

## "How to Make ENIAC's Accumulators Addressable Using a Subroutine"

Mark Priestley & Thomas Haigh, January 2016.

[www.EniacInAction.com](http://www.EniacInAction.com)

Discussion of early computer architecture has often been more concerned with computability than with practicality. Inspired by work on Turing machines and “universality,” a kind of game has arisen: what could an early computer have done with unlimited time and storage but no other changes to its hardware. This has most often been done by proving the equivalence of a novel or early computer to a Turing machine. As we note in *ENIAC in Action*, ENIAC in its original 1946 is almost universally recognized to be “Turing complete” although to the best of our knowledge no formal proof has been published.<sup>1</sup>

ENIAC’s original control method was modified in 1948, after which point its wires and switches were left mostly untouched while it ran only a single (but slowly evolving) program: a microcoded interpreter for a virtual von Neumann architecture machine. ENIAC’s application programs were written as a series of two digit instruction codes for this virtual machine and loaded into its read-only function table memory by turning knobs to set digits. The question of whether this virtual machine was a true “stored program computer” is less clear in the existing literature. So is the distinct, but often conflated, question of whether the virtual machine implemented on ENIAC in 1948 was itself a “universal” or “Turing complete” computer. There is, of course, no guarantee that a program running on a universal machine is itself an interpreter for another universal machine language. In fact most programs are not.

As we discuss in *ENIAC in Action*, it has sometimes been asserted that ENIAC was not a true “stored program computer” because it could not modify code held in memory and so could not carry out a full range of computations. We are skeptical about the usefulness of “stored program computer” as an analytical category.<sup>2</sup> If one defines it as requiring the storage of code in writable memory then this definition would clearly, if arbitrarily, exclude ENIAC, or any computer that held code in read only memory. However we are confident that ENIAC’s computational universality in its post-1948 configuration was in no way limited by the kind of memory in which it held its code.

This question directed us to the specialized literature addressing the minimal requirements for a von Neumann architecture computer to be “universal.” Raúl Rojas has drawn on this work in several publications, highlighting the requirement that the computer provide *either* the capability to modify addresses held in memory *or* include an indirect addressing mode for its store and load instructions.<sup>3</sup> This drew on the influential work of Calvin C. Elgot and Abraham Robinson to create the Random Access Stored Program (RASP) machine model, as an alternative to the Turing Machine based on features of real computer architecture such as the storage of programs and data in random access memory. They

---

<sup>1</sup> Thomas Haigh, Mark Priestley, and Crispin Rope, *ENIAC In Action: Making and Remaking the Modern Computer* (Cambridge, MA: MIT Press, 2016).

<sup>2</sup> T. Haigh, M. Priestley, and C. Rope, "Reconsidering the Stored-Program Concept," *Ieee Annals of the History of Computing* 36, no. 1 (Jan-Mar 2014):4-17.

<sup>3</sup> Raúl Rojas, "Who Invented the Computer? The Debate from the Viewpoint of Computer Architecture," *Proceedings of Symposia in Applied Mathematics* 48 (1994):361-36, Raúl Rojas, "How to Make Zuse's Z3 a Universal Computer," *IEEE Annals of the History of Computing* 20, no. 3 (Jul-Sep 1998):51-54.

wrote that the RASP “in a sense that is made precise, can compute all recursive functions (where the initial content of the machine, including the stored program, depends on the particular function which is to be computed and, to a lesser extent, on the values of the argument or arguments).”<sup>4</sup>

We also note that John von Neumann’s original “First Draft” design for the EDVAC explicitly prevented code modification except to address fields. This design is generally accepted as the key text in the codification and dissemination of modern, “stored program” computer architecture (though not everyone believes that von Neumann himself came up with its most important ideas). Just like the post-1948 ENIAC, von Neumann’s EDVAC would have been able to change the addresses used by jumps or arithmetic operations at runtime but not to substitute one operation code for another. Later computers frequently ran programs from read only memory (or from RAM that is protected by hardware from being overwritten), and any computation carried out by a program held in RAM could also be carried out by an equivalent program held in ROM, providing there is enough RAM for storage of variables and working data.

