

# The Icon Memory Monitoring System

Gregg M. Townsend

Department of Computer Science, The University of Arizona

## Introduction

The Icon memory monitoring system ("MemMon") provides tools for displaying Icon's allocated data regions [.iibook.]. It consists of instrumentation that produces *allocation history files* (Appendix D) and visualization programs that convert allocation history files to displays that show the sizes, types, and locations of strings as they are allocated. The garbage collection process is shown in detail.

There are several visualization programs; most of them are specific to the University of Arizona environment. Appendix A describes *xmemmon*, a program for displaying history files under the X Window System. Appendix B describes *mmps*, a program for producing displays that can be printed on any PostScript printer.

An allocation history file is produced by setting the environment variable `MEMMON` to the name of the desired file. No change in the Icon program is necessary and the production of an allocation history file does not change program behavior (except for increasing run time somewhat). On Unix systems, if the value of the `MEMMON` environment variable begins with '|', the rest of the value is interpreted as a shell command into which the history is piped.

## The Display

Icon has two primary allocated data regions: a string region and a block region. On implementations that support region expansion [.ipd51.], there is also a static region. The display shows the regions as if they were contiguous, which they are on implementations that support region expansion. The static region, if it exists, comes first, followed by the string and block regions. The choice of regions that are displayed can be specified. By default, the static region is not displayed, but the string and block regions are.

Color distinguishes the various uses of memory, and by inference the region boundaries. A legend at the top of the screen gives the meaning of each color. Default colors can be changed by providing an alternate specification file.

Colors and boundaries are set at allocation time; subsequent changes are not reflected until garbage collection occurs. For example, a string constructed in pieces may not show as a concatenated whole.

The line above the color legend gives the program status at the left, the name of the allocation history file in the center, and storage information at the right, such as

```
60480 + 25600 + 51200 (0+0+1+0)
```

The first three numbers give the current sizes (in bytes) of the static, string, and block regions respectively. The ordering reflects that of the display. The first three numbers in parentheses count the garbage collections caused by exhaustion of the regions, with the fourth number counting garbage collections initiated by calls to `collect(0)`.

## Garbage Collection

The garbage collection process is shown in detail by producing snapshots at critical points. The first comes at the beginning of garbage collection, and indicates the reason the garbage collection is required.

The next snapshot follows the marking phase. It shows the marked (active) data in color, before compaction, with the garbage painted black.

The next snapshot shows active data in dark gray; the garbage to be discarded shows in color. This is the inverse of the previous display.

The last snapshot shows the state of memory at the end of garbage collection, after compacting the active data. All garbage is gone, and the string region shows a single unbroken string. At this point the image may be rescaled to handle region expansion.

Alien blocks (such as I/O buffers) in the static region are not subject to marking or garbage collection and instead remain on constant display throughout. Obsolete co-expression blocks are freed during the marking phase, but they are displayed in a manner similar to other blocks so that their disappearance can be noted.

### The Programmer's Interface

Three built-in Icon functions write to the allocation history to control a subsequent MemMon run.

`mmpause(s)` generates a snapshot similar to those during garbage collection; the name comes from its effect on interactive visualization programs, which pause at this point. `s`, if supplied, is displayed to identify the pause. The default for `s` is "programmed pause".

`mmshow(x,s)` redraws `x` if `x` is in the managed memory region. This can be used to identify one or more particular data objects on the display. `s` determines the color of `x`:

"w"	white
"g"	gray
"b"	black
"h"	highlight: blinking white and black
"r"	redraw in normal color (the default)

The altered display persists until the next garbage collection. If `x` is outside the managed memory region, no action is taken.

`mmout(s)` writes `s` (without further interpretation) as a separate line in the history file. This can be used to insert comments (beginning with #).

All three functions return the null value.

### Accessing the Monitoring System

The `xmemmon` and `mmmps` programs are built in the `v8/src/memmon` subdirectory of the Icon distribution. Instructions for running these programs appear as appendices.

