




48k dynamic RAM-card

universal for 1802


H.B. Stuurman



If you want to work with Basic, Forth, editors etc., you need a lot of memory. It makes no sense to build this up from 4K memory cards. A cheaper and better solution is to use dynamic RAM ICs.

The industry-standard 4116 is a $16K \times 1$ dynamic RAM chip. With eight of these 16-pin ICs you get 16 Kbyte, and with 24 ICs 48 Kbyte. The dynamic RAM card for the 1802 fits into a bus connector of the Cosmicos computer. The card is also easy to use with other 1802 systems; it is then connected with 24 wires.

What is remarkable is the simplicity of the dynamic RAM card. This is because a special IC has been used for timing and refresh: the Intel 8202. Since this controller and the memory ICs functionally form a self-contained unit, use with e.g. KIM or Junior computers is certainly possible by adapting the addressing logic. You don't need to design a PCB for that part – no small job saved.




What actually is a dynamic RAM?

A dynamic RAM is a memory IC that has to be addressed continuously. Fig. 1 shows the memory element on which the dynamic RAM is based. Indeed, it is a capacitor, and a very small one at that, both in size and capacitance. It is so small that you cannot see it without a microscope, and its capacitance is only 0.03 to 0.05 pF, depending on the manufacturer. It will come as no surprise that such a capacitor is

very cheap, much cheaper than, for example, the flip-flop memory element in a static RAM such as the 2114. A capacitor also consumes no energy. One may therefore expect dynamic RAMs to be cheaper than static ones, to use less power, and to have more memory elements per chip (higher "density"). Indeed, a memory built with dynamic RAMs is cheaper, uses 10× less power, and has 4× the chip density of an equally sized memory built with static RAMs.

So why aren't they always used? Well, there are drawbacks too! Drawbacks so large that for memories up to about 16K people prefer static RAMs.

To see why, you need to understand a little of how such a dynamic RAM works.



How does a dynamic RAM work?

In fig. 2, memory capacitor C is connected to the input of an amplifier via an electronic switch, also called a transmission gate (TG). The TG can be closed or opened by a switching signal on S. When the switch is closed and there is a voltage across C as a result of a charge, the amplifier output assumes a logic "1" state. If there is no voltage across C, the amplifier output is "0".

The amplifier thus senses the state of the capacitor and converts it into a usable signal. This amplifier is therefore called a "sense" amplifier.

When the switch is open, the top of the capacitor "floats". No capacitor is leak-free, and this one is no exception. Without further measures the charge would quickly leak away – and with it the information we want to keep. A provision is needed that keeps the charge of C at its proper level, or at least restores it again and again. Fig. 3 shows such a provision: a diode between the output and input of the amplifier. When the switch is closed and the capacitor is charged, the amplifier output becomes "1" and C is charged to full voltage via the diode. If the capacitor has no charge, nothing happens, of course.

The switching threshold of the amplifier input plays an important role in all this.

In practice the voltage difference between the "0" and "1" levels of the capacitor is

about 70 mV, so heavy demands are placed on the sense amplifier.

If the switch is briefly closed each time before the capacitor voltage has dropped below the switching threshold of the sense amplifier, the capacitor is topped up and retains its information. This applies to every capacitor in a dynamic RAM IC. The period within which this must happen is about 2 ms. Maintaining the capacitor charge in this way is called "refreshing" or refresh.

The need for refresh is the main reason for not using dynamic RAMs.

Structure of a dynamic RAM

In a dynamic RAM the memory capacitors are arranged in rows and columns. There are as many rows as columns: in a 4K IC (4096), for example, 64 of each, and in a 16K IC 128. This keeps the number of address lines as small as possible.

Fig. 4 shows the circuit of a 16×1 dynamic memory. Very small indeed, but adequate for our purpose. There are 16 memory capacitors with associated switches, divided over 4 rows and 4 columns.

The capacitors can be identified by stating the row and column they are in. Thus C2,3 sits at the intersection of row 2 and column 3.

The rows are addressed by the row address and the columns by a column address. For this, separate latch/decoders are provided for both the row and the column address. To put, say, the data of memory capacitor C2,1 on the output, you must:

1. Put row address 2 on the address input (10B).
2. Give a clock pulse on the RAS input.
3. Put column address 1 on the address input (01B).
4. Give a clock pulse on the CAS input.

