memory. During the first data byte, the internal memory location has to be written. After the START condition and slave address is repeated, data can be transferred.
2. All decisions on auto-increment or decrement of previously accessed memory locations etc. are taken by the designer of the device.
3. Each byte is followed by an acknowledgement bit as indicated by the A or \bar{A} blocks in the sequence.
4. I²C-bus compatible devices must reset their bus logic on receipt of a START or repeated START condition such that they all anticipate the sending of a slave address.

14.0 GENERAL CALL ADDRESS AND START BYTE

The 10-bit addressing procedure for the I²C-bus is such that the first two bytes after the START condition (S) usually determine which slave will be selected by the master. The exception is the 'general call' address 00000000 (H'00'). Slave devices with 10-bit addressing will react to a 'general call' in the same way as slave devices with 7-bit addressing (see Section 9.1.1).

Hardware masters can transmit their 10-bit address after a 'general call'. In this case, the 'general call' address byte is followed by two successive bytes containing the 10-bit address of the master-transmitter. The format is as shown in Figure 17 where the first DATA byte contains the eight least-significant bits of the master address.

The START byte 00000001 (H'01') can precede the 10-bit addressing in the same way as for 7-bit addressing (see Section 9.1.2).

The I²C-bus and how to use it (including specifications)

15.0 ELECTRICAL SPECIFICATIONS AND TIMING FOR I/O STAGES AND BUS LINES

The I/O levels, I/O current, spike suppression, output slope control and pin capacitance for I²C-bus devices are given in Table 3. The I²C-bus timing is given in Table 4. Figure 34 shows the timing definitions for the I²C-bus.

The noise margin for HIGH and LOW levels on the bus lines for fast-mode devices are the same as those specified in Section 10.0 for standard-mode I²C-bus devices.

The minimum HIGH and LOW periods of the SCL clock specified in Table 4 determine the maximum bit transfer rates of 100 kbit/s for standard-mode devices and 400 kbit/s for fast mode devices. Standard-mode and fast-mode I²C-bus devices must be able to follow transfers at their own maximum bit rates, either by being able to transmit or receive at that speed or by applying the clock synchronization procedure described in Section 7.0 which will force the master into a wait state and stretch the LOW period of the SCL signal. Of course, in the latter case the bit transfer rate is reduced.

Table 3. Characteristics of the SDA and SCL I/O stages for I²C-bus devices

PARAMETER	SYMBOL	STANDARD-MODE DEVICES		FAST-MODE DEVICES		UNIT
		Min.	Max.	Min.	Max.	
LOW level input voltage: fixed input levels V _{DD} -related input levels	V _{IL}	-0.5 -0.5	1.5 0.3V _{DD}	-0.5 -0.5	1.5 0.3V _{DD}	V
HIGH level input voltage: fixed input levels V _{DD} -related input levels	V _{IH}	3.0 0.7V _{DD}	*1) *1)	3.0 0.7V _{DD}	*1) *1)	V
Hysteresis of Schmitt trigger inputs: fixed input levels V _{DD} -related input levels	V _{hys}	n/a n/a	n/a n/a	0.2 0.05V _{DD}	- -	V
Pulse width of spikes which must be suppressed by the input filter	t _{SP}	n/a	n/a	0	50	ns
LOW level output voltage (open drain or open collector): at 3 mA sink current at 6 mA sink current	V _{OL1} V _{OL2}	0 n/a	0.4 n/a	0 0	0.4 0.6	V
Output fall time from V _{IHmin} to V _{ILmax} with a bus capacitance from 10 pF to 400 pF: with up to 3 mA sink current at V _{OL1} with up to 6 mA sink current at V _{OL2}	t _{of}	- n/a	250 ³⁾ n/a	20 + 0.1C _b ²⁾ 20 + 0.1C _b ²⁾	250 250 ³⁾	ns
Input current each I/O pin with an input voltage between 0.4 V and 0.9V _{DDmax}	I _i	-10	10	-10 ⁴⁾	10 ⁴⁾	μA
Capacitance for each I/O pin	C _i	-	10	-	10	pF

n/a = not applicable

1. Maximum V_{IH} = V_{DDmax} + 0.5 V

2. C_b = capacitance of one bus line in pF.

3. The maximum t_f for the SDA and SCL bus lines quoted in Table 4 (300 ns) is longer than the specified maximum t_{of} for the output stages (250 ns). This allows series protection resistors (R_s) to be connected between the SDA/SCL pins and the SDA/SCL bus lines as shown in Figure 37 without exceeding the maximum specified t_f.