But did ENIAC’s 1948 instruction set really have indirect addressing? The instruction set always used indirect addressing when reading code and constant data from its large read only function table memory and to set the destination of jump instructions. However, as we discuss on page 256 of *ENIAC In Action*, its standard 1948 instruction set did not provide an indirect addressing mode, or indeed any form of addressing, for variables held in its writable accumulator memory (the ENIAC equivalent of RAM). Instead of using the same method as the function table instructions, where generalized load instructions (FTN, FTC) took the specific memory location to work from the value stored in a register-like location of accumulator memory, the conversion team defined separate two digit load and store instruction codes for each of its twenty accumulators.

For example, machine language instruction 23 loaded the contents of accumulator 3 into the virtual machine accumulator register (accumulator 15), while machine language instruction 03 stored the contents of the virtual machine accumulator in accumulator 3. This worked well in practice, but imposed certain theoretical limitations on dealing with data structures held in accumulator memory. For example, without indirect addressing or code modification to change the location acted on one could not write a loop to add a series of numbers held in consecutive accumulators.<sup>5</sup> One could avoid a loop and instead write a series of accumulator-specific instructions, but this would not work with a list of unknown length.

### **Two Ways to Implement Indirect Accumulator Addressing on ENIAC**

We suggest in *ENIAC In Action* that the configuration used to implement its post-1948 instruction set could easily have been modified to provide the accumulator memory with the same indirect addressing capability as the function table memory, if any practical need for such a feature existed. We also note that just such a change was planned for the new “register code” that was to be implemented once a larger writable delay line memory was installed. Our experiments have shown that even in ENIAC’s

---

<sup>4</sup> Calvin C Elgot and Abraham Robinson, "Random-Access Stored-Program Machines, an Approach to Programming Languages," *Journal of the ACM* 11, no. 4 (October 1964):365-399.

<sup>5</sup> As the function table memory used indirect addressing there would be no problem in using a loop to sum a series of numbers held there.

original 1946 control mode, with no changes to its hardware at all, clever use of the master programmer would have made it possible to work with data structures of indeterminate length held in accumulators.<sup>6</sup>

So if any pressing reason had arisen to make ENIAC's accumulator memory addressable in the 1948 instruction set, such as the addition of more storage capacity, then the historical evidence is quite clear that its users would have known exactly how to do this. The fact that they did not indicates that the penalties incurred by doing so outweighed the advantages. The costs of making the accumulators addressable included the commitment of precious logic circuits that would no longer be available to implement other instructions, an increase in the size of instructions if inline parameters were used for direct accessing, an increase in the number of instructions for programs using indirect addressing, a loss three digits of memory space to implement indirect addressing, and a loss of performance when the code was executed.

Although addressability would have been added in practice by changing the ENIAC instruction set this might be seen as cheating by those interested in computability. Their game is traditionally played by assuming that storage can be scaled up as needed and that unlimited time is available but that other changes are forbidden.

This brings us to a less efficient and less historically plausible way of adding indirect addressing to ENIAC's accumulators: writing subroutines. In *ENIAC in Action* we address by briefly asserting that "Even if such changes were, for some reason, ruled out, then, accumulator memory could have been made addressable by writing parameterized storage subroutines." Footnote 74, on page 334 elaborates this a little:

What if one plays the theorists' traditional game of exploring what a machine could do if granted vast amounts of time and storage but left otherwise unchanged? To make the accumulators addressable, one could simply write a pair of subroutines: one to store and one to load. Each would take as a parameter an accumulator number, which would be used to calculate a jump to the function-table address at which the appropriate "listen" or "talk" instruction has been placed. That would tie up two rows of a function table for each accumulator rendered addressable, but unlimited storage would already have been assumed.

We have produced just such a pair of subroutines and tested them using an emulator. To understand the code pieces below you might find it helpful to examine our summary of the 1948 ENIAC instruction set on pages 168-170 of *ENIAC in Action* or the detailed descriptions of instructions given in the 1949 document "Description and Use of the ENIAC Converter Code" (available from <http://eniacinaction.com/docs/DescriptionandUse1949.pdf>).

