119

How to Write a Long Formula

Leslie Lamport

December 25, 1993
Minor correction: January 18, 1994

Systems Research Center
130 Lytton Avenue
Palo Alto, California 94301

Systems Research Center

The charter of SRC is to advance both the state of knowledge and the state of the art in computer systems. From our establishment in 1984, we have performed basic and applied research to support Digital's business objectives. Our current work includes exploring distributed personal computing on multiple platforms, networking, programming technology, system modelling and management techniques, and selected applications.

Our strategy is to test the technical and practical value of our ideas by building hardware and software prototypes and using them as daily tools. Interesting systems are too complex to be evaluated solely in the abstract; extended use allows us to investigate their properties in depth. This experience is useful in the short term in refining our designs, and invaluable in the long term in advancing our knowledge. Most of the major advances in information systems have come through this strategy, including personal computing, distributed systems, and the Internet.

We also perform complementary work of a more mathematical flavor. Some of it is in established fields of theoretical computer science, such as the analysis of algorithms, computational geometry, and logics of programming. Other work explores new ground motivated by problems that arise in our systems research.

We have a strong commitment to communicating our results; exposing and testing our ideas in the research and development communities leads to improved understanding. Our research report series supplements publication in professional journals and conferences. We seek users for our prototype systems among those with whom we have common interests, and we encourage collaboration with university researchers.

Robert W. Taylor, Director

How to Write a Long Formula

Leslie Lamport

December 25, 1993
Minor correction: January 18, 1994

Author's Abstract

Standard mathematical notation works well for short formulas, but not for the longer ones often written by computer scientists. Notations are proposed to make one or two-page formulas easier to read and reason about.

Introduction

Mathematicians seldom write formulas longer than a dozen or so lines. Computer scientists often write much longer formulas. For example, an invariant of a concurrent algorithm can occupy more than a page, and the specification of a real system can be a formula dozens or even hundreds of pages long. Standard mathematical notation works well for short formulas, but not for long ones. I propose a few simple notations for writing formulas of up to a couple of pages. These notations can make formulas much easier to read and reason about.

Formulas significantly longer than two pages require hierarchical structuring. Methods for structuring long programs can be used to structure long formulas. Programs of less than a dozen or so pages can be adequately structured with procedures; longer programs require some method of grouping procedures into modules. The definition is the mathematical analog of the procedure. Definitions suffice for structuring formulas of up to about a dozen pages. For longer formulas, some form of module structure is also needed.

Any formula can be written with a hierarchy of definitions, each only a few lines long. However, just as programs become hard to read if broken into too many procedures, formulas are hard to read if broken into definitions that are too small. In my experience, the best way to structure a long formula is in terms of individual formulas of up to a page or two.

Writing Formulas

Consider the following definition, written with standard mathematical conventions. (The examples come from the invariant of an unpublished correctness proof for a cache coherence algorithm; the reader is not expected to understand them.)

$$
\begin{aligned}
memQLoc(a) \;\equiv\; & \begin{cases} \text{“None”} & \text{if } Locs = \emptyset \\ \max(Locs) & \text{otherwise} \end{cases} \\
\text{where } Locs \;\equiv\; & \{ i \in \{0 \ldots |memQ| - 1\} : \\
& \quad (memQ[i].req.type = \text{“Write”}) \\
& \qquad \wedge \; (memQ[i].req.adr = a) \quad \}
\end{aligned}
$$

This definition is easy to read because it is short. However, suppose that “None” and $\max(Locs)$ were replaced by much longer expressions. We would then see that the “where” construct is bad because it forces us to read the

entire definition of $memQLoc(a)$ before we learn what $Locs$ is. A structure that scales better to large formulas is

$$\begin{array}{l}\textbf{let}\ \ Locs \equiv \{i \in \{0 \ldots |memQ| - 1\} : \\ \qquad\qquad (memQ[i].req.type = \text{“Write”}) \\ \qquad\qquad\quad \wedge\ (memQ[i].req.adr = a) \quad \} \\ \textbf{in}\quad memQLoc(a) \equiv \begin{cases} \text{“None”} & \text{if } Locs = \emptyset \\ \max(Locs) & \text{otherwise} \end{cases}\end{array}$$

