

Contents

| | |
|---|----|
| Preface..... | 4 |
| Standardisation..... | 4 |
| So what is "standard"?..... | 5 |
| Notation..... | 6 |
| Where are all the grammar tables?..... | 6 |
| Acknowledgements..... | 7 |
| Overview of Z-machine architecture..... | 8 |
| 1. The memory map..... | 11 |
| 1.1 Regions of memory..... | 11 |
| 1.2 Addresses..... | 12 |
| Remarks..... | 13 |
| 2. Numbers and arithmetic..... | 14 |
| 2.1 Numbers..... | 14 |
| 2.2 Signed operations..... | 14 |
| 2.3 Arithmetic errors..... | 14 |
| 2.4 Random number generator..... | 14 |
| Remarks..... | 15 |
| 3. How text and characters are encoded..... | 16 |
| 3.1 Text..... | 16 |
| 3.2 Alphabets..... | 16 |
| 3.3 Abbreviations..... | 17 |
| 3.4 ZSCII escape..... | 17 |
| 3.5 Alphabet table..... | 17 |
| 3.6 Padding and incompleteness..... | 19 |
| 3.7 Dictionary truncation..... | 19 |
| 3.8 Definition of ZSCII and Unicode..... | 20 |
| Remarks..... | 26 |
| 4. How instructions are encoded..... | 28 |
| 4.1 Instructions..... | 28 |
| 4.2 Operand types..... | 28 |
| 4.3 Form and operand count..... | 29 |
| 4.4 Specifying operand types..... | 29 |
| 4.5 Operands..... | 30 |
| 4.6 Stores..... | 30 |
| 4.7 Branches..... | 30 |
| 4.8 Text opcodes..... | 31 |
| Remarks..... | 31 |
| 5. How routines are encoded..... | 33 |
| 5.1 Start position..... | 33 |
| 5.2 Header..... | 33 |
| 5.3 First instruction..... | 33 |
| 5.4 Main routine (V6)..... | 33 |
| 5.5 Initial execution point (other versions)..... | 33 |
| Remarks..... | 33 |
| 6. The game state: storage and routine calls..... | 34 |
| 6.1 Saved states..... | 34 |
| 6.2 Storage of global variables..... | 35 |
| 6.3 The stack..... | 35 |
| 6.4 Routine calls..... | 36 |
| 6.5 Stack frames..... | 37 |

| | |
|---|----|
| 6.6 User stacks (V6)..... | 37 |
| Remarks..... | 37 |
| 7. Output streams and file handling..... | 39 |
| 7.1 Output streams..... | 39 |
| 7.2 Buffering..... | 40 |
| 7.3 Selection (V1 and V2)..... | 40 |
| 7.4 Selection (later versions)..... | 41 |
| 7.5 Dealing with Unicode or invalid characters..... | 41 |
| 7.6 File handling..... | 41 |
| Remarks..... | 43 |
| 8. The screen model..... | 44 |
| 8.1 Fonts..... | 44 |
| 8.2 Status line..... | 45 |
| 8.3 Text colours..... | 46 |
| 8.4 Screen dimensions..... | 49 |
| 8.5 Screen model (V1, V2)..... | 49 |
| 8.6 Screen model (V3)..... | 49 |
| 8.7 Screen model (V4, V5)..... | 50 |
| 8.8 Screen model (V6)..... | 53 |
| Remarks..... | 59 |
| 9. Sound effects..... | 63 |
| 9.1 Sound effects..... | 63 |
| 9.2 Numbering of..... | 63 |
| 9.3 Volume..... | 63 |
| 9.4 Sound playing autonomously..... | 63 |
| Remarks..... | 64 |
| 10. Input streams and devices..... | 66 |
| 10.1 Keyboard only in V1..... | 66 |
| 10.2 Input streams..... | 66 |
| 10.3 Mouse support..... | 66 |
| 10.4 Menu support..... | 67 |
| 10.5 Terminating characters and timed input..... | 68 |
| 10.6 Single keypresses..... | 68 |
| 10.7 Reading ZSCII from the keyboard..... | 68 |
| Remarks..... | 69 |
| 11. The format of the header..... | 70 |
| 11.1 Header format..... | 70 |
| Remarks..... | 74 |
| 12. The object table..... | 75 |
| 12.1 Storage..... | 75 |
| 12.2 Property defaults table..... | 75 |
| 12.3 Object tree..... | 75 |
| 12.5 Well-foundedness of the tree..... | 77 |
| Remarks..... | 77 |
| 13. The dictionary and lexical analysis..... | 78 |
| 13.1 Storage..... | 78 |
| 13.2 Header..... | 78 |
| 13.3 Entries (V1 to V3)..... | 78 |
| 13.4 Entries (later versions)..... | 78 |
| 13.5 Ordering..... | 79 |
| 13.6 Lexical analysis..... | 79 |
| Remarks..... | 79 |

| | |
|--|-----|
| 14. Complete table of opcodes..... | 80 |
| 14.1 Contents..... | 84 |
| 14.2 Out of range opcodes..... | 84 |
| Reading the opcode tables..... | 85 |
| Inform assembly language..... | 85 |
| Remarks..... | 88 |
| 15. Dictionary of opcodes..... | 89 |
| 15.1..... | 89 |
| 15.2..... | 89 |
| 15.3..... | 89 |
| 16. Font 3 and character graphics..... | 113 |
| 16.1..... | 113 |
| Remarks..... | 117 |
| Appendix A. Error messages and debugging..... | 119 |
| Appendix B. Conventional contents of the header..... | 120 |
| Appendix C. Resources available..... | 122 |
| Public Interpreters..... | 122 |
| Testing compliance..... | 122 |
| Compilers..... | 122 |
| Utility programs..... | 123 |
| Story files..... | 123 |
| Documents..... | 124 |
| Appendix D. A short history of the Z-machine..... | 125 |
| Appendix E. Statistics..... | 129 |
| Appendix F. Canonical Story Files..... | 132 |

Preface

The Z-machine was created on a coffee table in Pittsburgh in 1979. It is an imaginary computer whose programs are adventure games, and is well-adapted to its task, implementing complex games remarkably compactly. They were still perhaps 100K long, too large for the memory of the home computers of their day, and the Z-machine seems to have made the first usage of virtual memory on a microcomputer. Further ahead of its time was the ability to efficiently save and restore the entire execution state.

The design's cardinal principle is that any game is 100% portable to different computers: that is, any legal program exactly determines its behaviour. This portability is largely made possible by a willingness to constrain maximum as well as minimum levels of performance (for instance, dynamic memory allocation is impossible).

Infocom's catalogue continues to be sold and to be played under interpreter programs, either original Infocom ones or more recent and generally better freeware ones. About 130 story files compiled by Infocom's compiler **Zilch** survive and since 1993 very many more story files have been created with the Inform design system.

Eight Versions of the Z-machine exist, and the first byte of any "story file" (that is: any Z-machine program) gives the Version number it must be interpreted under.

Standardisation

The majority of opcode names used in this document were agreed between 1994 and 1995 as a standard set by Mark Howell, author of the disassembler **Txd** (part of the **Ztools** suite of utility programs), and Graham Nelson, author of the assembly level of Inform. They do not correspond to Infocom's opcode names.

The first version of this Standard was drawn up in November 1995, drawing on a rougher description written in 1993 and, before that, sketches of table formats by Mike Threepoint and others. It formalised what different interpreter writers regard as the Z-machine, guaranteeing a reliable and well-featured platform for writers of new games. The initial 0.2 Standard was followed by the 1.0 Standard, which contained mostly corrections and clarifications, but also added some new features.

This third version of the Standard was initially put together by Kevin Bracey and Jason C. Penney and proposed on the Z-Machine Mailing List in December 2001. After much discussion, and nine drafts, the final list of changes to the Standard was uploaded to the if-archive in May 2006. This 2014 document adds no new material to that final draft, and is merely a merging of those changes to the main Standard.

The main additions in the 1.1 Standard are:

- ◆ Better colour support, giving games access to 32,768 colours as opposed to the previous 11. Version 6 games may also make use of the new 'transparent' background colour.

- ◆ Version 6 games may now hint to the interpreter that it is safe to make changes to the display in a backing store and flush them to the screen later, rather than making changes directly to the screen. In a program carrying out a complex layered graphical composition, this may speed up the process.
- ◆ The save and restore opcodes have a new optional operand that allow the game to control whether the interpreter prompts for a filename or executes the instructions silently.
- ◆ The optional operand removed from set_font in the previous Standard has been reinstated.

There are three companion documents are attached to this Standard.

- ◆ **Quetzal** by Martin Frost, defines a standard format for saved-game files, the purpose of which is to allow a player to save a game using one interpreter, and then restore and continue playing on a new interpreter, or even an entirely new machine.
- ◆ **Blorb** by Andrew Plotkin is a standard for a "resources" file to accompany or encapsulate a Z-machine game, neatly packaging up sound and graphics in modern formats.
- ◆ **The Treaty of Babel** is a standard for bibliographic information for interactive fiction games. The aim is to provide a unified way to identify and describe story files of many different internal formats. The Treaty is not "owned" by any individual, but by a committee of representatives from each of the pieces of software which have signed up.

Standard interpreters are not required to support these standards, since they do not affect Z-Machine behaviour, but interpreter-writers are strongly encouraged to consider it.

So what is "standard"?

To call itself "Standard", an interpreter should (as far as anyone knows) obey this document exactly for every Version of the Z-machine it claims to interpret. Interpreters need not provide optional features suggested in the "remarks" sections, and need not make their source code public. Each edition of this document has a Revision number, somewhat like the JFIF identification number used by the JPEG standard. A standard interpreter should communicate its revision number in three ways:

- ◆ To someone downloading it from the Internet: by including it in its filename.
- ◆ To the player: for instance by means of an "information" option on a menu, or in an initialisation sequence.
- ◆ To the game: by writing it into bytes in the header which were always left zero before this standard was devised (see S11). A game compiled with Inform library 5/12 or later prints the revision number in its banner (if this isn't 0.0).

Few arbitrary choices have been made in writing this document. Where Infocom's own shipped interpreters disagree, or contain manifest bugs, it has usually been possible to decide which was "correct". Elsewhere, minimum levels of performance have been invented where necessary. (For example, a minimum call-stack size is needed for programmers to be sure of what level of recursion is safe.)

Those paragraphs which genuinely extend the Infocom format are marked *****[n.m]**, where **n** and **m** are the major and minor version numbers for the Standard in which the feature was added. In any event, Infocom's original shipped interpreters do not conform to this standard document, because of bugs or because of slight variations between the Inform output format and Infocom's.

Notation

Hexadecimal numbers are written with an initial dollar, as in **\$ff**, while binary numbers are written with a double-dollar as in **\$\$11011**, according to Inform conventions. The bits in a byte are numbered 0 to 7, 0 being the least significant and the top bit, 7, the most.

Story files are mechanically best identified by their release number and serial code, which are written into the header information at the bottom of Z-machine memory. The release number can be anything between 0 and 65535 but is usually between 1 and 100. The serial code can consist of any six textual characters but is usually the date of compilation, arranged YYMMDD: thus **970619** refers to June 19th, 1997.

Paul David Doherty, in his extensive investigations into Infocom's released games, introduced the notation

Release number.Serial code

to identify particular story files: for example the first production copy of 'Enchanter' is 10.830810. This notation is used throughout the Standard when individual Infocom files need to be referred to.

Where are all the grammar tables?

The Z-machine has some lexical acuity but it doesn't contain a full parser: it's like a computer without an operating system. A game program has to contain its own parser and the tables this uses are not part of the formal Z-machine specification. (Many Infocom games have similar parsing table formats simply because, until Version 6, they used an evolving version of the 'Zork I' parser. A quite different parser was used in Version 6.) Inform's parsing table formats are documented in the *Inform Technical Manual*. For the usual format of Infocom's parsing tables, see the **Ztools** utility **Infodump**.

Acknowledgements

There is an obvious resemblance between an unreadable script and a secret code; similar methods can be employed to break both. But the differences must not be overlooked. The code is deliberately designed to baffle the investigator; the script is only puzzling by accident.

John Chadwick, *The Decipherment of Linear B*

The Z-machine was originally devised by Joel Berez and Marc Blank in 1979. Marc Blank made most of the Version 4 extensions, and Version 5 was created by Dave Lebling (with contributions from others including Brian Moriarty, Duncan Blanchard and Linde Dynneson). Version 6 was largely the work of Tim Anderson and Dave Lebling.

In the reverse direction, decipherment is mostly due to the InfoTaskForce (David Beazley, George Janczuk, Peter Lisle, Russell Hoare and Chris Tham), Matthias Pfaller, Mike Threepoint, Mark Howell, Paul David Doherty and Stefan Jokisch. Only a few of the pieces in the jigsaw were placed by myself.

I gratefully acknowledge the help of Paul David Doherty and Mark Howell, who each read drafts of this paper and sent back detailed corrections; also, of Stefan Jokisch and Marnix Klooster who have put a great deal of work into the fine detail of the specification; and of all those who commented on the circulated draft. Mistakes and misunderstandings remain my own.

Graham Nelson

15 November 1995

Kevin Bracey and Stefan Jokisch discovered most of the mistakes in Standard 0.2, in developing the first Version 6 interpreters of the modern age: **Zip2000** and **Frotz**. Matthew Russotto and Mark Knibbs supplied helpful information about Infocom's own Version 6 interpreters. Stefan also kindly read and commented on numerous drafts of the present revision. Finally, discussion about this document was greatly assisted by the Z-Machine Mailing List, organised by Marnix Klooster.

Graham Nelson

22 June 1997

The majority of the clarifications and updates in this latest revision are the work of Kevin Bracey and Jason C. Penney. Thanks go also to the members of the (now defunct) Z-Machine Mailing List, and those of the intfiction.org forum, especially Dannii Willis, for bringing to light issues with my initial revision. Special thanks to Andrew Plotkin for his notes, advice and general help while working on this revised document.

David Fillmore

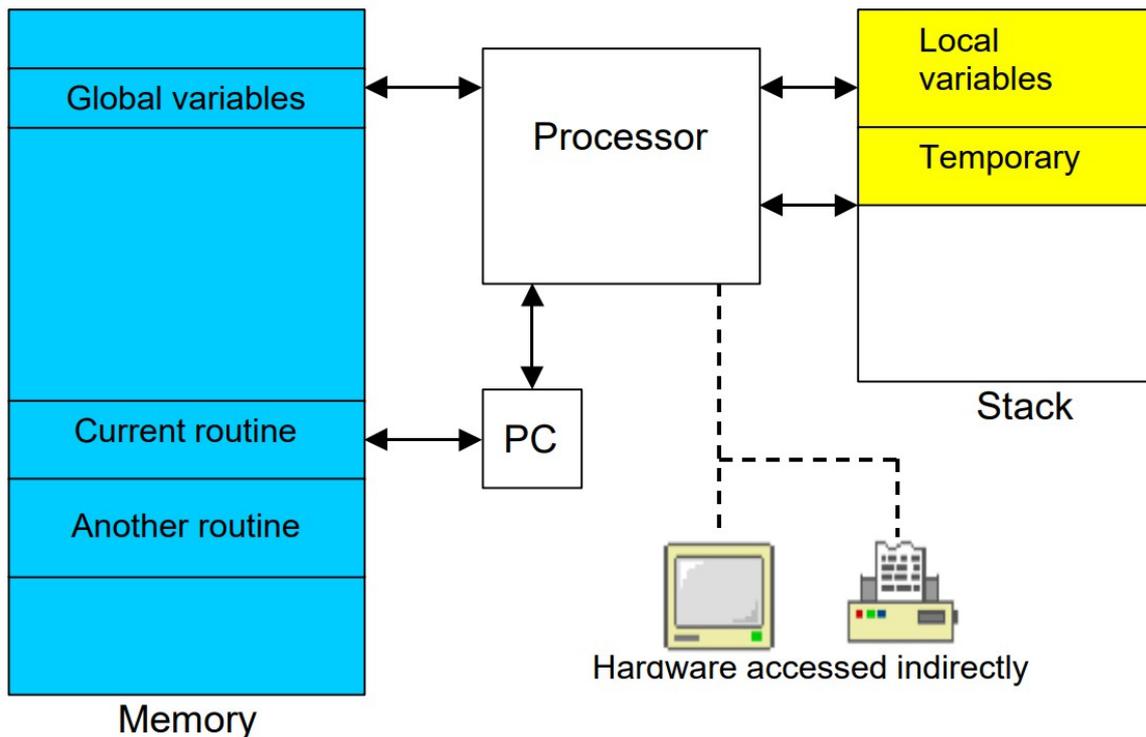
21 February 2014

The Z-Machine Standard Version 1.1 was the work of Kevin Bracey & Jason C. Penney. The initial document went through several drafts before arriving at the finished document, thanks to the comments and advice of the members of the Z-Machine Mailing List.

David Fillmore

24 February 2014

Overview of Z-machine architecture



The Z-machine is a design for an imaginary computer: Z is for 'Zork', the adventure game it was originally designed to play. Like any computer, it stores its information (mostly) in an array of variables numbered from 0 up to some large number: this is called its **memory**. A stock of some 240 memory locations are set aside for easy and quick access, and these are called **global variables** (since they are available to any part of the program which is running, at any time).

The two important pieces of information not stored in memory are the **program counter** (PC) and the **stack**. The Z-machine continuously runs a program by getting the instruction stored at position PC in memory, acting on the instruction and then moving the PC forward to the next. The **instruction set** of the Z-machine (the range of possible actions and how they are encoded as numbers in memory) occupies much of this document.

Programs are divided into **routines**: the Z-machine is always executing a particular routine, the one which the PC currently points inside. However, some instructions cause the Z-machine to **call** a new routine and then to return where the first routine left off. The Z-machine therefore needs to remember details of where to go back, and it stores these on the stack.

The stack is a second bank of memory, quite separate from the main one, which has variable size: initially it is empty. From time to time values are added to, or taken from, the top of the stack. As well as being used to keep return details, the stack is also used to store **local variables** (values needed only by a particular routine) and, for short periods only, the partial results of calculations.

Thus, whereas most physical processors (e.g. Z80 or 6502) have a number of quick-access variables outside of memory (called "registers") and a stack inside memory, the Z-machine has the reverse: it has global variables inside memory and a stack kept outside.

There is no access to hardware except by executing particular Z-machine instructions. For instance, **read** and **read_char** allow use of the keyboard; **print** and **draw_picture** allow use of the screen. The screen's image is not stored anywhere in memory. Conversely, hardware can cause the Z-machine to **interrupt**, that is, to make a spontaneous call to a particular routine, interrupting what it was previously working on. This happens only if the program has previously requested it: for example, by setting a sound effect playing and asking for a routine to be called when it finishes; or by asking for an interrupt if thirty seconds pass while the player is thinking what to type.

This simple architecture is overlaid by a number of special structures which the Z-machine maintains inside memory. There are around a dozen of these but the most important are:

- the **header**, at the bottom of memory, giving details about the program and a map of the rest of memory;
- the **dictionary**, a list of English words which the program expects that it might want to read from the keyboard;
- the **object tree**, an arrangement of chunks of memory called **objects**.

The Z-machine is primarily used for adventure games, where the dictionary holds names of items and verbs that the player might type, and the objects tend to be the places and artifacts which make up the game. Each object in the tree may have a **parent**, a **sibling** and a **child**. For instance, in the start position of 'Zork I':

West of House

You are standing in an open field west of a white house, with a boarded front door.
There is a small mailbox here.

>open mailbox

Opening the small mailbox reveals a leaflet.

At this point (part of) the game's object tree looks like this:

```
[ 41] ""
. [ 68] "West of House"
. . [ 21] "you"
. . [239] "small mailbox"
. . . [ 80] "leaflet"
. . [127] "door"
```

Note that objects are numbered from 1 upward. (Object 41 is a dummy object being used by the game to contain all the "rooms" or locations, and it has many more children besides object 68.) The parent of the player is "West of House", whose parent is 41, which has no parent. The sibling of the player is the mailbox; the child of the mailbox is the leaflet; the sibling of the mailbox is the door and so on.

Objects are bundled-up collections of variables, which come in two kinds: **attributes** and **properties**. Attributes are simply flags, that is, they can be set or unset, but have no numerical value. Properties hold numbers, which may in turn represent pieces of text or other information. For instance, one of the properties of the mailbox object above contains the information that the English word "mailbox" refers to it. One of the attributes of the mailbox object is set to indicate that it's a container, whereas the same attribute for the leaflet object is unset. Here is a breakdown of the state of the mailbox:

```

239. Attributes: 30, 34
   Parent object: 68 Sibling object: 127 Child object: 80
   Property address: 2b53
   Description: "small mailbox"
   Properties:
     [49] 00 0a
     [46] 54 bf 4a c3
     [45] 3e c1
     [44] 5b 1c

```

So the only set attributes are 30 and 34: all others are unset. Values are given for properties 44, 45, 46 and 49. The Z-machine itself does not know or care what this information means: that is for the program to sort out.

As a final example, here is part of one of the routines in 'Zork I':

```

10006: print_ret      "Suicide is not the answer."
10007: je             g57 #84 ~10008
      je             g48 #15 ~rfalse
      print_ret     "Why don't you just walk like normal people?"
10008: je             g57 #63 ~10009
      print_ret     "How romantic!"
10009: je             g57 #3b ~rfalse
      get_parent    "mirror" local0
      get_parent    "mirror" sp
      je             g6b local0 sp ~10010
      print_ret     "Your image in the mirror looks tired."
10010: print_ret     "That's difficult unless your eyes are
prehensile."

```

Z-machine programs are stored on disc, or archived on the Internet, in what are called **story files**. (Since they were introduced to hold interactive stories.) A story file consists of a snapshot of main memory only. The processor begins to run a story file by starting with an empty stack and a PC value set according to some information in the story file's header. Note that the story file has to be set up with many of the structures in memory, such as the dictionary and the object tree, already created and with sensible contents.

The first byte of any story file, and so the byte at memory address 0, always contains the version number of the Z-machine to be used. The design was evolutionary over a period of a decade: as version number increases, the instruction set grows and tables are reformatted to allow more room for larger games. All of Infocom's games can be played using versions between 3 (the majority) and 6. Games compiled by Inform mainly use versions 5 or 8.

1. *The memory map*

1.1 Regions of memory

The memory map of the Z-machine is an array of bytes with "byte addresses" running from 0 upwards. This is divided into three regions: "dynamic", "static" and "high". Dynamic memory begins from byte address **\$00000** and runs up to the byte before the byte address stored in the word at **\$0e** in the header. (Dynamic memory must contain at least 64 bytes.) Static memory follows immediately on. Its extent is not defined in the header (or anywhere else), though it must end by the last byte of the story file or by byte address **\$0fff** (whichever is lower). High memory begins at the "high memory mark" (the byte address stored in the word at **\$04** in the header) and continues to the end of the story file. The bottom of high memory may overlap with the top of static memory (but not with dynamic memory).

1.1.1

Dynamic memory can be read or written to (either directly, using **loadb**, **loadw**, **storeb** and **storew**, or indirectly with opcodes such as **insert_obj** and **remove_obj**).

1.1.1.1

By tradition, the first 64 bytes are known as the "header". The contents of this are given later but note that games are not permitted to alter many bits inside it.

1.1.1.2

It is legal for games to alter any of the tables stored in dynamic memory above the header, provided they leave the tables in legal states.

1.1.2

Static memory can be read using the opcodes **loadb** and **loadw**. It is illegal for a game to attempt to write to static memory.

1.1.3

Except for its (possible) overlap with static memory, high memory cannot be directly accessed at all by a game program. It contains routines, which can be called, and strings, which can be printed using **print_paddr**.

1.1.4

The maximum permitted length of a story file depends on the Version, as follows:

| | | |
|------|------|------|
| V1-3 | V4-5 | V6-8 |
| 128 | 256 | 512 |

1.2 Addresses

There are three kinds of address in the Z-machine, all of which can be stored in a 2-byte number: byte addresses, word addresses and packed addresses.

1.2.1

A byte address specifies a byte in memory in the range 0 up to the last byte of static memory.

1.2.2

A word address specifies an even address in the bottom 128K of memory (by giving the address divided by 2). (Word addresses are used only in the abbreviations table.)

1.2.3

***[1.0] A packed address specifies where a routine or string begins in high memory. Given a packed address P, the formula to obtain the corresponding byte address B is:

| | |
|-----------|-------------------------------------|
| 2P | Versions 1, 2 and 3 |
| 4P | Versions 4 and 5 |
| 4P + 8R_0 | Versions 6 and 7, for routine calls |
| 4P + 8S_0 | Versions 6 and 7, for print_paddr |
| 8P | Version 8 |

R_0 and S_0 are the routine and strings offsets (specified in the header as words at \$28 and \$2a, respectively).

An example memory map of a small game

| | | |
|---------|-------|------------------------------------|
| Dynamic | 00000 | header |
| | 00040 | abbreviation strings |
| | 00042 | abbreviation table |
| | 00102 | property defaults |
| | 00140 | objects |
| | 002f0 | object descriptions and properties |
| | 006e3 | global variables |
| | 008c3 | arrays |
| Static | 00b48 | grammar table |
| | 010a7 | actions table |
| | 01153 | preactions table |
| | 01201 | adjectives table |
| | 0124d | dictionary |
| High | 01a0a | Z-code |
| | 05d56 | static strings |
| | 06ae6 | end of file |

Remarks

Inform never compiles any overlap between static and high memory (it places all data tables in dynamic memory). However, many Infocom games group tables of static data just above the high memory mark, before routines begin; some, such as 'Nord 'n' Bert...', interleave static data between routines, so that static memory actually overlaps code; and a few, such as 'Seastalker' release 15, even contain routines placed below the high memory mark. (The original idea behind the high memory mark was that everything below it should be stored in the interpreter's RAM, while what was above could reasonably be kept in "virtual memory", i.e., loaded off disc as needed.)

Note that the total of dynamic plus static memory must not exceed 64K. (In fact, 64K minus 2 bytes.) This is the most serious limitation on the Z-machine.

Throughout the specification, Versions 7 and 8 are identical to Version 5 except as stated at 1.1.4 and 1.2.3 above.

2. Numbers and arithmetic

2.1 Numbers

In the Z-machine, numbers are usually stored in 2 bytes (in the form most-significant-byte first, then least-significant) and hold any value in the range \$0000 to \$ffff (0 to 65535 decimal).

2.2 Signed operations

These values are sometimes regarded as signed, in the range -32768 to 32767. In effect -n is stored as 65536-n and so the top bit is the sign bit.

2.2.1

The operations of numerical comparison, multiplication, addition, subtraction, division, remainder-after-division and printing of numbers are signed; bitwise operations are unsigned. (In particular, since comparison is signed, it is unsafe to compare two addresses using simply **jl** and **kg**.)

2.3 Arithmetic errors

Arithmetic errors:

2.3.1

It is illegal to divide by 0 (or to ask for remainder after division by 0) and an interpreter should halt with an error message if this occurs.

2.3.2

Formally it has never been specified what the result of an out-of-range calculation should be. The author suggests that the result should be reduced modulo \$10000.

2.4 Random number generator

The Z-machine needs a random number generator which at any time has one of two states, "random" and "predictable". When the game starts or restarts the state becomes "random". Ideally the generator should not produce identical sequences after each restart.

2.4.1

When "random", it must be capable of generating a uniformly random integer in the range $1 \leq x \leq n$, for any value $1 \leq n \leq 32767$. Any method can be used for this (for instance, using the host computer's clock time in milliseconds). The uniformity of randomness should be optimised for low values of n (say, up to 100 or so) and it is especially important to avoid regular patterns appearing in remainders after division (most crudely, being alternately odd and even).

2.4.2

The generator is switched into "predictable" state with a seed value. On any two occasions when the same seed is sown, identical sequences of values must result (for an indefinite period) until the generator is switched back into "random" mode. The generator should cope well with very low seed values, such as 10, and should not depend on the seed containing many non-zero bits.

2.4.3

The interpreter is permitted to switch between these states on request of the player. (This is useful for testing purposes.)

Remarks

It is dangerous to rely on the older ANSI C random number routines (rand() and srand()), as some implementations of these are very poor. This has made some games (in particular, 'Balances') unwinnable on some Unix ports of **Zip**.

The author suggests the following algorithm:

1. In "random" mode, the generator uses the host computer's clock to obtain a random sequence of bits.

2. In "predictable" mode, the generator should store the seed value S. If $S < 1000$ it should then internally generate

1, 2, 3, ..., S, 1, 2, 3, ..., S, 1, ...

so that **random n** produces the next entry in this sequence modulo n. If $S \geq 1000$ then S is used as a seed in a standard seeded random-number generator.

(The rising sequence is useful for testing, since it will produce all possible values in sequence. On the other hand, a seeded but fairly random generator is useful for testing entire scripts.)

Note that version 0.2 of this standard mistakenly asserted that division and remainder are unsigned, a myth deriving from a bug in **Zip**. Infocom's interpreters do sign division (this is relied on when calculating pizza cooking times for the microwave oven in 'The Lurking Horror'). Here are some correct Z-machine calculations:

$$\begin{array}{lll} -11 / 2 = -5 & -11 / -2 = 5 & 11 / -2 = -5 \\ -13 \% 5 = -3 & 13 \% -5 = 3 & -13 \% -5 = -3 \end{array}$$

3. *How text and characters are encoded*

This technique is similar to the five-bit Baudot code, which was used by early Teletypes before ASCII was invented.

Marc S. Blank and S. W. Galley, *How to Fit a Large Program Into a Small Machine*

3.1 Text

Z-machine text is a sequence of ZSCII character codes (ZSCII is a system similar to ASCII: see S 3.8 below). These ZSCII values are encoded into memory using a string of Z-characters. The process of converting between Z-characters and ZSCII values is given in SS 3.2 to 3.7 below.

3.2 Alphabets

Text in memory consists of a sequence of 2-byte words. Each word is divided into three 5-bit 'Z-characters', plus 1 bit left over, arranged as

```
--first byte-----  --second byte---  
7   6 5 4 3 2   1 0   7 6 5  4 3 2 1 0  
bit  --first--  --second---  --third--
```

The bit is set only on the last 2-byte word of the text, and so marks the end.