RAS and CAS stand for Row Address Strobe and Column Address Strobe respectively. In good Dutch: rijadresklokpuls and kolomadresklokpuls.

Apart from the data of C2,1 appearing on the output, something else has happened. By clocking in row address 2, all switches in that row were closed; consequently all capacitors in that row were refreshed by the amplifiers in the column lines.

So to retain the information in our 16-bit dynamic memory we only need to address every row periodically. If we connect a binary counter to the row address input and give a RAS at each step it takes, we're done. This is called "RAS only refresh".


How high is the refresh frequency? Every capacitor must be refreshed at least once every 2 ms. However, all capacitors in the same row are refreshed simultaneously.

The counter must therefore go all the way round once every 2 ms. For our 16-bit memory the refresh rate is at least 2 kHz (for 2 ms and 4 rows).

For a 4K memory with 64 rows and columns this comes to 32 kHz, and for 16K with 128 rows and columns to 64 kHz.

Does entering the row and column addresses one after the other increase the access time? The access time is not increased by multiplexed addressing, because it takes

some time for the information in the rows to reach the outputs of the sense amplifiers. As long as CAS stays within this time interval, the access time is constant. The periodic recharging of the capacitors in a dynamic RAM causes considerable current surges on the supply lines. Heavy demands are therefore placed on the PCB design and the decoupling of the supply lines.



The 4116


The 4116 dynamic RAM IC is currently the industry standard. It is made by many manufacturers, and although the type numbers may differ, they are compatible.

Also notable is that supply from Japan in particular has been large lately. For example, Hitachi supplies the HM4716, Toshiba the TMM416, NEC the UPD416 and Fujitsu the MB8116. Besides the MK4116, Mostek makes a selected version for what it calls "personal computing and home entertainment memories", namely the MK4215. "Selected" here means: not meeting the MK4116 specifications. For us they are perfectly usable. Intel supplies the 4116 under type

number 2117.

Although there is a strong trend towards using a single 5 V supply (e.g. Intel 2118, Mostek MK4516), this is not attractive for us because of the price, which is (still) a multiple of that of the three-voltage version. All the more so since, thanks to converter circuits, the 48K dynamic RAM card needs only one unregulated supply voltage.

Fig. 5 shows the pinout of the 4116. Three supply voltages are needed: -5, +5 and +12 V.



The negative voltage serves to isolate the substrate and draws very little current. The +5 V supply also draws little current; it is only needed as a reference level for the input and output voltages. The "real" supply comes from the 12 V. Per IC, the current drawn is about 2 mA in standby and about 30 mA during addressing or refresh.


The current drawn from +5 and -5 V is measured in μA .

There are seven address pins, A0-A6. The row and column addresses are applied to

these multiplexed, with RAS and CAS as clock signals (RAS, CAS and WE are active low). Effectively there are thus 14 address bits ($2^{14} = 16384$ bits).

The timing diagram in fig. 6 shows the relationship between the applied address and RAS and CAS. Shortly after the CAS clock, the data appears on Dout. It stays there until CAS goes high again. To write, a negative pulse is put on Write while RAS and CAS are low. The data on Din is then stored at the selected memory location. Although not relevant to the dynamic RAM card, for completeness we mention the possibility of so-called early-write and page-access cycles. In the first case Write is made low before CAS. Dout then stays floating, which makes it possible to connect Din and Dout together (common I/O).

With page access, by keeping RAS low and entering only column addresses with CAS, different columns can be addressed quickly.




Summary

Although a dynamic RAM such as the 4116 is designed to be user-friendly, the multiple timing signals and the multiplexing of row and column addresses place a considerable burden on the surrounding hardware for connection to the processor.

Besides that, the RAS addresses must be cycled at a frequency of more than 64 kHz to take care of refresh. Read or write operations must not be affected by this. The memory must also be refreshed while the processor is in the reset state. It must always be refreshed while it is powered. It is certainly possible to build the necessary circuitry with, e.g., TTL ICs. However, guaranteeing a completely safe situation under all conditions is not easy.

