

MSORT Reference Manual

William J. Poser
billposer@alum.mit.edu

Contents

1. Introduction
2. What *msort* Does
 1. Parsing Into Records
 2. Extracting Keys
 1. Parsing Records into Fields
 2. Specifying Key Fields
 3. Reversing Keys
 4. Multiple Keys
 5. Optional Keys
 3. Sorting
 1. Numeric Comparisons
 2. Lexicographic Comparisons
 3. Date comparisons
 4. Time comparisons
 5. Size comparisons
 4. Writing Out the Sorted Records
3. Exclusions
4. Overview of the Sorting Process
5. Backslash Escapes
6. The Command Line
7. Examples
8. Logging
9. Speed
10. Character Set
11. Limits
12. Defaults
13. Summary of Command Line Options
14. Exit Status

1. Introduction

msort is a program for sorting text files in sophisticated ways, intended especially for linguistic databases. It allows you to specify arbitrary sort orders, to sort blocks of text delimited in a number of ways rather than just lines, and to specify particular fields of a record as sort keys. *msort* is capable of sorting on several keys, so that when two records tie on one key, the tie may be broken on another. Any or all keys

may be optional. It can sort lexicographically, on numerical values, on dates and times, and on key sizes.

msort is available in binary form for GNU/Linux. It is written in standard C and can be compiled and run on virtually any variant of UNIX. It should work on any POSIX-compliant system, but has not yet been tested on non-UNIX systems.

2. What *msort* does

The use of *msort* is best understood in terms of the operations it performs. It performs four tasks:

- a. parsing the input into records;
- b. extracting the sort keys from each record;
- c. sorting the records;
- d. writing out the sorted records;

2.1. Parsing into Records

A record is the basic unit on which *msort* operates. When it reads its input, it breaks it up into a sequence of records. A record may consist of a single line of text, as in typical sorting programs. To specify that a record is a single line, give the command line option *-l*.

However, for many applications it is desirable for a record to consist of several lines. For example, a lexical database might contain records consisting of a pronunciation field, a gloss, and a category, such as this:

```
P:ditnikwun
G:dandelion
C:N
```

msort therefore recognizes a second type of record, one consisting of a block of text.

Sometimes records are separated by a special separator character. (Records consisting of a single line are really just the special case of this where the separator is a newline character.) *msort* recognizes such records if the separator character is specified by the *-r* option. Backslash escapes are recognized and may be used to supply values otherwise difficult to enter on the command line.

msort also recognizes blocks terminated by one or more blank lines. *msort* treats such blocks as records by default. This may be explicitly specified by giving the command line option *-b*.

2.2. Extracting Keys

The information used for sorting is known as the sort key. For example, if we sort a mailing list by postal code, ignoring the name and address, the postal code is the sort key. By default *msort* uses the entire record as the sort key. (This default may be explicitly specified by giving the command line option *-w*.)

2.2.1. Parsing Records into Fields

It is also possible to use a selected field as the sort key. A field is a portion of a record beginning at the beginning of a record or at the end of the previous field, and ending when a terminator character is encountered. For example, if the record consists of blocks of text and the field terminator is the newline character, then each line of the record will be a field. If a record consists of a single line of text and the field terminators are space and tab (collectively known as “whitespace”), then each “word” of the line will constitute a field.

msort sets default values for the field terminators which depend on whether the record is a single line or a block of text. If the record is a single line, the field terminators are space and tab. If the record is a block of text, the field terminator is a newline. These defaults may be over-ridden by using the command line option *-d*. Backslash escapes are recognized and may be used to supply values otherwise difficult to enter on the command line.

2.2.2. Specifying Key Fields

msort provides two ways of specifying which field is the sort key. One is by the order of fields. You may specify that the first field is the key, or the second, the third, etc. by using the command line option *-n*. Field *1* is the first field, field *2* the second field, etc. Field numbers may be negative, in which case they are interpreted relative to the end of the record. That is, field *-1* is the last field, field *-2* the next-to-last, etc. A field number of *0* is meaningless and is treated as an error.

The other way of specifying the field to use as a sort key is by its tag, the characters with which the field begins. For example, if we specify the sort key as the field with tag *P:*, the field beginning with the characters *P:* will be used as the sort key. Only the characters following the tag form part of the key; the tag is used only to identify the key field. Therefore, in the example above, if we give the command line option *-t P:* to specify the key as the field with tag *P:*, the field *P:ditnikwun* will be the key field, and the actual key will be *ditnikwun*.

Although tags that are simply strings are sufficient most of the time, *msort*’s tag specifications are actually regular expressions. This means that they may contain wildcards, character sets, and disjunctions. The syntax used is the POSIX standard “extended regular expression” syntax, familiar to users of the UNIX *egrep* command. If your tags are plain strings you need not concern yourself with the more general regular expression facility except in the case in which characters in your strings have special meanings. The characters with special meaning are: `^ $. * [] () + ? -` and `\`. The character “`-`” has special meaning only when enclosed by square brackets. In general, you may prevent these characters from taking on their special meaning by quoting them with a preceding `\`. To specify a literal backslash, use two backslashes: `\\`.

