

# MGPP: A search engine for XML documents User guide

Katherine Don  
Department of Computer Science  
University of Waikato  
Hamilton  
New Zealand  
kjdon@cs.waikato.ac.nz

MGPP (or MG++) is a reimplementation of MG, the ‘Managing Gigabytes’ text compression and indexing system. MG is described in the book with the same name, ‘Managing Gigabytes’<sup>1</sup>, and is the default indexing tool used in the Greenstone digital library software. MG provides document level indexes, and compression of the source documents. MGPP extends MG to also provide word-level indexes, enabling proximity and fielded searching. A small change to the configuration file for a collection is all that is required to use MGPP. Once the collection is online, Greenstone provides form based search interfaces (accessed through the preferences page) enabling fielded search on any specified metadata.

This documentation describes how to use MGPP, both as part of Greenstone, and as a stand-alone indexer.

## 1 Where to get MGPP

There is no stand-alone distribution for MGPP. However, it can be obtained in two ways: from a Greenstone distribution, or from CVS.<sup>2</sup>

1. Greenstone distributions contain both MG and MGPP, and provide both the source code and binaries. Greenstone is available from

<http://www.greenstone.org/english/download.html>.

The source code for MGPP is located in `gsdl/src/mgpp`, and the executables will be in `gsdl/bin/linux` or `gsdl/bin/windows`, depending on the distribution. The configure scripts and Makefiles have been adapted for use within Greenstone, and therefore may require some modification if MGPP is to be used on its own.

---

<sup>1</sup>Managing Gigabytes: Compressing and Indexing Documents and Images (1999), Witten, I.H., Moffat, A, and Bell, T.C. Morgan Kaufmann. <http://www.cs.mu.OZ.AU/mg/>

<sup>2</sup>Concurrent Versions System, <http://www.cvshome.org>

2. Alternatively, the source code for MGPP alone can be obtained through the Greenstone CVS server in the same way that Greenstone can. The configure scripts and Makefiles are set up for use of MGPP as a stand-alone system.

The CVS package is called `mgpp`, and can be obtained through a CVS checkout using anonymous access:

```
cvscvs -d :pserver:cvscvs_anon@cvscvs.scms.waikato.ac.nz:2402/usr/local/global-cvscvs/gsd1-src co mgpp
```

## 2 Using MGPP in Greenstone

There are two parts to using MGPP within Greenstone: collection building, and collection querying. Greenstone can build collections using MG or MGPP. The default is MG, but you can use MGPP by editing the collection configuration file (`collect.cfg`, found in the `etc` directory of a collection). The configuration file syntax has changed slightly between the 2.39 and 2.40 releases of Greenstone, and both styles are described here.

### 2.1 Collection building A: gsd1-2.40 or later

Figure 1 shows a sample configuration file, with the changes required to use MGPP instead of MG shown in the box.

The line

```
searchtype form plain
```

determines what kind of search interface should be provided: `form`, for fielded searching, or `plain`, for basic searching similar to MG. If `searchtype` is specified, then MGPP will be used to build the collection instead of MG. (Note, that for backwards compatibility, specifying `buildtype mgpp` will have the same effect as specifying `searchtype plain form`.) Either or both of the options (`form`, `plain`) can be specified. The order is important: the first one will be the default type used on the search page. The other type (if present) will be presented as an option on the preferences page. If only one search type is specified, this will always be used, and there will be no option given to the user to change it.

The way indexes are specified is slightly different between MG and MGPP. MG uses a line like:

```
indexes document:text section:text section:Title
```

This specifies three indexes, a document level index containing the text of each document, and two section level indexes, one containing the text, the other the Title of each section. The document and section parts determine the granularity of the searching and of the items retrieved. The first index returns a list of document numbers while the second and third indexes return section numbers.

MGPP does things differently. The index definition is like the following.

```
indexes allfields text Title
```

```
creator me@myaddress.com
maintainer me@myaddress.com
public true
```

<pre>indexes section:text section:Title document:text defaultindex section:text  collectionmeta .section:text "chapters" collectionmeta .document:text "entire documents" collectionmeta .section:Title "chapter titles"</pre>	<pre>searchtype form plain indexes allfields text Title  levels document section  collectionmeta .allfields "entire documents" collectionmeta .text "text only" collectionmeta .Title "Titles" collectionmeta .document "book" collectionmeta .section "chapter"</pre>
--	--

```
plugin GAPlug
plugin TEXTplug
plugin HTMLPlug
...
plugin RecPlug
```

```
classify AZList -metadata Title
classify AZList -metadata Source
```

```
collectionmeta collectionname "my collection"
collectionmeta collectionextra "This is a sample collection."
```

Figure 1: A sample collection configuration file (gsdl-2.40 and later), with differences between MG (left) and MGPP (right) shown in the box.