Intel's 8202 dynamic RAM controller

Intel's 8202 is an IC that takes the "dynamics" out of dynamic RAMs. It makes the dynamic memory look to the processor as if it were static. The 8202 independently handles refresh, address multiplexing, and CAS and RAS timing. It is a bipolar IC, meaning the circuit uses (bipolar) transistors. This stems from the fact that a controller must be faster than the circuits it drives. The price to be paid is more power – literally and figuratively. The 8202 is not cheap, but considering that only one controller is needed for 48K of RAM and that it relieves us of many worries, it can



be considered a justified expense. Incidentally, if you don't mind making the PCB, the 8202 can drive 64K. For practical reasons we did not do this, since the EPROMs have to go somewhere too.

The 8202 draws about 250 mA from its single 5 V supply; not little, but fortunately there is only one. All that effort makes it rather warm!

Operation of the 8202

Fig. 7 shows the pinout of the 8202 and fig. 8 the block diagram. Using this block diagram we will briefly discuss its operation.

With dynamic RAMs, a precise relationship between the various timing signals is of great importance.

In the 8202 these are obtained from taps on a shift register. This shift register is driven by a crystal oscillator. The recommended frequency may be between 18 and 22 MHz. A commonly used value is 18.432 MHz, because cheap crystals exist for it. There are 14 address inputs, divided into two groups of seven: a high group and a low group. The two groups are put on the address outputs one after the other. The fact that these outputs are inverted has no further consequences for operation. The high group forms the column address and is clocked into the dynamic RAMs by CAS. The low group forms the row address. This is clocked into a bank of eight dynamic RAMs by one of the four RAS outputs. Which RAS this is, is determined by the bank-select inputs B0 and B1. These inputs are decoded internally and combined with the RAS clock. Usually B0 is connected to A14 and B1 to A15. In this way RAS0 corresponds to address bank 0000 to 3FFF; RAS1 to 4000–7FFF; RAS2 to 8000–BFFF and RAS3 to C000–FFFF.


In the 48K dynamic RAM card, RAS3 is not connected. The dynamic memory can be enabled and disabled via the chip-select input. By decoding addresses and using the result as the chip-select signal, memory areas can be kept free. On the 48K RAM card this has been done for the memory area C000 to FFFF.

The 8202 further contains a refresh timer, a refresh counter and an arbiter. The refresh counter is simply a 7-bit binary counter whose value is put on the address outputs during a refresh operation. The four RAS outputs are then clocked simultaneously, after which the refresh counter is incremented by one.

The refresh timer is a frequency divider that brings the oscillator frequency down to the value required for periodic refresh. Every 2 ms, 128 RAS addresses must be refreshed: the timer period is about 15 μ s. The timer has a reset facility.

As long as no read or write requests arrive, the refresh timer and counter perform the necessary refresh on their own. It makes no difference whether the processor is on or off.

The data in the dynamic RAMs is safe as long as the controller receives supply voltage!



Arbiter

A key function in the 8202 is performed by the arbiter. An arbiter is a referee, and in the 8202 the term denotes the circuit that decides which operation takes priority. There are three possible operations: a request from the processor for a read, done by putting a logic "0" on the RD input; a request for a write, done by putting a logic "0" on the WR input; or a request for a refresh. The latter can come either from the internal timer or from the Refresh Request input (REFRQ) going "0".

If only one request comes in, it is granted immediately. However, if two requests come in at the same time – these will be a read or write request and a refresh request – the arbiter decides as follows. The read or write request is handled first, immediately

followed by the refresh. After a refresh, regardless of whether the request came from the timer or the refresh input, the refresh timer is reset. If a read or write request arrives while a refresh is in progress, the refresh is completed first and only then is the requested operation carried out (fig. 9). The processor must therefore be informed that the requested operation cannot take place yet. For this the 8202 has a system-acknowledge output (SACK). Normally SACK goes "0" immediately after a read or write request. However, if a refresh is in progress when a read or write request arrives, SACK stays "1" until the request is granted. This is called delayed-SACK mode. SACK can be used to make the processor wait by connecting it to the Wait input.

Fig. 10 shows the timing diagram of the 8202. It would of course be nicest if the processor never had to wait, in other words if the refresh happened at moments when no read or write is requested. This is called "hidden" refresh. Well, that is what the refresh-request input is for. It is often connected so that a refresh is also requested during an "instruction-fetch" cycle. The refresh then takes place while the instruction is being decoded (the processor is not using the memory then).