Regular expressions as tags have several uses. First, they can be used to absorb characters that are properly speaking part of the beginning of the key but that we wish to ignore in sorting. For example, the regular expression *P:-?* has the effect of sorting on the field with tag *P:* and of ignoring a hyphen at the beginning of the key.

Second, character sets can be used when, for example, the database is inconsistent in capitalization. For example, the tag specification `[pP]:` has the effect of sorting on the field with tag `P:` or the tag `p:`, whichever comes first within the record if both are present.

Third, some databases using tagged fields separate the field name from the data in the field by a separator consisting of an arbitrary amount of whitespace (spaces or tabs). This whitespace can be absorbed into the tag by using appropriate wildcards. For example, to specify a key of `P:` followed by any amount of whitespace, use the tag specification `P:[]*` where the brackets contain a space and a tab.

Here are some examples of more complex regular expressions. To select a tag of either `p:` or `P:` followed by any amount of whitespace, use the tag specification `[pP]:[]*` where as above the second pair of brackets contain a space and a tab.

To select a tag of `P:` or `p:` followed by any amount of whitespace and to ignore leading hyphens, use the tag specification `[pP]:[]*-?` where as above the second pair of brackets contain a space and a tab.

Occasionally, you may wish to sort on one of two quite different tags. For example each entry in a lexical database might have either a gloss, tagged `g:`, or a definition, tagged `df:`. To sort on whichever one is present, specify the tag as `df:g:`.

Note that what you consider the tag need not be what was intended to be the field label when the database was created. For example, suppose that you use the label `S:` to indicate the source of the material in an entry, and that some entries come from a written source abbreviated `M`, which is followed by the section or page number, e.g. `M82`. You might extract the entries obtained from this source and then sort them into their original order (perhaps in order to check them against the written source) by specifying the tag as `S:M` and performing a numerical sort. Treating the `M` as part of the field label has the effect of stripping it from the key and so preventing it from interfering with the numerical sort.

It is possible for a record not to contain a key field, in the case of numerical specification because there are not enough fields, in the case of specification by tag because there is no field with the specified tag. In this case, if the key is not optional, `msort` reports the number of the offending record, writes a copy of the offending record into the log, and aborts.

If, as is possible when keys are specified by tag, a record contains more than one key field, `msort` uses the first and ignores the remainder, without reporting an error.

Since `msort` has to start at the beginning of each record and search for key fields when key fields are specified by tags, sorting is faster if the key fields are near the beginning of the record. If you know that you are likely to want to sort on some fields and unlikely to want to sort on others, you can improve the speed of the sort by putting the fields on which you are likely to sort at the beginning of the record. With current computers this is not very important.

2.2.3. Reversing Keys

The command line option `-R` causes the characters of the associated key to be reversed if it is lexicographic. (If a numeric, time, or date comparison is specified, this

option has no effect.) That is, the last character of the key is treated as the first, the next to last character as the second, etc. This is useful for the generation of reverse dictionaries. Key reversal is performed after the interpretation of multigraphs.

2.2.4. Multiple Keys

When sorting on multiple keys, you specify the keys to sort on in the order in which they are to be used. For each key field, you may also specify the nature of the comparison to be used, that is, whether textual or numeric. If it is textual, you may specify the sort order to use. The specification of the field to use as the key and the nature of the comparison may be separated by other options, but each comparison specification is paired with the nearest preceding key field specification. Options that are not key-specific may appear anywhere. For example, in order to sort first on the *P:* field using the sort order in the file *c2.ord* and then on the *C:* field using the sort order in the file *cats.ord*, you would give a command beginning like this:

```
msort -t "P:" -s c2.ord -t "C:" -s cats.ord
```

If you give a non-key-specific option, such as *-l*, which tells *msort* to treat lines as records, it may appear in any position.

For example, the following is an acceptable command line:

```
msort -t "P:" -l -s c2.ord -t "C:" -s cats.ord
```

msort can still determine that *c2.ord* contains the sort order to be used on the key obtained from the field with tag "P:" because *-t "P:"* is the nearest preceding key field selector.

To sort on the section number following the letter *M* in the field labelled *S:* and then on the pronunciation in the *P:* field using the sort order in the file *c2.ord*, you would give a command beginning like this:

```
msort -t "S:M" -n -t "P:" -s c2.ord
```

To sort using ASCII order on the second field and then in decreasing numerical order on the first field, you would give a command beginning like this:

```
msort -n 2 -n 1 -i
```

Here the *-i* option reverses numerical comparisons and inverts the sort order for lexicographic comparisons. Note the omission of a comparison specification after the first field specification. This causes *msort* to use the default of textual comparison using the ASCII order, assuming that the machine collating sequence is ASCII.