That subdirectory also includes some sample color specification files as well as code for building other visualization programs. Further information about this appears in [ipd149.].

## Appendix A: Running xmemmon

**xmemmon** displays an allocation history file interactively under the X Window System. The display is designed for a color screen but **xmemmon** uses textures to represent colors if run in monochrome.

Action stops for acknowledgement at certain *pausepoints*. The RUN button on the display, the leftmost mouse button, or a carriage return acknowledges a pause. A g typein clears all pausepoints.

The STOP button, the center mouse button, or the space bar causes **xmemmon** to pause at any time. During a pause the STEP button, center mouse button, or space bar causes a "single-step" display of the next memory block.

The RESTART button or an r typein rewinds the history file and pauses; then the history file is replayed from the beginning.

The QUIT button or a q, DEL, or ^C exits **xmemmon**.

The rightmost mouse button pulls down a menu controlling the regions that are displayed and the pausepoints that are active. Changes to the region list are not reflected until after the next garbage collection. loop mode and showMarking (both described below) can also be changed.

### Command format

**xmemmon** [ options ] [ file ]

### Options

**xmemmon** accepts the standard set of X-window command options such as `-display`, `-geometry`, and `-font`. Colors are best controlled using a color specification file as described in appendix C.

Several additional options are particular to **xmemmon**. These may be abbreviated to a unique prefix or to the single character indicated:

- |   |  |
|---|--|
| <code>-regions</code> <i>regions</i> (or <code>-r</code> )    | Display the indicated memory regions:<br>f static (fixed) region<br>s string region<br>b block region<br>The default is <code>-regions sb</code> .   |
| <code>-pausePoints</code> <i>when</i> (or <code>-p</code> )   | Pause the display at the indicated points:<br>g showing garbage remaining after marking<br>f memory full (beginning of garbage collection)<br>a showing active blocks after marking<br>g showing garbage remaining after marking<br>c after compaction (end of garbage collection)<br>p explicit <code>mmpause()</code> calls<br>r after a restart<br>n never<br>The default is <code>-pause fagcpr</code> . |
| <code>-showMarking</code> (or <code>-m</code> )               | Run through the marking phase even when not pausing to display the results. Normally, marking is bypassed if neither <code>g</code> or <code>a</code> is selected as a pausepoint.   |
| <code>-loop</code> (or <code>-l</code> )                      | Loop repeatedly through the history without pausing. This overrides other options to set <code>-pause n</code> and <code>-showMarking</code> .   |
| <code>-title</code> <i>title</i> (or <code>-t</code> )        | Set the display title. The default is the input file name.   |
| <code>-colorSpec</code> <i>filename</i> (or <code>-c</code> ) | Use an alternate color specification file (see Appendix C).  |
| <code>-width</code> <i>n</i> (or <code>-w</code> )            | Set the width of the memory subwindow to <i>n</i> pixels.  |

<code>-height <i>n</i></code> (or <code>-h</code> )	Set the height of the memory subwindow to <i>n</i> pixels.
<code>-granularity <i>n</i></code> (or <code>-b</code> )	Represent <i>n</i> bytes of memory in each displayed pixel. The default is 4.
<code>-textHeight <i>n</i></code> (or <code>-L</code> )	Make the legend and status lines <i>n</i> pixels high.
<code>-rowHeight <i>n</i></code> (or <code>-M</code> )	Limit the memory region lines to a maximum of <i>n</i> pixels in height.
<code>-skipGC <i>n</i></code> (or <code>-g</code> )	Skip to the end of the <i>n</i> th garbage collection before displaying anything.

`xmemmon` can be customized through the X Resource Manager, for instance by placing values in an `.Xdefaults` file. The most useful resources are those corresponding to command options, such as:

- `xmemmon.colorSpec`
- `xmemmon.font`
- `xmemmon.geometry`
- `xmemmon.granularity`
- `xmemmon.pausePoints`
- `xmemmon.regions`

## Appendix B: Running mmmps

**mmmps** generates Encapsulated Postscript [.epsf.] displays of Icon memory. One or more images are produced under control of command options. Output files are Encapsulated PostScript documents suitable either for direct printing or for incorporation into other documents. The output includes complete color information, though most PostScript devices print only the black-and-white equivalent.

