Mysteries of
Wang Magnetic Memory

Douglas Miller <durgadas311@gmail.com>

August 12, 2026

Table of Contents

Wire-weave Microcode ROM

This journey began when I noticed something odd about the timing of signals used to access the Wang 500/600 microcode ROM. The signal timing did not match my understanding of how the ROM worked.

As background, the basic “wire-weave” ROM technology is as follows. Each word of micrcode is represented by a single wire that is woven through transformers (one for each bit) such that a “1” is represented by the wiring passing inside the given transformer. When a wire is activated (energized), it produces a pulse in the transformer, using the basic principals of electro-magnetism (Faraday’s Law of Induction). When the wire is first energized, it produces a negative pulse, and when the wire is de-energized it produces a positive pulse. The negative pulse is destructive to the digital circuitry and so must be shunted through a diode to protect that. The Wang 700 and 500/600 ROMs use essentially the same circuitry, with the main difference being the disable of transformers on the 500/600. The 700 transformers drive the base of an NPN transistor, however the SP380 NOR gates (driven by the transformers) on the 500/600 are RTL and the inputs are the base of an NPN transistor, so not really different. However, the two families use the ROM circuitry (signals and timing) slightly differently.

The Wang 700 ROM signals match my expectations (see simplified Wang 700 schematic). The leading edge of the energizing signal is “softened” by slowing the rise time, such that a minimal negative pulse is generated through the shunt diode. The falling edge of the signal produces the positive pulse used to set a “1” in the respective bit positions. The timing of the signals is as follows:

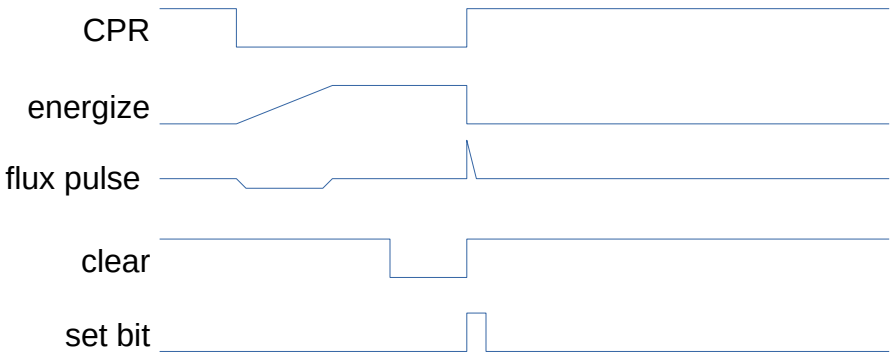

Note how the “clear” signal, which is required to clear the data latches before setting the “1” bits, precedes the end of the CPR (energizing) signal, such that it does not conflict with the pulses that set “1”.

The Wang 500/600 signals are different (see simplified Wang 600 schematic). The CUP signal, used to energize the wire, and the CL (clear) signals begin at the same time. The transformer outputs (and SET latch inputs) are normally held to ground through a diode by the inactive state of the ST signal, presumably to prevent noise from corrupting the results during the rest of the instruction cycle. But, the ST signal only enables the latches before the end of the energizing signal, and so the positive pulse resulting from the trailing edge of CUP cannot set the latches. Here are the timing signals:

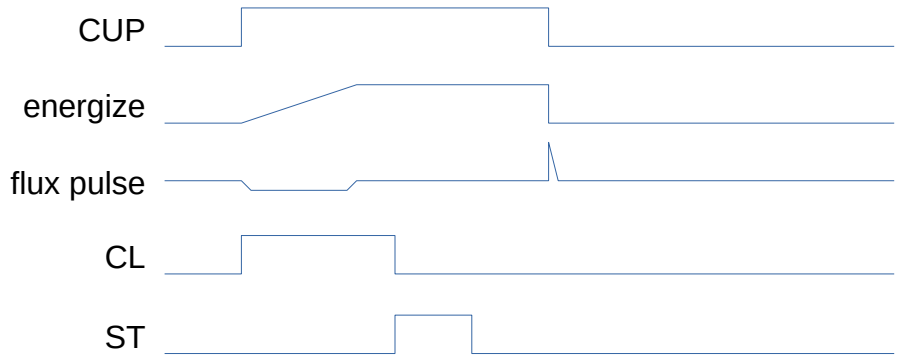

So, the question is how can the “normal” pulse generated at the end of the CUP signal be used to set latches if those are only enabled during the ST signal which precedes that event. Thanks to some scope traces from Andrew at https://dopecc.net/, I’ve confirmed that my interpretation of the signals timing and usage is correct.

