

The “M” and the “4” in “M4”

Chris Thornton

bitsbelow.us

published 2020-07-15

I always thought M4 was named with the “M” for “macro” and the “4” for being the fourth in a series of macro processors for UNIX. I thought this even though I wasn’t aware of any prior programs named M3, M2, M1, or M, or whatever this initial version would have been called. I’m not the only one to think this,¹ and I wanted to know, is this folklore or fact?

So, let’s investigate if (1) the “M” in “M4” does stand for “macro”, and if (2) the “4” in “M4” is for being the fourth in a series of macro processors. Like any serious inquiry, let us begin with Wikipedia; the M4 entry says:²

*[M4] is an extension of an earlier macro processor, written by D. M. Ritchie for an unknown AP-3 minicomputer.*³

The entry’s History section⁴ mentions the M6 macro processor⁵:

In [the] 1960s,⁶ an early general-purpose macro processor [M6] was in use at AT&T Bell Laboratories[...]

However, this statement does not place M6 anywhere within the chronology of M4 predecessors. M6 simply existed as an early macro processor at Bell Laboratories but otherwise appears disconnected from M4 and M3 lineage.

The existence of an early macro processor named M6 does not fit the M4, M3, M2, and M narrative. In order to answer our two questions we must dig into the histories and influences behind M4 and each M4 predecessor recursively.

But uncovering all influences is impractical. Some influences exist only in the social sphere of the authors at the time and place the software was written and would require an extensive investigation to determine. Our search will be limited to stated influences only, and consequently reduce the scope of research. The term

“stated influences” refers to the influences explicitly stated by the principal authors, irrespective of where the claim was made—in technical reports, mailing lists, books, memoirs, interviews, or anywhere else. We are assuming that stated influences have greater bearing over naming, at least in the cases where program names are increments of prior major versions—as seems to be the case with M4 and its immediate predecessor, M3. If the validity of this assumption does not hold then uncovering stated influences still serves as a starting point for further investigation.

The first manual for M4,⁷ Bell Laboratories Technical Report, The M4 Macro Processor, by B. W. Kernighan and D. M. Ritchie states:⁸

The M4 macro processor is an extension of a macro processor called M3 which was written by D. M. Ritchie for the AP-3 minicomputer[...]

This is the source for the paraphrased sentence in the M4 Wikipedia entry.⁹

Little is known about M3. We do know D. M. Ritchie was the principal author of M3.⁸ We do not know if M3 was written in C—as M4 and many other programs by D. M. Ritchie. No known record of why D. M. Ritchie wrote M3 exists. No known record of exactly when D. M. Ritchie wrote M3 exists. No known record of who owned the AP-3 minicomputer exists.¹⁰ We don’t know if the “3” in “M3” is for being the third in a series of macro processors, or if it stood for AP-3, or something else.

Continuing onward, we uncover M3’s stated influence:⁸

[...] M3 was in turn based on a macro processor implemented for [the book Software Tools]

From Software Tools, we can see that M3 was based on the macro processor detailed in chapter 8, “Macro Processing”.¹¹ The authors, B. W. Kernighan and P. J. Plauger, refer to this macro processor as macro.¹²

Interestingly, on the last page of chapter 8, B. W. Kernighan and P. J. Plauger write:¹³

The macro processor described in this chapter was originally designed and implemented in the language C by [D. M. Ritchie]; we are grateful to him for letting us steal it.

Thus, macro’s stated influence is another macro processor by D. M. Ritchie. Little is known about this macro processor. No known record of this macro processor exists beyond this single sentence in Software Tools. We don’t even know this macro processor’s name. I refer to D. M. Ritchie’s pre-macro macro pro-

cessor as `macro'` (pronounced macro-prime). Since there is no known record detailing the stated influences of `macro'`, we cannot recursively trace M4 and its predecessors's stated influences any further. We do not know the stated influences of `macro'`—it may state M6, it may state several influences, or it may state no influences at all.

As we've traced M4's stated influences this far, I think we have sufficient evidence to substantiate at least one claim: the "M" in "M4" likely stands for "macro". We can claim this with a reasonable level of confidence since these programs process macros as their primary mode of operation, and because the M4 predecessor implemented in Software Tools has the name "macro".

However, I don't think the available evidence supports the claim that the "4" in "M4" is for being the fourth in a series of macro processors for UNIX. As we've traced the stated influences from M4 to M3 to `macro` then to `macro'`, we count four macro processors in this series. It is tempting to conclude that these four macro processors make up the entirety of the series, but we know almost nothing of M3 and `macro'`. Moreover, if these four macro processors make up the series, then we exclude the earlier Bell Labs macro processor, M6, with deep similarities to M4.

Rather, the available evidence leads to the idea that the "4" in "M4" follows a progressional naming pattern similar to that of the C programming language. The names of both languages—C and M4—are increments of prior major revisions. "C" is an increment of its predecessor's name, "B".¹⁴ However, the predecessor of the B programming language is BCPL, not the (fictitious) A programming language.¹⁵ The progressional naming pattern started with C. Likewise, the name "M4" is a increment of its predecessor's name, "M3", and M3 does not have a predecessor named M2. The progressional naming pattern started with M4.

Thanks to Andrew Healey <healeycodes.com> for reviewing an earlier draft.

-
1. Well, at least one other person asked what the "M" and the "4" stood for <stackoverflow.com/q/47503109>.
 2. M4 Wikipedia entry <wikipedia.org/wiki/M4_(computer_language)>, "m4 (computer language)", 2, retrieved April 28, 2020: "It is an extension of an earlier macro processor m3, written by Ritchie for an unknown AP-3 minicomputer."