3.2.1

There are three 'alphabets', A0 (lower case), A1 (upper case) and A2 (punctuation) and during printing one of these is current at any given time. Initially A0 is current. The meaning of a Z-character may depend on which alphabet is current.

3.2.2

In Versions 1 and 2, the current alphabet can be any of the three. The Z-characters 2 and 3 are called 'shift' characters and change the alphabet for the next character only. The new alphabet depends on what the current one is:

| | from A0 | from A1 | from A2 |
|----------|---------|---------|---------|
| Z-char 2 | A1 | A2 | A0 |
| Z-char 3 | A2 | A0 | A1 |

Z-characters 4 and 5 permanently change alphabet, according to the same table, and are called 'shift lock' characters.

3.2.3

In Versions 3 and later, the current alphabet is always A0 unless changed for 1 character only: Z-characters 4 and 5 are shift characters. Thus 4 means "the next character is in A1" and 5 means "the next is in A2". There are no shift lock characters.

3.2.4

An indefinite sequence of shift or shift lock characters is legal (but prints nothing).

3.3 Abbreviations

In Versions 3 and later, Z-characters 1, 2 and 3 represent abbreviations, sometimes also called 'synonyms' (for traditional reasons): the next Z-character indicates which abbreviation string to print. If z is the first Z-character (1, 2 or 3) and x the subsequent one, then the interpreter must look up entry $32(z-1)+x$ in the abbreviations table and print the string at that word address. In Version 2, Z-character 1 has this effect (but 2 and 3 do not, so there are only 32 abbreviations).

3.3.1

Abbreviation string-printing follows all the rules of this section except that an abbreviation string must not itself use abbreviations and must not end with an incomplete multi-Z-character construction (see S 3.6.1 below).

3.4 ZSCII escape

Z-character 6 from A2 means that the two subsequent Z-characters specify a ten-bit ZSCII character code: the next Z-character gives the top 5 bits and the one after the bottom 5.

3.5 Alphabet table

The remaining Z-characters are translated into ZSCII character codes using the "alphabet table".

3.5.1

The Z-character 0 is printed as a space (ZSCII 32).

3.5.2

In Version 1, Z-character 1 is printed as a new-line (ZSCII 13).

3.5.3

In Versions 2 to 4, the alphabet table for converting Z-characters into ZSCII character codes is as follows:

| Z-char | 6789abcdef0123456789abcdef |
|---------|-------------------------------------|
| current | ----- |
| A0 | abcdefghijklmnopqrstuvwxy |
| A1 | ABCDEFGHIJKLMNPOQRSTUVWXYZ |
| A2 | ^0123456789.,!?'_#"/\<-:() ----- |

(Character 6 in A2 is printed as a space here, but is not translated using the alphabet table: see S 3.4 above. Character 7 in A2, written here as a circumflex ^, is a new-line.) For example, in alphabet A1 the Z-character 12 is translated as a capital G (ZSCII character code 71).

3.5.4

Version 1 has a slightly different A2 row in its alphabet table (new-line is not needed, making room for the < character):

| | 6789abcdef0123456789abcdef |
|----|---|
| A2 | ----- 0123456789.,!?'_#"/\<-:() ----- |

3.5.5

In Versions 5 and later, the interpreter should look at the word at **\$34** in the header. If this is zero, then the alphabet table drawn out in S 3.5.3 continues in use. Otherwise it is interpreted as the byte address of an alphabet table specific to this story file.

3.5.5.1

Such an alphabet table consists of 78 bytes arranged as 3 blocks of 26 ZSCII values, translating Z-characters 6 to 31 for alphabets A0, A1 and A2. Z-characters 6 and 7 of A2, however, are still translated as escape and newline codes (as above).

3.6 Padding and incompleteness

Since the end-bit only comes up once every three Z-characters, a string may have to be 'padded out' with null values. This is conventionally achieved with a sequence of 5's, though a sequence of (for example) 4's would work equally well.

3.6.1

It is legal for the string to end while a multi-Z-character construction is incomplete: for instance, after only the top half of an ASCII value has been given. The partial construction is simply ignored. (This can happen in printing dictionary words which have been guillotined to the dictionary resolution of 6 or 9 Z-characters.)

3.7 Dictionary truncation

When an interpreter is encrypting typed-in text to match against dictionary words, the following restrictions apply. Text should be converted to lower case (as a result A1 will not be needed unless the game provides its own alphabet table). Abbreviations may not be used. The pad character, if needed, must be 5. The total string length must be 6 Z-characters (in Versions 1 to 3) or 9 (Versions 4 and later): any multi-Z-character constructions should be left incomplete (rather than omitted) if there's no room to finish them. For example, "i" is encrypted as:

```
14, 5, 5, 5, 5, 5, 5, 5, 5  
$48a5 $14a5 $94a5
```

3.7.1

In Versions 1 and 2 only, when encoding text for dictionary words, shift-lock Z-characters 4 and 5 are used instead of the single-shift Z-characters 2 and 3 when the next two characters come from the same alphabet.

3.8 Definition of ZSCII and Unicode

The character set of the Z-machine is called ZSCII (Zork Standard Code for Information Interchange; pronounced to rhyme with "xyzzzy"). ZSCII codes are 10-bit unsigned values between 0 and 1023. Story files may only legally use the values which are defined below. Note that some values are defined only for input and some only for output.

Table 2: summary of the ZSCII rules

| | | |
|----------|-------------------------|--------------|
| 0 | null | Output |
| 1-7 | ---- | |
| 8 | delete | Input |
| 9 | tab (V6) | Output |
| 10 | ---- | |
| 11 | sentence space (V6) | Output |
| 12 | ---- | |
| 13 | newline | Input/Output |
| 14-26 | ---- | |
| 27 | escape | Input |
| 28-31 | ---- | |
| 32-126 | standard ASCII | Input/Output |
| 127-128 | ---- | |
| 129-132 | cursor u/d/l/r | Input |
| 133-144 | function keys f1 to f12 | Input |
| 145-154 | keypad 0 to 9 | Input |
| 155-251 | extra characters | Input/Output |
| 252 | menu click (V6) | Input |
| 253 | double-click (V6) | Input |
| 254 | single-click | Input |
| 255-1023 | ---- | |

3.8.1

The codes 256 to 1023 are undefined, so that for all practical purposes ZSCII is an 8-bit unsigned code.

3.8.2

The codes 0 to 31 are undefined except as follows:

3.8.2.1

ZSCII code 0 ("null") is defined for output but has no effect in any output stream. (It is also used as a value meaning "no character" when reporting terminating character codes, but is not formally defined for input.)

3.8.2.2

ZSCII code 8 ("delete") is defined for input only.

3.8.2.3

ZSCII code 9 ("tab") is defined for output in Version 6 only. At the start of a screen line this should print a paragraph indentation suitable for the font being used: if it is printed in the middle of a screen line, it should be converted to a space (Infocom's own interpreters do not do this, however).

3.8.2.4

ZSCII code 11 ("sentence space") is defined for output in Version 6 only. This should be printed as a suitable gap between two sentences (in the same way that typographers normally place larger spaces after the full stops ending sentences than after words or commas).

3.8.2.5

ZSCII code 13 ("carriage return") is defined for input and output.

3.8.2.6

ZSCII code 27 ("escape" or "break") is defined for input only.

3.8.3

ZSCII codes between 32 ("space") and 126 ("tilde") are defined for input and output, and agree with standard ASCII (as well as all of the ISO 8859 character sets and Unicode). Specifically:

```
0123456789abcdef0123456789abcdef
-----
$20  !"#%&'()*+,-./0123456789:;<=>?
$40  @ABCDEFGHIJKLMNPOQRSTUVWXYZ[\]^_
$60  'abcdefghijklmnopqrstuvwxy{!}~
-----
```

Note that code **\$23** (35 decimal) is a hash mark, not a pound sign. (Code **\$7c** (124 decimal) is a vertical stroke which is shown as ! here for typesetting reasons.)

3.8.3.1

ZSCII codes 127 ("delete" in some forms of ASCII) and 128 are undefined.

3.8.4

ZSCII codes 129 to 154 are defined for input only:

| | | | |
|----------------|------------------|------------------|-------------------|
| 129: cursor up | 130: cursor down | 131: cursor left | 132: cursor right |
| 133: f1 | 134: f2 | | 144: f12 |
| 145: keypad 0 | 146: keypad 1 | | 154: keypad 9 |

3.8.5

The block of codes between 155 and 251 are the "extra characters" and are used differently by different story files. Some will need accented Latin characters (such as French E-acute), others unusual punctuation (Spanish question mark), others new alphabets (Cyrillic or Hebrew); still others may want dingbat characters, mathematical or musical symbols, and so on.

3.8.5.1

***[1.0] To define which characters are required, the Unicode (or ISO 10646-1) Basic Multilingual Plane character set is used: characters are specified by unsigned 16-bit codes. These values agree with ISO 8859 Latin-1 in the range 0 to 255, and with ASCII and ZSCII in the range 32 to 126. The Unicode standard leaves a range of values, the Private Use Area, free: however, an Internet group called the ConScript Unicode Registry is organising a standard mapping of invented scripts (such as Klingon, or Tolkien's Elvish) into the Private Use Area, and this should be considered part of the Unicode standard for Z-machine purposes.

The Z-machine does not provide access to non-BMP characters (ie characters outside the range U+0000 to U+FFFF).

3.8.5.2

***[1.0] The story file chooses its stock of extra characters with a "Unicode translation table" as follows. Under Versions 1 to 4, the "default table" is always used (see below). In Version 5 or later, if Word 3 of the header extension table is present and non-zero then it is interpreted as the byte address of the Unicode translation table. If Word 3 is absent or zero, the default table is used.

3.8.5.2.1

The table consists of one byte giving a number N, followed by N two-byte words.

3.8.5.2.2

This indicates that ZSCII characters 155 to 155+N-1 are defined for both input and output. (It's possible for N to be zero, leaving the whole range 155 to 251 undefined.)

3.8.5.2.3

The words in the table give Unicode character codes for each of the ZSCII characters 155 to

155+N-1 in turn.

3.8.5.3

The default table is as shown in Table 1.

3.8.5.4

The defined extra characters are entirely normal ZSCII characters. They can appear in a story file's alphabet table, in an array created by print stream 3 and so on.

3.8.5.4.1

***[1.0] The interpreter is required to be able to print representations of every defined Unicode character under **\$0100** (i.e. of every defined ISO 8859-1 Latin1 character). If no suitable letter forms are available, textual equivalents may be used (such as "ss" in place of German sharp "s").

3.8.5.4.2

Normally, and where sensibly possible, all punctuation and letter characters in ISO 8859-1 Latin1 should be readable from the interpreter's keyboard. (However, some interpreters may want to provide alternative keyboard mappings, or to run in a different ISO 8859 set: Cyrillic, for example.)

3.8.5.4.3

***[1.0] An interpreter is not required to have suitable letter-forms for printing Unicode characters **\$0100** to **\$FFFF**. (It may, if it chooses, allow the user to configure certain fonts for certain Unicode ranges; but this is not required.) If a Unicode character must be printed which an interpreter has no letter-form for, a question mark should be printed instead.

3.8.5.4.4

The Z-machine is not required to handle complex Unicode formatting like combining characters, bidirectional formatting and unusual line-wrapping rules.

In Versions other than 6, interpreters may either handle these features, or not, in window 0. In window 1, and all version 6 windows, they should be ignored.

3.8.5.4.5

Unicode characters U+0000 to U+001F and U+007F to U+009F are control codes, and must not be used.

3.8.6

ZSCII codes 252 to 254 are defined for input only:

252: menu click 253: mouse double-click 254: mouse single-click

Menu clicks are available only in Version 6. A single click, or the first click of a double-click, is passed in as 254. The second click of a double-click is passed in as 253. In Versions 5 and later it is recommended that an interpreter should only send code 254, whether the mouse is clicked once or twice.

3.8.7

ZSCII code 255 is undefined. (This value is needed in the "terminating characters table" as a wildcard, indicating "any Input-only character with code 128 or above." However, it cannot itself be printed or read from the keyboard.)

Table 1: default Unicode translations (see S 3.8.5.3)

| ZSCII code (dec) | Unicode code (hex) | Name | Character | Textual Equivalent |
|------------------|--------------------|-------------|-----------|--------------------|
| 155 | 0e4 | a-diaeresis | ä | ae |
| 156 | 0f6 | o-diaeresis | ö | oe |
| 157 | 0fc | u-diaeresis | ü | ue |
| 158 | 0c4 | A-diaeresis | Ä | Ae |
| 159 | 0d6 | O-diaeresis | Ö | Oe |
| 160 | 0dc | U-diaeresis | Ü | Ue |
| 161 | 0df | sz-ligature | ß | ss |
| 162 | 0bb | quotation | » | >> or " |
| 163 | 0ab | marks | « | << or " |
| 164 | 0eb | e-diaeresis | ë | e |
| 165 | 0ef | i-diaeresis | ï | i |
| 166 | 0ff | y-diaeresis | ÿ | y |
| 167 | 0cb | E-diaeresis | Ë | E |
| 168 | 0cf | I-diaeresis | Ï | I |
| 169 | 0e1 | a-acute | á | a |
| 170 | 0e9 | e-acute | é | e |
| 171 | 0ed | i-acute | í | i |
| 172 | 0f3 | o-acute | ó | o |
| 173 | 0fa | u-acute | ú | u |
| 174 | 0fd | y-acute | ý | y |
| 175 | 0c1 | A-acute | Á | A |
| 176 | 0c9 | E-acute | É | E |

| | | | | |
|-----|-----|-----------------|---|----|
| 177 | 0cd | I-acute | Í | I |
| 178 | 0d3 | O-acute | Ó | O |
| 179 | 0da | U-acute | Ú | U |
| 180 | 0dd | Y-acute | Ý | Y |
| 181 | 0e0 | a-grave | à | a |
| 182 | 0e8 | e-grave | è | e |
| 183 | 0ec | i-grave | ì | i |
| 184 | 0f2 | o-grave | ò | o |
| 185 | 0f9 | u-grave | ù | u |
| 186 | 0c0 | A-grave | À | A |
| 187 | 0c8 | E-grave | È | E |
| 188 | 0cc | I-grave | Ì | I |
| 189 | 0d2 | O-grave | Ò | O |
| 190 | 0d9 | U-grave | Ù | U |
| 191 | 0e2 | a-circumflex | â | a |
| 192 | 0ea | e-circumflex | ê | e |
| 193 | 0ee | i-circumflex | î | i |
| 194 | 0f4 | o-circumflex | ô | o |
| 195 | 0fb | u-circumflex | û | u |
| 196 | 0c2 | A-circumflex | Â | A |
| 197 | 0ca | E-circumflex | Ê | E |
| 198 | 0ce | I-circumflex | Î | I |
| 199 | 0d4 | O-circumflex | Ô | O |
| 200 | 0db | U-circumflex | Û | U |
| 201 | 0e5 | a-ring | å | a |
| 202 | 0c5 | A-ring | Å | A |
| 203 | 0f8 | o-slash | ø | o |
| 204 | 0d8 | O-slash | Ø | O |
| 205 | 0e3 | a-tilde | ã | a |
| 206 | 0f1 | n-tilde | ñ | n |
| 207 | 0f5 | o-tilde | õ | o |
| 208 | 0c3 | A-tilde | Ã | A |
| 209 | 0d1 | N-tilde | Ñ | N |
| 210 | 0d5 | O-tilde | Õ | O |
| 211 | 0e6 | ae-ligature | æ | ae |
| 212 | 0c6 | AE-ligature | Æ | AE |
| 213 | 0e7 | c-cedilla | ç | c |
| 214 | 0c7 | C-cedilla | Ç | C |
| 215 | 0fe | Icelandic thorn | þ | th |

| | | | | |
|-----|-----|-----------------|---|----|
| 216 | 0f0 | Icelandic eth | ð | th |
| 217 | 0de | Icelandic Thorn | þ | Th |
| 218 | 0d0 | Icelandic Eth | Ð | Th |
| 219 | 0a3 | pound symbol | £ | L |
| 220 | 153 | oe-ligature | œ | oe |
| 221 | 152 | OE-ligature | Œ | OE |
| 222 | 0a1 | inverted ! | ¡ | ! |
| 223 | 0bf | inverted ? | ¿ | ? |

Remarks

In practice the text compression factor is not really very good: for instance, 155000 characters of text squashes into 99000 bytes. (Text usually accounts for about 75% of a story file.) Encoding does at least encrypt the text so that casual browsers can't read it. Well-chosen abbreviations will reduce total story file size by 10% or so.

The German translation of 'Zork I' uses an alphabet table to make accented letters (from the standard extra characters set) efficient in dictionary words. In Version 6, 'Shogun' also uses an alphabet table.

Unicode translation tables are new in Standard 1.0: in Standard 0.2, the extra characters were always mapped using the default Unicode translation table.

Note that if a random stretch of memory is accidentally printed as a string (due to an error in the story file), illegal ZSCII codes may well be printed using the 4-Z-character escape sequence. It's helpful for interpreters to filter out any such illegal codes so that the resulting on-screen mess will not cause trouble for the terminal (e.g. by causing the interpreter to print ASCII 12, clear screen, or 7, bell sound).

The continental European quotation marks << and >> should have spacing which looks sensible either in French style <<Merci!>> or in German style >>Danke!<<.

Ideally, an interpreter should be able to read time delays (for timed input) from stream 1 (i.e., from a script file). See the remarks in § 7.

The 'Beyond Zork' story file is capable of receiving both mouse-click codes (253 and 254), listing both in its terminating characters table and treating them equally.

The extant Infocom games in Versions 4 and 5 use the control characters 1 to 31 only as follows: they all accept 10 or 13 as equivalent, except that 'Bureaucracy' will only accept 13. 'Bureaucracy' needs either 127 or 8 to be a delete code. No other codes are used.

Curiously, 'Nord 'n' Bert Couldn't Make Head Nor Tail Of It' and 'A Mind Forever Voyaging' allow some letter characters to be typed in with the top bit set. That is, if reading an A, they would recognise 65 or 91 (upper or lower case) and also 193 or 219. Matthew Russotto suggests this was an accommodation for the Apple II, whose keyboard primitives returned the last key pressed in the bottom 7 bits of a byte, plus a top bit flag indicating whether or not the keyboard had been hit since last time.

In the past, not just in the Z-machine world, there has been general confusion over the rendering of ASCII/ZSCII/Latin-1/Unicode characters \$27 and \$60. For the Z-machine, the traditional interpretations of right-single-quote/apostrophe and left-single-quote are preferred over the modern neutral-single-quote and grave accent - see Table 2A of the Inform Designer's Manual. \$22 is a neutral double-quote.

An alternative rendering is to interpret both \$27 and \$60 as neutral quotes, but interpreting \$60 as a grave accent is to be avoided.

No doubt aware of this confusion, Infocom never used character \$60, and used \$27 almost exclusively as an apostrophe - hardly any single quotes appear in Infocom games. Modern authors would do well to follow their lead.

The few Infocom games that do use single quotes use \$27 for both opening and closing - but even on many of their interpreters this looked a little odd, so suggesting that \$27 be a right quote introduces no extra compatibility problems.

In Version 3 and later, many of Infocom's interpreters (and some subsequent interpreters, such as ITF's) treat two consecutive Z-characters 4 or 5 as shift locks, contrary to the Standard. As a result, story files should not use multiple consecutive 4 or 5 codes except for padding at the end of strings and dictionary words. In any case, these shift locks are not used in dictionary words, or any of Infocom's story files.

To handle languages like Arabic or Hebrew, text would have to be output "visually", with manual line breaks (possibly via an in-game formatting engine).

Far eastern languages are generally straightforward, except they usually use no spaces, and line wraps can occur almost anywhere. The easiest way to handle this would be for the game to turn off buffering. A more sophisticated game might include its own formatting engine. Also, fixed-space output is liable to be problematical with most Far Eastern fonts, which use a mixture of "full width" and "half width" forms - all half-width characters would have to be forced to full width.

4. *How instructions are encoded*

We do but teach bloody instructions

Which, being taught, return to plague th' inventor

Shakespeare, *Macbeth*

4.1 Instructions

A single Z-machine instruction consists of the following sections (and in the order shown):

| | |
|---------------------|---|
| Opcode | 1 or 2 bytes |
| (Types of operands) | 1 or 2 bytes: 4 or 8 2-bit fields |
| Operands | Between 0 and 8 of these: each 1 or 2 bytes |
| (Store variable) | 1 byte |
| (Branch offset) | 1 or 2 bytes |
| (Text to print) | An encoded string (of unlimited length) |

Bracketed sections are not present in all opcodes. (A few opcodes take both "store" and "branch".)

4.2 Operand types

There are four 'types' of operand. These are often specified by a number stored in 2 binary digits:

| | | |
|--------|-----------------------------|---------|
| \$\$00 | Large constant (0 to 65535) | 2 bytes |
| \$\$01 | Small constant (0 to 255) | 1 byte |
| \$\$10 | Variable | 1 byte |
| \$\$11 | Omitted altogether | 0 bytes |

4.2.1

Large constants, like all 2-byte words of data in the Z-machine, are stored with most significant byte first (e.g. **\$2478** is stored as **\$24** followed by **\$78**). A 'large constant' may in fact be a small number.

4.2.2

Variable number **\$00** refers to the top of the stack, **\$01** to **\$0f** mean the local variables of the current routine and **\$10** to **\$ff** mean the global variables. It is illegal to refer to local variables which do not exist for the current routine (there may even be none).

4.2.3

The type 'Variable' really means "variable by value". Some instructions take as an operand a "variable by reference": for instance, **inc** has one operand, the reference number of a variable to increment. This operand usually has type 'Small constant' (and Inform automatically assembles a line like **@inc turns** by writing the operand **turns** as a small constant with value the reference number of the variable **turns**).

4.3 Form and operand count

Each instruction has a form (long, short, extended or variable) and an operand count (0OP, 1OP, 2OP or VAR). If the top two bits of the opcode are **\$\$11** the form is variable; if **\$\$10**, the form is short. If the opcode is 190 (**\$BE** in hexadecimal) and the version is 5 or later, the form is "extended". Otherwise, the form is "long".

4.3.1

In short form, bits 4 and 5 of the opcode byte give an operand type as above. If this is **\$11** then the operand count is 0OP; otherwise, 1OP. In either case the opcode number is given in the bottom 4 bits.

4.3.2

In long form the operand count is always 2OP. The opcode number is given in the bottom 5 bits.

4.3.3

In variable form, if bit 5 is 0 then the count is 2OP; if it is 1, then the count is VAR. The opcode number is given in the bottom 5 bits.

4.3.4

In extended form, the operand count is VAR. The opcode number is given in a second opcode byte.

4.4 Specifying operand types

Next, the types of the operands are specified.

4.4.1

In short form, bits 4 and 5 of the opcode give the type.

4.4.2

In long form, bit 6 of the opcode gives the type of the first operand, bit 5 of the second. A value of 0 means a small constant and 1 means a variable. (If a 2OP instruction needs a large constant as operand, then it should be assembled in variable rather than long form.)

4.4.3

In variable or extended forms, a byte of 4 operand types is given next. This contains 4 2-bit fields: bits 6 and 7 are the first field, bits 0 and 1 the fourth. The values are operand types as above. Once one type has been given as 'omitted', all subsequent ones must be. Example: \$ **00101111** means large constant followed by variable (and no third or fourth opcode).

4.4.3.1

In the special case of the "double variable" VAR opcodes **call_vs2** and **call_vn2** (opcode numbers 12 and 26), a second byte of types is given, containing the types for the next four operands.

4.5 Operands

The operands are given next. Operand counts of 0OP, 1OP or 2OP require 0, 1 or 2 operands to be given, respectively. If the count is VAR, there must be as many operands as there were types other than 'omitted'.

4.5.1

Note that only **call_vs2** and **call_vn2** can have more than 4 operands, and no instruction can have more than 8.

4.5.2

Opcode operands are always evaluated from first to last - this order is important when the stack pointer appears as an argument. Thus

```
@sub sp sp -> i;
```

subtracts the second-from-top stack item from the topmost stack item.

4.6 Stores

"Store" instructions return a value: e.g., **mul** multiplies its two operands together. Such instructions must be followed by a single byte giving the variable number of where to put the result.

4.7 Branches

Instructions which test a condition are called "branch" instructions. The branch information is stored in one or two bytes, indicating what to do with the result of the test. If bit 7 of the first byte is 0, a branch occurs when the condition was false; if 1, then branch is on true. If bit 6 is set, then the branch occupies 1 byte only, and the "offset" is in the range 0 to 63, given in the bottom 6 bits. If bit 6 is clear, then the offset is a signed 14-bit number given in bits 0 to 5 of the first byte followed by all 8 of the second.

4.7.1

An offset of 0 means "return false from the current routine", and 1 means "return true from the current routine".

4.7.2

Otherwise, a branch moves execution to the instruction at address

$$\text{Address after branch data} + \text{Offset} - 2.$$

4.8 Text opcodes

Two opcodes, **print** and **print_ret**, are followed by a text string. This is stored according to the usual rules: in particular execution continues after the last 2-byte word of text (the one with top bit set).

Remarks

Some opcodes have type VAR only because the available codes for the other types had run out; **print_char**, for instance. Others, especially **call**, need the flexibility to have between 1 and 4 operands.

The Inform assembler can assemble branches in either form, though the programmer should always use long form unless there's a good reason. Inform automatically optimises branch statements so as to force as many of them as possible into short form. (This optimisation will happen to branches written by hand in assembler as well as to branches compiled by Inform.)

The disassembler **Txd** numbers locals from 0 to 14 and globals from 0 to 239 in its output (corresponding to variable numbers 1 to 15, and 16 to 255, respectively).

The branch formula is sensible because in the natural implementation, the program counter is at the address after the branch data when the branch takes place: thus it can be regarded as

$$\text{PC} = \text{PC} + \text{Offset} - 2.$$

If the rule were simply "add the offset" then, since the offset couldn't be 0 or 1 (because of the return-false and return-true values), we would never be able to skip past a 1-byte instruction (say, a OOP like **quit**), or specify the branch "don't branch at all" (sometimes useful to ignore the result of the test altogether). Subtracting 2 means that the only effects we can't achieve are

$$\text{PC} = \text{PC} - 1 \quad \text{and} \quad \text{PC} = \text{PC} - 2$$

and we would never want these anyway, since they would put the program counter somewhere back inside the same instruction, with horrid consequences.

On disassembly

Briefly, the first byte of an instruction can be decoded using the following table:

| | | | |
|--------------|-----------------|--------------------|---------------------------------|
| \$00 -- \$1f | long | 20P | small constant, small constant |
| \$20 -- \$3f | long | 20P | small constant, variable |
| \$40 -- \$5f | long | 20P | variable, small constant |
| \$60 -- \$7f | long | 20P | variable, variable |
| \$80 -- \$8f | short | 10P | large constant |
| \$90 -- \$9f | short | 10P | small constant |
| \$a0 -- \$af | short | 10P | variable |
| \$b0 -- \$bf | short | 00P | |
| except \$be | extended opcode | given in next byte | |
| \$c0 -- \$df | variable | 20P | (operand types in next byte) |
| \$e0 -- \$ff | variable | VAR | (operand types in next byte(s)) |

Here is an example disassembly:

```
@inc_chk c 0 label;    05 02 00 d4
    long form; count 20P; opcode number 5; operands:
        02    small constant (referring to variable c)
        00    small constant 0
    branch if true: 1-byte offset, 20 (since label is
    18 bytes forward from here).
@print "Hello.^";      b2 11 aa 46 34 16 45 9c a5
    short form; count 00P.
    literal string, Z-chars: 4 13 10 17 17 20 5 18 5 7 5 5.
@mul 1000 c -> sp;     d6 2f 03 e8 02 00
    variable form; count 20P; opcode number 22; operands:
        03 e8 long constant (1000 decimal)
        02    variable c
    store result to stack pointer (var number 00).
@call_1n Message;     8f 01 56
    short form; count 10P; opcode number 15; operand:
        01 56 long constant (packed address of routine)
.label;
```

5. How routines are encoded

5.1 Start position

A routine is required to begin at an address in memory which can be represented by a packed address (for instance, in Version 5 it must occur at a byte address which is divisible by 4).

5.2 Header

A routine begins with one byte indicating the number of local variables it has (between 0 and 15 inclusive).

5.2.1

In Versions 1 to 4, that number of 2-byte words follows, giving initial values for these local variables. In Versions 5 and later, the initial values are all zero.

5.3 First instruction

Execution of instructions begins from the byte after this header information. There is no formal 'end-marker' for a routine (it is simply assumed that execution eventually results in a return taking place).

5.4 Main routine (V6)

In Version 6, there is a "main" routine (whose packed address is stored in the word at \$06 in the header) called when the game starts up. It is illegal to return from this routine.

5.5 Initial execution point (other versions)

In all other Versions, the word at \$06 contains the byte address of the first instruction to execute. The Z-machine starts in an environment with no local variables from which, again, a return is illegal.

Remarks

Note that it is permissible for a routine to be in dynamic memory. Marnix Klooster suggests this might be used for compiling code at run time!

In Versions 3 and 4, Inform always stores 0 as the initial values for local variables.

6. The game state: storage and routine calls

6.1 Saved states

The "state of play" is defined as the following: the contents of dynamic memory; the contents of the stack; the value of the program counter (PC), and the "routine call state" (that is, the chain of routines which have called each other in sequence, and the values of their local variables). Note that the routine call state, the stack and the PC must be stored outside the Z-machine memory map, in the interpreter's private memory.

6.1.1

The entire state of play must be stored when the game is saved.

6.1.1.1

The format of a saved game file is not specified.

6.1.1.2

An internal saved game for "undo" purposes (if there is one) is not part of the state of play. This is important: if a saved game file also contained the internal saved game at the time of saving, it would be impossible to undo the act of restoration. It also prevents internal saved games from growing larger and larger as they include their predecessors.

6.1.1.3

It is illegal to save the game (either with **save** or **save_undo**) during an "interrupt routine" (one coming about through timed input, sound effect termination or newline interrupts). Therefore saved games need not store information capable of restoring such a position.

6.1.2

On a "restore" or "undo" (which restores a game saved into internal memory), the entire state of play is written back except that 'Flags 2' in the header is preserved. (This information includes whether the game is being transcribed to printer and whether a fixed-pitch font is being used.)