2.2.5. Optional Keys

Any or all keys may be specified as optional except when sorting on the entire record. A key is declared optional by giving the command line option *-o*. The argument to this option determines how records with missing values for this key are handled.

There are three possibilities. In some cases, it is desirable to have all records that lack a key sort before or after all records that possess it. If, for example, we are editing a database and wish to add the information in question to records that lack

it, we may wish to arrange for them to come first. On the other hand, if records that lack the information in question are of no immediate interest, we may wish to move them to the end. The third possibility is that we wish the missing key to have no effect at all, which will result if we arrange for missing keys to compare as equal to whatever they are compared to. Thus, for each optional key we specify whether missing keys are to compare as less than (that is, sort ahead of), greater than (that is, sort after), or equal to records that possess the key. The possible arguments to this option are therefore:

less than	l
	L
	<
greater than	g
	G
	>
equal to	e
	E
	=

Only the first character of the argument is checked, so if you wish you may spell these out or use familiar abbreviations, e.g. *less than* or *lt*. If the argument given is not recognized, the default is equality.

The optional key facility is useful for purposes other than dealing with missing values. It can also be used to put an entire class of key values into a certain category. For example, my database for the Stuart/Trembleur Lake dialect of Carrier contains entries taken from Father Adrien-Gabriel Morice’s book *The Carrier Language*. In the source field, these are identified by the letter “M” followed by the section number. The source field may contain multiple entries, separated by slashes, e.g.:

```
%S:D/ROHA/YVPI/M454
```

This field indicates that the entry is to be found in the *Central Carrier Bilingual Dictionary* (“D”) and in Father Morice’s book in section 454, and that it has also been recorded from two Carrier speakers whose names are abbreviated “ROHA” and “YVPI”.

Someone working on another Carrier dialect wanted to use this lexical database, which is by far the largest for any Carrier dialect, to suggest words to look for in her own dialect. She was particularly interested in old words such as those to be found in Father Morice’s book. I therefore wanted to sort the lexicon so that words found in Father Morice’s work would be listed first. I used the following command to do this:

```
msort -t "%S:(M[0-9]+|.*M[0-9]+).*" -o g -t "%P:"
```

The first tag specification is a regular expression that selects those source fields in which either the letter M plus one or more digits immediately follow either the source tag or a slash. The effect is to pick out all of the source fields which contain a reference to Father Morice’s book. By making this field optional and specifying that if it is not present the record follows one in which it is present, we arrange for

all of the records containing an entry found in Father Morice’s book to precede the others. Since the tag exhausts the field, the source field has no other effect on the sort. The second key is the Carrier headword, which results in the records being ordered alphabetically within the two major categories.

2.3. Sorting

Once *msort* has extracted the keys from the records, it performs the sort. Each time it compares two records, it compares the keys in order. If the first key is sufficient to order the two records, no other keys are used. If, however, two records tie on a key and there is a further key available, *msort* moves on to the next key.

The *-i* option causes the sense of comparisons to be inverted for a particular key.

The *-c* option is used to specify the type of comparison. Comparisons may be numeric, dates, times, string lengths, or lexicographic.

2.3.1. Numeric Comparisons

msort can sort on the numerical value of a field rather than on its textual value. The numerical value will be used if the argument of the *-c* option begins with *n* or *N*. If the argument to the *-c* option is *x* or *X*, the key field will be interpreted as a hexadecimal integer.

2.3.2. Lexicographic Comparisons

Lexicographic comparison is done if the argument of the *-c* option begins with *l* or *L*.

Lexicographic comparisons are based on a sort order that is normally read from a file specified by the *-s* command line option. If no sort order file is specified, each character is ranked in accordance with its numerical value. This means, for characters with codes ranging from 0 to 127, that the default order is that of the 7-bit ASCII character set, provided that the machine collating sequence be ASCII, as is normally the case.

If a sort order is specified but values are omitted, characters for which no rank is specified will sort in ASCII order compared to each other but after all characters, including multigraphs, for which a rank is specified.

The order of the lines in the sort order file determines the sort order. Characters appearing on an earlier line sort before those appearing on a subsequent line. For example, if the sort order file consists of the lines:

```
a
b
c
```

a will sort before *b* which will sort before *c*.

Characters appearing on the same line, separated by one or more tabs, have the same sort rank. For example, if we wish no distinction to be made, for purposes of sorting, between upper-and lower-case letters, we would put them on the same line:

```
a    A
```

b	B
c	C

etc.

Sequences of characters are often used to represent single sounds, and in such cases we often want to treat them as if they were single letters for purposes of sorting. For example, in languages in which the digraph *lh* is used to represent the voiceless lateral fricative, *lh* is often alphabetized as a single letter ordered after *l*, so that, for example, *lho* is ordered after *lo*, not before it as it would be if the *l* and *h* were treated as separate letters and the usual order of the Roman alphabet followed.

msort allows such multigraphs to be defined in the sort order file. Any characters not separated by tabs are treated as a single letter. For example, to add *lh* to the Roman alphabet as above the relevant portion of the sort order file would look like this:

k
l
lh
m

Like ordinary letters, multigraphs may be assigned the same sort rank. To continue the above example, if we wanted upper- and lower-case letters to have the same sort rank, the relevant portion of the sort order file would look like this:

K	k
L	l
Lh	lh
M	m

msort imposes no limit on the length of multigraphs. The limit on the number of multigraphs is set so high as to be unlikely to be of any practical importance.

