7 System description

7.1 Introduction

This chapter describes and illustrates the catalogue systems used for live evaluation. The description includes some background design information, some indication of what is going on "behind the screen" and occasional evaluative remarks.

As described in Chapter 6, for the live testing we constructed two systems which differ in the extent of the automatic search aids they provide. These are the *experimental* system (*EXP*) and the *control* system (*CTL*). There was also a third system (*OSTEM*) which used the same method of term combination as the live systems but performed no stemming or spelling normalisation. This was only used by the experimenters for comparative repetition of live searches.

Both the *EXP* and the *CTL* catalogues offer only a subject search. This fact is clearly displayed on the introductory screen (Fig 7.1), and was emphasised while we were collecting the evaluation data by means of a notice on top of each terminal. Both catalogues accessed the same bibliographic files, and the record displays and screen layouts are identical. It is unlikely that more than a very few users noticed any difference between the two systems since the two terminals were on different floors of the library.

The *CTL* catalogue applies "weak" stemming and spelling standardisation (6.2) to all search words. The *EXP* catalogue uses, in addition, "strong" stemming, a phrase dictionary, some automatic cross-referencing and will sometimes (about half the time) suggest an alternative for a word which it is unable to find.

The general appearance of the systems is quite similar to the subject search in the original version of Okapi (Okapi '84) described in the first Okapi report [1, Ch 7]. The most noticeable difference is an additional screen. This is the lookup and search results screen illustrated Fig 7.4 etc. It appears as soon as a user has entered a search. This screen records the progress of the search and displays options as appropriate. In Okapi '84 the bottom half of the input screen was used for these purposes.

Fig 7.1 Introductory screen

```
**********************************************************************************
:                                                                                :
:                              ** O K A P I  '8 6 **                             :
:                                                                                :
:                    EXPERIMENTAL COMPUTER CATALOGUE FOR SUBJECT SEARCHES         :
:                                                                                :
:                                                                                :
:      This catalogue has been developed by the library research team at PCL     :
:                                                                                :
:                                                                                :
:            You might have used Okapi '84  before. This version has been        :
:             improved: Okapi '86 is easier to use and is more effective         :
:                                                                                :
:                                                                                :
:           Until this version is complete -                                     :
:                                                                                :
:                      THIS CATALOGUE WILL ONLY DO SUBJECT SEARCHES               :
:                                                                                :
:      If you want to look for a particular book - use the other catalogues      :
:                                                                                :
:                                                                                :
:           To start, press the GREEN KEY on the right of the keyboard ...       :
:                                                                                :
**********************************************************************************
```

Fig 7.2 Empty input screen

```
================================================================================
                                SUBJECT SEARCH                        ** OKAPI


The computer will look for books which include all (or most) of your words
in their titles or subject descriptions


Type a word or a phrase which describes the books you want:

.....................................................................


================================================================================
```

Fig 7.3 Input screen after user has started to type

```
================================================================================
                          SUBJECT SEARCH                        ** OKAPI


The computer will look for books which include all (or most) of your words
in their titles or subject descriptions:


Type a word or a phrase which describes the books you want:

arbitration, mediation, conciliation.....................................


Press the GREEN KEY when you have finished


WHITE KEY  to change what you have typed

BLUE KEY   to get rid of what you have typed
================================================================================
```

7.2 Keyboard and display

The user stations are Apple IIe microcomputers, with six of the peripheral keys painted for use as function keys. One of them, the yellow key, is not used in any of the systems described here. In Okapi '84 it was used to invoke help or advice. Examination of logs from the earlier system showed that it was rarely used, and this seems to be the case for most online catalogues. The type of system we devised was intended to be usable at sight, and the incorporation of help adds an extraneous dependent variable to the more important ones which the systems were designed to investigate. (The yellow key is retained because it has a use in an as yet untested relevance feedback system.)