### **Background: Subroutine Returns in ENIAC**

The Monte Carlo computations of 1948 made use of what is called a "closed" subroutine, namely one called from several points within a program. At the end of the subroutine, control must jump back to different places depending on the address from which the subroutine was. The return address was stored in the three digits of accumulator 6 used to hold the address that a conditional branch instruction

---

<sup>6</sup> Similar techniques allowed Adele Goldstine to devise an initial instruction set for ENIAC's conversion that used only the native 1946 ENIAC hardware. The use of specialized hardware in 1948 made the conversion more practical, but did not increase ENIAC's computational capabilities in principle.

would jump to (known in ENIAC notation as “6(6,5,4)”). At the end of the subroutine, executing a conditional transfer instruction with a positive or 0 value in accumulator 15 would cause the program to jump back to the required return address.

Below are the beginning and the end of the preserved code for the Monte Carlo subroutine. It ends with a “CT” to jump back to the stored return address. The penultimate instruction, COUNT, would have left 0 in accumulator 15.

FT	Symbol	Code	Effect on accumulator contents
171	CL	15	[15] = 0
	16T	36	[15] = abcdefghij = $\xi'$
	S'R1	38	[15] = 0abcdefghi ; [12] = j000000000
	12T	62	[15] = jabcdefghi = $\xi$
	S'L5	86	[12]=0 <sup>5</sup> jabcd=10 <sup>-5</sup> $\xi_0$ ; [15]=efghi0 <sup>5</sup> =10 <sup>5</sup> $\xi_1$
	11L	11	[11] = 10 <sup>5</sup> $\xi_1$ ; [15] = 0
172	12T	62	[15] = 10 <sup>-5</sup> $\xi_0$
....	....	....	....
....	....	....	....
176	S'L1	66	[12] = 000000000j, [15] = abcdefghi0
	12T	62	[15] = abcdefghij = $\xi$
	16L	16	[16] = $\xi$
	COUNT	<b>N</b>	Halt after 3000 squarings
	CT	69	→ $\omega$

We use the same mechanism to provide the return address to our load and store subroutines.

### Code to Implement Indirect Addressing

ENIAC had twenty accumulators, each storing ten digits and a sign, but all or part of many of them were not available for program data storage because they were used for special purposes. Page 167 of *ENIAC in Action* describes the use made of the different accumulators. So to create a continuous range of addressable memory locations we need to allocate a series of “logical addresses”. We might let logical address 1 denote physical accumulator 3, logical address 2 denote physical accumulator 4, logical address 3 denote physical accumulator 18, and so on. We denote the contents of logical address n by [n], so using this scheme [1] would be the contents of Accumulator 3 etc.

Store and load subroutines need two parameters: the logical address to use and the quantity to be stored or loaded. ENIAC customarily used accumulator 15 as what would typically be called “the accumulator” in a von Neumann architecture machine. That would be the most natural place to put the number to be stored or loaded, but unfortunately accumulator 15 will be needed within the store and load subroutines to compute the address we jump to and so would be overwritten. In the code below, we use ENIAC’s accumulators as follows:

Accumulator 1	"Logical address" of memory location to be accessed
Accumulator 2	Holds the number being moved in or out of a logical memory location
Accumulator 3	Contents of logical memory location [1]
Accumulator 4	Contents of logical memory location [2]
Etc.	Etc.

Our approach requires the storage on the function table of two short fragments of code for each logical memory location, one for the "talk" or "Load" instruction that retrieves data from the corresponding physical accumulator, and one for the "listen" or "Store" instruction that copies data into the corresponding physical accumulator. The load and store subroutines each perform a calculated jump to the location where the appropriate code fragment is stored. We will interleave these code fragments in a function table. If the code for this purpose is stored starting at FT 500 then it would be laid out as follows:

500: entry point for Load subroutine  
 501: entry point for Store subroutine  
 502: code to Load logical memory location 1  
 503: code to Store logical memory location 1  
 500 + 2n: code to Load logical memory location n  
 501 + 2n: code to Store logical memory location n

As only thirteen of ENIAC's accumulators were been available for entirely unrestricted data, and we have just used two of them for parameters, the real ENIAC could have had at most eleven logical memory locations, tying up 24 rows (500-523 in our example) of the function table to implement the subroutines and code fragments.