The default image is 468 × 624 points, or 6.5" × 8.7", centered on a standard page. One line represents 1872 bytes of memory. Smaller or larger images can be specified; larger ones are reduced to fit within the above bounds.

### Command format

**mmmps** [ options ] [ file ]

### Options

- r regions** Display the indicated memory regions:
  - f** static (fixed) region
  - s** string region
  - b** block regionThe default is **-r sb**.
- p when** Produce a snapshot at the indicated points:
  - f** memory full (beginning of garbage collection)
  - a** showing active blocks after marking
  - g** showing garbage remaining after marking
  - c** after compaction (end of garbage collection)
  - p** explicit **mmppause()** calls
  - d** ("done") when the program terminates
  - n** neverThe default is **-p fagcpd**.
- m** Run through the marking phase even when not pausing to display the results. Normally, marking is bypassed if neither of **-p ga** is selected.
- g n** Skip to the end of the *n*th garbage collection before displaying anything.
- q n** Quit after completing the *n*th garbage collection.
- Q n** Quit after the *n*th snapshot.
- b n** Make each horizontal point represent *n* bytes of memory. The default is 4.
- w n** Set the display width to *n* points. The default is 468.
- h n** Set the display height to *n* points. The default is 624.
- L n** Make the legend and status lines *n* points high. **-L 0** eliminates the header entirely. The default is 11.
- M n** Limit the memory region lines to a maximum of *n* points in height. The default is 20.
- t title** Set the display title. The default is the allocation history file name.
- c file** Use an alternate color specification file (see Appendix C).
- S n** Set the PostScript screen frequency to *n* lines per inch.

## Appendix C: Color Specification Files

A *color specification file* can be used to change some or all of the colors produced by **mmps**.

The environment variable **MMCOLORS** can be used to name a file containing color specifications. **mmps** reads this file and uses it to override the normal defaults. Then, if a file is passed by the **-c** option of **mmps**, it overrides both the built-in defaults and anything from an **MMCOLORS** file.

Lines in a color specification file contain two whitespace-separated fields, a label and a value, optionally followed by comments. Blank lines are ignored, as are lines beginning with **#**. For example:

```
# change the colors for sets
set 657 light purple for set headers
selem 637 medium purple for set elements
```

The label field matches either one of the types shown in the legend or one of these additional keywords:

background	background
bsep	block separator
ssep	string separator
marked	marked block
unmarked	unmarked block (when showing active data)
status	status message
prompt	prompt message
title	title field
regions	region sizes

The value field is a set of digits specifying a color. The possible formats are:

rgb	three octal digits
#rgb	three hexadecimal digits
#rrggbb	three pairs of hexadecimal digits
rrrr,ggggg,bbbb	three decimal values between 0 and 65535

The three components specify red, green, and blue values in that order; a minimum value is off and a maximum value is full on. Some examples:

000	#000	#000000	0,0,0	(black)
777	#FFF	#ffffff	65535,65535,65535	(white)
020	#040	#002c00	0,0,20000	(dark green)
447	#99f	#8F8FFF	40000,40000,65535	(light blue)

Unfortunately, the final colors are somewhat device-dependent because of different responses to the same specification.

## Appendix D: Allocation History Files

An allocation history file is composed of printable characters forming a sequence of commands that trace interpreter actions related to memory management. This section describes the overall structure of an allocation history file, using terms and commands that are described later in detail.

An allocation history file begins with a *refresh sequence*, which completely specifies the memory layout at a particular instant. The initial refresh sequence gives a snapshot of memory just before the start of execution. Within a refresh sequence, *item* commands enumerate all the objects within the three regions, as if they are being placed, in order, into initially empty regions.