Notice that you may wish to define as multigraphs for sorting purposes sequences of letters that are not intended to represent single sounds. For example, suppose that you wish to sort a lexical database by part of speech. You might define a sort order like this:

N	Noun		
V	Verb		
A	Adjective	Adj	ADJ
ADV	Adverb	Adv	
PP	Postposition		
INT	Interjection		

Notice how variants are placed on the same line so that they will not be distinguished for purposes of sorting, and how words are treated as multigraphs.

Similarly, the multigraph facility can be used to sort records according to a classification scheme having nothing to do with spelling or sound. For example, dictionary entries tagged for semantic field can be sorted by semantic field. For example, here is a fragment of the sort order for the topical index to a dictionary:

```
animals-land
animals-water
animals-domestic
animals-distant
animals-misc
```

Since every character in the sort order file counts, it is important to put into it exactly what you mean to put into it. Do not include blank lines or comments.

Unless your character set is very unusual, it is recommended that you include the space (blank) character explicitly in your sort order and that you rank it first. Doing this has two advantages. First, it makes sure that items intended to contain spaces will be sorted as they are usually desired to be. Second, it avoids problems with unintended trailing spaces. It is easy to add unintended, invisible space characters at the end of a line. If the space character is not ranked before the other characters, this can result in incorrect orderings whose cause is mysterious. For example, the word *an* should precede the word *and*. If, however, a space has inadvertently been added to the end of *an*, and you have not specified that space should precede all other characters, the default ranking will cause *an* to be ordered after *and*, and since the space is invisible, it will be difficult to detect the cause of the problem.

In addition to ordinary single-byte characters *msort* recognizes backslash escapes in the sort order file. These escapes may be used to specify characters not easily entered by the user's word-processor, or which are given a special interpretation in the sort order file. For example, the tab and space characters that serve to separate entries in the sort order file may be entered into the sort order by specifying them as `\011` and `\040` since the ASCII codes for tab and space are octal 011 (decimal 9) and octal 040 (decimal 32) respectively.

2.3.3. Comparisons of Dates

msort will treat the key as a date if the argument of the `-c` option begins with *d* or *D*. Dates consist of a numerical day specification (an integer ranging from 1 through 31), a numerical month specification (an integer ranging from 1 through 12), and a numerical year specification (an integer), with the three components separated by a designated separator. Positive years are interpreted as years C.E. Dates B.C.E. may be specified by making the year negative.

By default, the separators are slashes and the components are given in the order year-month-day. The *f* option may be used to set the date format used. The argument to this option should consist of the letters “y”, for year, “m”, for month, and “d” for day in the desired order, separated by the desired separators. For example, if the date *1999/1/3* is to be interpreted as January 3, 1999, the argument would be *y/m/d*. If it is to be interpreted as March 1, 1999, the argument would be *y/d/m*. If *12-3-1956* should be interpreted as December 3, 1956, the argument should be *m-d-y*. Although in practice the two separators will usually be the same,

msort allows them to be distinct. An argument of *m;d,y* would interpret a date of *12;3,1994* as December 3, 1994.

The default is the International Date Format, which is usual in Japan. It is equivalent to an argument of *y/m/d*. The usual American format requires the argument *m/d/y*. The format usual in most other western countries requires the argument *d/m/y*.

2.3.4. Comparisons of Times

msort will treat the key as a time if the argument of the *-c* option begins with *t* or *T*.

Times must be of the form *hours:minutes* or *hours:minutes.seconds*. Hours outside the range 0-24, and minutes and seconds outside the range 0-60 are detected and reported as errors. The *a.m./p.m* notation is not recognized. 24 hour “railroad time” is assumed.

2.3.5. Comparisons of Sizes

If the argument of the *-c* option begins with *s* or *S*, *msort* will sort on the size of the key. The size is the length of the key in characters, after preprocessing, including multigraph compression and deletion of excluded characters. Keys compared for size are treated exactly like lexicographic keys except that when comparisons are finally done on them, the comparison is a numeric comparison of the length of the key rather than a lexicographic comparison of the key as a string. If you want multigraphs to be treated as single characters for the purpose of size comparison (as you might if sorting by word length in phonemes, for example), provide a sort order definition for the key just as you would when sorting lexicographically. The sort order will be ignored, but the multigraph definitions will be used.

2.4. Writing Out the Sorted Records

Once *msort* has sorted the records using the keys it has extracted, it writes out the records in the sorted order. If the command line option *-I* is given, it will write them out in reverse order. The effect of this option is the same as that of the key-specific *-i* option except, of course, that it is global. Giving both the *-i* option and the *-I* option for a sort on a single key therefore has no effect; the two options cancel each other out.