6.1.2.1

Before a "restore", an interpreter should check that the file to be used has been saved from the same game currently being played. (See remark below.)

6.1.2.2

After a "restore" or "undo", an interpreter should reset the header values marked **Rst** in the header table of S 11. (It should not be assumed that the game was saved by the same interpreter.)

6.1.3

A "restart" is similar: the entire state is restored from the original story file, and the stack is emptied; but 'Flags 2' is preserved; and the interpreter should reset the **Rst** parts of the header.

6.1.4

In Versions 5 and later, an interpreter unable to save the game state into internal memory (for "undo" purposes) must clear bit 4 of 'Flags 2' in the header.

6.2 Storage of global variables

Global variables (variable numbers **\$10** to **\$ff**) are stored in a table in the Z-machine's dynamic memory, at a byte address given in word 6 of the header. The table consists of 240 2-byte words and the initial values of the global variables are the values initially contained in the table. (It is legal for a program to alter the table's contents directly in play, though not for it to change the table's address.)

6.3 The stack

Writing to the stack pointer (variable number **\$00**) pushes a value onto the stack; reading from it pulls a value off. Stack entries are 2-byte words as usual.

6.3.1

The stack is considered as empty at the start of each routine: it is illegal to pull values from it unless values have first been pushed on.

6.3.2

The stack is left empty at the end of each routine: when a return occurs, any values pushed during the routine are thrown away.

6.3.3

The absolute minimum standard for stack size is defined as:

let the 'usage' of a routine call be 4 plus the number of local variables it has. During a game the total of the usages for each routine in the recursive chain of routines being called, plus the game's own stack usage, can be assumed to never reach 1024.

However, more recent games have required a much larger stack size than this allows for. It is advised that interpreters allow for these games by having a larger stack size if at all possible.

Two examples of modern interpreters with increased stack size are **Windows Frotz**, with 32768, and **nfrotz** with 61440.

6.3.4

In the seven opcodes that take indirect variable references (**inc**, **dec**, **inc_chk**, **dec_chk**, **load**, **store**, **pull**), an indirect reference to the stack pointer does not push or pull the top item of the stack - it is read or written in place.

6.4 Routine calls

Routine calls occur in the following circumstances: when one of the **call...** opcodes is executed; in Versions 4 and later, when timed keyboard input is being monitored; in Versions 5 and later, when a sound effect finishes; in Version 6, when the game begins (to call the "main" routine); in Version 6, when a "newline interrupt" occurs.

6.4.1

A routine call may have any number of arguments, from 0 to 3 (in Versions 1 to 3) or 0 to 7 (Versions 4 and later). All routines return a value (though sometimes this value is thrown away afterward: for example by opcodes in the form **call_vn***).

6.4.2

Routine calls preserve local variables and the stack (except when the return value is stored in a local variable or onto the top of the stack).

6.4.3

A routine call to packed address 0 is legal: it does nothing and returns false (0). Otherwise it is illegal to call a packed address where no routine is present.

6.4.4

When a routine is called, its local variables are created with initial values taken from the routine header (Versions 1 to 4) or with initial value 0 (Versions 5 and later). Next, the arguments are written into the local variables (argument 1 into local 1 and so on).

6.4.4.1

It is legal for there to be more arguments than local variables (any spare arguments are thrown away) or for there to be fewer.

6.4.5

The return value of a routine can be any Z-machine number. Returning 'false' means returning 0; returning 'true' means returning 1.

6.5 Stack frames

A "stack frame" is an index to the routine call state (that is, the call-stack of return addresses from routines currently running, and values of local variables within them). This index is a Z-machine number. The interpreter must be able to produce the current value and to set a value further down the call-stack than the current one, effectively throwing away its recent history (see **catch** and **throw**).

6.6 User stacks (V6)

In Version 6, the Z-machine understands a third kind of stack: a "user stack", which is a table of words in dynamic memory. The first word in this table always holds the number of spare slots on the stack (so the initial value is the capacity of the stack). The Z-machine makes no check on stack under-flow (i.e., pulling more values than were pushed) which would over-run the length of the table if the program allowed it to happen.

Remarks

Some interpreters store the whole of dynamic memory to disc as part of their saved game files, which can make them as much as 45K or so long. A player making a serious attack on a game may end up wasting a whole megabyte, more than convenient without a hard disc. A technique invented by Bryan Scattergood, taken up by most modern interpreters, greatly reduces file size by only saving bytes of dynamic memory which differ from the initial state of the game.

It is unspecified how an interpreter should decide whether a saved game file belongs to the game currently being played. It is normal to insist that the release numbers, serial codes and checksums all match. The **Pinfocom** interpreter deliberately checks only the release number, so that saved games can be exchanged between different editions of 'Seastalker' (presumably compiled to handle the sonarscope differently).

These issues are taken up in great detail in Martin Frost's **Quetzal** standard for saved game files, created to allow different interpreters to exchange saved games. This Standard doesn't require compliance with **Quetzal**, but interpreter writers are urged to consider it: it can only help authors if players can send them saved games where bugs seem to have appeared.

The stack is stored in the interpreter's own memory, not anywhere in the Z-machine. The game program has no direct access to the stack memory or stack pointer; on some implementations the game's main stack is also used to store the routine call state (i.e. the game stack and the call-stack are the same) but this need not be true.

The stack size specification guarantees in particular that if the game itself never uses more than 32 stack entries at once then it can have a recursive depth of at least 90 routine calls. The author believes that old Infocom games will all run with a stack size of 512 words.

Note that the "state of play" does not include numerous input/output settings (the current window, cursor position, splitness or otherwise, which streams are selected, etc.): neither does it include the state of the random-number generator. (Games with elaborate status lines must redraw them after a restore has taken place.)

Zip provides "undo" but most versions of the **ITF** interpreter do not (and **save_undo** returns 0, unfortunately). This is probably its greatest failing. Some Infocom-written interpreters will only provide "undo" to a game which has bit 4 of 'Flags 2' set: but Inform 5.5 doesn't set this bit, so modern interpreters should be more generous.

Given the existence of Quetzal, a portable saved file format, it is quite possible that after loading, the game may be running on a different interpreter to that on which the game started. As a result, it is strongly advisable for games to recheck any interpreter capabilities (eg Standard version, unicode support, etc) after loading.

7. *Output streams and file handling*

7.1 Output streams

At any given time text is being output through a selection of "output streams" (possibly none, possibly several at once).

7.1.1

Two output streams are common to all Versions: number 1 (the screen) and 2 (the game transcript, usually printed to a printer or a file).

7.1.1.1

In Versions 1 to 5, the player's input to the **read** opcode should be echoed to output streams 1 and 2 (if stream 2 is active), so that text typed in appears in any transcript. In Version 6 input should be sent only to stream 1 and it is the game's responsibility to write to the transcript.

7.1.1.2

In Infocom's Version 4 game 'A Mind Forever Voyaging', which anticipated a printer rather than a file to receive the transcript, stream 2 is turned off and on again several times in quick succession. Thus if an interpreter decides where to send the transcript by asking the player for a filename, this question should only be asked once per game session, not every time stream 2 is selected.

7.1.2

Versions 3 and later supply these and two other output streams, numbered 3 (Z-machine memory) and 4 (a script file of the player's whole commands and of individual keypresses as read by **read_char**).

7.1.2.1

Output stream 3 writes to a table in dynamic memory. When the stream is selected, the table may have any contents (even the initial 'size' word will be ignored by the interpreter). While the stream is selected, the table's contents are unspecified (and a game cannot safely read or write to it). When the stream is deselected, the initial word of the table holds the number of characters printed and subsequent bytes hold those characters. Similarly, in Version 6, the total width of printing (in units) will then be stored in the word at **\$30** in the header. (It is the programmer's responsibility to make the table large enough: the interpreter performs no overflow checking.)

7.1.2.1.1

***[1.0] It is possible for stream 3 to be selected while it is already on. If this happens, the previous table address is remembered and the previous table is resumed when the new one is finished. This nesting can reach a depth of up to 16: if stream 3 is opened for a seventeenth time, the interpreter should halt with an error message.

7.1.2.2

Output stream 3 is unusual in that, while it is selected, no text is sent to any other output streams which are selected. (However, they remain selected.)

7.1.2.2.1

Newlines are written to output stream 3 as ZSCII 13. (A game should never **print_char** the value 10, or any other value which is undefined as a ZSCII output code.)

7.1.2.3

Output stream 4 is unusual in that, when it is selected, the only text printed to it is that of the player's commands and keypresses (as read by **read_char**). (Each command is written, in one go, when it has been finished. Time delays and mouse-clicks should ideally be recorded. For suggestions on how this might be achieved, see the remarks section below. Mistypes and uses of 'delete' are not written.)

7.2 Buffering

On output streams 1 and 2 (only), text printing may be "buffered" in that new-lines are automatically printed to ensure that no word (of length less than the width of the screen) spreads across two lines (if the interpreter is able to control this). (This process is sometimes called "word-wrapping".)

7.2.1

In Versions 1 to 3, buffering is always on. In Versions 4 and later it is on by default (at the start of a game) and a game can switch it on or off using the **buffer_mode** opcode.

7.2.2

In Version 6, each of the eight windows has its own "buffering flag". In Versions 3 to 5, the **buffer_mode** applies only to the lower window, and buffering never happens in the upper window.

7.3 Selection (V1 and V2)

In Versions 1 and 2, output stream 1 is always selected and stream 2 can be selected or deselected by the game, by setting or clearing bit 0 of 'Flags 2'.

7.4 Selection (later versions)

In Versions 3 and later, all four output streams can be selected or deselected using the **output_stream** opcode. In addition, stream 2 can be selected or deselected by setting or clearing bit 0 of 'Flags 2'. Whichever method is used, the interpreter must ensure that this flag holds the current status of stream 2. ('A Mind Forever Voyaging' requires this.)

7.5 Dealing with Unicode or invalid characters

***[1.0] Because of the **print_unicode** opcode, it is possible for arbitrary Unicode characters to be sent to the output streams: that is, for characters which are not in the ZSCII set at all, even in the "extra characters" range.

7.5.1

See S 3.8.5.4 for rules on printing Unicode to stream 1.

7.5.2

Interpreters are free to use any representation of non-ASCII Unicode characters in stream 2. For example, they might print "[1a05]" to signify Unicode character **\$1a05**; or they might be configurable to write transcript files which conform to any chosen ISO 8859 set.

7.5.3

When printed to stream 3, Unicode characters should be converted to ZSCII if possible. If this is not possible, a question mark should be printed to stream 3.

7.5.4

Non-ZSCII characters never need to be printed to stream 4.

7.6 File handling

***[1.0] In Versions 5 and later, the Z-machine has the ability to load and save files (using optional operands with the **save** and **restore** opcodes: these operands were not used in Infocom's Version 5 games, but I wish to specify them as in Version 5 anyway).

7.6.1

***[1.0] Filenames have the following format (approximately the MS-DOS 8.3 rule): one to eight alphanumeric characters, a full stop and zero to three alphanumeric characters (the "file extension").

7.6.1.1

The interpreter must convert all filenames to upper case before use. If no full stop is given, ".AUX" should be appended.

7.6.1.2

Games should avoid the extensions ".INF", ".H", ".Z" followed by a number or ".SAV": otherwise they may be in danger of erasing their own object code, source code or saved game files.

7.6.1.3

***[1.1] The interpreter should delete from the filename any characters illegal for a filename. This will include all of the following characters (and more, if the OS requires it): slash, backslash, angle brackets (less-than and greater-than), colon, double-quote, pipe (vertical bar), question-mark, asterisk. The library should also truncate the argument at the first full stop (delete the first full stop and any following characters). If the result is the empty string, change it to the string "NULL".

7.6.2

***[1.0] Saved files are not associated with any particular session of a game. They are not part of the "state of play".

7.6.3

***[1.0] A game may depend on having up to 32 auxiliary files (with different names).

7.6.4

File-handling errors such as "disc corrupt" and "disc full" should be reported directly to the player by the interpreter. The error "file not found" should only cause a failure return code from **restore**.

7.6.5

Interpreters are allowed to not support access to external files (such as with `output_stream 2`, or the extra features of save and restore), or to only support some methods of access. Interpreters should support these features if possible, as some games may rely on external files, and in any case transcripts are very useful for testing, but in some environments such access is not feasible.

7.6.5.1

An attempt by the game to use save or restore in a manner not supported by the interpreter should simply return 0 as with any failure, and the game should notice and take appropriate actions.

7.6.5.2

An attempt by the game to use streams to access external files which is not supported by the interpreter should ideally print a warning to the user that the functionality is not available, and otherwise do nothing.

Remarks

The **ITF** interpreter incorrectly applies buffering when printing to the upper window.

Note that the requirement 7.1.2.1.1, that usages of stream 3 can be 'nested', is new in Standard 1.0. This is potentially important for Inform games, as stream 3 is often used to examine text before printing, for instance to choose between the articles "a" and "an" in front of an object name. But the process of printing an object name may itself require a usage of stream 3, and so on.

An ambiguous point about output stream 4 is whether it should contain the answers to interpreter questions like "what file name should your saved game have?": it can actually be quite useful to be able to include such answers in test script files. (When running a long script, I often save the game at several places during it, in order to save time in re-running passages.)

An interpreter should be able to write time delays (for timed input), accented characters or mouse clicks into stream 4 (i.e., to a script file). One possible style to record this information might be:

| | |
|-------------------------------------|--|
| take lamp | an ordinary command |
| turn it on.[154] | command, full stop, then keypad 9 (which might abbreviate for NE) |
| look unde[0] | timed out input |
| look under the rock [254][10][6] | the same input continuing mouse-click at (10,6) |

A typical auxiliary file might be one containing the player's preferred choices. This would be created when he first changed any of the default settings, and loaded (if present) whenever the game started up.

8. *The screen model*

8.1 Fonts

Text may be printed in any font of the interpreter's choice, variable- or fixed-pitch: except that when bit 1 of 'Flags 2' in the header is set, font 4 has been selected, or when the text style has been set to Fixed Pitch, then a fixed-pitch font must be used. The fixed pitch-bit is deprecated in Version 5 and later, and interpreter support in Version 6 is optional.

8.1.1

In Version 5, the height and width of the current font (in units (see below)) should be written to bytes \$27 and \$26 of the header, respectively. In Version 6, these bytes are the other way round (height in \$27, width in \$26). The width of a font is defined as the width of its '0' character.

8.1.2

An interpreter should ideally provide 4 fonts, with ID numbers as follows:

- 1: the normal font
- 2: a picture font
- 3: a character graphics font
- 4: a Courier-style font with fixed pitch

(Selecting font ID 0 is legal, and does not change the current font.) Ideally all text styles should be available for each font (for instance, Courier bold should be obtainable) except that font 3 need only be available in Roman and Reverse Video. Each font should provide characters for character codes 32 to 126 (plus character codes for any accented characters with codes greater than 127 which are being implemented as single accented letters on-screen).

8.1.3

***[1.0] A game must not use fonts other than 1 unless allowed to by the interpreter: see the `set_font` opcode for how to give or refuse permission. (This paragraph is marked ***[1.0] because existing Infocom games determined the availability of font 3 for 'Beyond Zork' in a complicated and unsatisfactory way: see S 16.)

8.1.3.1

***[1.0] It is legal for a game to change font at any time, including halfway through the printing of a word. (This might be needed to introduce exotic foreign accents in the future.)

8.1.4

The specification of the "picture font" is unknown (conjecturally, it was intended to provide pictures before Version 6 was properly developed). Interpreters should not implement it, and should refuse permission to any game requesting it.

8.1.5

The specification of the character graphics font is given in S 16.

8.1.5.1

In Version 5 (only), an interpreter which cannot provide the character graphics font should clear bit 3 of 'Flags 2' in the header.

8.1.6

Any new fonts will have numbers higher than 4. Fonts 5-1023 are reserved for future Standards to specify. Local use may be made of higher font numbers.

8.2 Status line

In Versions 1 to 3, a status line should be printed by the interpreter, as follows. In Version 3, it must set bit 4 of 'Flags 1' in the header if it is unable to produce a status line.

8.2.1

In Versions 1 and 2, all games are "score games". In Version 3, if bit 1 of 'Flags 1' is clear then the game is a "score game"; if it is set, then the game is a "time game".

8.2.2

The short name of the object whose number is in the first global variable should be printed on the left hand side of the line.

8.2.2.1

Whenever the status line is being printed the first global must contain a valid object number. (It would be useful if interpreters could protect themselves in case the game accidentally violates this requirement.)

8.2.2.2

If the object's short name exceeds the available room on the status line, the author suggests that an interpreter should break it at the last space and append an ellipsis "...". There is no guaranteed maximum length for location names but an interpreter should expect names of length up to at least 49 characters.

8.2.3

If there is room, the right hand side of the status line should display:

8.2.3.1

For "score games": the score and number of turns, held in the values of the second and third global variables respectively. The score may be assumed to be in the range -99 to 999 inclusive, and the turn number in the range 0 to 9999.

8.2.3.2

For "time games": the time, in the form **hours:minutes** (held in the second and third globals). The time may be given on a 24-hour clock or the number of hours may be reduced modulo 12 (but if so, "AM" or "PM" should be appended). Either way the player should be able to see the difference between 4am and 4pm, for example. The hours global may be assumed to be in the range 0 to 23 and the minutes global in the range 0 to 59.

8.2.4

The status line is updated in exactly two circumstances: when a **show_status** opcode is executed, and just before the keyboard is read by **read**. (It is not displayed when the game begins.)

8.3 Text colours

Under Versions 5 and later, text printing has a current foreground and background colour. In Version 6, each window has its own pair. (Note that a Version 6 interpreter going under the Amiga interpreter number must use the same pair of colours for all windows when running Infocom's games. If either is changed, then the interpreter must change the colour of all text on the screen to match. This simulates the Amiga hardware, which used two logical colours for text and switched palette to change their physical colour. This behaviour should not occur when running non-Infocom games, and modern games should never expect it. An interpreter that does not wish to handle this behaviour at all should avoid using the Amiga interpreter number when running Infocom's Version 6 games.)

8.3.1

The following codes are used to refer to colours:

```
-1 = the colour of the pixel under the cursor (if any) (true -3)
    [V6 only]
0 = current      (true -2)
1 = default     (true -1)
2 = black       (true $0000, $$0000000000000000)
3 = red         (true $001D, $$00000000000011101)
4 = green       (true $0340, $$0000001101000000)
5 = yellow      (true $03BD, $$0000001110111101)
6 = blue        (true $59A0, $$0101100110100000)
7 = magenta     (true $7C1F, $$0111110000011111)
```

| | |
|------------------|-------------------------------------|
| 8 = cyan | (true \$77A0, \$\$0111011110100000) |
| 9 = white | (true \$7FFF, \$\$0111111111111111) |
| 10 = light grey | (true \$5AD6, \$\$0101101011010110) |
| 11 = medium grey | (true \$4631, \$\$0100011000110001) |
| 12 = dark grey | (true \$2D6B, \$\$0010110101101011) |
| 13 | reserved |
| 14 | reserved |
| 15 = transparent | (true -4) [V6 only] |

Colours 10, 11, 12, 15 and -1 are available only in Version 6.

8.3.1.1

***[1.1] The equivalences between the colour numbers and true colours are recommended. The interpreter may allow the user to change the mapping, but the given values should be the default. If necessary, the game can check what true colour is being used for a given colour number using window properties 17 and 18.

Interpreters may provide different colours (eg making colour 10 dark grey), but if and only if they can detect they are running an original Infocom story file.

8.3.2

If the interpreter cannot produce colours, it should clear bit 0 of 'Flags 1' in the header. In Version 6 it should write colours 2 and 9 (black and white), either way round, into the default background and foreground colours in bytes **\$2c** and **\$2d** of the header.

8.3.3

If the interpreter can produce colours, it should set bit 0 of 'Flags 1' in the header, and write its default background and foreground colours into bytes **\$2c** and **\$2d** of the header.

8.3.4

If a game wishes to use colours, it should have bit 6 in 'Flags 2' set in its story file. (However, an interpreter should not rule out the use of colours just because this has not been done.)

8.3.5

If a true colour, or an "under the cursor" colour has been requested by the game, then the foreground or background colour shown in window property 11 is implementation defined, with two exceptions:

8.3.5.1

If the colour selected was one of the standard set (2-15), then that colour is indicated in property 11.

8.3.5.2

If the colour selected was not one of the standard set (this can happen when using graphics, which may use many more colours), the colour shown in property 11 will be ≥ 16 .

8.3.6

***[1.1] In Version 6 only, colour 15 is defined as transparent. This is only valid as a background colour; an attempt to select it for the foreground should produce a diagnostic. Interpreters not supporting transparency must ignore any attempt to select colour 15.

If the current background colour is transparent, then printed text is superimposed on the current window contents, without filling the background behind the text. **erase_window**, **erase_line** and **erase_picture** become null operations. The intent is to make it possible to superimpose text on non-uniform images. Up until now, only uniform images could be satisfactorily written on by sampling the background colour - that in itself would be problematical if the interpreter used dithering.

Scrolling with the background set to transparent is not permitted, so transparent should only be requested in a non-scrolling window. It is not valid to use Reverse Video style with the background set to transparent. Instructions that prompt for user input, such as **read** and **save**, should be avoided when the background is set to transparent, as it will not generally be possible for text entry to take place satisfactorily in the absence of a defined background colour. Printing text multiple times on top itself with the background set to transparent should be avoided, as the interpreter may use anti-aliasing, resulting in the text getting progressively heavier.

8.3.7

***[1.1] Standard 1.1 adds the ability for games to select many more colours with **set_true_colour**, which uses 15-bit RGB colour values, with the following special values:

- (-1) = default setting
- (-2) = current setting
- (-3) = colour under cursor (V6 only)
- (-4) = transparent (V6 only)

8.3.7.1

***[1.1] The interpreter selects the closest approximations available to the requested colours. In V6, the interpreter may store the approximations in window properties 16 and 17, so the program can tell how close it got (although it is acceptable for the interpreter to just store the requested value).

In the minimal implementation, interpreters just need to match to the closest of the standard colours and internally call **set_colour** (although that would have to ensure window properties 16 and 17 were updated). In a full implementation this would be turned around and **set_colour** would internally call **set_true_colour**.

True colour specifications are in the sRGB colour space, **\$0000** being black and **\$7FFF** being white. Colours should be gamma adjusted if necessary. See the **PNG** specification for a good introduction to colour spaces and gamma correction.

8.4 Screen dimensions

The screen should ideally be at least 60 characters wide by 14 lines deep. (Old Apple II interpreters had a 40 character width and some modern laptop ones have a 9 line height, but implementors should seek to avoid these extremes if possible.) The interpreter may change the exact dimensions whenever it likes but must write the current height (in lines) and width (in characters) into bytes \$20 and \$21 in the header.

8.4.1

The interpreter should use the screen height for calculating when to pause and print "[MORE]". A screen height of 255 lines means "infinite height", in which case the interpreter should never stop printing for a "[MORE]" prompt. (In case, say, the screen is actually a teletype printer, or has very good "scrollback".)

8.4.2

Screen dimensions are measured in notional "units". In Versions 1 to 4, one unit is simply the height or width of one character. In Version 5 and later, the interpreter is free to implement units as anything from character sizes down to individual pixels.

8.4.3

In Version 5 and later, the screen's width and height in units should be written to the words at \$22 and \$24.

8.5 Screen model (V1, V2)

The screen model for Versions 1 and 2 is as follows:

8.5.1

The screen can only be printed to (like a teletype) and there is no control of the cursor.

8.5.2

At the start of a game, the screen should be cleared and the text cursor placed at the bottom left (so that text scrolls upwards as the game gets under way).

8.6 Screen model (V3)

The screen model for Version 3 is as follows:

8.6.1

The screen is divided into a lower and an upper window and at any given time one of these is selected. (Initially it is the lower window.) The game uses the **set_window** opcode to select one of the two. Each window has its own cursor position at which text is printed. Operations in the upper window do not move the cursor of the lower. Whenever the upper window is selected, its cursor position is reset to the top left. Selecting, or re-sizing, the upper window does not change the screen's appearance.

8.6.1.1

The upper window has variable height (of n lines) and the same width as the screen. This should be displayed on the n lines of the screen below the top one (which continues to hold the status line). Initially the upper window has height 0. When the lower window is selected, the game can split off an upper window of any chosen size by using the **split_window** opcode.

8.6.1.1.1

Printing onto the upper window overlays whatever text is already there.

8.6.1.1.2

When a screen split takes place in Version 3, the upper window is cleared.

8.6.1.2

An interpreter need not provide the upper window at all. If it is going to do so, it should set bit 5 of 'Flags 1' in the header to signal this to the game. It is only legal for a game to use **set_window** or **split_window** if this bit has been set.

8.6.1.3

Following a "restore" of the game, the interpreter should automatically collapse the upper window to size 0.

8.6.2

When text reaches the bottom right of the lower window, it should be scrolled upwards. The upper window should never be scrolled: it is legal for a character to be printed on the bottom right position of the upper window (but the position of the cursor after this operation is undefined: the author suggests that it stay put).

8.6.3

At the start of a game, the screen should be cleared and the text cursor placed at the bottom left (so that text scrolls upwards as the game gets under way).

8.7 Screen model (V4, V5)

The screen model for Versions 4 and later, except Version 6, is as follows:

8.7.1

Text can be printed in five different styles (modelled on the VT100 design of terminal). These are: Roman (the default), Bold, Italic, Reverse Video (usually printed with foreground and background colours reversed) and Fixed Pitch. The specification does not require the interpreter to be able to display more than one of these at once (e.g. to combine italic and bold), and most interpreters can't. If the interpreter is going to allow certain combinations, then note that changing back to Roman should turn off all the text styles currently active.

8.7.1.1

An interpreter need not provide Bold or Italic (even for font 1) and is free to interpret them broadly. (For example, rendering bold-face by changing the colour, or rendering italic with underlining.)

8.7.1.2

It is legal to change text style at any point, including in the middle of a word being printed.

8.7.1.3

***[1.1] Although a story file can determine which individual styles are available by inspecting the header, this gives no indication of which styles can be combined. To improve this situation, at least for Version 6, Standard 1.1 requires window property 10 to show the actual style combination currently in use; with this a story file can probe for the availability of particular combinations.

8.7.2

There are two "windows", called "upper" and "lower": at any given time one of these two is selected. (Initially it is the lower window.) The game uses the **set_window** opcode to select one of the two. Each window has its own cursor position at which text is printed. Operations in the upper window do not move the cursor of the lower. Whenever the upper window is selected, its cursor position is reset to the top left.

8.7.2.1

The upper window has variable height (of n lines) and the same width as the screen. (It is usual for interpreters to print the upper window on the top n lines of the screen, overlaying any text which is already there, having been printed in the lower window some time ago.) Initially the upper window has height 0. When the lower window is selected, the game can split off an upper window of any chosen size by using the **split_window** opcode.

8.7.2.1.1

It is unclear exactly what **split_window** should do if the upper window is currently selected. The author suggests that it should work as usual, leaving the cursor where it is if the cursor is still inside the new upper window, and otherwise moving the cursor back to the top left. (This is analogous to the Version 6 practice.)

8.7.2.2

In Version 4, the lower window's cursor is always on the bottom screen line. In Version 5 it can be at any line which is not underneath the upper window. If a split takes place which would cause the upper window to swallow the lower window's cursor position, the interpreter should move the lower window's cursor down to the line just below the upper window's new size.

8.7.2.3

When the upper window is selected, its cursor position can be moved with **set_cursor**. This position is given in characters in the form (row, column), with (1,1) at the top left. The opcode has no effect when the lower window is selected. It is illegal to move the cursor outside the current size of the upper window.

8.7.2.4

An interpreter should use a fixed-pitch font when printing on the upper window.

8.7.2.5

In Versions 3 to 5, text buffering is never active in the upper window (even if a game begins printing there without having turned it off).

8.7.3

Clearing regions of the screen:

8.7.3.1

When text reaches the bottom right of the lower window, it should be scrolled upwards. (When the text style is Reverse Video the new blank line should **not** have reversed colours.) The upper window should never be scrolled: it is legal for a character to be printed on the bottom right position of the upper window (but the position of the cursor after this operation is undefined: the author suggests that it stay put).

8.7.3.2

Using the opcode **erase_window**, the specified window can be cleared to background colour. (Even if the text style is Reverse Video the new blank space should not have reversed colours.)

8.7.3.2.1

In Versions 5 and later, the cursor for the window being erased should be moved to the top left. In Version 4, the lower window's cursor moves to its bottom left, while the upper window's cursor moves to top left.

8.7.3.3

Erasing window -1 clears the whole screen to the background colour of the lower screen, collapses the upper window to height 0, moves the cursor of the lower screen to bottom left (in Version 4) or top left (in Versions 5 and later) and selects the lower screen. The same operation should happen at the start of a game.

8.7.3.4

Using **erase_line** in the upper window should erase the current line from the cursor position to the right-hand edge, clearing it to background colour. (Even if the text style is Reverse

Video the new blank space should not have reversed colours.)

8.8 Screen model (V6)

The screen model for Version 6 is as follows:

8.8.1

The display is an array of pixels. Coordinates are usually given (in units) in the form (y,x), with (1,1) in the top left.

8.8.2

If the interpreter thinks the screen should be redrawn (e.g. because a menu window has been clicked over it), it may set bit 2 of 'Flags 2'. The game is expected to notice, take action and clear the bit. (However, a more efficient interpreter would handle redraws itself.)

8.8.3

There are eight "windows", numbered 0 to 7. The code -3 is used as a window number to mean "the currently selected window". This selection can be changed with the **set_window** opcode. Windows are invisible and usually lie on top of each other. All text and graphics plotting is always clipped to the current window, and anything showing through is plotted onto the screen. Subsequent movements of the window do not move what was printed and there is no sense in which characters or graphics 'belong' to any particular window once printed. Each window has a position (in units), a size (in units), a cursor position within it (in units, relative to its own origin), a number of flags called "attributes" and a number of variables called "properties".

8.8.3.1

There are four attributes, numbered as follows:

- 0: wrapping
- 1: scrolling
- 2: text copied to output stream 2 (the transcript, if selected)
- 3: buffered printing

Each can be turned on or off, using the **window_style** opcode.