Suppose once again that “None” were replaced by a long expression e, perhaps crossing onto the next page. The typographic difficulties posed by the resulting large left brace are daunting. Simply removing the brace still leaves us with the problem of where to put the condition $Locs = \emptyset$. If it goes after e, we have to read several lines before discovering the structure of the definition. If it goes at the end of the first line, we read the $Locs = \emptyset$ in the middle of reading e. A better notation is the **if/then/else** construct used in programming languages.

$$\begin{array}{l}\textbf{let}\ \ Locs \equiv \{i \in \{0 \ldots |memQ| - 1\} : \\ \qquad\qquad (memQ[i].req.type = \text{“Write”}) \\ \qquad\qquad\quad \wedge\ (memQ[i].req.adr = a) \quad \} \\ \textbf{in}\quad memQLoc(a) \equiv \textbf{if}\ Locs = \emptyset\ \textbf{then}\ \text{“None”} \\ \qquad\qquad\qquad\qquad\qquad\quad \textbf{else}\quad \max(Locs)\end{array}$$

The **if/then/else** makes the structure immediately clear, even for long formulas. The obvious analog of the **case** construct of programming languages works for definitions with more than two alternatives. The customary closing **end** (or **fi**) is unnecessary, because we can use parentheses and indentation to delimit the scope of an **if** or **case**.

The original version of the definition had an important feature that has been lost in these transformations: we could see at once that it was a definition of $memQLoc(a)$. One further change recovers this feature.

$$\begin{array}{l}memQLoc(a) \equiv \textbf{let}\ \ Locs \equiv \{i \in \{0 \ldots |memQ| - 1\} : \\ \qquad\qquad\qquad\qquad\qquad (memQ[i].req.type = \text{“Write”}) \\ \qquad\qquad\qquad\qquad\qquad\quad \wedge\ (memQ[i].req.adr = a) \quad \} \\ \qquad\qquad\quad \textbf{in}\quad \textbf{if}\ Locs = \emptyset\ \textbf{then}\ \text{“None”} \\ \qquad\qquad\qquad\qquad\qquad\quad \textbf{else}\quad \max(Locs)\end{array}$$

The basic problem with the “if ... otherwise” construct is shared by all infix operators: we discover the high-level structure only after reading to

the end of the first argument. Consider the following formula.

$$
\begin{array}{l}
(\forall\, c \in \mathit{CacheAddress} : \\
\quad \mathit{cache}[p, c] \in ([\![\mathit{adr} : \mathit{Address}, \mathit{val} : \mathit{Value}]\!] \cup \{\text{“Invalid”}\})) \\
\quad \wedge\ ((\mathit{request}[p] \in \mathit{Request}) \\
\qquad \vee\ ((\mathit{request}[p] = \text{“Ready”}) \wedge (\mathit{state}[p] = \text{“Idle”}))) \\
\quad \wedge\ (\mathit{response}[p] \in \mathit{Value})
\end{array}
$$

We have to read to the end of the second line, and count parentheses, before learning that the formula is a conjunction. One possible solution is prefix notation, writing $\wedge(A,\, B,\, C)$ instead of $A \wedge B \wedge C$.

$$
\begin{array}{l}
\wedge\ (\forall\, c \in \mathit{CacheAddress} : \\
\qquad \mathit{cache}[p, c] \in ([\![\mathit{adr} : \mathit{Address}, \mathit{val} : \mathit{Value}]\!] \cup \{\text{“Invalid”}\}), \\
\quad \vee\ (\mathit{request}[p] \in \mathit{Request}, \\
\qquad \wedge\ (\mathit{request}[p] = \text{“Ready”}, \\
\qquad\quad \mathit{state}[p] = \text{“Idle”})), \\
\quad \mathit{response}[p] \in \mathit{Value})
\end{array}
$$