On the other hand, using these two options is not the same as inverting the sort order, which can be done by using a sort order specification file. For example, if you put lines like:

z
y
x
w
...

into a sort order file, the effect is to invert the sort ranking of individual letters. For example, *z* will now sort before *y* instead of after. Reranking individual letters in this way does not affect the rule that nothing precedes something. Consider, for example, the following data, which are in standard alphabetical order:

```
bad
badger
bag
```

If we use the *-i* option:

```
msort -l -w -i
```

we obtain the following, in which the order is strictly reversed:

```
bag
badger
bad
```

Using the *-I* option would have the same effect since we have only one key, and using both will give us the original order. If, however, we put the inverted sort order beginning with *z* and ending with *a* in a file called *rev.ord* and make this the sort order for our single key thus:

```
msort -l -w -s rev.ord
```

we obtain the following output:

```
bag
bad
badger
```

The reason is that inverting the sort order causes *g* to sort before *d*, but *bad* still precedes *badger* because it is a prefix.

3. Exclusions

Sometimes it is desirable to exclude certain characters in the key field from consideration in sorting. For example, a leading hyphen, used to indicate that a morpheme must be preceded by a prefix, will generally be ignored so that bound and unbound forms will sort together. In order to accomodate this need, *msort* provides an exclusion facility. For each key, the user may specify a file, using the *-x* option, containing a list of characters to ignore when sorting. In order to facilitate the use of exclusions when doing simple, one-shot sorts, exclusions may also be specified directly on the command line, using the *-X* option.

Typical exclusion facilities do not meet the full range of needs, since we typically do not wish to ignore a character completely. To see why a simple exclusion facility, one that allows certain characters to be ignored everywhere they occur, will not suffice, consider the situations in which we may wish to ignore characters. First,

we often want to specify exclusion only in certain circumstances. For example, we may want to consider word-internal hyphens but to ignore them at the beginning of a stem, where they may be used to indicate that the stem must be preceded by a prefix.

Second, we may not really wish to ignore a character entirely but rather to treat it as white space. For example, if we want to treat hyphenated words as if the hyphen were not there, so that word-internal and wordinternal are of the same sort order, exclusion of hyphens will suffice, but we may wish to treat such hyphens as spaces, so that word-internal has the same sort order as word internal. In this case simple exclusion will not work, but replacement of word-internal hyphens will, as will specification of hyphen and space as having the same rank in the sort order.

Consequently, the most general approach to character exclusion is to use another program to create a new key field, attach this stripped key field to each entry in your database, and specify this field as the sort key.

For example, if you wish to ignore leading hyphens, you would convert an entry like this:

```
P:-gan
G:arm
```

into one like this:

```
P:-gan
G:arm
K:gan
```

and specify the *K*: field as the sort key.

However, since the most typical situation is one in which we wish to distinguish only between initial, medial, and final position, and since using another program like AWK can be time-consuming, *msort* provides a simple context-sensitive exclusion facility. You may specify that a character is to be ignored in any combination of field-initial, field-medial, or field-final position.

The exclusion file consists of one line per character excluded, with two fields, separated by whitespace, on each line. The first field contains a specification of the character to exclude, either as a single character or as a backslash escape. The second field consists of one to three of the letters “i”, “f”, and “m” (upper- or lower-case, in any order), indicating that the character is to be ignored in initial, final, or medial position. For example, the exclusion file:

```
-  if
'  i
```

will cause hyphens to be ignored in initial and final position, and apostrophes to be ignored only in initial position.

Where the material that you wish to ignore is always at the beginning of a tagged field, there is a simpler alternative, namely treating it as part of the tag and making use of the fact that, using regular expressions, you may specify tags that are

disjunctive or contain optional material. For example, suppose that you wish to sort on the field with tag "P:" and that you wish to ignore leading hyphens. You can do this by defining the tag as "P:-|P:", that is, as the regular expression consisting of "P:-" or "P:". Whenever a leading hyphen is encountered it will be treated as part of the tag and thereby stripped from the key proper. When there is no leading hyphen, the second disjunct "P:" will match the tag specification. An alternative in this case is to specify the tag as "P:-?", where the question mark makes the hyphen optional. If a field begins with "P:-", the hyphen will be treated as part of the tag and so effectively stripped from the content of the field. A field beginning with "P:" not followed by a hyphen will also be recognized as the tag.