The other keys have the following general meanings:

GREEN	proceed, browse forwards
BLUE	go back one step, browse backwards
RED	stop, go back (more drastic than BLUE)
WHITE	delete last character (only works during input of search terms)
BLACK	return to introductory screen (end session)

None of the screens indicates the function of the black key, but there is a label next to the key which says PRESS WHEN FINISHED. There is, of course, a time-out. This returns to the introductory screen if no key has been pressed for 80 seconds - except during record display when the time-

out has a value of four minutes. Experience has shown that these values are about right for normal use.

The space bar is labelled with the message PRESS TO MAKE A SPACE, because logs show that an appreciable proportion of users do not separate the words of their search.

The Apple IIe keyboard is very good for people without keyboard experience because it has only 62 keys. Besides the normal keys which make displayable characters, blank and shift there are about 10 keys, all peripheral, which can be used as command or function keys. Key spacing is standard, although some users familiar with computer terminals perceive it as 'small'. Inexperienced keyboarders find many of the more modern computer terminals confusing; some terminals are easily vandalised by people who play with 'setup' and 'no scroll' keys. Suppliers of online catalogue systems would be well advised to create a demand for pleasantly simple keyboards before they become a thing of the past.

7.3 User input and preprocessing

When the green key is pressed from the introductory screen, the screen of Fig 7.2 appears. As soon as the user types anything this becomes the screen of Fig 7.3. This is a 'feature', and is perhaps carrying literalness to a fault (the argument is that since the green, white and blue keys have no effect until something has been typed, there should be no messages about their use). Its implementation is dependent on a very fast and predictable character output to the screen. This cannot be achieved without local storage and processing power.

There is room for the user to type about 76 characters. Only in one or two of the 1100 searches studied in detail did the user run out of typing space. The dots are of course meant to suggest that this is all the space which is available.

When the the user presses the green key the input string is processed to produce a search statement which, while being in a form displayable to the user, may not be exactly what the user typed. A precise description is out of place here, but the object is to transform the input into a form in which it can readily be parsed into the type of objects which are stored in the indexes.

> No distinction is made between upper and lower case, and punctuation is ignored except when it looks as though it is intended to separate words or denote an 'initialism'.

It is rare for general users to use typographical devices meaningfully (initial capitalisation to suggest a proper name, ' - ', or parentheses to introduce a subordinate description or subdivision of a subject).

Initialisms are recognised as any sequence of two or more single letters separated from each other by a single dot or a blank. Thus B B C = B.B.C. = BBC, and no distinction can be made between IT and 'it' (nor can 'it' be a stop word).

Okapi users rarely type *hyphens*. They either omit them or use blanks instead (NON VERBAL COMMUNICATION). This may be because it is not particularly obvious from the keyboard how to make a hyphen. If a word of the input string contains one hyphen, with an alphanumeric character on each side, the hyphen is removed (squashed) if it is in the second, third or fourth position in a word; otherwise is is replaced by a blank. This is a rule of thumb, probably due to Gill Venner, and makes the assumption that a short prefix is usually grammatical rather than part of a compound word. It works fairly satisfactorily because hyphenated pairs each make three index entries in Okapi '86 (GEO-PHYSICS makes GEO and PHYSICS and GEOPHYSICS). If a word contains more than one hyphen all the hyphens are replaced by blanks (before the space bar was labelled we occasionally had searches like TYPING-FOR-BEGINNERS).

After it has been pre-processed the input contains only alphanumeric characters and blanks.

In both *EXP* and *CTL* systems a second stage of processing ("weak stemming") is then applied: each "word" of the input is subjected to the weak stemming and spelling standardisation procedure, and the phrase is reassembled. This procedure is described in 6.5.1.

Example of input processing

[User]	post-war graphic designers in the U.S.A.
[Preprocess]	post war graphic designers in the usa
[Weak-stem]	post war grafic designer in the usa

Notice that the weak-stemmed string is not suitable for display to a user who doesn't know in some detail how the system works. The "searching" screen (Fig 7.4) displays the preprocessed string. Some users may come to appreciate that upper and lower case are not distinguished, that initialisms are recognised and that hyphens are squashed or replaced.

7.4 The search

This proceeds in two phases: *look-up* followed by *combination*.

The screen clears and the top line of Fig 7.4 appears.