8.8.3.1.1

"Wrapping" is the continuation of printed text from one line to the next. Text running up to the right margin will continue from the left margin of the following line. If "wrapping" is off then characters will be printed until no more can be fitted in without hitting the right margin, at which point the cursor will move to the right margin and stay there, so that any further text will be ignored.

8.8.3.1.2

"Buffered printing" means that text to be printed in the window is temporarily stored in a buffer and only flushed onto the screen at intervals convenient for the interpreter.

8.8.3.1.2.1

"Buffered printing" has two practical effects: firstly it causes a delay before printed text actually appears.

8.8.3.1.2.2

Secondly it affects the way "wrapping" is done. If "buffered printing" is on, then text is wrapped after the last word which could fit on a line. If not, then text is wrapped after the last character that could fit.

Example: suppose the text "Here is an abacus" is printed in a narrow window. The appearance (after the buffer has been flushed, if there is buffered printing) might be:

| | | |
|-------------|---------------|------------------------|
| | | ...margins.... |
| wrapping on | buffering on | Here is an abacus^ |
| off | buffering on | Here is an aba^ |
| wrapping on | buffering off | Here is an aba cus^ |
| off | buffering off | Here is an aba^ |

where the caret denotes the final position of the cursor. (Games often alter "wrapping": it would normally be on for a window holding running text but off for a status-line window, which is why window 0 has "wrapping" on by default but all other windows have "wrapping" off by default. On the other hand all windows have "buffered printing" on by default and games only alter this in rare circumstances to avoid delays in the appearance of individual printed characters.)

8.8.3.2

There are 16 properties, numbered as follows:

| | |
|----------------|-----------------------|
| 0 y coordinate | 9 interrupt countdown |
| 1 x coordinate | 10 text style |

| | |
|-----------------------------|---------------------------|
| 2 y size | 11 colour data |
| 3 x size | 12 font number |
| 4 y cursor | 13 font size |
| 5 x cursor | 14 attributes |
| 6 left margin size | 15 line count |
| 7 right margin size | 16 true foreground colour |
| 8 newline interrupt routine | 17 true background colour |

Each property is a standard Z-machine number and is readable with **get_wind_prop**. Properties 0 through 15 are writeable with **put_wind_prop**. However, a game should only use **put_wind_prop** to set the newline interrupt routine, the interrupt countdown and the line count: everything else is either set by the interpreter or by specialised opcodes (such as **set_font**). The true foreground and true background properties must not be written by **put_wind_prop**.

8.8.3.2.1

If a window has character wrapping, then text is clipped to stay inside the left and right margins. After a new-line, the cursor moves to the left margin on the next line. Margins can be set with **set_margins** but this should only be done just after a newline or just after the window has been selected. (These values are margin sizes in pixels, and are by default 0.)

8.8.3.2.2

If the interrupt countdown is set to a non-zero value (which by default it is not), then the line count is decremented on each new-line, and when it hits zero the routine whose packed address is stored in the "newline interrupt routine" property is called before text printing resumes. (This routine may, for example, meddle with margins to roll text around a crinkly-shaped picture.) The interrupt routine should not attempt to print anything.

8.8.3.2.2.1

Because of an Infocom bug, if the interpreter number is 6 (for MSDOS) and the story file is 'Zork Zero' release 393.890714, but in no other case, the interpreter must do the following instead: (1) move to the new line, (2) put the cursor at the current left margin, (3) call the interrupt routine (if it's time to do so). This is the least bad way to get around a basic inconsistency in existing Infocom story files and interpreters.

8.8.3.2.2.2

Note that the **set_margins** opcode, which is often used by newline interrupt routines (to adjust the shape of a margin as it flows past a picture), automatically moves the cursor if the change in margins would leave the cursor outside them. The effect will depend, unfortunately, on which sequence of events above takes place.

8.8.3.2.2.3

A line count is never decremented below -999.

8.8.3.2.3

The text style is set just as in Version 4, using **set_text_style** (which sets that for the current window). The property holds the operand of that instruction (e.g. 4 for italic).

8.8.3.2.4

The foreground colour is stored in the lower byte of the colour data property, the background colour in the upper byte.

8.8.3.2.5

The font height (in pixels) is stored in the upper byte of the font size property, the font width (in pixels) in the lower byte.

8.8.3.2.6

The interpreter should use the line count to see when it should print "[MORE]". A line count of -999 means "never print [MORE]". (Version 6 games often set line counts to manipulate when "[MORE]" is printed.)

8.8.3.2.7

If an attempt is made by the game to read the cursor position at a time when text is held unprinted in a buffer, then this text should be flushed first, to ensure that the cursor position is accurate before being read.

8.8.3.2.8

***[1.1] The true foreground and background colours show the actual colour being used for the foreground and background, whether it was set using **set_colour** or **set_true_colour**. Transparent is indicated as -4. If the colour was sampled from a picture then the value shown may be a 15-bit rounding of a more precise colour, leading to a slight inaccuracy if the colour is read and then written back.

8.8.3.3

All eight windows begin at (1,1). Window 0 occupies the whole screen and is initially selected. Window 1 is as wide as the screen but has zero height. Windows 2 to 7 have zero width and height. Window 0 initially has attribute 1 off and 2, 3 and 4 on (scrolling, copy to printer transcript, buffering). Windows 1 to 7 initially have attribute 4 (buffering) on, and the other attributes off.

8.8.3.4

A window can be moved with **move_window** and resized with **window_size**. If the window size is reduced so that its cursor lies outside it, the cursor should be reset to the left margin on the top line.

8.8.3.5

Each window remembers its own cursor position (relative to its own coordinates, so that the position (1,1) is at its top left). These can be changed using **set_cursor** (and it is legal to move the cursor for an unselected window). It is illegal to move the cursor outside the current window.

8.8.3.6

Each window can be scrolled vertically (up or down) any number of pixels, using the **scroll_window** opcode.

8.8.4

To some extent windows 0 and 1 mimic the behaviour of the lower and upper windows in the Version 4 screen model:

8.8.4.1

The **split_screen** opcode tiles windows 0 and 1 together to fill the screen, so that window 1 has the given height and is placed at the top left, while window 0 is placed just below it (with its height suitably shortened, possibly making it disappear altogether if window 1 occupies the whole screen).

8.8.4.2

An "unsplit" (that is, a **split_screen 0**) takes place when the entire screen is cleared with **erase_window -1**, if a "split" has previously occurred (meaning that windows 0 and 1 have been set up as above).

8.8.5

Screen clearing operations:

8.8.5.1

Erasing a picture is like drawing it (see below), except that the space where it would appear is painted over with background colour instead.

8.8.5.2

The current line can be erased using **erase_line**, either all the way to the right margin or by any positive number of pixels in that direction. The space is painted over with background colour (even if the current text style is Reverse Video).

8.8.5.3

Each window can be erased using **erase_window**, erasing to background colour (even if the current text style is Reverse Video).

8.8.5.3.1

Erasing window number -1 erases the entire screen to the background colour of window 0, unplits windows 0 and 1 (see S 8.7.3.3 above) and selects window 0.

8.8.5.3.2

Erasing window -2 erases the entire screen to the current background colour. (It doesn't perform **erase_window** for all the individual windows, and it doesn't change any window attributes or cursor positions.)

8.8.6

Pictures may accompany the game. They are not stored in the story file (or the Z-machine) itself, and the interpreter is simply expected to know where to find them.

8.8.6.1

Pictures are numbered from 1 upwards (not necessarily contiguously). They can be "drawn" or "erased" (using **draw_picture** and **erase_picture**). Before attempting to do so, a game may ask the interpreter about the picture (using **picture_data**): this allows the interpreter to signal that the picture in question is unavailable, or to specify its height and width.

8.8.6.2

The game may, if it wishes, use the **picture_table** opcode to give the interpreter advance warning that a group of pictures will soon be needed (for instance, a collection of icons making up a control panel). The interpreter may want to load these pictures off disc and into a memory cache.

8.8.7

***[1.1] Interpreters may use a backing store to store the Z-machine screen state, rather than plotting directly to the screen. This would normally be the case in a windowed operating system environment. If a backing store is in use, display changes executed by the Z-machine may not be immediately made visible to the user. Standard 1.1 adds the new opcode **buffer_screen** to Version 6 to control screen updates. An interpreter is free to ignore the opcode if it doesn't fit its display model (in which case it must act as if **buffer_screen** is always set to 0).

8.8.7.1

***[1.1] When **buffer_screen** is set to 0 (the default), all display changes are expected to become visible to the user either immediately, or within a short period of time, at the interpreter's discretion. At a minimum, all updates become visible before waiting for input. Any intermediate display states between input requests may not be seen; for example when printing a large amount of new text into a scrolling window, all the intermediate scroll positions may or may not be shown.

When **buffer_screen** is set to 1, the interpreter need not change the visible display at all. Any display changes can be done purely in the backing store. A program may set **buffer_screen** to 1 before carrying out a complex layered graphical composition, to indicate that the intermediate states are not worth showing. It would be extremely ill-advised to prompt for input with **buffer_screen** set to 1.

When **buffer_screen** is set back to 0, the display is not necessarily updated immediately. If this is required, the game must request it separately (see S 8.8.7.2 below).

8.8.7.2

***[1.1] With **buffer_screen** in either state, an update of the visible display can be forced immediately by issuing **buffer_screen -1**, without altering the current buffering state. Note that **buffer_screen -1** does not flush the text buffer.

Remarks

See S 16 for comment on how 'Beyond Zork' uses fonts.

Some interpreters print the status line when they begin running a Version 3 game, but this is incorrect. (It means that a small game printing text and then quitting cannot be run unless it includes an object.) The author's preferred status line formats are:

| | |
|------------------|----------|
| Hall of Mists | 80/733 |
| Lincoln Memorial | 12:03 PM |

Thus the score/turns block always fits in $3+1+4=8$ characters and the time in $2+1+2+1+2=8$ characters. (Games needing more exotic time lines, for example, should not be written in Version 3.)

The only existing Version 3 game to use an upper window is 'Seastalker' (for its sonarscope display).

Some ports of **ITF** apply buffering (i.e. word-wrapping) and scrolling to the upper window, with unfortunate consequences. This is why the standard Inform status line is one character short of the width of the screen.

The original Infocom files seldom use **erase_window**, except with window -1 (for instance 'Trinity' only uses it in this form). **ITF** does not implement it in any other case.

The Version 5 re-releases of older games make use of consecutive **set_text_style** instructions to attempt to combine boldface reverse video (in the hints system).

None of Infocom's Version 4 or 5 files use **erase_line** at all, and **ITF** implements it badly (with unpredictable behaviour in Reverse Video text style). (It's interesting to note that the Version 5 edition of 'Zork I' - one of the earliest Version 5 files -- blanks out lines by looking up the screen width and printing that many spaces.)

It's recommended that a Version 5 interpreter always use units to correspond to characters: that is, characters occupy 1×1 units. 'Beyond Zork' was written in the expectation that it could be using either 1×1 or 8×8 , and contains correct code to calculate screen

positions whatever units are used. (Infocom's Version 5 interpreter for MSDOS could either run in a text mode, 1x1, or a graphics mode, 8x8.) However, the German translation of 'Zork I' contains incorrect code to calculate screen positions unless 1x1 units are used.

Note that a minor bug in **Zip** writes bytes \$22 to \$25 in the header as four values, giving the screen dimensions in the form left, right, top, bottom: provided units are characters (i.e. provided the font width and height are both 1) then since "left" and "top" are both 0, this bug has no effect.

Some details of the known IBM graphics files are given in Paul David Doherty's "Infocom Fact Sheet". See also Mark Howell's program "pix2gif", which extracts pictures to GIF files. (This is one of his "Ztools" programs.)

Although Version 6 graphics files are not specified here, and were released in several different formats by Infocom for different computers, a consensus seems to have emerged that the MCGA pictures are the ones to adopt (files with filenames *.MG1). These are visually identical to Amiga pictures (whose format has been deciphered by Mark Knibbs). However, some Version 6 story files were tailored to the interpreters they would run on, and use the pictures differently according to what they expect the pictures to be. (For instance, an Amiga-intended story file will use one big Amiga-format picture where an MSDOS-intended story file will use several smaller MCGA ones.)

The easiest option is to interpret only DOS-intended Version 6 story files and only MCGA pictures. But it may be helpful to examine the **Frotz** source code, as **Frotz** implements **draw_picture** and **picture_data** so that Amiga and Macintosh forms of Version 6 story files can also be used.

It is generally felt that newly-written graphical games should not imitate the old Infocom graphics formats, which are very awkward to construct and have been overtaken by technology. Instead, the **Blorb** proposal for packaging up resources with Z-machine games calls for PNG format graphics glued together in a fairly simple way. The graphics for Infocom's Version 6 games have been made available in **Blorb** format, so that understanding Infocom's picture-sets is no longer very useful.

The line count of -999 preventing "[MORE]" is a device used by the demonstration mode of 'Zork Zero'.

Interpreter authors are advised that all 8 windows in Version 6 must be treated identically. The only ways in which they are distinguished are:

- ◆ Different default positions + sizes
- ◆ Different default attributes
- ◆ **split_window** manipulates windows 0 and 1 specifically
- ◆ Window 1 is the default mouse window

Differences in interpreter behaviour must only arise from differences in window attributes and properties.

In V6, it is legal to position the cursor up against the right or bottom of a window - eg at (1,1) in a zero-sized window or at (641,401) in 640x400 window. Indeed, this is the default state of windows 1 to 7, and the cursor may be left at the right-hand side of a window when wrapping is off.

Attempting to print text (including new-lines) when the cursor is fewer than `font_height` units from the bottom of the window results in undefined behaviour - this precludes any printing in windows less than `font_height` units high.

It is legal for interpreters to always show the same value in property 11 if a true or sampled colour is in use. As a result, story files cannot assume that setting a value that was read from property 11 will give the same colour, if `set_colour -1` has been used in that window.

The same rules apply if an interpreter offers non-standard default colours although in this case it would be ill-advised to show the same colour numbers for foreground and background - unless they can be distinguished, non-standard default colours should probably not be offered.

If the interpreter offers a limited palette, then there is no problem, as it can be arranged for there to be fewer than 240 distinct non-standard colours. In an interpreter with a higher colour-depth, a good implementation would be to use colours 16-255 to represent the last 240 distinct non-standard colours used, re-using numbers after 240 colours have been used. This will minimize potential problems caused by non-standard colours, particularly when set as defaults.

Regardless of the limitations on colour numbers, in Version 6 each window must remember accurately the colour pair selected, so it is preserved across window switches.

S 8.7.2.3 states that it is illegal to move the cursor outside the current size of the upper window. S 8.8.3.5 gives the equivalent rule for Version 6.

Many modern games have been lax in obeying this rule; in particular some of the standard Inform menu libraries have violated it. Infocom's Sherlock also miscalculated the size of the upper window to use for box quotes.

It is recommended that if the cursor is moved below the split position in V4/V5, interpreters should execute an implicit "split_window" to contain the requested cursor position, if possible. Diagnostics should be produced, but should be suppressable.

Some modern Z-Machine interpreters (mainly those using Andrew Plotkin's **Glk** interface standard) use a separate text windows for the status line. While this is not Standard behaviour, it largely causes no problems. However Trinity, and many more recent Inform games, print quote boxes using a technique that is not compatible with this implementation.

Andrew Plotkin has written up some notes on the issue, including a workaround.

Infocom's Version 6 interpreters and story files disagree on the meaning of window attributes 0 and 3 and the opcode **buffer_mode**, in such a way that the original specification is hard to deduce from the final behaviour. If we call the three possible ways that text can appear "word wrap", "char wrap" and "char clip":

```

                |...margins....|
word wrap      Here is an
                abacus^
char wrap      Here is an aba
                cus^
char clip      Here is an aba^

```

then Infocom's interpreters behave as follows:

| | Apple II | MSDOS | Macintosh | Amiga |
|-----------------|---------------|---------------|-----------|--------------|
| A0 off, A3 off | char clip(LR) | char clip() | --- | --- |
| A0 off, A3 on | char clip(LR) | char clip(LR) | --- | --- |
| A0 on, A3 off | word wrap | char wrap | --- | --- |
| A0 on, A3 on | word wrap | word wrap | --- | --- |
| buffer_mode off | --- | --- | char wrap | char clip(L) |
| buffer_mode on | --- | --- | word wrap | word wrap |

Here "---" means that the interpreter ignores the given state, and the presence of L, R or both after "char clip" indicates which of the left and right margins are respected. The Amiga behaviour may be due to a bug and two bugs have also been found in the MSDOS implementation. Under this standard, the appearance is as follows:

```

                Standard
A0 off, A3 off  char clip(LR)
A0 off, A3 on  char clip(LR)
A0 on,  A3 off  char wrap
A0 on,  A3 on  word wrap
buffer_mode off ---
buffer_mode on  ---

```

Due to a bug or an oversight, the V6 story files for all interpreters use **buffer_mode** once: to remove buffering while printing "Please wait..." with a row of full stops trickling out during a slow operation. Buffering would frustrate this, but fortunately on modern computers the operation is no longer slow and so the bug does not cause trouble.

9. Sound effects

9.1 Sound effects

Some games, from Version 3 onward, have sound effects attached. These are not stored in the story files (or the Z-machine) itself, and the interpreter is simply expected to know where to find them. Other games have only one sound effect, usable in a much more restricted way: a beep or bell sound, which we shall call a "bleep".

9.1.1

In Version 6, the interpreter should set bit 5 of 'Flags 1' if it can provide sound effects beyond a bleep.

9.1.2

In Version 5 and later, a game should have bit 7 of 'Flags 2' set in its story file if it wants to use sound effects beyond a bleep. The interpreter should then clear this bit if it cannot oblige.

9.2 Numbering of

Sound effects are numbered upwards from 1. Number 1 is a high-pitched bleep, number 2 a low-pitched one and effects from 3 upward are supplied by the interpreter somehow for the particular game in question.

9.2.1

***[1.1] Aside from bleeps, there are two types of sound effect, samples and music. The game has no way of telling the type of a given sound effect.

9.3 Volume

Sound effects (other than bleeps) can be played at any volume level from 1 to 8 (8 being loudest of these). The volume level -1 should be implemented as "loudest possible".

9.4 Sound playing autonomously

Bleeps are immediate and brief. Other sound effects take place in the background, while normal operation of the Z-machine is going on. Control is via the **sound_effect** opcode, allowing the game to prepare, start, stop or finish with an effect.

9.4.1

The game may (but need not) "prepare" a sound effect before use. This would indicate to the interpreter that the game intends to use the effect soon: an interpreter might act on this information by loading the sampled sound off disc and into a memory cache.

9.4.2

***[1.1] A sound effect (other than a bleep) can then be "stopped" or "started". Only one sound effect of each type can play at any given time, so that starting a new music sound effect stops any current music playing, and starting any new sample sound effect stops any current sample sound playing. Samples and music do not interrupt each other.

9.4.3

In Versions 5 and later, a sound effect may repeat any specified number of times, or repeat forever (until stopped).

9.4.4

Eventually, though, if it has not been stopped, it may end by itself. A routine (specified at start time) can then be called. The intention is that this routine may implement effects such as fading in and out, by replaying the sound effect at a different volume. (A game should not place any important code in such a routine.) The routine is only called when the sound has played the requested number of times. If manually stopped or interrupted by another sound, the routine is not called.

9.4.5

The game may, but need not, explicitly "finish with" any sound effect which is not likely to occur again for a while: the interpreter can then throw it out of memory.

Remarks

The safest way an Inform program can try to produce a bleep is by executing `@sound_effect 1`. Some ports of **Zip** believe that the first operand of this is the number of bleeps to make (so that `@sound_effect 2` bleeps twice), but this is incorrect.

Several Infocom games bleep (using `sound_effect` with only one operand, always equal to 1 or 2). Two provided sampled sound effects but did not bleep: 'The Lurking Horror' and 'Sherlock'. Their story files contain the following usages of `sound_effect`:

```
sound_effect number 2 volume                (in TLH)
sound_effect number 2 volume/repeats function (in Sherlock)
sound_effect 0 3
sound_effect number 3
sound_effect 0 4
```

except that, probably due to a bug in its own code, 'TLH' can also generate

```
sound_effect 4 8
sound_effect 4095 2 15
```

A further difficulty with 'TLH' is that it assumes the interpreter is as slow as Infocom's

Amiga interpreter was: it fires off several sound effects in one game round, assuming there will be time for it to play most of each one. To simulate this, **sound_effect** must be rewritten to pause sometimes:

if a new sample sound effect is begun while there is still one playing which was started since the last keyboard input, then wait until that earlier one finishes one cycle before replacing it with the new sound effect. Music sound effects are not affected by this. New music should interrupt old music immediately at all times.

Infocom's MS-DOS interpreters for V4 to V6 set bit 5 of 'Flags 1' in all circumstances (i.e., whether or not sound effects are available). This would be incorrect behaviour for a standard interpreter.

Infocom implemented sound effects differently on different machines. The format of Infocom's shipped sound effects files has been documented by Stefan Jokisch and his notes are available from www.ifarchive.org.

However, modern interpreters are strongly encouraged to support Andrew Plotkin's **Blorb** format, which is a more modern way to make sound effects available to newer games. Blorb files have been made available for the Infocom sound effects, so that modern interpreters need no longer support the Infocom format.

When using Blorb resources, the default interpreter behaviour (unless overridden by the player) should be for samples played at maximum volume (64), in one channel of a SONG or MOD played at volume 8, to be of equal volume to samples played at maximum volume (8) as an effect. This will be the natural behaviour if effects use one physical channel and MODs/SONGs use four physical channels.

Ideally, a sound played at volume n in a SONG played at volume m should sound the same as when played as an effect at volume $n*m/64$. This mandates that the volume scale for effects be equivalent to the scale defined for samples in the MOD specification.

If multi-channel effects are used, the overall volume should be independent of the number of channels used in the sound. Thus a stereo AIFF containing the same samples for left and right should sound as loud as a mono AIFF containing the same data. This will need adjustment of volume if stereo AIFFs use two physical channels and mono AIFFs use one. No adjustment would be required if an interpreter reduced all AIFFs to mono.

10. Input streams and devices

10.1 Keyboard only in V1

In Versions 1 and 2, the player's commands can only be drawn from the keyboard.

10.2 Input streams

In Versions 3 and later, the player's keypresses are drawn from the current "input stream". There are two input streams: numbered 0 (the keyboard) and 1 (a file containing commands). Other inputs (mouse clicks or menu selections), if available, are also implemented as keypresses (see below).

10.2.1

The format of a file containing commands must be the same as that written in output stream 4.

10.2.2

The game can change the current input stream itself, using the opcode **input_stream**. It has no way of finding out which input stream is currently in use. An interpreter is free to change the input stream whenever it likes (e.g. at the player's request) or, indeed, to run the entire game under input stream 1 (for testing purposes).

10.2.3

When input stream 1 is first selected, the interpreter may use any method of choosing a file name for the file of commands. (Good practice is to use the same conventions as when choosing a filename for output to stream 4.)

10.2.4

When the the current stream is stream 1, the interpreter should not hold up long passages of text (by printing "[MORE]" and waiting for a keypress, for instance).

10.3 Mouse support

Mouse support is optional but can be provided in Versions 5 and later.

10.3.1

In a game which wishes to use the mouse, bit 5 of 'Flags 2' in the header should be set in the story file. If it wishes to read the mouse position after clicks, it must provide at least the first two words of a header extension table. (Note that Inform 6.12 and later always provide a header extension table at least this large, but Inform 6.11 and earlier never provide an extension table at all.)

10.3.1.1

If the interpreter cannot offer mouse support, then it should clear bit 5 of 'Flags 2' to signal this to the game.

10.3.2

Whenever a mouse click takes place (and provided the header extension table exists and contains at least 2 words) the interpreter should update the coordinates as follows:

Word 1: x coordinate where click took place

Word 2: y coordinate where click took place

10.3.3

The mouse is presumed to have between 0 and 16 buttons. The state of these buttons can be read by the **read_mouse** opcode in Version 6. Otherwise, mouse clicks are treated as keyboard input codes (see below).

10.3.4

In Version 6, the mouse can either be free or constrained to one of the 8 windows: if so, clicks outside the 'mouse window' must be ignored, and the interpreter is at liberty to confine the mouse's movement to the boundary of its window.

10.4 Menu support

Menu support can optionally be provided in Version 6.

10.4.1

In a game which wishes to use menus, bit 8 of 'Flags 2' in the header should be set in the story file.

10.4.1.1

If the interpreter cannot offer menu support, then it should clear bit 8 of 'Flags 2' to signal this to the game.

10.4.2

Menus are numbered from 0 upwards. 0, 1 and 2 are reserved for the interpreter to manage (this system has only been implemented on the Macintosh, wherein 0 is the Apple menu, 1 the File menu and 2 the Edit menu). Menus numbered 3 and upwards can be created or removed with the **make_menu** opcode.

10.4.3

Menu selection is reported to the game as a keypress (see below). Details of what selection has been made are read with **read_mouse**.

10.5 Terminating characters and timed input

Whole commands are read from the input stream using the **read** opcode. (Note that this has two different internal names in Inform, **sread** for Versions 1 to 4 and **aread** subsequently.)

10.5.1

In Versions 1 to 3, the interpreter must redisplay the status line before it begins accepting input.

10.5.2

Commands are normally terminated by a new-line (a carriage return or a line feed as appropriate for the machine's keyboard or file format).

10.5.2.1

In Versions 5 and later, the game may provide a "terminating characters table" by giving its byte address in the word at **\$2e** in the header. This table is a zero-terminated list of input character codes which cause **aread** to finish the command (in addition to new-line). Only function key codes are permitted: these are defined as those between 129 and 154 inclusive, together with 252, 253 and 254. The special value 255 means "any function key code is terminating".

10.5.3

***[1.0] In Versions 4 and later, an interpreter should ideally be able to time input and to call a (game) routine at periodic intervals: see the **read** opcode. If it is able to do this, it should set bit 7 of 'Flags 1' in the header.

10.6 Single keypresses

In Versions 4 and later, individual characters can be read from the current input stream, using **read_char**. Again, the interpreter should ideally be able to time input and to call a (game) routine at periodic intervals. If it is able to do this, it should set bit 7 of 'Flags 1' in the header.

10.7 Reading ZSCII from the keyboard

The only characters which can be read from the keyboard are ZSCII characters defined for input (see S 3).

10.7.1

Every ZSCII character defined for input can be returned by **read_char**.

10.7.2

Only ZSCII characters defined for both input and output can be stored in the text buffer supplied to the **read** opcode.

10.7.3

The "escape" code is optional: that is, an interpreter need not provide an escape key. (The Inform library clears and quits menus if this code is returned to **read_char**.)

Remarks

Menus in 'Beyond Zork' define cursor up and cursor down as terminating characters, and make use of **read** in the upper window.

Mouse co-ordinates, whether returned by **read_mouse** or written into the header during input, are always relative to the top of the display at (1,1), regardless of the position of the current mouse window.

read_mouse is realtime. When called it must read the current mouse location, whether or not the mouse is inside the current mouse window. Interpreters are allowed to show positions and button states outside the Z-machine screen if the pointer is outside the interpreter's own user interface (using negative values if needed).

Programs must be prepared to cope with this. For example in a painting program you might want to ignore all buttons down outside the screen. When dragging something you might want to keep trying to follow the pointer, even outside the screen, until the buttons are released.

Interpreters may constrain the pointer to the screen as long as buttons are held down - this might aid dragging operations - although this is not required.

11. The format of the header

11.1 Header format

The header table summarises those locations in the Z-machine's header which an interpreter must deal with. (For further notes on traditional usage, see Appendix B.) "Hex" means the address, in hexadecimal; "V" the earliest Version to which the rule is applicable; "Dyn" means that the byte or bit may legally be changed by the game during play; "Int" means that the interpreter may change it; "Rst" means that the interpreter must set it correctly after loading the game, after a restore or after a restart.

Header format

| Hex | V | Dyn | Int | Rst | Cont |
|-----|---|-----|-----|-----|--|
| 0 | 1 | | | | Version number (1 to 6) |
| 1 | 3 | | | | <i>Flags 1 (in Versions 1 to 3):</i> |
| | | | | | Bit 1: Status line type: 0=score/turns, 1=hours:mins |
| | | | | | 2: Story file split across two discs? |
| | | | * | * | 4: Status line not available? |
| | | | * | * | 5: Screen-splitting available? |
| | | | * | * | 6: Is a variable-pitch font the default? |
| | 4 | | | | <i>Flags 1 (from Version 4):</i> |
| | 5 | | * | * | Bit 0: Colours available? |
| | 6 | | * | * | 1: Picture displaying available? |
| | 4 | | * | * | 2: Boldface available? |
| | 4 | | * | * | 3: Italic available? |
| | 4 | | * | * | 4: Fixed-space style available? |
| | 6 | | * | * | 5: Sound effects available? |
| | 4 | | * | * | 7: Timed keyboard input available? |
| 4 | 1 | | | | Base of high memory (byte address) |
| 6 | 1 | | | | Initial value of program counter (byte address) |
| | 6 | | | | Packed address of initial "main" routine |
| 8 | 1 | | | | Location of dictionary (byte address) |
| A | 1 | | | | Location of object table (byte address) |
| C | 1 | | | | Location of global variables table (byte address) |
| E | 1 | | | | Base of static memory (byte address) |
| 10 | 1 | | | | <i>Flags 2:</i> |
| | 1 | * | * | * | Bit 0: Set when transcribing is on |
| | 3 | * | | * | 1: Game sets to force printing in fixed-pitch font |
| | 6 | * | * | | 2: Int sets to request screen redraw: game clears |
| | 5 | | * | * | 3: If set, game wants to use pictures |

| | | | | | |
|----|----|--|---|---|--|
| | 5 | | * | * | 4: If set, game wants to use the UNDO opcodes |
| | 5 | | * | * | 5: If set, game wants to use a mouse |
| | 5 | | | | 6: If set, game wants to use colours |
| | 5 | | * | * | 7: If set, game wants to use sound effects |
| | 6 | | * | * | 8: If set, game wants to use menus |
| | | | | | (For bits 3,4,5,7 and 8, Int clears again if it cannot |
| 18 | 2 | | | | Location of abbreviations table (byte address) |
| 1A | 3+ | | | | Length of file (see note) |
| 1C | 3+ | | | | Checksum of file |
| 1E | 4 | | * | * | Interpreter number |
| 1F | 4 | | * | * | Interpreter version |
| 20 | 4 | | * | * | Screen height (lines): 255 means "infinite" |
| 21 | 4 | | * | * | Screen width (characters) |
| 22 | 5 | | * | * | Screen width in units |
| 24 | 5 | | * | * | Screen height in units |
| 26 | 5 | | * | * | Font width in units (defined as width of a '0') |
| | 6 | | * | * | Font height in units |
| 27 | 5 | | * | * | Font height in units |
| | 6 | | * | * | Font width in units (defined as width of a '0') |
| 28 | 6 | | | | Routines offset (divided by 8) |
| 2A | 6 | | | | Static strings offset (divided by 8) |
| 2C | 5 | | * | * | Default background colour |
| 2D | 5 | | * | * | Default foreground colour |
| 2E | 5 | | | | Address of terminating characters table (bytes) |
| 30 | 6 | | * | | Total width in pixels of text sent to output stream 3 |
| 32 | 1 | | * | * | Standard revision number |
| 34 | 5 | | | | Alphabet table address (bytes), or 0 for default |
| 36 | 5 | | | | Header extension table address (bytes) |

Some early Version 3 files do not contain length and checksum data, hence the notation 3+.

