

MakeIndex: An Index Processor For L^AT_EX

Leslie Lamport

17 February 1987

1 How to Use *MakeIndex*

MakeIndex is a program for making an index in a document generated with L^AT_EX. The first step in producing the index is to put the necessary `\index` commands in your document, as described in the next section. Here, I describe how to generate the index after the `\index` commands are in place.

Let's suppose that the root file of your document is `myfile.tex`. You must make the following changes to your document:

- Add the `makeidx` document-style option to the list of options in the `\documentstyle` command. (See page 21 of the L^AT_EX manual.)
- Put a `\makeindex` command in the preamble (between the `\documentstyle` and `\begin{document}` commands).
- Put a `\printindex` command where you want the index to appear—usually at the end, right before the `\end{document}` command.

You then run L^AT_EX on your entire document, causing it to generate the file `myfile.idx`, which I will call the `idx` file. Next, run the *MakeIndex* program by typing the following Unix command:

```
makeindex myfile.idx
```

This produces the file `myfile.ind`, which I will call the `ind` file. If *MakeIndex* generated no error messages, you can now rerun L^AT_EX on your document and the index will appear. (You can remove the `\makeindex` command first so the `idx` file is not regenerated.) If there were error messages, see Section 3.

By reading the index, you may discover additional mistakes. These should be corrected by changing the appropriate `\index` commands in the document and regenerating the `ind` file. If there are problems that cannot

be corrected in this way, you can always edit the `ind` file directly. However, such editing is to be avoided because it must be repeated every time you generate a new version of the index.

2 How to Generate Index Entries

2.1 When, Why, What, and How to Index

It's tempting to generate the index as you write the document. Resist the temptation. It is virtually impossible to obtain any consistency in an index that is generated in this way.

An index is there to help the reader find what he's looking for. With this in mind, common sense can help in figuring out what should be in the index and how it should be organized. Since it's often hard to distinguish common sense from equally common nonsense, professional advice is useful. Many style guides discuss indexing; the pamphlet *Indexing Your Book* by Sina Spiker (The University of Wisconsin Press, 1954) is, according to its subtitle, "A Practical Guide for Authors".

Unfortunately, these guides to indexing seem to have been written when high tech meant using a ball-point pen instead of a quill, so their advice on the mechanics of creating an index revolve around how to stack your 3×5 index cards. You'll have to figure out your own method of using the computer to lighten the chore. An alphabetical list of every word in your document, with duplicates removed, is not a bad place to start. If your system has Howard Trickey's `delatex` program, then the following Unix command generates such a list from the file `myfile.tex` and puts it on the file `foo`:

```
delatex myfile.tex | sort -uf > foo
```

The computer is only a tool; it can't write the index for you. It may be easy to choose which words are important and mechanically generate an index citing every occurrence of those words, but the resulting index will not be as useful to the reader as one prepared with more care.

2.2 The Basics

Each `\index` command causes L^AT_EX to write an entry on the `idx` file. The following example shows some simple `\index` commands and the index en-

tries that they produce. The page number refers to the page containing the text where the `\index` command appears.¹

Page ii	<code>\index{Alpha}</code>	Alpha, ii
Page viii:	<code>\index{alpha}</code>	alpha, viii, ix, 22
Page ix:	<code>\index{alpha}</code>	alpha bet, 24
	<code>\index{Alphabet}</code>	Alphabet, ix
Page 22:	<code>\index{alpha}</code>	alphabet, 23
	<code>\index{alphas}</code>	alphas, 22
Page 23:	<code>\index{alphabet}</code>	
	<code>\index{alphabet}</code>	
Page 24:	<code>\index{alpha bet}</code>	

Note that the duplicate `\index{alphabet}` commands on page 23 produce only one “23” in the index.

To produce a subentry, the argument of the `\index` command should contain both the main entry and the subentry, separated by a `!` character.

Page 7:	<code>\index{gnat!size of}</code>	gnat, 32
Page 32:	<code>\index{gnat}</code>	anatomy, 35
Page 35:	<code>\index{gnat!anatomy}</code>	size of, 7
	<code>\index{gnus!good}</code>	gnus
Page 38:	<code>\index{gnus!bad}</code>	bad, 38
		good, 35

You can also have subsubentries.