Fig 7.4 Screen during lookup (*EXP* system)

```
================================================================================
Your search: 'arbitration mediation conciliation in industrial relations'

Looking up these words

  116 books under 'arbitration'
    9 books under 'mediation'
   29 books under 'conciliation'
  464 books under 'industrial relations'




(RED KEY  to abandon this search)
================================================================================
```

Figs 7.3 and 7.4 show a search which needs a boolean OR. The user has signalled this by using commas - which is unusual, and suggests a thoughtful or experienced user or one with a high expectation of the system. This search would fail to find any records on most online catalogues (using an implicit AND). It succeeds moderately well on both the *EXP* and the *CTL* systems, although not as well as it would if the parser were to interpret the commas or if the system knew that the first three words are quasi-synonymous.

It is here that the control and experimental systems diverge. The *EXP* system regards "industrial relations" as a single term, and will consider records indexed under "arbitrate", "arbitrating", etc. as well as those under "arbitration" (and similarly for the other words).

7.4.1 Control system

Each "word" of the weak-stemmed string is looked up in the index, unless it is in the stop list or has already been looked up. The control system look-up procedure also rejects any index entries which are flagged as being members of a cross-reference class (see 6.3). USA, for

example, will find fewer postings in the control system, because it is not equivalenced to UNITED STATES, US etc.

As each word is looked up, a message is displayed.

If the weak stem has not already been looked up there is either an 'N books under "<word>"' message, or a 'CAN'T FIND "<word>"' message. Since the components of the weak-stemmed string are not, in general, suitable for display the system shows the portion of the (preprocessed) source string which gave rise to the stem which it is looking up.

In the case of a "CAN'T FIND", the user has to take some action - replace the word, tell the system to ignore it or abort the search. The display is illustrated in Fig 7.5.

If the weak stem has already been looked up, but the source word is not the same (example: EFFECTIVE COST AND SOCIAL COSTS) there is a message of the form

('costs' included under 'cost')

Fig 7.5 Display while looking up - word not found

```
================================================================================
Your search: 'aristotles poegics'

Looking up these words

   25 books under 'aristotles'
CAN'T FIND 'poegics'

BLUE KEY  to change this word
GREEN KEY to continue without this word




(RED KEY  to abandon this search)
================================================================================
```

Fig 7.6 shows the user being prompted for a replacement search term. If the user substitutes another word or words, the top line of the screen alters accordingly (Fig 7.7), the replacement string is weak-stemmed, the "CAN'T FIND" message disappears and look-up continues.

The system assumes that the most usual and appropriate action will be to replace the word, and so this is the first option. When the user does choose to have the word ignored we were unable to arrive at a wholly satisfactory way of indicating this in the top line of the screen. If the word is simply deleted from the displayed string one is often left with a non-meaningful phrase or sentence. For example, in the search POST WAR GRAPHIC DESIGNS NI THE USA, the user might quite sensibly tell the computer to ignore NI, but

Your search: 'post war graphic designs the usa'

does not look very sensible. The easiest compromise was to replace the ignored word by something which might remind the user that a word has been omitted: at the risk of looking prudish the system displays an asterisk, and shows "post war graphic designs * the usa".

Whether the word is replaced or ignored, the "CAN'T FIND" message disappears. This is mainly to avoid recording the progress of a search overflowing the screen. Also because of screen capacity, if the number of words with non-zero hits reaches 12, the remainder of the search statement is ignored.

Fig 7.6 Retyping a misspelt word

```
================================================================================
Your search: 'aristotles poegics'

Looking up these words

   25 books under 'aristotles'
CAN'T FIND 'poegics'

Type your new word : poetics.....................................

GREEN KEY when you have finished




(RED KEY  to abandon this search)
================================================================================
```

Fig 7.7 Display during and after merging

```
================================================================================
Your search: 'aristotles poetics'

Looking up these words

   25 books under 'aristotles'
   67 books under 'poetics'

Looking for books described by your search

1 book matches your search exactly
(91 books found altogether)

GREEN KEY to look at the book(s) found
(the most similar books should appear first)



BLUE KEY  to correct or change your search

RED KEY   to do a different search
================================================================================
```

The look-up procedure is quite fast - half a second or less for each term - and it is independent of the number of postings. When look-up is complete, the system assigns weights to the terms and calculates maximum, "good" and "minimum acceptable" weights for records as described in 6.5.2.