XACK (transfer acknowledge) going low indicates that the information to or from the memory is valid. XACK is usually used during a read to clock the data coming from memory into a holding latch.

There it sits quietly while the 8202 gives the dynamic RAMs a little refresher.

Part 1 captions:

Fig. 1 Memory element of a dynamic RAM as seen through the eyes of an MK artist.

Fig. 2 The amplifier converts the weak capacitor voltage into a usable signal.

Fig. 3 By connecting the amplifier output to the capacitor via a diode and briefly closing the switch periodically, the charge is restored each time – refreshing.

Fig. 4 Possible structure of a 16×1 bit dynamic RAM. Note the division into rows and columns. The write option, omitted here for clarity, can easily be implemented with transmission gates.

Fig. 5 Pinout of the 4116 dynamic RAM, the industry standard.

Fig. 6 The row and column addresses are applied multiplexed. The RAS clock signal serves for the rows and the CAS clock signal for the columns.

Fig. 7 Pinout of Intel's dynamic RAM controller, the 8202.

Fig. 8 Internal block diagram of the 8202.


Fig. 9 Flow chart of the "arbiter" in the 8202.

Fig. 10 Timing diagram of the 8202. "Delayed-SACK mode" occurs when a read or write is requested while a refresh is in progress.



RAM-kaart

Afb. 11 Schakeling van de 48K dynamische RAM-kaart. Zie voor RX de tekst.




The circuit of the 48K dynamic RAM card is shown in fig. 11 and its power supply in fig. 12. At its heart is the 8202 dynamic RAM controller, which generates the timing and address signals for the memory block of 24 ICs. The seven address lines OUT 0–6, Column Address Strobe and Write Enable go in common to all RAM ICs.

The Row Address Strobe selects one bank of eight ICs. Three RAS outputs are used for the three banks of 16K each. RAS-3 is not connected here.

Data bytes to be written into memory are fed to the RAM data inputs via an octal buffer (IC2). Data bytes coming from memory are put into a holding latch (IC3). The controller doesn't have to wait until the CPU takes this data, but can start a refresh cycle immediately. The holding latch is an octal latch with tri-state outputs. The data is latched by means of the XACK output. XACK indicates that the transfer is valid (transfer acknowledge). For speed, low-power Schottky TTL ICs are used for the data bus buffer and the holding latch. For proper matching to the CMOS data inputs of the 1802 microprocessor, eight 10 k Ω pull-up resistors are included. The load the 1802 "sees" on its data bus is one low-power TTL input in parallel with a 10 k Ω resistor to +5 V.

This is no problem for the 1802, which can drive one standard TTL input per output.

The circuitry of the controller, the dynamic RAM block, the data buffer and the holding latch is standard and can be used as-is with many other processors, including the KIM and the Junior. The difference lies in how the 8202 is driven by the microprocessor. The 1802, as will become clear, is not the easiest in this respect.



Driving the 8202 controller

Let's get straight to the point: the problems are caused by the multiplexed address bus of the 1802 microprocessor. Looking at the 1802 timing diagram (fig. 13), we see that early in the machine cycle the high address byte is put on the address bus. It is valid at the trailing edge of TPA and can be latched with it. Immediately afterwards the low byte is put on the address bus, and it stays on the bus for the rest of the machine cycle. In a

write operation the write pulse is given when the full address is valid, around TPB. In a read operation the data is copied into the CPU at the start of TPB. A read is initiated by Read going low. That happens very early in the machine cycle. The 1802 datasheet therefore speaks not of a read "pulse" but of a read "level". The controller, however, sees Read going low as a command to read and responds immediately. Read must therefore be delayed until the full address is present. On the other hand, the delay must not be too long, because the controller needs time to output the column and row addresses and the RAMs have a certain access time.

The CPU copies the data byte at the start of TPB. The read request must start well after TPA – so that the full address is available – and well before TPB, to give the controller and the RAMs enough time; see fig. 14.

The 8202 has 14 address inputs, A0–A6 low and A0–A6 high. The high address byte is latched by two latches, IC7 and IC8, and address bits 0–13 are applied directly to the inputs mentioned. The remaining address bits, A14 and A15, are used for bank select (see the table).