11.1.1

It is illegal for a game to alter those fields not marked as "Dyn". An interpreter is therefore free to store values of such fields in its own variables.

11.1.2

The state of the transcription bit (bit 0 of Flags 2) can be changed directly by the game to turn transcribing on or off (see S 7.3, S 7.4). The interpreter must also alter it if stream 2 is turned on or off, to ensure that the bit always reflects the true state of transcribing. Note that the interpreter ensures that its value survives a restart or restore.

11.1.3

Infocom used the interpreter numbers:

| | | | | | |
|---|---------------------------|---|---------------|----|-------------|
| 1 | DECSys tem -20 | 5 | Atari ST | 9 | Apple Iic |
| 2 | Apple IIe | 6 | IBM PC | 10 | Apple Iigs |
| 3 | Macintosh | 7 | Commodore 128 | 11 | Tandy Color |
| 4 | Amiga | 8 | Commodore 64 | | |

(The DECSys~~tem~~-20 was Infocom's own in-house mainframe.) An interpreter should choose the interpreter number most suitable for the machine it will run on. In Versions up to 5, the main consideration is that the behaviour of 'Beyond Zork' depends on the interpreter number (in terms of its usage of the character graphics font). In Version 6, the decision is more serious, as existing Infocom story files depend on interpreter number in many ways: moreover, some story files expect to be run only on the interpreters for a particular machine. (There are, for instance, specifically Amiga versions.)

11.1.3.1

Interpreter versions are conventionally ASCII codes for upper-case letters in Versions 4 and 5 (note that Infocom's Version 6 interpreters just store numbers here).

Modern games are strongly discouraged from testing the interpreter number or interpreter version header information for any game-changing behaviour. It is rarely meaningful, and a Standard interpreter provides many better ways to query the interpreter for information.

11.1.4

***[1.0] The use of bit 7 in 'Flags 1' to signal whether timed input is available was new in the 1.0 document: see the preface.

11.1.5

***[1.0] If an interpreter obeys Revision **n.m** of this document *perfectly*, as far as anyone knows, then byte **\$32** should be written with **n** and byte **\$33** with **m**. If it is an earlier (non-standard) interpreter, it should leave these bytes as 0.

11.1.6

The file length stored at **\$1a** is actually divided by a constant, depending on the Version, to make it fit into a header word. This constant is 2 for Versions 1 to 3, 4 for Versions 4 to 5 or 8 for Versions 6 and later.

11.1.7

The header extension table provides potentially unlimited room for further header information. It is a table of word entries, in which the initial word contains the number of words of data to follow.

11.1.7.1

If the interpreter needs to read a word which is beyond the length of the extension table, or the extension table doesn't exist at all, then the result is 0.

11.1.7.2

If the interpreter needs to write a word which is beyond the length of the extension table, or the extension table doesn't exist at all, then the result is that nothing happens.

11.1.7.3

***[1.0][1.1] Words in the header extension table have been allocated as follows:

Header extension format

| Hex | V | Dyn | Int | Rst | Cont |
|-----|---|-----|-----|-----|--|
| 0 | 5 | | | | Number of further words in table |
| 1 | 5 | | * | | X-coordinate of mouse after a click |
| 2 | 5 | | * | | Y-coordinate of mouse after a click |
| 3 | 5 | | | | Unicode translation table address (optional) |
| 4 | 5 | | | | <i>Flags 3:</i> |
| | 6 | | * | * | 0: If set, game wants to use transparency |
| 5 | 5 | | * | * | True default foreground colour |
| 6 | 5 | | * | * | True default background colour |

11.1.7.4

***[1.1] The bits in Flags 3 are set by the game to request use of a feature. If the interpreter cannot provide a feature, it must clear the relevant bit.

11.1.7.4.1

***[1.1] All unused bits in Flags 3 must be cleared by the interpreter.

Remarks

In the Infocom period, the larger Version 3 story files would not entirely fit on a single Atari 800 disc (though they would fit on a single Apple II, or a single PC disc). Atari versions were therefore made which were identical to the normal ones except for having Flags 1 bit 2 set, and were divided into the resident part on one disc and the rest on another. (This discovery was announced by Stefan Jokisch on 26 August 1997 and sees the end of one of the very few Z-machine mysteries left when Standard 1.0 was first published.)

See the "Infocom fact sheet" for numbers and letters of the known interpreters shipped by Infocom. Interpreter versions are conventionally the upper case letters in sequence (A, B, C, ...). At present most ports of **Zip** use interpreter number 6, and most of **ITF** use number 2.

The unusual behaviour of 'Beyond Zork' concerns its character graphics: see the remarks to **S 16**.

The Macintosh story file for 'Zork Zero' erroneously does not set the pictures bit (Flags 2, bit 3).

The bit in the header described as "requesting screen redraw" may be set by modern interpreters after, for example, resizing the "screen"; games should ideally redraw the screen if they see this bit set. This will usually mean the game clears the screen contents and rearranges borders, etc, so the bit should not be set except when necessary.

The (Version 6) sound and picture bits in Flags 1 indicate general availability of sound and graphics - ie whether the associated opcodes are available and functional.

The bits in Flags 2 should ideally be set reflecting current availability, rather than general support. In other words, if no Blorb (or other) resources for this story file have been found, or if the Blorb file contains no graphics or no sound, the corresponding bits should be cleared.

Also, it is recommended that interpreters that would prompt for an auxiliary Blorb file should do so immediately on start up if any of the "game wants to use sound/music/graphics" bits are set; this allows the bits to be cleared if no file is forthcoming, before the game starts execution. The game can then take appropriate action.

12. The object table

12.1 Storage

The object table is held in dynamic memory and its byte address is stored in the word at **\$0a** in the header. (Recall that objects have flags attached called attributes, numbered from 0 upward, and variables attached called properties, numbered from 1 upward. An object need not provide every property.)

12.2 Property defaults table

The table begins with a block known as the property defaults table. This contains 31 words in Versions 1 to 3 and 63 in Versions 4 and later. When the game attempts to read the value of property *n* for an object which does not provide property *n*, the *n*-th entry in this table is the resulting value.

12.3 Object tree

Next is the object tree. Objects are numbered consecutively from 1 upward, with object number 0 being used to mean "nothing" (though there is formally no such object). The table consists of a list of entries, one for each object.

12.3.1

In Versions 1 to 3, there are at most 255 objects, each having a 9-byte entry as follows:

```
the 32 attribute flags      parent      sibling      child      properties
---32 bits in 4 bytes---   ---3 bytes-----   ---2 bytes--
```

parent, **sibling** and **child** must all hold valid object numbers. The **properties** pointer is the byte address of the list of properties attached to the object. Attributes 0 to 31 are flags (at any given time, they are either on (1) or off (0)) and are stored topmost bit first: e.g., attribute 0 is stored in bit 7 of the first byte, attribute 31 is stored in bit 0 of the fourth.

12.3.2

In Version 4 and later, there are at most 65535 objects, each having a 14-byte entry as follows:

```
the 48 attribute flags      parent      sibling      child      properties
---48 bits in 6 bytes---   ---3 words, i.e. 6 bytes----   ---2 bytes--
```

12.4 Property tables

Each object has its own property table. Each of these can be anywhere in dynamic memory (indeed, a game can legally change an object's properties table address in play, provided the new address points to another valid properties table). The header of a property table is as follows:

```
text-length      text of short name of object
-----byte----- --some even number of bytes---
```

where the **text-length** is the number of 2-byte words making up the text, which is stored in the usual format. (This means that an object's short name is limited to 765 Z-characters.) After the header, the properties are listed in descending numerical order. (This order is essential and is not a matter of convention.)

12.4.1

In Versions 1 to 3, each property is stored as a block

```
size byte      the actual property data
                ---between 1 and 8 bytes---
```

where the **size byte** is arranged as 32 times the number of data bytes minus one, plus the property number. A property list is terminated by a size byte of 0. (It is otherwise illegal for a size byte to be a multiple of 32.)

12.4.2

In Versions 4 and later, a property block instead has the form

```
size and number      the actual property data
--1 or 2 bytes----   --between 1 and 64 bytes--
```

The property number occupies the bottom 6 bits of the first size byte.

12.4.2.1

If the top bit (bit 7) of the first size byte is set, then there are two size-and-number bytes as follows. In the first byte, bits 0 to 5 contain the property number; bit 6 is undetermined (it is clear in all Infocom or Inform story files); bit 7 is set. In the second byte, bits 0 to 5 contain the property data length, counting in bytes; bit 6 is undetermined (it is set in Infocom story files, but clear in Inform ones); bit 7 is always set.

12.4.2.1.1

***[1.0] A value of 0 as property data length (in the second byte) should be interpreted as a length of 64. (Inform can compile such properties.)

12.4.2.2

If the top bit (bit 7) of the first size byte is clear, then there is only one size-and-number byte. Bits 0 to 5 contain the property number; bit 6 is either clear to indicate a property data length of 1, or set to indicate a length of 2; bit 7 is clear.

12.5 Well-foundedness of the tree

It is the game's responsibility to keep the object tree well-founded: the interpreter is not required to check. "Well-founded" means the following:

- (a) An object with a sibling also has a parent.
- (b) An object is the parent of exactly those objects in the sibling list of its child.
- (c) Each object can be given a level n , such that parentless objects have level 0 and all children of a level n object have level $n+1$.

Remarks

The largest valid object number is not directly stored anywhere in the Z-machine. Utility programs like **Infodump** deduce this number by assuming that, initially, the object entries end where the first property table begins.

Infocom's 'Sherlock' contains a bug making it try to set and clear attribute 48.

The reason why the second property size byte needs to have top bit set is that the size field must be parsable either forwards or backwards -- the Z-machine needs to be able to reconstruct the length of a property given only the address of the first byte of its data. (There are very many (e.g. 2000) property entries in a story file, so optimising size into one byte most of the time is worthwhile.)

Bit 6 in the second byte is presently wasted, which is a pity as it could be used to allow up to 128 bytes of property data. But such a change would cause Infocom's story files to fail (since they set this bit, unlike Inform story files).

Inform can only construct well-founded object trees as the initial game state, but it is easy to compile sequences of code like "move red box to blue box" followed by "move blue box to red box" which leave the object tree in an ill-founded state. (The Inform library protects the standard object-movement verbs against this.)

13. The dictionary and lexical analysis

13.1 Storage

The dictionary table is held in static memory and its byte address is stored in the word at \$08 in the header.

13.2 Header

The table begins with a short header:

| | | | |
|------|------------------------------|--------------|-------------------|
| n | list of keyboard input codes | entry-length | number-of-entries |
| byte | -----n bytes----- | byte | 2-byte word |

The keyboard input codes are "word-separators": typically (and under Inform mandatorily) these are the ZSCII codes for full stop, comma and double-quote. Note that a space character (32) should never be a word-separator. The "entry length" is the length of each word's entry in the dictionary table. (It must be at least 4 in Versions 1 to 3, and at least 6 in later Versions.)

13.2.1

Note that the word-separators table can only contain codes which are defined in ZSCII for both input and output.

13.3 Entries (V1 to V3)

In Versions 1 to 3, each word has an entry in the form

| | |
|----------------------|------------------------|
| encoded text of word | bytes of data |
| ----- 4 bytes ----- | (entry length-4) bytes |

The interpreter ignores the bytes of data (presumably the game's parser will use them). The encoded text contains 6 Z-characters (it is always padded out with Z-character 5's to make up 4 bytes: see S 3). The text may include spaces or other word-separators (though, if so, the interpreter will never match any text to the dictionary word in question: surprisingly, this can be useful and is a trick used in the Inform library).

13.4 Entries (later versions)

In Versions 4 and later, the encoded text has 6 bytes and always contains 9 Z-characters.

13.5 Ordering

The word entries follow immediately after the dictionary header and must be given in numerical order of the encoded text (when the encoded text is regarded as a 32 or 48-bit binary number with most-significant byte first). It must not contain two entries with the same encoded text.

13.6 Lexical analysis

Lexical analysis takes place in two circumstances: on request of a **tokenise** opcode (in which case it can use any dictionary table it likes, in the format above) and during acceptance of a game command (in which case the standard dictionary is used).

13.6.1

First, the text is broken up into words. Spaces divide up words and are otherwise ignored. Word separators also divide words, but each one of them is considered a word in its own right. Thus, the erratically-spaced text "fred,go fishing" is divided into four words:

```
fred / , / go / fishing
```

13.6.2

Each word is then encoded as a Z-machine string in dictionary form, and searched for in the dictionary.

13.6.3

A "parse table" is then written, recording the number of words, the length and position of each word and the dictionary address of each word which is recognised. For the format, see the **read** opcode.

Remarks

Usually (under Inform, mandatorily) there are three bytes of data in the word entries, so that dictionary entry lengths are 7 and 9 in the early and late Z-machine, respectively.

It is essential that dictionary entries are in numerical order of the bytes of encrypted text so that interpreters can search the dictionary efficiently (e.g. by a binary-chop algorithm). Because the letters in A0 are in alphabetical order, because the bits are ordered in the right way and because the pad character 5 is less than the values for the letters, the numerical ordering corresponds to normal English alphabetical order for ordinary words. (For instance "an" comes before "anaconda".)

Both Infocom and Inform-compiled games contain words whose initial character is not a letter (for instance, "#record").

Linards Ticmanis reports that some of Infocom's interpreters convert question marks to spaces before lexical analysis. This is not Standard behaviour. (Thus, typing "What is a grue?" into 'Zork I' no longer works: the player must type "What is a grue" instead.)

14. Complete table of opcodes

Two-operand opcodes 2OP

| St | Br | Opcode | Hex | V | Inform name and syntax | Link |
|---|----|--------|-----|-----|---|---------------|
| | | ----- | 0 | --- | --- | |
| * | | 2OP:1 | 1 | | je a b?(label) | je |
| * | | 2OP:2 | 2 | | j1 a b?(label) | j1 |
| * | | 2OP:3 | 3 | | jg a b?(label) | jg |
| * | | 2OP:4 | 4 | | dec_chk (variable) value?(label) | dec_chk |
| * | | 2OP:5 | 5 | | inc_chk (variable) value?(label) | inc_chk |
| * | | 2OP:6 | 6 | | jin obj1 obj2?(label) | jin |
| * | | 2OP:7 | 7 | | test bitmap flags?(label) | test |
| * | | 2OP:8 | 8 | | or a b -> (result) | or |
| * | | 2OP:9 | 9 | | and a b -> (result) | and |
| * | | 2OP:10 | A | | test_attr object attribute?(label) | test_attr |
| | | 2OP:11 | B | | set_attr object attribute | set_attr |
| | | 2OP:12 | C | | clear_attr object attribute | clear_attr |
| | | 2OP:13 | D | | store (variable) value | store |
| | | 2OP:14 | E | | insert_obj object destination | insert_obj |
| * | | 2OP:15 | F | | loadw array word-index -> (result) | loadw |
| * | | 2OP:16 | 10 | | loadb array byte-index -> (result) | loadb |
| * | | 2OP:17 | 11 | | get_prop object property -> (result) | get_prop |
| * | | 2OP:18 | 12 | | get_prop_addr object property -> (result) | get_prop_addr |
| * | | 2OP:19 | 13 | | get_next_prop object property -> (result) | get_next_addr |
| * | | 2OP:20 | 14 | | add a b -> (result) | add |
| * | | 2OP:21 | 15 | | sub a b -> (result) | sub |
| * | | 2OP:22 | 16 | | mul a b -> (result) | mul |
| * | | 2OP:23 | 17 | | div a b -> (result) | div |
| * | | 2OP:24 | 18 | | mod a b -> (result) | mod |
| * | | 2OP:25 | 19 | 4 | call_2s routine arg1 -> (result) | call_2s |
| | | 2OP:26 | 1A | 5 | call_2n routine arg1 | call_2n |
| | | 2OP:27 | 1B | 5 | set_colour foreground background | set_color |
| | | | | 6 | set_colour foreground background window | set_color |
| | | 2OP:28 | 1C | 5/6 | throw value stack-frame | throw |
| | | ----- | 1D | --- | --- | |
| | | ----- | 1E | --- | --- | |
| | | ----- | 1F | --- | --- | |
| <i>Opcode numbers 32 to 127: other forms of 2OP with different types.</i> | | | | | | |

One-operand opcodes 1OP

| St | Br | Opcode | Hex | V | Inform name and syntax | Link |
|--|----|---------|-----|-----|---|--------------|
| | * | 1OP:128 | 0 | | jz a?(label) | jz |
| * | * | 1OP:129 | 1 | | get_sibling object -> (result)?(label) | get_sibling |
| * | * | 1OP:130 | 2 | | get_child object -> (result)?(label) | get_child |
| * | | 1OP:131 | 3 | | get_parent object -> (result) | get_parent |
| * | | 1OP:132 | 4 | | get_prop_len property-address -> (result) | get_prop_len |
| | | 1OP:133 | 5 | | inc (variable) | inc |
| | | 1OP:134 | 6 | | dec (variable) | dec |
| | | 1OP:135 | 7 | | print_addr byte-address-of-string | print_addr |
| * | | 1OP:136 | 8 | 4 | call_1s routine -> (result) | call_ls |
| | | 1OP:137 | 9 | | remove_obj object | remove_obj |
| | | 1OP:138 | A | | print_obj object | print_obj |
| | | 1OP:139 | B | | ret value | ret |
| | | 1OP:140 | C | | jump?(label) | jump |
| | | 1OP:141 | D | | print_paddr packed-address-of-string | print_paddr |
| * | | 1OP:142 | E | | load (variable) -> (result) | load |
| * | | 1OP:143 | F | 1/4 | not value -> (result) | not |
| | | | | 5 | call_1n routine | call_1n |
| <i>Opcode numbers 144 to 175: other forms of 1OP with different types.</i> | | | | | | |

Zero-operand opcodes 0OP

| St | Br | Opcode | Hex | V | Inform name and syntax | Link |
|----|----|---------|-----|-----|----------------------------|------------|
| | | 0OP:176 | 0 | | rtrue | rtrue |
| | | 0OP:177 | 1 | | rfalse | rfalse |
| | | 0OP:178 | 2 | | print (literal-string) | print |
| | | 0OP:179 | 3 | | print_ret (literal-string) | print_ret |
| | | 0OP:180 | 4 | 1/- | nop | nop |
| * | | 0OP:181 | 5 | 1 | save?(label) | save |
| | | | | 4 | save -> (result) | save |
| | | | | 5 | [illegal] | |
| * | | 0OP:182 | 6 | 1 | restore?(label) | restore |
| | | | | 4 | restore -> (result) | restore |
| | | | | 5 | [illegal] | |
| | | 0OP:183 | 7 | | restart | restart |
| | | 0OP:184 | 8 | | ret_popped | ret_popped |
| | | 0OP:185 | 9 | 1 | pop | pop |
| * | | | | 5/6 | catch -> (result) | catch |
| | | 0OP:186 | A | | quit | quit |

| | | | | | | |
|---|---|---------|---|-----|---------------------------------|-------------|
| | | 0OP:187 | B | | new_line | new_line |
| | | 0OP:188 | C | 3 | show_status | show_status |
| | | | | 4 | [illegal] | |
| | * | 0OP:189 | D | 3 | verify?(label) | verify |
| | | 0OP:190 | E | 5 | [first byte of extended opcode] | extended |
| | * | 0OP:191 | F | 5/- | piracy?(label) | piracy |
| <i>Opcode numbers 192 to 223: VAR forms of 2OP:0 to 2OP:31.</i> | | | | | | |

Variable-operand opcodes VAR

| St | Br | Opcode | Hex | V | Inform name and syntax | Link |
|----|----|---------|-----|-----|--|----------------|
| * | | VAR:224 | 0 | 1 | call routine ...0 to 3 args... -> (result) | call |
| | | | | 4 | call_vs routine ...0 to 3 args... -> (result) | call_vs |
| | | VAR:225 | 1 | | storew array word-index value | storew |
| | | VAR:226 | 2 | | storeb array byte-index value | storeb |
| | | VAR:227 | 3 | | put_prop object property value | put_prop |
| | | VAR:228 | 4 | 1 | sread text parse | sread |
| | | | | 4 | sread text parse time routine | sread |
| * | | | | 5 | aread text parse time routine -> (result) | aread |
| | | VAR:229 | 5 | | print_char output-character-code | print_char |
| | | VAR:230 | 6 | | print_num value | print_num |
| * | | VAR:231 | 7 | | random range -> (result) | random |
| | | VAR:232 | 8 | | push value | push |
| | | VAR:233 | 9 | 1 | pull (variable) | pull |
| * | | | | 6 | pull stack -> (result) | pull |
| | | VAR:234 | A | 3 | split_window lines | split_window |
| | | VAR:235 | B | 3 | set_window window | set_window |
| * | | VAR:236 | C | 4 | call_vs2 routine ...0 to 7 args... -> (result) | call_vs2 |
| | | VAR:237 | D | 4 | erase_window window | erase_window |
| | | VAR:238 | E | 4/- | erase_line value | erase_line |
| | | | | 6 | erase_line pixels | erase_line |
| | | VAR:239 | F | 4 | set_cursor line column | set_cursor |
| | | | | 6 | set_cursor line column window | set_cursor |
| | | VAR:240 | 10 | 4/6 | get_cursor array | get_cursor |
| | | VAR:241 | 11 | 4 | set_text_style style | set_text_style |
| | | VAR:242 | 12 | 4 | buffer_mode flag | buffer_mode |
| | | VAR:243 | 13 | 3 | output_stream number | output_stream |
| | | | | 5 | output_stream number table | output_stream |
| | | | | 6 | output_stream number table width | output_stream |
| | | VAR:244 | 14 | 3 | input_stream number | input_stream |

| | | | | | | |
|---|---|---------|----|-----|---|-----------------|
| | | VAR:245 | 15 | 5/3 | sound_effect number effect volume routine | sound_effect |
| * | | VAR:246 | 16 | 4 | read_char 1 time routine -> (result) | read_char |
| * | * | VAR:247 | 17 | 4 | scan_table x table len form -> (result) | scan_table |
| * | | VAR:248 | 18 | 5/6 | not value -> (result) | not |
| | | VAR:249 | 19 | 5 | call_vn routine ...up to 3 args... | call_vn |
| | | VAR:250 | 1A | 5 | call_vn2 routine ...up to 7 args... | call_vn2 |
| | | VAR:251 | 1B | 5 | tokenise text parse dictionary flag | tokenise |
| | | VAR:252 | 1C | 5 | encode_text zscii-text length from coded-text | encode_text |
| | | VAR:253 | 1D | 5 | copy_table first second size | copy_table |
| | | VAR:254 | 1E | 5 | print_table zscii-text width height skip | print_table |
| | * | VAR:255 | 1F | 5 | check_arg_count argument-number | check_arg_count |

Extended opcodes EXT

| St | Br | Opcode | Hex | V | Inform name and syntax | Link |
|----|----|--------|-----|-----|--|----------------|
| * | | EXT:0 | 0 | 5 | save table bytes name prompt -> (result) | save |
| * | | EXT:1 | 1 | 5 | restore table bytes name prompt -> (result) | restore |
| * | | EXT:2 | 2 | 5 | log_shift number places -> (result) | log_shift |
| * | | EXT:3 | 3 | 5/- | art_shift number places -> (result) | art_shift |
| * | | EXT:4 | 4 | 5 | set_font font -> (result) | set_font |
| * | | | | 6/- | set_font font window -> (result) | set_font |
| | | EXT:5 | 5 | 6 | draw_picture picture-number y x | draw_picture |
| | * | EXT:6 | 6 | 6 | picture_data picture-number array?(label) | picture_data |
| | | EXT:7 | 7 | 6 | erase_picture picture-number y x | erase_picture |
| | | EXT:8 | 8 | 6 | set_margins left right window | set_margins |
| * | | EXT:9 | 9 | 5 | save_undo -> (result) | save_undo |
| * | | EXT:10 | A | 5 | restore_undo -> (result) | restore_undo |
| | | EXT:11 | B | 5/* | print_unicode char-number | print_unicode |
| | | EXT:12 | C | 5/* | check_unicode char-number -> (result) | check_unicode |
| | | EXT:13 | D | 5/* | set_true_colour foreground background | set_true_color |
| | | | | 6/* | set_true_colour foreground background window | set_true_color |
| | | ----- | E | --- | --- | |
| | | ----- | F | --- | --- | |
| | | EXT:16 | 10 | 6 | move_window window y x | move_window |
| | | EXT:17 | 11 | 6 | window_size window y x | window_size |
| | | EXT:18 | 12 | 6 | window_style window flags operation | window_style |
| * | | EXT:19 | 13 | 6 | get_wind_prop window property-number -> (result) | get_wind_prop |
| | | EXT:20 | 14 | 6 | scroll_window window pixels | scroll_window |
| | | EXT:21 | 15 | 6 | pop_stack items stack | pop_stack |
| | | EXT:22 | 16 | 6 | read_mouse array | read_mouse |

| | | | | | | |
|--|---|--------|----|-----|--|----------------|
| | | EXT:23 | 17 | 6 | mouse_window window | mouse_window |
| | * | EXT:24 | 18 | 6 | push_stack value stack ?(label) | push_stack |
| | | EXT:25 | 19 | 6 | put_wind_prop window property-number value | put_wqind_prop |
| | | EXT:26 | 1A | 6 | print_form formatted-table | print_form |
| | * | EXT:27 | 1B | 6 | make_menu number table ?(label) | make_menu |
| | | EXT:28 | 1C | 6 | picture_table table | picture_table |
| | * | EXT:29 | 1D | 6/* | buffer_screen mode -> (result) | buffer_screen |

14.1 Contents

This table contains all 119 opcodes and, taken with the dictionary in S 15, describes exactly what each should do. In addition, it lists which opcodes are actually used in the known Infocom story files, and documents the Inform assembly language syntax.

14.2 Out of range opcodes

Formally, it is illegal for a game to contain an opcode not specified for its version. An interpreter should normally halt with a suitable message.

14.2.1

However, extended opcodes in the range EXT:29 to EXT:255 should be simply ignored (perhaps with a warning message somewhere off-screen).

14.2.2

***[1.0][1.1] EXT:11 and EXT:12 were opcodes added in Standard 1.0 and can be generated in code compiled by Inform 6.12 or later. EXT:13 and EXT:29 are new in Standard 1.1. EXT:14 to EXT:15, and EXT:30 to EXT:127, are reserved for future versions of this document to specify.

14.2.3

Designers who wish to create their own "new" opcodes, for one specific game only, are asked to use opcode numbers in the range EXT:128 to EXT:255. It is easy to modify Inform to name and assemble such opcodes. (Of course the game will then have to be circulated with a suitably modified interpreter to run it.)

14.2.4

Interpreter-writers should ideally make this easy by providing a routine which is called if EXT:128 to EXT:255 are found, so that the minimum possible modification to the interpreter is needed.

Reading the opcode tables

The two columns "St" and "Br" (store and branch) mark whether an instruction stores a result in a variable, and whether it must provide a label to jump to, respectively.

The "Opcode" is written **TYPE:Decimal** where the **TYPE** is the operand count (2OP, 1OP, 0OP or VAR) or else EXT for two-byte opcodes (where the first byte is (decimal) 190). The decimal number is the lowest possible decimal opcode value. The hex number is the opcode number within each **TYPE**.

The "V" column gives the Version information. If nothing is specified, the opcode is as stated from Version 1 onwards. Otherwise, it exists only from the version quoted onwards. Before this time, its use is illegal. Some opcodes change their meanings as the Version increases, and these have more than one line of specification. Others become illegal again, and these are marked **[illegal]**.

In a few cases, the Version is given as "3/4" or some such. The first number is the Version number whose specification the opcode belongs to, and the second is the earliest Version in which the opcode is known actually to be used in an Infocom-produced story file. A dash means that it seems never to have been used (in any of Versions 1 to 6). The notation "5/*" or "6/*" means that the opcode was introduced in this Standards document long after the Infocom era.

The table explicitly marks opcodes which do not exist in any version of the Z-machine as -----: in addition, none of the extended set of codes after EXT:29 have been used.

Inform assembly language

This section documents Inform 6 assembly language, which is richer than that of Inform 5. The Inform 6 assembler can generate every legal opcode and automatically sets any consequent header bits (for instance, a usage of **set_colour** will set the "colours needed" bit).

One way to get a picture of Inform assembly language is to compile a short program with tracing switched on (using the **-a** or **-t** switches).

1. An Inform statement beginning with an **@** is sent directly to the assembler. In the syntax below, **(variable)** and **(result)** must be variables (or **sp**, a special variable name available only in assembly language, and meaning the stack pointer); **(label)** a label (not a routine name).

(literal-string) must be literal text in quotation marks "thus".

routine should be the name of a routine (this assembles to its packed address). Otherwise any Inform constant term (such as **'** or **'beetle'**) can be given as an operand.

2. It is optional, but sensible, to place a -> sign before a store-variable. For example, in

```
@mul a 56 -> sp;
```

("multiply variable **a** by 56, and put the result on the stack") the -> can be omitted, but should be included for clarity.

