version date: 06/05/2001 - Added title bar to Colossus
version date: 12/02/2004 - Added M and Psi wheels patterns and Multitesting
version date 12/03/2004 - Added "count less than"

An interactive computer simulation of the World War II Colossus computer.
by Tony Sale.

1. Introduction.

This simulation of Colossus has been created to show how the real Colossus helped in breaking the German Lorenz cipher in Bletchley Park in World War II.

Tony Sale had started the rebuilding of a Colossus in 1993. Only minimal information had been declassified and was available at that time, but a start was made which resulted in the basic functions of Colossus being demonstrated in 1996. But there were still large areas of the code breaking work which were classified. In 2000 the Newmanry report was finally declassified. Although many more circuit boards had been identified and constructed, how they were connected together remained a mystery. Solving these mysteries has now been achieved by using Virtual Colossus to reproduce the code breaking procedures now revealed in the Newmanry report, the report by Albert Small, and the Fish Notes by Walter Fried, Americans seconded to Bletchley Park during the war.

Producing Virtual Colossus has been a major computer software task and the current version is still has a few problems. It does, however, demonstrate the setting of the Chi wheels , Motor wheels and Psi wheels as reported in Newmanry, Small and Fried.

The major omissions from this version is spanning and the rectangling gadget..

Only the version for Internet Explorer is currently available.

You might like to read the Appendix giving technical details of the simulation.

2. Getting started with Virtual Colossus.

Accessing colossus.htm from Internet Explorer brings up Colossus. Three frames appear.

The frame across the top shows the title bar. Scrolling this up reveals the test and monitoring fields and the counter outputs. It also reveals the pull down menus on the left for selecting the cipher text, the wheel start positions and the wheel patterns.

The next two frames contain scrollable pictures of Colossus. These pictures have to be scrollable to achieve a large enough picture size to make animation possible. Everything is viewable at 800x600 screen size. 1024x800 size will show more but may make operation of Colossus more difficult.

The left hand frame holds the Colossus switch panel. (see Small page 109 and Newmanry 53J page 340). Its details will be revealed as the tests progress.

The right hand panel contains a montage of various parts of Colossus. The top left hand area shows the Plug Panel. (Newmanry 53K page 343). Above the Plug Panel are the Spanning switches. (Not yet implemented) . Scrolling horizontally to the right shows first of all the big black ganged switches for setting signals to the left hand big switch panel. Next comes the Master Switch Panel and below it the wheel setting jacks for setting the wheel start positions. On the far right hand side are the patch panels for setting up Psi and Motor wheel patterns. Returning to the left hand side, below the Plug Panel is the wheel breaking panel which is set up with the Chi wheel patterns. then below the wheel starts panel is the set totals panel.

The paper tape reader is not shown. The data on the paper tape is held as a series of string constant of lengths from 2,000 to 9,000 characters. The character read from the tape appears in the Zin box on the top left hand of the computer screen. The paper tape data held in the simulator is either of original cipher texts or of original German texts re-enciphered.

The first thing to do is to load a cipher text. Click on the "cipher?" pull down menu on the far left side of the top frame. Click on "KHcipher" for starters. On the next pull down menu, "starts?" click on "cipher" This will set all the wheel starts to the correct ones for this cipher text. Then on the next pull down menu, "pattern?", click on "KHpattern". (There are 501 links to be set or cleared so this takes about 20 seconds on a 600 MHz lap top). the " Counts" box showing 0 0 0 0 0 indicates that all the links have been set.

The first thing to do is to 'read a character from the paper tape".

Now scroll the right hand image right to show the Master Switch Panel. This has two rows of mixed black and red keys on it.

Place the mouse pointer under the black switch with "T" above it and click to put the switch down.

This is the "Test and Trace" switch (not actually on the original Colossus)

Activating this switch simulates reading the next character from the paper tape and it appears in the Zin box which should now show J and its bit pattern, as the first character read. Activating it again gets the next character, 3.

Now we need to route these input characters to the main big switch panel.

Scroll the right hand panel until three large ganged switches show. These have "Q" written above them and Z, DeltaZ at the left had side of the top switch.

Place the mouse pointer just above the top horizontal black bar and when the little hand appears, click the left hand mouse button. The switch will be set up which connects the Z input (from the paper tape) to the Q connection at the top of the big switch panel shown in the left hand panel on the computer screen.