I am faced with the possibility that Wang was using some property of electro-magnetism that I don’t understand. This led me to take a detour into magnetic core memory, which Wang used on the 700 models. In conventional core memory (see simplified Wang 700 schematic), the reading of a bit involves setting the bit to “0” and then detecting (sensing) whether the core changed state in order to do that (i.e. 1→0 involves a flux change, while 0→0 does not). This is done by writing a “0” to the bit “in the normal way” while using the “sense” wire to detect any induced current from a 1→0 transition. The timing of a read from Wang 700 core memory is as follows:

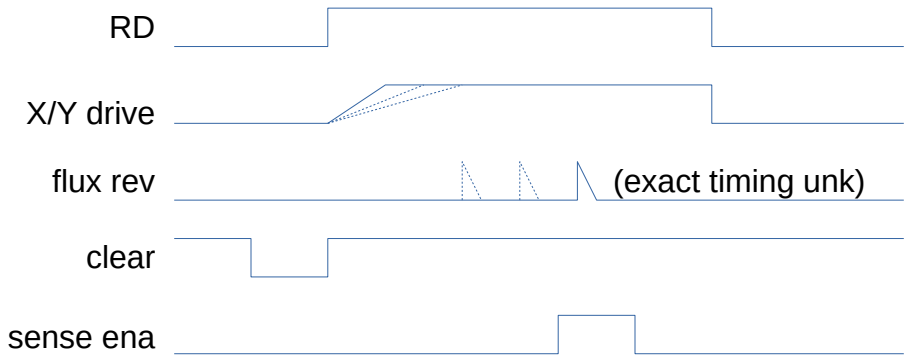

Note how the output latches are cleared early, then the X/Y wires are energized (RD) in order to write a “0”. The rise time of energizing current (X/Y wires) is intentionally slowed, as was done for ROM. However, this rise time is adjustable and appears to be matched to the core access time such that flux

reversal pulses coincide with the sense enable signal. In addition, this rise time seems to be temperature compensated, as the diodes used to establish the target voltage are physically placed (remote to the PCB) at the top of the core-planes (adjacent) PCB. Notice that the sense amplifier outputs are enabled before the end of energize signals, similar to what is done on the Wang 500/600 ROM. In this case, that is because the core "flips" state some time after the beginning of the energize signal, but is not tied to the end of the energize signal (instead, it is based on the access time of the memory, which depends on design details and materials). It is not clear to what degree the sense amplifiers might hold a pulse long enough for the sense enable to allow it to pass along to the latches, making the timing less critical. There is no ability to tune or vary the timing of the sense enable signal relative to the RD signal (only the rise time of the energize signal).

One more difference between the wire-weave ROM and magnetic core memory is that the ROM transformers have a true coil as the secondary/sense, which magnifies the effects of the flux changes. Magnetic core memory uses simple sense wires. Also, the cores retain their flux direction after removal of the drive current, so there is no collapsing magnetic field to induce current when the energize signal turns off. The only time the core induces a current is when it changes state (flips direction).

The Wang wire-weave ROMs, supposedly, do not use the same type of materials in the transformer cores as are used for magnetic core memory (so may not have any "memory" side-effect). Although, this is not confirmed. The drive circuitry for the ROM does not use negative current (unlike core memory, which uses +/- current for 0/1), so the transformers (supposedly) could not be detecting a flux-reversal as is done in core memory. However, on the 500/600, the transformer secondary (sense) coils are kept grounded except for during the ST pulse. The magnetic field from the energized wire is present at the time that the coils are ungrounded. So perhaps there is some special effect, such that the existing and steady magnetic field on the energized wire appears to the recently-ungrounded coil as a changing magnetic field, thus inducing a pulse. It is not clear if this has any relationship to "The Hall Effect".