4. I/O pins of fast-mode devices must not obstruct the SDA and SCL lines if V_{DD} is switched off.

The I²C-bus and how to use it (including specifications)

Table 4. Characteristics of the SDA and SCL bus lines for I²C-bus devices

PARAMETER	SYMBOL	STANDARD-MODE I ² C-BUS		FAST-MODE I ² C-BUS		UNIT
		Min.	Max.	Min.	Max.	
SCL clock frequency	f _{SCL}	0	100	0	400	kHz
Bus free time between a STOP and START condition	t _{BUF}	4.7	–	1.3	–	μs
Hold time (repeated) START condition. After this period, the first clock pulse is generated	t _{HD;STA}	4.0	–	0.6	–	μs
LOW period of the SCL clock	t _{LOW}	4.7	–	1.3	–	μs
HIGH period of the SCL clock	t _{HIGH}	4.0	–	0.6	–	μs
Set-up time for a repeated START condition	t _{SU;STA}	4.7	–	0.6	–	μs
Data hold time: for CBUS compatible masters (see NOTE, Section 9.1.3) for I ² C-bus devices	t _{HD;DAT}	5.0 0 ¹⁾	–	– 0 ¹⁾	– 0.9 ²⁾	μs
Data set-up time	t _{SU;DAT}	250	–	100 ³⁾	–	ns
Rise time of both SDA and SCL signals	t _r	–	1000	20 + 0.1C _b ⁴⁾	300	ns
Fall time of both SDA and SCL signals	t _f	–	300	20 + 0.1C _b ⁴⁾	300	ns
Set-up time for STOP condition	t _{SU;STO}	4.0	–	0.6	–	μs
Capacitive load for each bus line	C _b	–	400	–	400	pF

All values referred to V_{IHmin} and V_{ILmax} levels (see Table 3).

1. A device must internally provide a hold time of at least 300 ns for the SDA signal (referred to the V_{IHmin} of the SCL signal) in order to bridge the undefined region of the falling edge of SCL.
2. The maximum t_{HD;DAT} has only to be met if the device does not stretch the LOW period (t_{LOW}) of the SCL signal.
3. A fast-mode I²C-bus device can be used in a standard-mode I²C-bus system, but the requirement t_{SU;DAT} ≥ 250 ns must then be met. This will automatically be the case if the device does not stretch the LOW period of the SCL signal. If such a device does stretch the LOW period of the SCL signal, it must output the next data bit to the SDA line t_{r max} + t_{SU;DAT} = 1000 + 250 = 1250 ns (according to the standard-mode I²C-bus specification) before the SCL line is released.
4. C_b = total capacitance of one bus line in pF.




Figure 34. Definition of timing on the I²C-bus

The I²C-bus and how to use it (including specifications)

16.0 APPLICATION INFORMATION

16.1 Slope-controlled output stages of fast-mode I²C-bus devices

The electrical specifications for the I/Os of I²C-bus devices and the characteristics of the bus lines connected to them are given in Tables 3 and 4 in Section 15.0.

Figures 35 and 36 show examples of output stages with slope control in CMOS and bipolar technology. The slope of the falling edge is defined by a Miller capacitor (C1) and a resistor (R1). The typical values for C1 and R1 are indicated on the diagrams. The wide tolerance for output fall time t_{of} given in Table 3 means that the design is not critical. The fall time is only slightly influenced by the external bus load (C_b) and external pull-up resistor (R_p). However, the rise time (t_r) specified in Table 4 is mainly determined by the bus load capacitance and the value of the pull-up resistor.

16.2 Switched pull-up circuit for fast-mode I²C-bus devices

The supply voltage (V_{DD}) and the maximum output LOW level determine the minimum value of pull-up resistor R_p (see Section 10.1). For example, with a supply voltage of V_{DD} = 5V ± 10% and V_{OLmax} = 0.4V at 3mA, R_{p min} = (5.5 - 0.4)/0.003 = 1.7 kΩ. As shown in Figure 38, this value of R_p limits the maximum bus capacitance to about 200pF to meet the maximum t_r requirement of 300 ns. If the bus has a higher capacitance than this, a switched pull-up circuit as shown in Figure 37 can be used.