Here is the actual ENIAC code required to implement this for memory location 1, corresponding to physical accumulator 3. Rows 502 and 503 would be repeated, with the necessary minor changes, for the each logical memory location.

F.T. Row	Symbol	Code	Comment
500:	N4D15	73	
		05	
		00	store 500 in Acc 15
	1T	21	Copy logical address n from Acc 1 to Acc 15 (twice)
	1T	21	Acc 15 = 500 + 2n
	6R3	78	Right 3 digits of Acc 15 to Acc 6, ie jump to 500 + 2n
501	N4D15	73	
		05	
		01	store 501 in Acc 15
	1t	21	Add logical address n from Acc 1 to Acc 15 (twice)
	1t	21	Acc 15 = 501 + 2n
	6R3	78	Jump to 501 + 2n
502:	3T	23	Move [1] ( i.e. contents of Acc 3) to Acc 15
	2L	02	Move contents of Acc 15 to Acc 2
	CT	69	Jump to return address as Acc 15 is now 0

		00	<i>(Rest of row unused)</i>
		00	
		00	
503:	2T	23	Move contents of Acc 2 to Acc 15
	3L	03	Move contents of Acc 15 to [1] (Acc 3). Acc 15 clears.
	2L	02	Copy 0 from Acc 15 to Acc 2 (i.e. clear Acc 2)
	CT	69	Jump to return address as Acc 15 is now 0
		00	<i>(Rest of row unused)</i>
		00	

Accumulator 2 is cleared after a "Store" operation. This is consistent with the behavior of the converter code, which clears Accumulator 15 after its value is transferred to another accumulator.

A typical calling sequence for a Load operation would then be:

F.T. Row	Symbol	Code	Comment
101	N2D15	72	
		01	Place logical address (here 1) in Acc 15
	1L	01	Copy logical address number from Acc 15 to Acc 1
	N6D6	84	Set up return address 103
		10	
		35	
102		00	Jump to row 500.
		00	<i>(Rest of row unused)</i>
		00	
		00	
		00	
		00	
103	..		Return point. Value of [1] is now in Accumulator 2. Carry on with the program.

Assuming that the value to be stored has been computed and placed in Accumulator 2, a typical calling sequence for a Store operation would be:

F.T. Row	Symbol	Code	Comment
101	N2D15	72	
		01	Place logical address (here 1) in Acc 15
	1L	01	Copy logical address number from Acc 15 to Acc 1
	N6D6	84	Set up return address 103
		10	
		35	
102		00	Jump to row 501
		00	<i>(Rest of row unused)</i>
		00	
		00	
		00	
		00	

103	....		Return point. Value in Accumulator 2 is now stored as [1], and Accumulator 2 has been cleared. Carry on with the program
-----	------	--	-----------------------------------------------------------------------------------------------------------------------------

This method develops an idea first mentioned by Crispin Rope, our collaborator on this project, in a 2007 article. Disputing a characterization of ENIAC by Brian Randell, Rope observed that ENIAC

could perform an unconditional jump to one of a series of orders, each of which addressed a different memory location. This selection of the “jump to” instruction was determined by a number stored in an accumulator which could be based on results so far computed. The extent of the read-write memory was small, but in principle the ENIAC in stored-program mode could be said to have “general-purpose programming ability” in the terms set by Randell.<sup>7</sup>

### Test Case – Summing a Series

As a test case for these subroutines, we show below a routine to sum a series of numbers stored in the logical accumulators. The length of the series is not known in advance, but the end of the series is marked by a negative number, the other numbers being assumed positive or zero. This example demonstrates that with the 1948 conversion code, ENIAC could manipulate dynamically sized data structures. On a conceptual level, it implements a very simple example of the class of “sequential functions” defined by Elgot and Robinson in their landmark paper (“Now, digital computers may well be called upon to compute a function, say, one which takes a finite sequence of numbers of “arbitrary” length into their sum ...”p.393)

The code below assumes that the series to be summed is already set up in the logical accumulators. It forms the sum of the series in physical accumulator 10 (which therefore can’t be used to hold a logical accumulator). When a negative value is found in a logical accumulator, the code jumps out of the loop and control is transferred to address 099.