This formula is easy to read only because of the way it is indented. If one needs indentation anyway, why not use it to eliminate the parentheses and commas required by a prefix notation? We write the formula $A_1 \wedge A_2 \wedge \ldots \wedge A_n$ as the aligned list

$$
\begin{array}{l}
\wedge\ A_1 \\
\wedge\ A_2 \\
\quad \ldots \\
\wedge\ A_n
\end{array}
$$

and write disjunctions similarly. We can then use indentation to eliminate parentheses, writing the formula above as

$$
\begin{array}{l}
\wedge\ \forall\, c \in \mathit{CacheAddress} : \\
\qquad \mathit{cache}[p, c] \in ([\![\mathit{adr} : \mathit{Address}, \mathit{val} : \mathit{Value}]\!] \cup \{\text{“Invalid”}\}) \\
\wedge\ \vee\ \mathit{request}[p] \in \mathit{Request} \\
\quad \vee\ \wedge\ \mathit{request}[p] = \text{“Ready”} \\
\qquad \wedge\ \mathit{state}[p] = \text{“Idle”} \\
\wedge\ \mathit{response}[p] \in \mathit{Value}
\end{array}
$$

We continue to use $\wedge$ and $\vee$ as infix operators in subformulas. For example, the second conjunct of this formula can also be written

$$
\begin{array}{l}
\wedge\ \vee\ \mathit{request}[p] \in \mathit{Request} \\
\quad \vee\ (\mathit{request}[p] = \text{“Ready”}) \wedge (\mathit{state}[p] = \text{“Idle”})
\end{array}
$$

The list convention for conjunction and disjunction can be used for other associative operators, including addition and multiplication. However, it does not work for the nonassociative boolean operator $\Rightarrow$ (implies). I have not found a good general method of writing $A \Rightarrow B$ when A and B are long formulas. When A and B are conjunctions or disjunctions, the format

$$\begin{array}{l} \wedge\ A_1 \\ \quad \ldots \\ \wedge\ A_m \\ \Rightarrow \wedge\ B_1 \\ \qquad \ldots \\ \qquad \wedge\ B_n \end{array}$$

works fairly well if $A_1 \wedge \ldots \wedge A_m$ is only a few lines long.

Writing conjunctions and disjunctions as lists lets us take full advantage of indentation to eliminate parentheses. Indentation has meaning; shifting an expression to the left or right changes the way a formula is parsed. It is not hard to devise precise rules for parsing these two-dimensional formulas. However, there is some question about what formulas should be allowed. For example, should it be legal to write $(A_1 \vee A_2) \wedge B$ as follows?

$$\begin{array}{l} \quad \vee\ A_1 \\ \quad \vee\ A_2 \\ \wedge\ B \end{array}$$

Answers to these questions will evolve as people use the notation.

Numbering Parts of Formulas

We don't just write formulas, we also reason about them. Reasoning about a large formula requires a convenient way of referring to its components. With the list convention, we can name individual conjuncts and disjuncts by numbering them. The ith conjunct or disjunct of a formula named F is called $F.i$. A universally quantified formula can be viewed as a conjunction, where the yth conjunct of $\forall\, x : Q$ is $Q[y/x]$, the formula obtained by substituting y for x in Q. If F is the name of the formula $\forall\, x : Q$, then we take $F(y)$ to be the name of the formula $Q[y/x]$. A similar convention applies to existential quantification.

Figure 1 illustrates the use of these structuring and naming conventions in a real example—the definition of an invariant I for a cache coherence algorithm. For simplicity, only the outermost three levels of con-