It is often not desirable to exclude a character completely. For example, if we ignore a leading hyphen so as to sort on the first regular character, forms with and without leading hyphens may be interspersed (depending on their order in the input). If the input contains two tokens of *ba* and two of *-ba*, they may end up ordered:

```
ba
-ba
ba
-ba
```

If, instead, we want the forms with leading hyphens to be grouped together, like this:

```
ba
ba
-ba
-ba
```

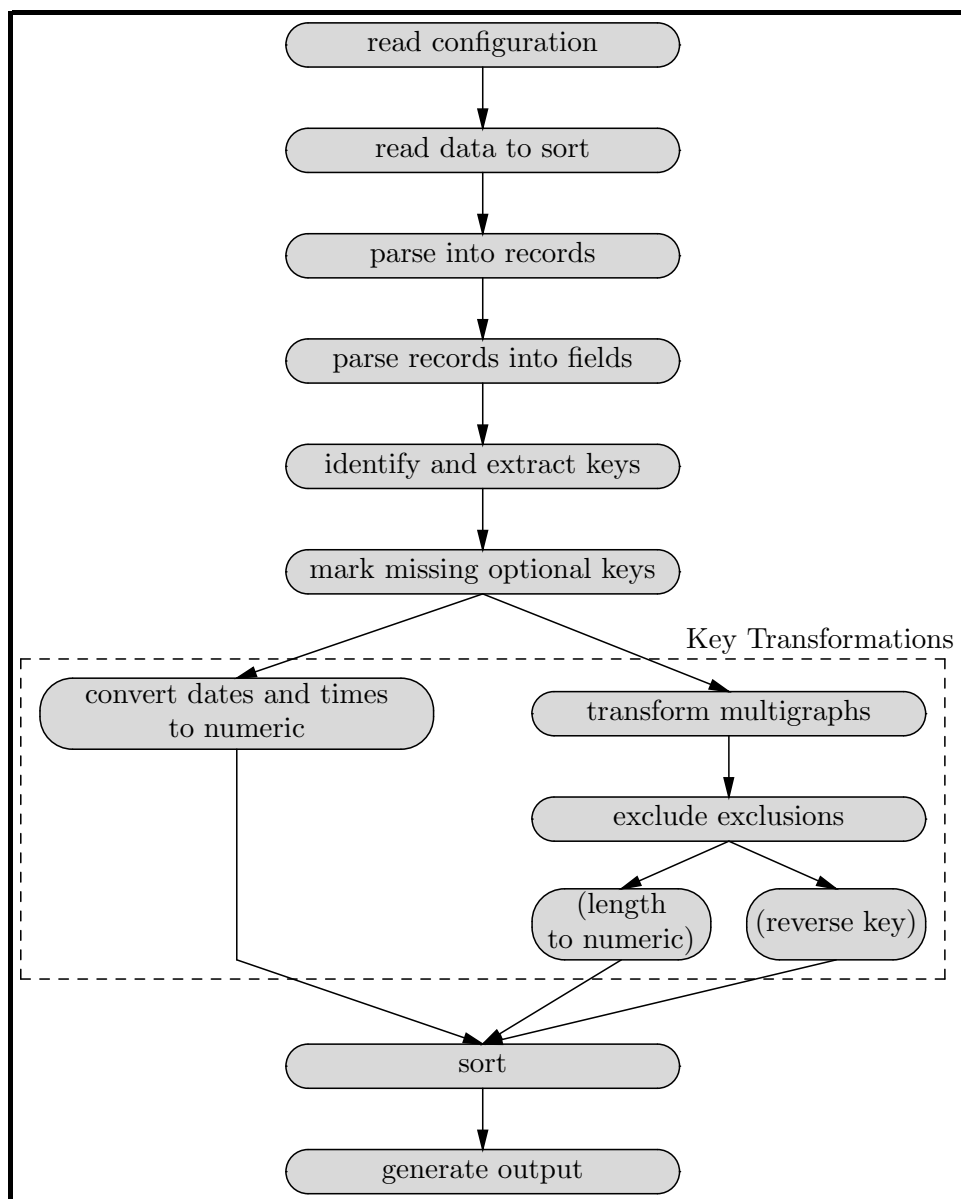
we must take into account the hyphen. To obtain this effect, use the same key field twice, in the primary key excluding leading hyphens, in the secondary key, including them, and defining the sort order so that they follow the alphabetic characters. On the first key, the forms with and without leading hyphens will tie. The tie will be broken on the second key, and the forms with hyphens will be made to follow those without.

When exclusions are specified directly on the command line using the `-X` option, the characters listed are excluded in all positions. If you wish to make the exclusion context-sensitive, you must use the `-x` option instead.

If a key is present but becomes empty as a result of exclusions, it is treated just as if the field had been missing from the outset.

4. Overview of the Sorting Process

The overall process of sorting may usefully be schematized as follows:



5. Backslash Escapes

The backslash escapes recognized by *msort* are as follows. A backslash followed by one of the digits 0-3 followed by two of the digits in the range 0-7 is treated as an octal number in the range 0-377 (0-255 decimal). The escapes `\n` for newline, `\t` for tab, and `\ space` for space are also recognized. A sequence of two backslashes is interpreted as a single backslash. Any other sequence beginning with a backslash is treated literally. Examples:

<code>\ t</code>	a tab
<code>\\</code>	a backslash
<code>\\\\</code>	two backslashes
<code>\\t</code>	a backslash followed by a tab
<code>\a</code>	a backslash followed by an “a”
<code>\n</code>	a newline
<code>\011</code>	octal 011 (decimal 9), the ASCII code for a tab
<code>\411</code>	the four characters backslash, “4”, “1”, “1”.
<code>\\011</code>	a backslash followed by a tab

Take care to understand that this description of backslash escapes is in terms of what *msort* is passed by the operating system. If the backslash character has special meaning to your operating system or command interpreter, it may be necessary to quote backslashes in order to present the desired input to *msort*. On a DOS system, you must type two backslashes (`\\`) to pass a single backslash to a program. On a UNIX system, you can put the argument within double quotes to prevent the shell from interpreting the backslash.