3. A label to branch to should be prefaced with a question mark ?, as in

```
@je a b ?Equal;      ! Branch to Equal if a == b
```

(If the question mark is omitted, the branch is compiled in the short form, which will only work for very nearby labels and is very seldom useful in code written by hand.) Note that the effect of any branch instruction can be negated using a tilde ~:

```
@je a b ?~Different; ! Branch to Different if a ~= b
```

4. Labels are assembled using full stops:

```
.MyLabel;
```

All branches must be to such a label within the same routine. (The Inform assembler imposes the same-routine restriction.)

5. Most operands are assembled in the obvious way: numbers and constant values (like characters) as numbers, variables as variables, **sp** as the value on top of the stack. There are two exceptions. "Call" opcodes expect as first operand the name of a routine to call:

```
@call_1n MyRoutine;
```

but one can also give an indirect address, as a constant or variable, using square brackets:

```
@call_1n [x];      ! Call routine whose address is in x
```

Secondly, seven Z-machine opcodes access variables but by their numbers: thus one should write, say, the constant 0 instead of the variable **sp**. This is inconvenient, so the Inform assembler accepts variable names instead. The operands affected are those marked as **(variable)** in the syntax chart; Inform translates the variable name as a "small constant" operand with that variable's number as value. The affected opcodes are:

```
inc, dec, inc_chk, dec_chk, store, pull, load.
```

This is useful, but there is another possibility, of genuinely giving a variable operand. The Inform notation for this involves square brackets again:

```
@inc frog;           ! Increment var "frog"
@inc [frog];         ! Increment var whose number is in "frog"
```

Infocom story files often use such instructions.

6. The Inform assembler is also written with possible extensions to the Z-machine instruction set in mind. (Of course these can only work if a customised interpreter is used.) Simply give a specification in double-quotes where you would normally give the opcode name. For example,

```
@"10P:4S" 34 -> i;
@get_prop_len 34 -> i;
```

are equivalent instructions, since **get_prop_len** is instruction 4 in the 1OP (one-operand) set, and is a Store opcode. The syntax is:

| | | | | | | |
|---|----------|---|----------------|-------|---|-----------------|
| " | 0OP | : | decimal-number | flags | " | (range 0 to 15) |
| | 1OP | | | | | 0 15 |
| | 2OP | | | | | 0 15 |
| | VAR | | | | | 32 63 |
| | VAR_LONG | | | | | 32 63 |
| | EXT | | | | | 0 255 |
| | EXT_LONG | | | | | 0 255 |

(**EXT_LONG** is a logical possibility but has not been used in the Z-machine so far: the assembler provides it in case it might be useful in future.) The possible flags are:

```
S    Store opcode
B    Branch opcode
T    Text in-line instead of operands
      (as with "print" and "print_ret")
I    "Indirect addressing": first operand is a (variable)
Fnn  Set bit nn in Flags 2 (signalling to the interpreter that an
      unusual feature has been called for): the number is in decimal
```

For example,

```
"EXT:128BSF14"
```

is an exotic new opcode, number 128 in the extended range, which is both Branch and Store, and the assembly of which causes bit 14 to be set in "Flags 2". See § 14.2 below for rules on how to number newly created opcodes.

Remarks

The opcodes EXT:5 to EXT:8 were very likely in Infocom's own Version 5 specification (documentary records of which are lost): they seem to have been partially implemented in existing Infocom interpreters, but do not occur in any existing Version 5 story file. They are here left unspecified.

The notation "5/3" for **sound_effect** is because this plainly Version 5 feature was used also in one solitary Version 3 game, 'The Lurking Horror' (the sound version of which was the last Version 3 release, in September 1987).

The 2OP opcode 0 was possibly intended for setting break-points in debugging (and may be used for this again). It was not **nop**.

read_mouse and **make_menu** are believed to have been used only in 'Journey' (based on a check of 11 Version 6 story files).

picture_table is used once by 'Shogun' and several times by 'Zork Zero'.

15. Dictionary of opcodes

The highest ideal of a translation... is achieved when the reader flings it impatiently into the fire, and begins patiently to learn the language for himself.

Philip Vellacott

15.1

The dictionary below is alphabetical and includes entries on every opcode listed in the table above, as well as brief notes on a few opcodes once thought to exist but now disproved.

15.2

The Z-machine has the same concept of "table" (as an internal data structure) as Inform. Specifically, a table is an array of words (in dynamic or static memory) of which the initial entry is the number of subsequent words in the table. For example, a table with three entries occupies 8 bytes, arranged as the words 3, x, y, z.

15.3

In all cases below where one operand is supposed to be in a particular range, behaviour is undefined if it is not. For instance an interpreter complies with the Standard even if it crashes when an illegal object number (including 0) is given for an object operand. However, see S A for guidelines on detecting and dealing with errors.

add

2OP:20 14 add a b -> (result)

Signed 16-bit addition.

and

2OP:9 9 and a b -> (result)

Bitwise AND.

aread

This is the Inform name for the keyboard-reading opcode under Version 5 and later. (Inform calls the same opcode **sread** under Versions 3 and 4.) See **read** for the specification.

art_shift

EXT:3 3 5/- art_shift number places -> (result)

Does an arithmetic shift of **number** by the given number of places, shifting left (i.e. increasing) if places is positive, right if negative. In a right shift, the sign bit is preserved as well as being shifted on down. (The alternative behaviour is **log_shift**.)

The "places" operand must be in the range -15 to +15, otherwise behaviour is undefined.

buffer_mode

VAR:242 12 4 buffer_mode flag

If set to 1, text output on the lower window in stream 1 is buffered up so that it can be word-wrapped properly. If set to 0, it isn't.

In Version 6, this opcode is redundant (the "buffering" window attribute can be set instead). It is used twice in each of Infocom's Version 6 story files, in the **\$verify** routine. **Frotz** responds by setting the current window's "buffering" attribute, while Infocom's own interpreters respond by doing nothing. This standard leaves the result of **buffer_mode** undefined in Version 6.

buffer_screen

EXT:29 1D 6/* buffer_screen mode -> (result)

Tells the interpreter how to handle display buffering. If **mode** is 0, updates must be made as soon as possible. If **mode** is 1, the interpreter may make changes to a backing store, and need not update the screen. The interpreter is free to ignore the advice, but if so must always act as though the **mode** is 0 (update the screen as soon as possible).

With **buffer_screen** in either state, an update of the visible display can be forced immediately by issuing **buffer_screen -1**, without altering the current buffering state. Note that **buffer_screen -1** does not flush the text buffer.

The return value is the old **buffer_screen** state.

See **S 8** for more details.

*****[1.1]** This opcode will only be present in interpreters obeying Standard 1.1 or later, so story files should check the standard number of the interpreter before executing this opcode.

call

VAR:224 0 1 call routine ...up to 3 args... -> (result)

The only call instruction in Version 3, Inform reads this as **call_vs** in higher versions: it calls the routine with 0, 1, 2 or 3 arguments as supplied and stores the resulting return value. (When the address 0 is called as a routine, nothing happens and the return value is false.)

call_1n

1OP:143 F 5 call_1n routine

Executes **routine()** and throws away result.

call_1s

1OP:136 8 4 call_1s routine -> (result)

Stores **routine()**.

call_2n

2OP:26 1A 5 call_2n routine arg1

Executes **routine(arg1)** and throws away result.

call_2s

2OP:25 19 4 call_2s routine arg1 -> (result)

Stores **routine(arg1)**.

call_vn

VAR:249 19 5 call_vn routine ...up to 3 args...

Like **call**, but throws away result.

call_vs

VAR:224 0 4 call_vs routine ...up to 3 args... -> (result)

See **call**.

call_vn2

VAR:250 1A 5 call_vn2 routine ...up to 7 args...

Call with a variable number (from 0 to 7) of arguments, then throw away the result. This (and **call_vs2**) uniquely have an extra byte of opcode types to specify the types of arguments 4 to 7. Note that it is legal to use these opcodes with fewer than 4 arguments (in which case the second byte of type information will just be **\$ff**).

call_vs2

VAR:236 C 4 call_vs2 routine ...up to 7 args... -> (result)

See **call_vn2**.

catch

0OP:185 9 5/6 catch -> (result)

Opposite to **throw** (and occupying the same opcode that **pop** used in Versions 3 and 4). **catch** returns the current "stack frame".

check_arg_count

VAR:255 1F 5 check_arg_count argument-number

Branches if the given argument-number (counting from 1) has been provided by the routine call to the current routine. (This allows routines in Versions 5 and later to distinguish between the calls **routine(1)** and **routine(1,0)**, which would otherwise be impossible to tell apart.)

check_unicode

EXT:12 C 5/* check_unicode char-number -> (result)

Determines whether or not the interpreter can print, or receive from the keyboard, the given Unicode character. Bit 0 of the result should be set if and only if the interpreter can print the character; bit 1 if and only if the interpreter can receive it from the keyboard. Bits 2 to 15 are undefined.

***[1.0] This opcode will only be present in interpreters obeying Standard 1.0 or later, so story files should check the standard number of the interpreter before executing this opcode.

clear_attr

2OP:12 C clear_attr object attribute

Make **object** not have the attribute numbered **attribute**.

copy_table

VAR:253 1D 5 copy_table first second size

If **second** is zero, then **size** bytes of **first** are zeroed.

Otherwise **first** is copied into **second**, its length in bytes being the absolute value of **size** (i.e., **size** if **size** is positive, **-size** if **size** is negative).

The tables are allowed to overlap. If **size** is positive, the interpreter must copy either forwards or backwards so as to avoid corrupting **first** in the copying process. If **size** is negative, the interpreter must copy forwards even if this **corrupts** first. ('Beyond Zork' uses this to fill an array with spaces.)

(Version 0.2 of this document wrongly specified that if **size** is positive then copying should always run backward. This results in the player being unable to cross the river near the start of 'Journey', as the game uses **copy_table** to shuffle menu options, and the menu "Downstream, Upstream, Cross, Return" is changed to "Return, Return, Return".)

dec

1OP:134 6 dec (variable)

Decrement variable by 1. This is signed, so 0 decrements to -1.

dec_chk

2OP:4 4 dec_chk (variable) value?(label)

Decrement variable, and branch if it is now less than the given value.

div

2OP:23 17 div a b -> (result)

Signed 16-bit division. Division by zero should halt the interpreter with a suitable error message.

draw_picture

EXT:5 5 6 draw_picture picture-number y x

Displays the picture with the given number. (y,x) coordinates (of the top left of the picture) are each optional, in that a value of zero for y or x means the cursor y or x coordinate in the current window. It is illegal to call this with an invalid picture number.

encode_text

VAR:252 1C 5 encode_text zscii-text length from coded-text

Translates a ZSCII word to Z-encoded text format (stored at **coded-text**), as if it were an entry in the dictionary. The text begins at **from** in the **zscii-text** buffer and is **length** characters long. (Some interpreters ignore this and keep translating until they hit a 0 character anyway, or have already filled up the 6-byte Z-encoded string.)

erase_line

VAR:238 E 4/6 erase_line value

Versions 4 and 5: if the value is 1, erase from the current cursor position to the end of its line in the current window. If the value is anything other than 1, do nothing.

Version 6: if the value is 1, erase from the current cursor position to the end of its line in the current window. If not, erase the given number of pixels minus one across from the cursor (clipped to stay inside the right margin). The cursor does not move.

erase_picture

EXT:7 7 6 erase_picture picture-number y x

Like **draw_picture**, but paints the appropriate region to the background colour for the given window. It is illegal to call this with an invalid picture number.

erase_window

VAR:237 D 4 erase_window window

Erases window with given number (to background colour); or if -1 it unspits the screen and clears the lot; or if -2 it clears the screen without unsplitting it. In cases -1 and -2, the cursor may move (see S 8 for precise details).

"extended"

This byte (decimal 190) is not an instruction, but indicates that the opcode is "extended": the next byte contains the number in the extended set.

get_child

1OP:130 2 get_child object -> (result) ?(label)

Get first object contained in given object, branching if this exists, i.e. is not **nothing** (i.e., is not 0).

get_cursor

VAR:240 10 4/6 get_cursor array

Puts the current cursor row into the word 0 of the given array, and the current cursor column into word 1. (The array is not a table and has no size information in its initial entry.)

get_next_prop

2OP:19 13 get_next_prop object property -> (result)

Gives the number of the next property provided by the quoted object. This may be zero, indicating the end of the property list; if called with zero, it gives the first property number present. It is illegal to try to find the next property of a property which does not exist, and an interpreter should halt with an error message (if it can efficiently check this condition).

get_parent

1OP:131 3 get_parent object -> (result)

Get parent object (note that this has no "branch if exists" clause).

get_prop

2OP:17 11 get_prop object property -> (result)

Read property from object (resulting in the default value if it had no such declared property). If the property has length 1, the value is only that byte. If it has length 2, the first two bytes of the property are taken as a word value. It is illegal for the opcode to be used if the property has length greater than 2, and the result is unspecified.

get_prop_addr

2OP:18 12 get_prop_addr object property -> (result)

Get the byte address (in dynamic memory) of the property data for the given object's property. This must return 0 if the object hasn't got the property.

get_prop_len

1OP:132 4 get_prop_len property-address -> (result)

Get length of property data (in bytes) for the given object's property. It is illegal to try to find the property length of a property which does not exist for the given object, and an interpreter should halt with an error message (if it can efficiently check this condition).

@get_prop_len 0 must return 0. This is required by some Infocom games and files generated by old versions of Inform.

get_sibling

1OP:129 1 get_sibling object -> (result) ?(label)

Get next object in tree, branching if this exists, i.e. is not 0.

get_wind_prop

EXT:19 13 6 get_wind_prop window property-number -> (result)

Reads the given property of the given window (see S 8).

inc

1OP:133 5 inc (variable)

Increment variable by 1. (This is signed, so -1 increments to 0.)

inc_chk

2OP:5 5 inc_chk (variable) value?(label)

Increment variable, and branch if now greater than value.

input_stream

VAR:244 14 3 input_stream number

Selects the current input stream.

insert_obj

2OP:14 E insert_obj object destination

Moves object O to become the first child of the destination object D. (Thus, after the operation the **child** of D is O, and the **sibling** of O is whatever was previously the **child** of D.) All children of O move with it. (Initially O can be at any point in the object tree; it may legally have **parent** zero.)

je

2OP:1 1 je a b c d?(label)

Jump if **a** is equal to any of the subsequent operands. (Thus **@je a** never jumps and **@je a b** jumps if **a = b**.)

je with just 1 operand is not permitted.

jg

2OP:3 3 jg a b?(label)

Jump if **a > b** (using a signed 16-bit comparison).

jin

2OP:6 6 jin obj1 obj2?(label)

Jump if object **a** is a direct child of **b**, i.e., if **parent** of **a** is **b**.

jl

2OP:2 2 jl a b ?(label)

Jump if **a < b** (using a signed 16-bit comparison).

jump

1OP:140 C jump ?(label)

Jump (unconditionally) to the given label. (This is not a branch instruction and the operand is a 2-byte signed offset to apply to the program counter.) It is legal for this to jump into a different routine (which should not change the routine call state), although it is considered bad practice to do so and the **Txd** disassembler is confused by it.

The destination of the jump opcode is:

address after instruction + Offset - 2

This is analogous to the calculation for branch offsets.

jz

1OP:128 0 jz a ?(label)

Jump if **a = 0**.

load

1OP:142 E load (variable) -> (result)

The value of the variable referred to by the operand is stored in the result. (Inform doesn't use this; see the notes to S 14.)

loadb

2OP:16 10 loadb array byte-index -> (result)

Stores **array->byte-index** (i.e., the byte at address **array+byte-index**, which must lie in static or dynamic memory).

loadw

2OP:15 F loadw array word-index -> (result)

Stores **array-->word-index** (i.e., the word at address **array+2*word-index**, which must lie in static or dynamic memory).

log_shift

EXT:2 2 5 log_shift number places -> (result)

Does a logical shift of **number** by the given number of **places**, shifting left (i.e. increasing) if **places** is positive, right if **negative**. In a right shift, the sign is zeroed instead of being shifted on. (See also **art_shift**.)

The "places" operand must be in the range -15 to +15, otherwise behaviour is undefined.

make_menu

EXT:27 1B 6 make_menu number table ?(label)

Controls menus with numbers greater than 2 (i.e., it doesn't control the three system menus). If the table supplied is 0, the menu is removed. Otherwise it is a table of tables. Each table is a ZSCII string: the first item being a menu name, subsequent ones the entries.

mod

2OP:24 18 mod a b -> (result)

Remainder after signed 16-bit division. Division by zero should halt the interpreter with a suitable error message.

mouse_window

EXT:23 17 6 mouse_window window

Constrain the mouse arrow to sit inside the given window. By default it sits in window 1. Setting to -1 takes all restriction away. (The mouse clicks are not reported if the arrow is outside the window and interpreters are presumably supposed to hold the arrow there by hardware means if possible.)

move_window

EXT:16 10 6 move_window window y x

Moves the given window to pixels (y,x): (1,1) being the top left. Nothing actually happens (since windows are entirely notional transparencies): but any future plotting happens in the new place.

mul

2OP:22 16 mul a b -> (result)

Signed 16-bit multiplication.

new_line

0OP:187 B new_line

Print carriage return.

nop

0OP:180 4 1/- nop

Probably the official "no operation" instruction, which, appropriately, was never operated (in any of the Infocom datafiles): it may once have been a breakpoint.

not

1OP:143 F 1/4 not value -> (result)

VAR:248 18 5/6 not value -> (result)

Bitwise NOT (i.e., all 16 bits reversed). Note that in Versions 3 and 4 this is a 1OP instruction, reasonably since it has 1 operand, but in later Versions it was moved into the extended set to make room for **call_1n**.

or

2OP:8 8 or a b -> (result)

Bitwise OR.

output_stream

VAR:243 13 3 output_stream number table

6 output_stream number table width

If **stream** is 0, nothing happens. If it is positive, then that stream is selected; if negative, deselected. (Recall that several different streams can be selected at once.)

When stream 3 is selected, a **table** must be given into which text can be printed. The first word always holds the number of characters printed, the actual text being stored at bytes **table+2** onward. It is not the interpreter's responsibility to worry about the length of this table being overrun.

In Version 6, a **width** field may optionally be given: text will then be justified as if it were in the window with that number (if width is zero or positive) or a box **-width** pixels wide (if negative). Then the table will contain not ordinary text but formatted text: see **print_form**.

picture_data

EXT:6 6 6 picture_data picture-number array?(label)

Asks the interpreter for data on the picture with the given number. If the picture number is valid, a branch occurs and information is written to the array: the height in word 0, the width in word 1, in pixels. (This is an array, not a "table" with initial size information.)

Otherwise, if the picture number is zero, the interpreter writes the number of available pictures into word 0 of the array and the release number of the picture file into word 1, and branches if any pictures are available. (Infocom's first Version 6 Amiga interpreter did not handle this case properly, and early releases of 'Zork Zero' did not use it. The feature may have been added on the MSDOS release of 'Zork Zero'.)

Otherwise, nothing happens.

picture_table

EXT:28 1C 6 picture_table table

Given a table of picture numbers, the interpreter may if it wishes load or unpack these pictures from disc into a cache for convenient rapid plotting later. 'Zork Zero' makes frequent use of this, for instance for its peggleboard display. Moreover, it expects rapid plotting only for those images listed in the last call to **picture_table**. In other words, any images still in the cache when **picture_table** is called can safely be thrown away. (The Amiga interpreter 6.14 uses a cache of size 5K and never caches any individual image larger than 1K.)

piracy

0OP:191 F 5/- piracy?(label)

Branches if the game disc is believed to be genuine by the interpreter (which is assumed to have some arcane way of finding out). Interpreters are asked to be gullible and to unconditionally branch.

pop

0OP:185 9 1 pop

Throws away the top item on the stack. (This was useful to lose unwanted routine call results in early Versions.)

pop_stack

EXT:21 15 6 pop_stack items stack

The given number of items are thrown away from the top of a stack: by default the system stack, otherwise the one given as a second operand.

print

0OP:178 2 print <literal-string>

Print the quoted (literal) Z-encoded string.

print_addr

1OP:135 7 print_addr byte-address-of-string

Print (Z-encoded) string at given byte address, in dynamic or static memory.

print_char

VAR:229 5 print_char output-character-code

Print a ZSCII character. The operand must be a character code defined in ZSCII for output (see S 3). In particular, it must certainly not be negative or larger than 1023.

print_form

EXT:26 1A 6 print_form formatted-table

Prints a formatted table of the kind written to output stream 3 when formatting is on. This is an elaborated version of **print_table** to cope with fonts, pixels and other impedimenta. It is a sequence of lines, terminated with a zero word. Each line is a word containing the number of characters, followed by that many bytes which hold the characters concerned.

print_num

VAR:230 6 print_num value

Print (signed) number in decimal.

print_obj

1OP:138 A print_obj object

Print short name of object (the Z-encoded string in the object header, not a property). If the object number is invalid, the interpreter should halt with a suitable error message.

print_paddr

1OP:141 D print_paddr packed-address-of-string

Print the (Z-encoded) string at the given packed address in high memory.

print_ret

0OP:179 3 print_ret <literal-string>

Print the quoted (literal) Z-encoded string, then print a new-line and then return true (i.e., 1).

print_table

VAR:254 1E 5 print_table zscii-text width height skip

Print a rectangle of text on screen spreading right and down from the current cursor position, of given **width** and **height**, from the table of ZSCII text given. (Height is optional and defaults to 1.) If a **skip** value is given, then that many characters of text are skipped over in between each line and the next. (So one could make this display, for instance, a 2 by 3 window onto a giant 40 by 40 character graphics map.)

print_unicode

EXT:11 B 5/* print_unicode char-number

Print a Unicode character. See S 3.8.5.4 and S 7.5 for details. The given character code must be defined in Unicode.

***[1.0] This opcode will only be present in interpreters obeying Standard 1.0 or later, so story files should check the standard number of the interpreter before executing this opcode.

pull

VAR:233 9 1 pull (variable)
6 pull stack -> (result)

Pulls value off a stack. (If the stack underflows, the interpreter should halt with a suitable error message.) In Version 6, the stack in question may be specified as a user one: otherwise it is the game stack.

push

VAR:232 8 push value

Pushes value onto the game stack.

push_stack

EXT:24 18 6 push_stack value stack ?(label)

Pushes the value onto the specified user stack, and branching if this was successful. If the stack overflows, nothing happens (this is not an error condition).

put_prop

VAR:227 3 put_prop object property value

Writes the given value to the given property of the given object. If the property does not exist for that object, the interpreter should halt with a suitable error message. If the property length is 1, then the interpreter should store only the least significant byte of the value. (For instance, storing -1 into a 1-byte property results in the property value 255.) As with **get_prop** the property length must not be more than 2: if it is, the behaviour of the opcode is undefined.

put_wind_prop

EXT:25 19 6 put_wind_prop window property-number value

Writes a window property (see **get_wind_prop**). This should only be used when there is no direct command (such as **move_window**) to use instead, as some such operations may have side-effects.

quit

00P:186 A quit

Exit the game immediately. (Any "Are you sure?" question must be asked by the game, not the interpreter.) It is not legal to return from the main routine (that is, from where execution first begins) and this must be used instead.

random

VAR:231 7 random range -> (result)

If **range** is positive, returns a uniformly random number between 1 and **range**. If **range** is negative, the random number generator is seeded to that value and the return value is 0. Most interpreters consider giving 0 as **range** illegal (because they attempt a division with remainder by the **range**), but correct behaviour is to reseed the generator in as random a way as the interpreter can (e.g. by using the time in milliseconds).

(Some version 3 games, such as 'Enchanter' release 29, had a debugging verb **#random** such that typing, say, **#random 14** caused a call of **random** with -14.)

read

VAR:228 4 1 sread text parse

4 sread text parse time routine

5 aread text parse time routine -> (result)

(Note that Inform internally names the **read** opcode as **aread** in Versions 5 and later and **sread** in Versions 3 and 4.)

This opcode reads a whole command from the keyboard (no prompt is automatically displayed). It is legal for this to be called with the cursor at any position on any window.

In Versions 1 to 3, the status line is automatically redisplayed first.

A sequence of characters is read in from the current input stream until a carriage return (or, in Versions 5 and later, any terminating character) is found.

In Versions 1 to 4, byte 0 of the **text-buffer** should initially contain the maximum number of letters which can be typed, minus 1 (the interpreter should not accept more than this). The text typed is reduced to lower case (so that it can tidily be printed back by the program if need be) and stored in bytes 1 onward, with a zero terminator (but without any other terminator, such as a carriage return code). (This means that if byte 0 contains n then the buffer must contain n+1 bytes, which makes it a **string** array of length n in Inform terminology.)

In Versions 5 and later, byte 0 of the **text-buffer** should initially contain the maximum number of letters which can be typed (the interpreter should not accept more than this). The interpreter stores the number of characters actually typed in byte 1 (not counting the terminating character), and the characters themselves (reduced to lower case) in bytes 2 onward (not storing the terminating character). (Some interpreters wrongly add a zero byte after the text anyway, so it is wise for the buffer to contain at least n+3 bytes.)

Moreover, if byte 1 contains a positive value at the start of the input, then **read** assumes that number of characters are left over from an interrupted previous input, and writes the new characters after those already there. Note that the interpreter does not redisplay the characters left over: the game does this, if it wants to. This is unfortunate for any interpreter wanting to give input text a distinctive appearance on-screen, but 'Beyond Zork', 'Zork Zero' and 'Shogun' clearly require it. ("Just a tremendous pain in my butt" -- Andrew Plotkin; "the most unfortunate feature of the Z-machine design" -- Stefan Jokisch.)

In Version 4 and later, if the operands **time** and **routine** are supplied (and non-zero) then the routine call **routine()** is made every **time/10** seconds during the keyboard-reading process.

If this routine returns true, all input is erased (to zero) and the reading process is terminated at once. (The terminating character code is 0.) The **routine** is permitted to print to the screen even if it returns false to signal "carry on": the interpreter should notice and redraw the input line so far, before input continues. (**Frotz** notices by looking to see if the cursor position is at the left-hand margin after the interrupt routine has returned.)

If input was terminated in the usual way, by the player typing a carriage return, then a carriage return is printed (so the cursor moves to the next line). If it was interrupted, the cursor is left at the rightmost end of the text typed in so far.

Next, lexical analysis is performed on the text (except that in Versions 5 and later, if **parse-buffer** is zero then this is omitted). Initially, byte 0 of the **parse-buffer** should hold the maximum number of textual words which can be parsed. (If this is n , the buffer must be at least $2 + 4*n$ bytes long to hold the results of the analysis.)

The interpreter divides the text into words and looks them up in the dictionary, as described in **S 13**. The number of words is written in byte 1 and one 4-byte block is written for each word, from byte 2 onwards (except that it should stop before going beyond the maximum number of words specified). Each block consists of the byte address of the word in the dictionary, if it is in the dictionary, or 0 if it isn't; followed by a byte giving the number of letters in the word; and finally a byte giving the position in the **text-buffer** of the first letter of the word.

In Version 5 and later, this is a store instruction: the return value is the terminating character (note that the user pressing his "enter" key may cause either 10 or 13 to be returned; the interpreter must return 13). A timed-out input returns 0.

(Versions 1 and 2 and early Version 3 games mistakenly write the parse buffer length 240 into byte 0 of the parse buffer: later games fix this bug and write 59, because $2+4*59 = 238$ so that 59 is the maximum number of textual words which can be parsed into a buffer of length 240 bytes. Old versions of the Inform 5 library commit the same error. Neither mistake has very serious consequences.)

(Interpreters are asked to halt with a suitable error message if the text or parse buffers have length of less than 3 or 6 bytes, respectively: this sometimes occurs due to a previous array being overrun, causing bugs which are very difficult to find.)

read_char

VAR:246 16 4 read_char 1 time routine -> (result)

Reads a single character from input stream 0 (the keyboard). The first operand must be 1 (presumably it was provided to support multiple input devices, but only the keyboard was ever used). **time** and **routine** are optional (in Versions 4 and later only) and dealt with as in **read** above.

read_mouse

EXT:22 16 6 read_mouse array

The four words in the **array** are written with the mouse y coordinate, x coordinate, button bits, and a menu word.

The buttons bits are arranged so that the "primary" button is the lowest bit, the "secondary" (if present) is the next lowest bit, and so on, up to a potential 16 buttons. The ordering of buttons should be that which is most natural for the host system. Here are some suggested assignments:

| Platform | Button assignments | | |
|----------|--------------------|-----------|--------------|
| | Bit 0 (low) | Bit 1 | Bit 2 |
| RISC OS | Select | Adjust | Menu |
| MacOS | Primary/only | Secondary | Tertiary ... |
| Windows | Left | Right | Middle |
| X | Left | Right | Middle |

In the menu word, the upper byte is the menu number and the lower byte is the item number (from 0). (Note that the array isn't a table and has no initial size information. The data is written to words 0 to 3 in the array.)

remove_obj

1OP:137 9 remove_obj object

Detach the object from its parent, so that it no longer has any parent. (Its children remain in its possession.)

restart

0OP:183 7 1 restart

Restart the game. (Any "Are you sure?" question must be asked by the game, not the interpreter.) The only pieces of information surviving from the previous state are the "transcribing to printer" bit (bit 0 of 'Flags 2' in the header, at address \$10) and the "use fixed pitch font" bit (bit 1 of 'Flags 2').

In particular, changing the program start address before a restart will not have the effect of restarting from this new address.

restore

0OP:182 6 1 restore ?(label)

0OP:182 5 4 restore -> (result)

EXT:1 1 5 restore table bytes name prompt-> (result)

See **save**. In Version 3, the branch is never actually made, since either the game has successfully picked up again from where it was saved, or it failed to load the save game file.

As with **restart**, the transcription and fixed font bits survive. The interpreter gives the game a way of knowing that a restore has just happened (see **save**).

***[1.0] From Version 5 it can have optional parameters as **save** does, and returns the number of bytes loaded if so. (Whether Infocom intended these options as part of Version 5 is doubtful, but it's too useful a feature to exclude from this Standard.)