$$
\begin{array}{lll}
I & \equiv & \textbf{let}\ cacheLocs(p,a) \equiv \{c \in CacheAddress : \wedge\ cache[p,c] \neq \text{“Invalid”} \\
& & \qquad\qquad\qquad\qquad\qquad\qquad \wedge\ cache[p,c].adr = a \qquad \} \\
& & \quad inCache(p,a) \equiv cacheLocs(p,a) \neq \emptyset \\
& & \quad memQLoc(a) \equiv \textbf{let}\ Locs \equiv \{i \in \{0 \ldots |memQ|-1\} : \\
& & \qquad\qquad\qquad\qquad\qquad \wedge\ memQ[i].req.type = \text{“Write”} \\
& & \qquad\qquad\qquad\qquad\qquad \wedge\ memQ[i].req.adr = a \qquad \} \\
& & \qquad\qquad \textbf{in}\ \ \textbf{if}\ Locs = \emptyset\ \textbf{then}\ \text{“None”} \\
& & \qquad\qquad\qquad\qquad\ \ \textbf{else}\ \ \max(Locs) \\
& & \quad memVal(a) \equiv \textbf{if}\ memQLoc(a) = \text{“None”} \\
& & \qquad\qquad\qquad\ \ \textbf{then}\ mainMemory[a] \\
& & \qquad\qquad\qquad\ \ \textbf{else}\ \ memQ[memQLoc(a)].req.val \\
& & \textbf{in}\ \ 1.\wedge\ \forall p \in Process : \\
& & \qquad 1.\wedge\ \forall a \in Address : \\
& & \qquad\qquad 1.\wedge\ \#cacheLocs(p,a) \leq 1 \\
& & \qquad\qquad 2.\wedge\ inCache(p,a) \Rightarrow (cacheVal(p,a) = memVal(a)) \\
& & \qquad\qquad 3.\wedge\ mainMemory[a] \in Value \\
& & \qquad 2.\wedge\ \forall c \in CacheAddress : \\
& & \qquad\qquad cache[p,c] \in ([\![adr : Address,\ val : Value]\!] \cup \{\text{“Invalid”}\}) \\
& & \qquad 3.\wedge\ \text{a.}\vee\ request[p] \in Request \\
& & \qquad\qquad \text{b.}\vee\ \wedge\ request[p] = \text{“Ready”} \\
& & \qquad\qquad\qquad\ \wedge\ state[p] = \text{“Idle”} \\
& & \qquad 4.\wedge\ response[p] \in Value \\
& & \qquad 5.\wedge\ 1.\wedge\ state[p] \in \{\text{“RdCache”}, \text{“MemWait”}, \\
& & \qquad\qquad\qquad\qquad\qquad \text{“BusWait”}, \text{“WrDone”}, \text{“Idle”}\} \\
& & \qquad\qquad\ 2.\wedge\ (state[p] = \text{“RdCache”}) \Rightarrow \wedge\ request[p].type = \text{“Read”} \\
& & \qquad\qquad\qquad\qquad\qquad\qquad\qquad \wedge\ inCache(p, request[p].adr) \\
& & \qquad\qquad\ 3.\wedge\ (state[p] = \text{“MemWait”}) \\
& & \qquad\qquad\qquad \Rightarrow \wedge\ \neg inCache(p, request[p].adr) \\
& & \qquad\qquad\qquad\qquad \wedge\ \#\{i \in \{0 \ldots |memQ|-1\} : \\
& & \qquad\qquad\qquad\qquad\qquad\qquad \wedge\ p = memQ[i].proc \\
& & \qquad\qquad\qquad\qquad\qquad\qquad \wedge\ memQ[i].req.type = \text{“Read”}\} = 1 \\
& & \qquad\qquad\ 4.\wedge\ (state[p] = \text{“BusWait”}) \wedge (request[p].type = \text{“Read”}) \\
& & \qquad\qquad\qquad\ \Rightarrow \neg inCache(p, request[p].adr) \\
& & \qquad\qquad\ 5.\wedge\ (state[p] = \text{“WrDone”}) \Rightarrow (request[p].type = \text{“Write”}) \\
& & \quad 2.\wedge\ memQ \in SequenceOf([\![proc : Process,\ req : Request]\!]) \\
& & \quad 3.\wedge\ \forall i \in \{0 \ldots |memQ|-1\} : \\
& & \qquad memQ[i].req.type = \text{“Read”} \\
& & \qquad\ \Rightarrow 1.\wedge\ state[memQ[i].proc] = \text{“MemWait”} \\
& & \qquad\qquad\ 2.\wedge\ request[memQ[i].proc] = memQ[i].req
\end{array}
$$

Figure 1: An invariant of a cache coherence algorithm.

juncts and disjuncts are labeled. (I like to label conjuncts with numbers and disjuncts with letters.) The naming convention implies that $I.2$ is the formula $memQ \in SequenceOf(\ldots)$, and $I.1(q).3.a$ is the formula $request[q] \in Request$.

Conclusion

The notations introduced here will be unfamiliar to most readers, and unfamiliar notation usually seems unnatural. I have used the notations for several years, and I now find them indispensable. I urge the reader to rewrite formula I of Figure 1 in conventional notation and compare it with the original. Having to keep track of six or seven levels of parentheses reveals the advantage of using indentation to eliminate parentheses.