6. The Command Line

There are two types of command line options; some apply to a particular key; others are global. You will probably find it less confusing to group together options pertaining to a particular key, but *msort* does not require you to do this; global options may be interspersed among key-specific options.

msort interprets its command line from left to right. Whenever it encounters one of the two options that specify keys, namely `-n` and `-t`, it terminates processing of the previous key, if any, and begins constructing the specification for a new key. All key-specific options are interpreted as pertaining to the current key.

For example, suppose that the command line options are:

```
-n 3 -l -c n -n 1 -d ; -c l
```

This tells *msort* to sort on two keys. The primary key will be the third field in the record. The secondary key will be the first field in the record. The primary key is numeric. The second is lexicographic. Records consist of single lines, with fields delimited by semi-colons.

msort knows that specification that sort is to be lexicographic applies to the secondary key because it follows the field specifier for that key and no other field specifier intervenes. In effect, *msort* parses the command line like this:

```
[-n 3 -l -c n] [-n 1 -d ; -c l]
```

Any key specific-options within a group apply to that key. As a result, global options may be placed anywhere, but key-specific options must occur within the domain of their field-specifier.

7. Examples

Here are a number of examples of the use of *msort* with a variety of databases.

7.1. Example 1

Records consist of single lines with fields delimited by whitespace. A typical database fragment looks like this:

```
datasan crow N
khuni word N
```

To sort on the second field we would give the following command:

```
msort -l -n 2 db.txt
```

No specification of the field delimiter is necessary since whitespace is the default when records consist of single lines.

7.2. Example 2

Records consist of single lines with fields delimited by whitespace. A typical database fragment looks like this:

```
datasan 21
khuni 18
```

To sort in numerical order on the second field, we would give the following command:

```
msort -l -n 2 -c n db.txt
```

7.3. Example 3

Records consist of single lines with fields delimited by colons. A typical database fragment looks like this:

```
datasan:crow:N
khuni:word:N
```

To sort on the second field we would give the following command:

```
msort -l -n 2 -d : db.txt
```


7.4. Example 4

Records consist of single lines with fields delimited by whitespace. Each field is marked by a tag. A typical database fragment looks like this:

```
P:datsan G:crow C:N
P:khuni G:word C:N
```

To sort on the field with tag *P*: we would give the following command:

```
msort -l -t P: db.txt
```

7.5. Example 5

Records consist of blocks of text with each field a separate line. A typical database fragment looks like this:

```
P:datsan
G:crow
C:N

P:khuni
G:word
C:N
```

To sort on the field with tag *P*: we would give the following command:

```
msort -b -t P: db.txt
```

7.6. Example 6 — *Shoebox* Format

The Summer Institute of Linguistics uses a format in its *SHOEBOX* program and other software in which records consist of blocks of text. Each field begins with a tag consisting of a backslash (\) immediately followed by a field name of up to four characters, separated by a space from the content of the field. A field may consist of more than one line as newline characters do not serve as delimiters of fields. A typical database fragment looks like this:

```
\le datsan
\df crow
\ps N

\le khuni
\df word
\ps N
```

From *msort*'s perspective the fields of SIL-format databases are **terminated** by the backslash that begins the tag. The field name and the space that separates it

from the content of the field must be treated as the tag. To sort on the field with field name *le*, we would therefore use the following command on a UNIX system:

```
msort -b -t "le " -d "\" db.txt
```

The quotes around the tag are necessary to include the space. Recall that the space is, from *msort*'s point of view, part of the tag.

On a DOS system the command would be:

```
msort -b -t "le " -d "\\" db.txt
```

The double backslash is necessary in DOS to pass a backslash to a program on the command line.

If the field label is separated from the content of the field by arbitrary amounts of whitespace, the regular expression facility allows this to be handled as follows:

```
msort -b -t "le[040011]+" -d "\" db.txt
```

Here the brackets after "le" contain the octal character codes for space and tab. The + causes the regular expression to match one or more tokens of space or tab.

8. Logging

msort writes a record of its activity into a log file called *msort.log*. The information written is information useful when something goes wrong. In particular, ill-formed records are copied to the log file to make it easier for you to find them and repair them. The result of running *msort* with command line options that produce information about *msort* is also copied to the log file. The file is overwritten each time *msort* is run.