The switched pull-up circuit in Figure 37 is for a supply voltage of V_{DD} = 5V ± 10% and a maximum capacitive load of 400pF. Since it is controlled by the bus levels, it needs no additional switching control signals. During the rising/falling edges, the bilateral switch in the HCT4066 switches pull-up resistor R_{p2} on/off at bus levels between 0.8 V and 2.0 V. Combined resistors R_{p1} and R_{p2} can pull-up the bus line within the maximum specified rise time (t_r) of 300 ns. The maximum sink current for the driving I²C-bus device will not exceed 6 mA at V_{OL2} = 0.6 V, or 3 mA at V_{OL1} = 0.4 V.

Series resistors R_s are optional. They protect the I/O stages of the I²C-bus devices from high-voltage spikes on the bus lines, and minimize crosstalk and undershoot of the bus line signals. The maximum value of R_s is determined by the maximum permitted voltage drop across this resistor when the bus line is switched to the LOW level in order to switch off R_{p2}.




Figure 35. Slope-controlled output stage in CMOS technology




Figure 36. Slope-controlled output stage in bipolar technology

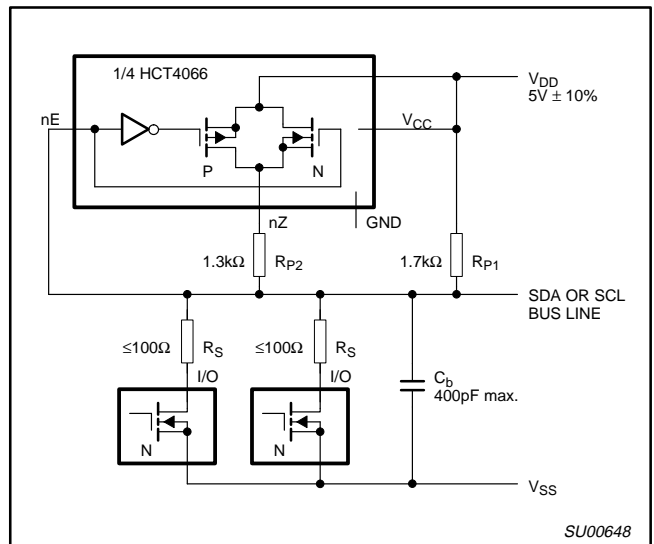


Figure 37. Switched pull-up circuit




Figure 38. Maximum value of R_p as a function of bus capacitance for meeting the $t_{R MAX}$ requirement for a fast-mode I²C-bus

The I²C-bus and how to use it (including specifications)

16.3 Wiring pattern of the bus lines

In general, the wiring must be so chosen that crosstalk and interference to/from the bus lines is minimized. The bus lines are most susceptible to crosstalk and interference at the HIGH level because of the relatively high impedance of the pull-up devices.

If the length of the bus lines on a PCB or ribbon cable exceeds 10cm and includes the V_{DD} and V_{SS} lines, the wiring pattern must be:

SDA _____
 V_{DD} _____
 V_{SS} _____
 SCL _____

If only the V_{SS} line is included, the wiring pattern must be:

SDA _____
 V_{SS} _____
 SCL _____

These wiring patterns also result in identical capacitive loads for the SDA and SCL lines. The V_{SS} and V_{DD} lines can be omitted if a PCB with a V_{SS} and/or V_{DD} layer is used.

If the bus lines are twisted-pairs, each bus line must be twisted with a V_{SS} return. Alternatively, the SCL line can be twisted with a V_{SS} return, and the SDA line twisted with a V_{DD} return. In the latter case, capacitors must be used to decouple the V_{DD} line to the V_{SS} line at both ends of the twisted pairs.

If the bus lines are shielded (shield connected to V_{SS}), interference will be minimized. However, the shielded cable must have low capacitive coupling between the SDA and SCL lines to minimize crosstalk.

16.4 Maximum and minimum values of resistors R_p and R_s for fast-mode I²C-bus devices

The maximum and minimum values for resistors R_p and R_s connected to a fast-mode I²C-bus can be determined from Figure 25, 26 and 28 in Section 10.1. Because a fast-mode I²C-bus has faster rise times (t_r) the maximum value of R_p as a function of bus capacitance is less than that shown in Figure 27. The replacement graph for Figure 27 showing the maximum value of R_p as a function of bus capacitance (C_b) for a fast mode I²C-bus is given in Figure 38.