If the restore fails, 0 is returned, but once again this necessarily happens since otherwise control is already elsewhere.

restore_undo

EXT:10 A 5 restore_undo -> (result)

Like **restore**, but restores the state saved to memory by **save_undo**. (The optional parameters of **restore** may not be supplied.) The behaviour of **restore_undo** is unspecified if no **save_undo** has previously occurred (and a game may not legally use it): an interpreter might simply ignore this.

ret

1OP:139 B ret value

Returns from the current routine with the value given.

ret_popped

0OP:184 8 ret_popped

Pops top of stack and returns that. (This is equivalent to **ret sp**, but is one byte cheaper.)

rfalse

0OP:177 1 rfalse

Return false (i.e., 0) from the current routine.

rtrue

0OP:176 0 rtrue

Return true (i.e., 1) from the current routine.

save

0OP:181 5 1 save ?(label)

0OP:181 5 4 save -> (result)

EXT:0 0 5 save table bytes name prompt -> (result)

On Versions 3 and 4, attempts to save the game (all questions about filenames are asked by interpreters) and branches if successful. From Version 5 it is a store rather than a branch instruction; the store value is 0 for failure, 1 for "save succeeded" and 2 for "the game is being restored and is resuming execution again from here, the point where it was saved".

It is illegal to use this opcode within an interrupt routine (one called asynchronously by a sound effect, or keyboard timing, or newline counting).

***[1.0] The extension also has (optional) parameters, which save a region of the save area, whose address and length are in bytes, and provides a suggested filename: **name** is a pointer to an array of ASCII characters giving this name (as usual preceded by a byte giving the number of characters). See S 7.6. (Whether Infocom intended these options as part of Version 5 is doubtful, but it's too useful a feature to exclude from this Standard.)

***[1.1] As of Standard 1.1 an additional optional parameter, **prompt**, is allowed on Version 5 extended save/restore. This allows a game author to tell the interpreter whether it should ask for confirmation of the provided file name (**prompt** is 1), or just silently save/restore using the provided filename (**prompt** is 0). If the parameter is not provided, whether to prompt or not is a matter for the interpreter - this might be globally user-configurable. Infocom's interpreters do prompt for filenames, many modern ones do not.

save_undo

EXT:9 9 5 save_undo -> (result)

Like save, except that the optional parameters may not be specified: it saves the game into a cache of memory held by the interpreter. If the interpreter is unable to provide this feature, it must return -1: otherwise it returns the **save** return value.

It is illegal to use this opcode within an interrupt routine (one called asynchronously by a sound effect, or keyboard timing, or newline counting).

(This call is typically needed once per turn, in order to implement "UNDO", so it needs to be quick.)

scan_table

VAR:247 17 4 scan_table x table len form -> (result)

Is **x** one of the words in **table**, which is **len** words long? If so, return the address where it first occurs and branch. If not, return 0 and don't.

The **form** is optional (and only used in Version 5?): bit 7 is set for words, clear for bytes: the rest contains the length of each field in the table. (The first word or byte in each field being the one looked at.) Thus **\$82** is the default.

scroll_window

EXT:20 14 6 scroll_window window pixels

Scrolls the given window by the given number of pixels (a negative value scrolls backwards, i.e., down) writing in blank (background colour) pixels in the new lines. This can be done to any window and is not related to the "scrolling" attribute of a window.

set_attr

2OP:11 B set_attr object attribute

Make **object** have the attribute numbered **attribute**.

set_colour

2OP:27 1B 5 set_colour foreground background 6 set_colour foreground background window

If coloured text is available, set text to be foreground-against-background. (Flush any buffered text to screen, in the old colours, first.) In version 6, the **window** argument is optional and is by default the current window. (This option is supported in Infocom's Amiga and DOS interpreters.)

(One Version 5 game uses this: 'Beyond Zork' (Paul David Doherty reports it as used "76 times in 870915 and 870917, 58 times in 871221") and from the structure of the table it clearly logically belongs in version 5.)

set_cursor

VAR:239 F 4 set_cursor line column 6 set_cursor line column window

Move cursor in the current window to the position (x,y) (in units) relative to (1,1) in the top left. (In Version 6 the window is supplied and need not be the current one. Also, if the cursor would lie outside the current margin settings, it is moved to the left margin of the current line.)

In Version 6, **set_cursor -1** turns the cursor off, and either **set_cursor -2** or **set_cursor -2 0** turn it back on. It is not known what, if anything, this second argument means: in all known cases it is 0.

set_font

EXT:4 4 5 set_font font -> (result)

EXT:4 4 6 set_font font window -> (result)

If the requested font is available, then it is chosen for the current window, and the store value is the font ID of the previous font (which is always positive). If the font is unavailable, nothing will happen and the store value is 0.

If the font ID requested is 0, the font is not changed, and the ID of the current font is returned.

(Infocom's old interpreters did not store 0 for an unavailable font, but the feature is clearly useful and so was introduced in release 0.2 of this Standard.

***[1.1] In Version 6, **set_font** has an optional window parameter, as for **set_colour**. This was part of the original Infocom design, but omitted by earlier Standards. It is reinstated here, as it is useful to be able to measure a font that is about to be used in another window, so that window can be sized before attempting to place the cursor in it. A window number of -3 signifies "the currently selected window"

set_margins

EXT:8 8 6 set_margins left right window

Sets the margin widths (in pixels) on the left and right for the given window (which are by default 0). If the cursor is overtaken and now lies outside the margins altogether, move it back to the left margin of the current line (see S 8.8.3.2.2.1).

set_text_style

VAR:241 11 4 set_text_style style

Sets the text style to: Roman (if 0), Reverse Video (if 1), Bold (if 2), Italic (4), Fixed Pitch (8). In some interpreters (though this is not required) a combination of styles is possible (such as reverse video and bold). In these, changing to Roman should turn off all the other styles currently set.

***[1.1] As of Standard 1.1, it is legal to request style combinations in a single **set_text_style** opcode by adding the values (which are powers of two) together. If the parameter is non-zero, then all the styles given are activated. If the parameter is zero, then all styles are deactivated. If the interpreter is unable to provide the requested style combination, it must give precedence first to the styles requested in the most recent call to **set_text_style**, and within that to the highest bit, making the priority Fixed, Italic, Bold, Reverse.

set_true_colour

EXT:13 D 5/* set_true_colour foreground background

EXT:13 D 6/* set_true_colour foreground background window

The foreground and background are 15-bit colour values:

```
bit 15 = 0
bits 14-10 blue
bits 9-5 green
bits 4-0 red
```

The optional window parameter is only allowed in V6, and operates the same as in **set_colour**.

***[1.1] This opcode will only be present in interpreters obeying Standard 1.1 or later, so story files should check the standard number of the interpreter before executing this opcode.

set_window

VAR:235 B 3 set_window window

Selects the given window for text output.

show_status

0OP:188 C 3 show_status

(In Version 3 only.) Display and update the status line now (don't wait until the next keyboard input). (In theory this opcode is illegal in later Versions but an interpreter should treat it as **nop**, because Version 5 Release 23 of 'Wishbringer' contains this opcode by accident.)

sound_effect

VAR:245 15 5/3 sound_effect number effect volume routine

The given effect happens to the given sound number. The low byte of **volume** holds the volume level, the high byte the number of repeats. (The value 255 means "loudest possible" and "forever" respectively.) (The "repeats" parameter indicates the total number of times to play the sound, not the number of times to repeat it after the first play.) (In Version 3, repeats are unsupported and the high byte must be 0.)

Note that sound effect numbers 1 and 2 are bleeps (see S 9) and in these cases the other operands must be omitted. Conversely, if any of the other operands are present, the sound effect number must be 3 or higher.

The **effect** can be: 1 (prepare), 2 (start), 3 (stop), 4 (finish with).

In Versions 5 and later, the **routine** is called (with no parameters) after the sound has been finished (it has been playing in the background while the Z-machine has been working on other things). (This is used by 'Sherlock' to implement fading in and out, which explains why mysterious numbers like **\$34FB** were previously thought to be to do with fading.) The routine is not called if the sound is stopped by another sound or by an effect 3 call.

See the remarks to S 9 for which forms of this opcode were actually used by Infocom.

In theory, **@sound_effect**; (with no operands at all) is illegal. However interpreters are asked to beep (as if the operand were 1) if possible, and in any case not to halt.

Setting repeats to zero in V5 is illegal - it is suggested that interpreters treat this as a request to play the sound once, and maybe issue a warning.

To clarify:

```
@sound_effect number 3/4
```

will stop (and optionally unload) sound "number" if it is currently playing (or loaded). Otherwise it is ignored.

```
@sound_effect 0 3/4
```

will stop (and unload) all sounds - music and effects.

split_window

VAR:234 A 3 split_window lines

Splits the screen so that the upper window has the given number of lines: or, if this is zero, unsplit the screen again. In Version 3 (only) the upper window should be cleared after the split.

In Version 6, this is supposed to roughly emulate the earlier Version 5 behaviour (see **S 8**), though the line count is in units rather than lines. (Existing Version 6 games seem to use this opcode only for bounding cursor movement. 'Journey' creates a status region which is the whole screen and then overlays it with two other windows.)

Windows 0 and 1 are tiled together to fill the screen, so that window 1 has the given height and is placed at the top left, while window 0 is placed just below it (with its height suitably shortened, possibly making it disappear altogether if window 1 occupies the whole screen).

A cursor remains in the same absolute screen position (which means that its y-coordinate will be different relative to the window origin, since this origin will have moved) unless this position is no longer in the window at all, in which case it is moved to the window origin (at the top left of the window).

sread

This is the Inform name for the keyboard-reading opcode under Versions 3 and 4. (Inform calls the same opcode **aread** in later Versions.) See **read** for the specification.

store

2OP:13 D store (variable) value

Set the **variable** referenced by the operand to **value**.

storeb

VAR:226 2 storeb array byte-index value

array->byte-index = value, i.e. stores the given value in the byte at **address array+byte-index** (which must lie in dynamic memory). (See **loadb**.)

storew

VAR:225 1 storew array word-index value

array-->word-index = value, i.e. stores the given value in the word at **address array+2*word-index** (which must lie in dynamic memory). (See **loadw**.)

sub

2OP:21 15 sub a b -> (result)

Signed 16-bit subtraction.

test

2OP:7 7 test bitmap flags?(label)

Jump if all of the flags in bitmap are set (i.e. if **bitmap & flags == flags**).

"test_array"

See **clear_flag**. (ITF implements this as unconditionally false.)

test_attr

2OP:10 A test_attr object attribute?(label)

Jump if **object** has **attribute**.

throw

2OP:28 1C 5/6 throw value stack-frame

Opposite of **catch**: resets the routine call state to the state it had when the given stack frame value was 'caught', and then returns. In other words, it returns as if from the routine which executed the **catch** which found this stack frame value.

tokenise

VAR:251 1B 5 tokenise text parse dictionary flag

This performs lexical analysis (see **read** above).

The dictionary and flag operands are optional.

If a non-zero **dictionary** is supplied, it is used (if not, the ordinary game dictionary is). If the **flag** is set, unrecognised words are not written into the parse buffer and their slots are left unchanged: this is presumably so that if several **tokenise** instructions are performed in a row, each fills in more slots without wiping those filled by the others.

Parsing a user dictionary is slightly different. A user dictionary should look just like the main one but need not be alphabetically sorted. If the number of entries is given as -n, then the interpreter reads this as "n entries unsorted". This is very convenient if the table is being altered in place: if, for instance, the player is naming things.

verify

00P:189 D 3 verify ?(label)

Verification counts a (two byte, unsigned) checksum of the file from **\$0040** onwards (by taking the sum of the values of each byte in the file, modulo **\$10000**) and compares this against the value in the game header, branching if the two values agree. (Early Version 3 games do not have the necessary checksums to make this possible.)

The interpreter must stop calculating when the file length (as given in the header) is reached. It is legal for the file to contain more bytes than this, but if so the extra bytes should all be 0. (Some story files are padded out to an exact number of virtual-memory pages.) However, many Infocom story files in fact contain non-zero data in the padding, so interpreters must be sure to exclude the padding from checksum calculations.

window_size

EXT:17 11 6 window_size window y x

Change size of window in pixels. (Does not change the current display.)

window_style

EXT:18 12 6 window_style window flags operation

Changes attributes for a given window. A bitmap of attributes is given, in which the bits are: 0 -- keep text within margins, 1 -- scroll when at bottom, 2 -- copy text to output stream 2 (the printer), 3 -- buffer text to word-wrap it between the margins of the window.

The operation, by default, is 0, meaning "set to these settings". 1 means "set the bits supplied". 2 means "clear the ones supplied", and 3 means "reverse the bits supplied" (i.e. eXclusive OR).

16. Font 3 and character graphics

16.1

The following table of 8x8 bitmaps gives a suitable appearance for font 3. The font must have a fixed pitch and characters must be printed immediately next to each other in all four directions.

| | | | |
|------------------|-----------------|-----------------|-----------------|
| 32(): 76543210 | 33(!): 76543210 | 34("): 76543210 | 35(#): 76543210 |
| 0 | 0 | 0 | 0 # |
| 1 | 1 | 1 | 1 # |
| 2 | 2 # | 2 # | 2 # |
| 3 | 3 ## | 3 ## | 3 # |
| 4 | 4##### | 4##### | 4 # |
| 5 | 5 ## | 5 ## | 5 # |
| 6 | 6 # | 6 # | 6 # |
| 7 | 7 | 7 | 7# |
| 36(\$): 76543210 | 37(%): 76543210 | 38(&): 76543210 | 39('): 76543210 |
| 0# | 0 | 0 | 0 |
| 1 # | 1 | 1 | 1 |
| 2 # | 2 | 2 | 2 |
| 3 # | 3 | 3 | 3##### |
| 4 # | 4 | 4##### | 4 |
| 5 # | 5 | 5 | 5 |
| 6 # | 6 | 6 | 6 |
| 7 # | 7 | 7 | 7 |
| 40(:): 76543210 | 41(): 76543210 | 42(*): 76543210 | 43(+): 76543210 |
| 0 # | 0 # | 0 # | 0 |
| 1 # | 1 # | 1 # | 1 |
| 2 # | 2 # | 2 # | 2 |
| 3 # | 3 # | 3##### | 3 |
| 4 # | 4 # | 4 | 4##### |
| 5 # | 5 # | 5 | 5 # |
| 6 # | 6 # | 6 | 6 # |
| 7 # | 7 # | 7 | 7 # |
| 44(,): 76543210 | 45(-): 76543210 | 46(.): 76543210 | 47(/): 76543210 |
| 0 # | 0 # | 0 # | 0 |
| 1 # | 1 # | 1 # | 1 |
| 2 # | 2 # | 2 # | 2 |
| 3 # | 3 # | 3 # | 3 ##### |
| 4 #### | 4#### | 4 ##### | 4 # |
| 5 # | 5 # | 5 | 5 # |
| 6 # | 6 # | 6 | 6 # |
| 7 # | 7 # | 7 | 7 # |
| 48(0): 76543210 | 49(1): 76543210 | 50(2): 76543210 | 51(3): 76543210 |
| 0 | 0 # | 0 # | 0# |
| 1 | 1 # | 1 # | 1 # |
| 2 | 2 # | 2 # | 2 # |
| 3##### | 3 # | 3 # | 3 ##### |
| 4 # | 4##### | 4 ##### | 4 # |
| 5 # | 5 | 5 # | 5 # |
| 6 # | 6 | 6 # | 6 # |
| 7 # | 7 | 7# | 7 # |

| | | | |
|-----------------|-----------------|-----------------|-----------------|
| 52(4): 76543210 | 53(5): 76543210 | 54(6): 76543210 | 55(7): 76543210 |
| 0 # | 0 # | 0##### | 0##### |
| 1 # | 1 # | 1##### | 1##### |
| 2 # | 2 # | 2##### | 2##### |
| 3##### | 3 # | 3##### | 3##### |
| 4 # | 4##### | 4##### | 4##### |
| 5 # | 5 # | 5##### | 5 |
| 6 # | 6 # | 6##### | 6 |
| 7 # | 7 # | 7##### | 7 |
| 56(8): 76543210 | 57(9): 76543210 | 58(:): 76543210 | 59(;): 76543210 |
| 0 | 0##### | 0 ##### | 0 # |
| 1 | 1##### | 1 ##### | 1 # |
| 2 | 2##### | 2 ##### | 2 # |
| 3##### | 3##### | 3 ##### | 3##### |
| 4##### | 4##### | 4 ##### | 4##### |
| 5##### | 5##### | 5 ##### | 5##### |
| 6##### | 6##### | 6 ##### | 6##### |
| 7##### | 7##### | 7 ##### | 7##### |
| 60(<): 76543210 | 61(=): 76543210 | 62(>): 76543210 | 63(?): 76543210 |
| 0##### | 0##### | 0 ##### | 0 ##### |
| 1##### | 1##### | 1 ##### | 1 ##### |
| 2##### | 2##### | 2 ##### | 2 ##### |
| 3##### | 3##### | 3 ##### | 3 ##### |
| 4##### | 4##### | 4##### | 4 ##### |
| 5 # | 5##### | 5 ##### | 5 |
| 6 # | 6##### | 6 ##### | 6 |
| 7 # | 7##### | 7 ##### | 7 |
| 64(@): 76543210 | 65(A): 76543210 | 66(B): 76543210 | 67(C): 76543210 |
| 0 | 0 | 0##### | 0 ##### |
| 1 | 1 | 1##### | 1 ##### |
| 2 | 2 | 2##### | 2 ##### |
| 3 ##### | 3##### | 3##### | 3 ##### |
| 4 ##### | 4##### | 4##### | 4 ##### |
| 5 ##### | 5##### | 5 | 5 # |
| 6 ##### | 6##### | 6 | 6 # |
| 7 ##### | 7##### | 7 | 7# |
| 68(D): 76543210 | 69(E): 76543210 | 70(F): 76543210 | 71(G): 76543210 |
| 0# | 0 # | 0##### | 0 # |
| 1 # | 1 # | 1##### | 1 |
| 2 # | 2 # | 2##### | 2 |
| 3 ##### | 3##### | 3##### | 3 |
| 4 ##### | 4##### | 4##### | 4 |
| 5 ##### | 5##### | 5 # | 5 |
| 6 ##### | 6##### | 6 # | 6 |
| 7 ##### | 7##### | 7 # | 7 |
| 72(H): 76543210 | 73(I): 76543210 | 74(J): 76543210 | 75(K): 76543210 |
| 0 | 0 | 0# | 0##### |
| 1 | 1 | 1 | 1 |
| 2 | 2 | 2 | 2 |
| 3 | 3 | 3 | 3 |
| 4 | 4 | 4 | 4 |
| 5 | 5 | 5 | 5 |
| 6 | 6 | 6 | 6 |
| 7 # | 7# | 7 | 7 |

| | | | |
|-----------------|-----------------|-----------------|-----------------|
| 76(L): 76543210 | 77(M): 76543210 | 78(N): 76543210 | 79(O): 76543210 |
| 0 | 0# | 0 # | 0 |
| 1 | 1# | 1 # | 1##### |
| 2 | 2# | 2 # | 2 |
| 3 | 3# | 3 # | 3 |
| 4 | 4# | 4 # | 4 |
| 5 | 5# | 5 # | 5 |
| 6 | 6# | 6 # | 6##### |
| 7##### | 7# | 7 # | 7 |
| 80(P): 76543210 | 81(Q): 76543210 | 82(R): 76543210 | 83(S): 76543210 |
| 0 | 0 | 0 | 0 |
| 1##### | 1##### | 1##### | 1##### |
| 2# | 2## | 2### | 2#### |
| 3# | 3## | 3### | 3#### |
| 4# | 4## | 4### | 4#### |
| 5# | 5## | 5### | 5#### |
| 6##### | 6##### | 6##### | 6##### |
| 7 | 7 | 7 | 7 |
| 84(T): 76543210 | 85(U): 76543210 | 86(V): 76543210 | 87(W): 76543210 |
| 0 | 0 | 0 | 0 |
| 1##### | 1##### | 1##### | 1##### |
| 2##### | 2##### | 2##### | 2##### |
| 3##### | 3##### | 3##### | 3##### |
| 4##### | 4##### | 4##### | 4##### |
| 5##### | 5##### | 5##### | 5##### |
| 6##### | 6##### | 6##### | 6##### |
| 7 | 7 | 7 | 7 |
| 88(X): 76543210 | 89(Y): 76543210 | 90(Z): 76543210 | 91([): 76543210 |
| 0 | 0 | 0# # | 0 # |
| 1 # | 1# | 1 # # | 1 # |
| 2 # | 2# | 2 # # | 2 # |
| 3 # | 3# | 3 ## | 3 # |
| 4 # | 4# | 4 ## | 4##### |
| 5 # | 5# | 5 # # | 5 # |
| 6 # | 6# | 6 # # | 6 # |
| 7 | 7 | 7# # | 7 # |
| 92(\): 76543210 | 93(]): 76543210 | 94(^): 76543210 | 95(_): 76543210 |
| 0 ## | 0 ## | 0 ## | 0##### |
| 1 #### | 1 ## | 1 #### | 1# # |
| 2## ## ## | 2 ## | 2## ## ## | 2# # |
| 3 ## | 3 ## | 3 ## | 3# # |
| 4 ## | 4## ## ## | 4## ## ## | 4# # |
| 5 ## | 5 ##### | 5 ##### | 5# # |
| 6 ## | 6 ## | 6 ## | 6# # |
| 7 | 7 | 7 | 7##### |
| 96('): 76543210 | 97(a): 76543210 | 98(b): 76543210 | 99(c): 76543210 |
| 0 #### | 0## # | 0 ## | 0 # |
| 1 ## ## | 1# # # | 1 # # | 1 ## |
| 2 ## | 2# # | 2 # # | 2 # # |
| 3 ## | 3## | 3 ### | 3# # # |
| 4 ## | 4# # | 4 # # | 4 # # |
| 5 | 5# # | 5 # # | 5 ## |
| 6 ## | 6# | 6 ## | 6 # |
| 7 | 7 | 7 | 7 |

| | | | |
|------------------|------------------|------------------|------------------|
| 100(d): 76543210 | 101(e): 76543210 | 102(f): 76543210 | 103(g): 76543210 |
| 0# # | 0# # | 0# # # | 0# # |
| 1## ## | 1## ## | 1# # # | 1 # # |
| 2# # # # | 2# # # # | 2## # | 2 # # |
| 3# # # | 3# # # | 3# # | 3 # |
| 4# # # # | 4# # | 4## | 4 # # |
| 5## ## | 5# # | 5# | 5 # # |
| 6# # | 6# # | 6# | 6# # |
| 7 | 7 | 7 | 7 |
| 104(h): 76543210 | 105(i): 76543210 | 106(j): 76543210 | 107(k): 76543210 |
| 0## # | 0 # | 0 # | 0 # |
| 1# # # | 1 # | 1 ### | 1 # |
| 2## # # | 2 # | 2 # # # | 2 # |
| 3# # # # | 3 # | 3# # # | 3 ### |
| 4# # ## | 4 # | 4 # # # | 4 # # # |
| 5# # # | 5 # | 5 ### | 5# # # |
| 6# ## | 6 # | 6 # | 6# # # |
| 7 | 7 | 7 | 7 |
| 108(l): 76543210 | 109(m): 76543210 | 110(n): 76543210 | 111(o): 76543210 |
| 0 # | 0## ## | 0# # | 0## # |
| 1 ## | 1# # # # | 1 # # | 1# # ## |
| 2 # # | 2# # # | 2 ### | 2## # # |
| 3 # # | 3# # # # | 3 # # | 3# # # |
| 4 # | 4## ## | 4 # # # | 4# # |
| 5 # | 5# # | 5 # | 5# |
| 6 # | 6# # | 6 # | 6# |
| 7 | 7 | 7 | 7 |
| 112(p): 76543210 | 113(q): 76543210 | 114(r): 76543210 | 115(s): 76543210 |
| 0# | 0 # | 0 ## | 0 # |
| 1# | 1 # | 1 # # | 1 # # |
| 2# | 2 # | 2 # # | 2 # ## |
| 3# # | 3 ##### | 3 # # | 3 # # # |
| 4# # # | 4 # # | 4 ## | 4 ## # |
| 5## # | 5 # # | 5 # # | 5 # # |
| 6# # | 6 # # | 6 # # | 6 # # |
| 7 | 7 | 7 | 7 |
| 116(t): 76543210 | 117(u): 76543210 | 118(v): 76543210 | 119(w): 76543210 |
| 0 # | 0 ## | 0 # | 0 ## |
| 1 ### | 1 # # | 1# ### # | 1 # # |
| 2 # # # | 2 # # | 2 # # # | 2 # # |
| 3# # # | 3 # # | 3 # | 3 # # |
| 4 # | 4 # # | 4 # | 4 ## |
| 5 # | 5 # # | 5 # | 5 # |
| 6 # | 6 # # | 6 # | 6 # |
| 7 | 7 | 7 | 7 |
| 120(x): 76543210 | 121(y): 76543210 | 122(z): 76543210 | 123({): 76543210 |
| 0# # # | 0### | 0 # | 0### ### |
| 1 # # # | 1## # | 1 # # | 1## ## |
| 2 ### | 2# # # | 2 # # | 2 # # |
| 3 # | 3# # # | 3 # # | 3### ### |
| 4 # | 4# ## # | 4 # | 4### ### |
| 5 # | 5# # ## | 5 # # | 5### ### |
| 6 # | 6# # # | 6 # # | 6### ### |
| 7 | 7 | 7 | 7##### |

| | | | | | |
|---------|----------|----------|----------|---------|-----------|
| 124(): | 76543210 | 125({}): | 76543210 | 126(~): | 76543210 |
| | 0### ### | | 0### ### | | 0## ## |
| | 1### ### | | 1## ## | | 1# ## # |
| | 2### ### | | 2 # # | | 2##### # |
| | 3### ### | | 3### ### | | 3##### ## |
| | 4 # # | | 4 # # | | 4### ### |
| | 5## ## | | 5## ## | | 5##### |
| | 6### ### | | 6### ### | | 6### ### |
| | 7##### | | 7##### | | 7##### |

Remarks

Two different versions of font 3 were supplied by Infocom, which we shall call the Amiga and PC forms (the Atari form is the same as for the PC). The arrow shape differed slightly and so did the rune alphabet. Each was an attempt to map the late Anglian ("futhorc") runic alphabet, which has 33 characters, onto our Latin alphabet. The drawings above are from the Amiga set.

Most of the mappings are straightforward (e.g., Latin A maps to Anglian a), except that: Latin C is mapped to Anglian eo; K to "other k" (previously a z sound); Q to Anglian k (the same rune as c); V to ea; X to z and Z to oe. The PC runes differ as follows: G has an ornamental circle making it more look like "other z"; K maps to Anglian k (or c); Q is an Anglian ea (which resembles the late Anglian q); V is an oe; X is an "other k" and Z is a symbol Infocom seem to have invented themselves. (Though less well drawn the PC runes arguably have a better sound-mapping.)

The font behaviour of 'Beyond Zork', which does have bit 3 of 'Flags 2' set, is rather complicated and depends on the interpreter number it finds in the header (see S 11). Specifically:

1. (a Digital terminal) 'BZ' asks whether the player has a VT220 terminal (a model capable of character graphics) and uses font 3 if and only if the answer is yes. (An in-house convenience: Infocom used a Digital mainframe.)
2. (Apple IIe) 'BZ' never uses font 3.
3. (Macintosh) 'BZ' always uses font 3.
4. (Amiga) 'BZ' always uses font 3.
5. (Atari ST) 'BZ' always uses font 3.
6. (MSDOS) 'BZ' uses font 3 if it finds bit 3 of 'Flags 2' set (indicating that a graphical screen mode is in use) and otherwise uses IBM PC graphics codes. These need to be converted back into ASCII. The conversion process used by the Zip interpreter is as follows:

| | | |
|--|---------|-------------------------------|
| 179 | becomes | a vertical stroke (ASCII 124) |
| 186 | | a hash (ASCII 35) |
| 196 | | a minus sign (ASCII 45) |
| 205 | | an equals sign (ASCII 61) |
| all others in the range 179 to 218 become a plus sign (ASCII 43) | | |

7. (Commodore 128) 'BZ' always uses font 3.
8. (Commodore 64) 'BZ' always uses font 3.
9. (Apple IIc) 'BZ' uses Apple character graphics (possibly "Mousetext"), but has problems when the units used are not \$1\times 1\$.

10. (Apple IIgs) 'BZ' always uses font 3.
11. (Tandy) 'BZ' crashes on the public interpreters.

A similarly tangled process is used in 'Journey'. It is obviously highly unsatisfactory to have to make the decision in the above way, which is why **set_font** is now required to return 0 indicating non-availability of a font.

Stefan Jokisch suggests that Infocom originally intended the graphics bit as a way to develop Version 5 to allow a graphical version in parallel with the normal text one. For instance, when the Infocom MSDOS interpreter starts up, it looks at the graphics flag and:
if clear, it sets the font width/height to 1/1 (so that screen units are character positions);
if set, it enters MGCA, a graphical screen mode and sets the font width/height to 8/8 (so that screen units are pixels).

The "COLOR" command in 'BZ' (typed at the keyboard) also behaves differently depending on the interpreter number, which is legal behaviour and has no impact on the specification.

Appendix A. Error messages and debugging

Older interpreters, such as **ITF**, are extremely curt when an error condition is reached (for example, an illegal opcode). It was assumed that Infocom's shipped story files were bug-free, which is mostly true, so that errors could only arise through a bug elsewhere in the interpreter.

In debugging Inform games, though, many error conditions can arise and it is extremely helpful to report these as fully as possible. These include:

1. An illegal opcode being hit;
2. A call to what can't be a routine (because the initial byte is not between 0 and 15);
3. A jump or call to an address beyond the size of the story file;
4. An attempt to **print_obj**, or otherwise access, an object which doesn't exist, such as object number 0.
5. An attempt to write to, or get the property length of, a nonexistent property.
6. An attempt to access an attribute outside the range 0 to 31 or 0 to 47 (depending on Version). (But note that Infocom's 'Sherlock' contains a bug causing it to try setting or clearing attribute number 48.)
7. Division by zero. The player sometimes then has the annoying task of working out where the error took place in source code. Providing a stack back-trace would be a help.