The level information is not included in the index definition. The index specification can include the keywords `allfields`, `text` or `metadata`, along with any individual metadata elements that may be found in the collection, for example, Title, Subject, Organization etc. The items are separated by spaces. `metadata` will result in indexes being built on all the metadata that is found in the documents. `allfields` is a pseudo-index that enables searching everything that is indexed.

The order of the indexes determines the order that they will be presented in the index list, with the first entry being the default index.

The main advantage to using MGPP is that searches across multiple indexes can be carried out—“fielded searching”. For example, a search such as “Smith in Author and snail in Title” can be done.

All the indexes have the same levels. The default level is `document`, but this can be changed using a line like

```
levels document section
```

This provides search and retrieval at `document` and `section` level. `paragraph` can also be specified, but will only be available for searching, not retrieval. All of the specified indexes will be searchable at all of the specified levels. However, `paragraph` will not be available for `form` search, only for `plain` search. (As all metadata items are placed in separate paragraphs to be indexed, it doesn’t make sense to do fielded searching at paragraph level.)

`collectionmeta` can be specified for each index, as for MG collections, and also for levels. These can be in different languages, by using the `[l=xx]` syntax. For example,

```
collectionmeta .section ``chapters``  
collectionmeta .section [l=fr] ``chapitres``
```

This `collectionmeta` is optional for MGPP collections. If it is not specified, then default values will be used. For metadata, this is the original metadata name, eg Title, Subject. This will be the same for all languages. The other indexes and levels (`allfields`, `text`, `document`, `section`, `paragraph`) will default to macros which are language dependent (`_textallfields_`, `_texttextonly_`, `_textdocument_`, `_textsection_`, `_textparagraph_`). The definition of these macros can be found in `english.dm` and other language macro files.

By default, only the text is compressed, not the metadata. To change this, you can add a line to the configuration file like

```
textcompress text,Title
```

This will add Title metadata to the text that gets passed to the compressor. Whatever is specified here is what Greenstone will treat as the ‘text’ of the document. If Titles have been added to the compressed text, then they will appear anywhere the text of the document is displayed.

After the building has finished, the `build.cfg` file in the `building` directory of the collection has a list of what metadata it has found and indexed, for example

```
indexfieldmap allfields->ZZ text->TX Title->TI  
indexfields allfields text Title
```

```
indexlevels Doc Sec
```

The index names are passed to MGPP during building as two letter codes—`indexfieldmap` specifies what codes were used. These are the codes that need to be used when carrying out fielded search on the collection. `indexlevels` specifies what actual level names were used during building.

## 2.2 Collection building B: gsdl-2.39 or earlier

Greenstone versions 2.39 and earlier had a slightly different configuration file format. The differences with the later versions are explained here. Figure 2 shows the differences between the two versions, for the same collection as in Figure 1.

<pre>searchtype plain form indexes allfields text Title  levels document section  collectionmeta .allfields "entire documents" collectionmeta .text "text only" collectionmeta .Title "Titles" collectionmeta .document "book" collectionmeta .section "chapter"</pre>	<pre>buildtype mgpp indexes text,Title defaultindex text,Title levels Section  collectionmeta .text,Title "documents"</pre>
--	---

Figure 2: Differences in MGPP configuration between gsdl-2.40 (and later) (left) and gsdl-2.39 (and earlier)(right). Note that the rest of the configuration file is the same as shown in Figure 1.

The line `buildtype mgpp` specifies using MGPP to build the collection. Index specification was also slightly different. Index items were separated by commas, like `text,Title,Subject`. The keywords `text` and `metadata` were available, but not `allfields`—this was included by default. It was possible to specify more than one set of indexes, using a space for a separator. For example

```
indexes text Title,Subject,Organization
```

This would build indexes in these fields, but there would not be any cross-index searching between the two groups. A search could be done on `Title` combined with one on `Subject`, but not `Title` and `text`, for example.

(Technical note: the 2.40 specification results in only one physical index being built, with all the fields inside it. The 2.39 specification would build one physical index for each set of fields.)

Levels was also slightly different. `Document` level was always used. To include `Section` level, the line

```
levels Section
```

should be added to the configuration file. `Paragraph` was not available.