Page 8:	<code>\index{bites!animal!gnats}</code>	bites
Page 10:	<code>\index{bites!animal!gnus}</code>	animal
Page 12 :	<code>\index{bites!vegetable}</code>	gnats, 8
		gnus, 10
		vegetable, 12

\LaTeX and *MakeIndex* support only three levels of indexing; you can’t have subsubsubentries.

To specify a page range, put an `\index{...|}` command at the beginning of the range and an `\index{...|)}` command at the end of the range.

¹To avoid any ambiguity if a page break occurs right at an `\index` command, it’s a good idea to attach the command to a word instead of putting it between words.

Page vi:	<code>\index{gnat {}</code>	gnat, vi–x, 22
Page x:	<code>\index{gnat)}</code>	gnus
Page 22:	<code>\index{gnat}</code>	bad, 22
	<code>\index{gnus!bad {}</code>	good, 28–32
	<code>\index{gnus!bad)}</code>	
Page 28:	<code>\index{gnus!good {}</code>	
Page 30:	<code>\index{gnus!good}</code>	
Page 32:	<code>\index{gnus!good)}</code>	

Note that *MakeIndex* does the right thing when both ends of the range fall on the same page.²

Sometimes, you want to add a cross-reference with no page number. This is done as follows;

Page 2:	<code>\index{at}</code>	at, 2
Page 2:	<code>\index{at!bat see{bat, at}}</code>	bat, <i>see</i> bat, at

Since the “see” entry does not print any page number, it doesn’t matter where the `\index{...|see{...}}` command goes, so long as it follows the `\begin{document}` command. You might want to put all such cross-referencing commands in one place.

If you specify an entry of the form $\sigma@ \tau$, the string σ determines the alphabetical position of the entry, while the string τ produces the text of the entry.

Page 44:	<code>\index{twenty}</code>	twenty, 44
Page 46:	<code>\index{twenty-one}</code>	xx, 55
Page 55:	<code>\index{twenty@xx}</code>	twenty-one, 46

This feature is useful because the argument of the `\index` command provides the actual input string that L^AT_EX uses to generate the index entry. In the following example, the input α produces the symbol α ,

Page 12:	<code>\index{alphas}</code>	alpha, 13
Page 13:	<code>\index{alpha}</code>	α , 14
Page 14:	<code>\index{alpha@\$\alpha\$}</code>	alphas, 12

Similarly, the command `\index{gnu@{\bf gnu}}` produces a boldface **gnu** index entry.

In some indexes, certain page numbers are specially formatted—for example, an italic page number may indicate the primary reference, and an

²The use of “28ff.” rather than “28–32” is frowned upon by the experts.

n after a page number may denote that the item appears in a footnote on that page. *MakeIndex* makes it easy to format an individual page number any way you want. For any string of characters σ , the command `\index{...|\sigma}` produces a page number of the form $\sigma\{n\}$. Similarly, the command `\index{...|(\sigma}` may produce a page number of the form $\sigma\{n-m\}$. Thus, suppose the document contains the following command definitions:

```
\newcommand{\ii}[1]{\it #1}
\newcommand{\nn}[1]{#1n}
```

We could then have:

Page 3:	<code>\index{gnat ii}</code>	gnat, 3, 4n
Page 4:	<code>\index{gnat nn}</code>	gnu, 5, 44-46
Page 5:	<code>\index{gnu}</code>	
Page 44:	<code>\index{gnu (ii}</code>	
Page 46:	<code>\index{gnu)}</code>	

The “see” option is a special case of this facility, where the `\see` command is predefined by the `makeidx` document-style option.

2.3 The Fine Print

Commands in an index entry are expanded when the index is typeset, not when the `idx` file is written. Hence, the command `\index{\gnu}` produces an entry that is alphabetized by `\gnu`, regardless of how the `\gnu` command is defined.

Recall that special characters like `\` may appear in the argument of an `\index` command only if that command is not itself contained in the argument of another command. This is most likely to be a problem when indexing items in a footnote. Even in this case, robust commands can be placed in the “@” part of an entry, as in `\index{gnu@{\it gnu}}`, and fragile commands can be used if protected with the `\protect` command.³

Remember that the argument of an `\index` command must always have matching braces, where the brace in a `\{` or `\}` command counts.