The line "Looking for books described by your search" (Fig 7.7) appears. The system makes an estimate of the number of postings it is going to have to examine, and displays " - please wait.. " and a countdown if it is going to take more than a few seconds. A term combination, or merge, then follows. (This is, of course, trivial for single-term searches.)

7.4.2 Experimental system

Instead of looking for individual words, *EXP* submits the whole of the remainder of the weak-stemmed string for index look-up (after removal of any stop words from the beginning of the string). This look-up procedure looks for the longest match from the left of the string up to a word boundary, and returns either the number of characters matched, or failure.

For example, the search UNITED STATES AND THE WORLD WAR 2 ECONOMY matches on the 13 characters of UNITED STATES. The

stop words AND and THE are removed. Lookup is re-entered with WORLD WAR 2 ECONOMY. WORLD WAR 2 is matched and removed, and finally ECONOMY is processed.

The count displayed for each term is the number of postings for the weak stem, just as in the control system. In the case of a single-word search this can result in an inconsistency between the number of "books under" the word and the final number of books retrieved. No interviewed user mentioned this.

Whenever a single word which is not a member of an equivalence class is retrieved, it is subjected to "strong stemming" and the resulting strong stem is also looked up. In the example, UNITED STATES and WORLD WAR 2 are not strong-stemmed because they are "go phrases", but ECONOMY is looked up both as a weak and a strong stem. As a strong stem this will bring in postings for records containing "economics", "economic" and also "economical" and other potential causes of false drops.

When both the weak stem and the strong stem look-up return *failure* - that is, the index contains no match even on the strong stem of the first word of the sought string - the approximate matching technique described in 6.4 comes into play, if the word is long enough. If the system can find what it thinks is a near match, this is presented as shown in Figs 7.8 and 7.9.

If there is no candidate replacement, or if the user rejects what the system offers, the procedure is as for the control system.

If the weak-stem look-up fails, but the strong stem succeeds, the "CAN'T FIND" message takes the form

CAN'T FIND '<u>manageability</u>' - 3866 books under similar words

In the example the system will have found postings for stems arising from "manage", "manager", "management" etc. The user is not given the option of replacing the word (but can, of course, abort the search with the red key). Sometimes the strong stemming causes a user's misspelling to be matched, often to a misspelling in the source file or to a foreign word.

Fig 7.8 *EXP* system suggests a replacement

```
================================================================================
Your search: 'introductory sociolgy'

Looking up these words

  318 books under 'introductory'
CAN'T FIND 'sociolgy' - closest match found is 'sociology'

GREEN KEY  to use 'sociology' instead
BLUE KEY   to type a different word




(RED KEY  to abandon this search)
================================================================================
```

Fig 7.9 User accepts suggested replacement

```
================================================================================
Your search: 'introductory sociology'

Looking up these words

  318 books under 'introductory'
 3497 books under 'sociology'

Looking for books described by your search - please wait..

18 books match your search exactly

GREEN KEY to look at the book(s) found




BLUE KEY  to correct or change your search

RED KEY   to do a different search
================================================================================
```

The "included under" message is more frequent in the experimental system than in the control system: this message is displayed if two terms in the search belong to the same equivalence class (this is of course not known until the second term is looked up). An example is the search TERTIARY EMPLOYMENT IN UNITED KINGDOM U.K., giving

'uk' included under 'united kingdom'

which is presumably what the user intended. (It is rather rare for users to think of including synonyms.)

7.5 Term combination - the merge

There is a description of the process itself in 6.5.4. Here we show how it is presented to the user.

If the merge was always very quick (up to about three seconds, say), there would be no need to indicate anything further to the user except the number of books found.