## 2.3 Collection querying

A collection built with MGPP can be searched in the usual way through Greenstone. Search terms can be combined with boolean operators, and phrases are specified using double quotes (“...”). Because MGPP uses a word level index, it has some extended searching capability over MG. If metadata has been specified in the index, fielded search can also be done.

The current query syntax involves the following:

### *Boolean operators:*

& AND | OR ! NOT, with () for precedence

### *Term modifiers:*

#i #c #u #s /x—this is stemming, casefolding and weighting

#i case insensitive

#c case sensitive

#u unstemmed

#s stemmed

/x term weight (default = 1)

These can be combined like `snail#is/10`.

### *Proximity searching:*

“...”—phrase searching, a form of very strict proximity matching where the query terms must be in the exact order specified by the phrase.

NEAR<sub>x</sub>— this is used to specify the maximum distance apart (x words) two query terms must be for a document to match. NEAR by itself defaults to 20.

### *Fielded searching:*

[terms]:Field—specifies searching within a particular field or metadata element of the document. The field name needs to be the name of a metadata element in the collection. If the collection was built with Greenstone, metadata is named with the two letter codes found in the build.cfg file. Multiple terms inside the [] are ANDed together. Boolean operators cannot be used inside the []. However, the term modifiers (#i#c#u#s/x) can be used.

Some example queries are:

computer#is	search for computer, stemmed and casefolded
Bob & Dylan/10	search for Bob and Dylan, giving Dylan a 10-fold weighting compared to Bob (used in ranking the result list)
dog NEAR4 cat	search for cat within 4 words either side of dog
[Witten]:CR	search for Witten in the CR (creator) field
[snail farming]:TI	search for snail and farming in the TI (title) field
[Witten]:CR & [Gigabytes]:TI	search for Witten in CR and Gigabytes in TI

This syntax can be entered into the standard Greenstone search box. However, there are additional query pages using forms. These can be accessed through the preferences page - select form query, then simple/ advanced. These provide fielded searching without the user needing to learn a new query syntax.

### 3 Using MGPP outside Greenstone: Indexing

The programs for building are named mgpp\_\*, for example, mgpp\_passes, mgpp\_invf\_dict etc. They have man pages with them, named mgpp\_passes.1 etc. mgpp\_passes is the main program used for building and its man page gives an example of how to build a collection.

Here is a simple bash script that can be used to build a collection,

```
#!/bin/bash

# The arguments on the command line specify the
# source of the text
source=$@

# This is the name of the collection
text=demo

# Create *.text.stats, *.invf.dict, *.invf.level
# *.invf.chunk and *.invf.chunks.trans
cat ${source} | mgpp_passes -T1 -I1 -f ${text}

# Create *.text.dict
mgpp_compression_dict -f ${text}

# Create *.invf.dict.hash
mgpp_perf_hash_build -f ${text}

# Create *.text, *.text.idx, *.text.level
# *.invf and *.invf.idx
cat ${source} | mgpp_passes -T2 -I2 -f ${text}

# Create *.text.weight and *.weight.approx
mgpp_weights_build -f ${text}

# Create *.invf.dict.blocked
```

```

mgpp_invf_dict -f ${text}

# Create *.invf.dict.blocked.1
mgpp_stem_idx -s 1 -f ${text}

# Create *.invf.dict.blocked.2
mgpp_stem_idx -s 2 -f ${text}

# Create *.invf.dict.blocked.3
mgpp_stem_idx -s 3 -f ${text}

```

This builds a basic collection, using `Document` as the document level tag, with a word level index, and three stemmed indexes (stem, casefold, stem and casefold). Each of the programs used in the compression and indexing process has different options that can be specified. The man pages have details about what each program does, and the options available.

### 3.1 Document format

The input to MGPP requires a certain format. A sample document looks like the following:

```

<Document>
  <Section>
    <Title>Snail Farming</Title>
    <Creator>Joe Bloggs</Creator>
    Snail farming is a productive and interesting pastime.
  </Section>
  <Section>
    <Title>Chapter 1</Title>
    There are many types of snails. Make sure you only farm ones
    which are safe to eat.
  </Section>
</Document>

```

There must be a document level tag that starts (and optionally ends) each document. The default is `Document`. Text outside these tags is ignored. To change the document tag that MGPP looks for, add `-J <tagname>` to the `mgpp-passes` commands. There can only be one `-J` option.