F.T.Row	Symbol	Code	Comment
050	N2D15	72	Copy next 2 digits to Acc 15
		01	Acc 15 = 1, the initial value of the index i.
	1L	01	Move Acc 15 to Acc 1. Acc 1 = 1, Acc 15 = 0.
		99	
		99	
		99	
051	N6D6	84	Copy next 6 digits to Acc 6.
		05	“500” is the entry point for the Load
		25	subroutines. “052” is the return address.
		00	Transfer to the subroutine.
		00	<i>(Rest of row unused)</i>
		00	
052	N6D6	84	Store next 6 digits “099053” in Acc 6.
		09	“099” is the address the conditional jump will transfer

<sup>7</sup> Crispin Rope, "ENIAC as a Stored-Program Computer: A New Look at the Old Records," *IEEE Annals of the History of Computing* 29, no. 4 (Oct-Dec 2007):82-87.

		90	control to when a negative number is encountered
		53	Unconditional jump to 053
		00	<i>(Rest of row unused)</i>
		00	
053	2T	22	Copy contents of Acc 2 to Acc 15
	M	41	Change sign of Acc 15
	CT	69	If contents of Acc 15 >= 0 go to 099
	N4D6	83	
		00	
		54	Unconditional jump to 054
054	10T	30	Copy contents of Acc 10 to Acc 15
	2T	22	Add contents of Acc 2 to Acc 15
	10L	10	Move updated total from Acc 15 to 10
		99	
		99	
		99	
055	N2D15	72	
		01	Acc 15 = 1 (to increment index)
	1T	21	Add Acc 1 (index i) to Acc 15 = i + 1
	1L	01	Move i + 1 from Acc 15 to Acc 1
		99	
		99	
055	N4D6	83	
		00	
		51	Go to 051 for next iteration of loop

This code has been tested with our ENIAC emulator running the 1948 conversion code, and using a sequence of logical accumulators defined over a non-consecutive sequence of physical accumulators. The “99” codes are dummy instructions, inserted so that functionally coherent pieces of code fit into single rows in the function tables. More idiomatic programming in the style of 1948 would remove these and compress the code as much as possible to save space. Notice that the design of the instruction set meant that apparently redundant unconditional jumps to the next instruction did occur in various circumstances in the original Monte Carlo program.

#### Discussion – Use with Nested Subroutines

The technique used by the Monte Carlo programmers to store the subroutine return address would not work for a subroutine that called another subroutine, since the return address for the first subroutine would be overwritten when the second one was called. A similar problem faced von Neumann’s team at the Institute for Advanced Studies in the third volume of their Planning and Coding report, which dealt with subroutines.<sup>8</sup> They proposed a way of automatically processing a program to set the appropriate values. David Wheeler, of the Cambridge University EDSAC team, came up with the scheme known as

<sup>8</sup> Reproduced in William Aspray and Arthur W Burks, *Papers of John von Neumann on Computing and Computer Theory* (Cambridge, MA: MIT Press, 1987).



the “Wheeler jump” as a more elegant solution to this problem.<sup>9</sup> Eventually the introduction of stack capabilities into computer architecture made it easy to save the return addresses, variables, and parameters used by a subroutine safe when it called another subroutine. This also provided an efficient mechanism for recursive calls.

We know of no historical examples of nested subroutines in surviving ENIAC code, but there is no conceptual reason that they could not have been used provided the programmer took responsibility for using a different memory location to store the return address for each subroutine involved. They would certainly be pleasant to have in the utopian world inhabited by computability theorists, where storage and time are unconstrained. Here we present an example that demonstrates that the indexing subroutines can be called from both the main program and from within another subroutine.

The technique adopted here is for the subroutine that calls the indexed addressing subroutine to save its own return address somewhere else in memory before calling the indexed addressing subroutine. The main program fragment calls the indexing subroutine directly to store a value in logical accumulator 1. It then calls subroutine A, which itself calls the store and load subroutines to double the value held in logical accumulator 1. Before doing so, it stores its return address in accumulator 20, and then then retrieves it in order to return to the main program via an unconditional jump.