As mentioned in **S 3**, it's helpful to screen out any illegal ZSCII characters between 0 and 31 which are accidentally printed: crashes can be very mysterious when they cause interpreters to send control codes to the terminal.

In addition, an interpreter might provide options for keeping track of maximum stack usage and the typical number of opcodes executed between each read from the keyboard. (But watching these is a time-wasting activity, so they should be options.)

Finally, infinite loops fairly often happen, as in any programming language. On a system without pre-emptive multi-tasking, this may lock up the whole machine, as the usual way that porters implement multi-tasking is to return control to the host operating system only when the keyboard is read. This can be avoided by providing a point in the code which could return control to the OS from time to time (say, every 2000 instructions).

A number of post-Infocom games have been released which contain errors, most often trying to perform illegal operations on object 0. Many interpreters silently ignored these errors, which can make it very to notice and track down bugs.

It is desirable for modern interpreters to be able to notify players about these bugs, but this can also ruin gameplay. It is highly recommended, then, that interpreters have four levels of error checking, selectable by the user (through a command-line or menu option, or similar):

- Never report the bug.
- Report the first instance of each type of error.
- Report every error.
- Fatal error and close the interpreter on any error.

Appendix B. Conventional contents of the header

The header table in section 11 details everything the interpreter needs to know about the header's contents. A few other slots in the header are meaningful but only by convention (since existing Infocom and Inform games write them). These additional slots are described here.

As in S 11, "Hex" means the address, in hexadecimal; "V" the earliest version in which the feature appeared; "Dyn" means that the byte or bit may legally be changed by the game during play (no others may legally be changed by the game); "Int" means that the interpreter may (in some cases must) do so.

Conventional usage of unspecified header memory

| Hex | V | Dyn | Int | Contents |
|-----|---|-----|-----|---|
| 1 | 1 | | | <i>Flags 1:</i> |
| | 3 | | * | The legendary "Tandy" bit (see note) |
| 2 | 1 | | | Release number (word) |
| 10 | 1 | * | | <i>Flags 2:</i> |
| | 3 | | | Bit 4: Set in the Amiga version of The Lurking Horror so presumably to do with sound effects? |
| | ? | ? | * | 10: Possibly set by interpreter to indicate an error with the printer during transcription |
| 12 | 2 | | | Serial code (six characters of ASCII) |
| | 3 | | | Serial number (ASCII for the compilation date in the form YYMMDD) |
| 38 | 6 | | * | Infocom used this as 8 bytes of ASCII: the player's user-name on their own mainframe, useful to identify which person played a saved-game (however, the bytes are all 0 in shipped story files) |
| 3C | | | | Inform 6 stores 4 bytes of ASCII here, giving the version of Inform used to compile the file: for instance, "6.11". |

1. In Versions 1 to 3, bits 0 and 7 of 'Flags 1' are unused. (The meaning of bit 2 has recently been discovered: see section 11.) In later Versions, bits 0, 6 and 7 are unused. In 'Flags 2', bits 9 and 11-15 are unused. Infocom used up almost the whole header: only the bytes at **\$32** and **\$33** are unused in any Version, and those are now allocated for standard interpreters to give their Revision numbers.

2. Some early Infocom games were sold by the Tandy Corporation, who seem to have been sensitive souls. 'Zork I' pretends not to have sequels if it finds the Tandy bit set. And to quote Paul David Doherty:

In 'The Witness', the Tandy Flag can be set while playing the game, by typing **\$DB** and then **\$TA**. If it is set, some of the prose will be less offensive. For example, "private dicks" become "private eyes", "bastards" are only "idiots", and all references to "slanteyes" and "necrophilia" are removed.

We live in an age of censorship.

3. For comment on interpreter numbers, see S 11. Infocom's own interpreters were generally rewritten for each of versions 3 to 6. For instance, interpreters known to have been shipped with the Macintosh gave version letters B, C, G, I (Version 3), E, H, I (Version 4), A, B, C (Version 5) and finally 6.1 for Version 6. (Version 6 interpreters seem to have version numbers rather than letters.) See the "Infocom fact sheet" for fuller details.
4. Inform 6 story files are easily distinguished from all other story files by their usage of the last four header bytes. Inform 1 to 5 story files are best distinguished from Infocom ones by the serial code date: anything before 930000 is either an Infocom file, or a fake. (The author of 'Jigsaw' is tempted to compile a millennial version with serial code 991231 when the time comes, but then the next day serial codes will clock over to 000101. The decision of how to continue serial codes past the year 2079 is deferred to a future revision of this Standard.) Clearly there is no point going to any trouble to prevent fakes, but with a little practice it's easy to tell whether Zilch or Inform compiled a file from the style of code generated.

Appendix C. Resources available

...the dead hand of the academy had yet to stifle the unbridled enthusiasms of a small band of amateurs in Europe and America.

Michael D. Coe, *Breaking the Maya Code*

[Note: the hypertext links in this appendix are to WWW sites external to the standards document, and were correct as of 21th February 2014.]¹

The resources below are mainly available from the if-archive maintained by David Kinder and Stephen Granada.²

Public Interpreters

Since the 1.0 Standard was first introduced, a variety of interpreters have become available, for many different platforms and with a variety of features. The page on Playing Interactive Fiction at inform7.com is a good place to start when looking for an interpreter, as is the ifwiki page on the Z-Machine.

Testing compliance

A few story files exist to test interpreter for various Z-Machine features:

1. **Czech** is a general Z-Machine unit test by Evin Robertson and Amir Karger.
2. **Praxis** is a general Z-Machine unit test by Andrew Plotkin and Dannii Willis.
3. **TerpEtude** by Andrew Plotkin is a unit test with a focus on IO.
4. **Unicode** by David Kinder tests Unicode input and output.
5. **Strict Z Test** by Torbjorn Andersson tests an interpreter's ability to handle errors in story files.

Compilers

Infocom's original compiler **Zilch** no longer exists, but a PDF copy of Learning ZIL, Infocom's documentation for their **ZIL** language, is available. Since the emergence of this file, there have been a few attempts at making a new **ZIL** compiler, most notably **ZILF** by Jesse McGrew.³

No continuous part of the source code of any of Infocom's games is in the public domain [but see Stu Galley's chapter of an Infocom history article, and the IEEE article, for fragments].⁴

¹ There is no links in this document.

² 2020 the if-archive is reachable at www.ifarchive.org

³ ZILF can now (2020) compile most of Infocom's ZIL source.

⁴ The source to Inocom's games is, since 2019, publically available.

Inform remains the most popular compiler for Z-Machine games. In 2006, the first public release of Inform 7 was made available. Inform 7 is a very different language from previous versions of Inform, with a focus on natural language and a powerful IDE.

While Inform 7 is the main focus of development, Inform 6 is still widely used and the compiler and libraries continue to be maintained.

A few other compilers have been made available over the years.

Utility programs

Mark Howell has written a toolkit of **Ztools**, or utility programs (1991-5, updated 1997), which includes:

1. **Txd**, a disassembler for Versions 1 to 8. (Uses the same opcode names as Inform and this document, and has an option to disassemble in Inform assembly-language syntax.)
2. **Infodump**, capable of printing the header information, object tree (with properties and attributes), dictionary and grammar tables of any Infocom or Inform-compiled game. (Understands four varieties of grammar table: Infocom pre-Version 6, Infocom Version 6, Inform GV1 and GV2.)
3. **Pix2gif**, for converting Version 6 picture data to GIF files.
4. **Check**, for verifying Infocom or Inform story files.

These continue to be maintained (by Matthew Russotto) and the first two are extremely useful. **Infodump** largely supersedes Mike Threepoint's vocabulary dumper **Zorkword** (1991-2), which was important in its day (and which this author found extremely helpful when writing Inform 1).

Reform, by Ben Rudiak-Gould, is a decompiler for Z-Machine games which will attempt to convert them into Inform 6 source code.

Story files

1. Numerous Inform-compiled story files are publically available.
2. A few Infocom story files are public, notably two 4-in-1 sample games (released for advertising purposes: 55.850823 and 97.870601) and Minizork (a heavily abbreviated form of Zork I released with a Commodore magazine).
3. Almost all Infocom's games remain commercially available in anthologies published by Activision. Copyright resides in them and they should not be available by FTP from any site.
4. A few other Infocom story files have existed but are neither on sale nor released from copyright: this applies to several of the Version 6 games, those games involving literary rights or other legal issues ('Shogun', 'Hitch-Hiker's Guide To The Galaxy') and ephemera such as beta-test versions (notably the German version of 'Zork I') which have somehow passed into private circulation.

Most of the Infocom games exist in several different releases, and some were written for one Version and then ported to later ones. 'Zork I', for instance, has at least 11 releases, 2

early, 8 in Version 3 (with release numbers between 5 to 88 in chronological order) and one in Version 5 (release 52 -- the releases go back to 1 when the version changes).

Version 1 and 2 games are extinct, though there are a few fossils in the hands of collectors.

Documents

The definitive guide to all Infocom story files known to exist, and an indispensable reference for anyone interested in Infocom, is Paul David Doherty's Infocom fact sheet, which is regularly updated, concise and precise. This supersedes Paul Smith's "Infocom Game Information" file.

Stefan Jokisch has written a brief specification of Infocom-format sound effects files.

Martin Frost is the author of the **Quetzal** standard for saved-game files. Patches to adapt **Zip**-based interpreters to use **Quetzal** are available.

Andrew Plotkin is the author of the **Blorb** standard for packaging up images and sounds with Z-machine games.

The Inform Technical Manual documents the format of parsing tables used in Inform games.

The documentation for Infocom's games is available online (with permission from Activision) courtesy of the Infocom Documentation Project. In addition, an archive of the "samplers" is publically available, as is an interesting historical archive of magazine articles concerning Infocom, and articles from Infocom's own publicity magazine [indexed here].

Appendix D. A short history of the Z-machine

Infocom made six main Versions of the Z-machine and several minor variant forms. These are recognisably similar but with labyrinthine differences, like different archaic dialects of the same language. (The archaeological record stops sharply in 1989 when the civilisation in question collapsed.)

Broadly, these fall into two groups: early (Versions 1 to 3) and late (4 to 6). More fully:

| | |
|-----------|--|
| Version 1 | Early Apple II and TRS-80 Model I games |
| Version 2 | Early Apple II and TRS-80 Model I games |
| Version 3 | "Standard" series games |
| Version 4 | "Plus" series games |
| Version 5 | "Advanced" series games, or, as the marketing division would have it, "Solid Gold Interactive Fiction" -- a reference to the colour (though not composition) of the boxes they came in |
| Version 6 | Later games with graphics, mouse support, etc. |

Infocom called their own interpreters ZIP (versions 1 to 3), EZIP/LZIP (V4), XZIP (V5) and YZIP (V6). They speculated on the possibility of an interpreter capable of running all Versions, but never published one.

The original purpose of the Z-machine was simply to implement as much as possible of the mainframe game "Zork" on the first popular wave of home computers.

(Apparently "zork" was a nonsense word used at MIT for the current uninstalled program in progress, and stuck. Just as this document uses the term "Z-machine" for both the machine and its loaded program (which is also sometimes called the "story file"), so ZIP (Zork Implementation Program) was used to mean either the interpreter or the object code it interpreted. Code was written in ZIL (Zork Implementation Language), which was derived from MDL (informally called "muddle"), a particularly unhelpful form of LISP. It was then compiled by ZILCH to assembly code which was passed to ZAP to make the ZIP.)

The Z-machine as originally constructed was surprisingly similar to that in use today. Version 1 (by Joel Berez and Marc Blank, in Autumn 1979) contained essentially all of the main architecture: the header, the memory divided into three, the variables and stack, the object tree, the dictionary, the instruction format. It used "shift lock" characters (a text compression trick which did not survive, though it was more efficient on long sequences of capital letters or punctuation characters than the technique which replaced it). The first micro interpreters were for the TRS-80 Model I (by Scott Cutler) and the Apple II (by Bruce K. Daniels). (A TRS-80 Model II interpreter was written but never actually shipped.)

Version 2 was only a minor enhancement. Abbreviations (used to help text compression) appeared, but only in one 32-word bank, and the six-digit serial number appeared in the header, though it wasn't always the date in those days: Release 7 of 'Zork II', for instance, is numbered **UG3AU5**. (Other bizarre serial numbers, such as **000000**, appear on fakes or beta-test releases.)

In Version 3, the text encoding alphabets changed again, and the old "shift lock" codes were dropped in favour of expanding the abbreviations bank to 96 entries. The "verify" opcode and checksums appeared; and a new opcode to reprint the status line at the top of the screen was introduced. (Previously, this had been updated only when input was taken from the keyboard.) The earliest Version 3 releases ('Deadline', then reissues of 'Zork I' and 'II') were in March and April 1982; the last (the 'Minizork', a cassette-based Commodore-64 sample of 'Zork') was in November 1987.

The idea of widespread portability finally came of age as (between 1982 and 1985) interpreters were developed for the Atari 400/800, CP/M, the IBM PC, the TRS-80 Model III, the NEC APC, the DEC Rainbow, the Commodore 64, the TI Professional, the DECmate, the Tandy-2000, the Kaypro II, the Osborne 1, MS-DOS, the TI 99/4a, the Apple Macintosh, the Epson QX-10, the Apricot, the Atari ST and the Amiga. Infocom's middle period coincided with the bubble in home computers, before the market collapsed to its present apparently stable state (in which IBM and Apple share almost the entire market), and the Z-machine's portability gave Infocom a unique advantage over its competitors. Also, it was an expertly marketed quality brand at a time when standards of workmanship were very variable; and text-only games did not seem so dull at a time when graphics were on the whole crude and slow. These factors combined to give Infocom considerable (though never enormous) commercial success.

By 1982, then, the Z-machine had stabilised to a clean design which was to remain in use for six years. It was very portable, contained everything reasonably necessary and most of its complications were badly-needed space optimisations. (Although Version 3 can fit 128K of story file, the practical limit in 1982-4 was about 110K, that being the typical disc capacity on target machines.) The ZAP assembler was cleverly written to exploit these optimisations, though the Zilch compiler's code generator was much less efficient. (Interestingly, Infocom did not develop any generic central library, and Infocom's authors worked fairly independently of each other: each new game would inherit a small core of code from a previous one, but this would make up only about 10K of code (about a third of the size of the Inform library) and would end up being hacked about to suit the new game. Without a central library, Infocom games waste a fair amount of space in duplicating code for routine operations many times over. For this reason, Inform games tend to squash appreciably more design into the format.)

"Verify" and checksum data were quickly introduced. However, the first serious variant on Version 3 was made in 1984 when a primitive form of screen-splitting was invented to give 'Seastalker' a sonar display. This design (perhaps accidentally) became the foundation for the graphics systems of later versions.

Much later (in 1987) sound effects were added to Version 3 for 'The Lurking Horror', though by that time it was really a Version 5 feature being passed down to the old model (and only to the Amiga interpreter in any case). ('TLH' is contemporaneous with 'Sherlock' (in Version 5), the only other game to actually use the sound effects features.)

During 1983-5, Infocom poured resources into an ambitious pet project of its founders: 'Cornerstone', a database which used some of the same portable virtual machine ideas as the Z-machine. The business market, however, was not nearly as diverse as the home computer market: 'Cornerstone' probably was the best database available on the Atari ST, but it made no impression on the IBM PC market. The result was a commercial failure which compounded the company's over-expansion problems (driving it into a merger with Activision), though it certainly did not destroy Infocom's viability.

By 1985, Infocom had begun to write interpreters in C for the sake of portability (previously, a different assembly-language program had to be maintained for every model of computer). The main motivation to keep the format stable was therefore largely removed: it became possible to upgrade the Z-machine for every new game, if need be.

There were two basic pressures for change. One was that home computers were larger, and several fundamental restrictions (the game size being only 128K, the number of objects only 255, the attributes only 32, the properties only 31) were beginning to bite. The other was the drive for more gimmicks - character graphics, flashier status lines, sound effects, different typefaces, and so on. The former led to logical, easy to understand structural changes in the machine (designed by Marc Blank). The latter, in contrast, made a mess of the system of opcodes (designed by committee).

More does not mean better (halving the price of paper does not double the quality of the novel). The relieving of size restrictions only increased design time -- or endangered the quality of the designs being produced. The Version 3 games have a spare, concise literary style which is absent from the later games. (But Inform authors have certainly found Version 3 slightly too small for comfort, and it's useful to be able to spill over its boundaries.)

In August the first Version 4 game ('A Mind Forever Voyaging') reached production. Opinions vary as to whether it was brilliant or awful, but it was certainly a departure (and could not have been written under Version 3). In retrospect there is no doubt about 'Trinity', now generally considered the finest game written: it had previously been shelved as too ambitious for the Version 3 format. Still, most of the new 1985/6 games remained in Version 3: there were still plenty of 8-bit home computers around which were too small for Version 4 games. Despite critical acclaim, the new games consequently did not sell as well. (Brian Moriarty commented that 'Trinity' "sold tolerably well. Better than we'd hoped." But his previous game, the more modest 'Wishbringer', had sold rather better.)

Version 5 games began to appear in September 1987 with 'Beyond Zork' and 'Border Zone'. Both of these games needed new features -- character graphics run wild in the case of the former, and real-time keyboard interaction in the latter. The number of opcodes grew ever faster as a result.

Although five old games were re-released in Version 5 editions (with an in-game hints system added, and benefiting from 9-letter word dictionaries, but otherwise as written), the direction was all too clearly away from the old text game into graphics: 'Beyond Zork' can look like a parody of an early mainframe maze game, for instance. Version 6 completed the process during something of a hiatus in 1988, after which the last few increasingly-unrecognisable Infocom games appeared: 'Zork Zero', 'Shogun', 'Journey' and 'Arthur'.

It would be wrong, though, to suggest that Infocom regarded text and graphics as incompatible opposites. Infocom had never been puritanically opposed to graphics --

We have nothing against graphics per se. However, given the quality of graphics currently available on home computers, we would rather use that disk space for additional puzzles and richer descriptions.

The New Zork Times (Spring 1984)

(and, after all, the same author wrote both 'Trinity' and 'Beyond Zork'). Although the old Infocom parser was considered to have passed its sell-by date, Version 6 did not drop textual input in favour of some inane point-and-click interface. Instead, an entirely new parser was devised from scratch ("using the theory of computational linguistics", according to a puff by Stu Galley: broadly an LALR(1) parser).

Infocom gradually ceased to exist during 1987-9 as its financial problems grew. But its products were increasingly regarded as anachronistic and most of its staff had left since the middle years: if Infocom had not finally been wound up, whether it would have continued to release text games of the classical style is arguable.

Two new formats, versions 7 and 8, have recently been devised to cope with large Inform games.

Appendix E. Statistics

LORD DIMWIT FLATHEAD: "It must have two hundred thousand rooms, four million takeable objects, and understand a vocabulary of every single word ever spoken in every language ever invented."

The New York Times (Winter 1984)

[Note: the information below has not been updated since the 20th June 1997 revision of this document.]

To give some idea of the sizes found in typical story files, here are a few statistics, mostly gathered by Paul David Doherty, whose "Infocom fact sheet" file is the definitive reference.

(i) *Length*

The shortest files are those dating from the time of the 'Zork' trilogy, at about 85K; middle-period Version 3 games are typically 105K, and only the latest use the full memory map. In Versions 4 and 5, only 'Trinity', 'A Mind Forever Voyaging' and 'Beyond Zork' use the full 256K. 'Border Zone' and 'Sherlock', for instance, are about 180K. (The author's short story 'Balances' is about 50K, an edition of 'Adventure' takes 80K, and 'Curses' takes 256K (it's padded out to the maximum size with background information; the actual game comprises only about 245K). Under Inform, the library occupies about 35K regardless of the size of game.)

(ii) *Code size*

'Zork I' uses only about 5500 opcodes, but the number rises steeply with later games; 'Hollywood Hijinx' has 10355 and, e.g. 'Moonmist' has 15900 (both these being Version 3). Against this, 'A Mind Forever Voyaging' has only 18700, and only 'Trinity' and 'Beyond Zork' reach 32000 or so. (Inform games are more efficiently compiled and make better use of common code -- the library -- so perform much better here: the old Version 3, release 10 of 'Curses' (128K long, and a larger game than any Infocom Version 3 game) has only 6720 opcodes.)

(iii) *Objects and rooms*

This varies greatly with the style of game. 'Zork I' has 110 rooms and 60 takeable objects, but several quite complex games have as few as 30 rooms (the mysteries, or 'Hitch-hikers'). The average for Version 3 games is 69 rooms, 39 takeable objects.

'A Mind Forever Voyaging' contains many rooms (178) but few objects (30). 'Trinity', a more typical style of game, contains 134 rooms and 49 objects: the Version 5 'Curses' has a few more of each. Of the Version 6 games, only 'Zork Zero' scores highly here, with 215 rooms and 106 objects. The average for Version 4/5 games is 105 rooms and 54 objects.

The total number of objects tends to be close to the limit of 255 in Version 3 games. 'Curses' contains 508.

(iv) *Dictionary*

Early games such as 'Zork I' know about 600 words, but again this rises steeply to about 1000 even in Version 3. Later games know 1569 ('Beyond Zork') to the record, 2120 ('Trinity'). (This is achieved by heroic inclusion of unlikely synonyms: e.g. the Japanese lady with the umbrella can be called WOMAN, LADY, CRONE, MADAM, MADAME,

MATRON, DAME or FACE with any of the adjectives OLD, AGED, ANCIENT, JAP, JAPANESE, ORIENTAL or YELLOW.) Version 6 games have smaller dictionaries. So has 'Curses', at 1364.

(v) *Opcodes*

(a) Of the 1426854 opcodes in the shipped Infocom story files in Paul David Doherty's collection, here are the top and bottom ten most popular. (Leaving out those which never occur and so score 0: **nop**, **art_shift**, **piracy** and the two post-Infocom opcodes, **print_unicode** and **check_unicode**.)

| Top Ten Opcodes Chart | | | Bottom Ten Opcodes Chart | | |
|-----------------------|-----------|--------|--------------------------|---------------|----|
| 1. | je | 195959 | 1. | print_form | 2 |
| 2. | print | 142755 | 2. | erase_picture | 3 |
| 3. | jz | 112016 | 3. | read_mouse | 3 |
| 4. | call_vs | 104075 | 4. | encode_text | 7 |
| 5. | print_ret | 80870 | 5. | make_menu | 9 |
| 6. | store | 71128 | 6. | not | 14 |
| 7. | rtrue | 66125 | 7. | scroll_window | 16 |
| 8. | jump | 56534 | 8. | pop_stack | 17 |
| 9. | new_line | 52553 | 9. | restore_undo | 18 |
| 10. | test_attr | 46627 | 10. | mouse_window | 22 |

So about 2/3rds of all opcodes are those in the top ten; 1 in 8 opcodes is a je, and only 1 in 710000 is a print_form.

(b) An experiment (conducted with the help of Kevin Bracey) sheds some light on the opcodes most frequently interpreted in typical play. Two very different games ('Zork I', Version 5 "solid gold" edition; 'Museum of Inform', a complex Inform example game) were played for about 50000 cycles of the Z-machine (about 20 moves in 'Zork I', rather less in the 'Museum'). The following table records all opcodes with a frequency of at least 1% (i.e., 0.01):

| Zork I Solid Gold (Infocom) | | Museum of Inform (Inform) | |
|-----------------------------|-------------|---------------------------|---------------|
| 0.116110 | loadb | 0.104952 | je |
| 0.103990 | storeb | 0.101151 | jz |
| 0.101616 | jz | 0.092727 | jump |
| 0.074979 | dec_chk | 0.080985 | jg |
| 0.066375 | add | 0.079039 | jl |
| 0.066283 | je | 0.070550 | inc |
| 0.060760 | store | 0.070139 | store |
| 0.053867 | loadw | 0.047058 | loadw |
| 0.038095 | storew | 0.034137 | get_prop_addr |
| 0.036428 | mul | 0.024105 | jin |
| 0.032069 | inc_chk | 0.022734 | rtrue |
| 0.030243 | jump | 0.021583 | storew |
| 0.029170 | test_attr | 0.020075 | add |
| 0.020634 | call_vs | 0.018485 | call_vs |
| 0.011184 | get_sibling | 0.016731 | and |
| | | 0.016082 | loadb |
| | | 0.012061 | call_vn |
| | | 0.011879 | test_attr |
| | | 0.011824 | dec |
| | | 0.011687 | ret |

Adventure games spend most of the time parsing, and the differences between these tables reflect different parser designs (byte arrays versus word arrays and arrays stored in properties) as well as different compiler code generators (Inform does not use **inc_chk** or **dec_chk**, so it uses **inc**, **dec**, **jl** and **jl** correspondingly more). In the case of 'Zork I', about a third of all opcodes are branches: in the case of 'Museum', almost half.

Appendix F. Canonical Story Files

Story files are mechanically best identified by their release number and serial code, which are written into the header information at the bottom of Z-machine memory. The release number can be anything between 0 and 65535 but is usually between 1 and 100. The serial code can consist of any six textual characters but is usually the date of compilation, arranged **YYMMDD**: thus 970619 refers to June 19th, 1997. The notation

Release number.Serial code

identifies particular story files: for example the first production copy of 'Enchanter' is 10.830810.

Paul David Doherty's investigations into Infocom's released games have resulted in the following list of all known story files compiled by **Zilch**, the Infocom compiler:

Version 1 story file

| | |
|--------|---------------------|
| Zork I | 5. (no serial code) |
|--------|---------------------|

Version 2 story file

| | |
|---------|-----------|
| Zork I | 15.UG3AU5 |
| Zork II | 7.UG3AU5 |

Version 3 story file

| | |
|---------------------------------------|--|
| Ballyhoo | 97.851218 |
| Cutthroats | 23.840809 |
| Deadline | 18.820311, 19.820427, 21.820512, 26.821108, 27.831005 |
| Enchanter | 10.830810, 15.831107, 16.831118, 24.851118, 29.860820 |
| Four-In-One Sampler I | 26.840731, 53.850407, 55.850823 |
| Four-In-One Sampler II | 97.870601 |
| The Hitch Hiker's Guide To The Galaxy | 47.840914, 56.841221, 58.851002, 59.851108 |
| Hollywood Hijinx | 37.861215 |
| Infidel | 22.830916 |
| Leather Goddesses of Phobos | 118.860325? (beta), 50.860711?, 59.860730, 59.861114 |
| Lurking Horror | 203.870506, 219.870912 (s), 221.870918 (s) |
| Mini-Zork I | 34.871124 |
| Moonmist | 4.860918, 9.861022 |
| Planetfall | 20.830708, 29.840118, 37.851003 |
| Plundered Hearts | 26.870730 |
| Seastalker | 86.840320 (beta), 15.840501, 15.840522, 16.850515, 16.850603 |
| Sorcerer | 67.000000? (beta), 4.840131, 6.840508, 13.851021, 15.851108, 18.860904 |
| Spellbreaker | 63.850916, 87.860904 |
| Starcross | 15.820901, 17.821021 |
| Stationfall | 107.870430 |
| Suspect | 14.841005 |

| | |
|-------------|--|
| Suspended | 5.830222, 7.830419, 8.830521, 8.840521 |
| Wishbringer | 68.850501, 69.850920 |
| Witness | 13.830524, 18.830910, 20.831119, 21.831208, 22.840924 |
| Zork I | 23.820428, 25.820515, 26.820803, 28.821013, 30.830330, 75.830929, 76.840509, 88.840726 |
| Zork II | 17.820427, 18.820512, 18.820517, 19.820721, 22.830331, 23.830411, 48.840904, ?.841220? |
| Zork III | 10.820818, 15.830331, 16.830410, 15.840518, 17.840727 |

Note that the two samplers and the mini-Zork are in the public domain and may be downloaded from Internet archive sites. One form of 'Zork I' can be downloaded freely from Activision's Web pages promoting the Zork brand name.

The two problem children here are 'Seastalker', a submarine game which produces a sonar display across the top of the screen (and thus needs more sophisticated screen control features than the other Version 3 games) and 'The Lurking Horror', which uses sound effects (hence the "(s)" notation).

Version 4 story file

| | |
|---|-----------------------|
| A Mind Forever Voyaging | 77.850814, 79.851122 |
| Bureaucracy | 86.870212, 116.870602 |
| Nord and Bert Couldn't Make Head Nor Tail Of It | 19.870722 |
| Trinity | 11.860509, 12.860926 |

Version 5 story file

| | |
|--|--|
| Beyond Zork | 47.870915, 49.870917, 51.870923, 57.871221 |
| Border Zone | 9.871008 |
| The Hitch Hiker's Guide To The Galaxy SG | 31.871119 |
| Leather Goddesses of Phobos SG | 4.880405 |
| Planetfall SG | 10.880531 |
| Sherlock | 21.871214, 26.880127 (s) |
| Wishbringer SG | 23.880706 |
| Zork I SG | 52.871125 |
| Zork I German | 3.880113 (beta) |

The "SG" games were "solid gold" revisions of existing Version 3 games, adding on-line hints and an UNDO command. Regrettably these are not the versions distributed by Activision on their recent re-releases of the Infocom back catalogue.

One form of 'Sherlock' uses sound effects. 'Border Zone' introduces timed input. 'Beyond Zork' features a character-graphics font. But the most interesting file is the German translation of 'Zork I', which was never commercially released, introducing an alphabet table to the format.

Version 6 story file

| Game | Mac | Amiga | Apple II | IBM |
|-----------|------------|-------------|-------------|------------|
| Arthur | 54.890606 | same as Mac | 63.890622? | 74.890714 |
| Journey | 26.890316 | 30.890322 | 77.890616? | 83.890706 |
| Shogun | 292.890314 | 295.890321 | 311.890510? | 322.890706 |
| Zork Zero | 296.881019 | 366.890323 | 383.890602? | 393.890714 |

The rule that story files should be independent of their target computers was dropped for Version 6 games and this leads to copious footnotes and exceptions in sections 8.8 and 16 of the standard. Story files for a particular game are substantially similar to each other but use fonts, pictures and so on slightly differently.

Note that a new Infocom game, Zork: The Undiscovered Underground was published by Activision in 1997. 16.970828 is the only public version I know of: however, note that this file was compiled by Inform and not by Zilch. It is therefore not useful as a witness to Z-machine rules.