However, with multi-word searches containing some heavily posted terms, particularly on the experimental system where both strong and weak stems enter the merge, this process could take a minute or more. The computer displays "Looking for books described by your search" (Fig 7.8). If the total number of postings to be examined is more than a few hundred, "please wait" is appended. This is followed by a "countdown" which is actually the number of disk accesses needed to read all the postings which may need to be considered. It is certainly not obvious to users what the decreasing number represents - at least one person thought it was a time in seconds - but people can see that something is happening, and none of those interviewed found it off-putting. It would be easy to display an estimate of the time, or to provide a range of messages, but we did not think that this was necessary.

At the end of the merge, the system knows how many postings there are in its output list, and also how many of these are of maximum possible weight (i.e. would have been retrieved by a boolean AND of all the terms), and how many are of "good" weight.

There are about ten different results messages depending on the distribution of weights in the output list. Partly to avoid making the system sometimes look silly and partly to avoid cluttering the screen, the message is never more than two lines. Some of the possible combinations are illustrated in Figs 7.7, 7.9, 7.10 and 7.11.

Fig 7.10 A search for two common terms which do not co-occur

```
==========================================================================
Your search: 'psychology of plants'

Looking up these words

 3034 books under 'psychology'
 1278 books under 'plants'

Looking for books described by your search - please wait

No books match your search




BLUE KEY  to correct or change your search

RED KEY   to do a different search
==========================================================================
```

Fig 7.11 Two 'rare' terms which do not co-occur

```
==========================================================================
Your search: 'yachting and boating'

Looking up these words

   5 books under 'yachting'
   5 books under 'boating'

Looking for books described by your search

10 books found but none match your search very well

GREEN KEY  to look at the book(s) found




BLUE KEY  to correct or change your words

RED KEY   to do a different search
==========================================================================
```

The green (record display) option, the blue (correct or change your search) and the red (do a different search) are always in the same order and relative positions on the screen. In a production system this would be a design weakness, because there are some searches, such as HISTORY or SOCIOLOGY, which *may* be sensible (the user may want to see where in the library these general topics are shelved), but are more likely to be associated with misconceptions about the size of the collection or the way the catalogue works. When the number of postings with "good" weight is greater than, say, fifty, it is appropriate to suggest quite prominently that the user might like to "make the search more specific" (while not preventing the display of records if that is what the user wants).

7.6 Record display

At first we tried providing both a brief one-line record and a "full" record. It was clear that single line records are not generally appropriate for subject searching because it is not possible to display the subject headings, or even the complete title. Experience with users of previous Okapi systems suggested that subject headings, while not being very good sources of index terms, are felt to be useful in making relevance judgments.

On the other hand, brief records do display slightly more quickly, and at least one user asked for them. The most serious difficulty is that of the number of options which would have to be offered during record display. If brief display were the default, with full records selected by line number and browsing forwards and backwards provided at both levels, the red key could be used to return from full to brief display as in previous Okapi systems. However, there is little doubt that full records should be the default. It may, on balance, be worth providing an additional function key for switching between full and brief displays. We did not provide this because the systems described here also form the basis of an experimental "relevance feedback" (RF) system, in which there are more options and questions at the foot of the full display screen (LOOK FOR OTHER BOOKS LIKE THIS ONE; SEE BOOKS SHELVED NEAR THIS ONE; IS THIS THE SORT OF THING YOU ARE LOOKING FOR?).

Fig 7.12 shows the information displayed and the layout of the records. The display is almost identical to earlier Okapis. Records never occupy more than one screen. In the rare cases when there would be an overflow they are arbitrarily truncated.

Fig 7.12 Full record display

```
==========================================================================
SUBJECT SEARCH                  FULL DISPLAY                    Book 3 of 9

'sociology of unemployment in the uk'

..........................................................................
AUTHOR(S)    (THOMPSON K)
TITLE(S)     The unemployed
PUBLICATION  Harrap, May 1984

SUBJECT(S)   Great Britain. Industrial sociology. Employment. Sociological
             perspectives




Not in this branch
No. of copies in other PCL libraries : MRd (1)
                          Shelved at : 306.36 WOR
..........................................................................

RED KEY  to search again or to finish
BLUE KEY to see the PREVIOUS book again      GREEN KEY to see the NEXT book

==========================================================================
```

7.6.1 Highlighting of search terms in records

Where possible, words in records which match search terms are highlighted. This feature is almost universally liked, but it is not at all easy to implement economically and universally in a system where the records retrieved may not contain the actual words of the search.