17.0 DEVELOPMENT TOOLS

17.1 Development tools for 8048 and 8051-based systems

PRODUCT	DESCRIPTION
OM1016	I ² C-bus demonstration board with microcontroller, LCD, LED, Par. I/O, SRAM, EEPROM, Clock, DTMF generator, AD/DA conversion, infrared link.
OM1018	manual for OM1016
OM1020	LCD and driver demonstration board
OM4151	I ² C-bus evaluation board (similar to OM1016 above but without infrared link).
OM5027	I ² C-bus evaluation board for low-voltage, low-power ICs & software

17.2 Development tools for 68000-based systems

PRODUCT	DESCRIPTION
OM4160	Microcore-1 demonstration/evaluation board: SCC68070, 128K EPROM, 512K DRAM, I ² C, RS-232C, VSC SCC66470, resident monitor
OM4160/3	Microcore-3 demonstration/evaluation board: 128K EPROM, 64K SRAM, I ² C, RS-232C, 40 I/O (inc. 8051-compatible bus), resident monitor
OM4160/3QFP	Microcore-3 demonstration/evaluation board for 9XC101 (QFP80 package)

17.3 Development tools for all systems

PRODUCT	DESCRIPTION
OM1022	I ² C-bus analyzer. Hardware and software (runs on IBM or compatible PC) to experiment with and analyze the behaviour of the I ² C-bus (includes documentation)

The I²C-bus and how to use it (including specifications)

18.0 SUPPORT LITERATURE

DATA HANDBOOKS
Semiconductors for radio and audio systems IC01a 1995 IC01b 1995
Semiconductors for television and video systems IC02a 1995 IC02b 1995 IC02c 1995
Semiconductors for telecom systems IC03 1995
I ² C Peripherals IC12
8048-based 8-bit microcontrollers IC14 1994
Wireless communications IC17 1995
Semiconductors for in-car electronics IC18
80C51-based 8-bit microcontrollers IC20 1995
68000-based 16-bit microcontrollers IC21
Desktop video IC22 1995
Brochures/leaflets/lab. reports
I ² C-bus compatible ICs and support overview
I ² C-bus control programs for consumer applications
Microcontrollers, microprocessors and support overview
Application notes for 80C51-based 8-bit microcontrollers
OM5027 I ² C-bus evaluation board for low-voltage, low-power ICs & software
P90CL301 I ² C driver routines
User manual of Microsoft Pascal I ² C-bus driver (MICDRV4.OBJ)
User's guide to I ² C-bus control programs

The I²C-bus and how to use it (including specifications)

19.0 APPLICATION OF THE I²C-BUS IN THE ACCESS.bus SYSTEM

The ACCESS.bus (bus for connecting ACCESSory devices to a host system) is an I²C-bus based open-standard serial interconnect system jointly developed and defined by Philips and Digital Equipment Corporation. It is a lower-cost alternative to an RS-232C interface for connecting up to 14 inputs/outputs from peripheral equipment to a desk-top computer or workstation over a distance of up to eight metres. The peripheral equipment can be relatively low speed items such as keyboards, hand-held image scanners, cursor positioners, bar-code readers, digitizing tablets, card readers or modems.

All that's required to implement an ACCESS.bus is an 8051-family microcontroller with an I²C-bus interface, and a 4-wire cable carrying a serial data (SDA) line, a serial clock (SCL) line, a ground wire and a 12V supply line (500mA max.) for powering the peripherals.

Important features of the ACCESS.bus are that the bit rate is only about 20% less than the maximum bit rate of the I²C-bus, and the peripherals don't need separate device drivers. Also, the protocol allows the peripherals to be changed by 'hot-plugging' without re-booting.

As shown in Figure 39, the ACCESS.bus protocol comprises three levels: the I²C-bus protocol, the base protocol, and the application protocol.

The base protocol is common to all ACCESS.bus devices and defines the format of the ACCESS.bus message. Unlike the I²C-bus protocol, it restricts masters to sending and slaves to receiving data. One item of appended information is a checksum for reliability control. The base protocol also specifies seven types of control and status messages which are used in the system configuration which assigns unique addresses to the peripherals without the need for setting jumpers or switches on the devices.

The application protocol defines the message semantics that are specific to the three categories of peripheral device (keyboards, cursor locators, and text devices which generate character streams, e.g., card readers) which are at present envisaged.

Philips offers computer peripheral equipment manufacturers technical support, a wide range of I²C-bus devices and development kits for the ACCESS.bus. Hardware, software and marketing support is also offered by DEC.




Figure 39. ACCESS.bus protocol hierarchy