Smaller granularity is added using level tags. These must enclose all the text enclosed by the document tags. They can be specified to `mgpp-passes` by `-K <level tag>`. In the example above, `Section` would be a suitable level tag. There can be many `-K` options. However, they cannot be overlapping. You cannot have the following:

```
<Section>...<Paragraph>...</Section>...</Paragraph>
```

The smallest level of granularity (default is word level) can be changed by adding `-L <index level>` to `mgpp-passes`. This can be no larger than the smallest level that has been specified.



Metadata or tagged fields are specified like `<Title>the text of the title</Title>`. These are all automatically indexed and can be searched as fields.

The output files are placed in the directory in which the commands were run. To change where they go, use `-d <directory name>` with all of the commands.

There are four stages to `mgpp_passes`, specified by T1 and T2 (for text compression) and I1 and I2 (for indexing). T1 and I1 must be done before T2 and I2, respectively. As shown above, T1/I1 and T2/I2 can be run together, reducing the number of passes through the documents to 2. But all output files get placed in the same directory, and the same text is passed to the indexer and the compressor.

In Greenstone, we pass only the text to the compressor, but metadata and text to the indexer, and text files and index files are put in separate directories. In this case, `mgpp_passes` is run four times, like:

```
cat <text src> | mgpp_passes -d <text_dir> -T1 ...
cat <text src> | mgpp_passes -d <text_dir> -T2 ...
cat <index src> | mgpp_passes -d <index_dir> -I1 ...
cat <index src> | mgpp_passes -d <index_dir> -I2 ...
```

## 4 Using MGPP outside Greenstone: Querying

Queryer is a stand-alone program enabling search and retrieval of documents in an MGPP collection. It is written in C++, and is the equivalent to `mgquery` for MG. Even when using Greenstone, it is useful for testing collections. Queryer can be run using the following command

```
Queryer -f <index files> -t <text files>
```

`<index files>` and `<text files>` are the paths to these files, including the collection name. For example, if the text files are `demo/index/text/demo.t`, `demo/index/text/demo.ti` etc, then the text path is `demo/index/text/demo`, and similarly for the index files.

There is also a Java version of the Queryer: `org.greenstone.mgpp.Queryer`. It uses JNI (Java Native Interface) wrappers to talk to the C++ code. It is run using:

```
java org.greenstone.mgpp.Queryer <basedir> <indexdir> <textdir>
```

The two Queryer programs have a similar set of commands available:

.q	quit
.h	print the help message
.i	change the search level
.l	change the result level
.b	full text browse (enter a word or fragment at the prompt)
.r0/.r1	ranking off/on
.t0/.t1	query type some/all
.c0/.c1	casefolding off/on
.s0/.s1	stemming off/on
.o0/.o1	short output off/on
.p	print a document
query	a text query

Queries can be entered straight at the prompt. The query syntax is that described in Section 2.3.

.i and .l are used for changing the search and retrieval levels, respectively. The search level controls where two or more query terms must occur for a document or section to match, for example they must occur within a Section, or a Document. The retrieval level specifies what granularity of document/section numbers to return. These two levels need not be the same. MGPP can return document numbers where all of the query terms occur within a section. Set the search level to Section, and the retrieval level to Document. The default level used is Document. (Note that collections built with gsdI-2.40 or later use 'Doc', 'Sec' and 'Para' tags, so the levels will need to be changed before searching will work. To find out what level names were actually used, check the index/build.cfg file: the `indexlevels` line specifies what levels were indexed.)

.c and .s are used for setting global casefolding and stemming options. These settings will be applied to all of the query terms. Specifying stem and case for individual terms using the `#icus` syntax overrides these global settings.

.t specifies whether the query is a 'some' or 'all' query, i.e. specifies what the default boolean operator is. A 'some' query looks for documents containing any of the terms, corresponding to boolean OR, an 'all' query looks for documents containing all of the terms, corresponding to boolean AND. This can be overridden by specifying operators in the query. A 'some' query is most useful when the results are ranked.

.r specifies whether to return the results in ranked or natural order. Natural order returns the document numbers in increasing order. Ranked returns document numbers in order of relevance to the query—the more terms that feature in a document, the greater the relevance to the query.

.p prints a document. Document numbers are related to the retrieval level. If the retrieval level is set to Document, then 2 will specify the second document. If the level is set to Section, then 2 will specify the second section.

.b is used for full-text browsing. A word, or word fragment, is entered at the prompt, and the program returns a list of words from the index that occur next to the entered word or its closest match.

.o is used to change the output from a more descriptive result to a shorter one that doesn't list all the document numbers, just the total number of documents matched.