Bank 3 is not present. IC6C decodes this condition ($A_{14} = 1$ and $A_{15} = 1$) and uses it to make the 8202 "inactive" via the chip-select input. The chip-select signal is also used, in combination with Read, to enable the tri-state outputs of the holding latch. IC4B and IC4A serve this purpose.

The write pulse from the microprocessor can be connected directly to the controller.

Delaying Read

In a microprocessor all timing signals are derived from the clock. The clock can also be used to delay Read. This has the advantage that at other clock frequencies the relationship between the delayed Read and the other signals doesn't change – in contrast to a delay made with, e.g., a monostable multivibrator.

IC5 is a 4-bit shift register with which the Read signal can easily be delayed by one, two or three clock cycles. This IC has an input that switches the four outputs to their complements. On a read request the reset condition is released, and on the next positive clock edge a "1" is shifted into flip-flop 0. Its output becomes "0". Subsequent clock pulses make outputs 1, 2 and 3 go "0" respectively. In practice a delay of about three clock pulses (output 2) proves satisfactory. At clock frequencies above 2 MHz, output 1 is preferable because of the propagation delay of the shift register. With a 2 MHz clock the prototype of the dynamic card worked reliably, regardless of the chosen Read delay.

Hidden refresh?

The 8202 datasheet states that a so-called hidden refresh can be obtained by having a refresh immediately follow an instruction fetch. In practice read and refresh requests may arrive simultaneously; the arbiter gives the read request priority. By combining the delayed Read with state code line 0 (SC0), an instruction fetch is immediately followed by a refresh. IC4C serves this purpose.

Up to about 3 MHz the cycle time of the 1802 is long enough to "hide" the refresh in it. Conversely, from a processor clock frequency of about 1.6 MHz upwards, even with 3-cycle instructions (LBR etc.) the instruction-fetch cycles follow each other quickly enough to reset the refresh timer in the 8202. For the most common clock frequencies, 1.75 and 2 MHz, the timing works out.

Now suppose we use the facility of addressing memory directly (DMA). With DMA-out no problems are expected: it is a read operation, and moreover SC0 is low during a DMA operation. Every read is automatically followed by a refresh. With DMA-in, however, problems are to be expected, because in the absence of reads no periodic refresh is requested. The 8202's refresh timer now takes care of refresh, and occasionally memory locations will be skipped during writing.

There are more bottlenecks. If the processor is switched from reset to run and a refresh starts just before the first instruction is fetched, is there enough time left to fetch the instruction? And what if the pause facility of the 1802 is used? We mustn't leave this to chance, and the 8202 offers a solution: the System Acknowledge output (SACK). If no refresh is in progress, SACK responds immediately to a read or write request by going "0". If a refresh is in progress when a read or write request arrives, SACK only goes "0" once the request has been granted. SACK can therefore be used to make the processor wait in those cases where a "collision" threatens.

To make the processor wait, three conditions must be met:

1. The 8202 controller must be selected.
2. A read or write request must be present at the 8202.
3. SACK remains "1" (because of a refresh).

IC6A and IC6B serve this purpose. By including a diode in the connection to WAIT, single-stepping remains possible. Using the SACK output has one consequence: even when no refresh is in progress, SACK does not go "0" instantaneously; it takes about 100 ns. If this happens just at a negative clock edge of the processor, one clock period of waiting results. With the 1861 or 1864 video chip this can cause a jittery picture. In such a case you must accept the possible bottlenecks mentioned above and not use SACK (leave out D1).

Table – How the address space is divided into four banks by means of address lines A14 and A15.

A15	A14	Bank	Address range
0	0	0	0000 to 3FFF
0	1	1	4000 to 7FFF
1	0	2	8000 to BFFF
1	1	3	C000 to FFFF

Power supply

A 4116-type dynamic RAM needs three supply voltages: -5, +5 and +12 V. The -5 and +5

V carry only a light load. The +12 V is the actual working voltage. The negative supply serves to isolate the substrate. If this voltage is not sufficiently negative while the other voltages are present, these are partially short-circuited and the 4116 can be destroyed. Recommended power-up procedure: first ensure the negative voltage is present, then switch on the other voltages.

For power-down the reverse order applies.

Because the current drawn from the -5 V supply is so small, it is worthwhile to include a converter for it. In fig. 12 this is a multivibrator built with a 555 (IC10), followed by an inverting voltage doubler (C5, C6, D3 and D4).