MakeIndex assumes that all page numbers are either arabic or lowercase roman numerals; it assumes that pages numbered with roman numerals precede those numbered with arabic numerals.

³In versions of L^AT_EX released before 3 February 1987, you must use the `\string` command instead of `\protect`.

To put a !, @, or | character in an index entry, *quote* it by preceding the character with a ". More precisely, any a character is said to be quoted if it follows an unquoted " that is not part of a \ command. A quoted !, @, or | character is treated like an ordinary character rather than having its usual meaning. The " preceding a quoted character is deleted before the entries are alphabetized.

Page 2:	<code>\index{exclaim ("!)}</code>		exclaim (!), 2
Page 3:	<code>\index{exclaim ("!")!loudly}</code>		loudly, 3
Page 4:	<code>\index{fur@f\{"u}r}</code>		für, 4
Page 5:	<code>\index{quote (\verb+""+)}</code>		quote (""), 5

MakeIndex regards spaces as ordinary characters when alphabetizing the entries and deciding whether two entries are the same. Thus, letting “`␣`” denote a space character, the commands `\index{gnu}`, `\index{␣gnu}`, and `\index{gnu␣}` produce three separate entries, the first appearing near the beginning of the index, since `␣` comes before any letter in *MakeIndex*’s “alphabetical” order. All three entries look the same when printed, since L^AT_EX ignores extra spaces in the input. Similarly, `\index{a␣space}` and `\index{a␣␣space}` produce two different entries that look the same when printed. Do not split the argument of an `\index` command across lines in the input file.

3 Error Messages

MakeIndex types out on the terminal the number of lines read and written and how many errors were found. Message to identify the error are written on a file with extension `ilg`. There are two phases in which *MakeIndex* can produce error messages: when it is reading the `idx` file, and when it is writing the `ind` file. Each error message prints the nature of the error followed by a line number, identifying where in the file the error occurs. In the reading phase, the line number refers to the `idx` file; in the writing phase, it refers to the `ind` file.

3.1 Errors in Reading Phase

Extra ‘!’ at position ...

The `\index` command’s argument has more than two unquoted ! characters. Perhaps some of them should be quoted.

Extra ‘@’ at position ...

The `\index` command argument has two or more unquoted @ characters with no intervening !. Perhaps one of the @ characters should be quoted.

Extra ‘|’ at position ...

The `\index` command’s argument has more than one unquoted ! characters. Perhaps the extras should be quoted.

Illegal null field

The `\index` command argument doesn’t make sense because some string is null that shouldn’t be. The command `\index{!big}` will produce this error, since it specifies a subentry “big” with no entry. Similarly, the command `\index{@big}` is incorrect because it specifies a null string for alphabetizing.

Argument ... too long (max 1024).

Your document contained an `\index` command with a very long argument. You probably forgot the right brace that was supposed to delimit the argument.

Other errors

MakeIndex can produce a variety of other error messages indicating that something is seriously wrong with the `idx` file. If you get one, it probably means that the `idx` file was corrupted in some way. If \LaTeX did not generate any errors when it created the `idx` file, then it is highly unlikely to have produced a bad `idx` file. If it did, you’ll have to examine the `idx` file to figure out what went wrong.

3.2 Errors in Writing Phase

Unmatched range opening operator

An `\index{...|}` command was not followed by a matching `\index{...|}` command. The “...” in the two commands must be completely identical.

Unmatched range closing operator

An `\index{...|}` command was not preceded by a matching `\index{...|}` command.

Extra range opening operator

Two `\index{...|}` commands appear in the document with no intervening `\index{...|}` command.

Inconsistent page encapsulator ... within range

MakeIndex has been instructed to include a page range for an entry and a single page number within that range that is formatted differently—for example, by having a `\index{gnu|ii}` command between a `\index{gnu|}` and a `\index{gnu|}` command.

Conflicting entries

MakeIndex thinks it has been instructed to print the same page number twice in two different ways—for example, by the commands `\index{gnu}` and `\index{gnu|see{...}}` appearing on the same page.

Acknowledgements

MakeIndex is a C program written by Pehong Chen, with a little advice from me. It was inspired by a program written by Mike Urban, which was based on a program written by Marshall Rose. Perhaps someday someone will rewrite *MakeIndex* in Web, making it available to all \LaTeX users.