9. Speed

The time required for a sort depends on a variety of factors: the speed of the computer's processor, the size of the file, the number of records, storage media, the number and length of the keys, the complexity of the regular expressions defining tags, the number of multigraphs defined, and whether the sort is lexical or numerical.

In most cases almost all (over 99%) of *msort*'s run time is devoted to processing records on input, that is, to parsing them and extracting keys. The actual sort takes a small fraction of the time. This may be worthwhile if you are sorting very large databases which require minimal input processing.

The only situation in which the actual sort takes a large fraction of the time is when there are many keys that compare the same. In this case, providing an additional key on which to break ties is likely to speed up the sort.

Contrary to naive intuition, using more sort keys does not necessarily add to the time required for the sort. In fact, in some cases it can cut it considerably. When there are many records that have the same sort rank, the sort takes a long time

because a lot of comparisons are made. In such a situation, adding another key that will serve as a tie-braker reduces the number of comparisons made.

For example, sorting a list of 25,143 words by word length took 18 seconds on my laptop and required 37,762,748 comparisons. This is because there are only 22 different word lengths, so thousands of words have the same length and so compare as equal. When I added a second key, namely the same field compared lexicographically, the time dropped to less than one second because only 454,391 comparisons were made.

10. Character Set

There are two versions of *msort*, a 7-bit version and an 8-bit version, produced by conditional compilation of the same source code. When *msort* is run it will indicate which of the two versions it is along with the version number.

The 7-bit version assumes that the input is restricted to characters in which the high (8th) bit is 0. The 8-bit version permits the use of extended character sets. Since the 7-bit version uses only half the storage for keys that the 8-bit version does, it should be used when memory is at a premium, assuming that no more than 128 multigraphs are required and that the input is restricted to character codes with values less than or equal to 127.

11. Limits

msort imposes the following limits. These may also be obtained by giving the command line option *-L*.

number of keys	64
number of multigraphs per key	
7-bit version	128
8-bit version	65,279

Limiting the number of keys prevents some significant programming headaches, but nothing in the structure of the program imposes any particular limit. If you need to use more keys than are permitted and have the source code, just redefine the constant *MAXKEYS* in the file *limits.h* and recompile.

12. Defaults

These are the default values of parameters that may be set via command line options. These values may also be obtained by giving the command line option *-D*.

Field terminator characters	
record = line	whitespace

record = block	newline
Key specification	whole record
Initial maximum number of records	8192
Record type	double-newline terminated block
Sense of sort	forward

The phrase “initial maximum number of records” requires some explanation. *msort* imposes no intrinsic limit on the number of records it can handle; the only real limit is available memory. However, it allocates space for a certain number of records and then allocates more as necessary. The initial value used is the “initial maximum”. It has the default value shown above and may be set from the command line using the *-M* option. *msort* increases the maximum number of records by 50% each time it finds that it has run out. This means that the time taken for storage allocation and the sometimes necessary relocation of data in memory will be larger if you use a small initial maximum. On the other hand, using a large initial maximum and thereby a large increment may cause *msort* to allocate more storage than it really needs, conceivably using up memory needed for other purposes. If you know the approximate number of records in your database, the best strategy is to set the initial maximum a little bit larger than your estimate. If you have little idea how many records there are in a new database and want to maximize the chance of fitting into available memory, choose a small initial maximum and hence a small increment. *msort* will then fit the storage fairly closely to that actually required, at the expense of more time for repeated expansion of the record list. Most users should never need to pay attention to this option. It is provided to allow fine-tuning in special cases.

13. Summary of Command Line Options

–b	A record is terminated by two or more newlines.
–l	A record consists of a single line.
–r < separator>	A record is terminated by separator.
–d <character> ⁺	Fields are delimited by the named character(s).
–n <field number>	Sort on the specified field (counting from one).
–t <tag regexp>	Sort on the field with the specified tag.
–w	Sort on the entire text of the record.
–M <records>	Set initial maximum number of records.
–I	Invert sense of comparisons globally.
–q	Be quiet — do not chat while working.
–D	List defaults.
–F	List command line options.
–L	List limits.

Key Specific:

–c <key type>	l(exicographic), t(ime), d(ate), n(umeric), s(ize)
–f <date format>	Permutation of ymd with separators (e.g. y/m/d)
–i	Invert sense of comparisons on key.
–o <comparison>	Optional key: compare as (<, =, >) to present key if absent
–R	Reverse characters in key.
–s <file name>	Read the sort order from the named file.
–x <file name>	Read the exclusions from the named file.
–X <characters>	Exclude the specified characters.

14. Exit Status

msort returns the following status codes:

0	success
1	error opening file
2	input/output error
3	provided information and exited without sorting
4	limit exceeded
5	invalid command line option
6	invalid argument to command line option
7	ran out of memory
8	ill-formed record encountered
9	other error

15. Copyright and License

This program is copyrighted by William J. Poser. It is released under the terms of the GNU General Public License. (<http://www.gnu.org/licenses/gpl.txt>).