Regulation at -5 V is obtained by connecting the negative voltage to the reset input of the 555 via a 5.6 V zener diode.

The coil of a small relay is also connected to the -5 V; its switch contact is in the supply path for the +5 and +12 V. As soon as the negative voltage reaches about -3 V, the relay closes. The negative voltage is thus always present first. At power-down, the large capacitance of C6 makes the negative voltage fall slowly to 0 V, while because of the relatively low total capacitance the positive voltages drop quickly – faster than the negative voltage. The relay thus simply ensures the recommended power-up and power-down sequence.

The -5 V converter is connected to the output of the 7812 via a diode (D2). This limits the current to be switched by the relay to at most the short-circuit current of the 7812. For the relay used, that is more than the value specified for achieving its rated life. A life of 5×10^6 operations will therefore probably not be reached. Tests, however, showed a relay still working perfectly after 10,000 operations at this load, after which testing was simply stopped.

A LED is included to indicate the presence of the three supply voltages. At power-up it lights with a slight delay.

For decoupling the dynamic RAMs the manufacturer's guidelines were followed, resulting in a considerable number of 100 nF ceramic capacitors.

References

- Intel Component Data Catalog.
- Intel application note for the 8202.
- Take the "dynamics" out of dynamic RAM with an IC controller handling refresh. Electronic Design No. 26, December 1978.
- Memory System Design Seminar. Published by Intel Memory Components Division.
- Mostek memory databook and designer's guide.
- Toshiba MOS memory products Databook.


Part 2 captions:

Fig. 11 Circuit of the 48K dynamic RAM card. See the text for RX.

Fig. 12 Power supply for the dynamic RAM card. This circuit and the circuit of fig. 11 are on PCB 7561.

Fig. 13 On the 1802 the address is multiplexed onto the 8-bit address bus: first the high byte, which can be latched with TPA, then the low byte.

Fig. 14 The read pulse to the 8202 controller must be delayed until the full address is available. The delay must not be too long, so as to give the controller and the RAMs enough time.




Construction

A PCB has been designed for the dynamic RAM card. Because of the strict decoupling requirements, this board is rather complicated. It is a double-sided board with plated-through holes. Thanks to the through-plating, construction is simple, although there are many solder joints. Good soldering is essential! It is wise to use sockets for all ICs and the relay. These are fitted first. Some makes have a marking for pin 1 of the IC. Pay attention to this; it makes placing the ICs easier later. After the IC sockets, the remaining components can be fitted.

The crystal case is connected to ground; a pad for this is provided. Also don't forget the wire link for the delayed Read.



The parts list specifies relay type GR114AD5 V. Other types, with or without a diode, are also usable. If they have a built-in diode, some care is needed regarding polarity, as it can differ. The Electrol RA30461051, for example, is fitted the other way round. The pinouts of several relays are drawn in fig. 16.



The connector strip is not placed in the middle of the board. Thanks to the asymmetric placement, a card guide can be used on one side. With this long board that is no unnecessary luxury! If the dynamic RAM card is plugged upright into a connector, the recommended heatsinks are sufficient, provided the supply voltage is no higher than 15 V. It is also possible to mount the card underneath the computer. The computer is then placed on a wooden frame

about 25 mm high, in which a connector for the RAM card is mounted. Figs. 19 and 20 show one possibility. A good number of holes have been drilled in the frame to allow at least some air circulation. This, however, is too little for good cooling. The two regulators are therefore not mounted on the board but on a separate heatsink block at the back of the frame. The controller itself also gets warm; it too is fitted with a small heatsink. If the board is placed vertically in a connector, these measures are not needed. The total current consumption of a fully populated dynamic RAM card is 600 to 700 mA. The current consumption of the complete computer is about 1 A. Because a regulated 12 V is now used, the unregulated voltage should be about 14 to 15 V. The transformer used must be able to supply an AC voltage of 14 to 15 V at a current of 1.5 to 2 A. A 12 V transformer is therefore not usable unless some extra windings are added.







The mysterious RX

The prototype of the dynamic RAM card is fitted with the 8202. It has been working to full satisfaction for quite some time.