Now when the Test key is pressed the character read from the tape appears in the "Qin" box as well as the Zin box. So the input characters are appearing at the "Q" point on the big switch panel. The Q bit pattern is presented to all ten rows each of five switches, on the left hand side of the big switch panel. The convention is that bit 1 is at the left hand edge of a row and bit 5 on the right.

As explained in Nernanry and Small, these switches allow an input bit pattern to be tested true or false. The switches actually have three positions, up to test for a dot in BP terms (modern 0) or down to test for a cross (modern 1)
In the centre position (as shown when the Virtual Colossus is first loaded) no check is made on the input bit, i.e. either 0 or 1 gives true.

Thus only a bit pattern matching the pattern of the switches which are set on the row of switches will give true.

The result of this test feeds horizontally to the right to the Counter switches, labelled CNTRS on the real Colossus. These switches allow the true or false result to be switched to any of the five counters in Colossus. The counters are numbers one to five from left to right in each row.

So now lets count some characters coming off the tape. We will first count all the characters so no switches are set on the left hand top Q row.

We will count into counter 1 so use the mouse to set down counter 1 switch on the top row of Counters.

Now return to the Test switch on the Control panel and press it down. The next character will be read from the tape and this time the first count in the Counts box will go to 1. Pressing test again will result in two etc.

So now lets be really bold and count all the characters on the simulated tape.

On the Control panel press down the "C" key, (third from the right in the top row) . This clears the counters. Now press down the "LC" switch (Letter Count), third from the right on the bottom row. After about 6 seconds a count of 2042 will show as the number of cipher characters on the tape.

If you want to show that all five counters are working, set all five counter switches down in the top row and press LC again. All five counts will now show 2042.

The yellow switch to the left of the five counter switches is a "NOT" switch. It inverts the true or false from the test switches on the left, to false or true before the result gets counted. (remember the counters only register a true result) . So if the yellow switch is put down the true result from the switches will now be false so no counts. Press the LC switch to try it.

3. Getting outputs from the Chi (X) patterns.

The essence of Colossus, (and Tommy Flower's genius) is the production of the Chi patterns electronically in synchronism with the cipher text on the paper tape. So lets see how that was done in Colossus.

The first thing is the Chi wheel patterns, the mechanical lug positions round the Chi wheels used by the German cipher operator for this intercepted cipher text.

In this demonstration we assume that this has been worked out and is known.

The selecting of "KHcipher" in the "pattern?" pull down menu has put the patterns onto the wheel pattern board in the lower left part of the right hand image of Colossus.

On this board Chi wheel 1 is at the top and wheel 5 at the bottom.

Now we need to set the wheel start positions from which to test the match against the cipher text on the paper tape.

Scroll over to the Control panel. Below this are the wheel setting jacks, Wheel 1 setting is in the top row.

This shows the actual correct settings for this cipher text as put in by selecting "cipher" in the "starts?" pull down menu. These settings are K1=31, K2=1 K3=1, K4=15, K5=1.

Each setting block holds 20 jacks so 31 is just over half way along the top row in the second block. The jack hole positions are not very accurate. Look at the resultant setting in the "Strts" box at the top of the screen. If it is not correct clicking left or right of the plug will move it until 31 shows in the Strts box. The important thing is to end up with 31 1 1 15 1 in the Strts box.

Now scroll left to the big selector switches, click on the Z key to centre it and then click up the X switch to switch the generated Chi (X) patterns to the Q input to the big switch panel.

Press SU again to set the starts in, followed by Test, (Colossus needs input characters from the tape to get the sprocket hole signals which form the "clock" for Colossus and which advances the wheel readouts) the X character constructed from the Chi patterns will appear in the "Xin" box and in the Qin box (because X has been switched to Q). Note that the readings in the "Kpos" box advance from the set up position as the wheel rings precess round.

The first few characters of X are /FVJXU4O.Q.QOOLUF

4. An actual wheel setting run.

We have seen how to get both Z (cipher text) and X (generated wheel pattern characters) individually into Q. But they can be added together modulo two bit by bit by setting both the large Z switch and the large X switch.

However to do a wheel setting run we actually need DeltaZ and DeltaX (see Newmanry and Small).

So scroll to the large switches and set both Z and X down to give the Deltas.