	Symbol	Code	Comment
<b>Main Program in FT1</b>			
000	N2D15	72	Copy next 2 digits (“01”) to Acc 15
		01	
	1L	01	Move “01” (id of logical accumulator) to Acc 1
		99	
		99	
		99	
001	N6D15	74	Copy next 6 digits (“123456”) to Acc 15
		12	
		34	
		56	
	2L	02	Move to Acc 2
		99	
002	N6D6	84	Copy next 6 digits to Acc 6.
		00	“501” is the entry point for the Store
		35	subroutine. “003” is the return address.
		01	Transfer to the subroutine to store “123456” in logical acc [1]
		00	<i>(Rest of row unused)</i>
		00	
003	N6D6	84	Copy next 6 digits to Acc 6.
		00	“100” is the entry point for the subroutine to double
		41	logical acc [1]. “004” is the return address.
		00	Call subroutine A to double value in logical acc [1]
		00	<i>(Rest of row unused)</i>

<sup>9</sup> Described in Martin Campbell-Kelly, "Programming the EDSAC: Early Programming Activity at the University of Cambridge," *Annals of the History of Computing* 2, no. 1 (January 1980):7-36.

		00	
004	...	...	Logical acc [1] now holds $2 \times 123456 = 246912$
<b>Subroutine A in FT2</b>			
100	6T	26	Copy addresses from Acc 6 to Acc 15
	SR3	42	Shift right 3 places to extract return address
	20L	20	Save it in Acc 20
		99	
		99	
		99	
101	N6D6	84	Copy next 6 digits to Acc 6.
		10	"500" is the entry point for the Load
		25	subroutine. "102" is the return address.
		00	Transfer to the subroutine to store "123456" in logical acc [1]
		00	<i>(Rest of row unused)</i>
		00	
102	20T	22	Copy value in Acc 2 to Acc 15
	20T	22	Add value in Acc 2 to Acc 15
	20L	02	Move doubled value in Acc 15 back to Acc 2
		99	
		99	
		99	
103	N6D6	84	Copy next 6 digits to Acc 6.
		10	"501" is the entry point for the Store
		45	subroutine. "104" is the return address.
		01	Copy value in Acc 2 to logical acc [1]
		99	
		99	
104	20T	40	Copy saved return address from Acc 20 to Acc 15
	6R3	78	Copy address to Acc 6 and jump back to main program
		00	<i>(Rest of row unused)</i>
		00	
		00	
		00	
<b>Index Subroutines in FT3</b>			As above

This code has been tested with our ENIAC emulator running the 1948 conversion code. Its rather profligate use of storage would not make this technique viable for regular use on the real ENIAC, where accumulator memory was precious, but demonstrates that in principle ENIAC programmers could have made use of nested subroutines.

### Conclusion

We have shown here that a writable addressable data memory, with indirect addressing capabilities, could have been implemented on ENIAC as configured in 1948 without any changes to its instruction set.

All that is required is the use of subroutines for store and load operations and the sacrifice of two rows of function table memory for each ten digit addressable memory location created.

Coupled with ENIAC's existing instructions for conditional branching and arithmetic, this indicates that its 1948 instruction set would not have required any modification to satisfy the basic requirements identified by Rojas for a "universal computer." These were "CLR, INC, LOAD, STORE, conditional branches and indirect addressing (or equivalently self-modifying programs)."<sup>10</sup> This finding calls into question many previous characterizations made of ENIAC's computational power made on grounds that certain fundamental features were missing from its architecture.

For example, Arthur Burks, one of the original designers of ENIAC, later wrote that the address substitution capability introduced by John von Neumann in his "First Draft" design for EDVAC gave it a revolutionary capability lacking in ENIAC even after its 1948 conversion.<sup>11</sup> In fact ENIAC's 1948 instruction set had, within the limits of available storage, exactly the same computational power as the First Draft EDVAC thanks to its use of indirect addressing for program jumps and the potential we have demonstrated here to use subroutines to store and retrieve data from an addressable, writable, logical memory.

---

<sup>10</sup> Rojas, "Who Invented the Computer? The Debate from the Viewpoint of Computer Architecture".

<sup>11</sup> See discussion in *ENIAC in Action*, page 245.