Nevertheless, various users seem to have had problems with the 8202, specifically with its clock oscillator circuit. Its frequency may be between 18 and 25 MHz. At the higher frequency the crystal with its series capacitor turned out to be sensitive to interference coupled into the clock input. Instead of at its fundamental frequency, the crystal started oscillating on its third overtone. Intel temporarily got round this problem by issuing new datasheets that mention only the use of an external clock; the 8202 was then renamed 8202A. Meanwhile we have heard from a reliable source that in cases where the crystal oscillator does not work reliably, pin X0 should be connected to +5 V through

a 1 kΩ resistor. This resistor is then fitted on the top side of the board. Intel itself has also found a remedy that is easy to apply where needed: connect X1 (pin 37) to GND through a 1 kΩ resistor, and X0 (pin 36) to GND as well through a 680 Ω resistor. On our board the crystal is connected to the 8202 with extremely short connections, and ten to one the crystal oscillator will work without this provision.



Parts list

Resistors

R1 to R8	10 k Ω	e.g. network 8 $\times$ 10 k Ω , VR8-10 k, ITT Multicomponents, Zoetermeer, order no. 290 10 450, or DIL network 8 $\times$ 10 k Ω , Texim, order no. 916C103X2SR, or make your own from 8 resistors 10 k Ω , 1/8 W
R9, R11	2.2 k Ω , 1/8 W	
R10	15 k Ω , 1/8 W	
R12	330 Ω , 1/4 W	

Capacitors

C1	6.8 pF, ceramic (2.5 mm pitch)
C2, C9 to C12	1 μ F, tantalum (2.5 mm pitch)
C3	47 μ F, 16 V, tantalum (2.5 or 5 mm pitch), or radial electrolytic
C4	2.2 nF, MKM (7.5 mm pitch)
C5	1 μ F, MKM (7.5 or 10 mm pitch)
C6	220 μ F, 10 V ($\varnothing$ 9 $\times$ 18 mm)
C7	22 μ F, 25 V ($\varnothing$ 6 $\times$ 11 mm)
C8	4.7 μ F, 63 V ($\varnothing$ 6 $\times$ 12 mm)
C13 to C18	100 nF, 60 V (5 mm pitch), Siemens Sibatit
16 $\times$ CA	100 nF, 60 V (5 mm pitch), Siemens Sibatit
16 $\times$ CB	100 nF, 60 V (5 mm pitch), Siemens Sibatit

ICs

IC1	8202 or 8202A-3, dynamic RAM controller, Intel (Inelco, Aalsmeer)
IC2	74LS244 (octal buffer)
IC3	74LS373 (octal latch)
IC4	4001
IC5	4035 (4-stage shift register)
IC6	4023
IC7, IC8	4042 (quad latch)
IC9	7812, TO220
IC10	555 (DIL 8)
IC11	7805, TO220
Per 16K bank: 8 $\times$ 4116 (max. 24 = 48K)	

Diodes

D1 to D4	1N914 (1N4148)
Dz	zener 5.6 V, 250 mW

LED red LED (2.5 mm pitch)

Miscellaneous

Crystal 18.432 MHz, HC18U

Relay DIL reed relay, e.g. type GR114AD5-5V (Van Reijssen Elektronica, Delft)
(1 × make, 5 V, approx. 500 Ω)

3 IC sockets, 14 pins, low profile

27 IC sockets, 16 pins, low profile

2 IC sockets, 20 pins, low profile

1 IC socket, 40 pins, low profile

2 heatsinks for TO220, type KL-105/SW, Amroh, order no. 6.039 (see also text)

4 M3 × 5 nuts and bolts

thermal paste

1 double-sided plated-through PCB, MK order no. 7561

Part 3 captions:

Fig. 15 Component layout of the dynamic RAM card. Pay particular attention to the orientation of IC5! The numbering of A-low and A-high is symbolic; refer to the pin numbers when testing etc. (This is a double-sided board; for clarity only the copper side, i.e. the underside, is shown.)

Fig. 16 Pinouts of some components.

Fig. 17 Ready-to-use 48K dynamic RAM card. Super Basic or Forth is no longer a problem!

Fig. 18 Close-up of the controller with crystal. The new type number is D8202A-3.

Figs. 19 and 20 If the computer is placed on a wooden frame, the dynamic RAM card can go underneath. The regulators, IC9 and IC11, are then mounted on a separate heatsink at the back of the frame. It is also advisable to fit the 8202 with a heatsink fin because of the poor air circulation.