We are going to do an X1 and X2 run against successive positions of X1, so move the X1 start set up position left until 27 shows on the Strts box, leaving all the other starts set.

Now go over to the big switch panel and clear (centre) all negate and counter switches on the top row.

Scroll the big switch panel down to reveal the red plus keys.

We are going to run the basic wheel setting algorithm 1p2 = . (dot)

This says we want to count the number of times, throughout the cipher text where Q bit 1 plus (modulo 2) Q bit 2 = (modern 0). Remember by setting the big Z and X switches, Q bit 1 already is DeltaZ1 plus DeltaX1 and 0 bit 2 is DeltaZ2 plus DeltaX2. (all modulo 2). Thus we are executing the famous double delta test: DeltaZ1 plus DeltaX1 plus DeltaZ2 plus DeltaX2 and recording whenever this comes
to zero.

To actually do this on Colossus we need to use the red keys with "+" between them. So clicking down key 1 and key 2 (from the left) will result in Qbit 1 plus Qbit 2 going left along the first row. Now look at the yellow key marked with dot and cross. Clicking this up to dot will enforce a test for "dot" (0) as the result of the red key additions, so put this into counter 1 by clicking down the red counter key next right.

Now return to the master switch panel on the right hand picture. We need to tell Colossus which wheels we want to step on each complete rotation of the tape. We do this by clicking down the left hand black key on the lower row of keys (X1) and clicking up the next key (X2) . This tells Colossus we want X1 to step fast (every revolution of the tape) and X2 to only step when X1 has got back to its set up position. (Newman 53D page 335) But a full run of X1 and X2, 1271 positions, takes a long time, so just to show how it works, click down the "10" key , 4 from the right, bottom row. This only allows 10 positions to be explored before stopping, hence the setting back of the X1 start by 5 places.

Now lets do the run. First press the "SU" key to make sure everything is set up. Next press the C key to clear the counters and now press the "M" key, the Master key second from the right, top row.

A problem with Internet Explorer is the big grey slow running Alerts panel that comes up. You can click "no" to Abort? to carry on but the Printout window has now been pushed behind the Internet Explorer window and can't be seen. It will appear next time anything is printed onto it but if you want to bring it to view, click on the downsize button, the centre of the three on the right hand end of the blue top bar. This should reveal at least part of the Printout window and clicking on that will bring it to the front.

The Printout window will be headed "K1 K2 count" and after every two seconds it takes to scan the tape, the wheel positions and the resultant count will be printed.

You can see on this the obvious maximum score of 1087 against the correct position, 31 1. The average score would be half 2042 i.e. 1021.

5. Letter counts.

Small makes considerable use of letter counts in DeltaD to check possible wheel setting results. Virtual Colossus can do these too.

First return X1 wheel set up to 31. Now clear the switches in the red lower red area of the big switch board. Scroll up to the top left hand area of the switch panel and set the following patterns in the test switches under 0.

11001 for W, 11011 for 5,10011 for B and 00000 for / (remember 1 is down, 0 up)

Now set row one into counter 1, row 2 into counter 2 etc.

Go to the Control panel and press down SU and C, then press down LC.

After 6 seconds the counters show 60, 84, 50, 91.

That is in 2042 characters of DeltaD, (DeltaZ plus DeltaX) , W scores 60, S scores 84, B scores 50 and / scores 91.

Small page 5 quotes W as 89, 5 as 143, B as 82 and I as 128 in a 3,200 length of DeltaD. When our scores are corrected to 3,200 the agreement is very good.

6. Multiple stepping.

Colossus contained "Remembering" circuits and had four stages which could hold the previous four bits of any Chi, Mu or Psi stream. This allowed five comparisons to be made for each cipher character read, the appropriate wheel could then be stepped on five places resulting in much quicker runs.

The remembered bits appear on five channels under "R" on the big switch panel.

Virtual Colossus cannot yet do all the tricks that this allowed but the principle can be shown as follows.

First clear all the switches on the large switch panel.

Go dawn to the bottom left hand of the switch panel. Set down all the red switches second from the left (the Q2 channel) -

Set down all the staggered yellow keys in the R channel. Set up to dot, all the yellow test switches. Now put down counter 5 on the first row, counter 4 on the second etc. This is so that the printout comes out in the right order.

