

Speculating on the naming of M6

Chris Thornton

bitsbelow.us

published 2021-06-23

Given M6 is the *first* macro processor¹ in the series of Bell Laboratories macro processors leading up to M4, why does the series start with a “6”?

The M4 macro processor series carries an illusory naming convention. This apparent naming convention seemingly stems from M4 having the name “M4” and its immediate predecessor having the name “M3”, leading some² to incorrectly deduce that: (1) M4 is the fourth macro processor in this series; (2) M3 is the third macro processor in this series; and, in some cases, (3) that there exists an M2, M1 or M. These deductions are logical when reasoning from the apparent naming convention but they are falsehoods. The truth is M4 is not the fourth macro processor in this series and there are no M4 descendants named M2 or M1 or M.³ The first macro processor in this series is M6. There exists two macro processors in this series between the release of M6 and M3. M3 is the fourth macro processor in this series, and M4 is the fifth macro processor in this series.

The history of macro processors leading up to M4 is obscure and, sometimes, complicated. M6 is built upon prototypes developed at Bell Laboratories until its creation in c. 1969–1971.⁴ M6 itself, and at least one of the prototypes, was influenced by C. S. Strachey’s general-purpose macrogenerator (GPM).⁵ M6 was released as part of V2 UNIX.⁶ None of the prior prototypes were released to the public. Hence, M6 is the first macro processor in the series of macro processors leading up to M4. However, the available documentation does not indicate why the name “M6” was chosen and leaves room for speculation.

If we assume M6 took a conventional project naming approach then the “6” in M6 reflects *something* about the authors of the software or the software itself. To determine what that *something* is we should be able to read the release notes, the manual, or other documentation and take note of what the authors state as important or unique about M6. And indeed, the first paragraph of *The M6 Macro Processor* touts portability as a defining feature.⁷ Based on the writing structure of the documentation we can infer that A. D. Hall believed portability was such an important feature that he placed it next to the same sentences defining what a macro processor is.

I’m speculating here, and as I stated previously, the available evidence does not explicitly state why M6 is named “M6”. But I think it’s likely that the “6” in M6 refers to the six supported platforms when

M6 was released: the GE-635, IBM 360/65, CDC 6600, Univac 1108, PDP-10, and SIGMA 7.

Support for this claim can be found in the The Computing Science Technical Report #2 (CSTR2), The M6 Macro Processor, where A. D. Hall writes:⁷

*[M6] described here is a translation of an earlier experimental version written in MAD by [R. H. Morris] and has been written in FORTRAN IV in a way that is intended to be highly portable.**

and the asterisked footnote elaborates:⁸

** M6 has been compiled and executed on the GE-635, IBM 360/65, CDC 6600, Univac 1108, PDP-10 and SIGMA 7 without any changes in the source code except for input-output unit numbers⁹].*

This is a reasonable interpretation of what the “6” in M6 may refer to based on the available evidence. Keep in mind, however, this could be coincidental, and is not necessarily fact.

cmth@

-
1. When I state that M6 is the “the first macro processor” in this series, I specifically mean that M6 is the first macro processor in the series of macro processors leading up to M4 that was publicly released. The boundaries here are fuzzy. M6 is technically not *the first* Bell Laboratories macro processor if you include unreleased prototypes. Computing Science Technical Report #2, The M6 Macro Processor, notes that M6 is based on an experimental macro processor written in MAD by R. H. Morris, which is itself based on prototypes developed by M. D. McIlroy and R. H. Morris. Distinguishing which macro processor counts as “the first Bell Labs macro processor”, or “the first of the M4 series” is not simple. Several macro processors were written at Bell Labs before M6. What distinguishes M6 from the rest is that it was shipped as a part of V2 UNIX. The other macro processors were not released and seem to have existed only internally at Bell Labs. For this reason, I make a principled distinction to determine which macro processor is *the first* and then which subsequent macro processors count as part of the series. The principle for what counts as *the first* macro processor in the series is the first macro processor that is publicly released. The principle for what counts as a subsequent macro processor is any program design that is influenced by either the first publicly released macro processor, or by a subsequent macro processor.
 2. At least one other person, besides myself, thought the “4” in M4 was for being the fourth in a

series of macro processors for UNIX: <stackoverflow.com/questions/47503109>. But I believe this misconception is much more widespread.

3. I've written about M4, some of its history, and what the "4" means here <bitsbelow.us/p/the-m-and-the-4-in-m4>. In short, M4 is the fifth macro processor in the series of macro processors for UNIX. Going from M4 and back to M6, we have: M4, M3, macro, macro', then M6. For more information see *The M4 Macro Processor*, the book *Software Tools*, and *The M6 Macro Processor*.
4. Determining the exact year M6 was written is difficult. I've been scouring the The Unix Heritage Society's website <tuhs.org> and mailing list <minnie.tuhs.org/pipermail/tuhs/>, the UNIX manuals, and searching the web trying to uncover any facts that will narrow down when M6 was written. M6 was included in V2 UNIX, which was released June 12, 1972. The M6 Macro Processor, wierdly, dates each page with the year 1971 but the cover page says 1972. Its likely the document was first written in 1971 then updated in 1972 when released as a part of V2 UNIX. However, the V3 UNIX manual page for M6, under the "See Also" section, places "1969" inline with The M6 Macro Processor, indicating it was originally published in 1969. Interestingly, the V2 UNIX manual page for M6 also has a "See Also" section, and it references The M6 Macro Processor, however, the V2 UNIX manual does not include this date. I am still attempting to triangulate a more precise date based on other facts and events.
5. Computing Science Technical Report #2 (CSTR2), The M6 Macro Processor pp. 1, ¶ 1, § 1:
"M6 is a macro processor designed by M. D. McIlroy and [R. H. Morris] of Bell Telephone Laboratories and combined ideas from many sources [1,2,3,4]."
and page 13, source #1, is:
"1. [C. S. Strachey], A General Purpose Macrogenerator, *Comput. J.* 8, 3 (Oct. 1965) pp. 225-241."
continuing with The M6 Macro Processor pp. 1, ¶ 1, § 2 states:
"The version described here is a translation of an earlier experimental version written in MAD by [R. H. Morris] [...]"
Thus, M6 is a rewrite of a macro processor that was written in MAD by R. H. Morris, which combined ideas from many sources, including C. S. Strachey's GPM.
6. See V2 UNIX manual <web.archive.org/web/20161006034736/http://sunsite.icm.edu.pl/pub/unix/UnixArchive/PDP-11/Distributions/research/1972_stuff/unix_2nd_edition_manual.pdf> pp. vi, xvi, and the M6 man page within the V2 UNIX manual.
7. Computing Science Technical Report #2 (CSTR2), The M6 Macro Processor pp. 1, ¶ 1, § 2:
"The version described here is a translation of an earlier experimental version written in MAD by R. Morris and has been written in FORTRAN IV in a way that is intended to be highly portable.*"
8. Computing Science Technical Report #2 (CSTR2), The M6 Macro Processor pp. 1, ¶ 6.
9. *Input-output unit numbers* is FORTRAN jargon, however, the concept is similar to the UNIX mapping of `stdin`, `stdout`, and `stderr` to 0, 1, and 2 respectively.