After the initial refresh sequence, the rest of the file contains any number of the following components, in any order:

- item commands*
- interaction commands*
- garbage collection sequences*

Item commands, when outside other sequences, record new allocations of memory.

Interaction commands are generated by programmed calls to `mmshow()` and `mmppause()`.

A garbage collection sequence begins with a *marking sequence*, in which item commands identify live objects. The marking sequence is followed by a new refresh sequence giving the memory configuration after compaction. Then, a final marker signals the end of garbage collection.

*Comment* and *verification* commands may appear at any point in the file.

### Command Format

A command has several components:

[*addr+*] [*len*] *cmd* [*etc*]

*addr* is an optional address, given as a distance from the start of a region. If the address is omitted, the current end of the region is assumed. An address is always followed immediately by + .

*len* is a length. If a length is needed by a command, but none is supplied, then the most recent length specified for that particular command is used.

*cmd* is a single character identifying the command.

*etc* is additional information needed by a few particular commands.

Addresses and lengths are nonnegative decimal numbers. In the string region, the unit of measurement is a character; in the static and block regions, it is specified in the refresh sequence and usually is 4 bytes.

Whitespace between commands is optional; whitespace within a command is allowed only in *etc* data.

### Item Commands

Item commands identify individual allocated objects. The meaning of an item command depends on its context. Within a refresh sequence, item commands enumerate the existing allocations. Within a marking sequence, item commands mark live objects. Otherwise, item commands announce new allocations.

Except during marking, string and block region item commands do not include addresses. Addresses may be obtained by totaling the allocations made since the beginning of the last refresh sequence.

The item commands for the block region are:

[ <i>addr+</i> ] [ <i>len</i> ] C	cset
[ <i>addr+</i> ] [ <i>len</i> ] e	table element trapped variable
[ <i>addr+</i> ] [ <i>len</i> ] E	external block
[ <i>addr+</i> ] [ <i>len</i> ] f	file block
[ <i>addr+</i> ] [ <i>len</i> ] L	list header
[ <i>addr+</i> ] [ <i>len</i> ] l	list element

[ <i>addr+</i> ] [ <i>len</i> ] R	record
[ <i>addr+</i> ] [ <i>len</i> ] r	real number
[ <i>addr+</i> ] [ <i>len</i> ] S	set header
[ <i>addr+</i> ] [ <i>len</i> ] s	set element
[ <i>addr+</i> ] [ <i>len</i> ] T	table header
[ <i>addr+</i> ] [ <i>len</i> ] t	table element
[ <i>addr+</i> ] [ <i>len</i> ] u	substring trapped variable
[ <i>addr+</i> ] [ <i>len</i> ] x	co-expression refresh block

The string item command is:

[*addr+*] [*len*] " string

The static region item commands always include an address. They are:

<i>addr</i> + [ <i>len</i> ] X	co-expression block
<i>addr</i> + [ <i>len</i> ] A	alien block

### Deallocation Commands

[*len*] -

Deallocate from the end of the allocated portion of the block region.

[*len*] ^

Deallocate from the end of the allocated portion of the string region.

### Refresh Sequence Commands

[*units*] < *static-region string-region block-region*

Begin a refresh sequence. Each region specification has the form

*base* : *used* / *max*

where *base* is the beginning address of the region, *used* is the amount of memory used, and *max* is the amount of memory allocated to the region. For the static region, the *used* value is meaningless, and *base* and *max* are zero with fixed-region versions of Icon. All values in this command are given in bytes. The *units* parameter, if present, gives the size of a unit of measurement for other commands referencing the static and block regions. If the *units* parameter is absent, the unit of measurement is four bytes.

> End a refresh sequence

### Verification Command

= *static-region block-region string-region*

Region specifications are the same as for the < command. This command provides no new information but confirms the accumulated memory usage after a series of allocations.

### Garbage Collection Commands

*n* { Start a garbage collection and begin the marking phase

*n* indicates the reason for the collection:

0	collect(0) call
1	static region
2	string region
3	block region

} End the marking phase