Analyzing the Wang 600 ROM scope traces I was given, I see that the "sense" signal (input to the "set 1" latches) goes high right after the ST signal activates, and then slowly tapers off until the end of ST (in the case that the ROM word/bit is "1"). This would indicate that the ST signal does indeed allow some sort of sensing of a static magnetic field. While it is difficult to separate from noise, there is a distinct activation of the sense signal from the start of CUP, probably rising to the level of the germanium diode plus the 7404 inverter used to drive the low level of ST. Something along the lines of:

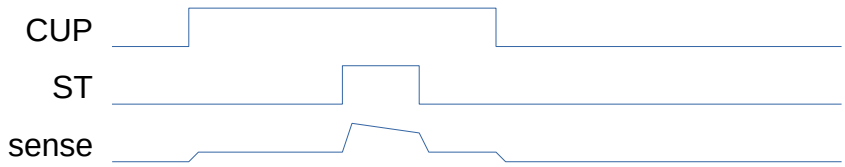

It could be the case that when the magnetic field begins at the start of CUP, the transformer coil presents a voltage (EMF) which is initially shunted by the 7404 and diode and limited to the voltage drop of the diode and 7404 transistor CE (less than 1V). Once ST releases the signal, it rises to somewhere around 2.5V where the input transistor of the SP380A turns on (setting the latch to "1") and then starts to decrease again as the current increases. Then at the end of ST the 7404 and diode again shunt the signal to ground, and it returns to less than 1V. At the end of CUP, the current through the wire stops and the magnetic field subsides as does the sense signal. It is possible that this level of detail cannot be inferred from the scope trace, but this would explain how it works.

None of the scope traces suggest that the “energize” signal is rising throughout. But if this was the case, it could explain why there is an EMF from the transformer coil during that time – the magnetic field is actually changing and so Faraday’s Law applies. But the available traces seem to show that Q33 simply turns on and the current into the selected wire is flat.

If the “energize” signal is not actually rising throughout, then I am tempted to think that an understanding of the Hall Effect might explain how it works.

Magnetic Core Memory

Wang used magnetic core memory in the 700 series, but switched to semiconductor memory in the 500/600. The semiconductor memory is fairly standard, with some models using dynamic RAM and others using static. The handling of dynamic RAM refresh was handled automatically by circuitry on the PCB.

The magnetic core memory has a fabric of X and Y wires that intersect at each unique bit (core). Half current is applied to X and Y, so that only the selected bit receives the full current required to flip its state. The direction of the current determines whether the core is set to “0” or “1”. Each bit-plane has it’s own sense wire and “bit” (suppress/inhibit) wire.

The memory is organized in 8-bit words. Each read or write involves all 8 bits.

Reading Core Memory

Reading of memory is destructive. The read cycle begins with a clearing (setting to 0) of the RA,RB S-R flip-flops, in preparation to receive the new data. Then a “force to 0” current is sent through the X and Y wires corresponding to the address specified in L,M,N. Any bit that transitioned from “1” to “0” will induce a current in its sense wire, which passes through the sense AMPs and is delivered to the RA,RB FFs “set” input, resulting in RA,RB containing the previous contents of the memory word by the end of the read cycle. The memory cores related to that word are now “0”, requiring a write-back in order to restore the previous contents. In cases where the intent was to clear the word, no write-back is required.

The only way to set a value in RA and RB is by memory read. Two successive reads of the same location (without any intervening write) will result in RA and RB being “0”.

Writing Core Memory

Writing of memory is accomplished by sending an opposite current through the X and Y wires for the word selected by L,M,N. This will (attempt to) force each bit to a “1”. The data supplied for the write will activate the bit wires for each “0” bit, suppressing the setting to “1”. Note that this suppression does not force the bit to “0” if it was already a “1”. If a write was done without a prior read, the new data is effectively (inclusive) ORed with the previous contents of the memory word. Writing a “0” word in such cases results in no change to the memory contents.

The RA and RB registers are used as the basis for the write data. However, one of those must be replaced with the output of the ALU. The two types of write data are thus “alu,RB” and “RA,alu”. The only way to set a value in RA and RB is by memory read. Two successive reads of the same location (without any intervening write) will result in RA and RB being “0”.

There are two main methods used in the 720C microcode to store an 8-bit value in a memory location. The “conventional” method is:

```
// KA,KB contain the new value
CA,CB=RA,RB=mem(LMN) // read current value
KA=KA; mem(LMN)=alu,RB // write new high nibble
      RA,RB=mem(LMN) // read latest value
KB=KB; mem(LMN)=RA,alu // write new low nibble
// CA,CB contain the previous value
```

An alternate (less obvious) method is:

```
// KA,KB contain the new value
CA,CB=RA,RB=mem(LMN) // read current value
      RA,RB=mem(LMN) // clear RA,RB
KA=KA; mem(LMN)=alu,RB // OR in new high nibble
KB=KB; mem(LMN)=RA,alu // OR in new low nibble
// CA,CB contain the previous value
```

Both methods take up the same amount of space in the microcode and use the same number of machine cycles. There does not appear to be any advantage of one over the other.

END