By far the easiest way to implement highlighting is by storing with each bibliographic record all the terms under which it is indexed, in the form in which they appear in the index - i.e. words and phrases, with spelling normalised and subjected to both weak and strong stemming. These index terms would be stored with pointers to the words or phrases in the body of the record which they represent. The storage overhead might be around 50%. (The stored index terms could also be used to compute measures of "similarity" between retrieved records and the query: the search WAR AND PEACE may be a better match with a record indexed only under WAR and PEACE than it is with one indexed additionally by AUSTRALIA 1914-1946.)

A slightly different approach to highlighting, avoiding the storage overhead, is to extract the index terms from a retrieved record before it is formatted for display, as if it were being indexed. (This is in fact done, in preparation for a relevance feedback system in which words extracted from relevant records would be considered for query expansion.) However, shortage of memory prevents online storage of the look-up table (go/see list) at run time, so this does not help with the highlighting of terms such as FIRST WORLD WAR which have a number of equivalent forms (WORLD WAR ONE etc).

7.6.2 Sequencing of displayed records

Records are displayed in decreasing order of weight, so that those containing all or most of the words of the query appear first, followed by those containing only some of the words, with rare words being worth more than common words in the scoring.

Where there is a "run" of records of equal weight the display sequence in these versions of Okapi is effectively arbitrary. Reverse date order is often felt to be more sensible. It is also desirable that different editions of the same work appear together. Both these objects can be achieved by storing some representation of the title and the publication date with each posting in the index. This is technically easy but represents a substantial storage overhead.

Where not all the retrieved records have the same weight, both the design team and a number of users would have liked some indication of the point in the decreasing weight sequence which has been reached. For example, if there are two records of maximum possible weight (2 BOOKS MATCH YOUR SEARCH EXACTLY), ten of "good" weight and 200 altogether, we could have provided "break" screens or half-screens which appear when the green key is pressed from the second and tenth record respectively. These could display something like

The next 8 books probably match your search less well

190 more books, but they don't match your search very well

7.6.3 Options following record display

In Okapi '88 the current search is lost on exit from record display - except that the user is reminded of it on the next input screen (Fig 7.13). There is a case for the red key returning the user to a screen showing the results of the current search. This screen would contain options to start a new search, to alter the current search or re-display the records (perhaps with a sub-option of returning to the record display at the point at which it was left).

This is the approach provided in the SWALCAP LIBERTAS online catalogue. It is difficult to think of a suitably concise and expressive message for the red key prompt.

We have no evidence as to the relative merits of the two approaches: the first has the virtue of simplicity, but the second provides more flexibility.

7.7 Second and subsequent input screens

When the user presses the red key from a record display, a modified input screen appears (Fig 7.13). Okapi '84 made it too easy for users to repeat their previous searches by always returning to an input screen with the previous search occupying the input field. In the new systems second and subsequent input screens have a blank input field but carry a brief reminder of up to three previous searches.

Fig 7.13 **Subsequent input screen showing the results of previous searches**

```
==========================================================================
Previous search(es)                                            No of books

'domestic violence'                                                      2
'battered women'                                                        22
'women'                                                      more than 500


Type a word or a phrase which describes the books you want:

█.....................................................................


==========================================================================
```

The "number of books" for previous searches is the total number retrieved, not the number which matched "exactly" or "quite well". It would be better to display something like

Fig 7.14 **Improved subsequent input screen**

```
==========================================================================
Previous search(es)                                     No of books found
                                                      good matches  total

'domestic violence' . . . . . . . . . . . . . . . . . .      2          2
'battered women'  . . . . . . . . . . . . . . . . . . .      1         22
'women' . . . . . . . . . . . . . . . . . . . . . . . .    more than 500
==========================================================================
```

Reference

1 MITEV N N, VENNER G M and WALKER S. *Designing an online public access catalogue : Okapi, a catalogue on a local area network* (Library and Information Research Report 39). London : the British Library, 1985.