3. I modified all quotations appearing in the body of this article to spell and capitalize program names, author names, and so on, in a consistent manner. Quotations appearing in the footnotes preserve the exact wording of sources.
4. M4 Wikipedia entry <[wikipedia.org/wiki/M4_\(computer_language\)#History](https://wikipedia.org/wiki/M4_(computer_language)#History)>, § “History”, ¶ 2, retrieved April 28, 2020: “In 1960s, an early general-purpose macro processor M⁶ was in use at AT&T Bell Laboratories, which was developed by Douglas McIlroy, Robert Morris and Andrew Hall.”
5. I’m not certain if M6 should be typeset as “M6” or “M⁶”. TheGNU M4 manual uses “M6” <gnu.org/savannah-checkouts/gnu/m4/manual/m4-1.4.18/m4.html#History>. The Wikipedia M4 entry uses “M⁶”. I suspect the “M⁶” style was used in Macro Processors, revised 2E, by A. J. Cole, M4 Wikipedia entry cites this as a source. I have not been able to get my hands on Macro Processors, revised 2E, from 1981 to verify this. I have seen an earlier Macro Processors edition but it did not mention M6. Furthermore, the original M6 technical report, written in 1971, does not use a superscripted six <plan9.io/cm/cs/cstr/2.pdf>. The entire document is typeset in monospace. M6 was released with V2 UNIX, and the V2 UNIX manual is also typeset in monospace. This indicates that the document publishing software at the time may have lacked the ability to typeset a superscripted six rather than this being a deliberate choice <web.archive.org/web/20161006034736/http://sunsite.icm.edu.pl/pub/unix/UnixArchive/PDP-11/Distributions/research/1972_stuff/unix_2nd_edition_manual.pdf>.
6. Claiming M6 was present at Bell Labs in the “1960s” is vague and possibly too early. Looking at just the publication dates of various documents: (1) The M6 Macro Processor was published in 1971—the cover page, which is typeset in a noticeably cleaner monospaced font than rest of the document, states 1972, but each page within the document is dated “June, 1971”; (2) Christopher Strachey’s GPM, which Hall states as one of M6’s influences, was published in 1965; and (3) Calvin Mooers’s TRAC, another M6 stated influence, was also published in 1965. There is a possibility that TRAC, GPM, and other influences, were known at Bell Labs before they were officially published. Furthermore, V1 UNIX was released November 3, 1971—and it did not include M6, but V2 UNIX, which was released June 12, 1972, did include M6. Based on this evidence, M6 may not have existed at Bell Labs until the mid-to-late 1960s.
7. Kernighan, B. W., Ritchie, D. M., The M4 Macro Processor. Bell Laboratories Technical Report, Murray Hill, New Jersey, USA, 1977 <wolfram.schneider.org/bsd/7thEdManVol2/m4/m4.pdf>. Last line on the cover page: “This paper is a user’s manual for M4”
8. Kernighan, B. W., Ritchie, D. M., The M4 Macro Processor. Bell Laboratories Technical Report, Murray Hill, New Jersey, USA, 1977 <wolfram.schneider.org/bsd/7thEdManVol2/m4/m4.pdf>. pp. 1, ¶ 2: “The M4 macro processor is an extension of a macro processor called M3 which was written by D. M. Ritchie for the AP-3 minicomputer; M3 was in turn based on a macro processor implemented for [1]” and from pp. 5, § “References”, “[1]” refers to Software Tools.
9. M4 Wikipedia entry <[wikipedia.org/wiki/M4_\(computer_language\)#cite_note-KR1977-1](https://wikipedia.org/wiki/M4_(computer_language)#cite_note-KR1977-1)>, § “Refer-

- ences”, retrieved April 28, 2020: “Brian W. Kernighan and Dennis M. Ritchie. The m4 macro processor. Technical report, Bell Laboratories, Murray Hill, New Jersey, USA, 1977.”
10. At least one attempt has been made to find out more information about this AP-3 minicomputer, but nothing seems to have come up: <minnie.tuhs.org/pipermail/tuhs/2019-December/019661.html> and <[wikipedia.org/wiki/Talk:M4_\(computer_language\)#AP-3_reference_as_”unknown”_mini](http://wikipedia.org/wiki/Talk:M4_(computer_language)#AP-3_reference_as_”unknown”_mini)>, retrieved May 20, 2020.
 11. I ordered Software Tools to do this research but you don’t have to. OpenLibrary carries it and many other old or out-of-print books <openlibrary.org/books/OL4934660M/Software_tools>. This truly seems like an immense resource for computing history that I wish I had known about before beginning this project.
 12. Technically, `macro` is the name of the main routine but B. W. Kernighan and P. J. Plauger use `macro` to refer to the program in its entirety.
 13. Kernighan, B. W., Plauger, P. J., Software Tools. Addison-Wesley Publishing Company, 1976. ISBN 0-201-03669-X. Chapter 8, pp. 283, ¶ 4.
 14. Ritchie, D. M., The Development of the C Language. Bell Labs/Lucent Technologies, Murray Hill, New Jersey, USA, 2003 <bell-labs.com/usr/dmr/www/chist.html>: “After creating the type system, the associated syntax, and the compiler for the new language, I felt that it deserved a new name; NB seemed insufficiently distinctive. I decided to follow the single-letter style and called it C, leaving open the question whether the name represented a progression through the alphabet or through the letters in BCPL.”
 15. Ritchie, D. M., The Development of the C Language. Bell Labs/Lucent Technologies, Murray Hill, New Jersey, USA, 2003 <bell-labs.com/usr/dmr/www/chist.html>: “I omit full descriptions of C itself, its parent B [Johnson 73] and its grandparent BCPL [Richards 79], and instead concentrate on characteristic elements of each language and how they evolved.”