Now go to the Master switch panel and put up the red key marked with a X1 above it. Check that on the lower row, X1 key is down and X2 key up.
Go to the setting plug for X1 and move it on to 33.

What this all means is that Colossus has been told to "remember" four bits back on X1 (the top red key), these remembered bits appear on the staggered yellow keys on the switch panel. The switching there means that in the third row, Qbit 2 is added to X1 bit (Rbit) two back and the result counted in counter 3.

Now press down SU and then LC. After 3 seconds the four counts at position back from X1 33 show in the Counts box. The high count at 31 shows again.

7. A longer run using multiple testing.

Check that the "10" key is down to limit the number of iterations round the tape to 10. Set the X1 start back to 28. Now press SU followed by M. The output window will show rows of five counts preceded by a lower case letter a to e. The wheel positions on the left of the printout correspond to the "T" count position and the other counts are successively BACK from this position.

This is where the "set totals" became so important. The default set by the vertical column of five switches on the right hand side is "greater than" only counts greater than the set total value are printed. By choosing a fairly high value, only a few counts will be shown. The top key, counter 1 key, can also be set left to give "less than".

To set the Set Total, scroll the right hand panel to its lower right hand corner. This will bring the set totals panel into view. Only the set total for counter 1, on the left hand side of the panel has been activated. By clicking on the clockwise and anti-clockwise arrows to the right of the setting knobs they can be rotated until their pointers show the required set total value. For this run select 1050 as the set total. The greater than test is the default and this set total applies to all counters.

Now pressing SU followed by M performs 50 iterations through the tape with 5 Chi wheel positions for each iteration.

8. Playing with Virtual Colossus.

You can now explore all sorts of other Colossus runs within the limitations of Virtual Colossus. The biggest limitations at present are:

a) the slow speed of the simulation (Javascript) which means that only runs of ten or twenty wheel positions are feasible as continuous runs unless your PC is over 800Mhz processor speed..

b) no Spanning

c) Multitest not yet working on Motor and Psi wheel setting.

All the above will be addressed and incorporated but what is already in Virtual Colossus allows the simulation of most of the basic functions described in Newmanry, Small and Fried.

Tony Sale, March 2004

Appendix.

Virtual Colossus, the software and simulation details.

1. Javascript. This simulation has been written, with great difficulty, entirely in Javascript so that the whole simulation can be downloaded and run from the Net.

As a result of this decision the programs are BIG and the simulation runs slower than the original Colossus. (about one second with a 600Mhz PC to scan and process 2,000 input characters, original Colossus, 5,000 characters per secondi) It may still be possible to optimise the core code further. Tony Sale defensive programming is used throughout, (KISS rules) . The Javascript is all there for anyone to see in the download.

2. Hit boxes. As you can imagine there are literally hundreds of hit boxes all over the image pictures of Colossus. Positioning these is a real pain and some of the alignments still need tweaking. Because of the small size of the Colossus switches, it has been found essential to use large images and scroll them.

Each switch key on Colossus has at least a hit box below it and some switches which can be set up or down have a hit box above as well. The cursor changes to a right hand with extended first digit when a hit box is entered. But sometimes this does not appear. The best thing seems to be to move the mouse pointer off the current image and come back and try again. A similar effect occurs when, after clicking on a hit box, the hand pointer remains even after moving off the hit box. Again flicking the pointer off the current image causes it to revert to just an arrow pointer.

A more annoying effect is that occasionally when the mouse left button is clicked the current image reverts to its unscrolled position. I'm still working on solving this.

3. Netscape version. Because of the different way Netscape handles small image manipulation, a completely different version of Virtual Colossus is required. Not really difficult, {I have done it for my Enigma and Bombe simulations) but it takes time. The codesandciphers access logs show 80% I Exp usage so 20% of you are disenfranchised for a short time.

4. A table of bit patterns for characters: (bits 1 to 5 left to right)

A 11000 B 10011 C 01110 D 10010 E 10000 F 10110 G 01011 H 00101

I 01100 J 11010 K 11110 L 01001 M 00111 N 00110 O 00011 P 01101

Q 11101 R 01010 S 10100 T 00001 U 11100 V 01111 W 11001 X 10111

Y 10101 Z 10001

The special characters:
3 00010 4 01000 8 11111 5 or + 11011 9 or . 00100 / 00000