

## 16. The Compiler

### 16.1 The Basic Operations of the Compiler

The purpose of the Lisp compiler is to convert Lisp functions into programs in the Lisp Machine's instruction set, so that they will run more quickly and take up less storage. Compiled functions are represented in Lisp by FEF's (Function Entry Frames), which contain machine code as well as various other information. The printed representation of a FEF is

```
#<DTP-FEF-POINTER address name>
```

If you want to understand the output of the compiler, refer to chapter 28, page 602.

There are three ways to invoke the compiler from the Lisp Machine. First, you may have an interpreted function in the Lisp environment that you would like to compile. The function `compile` is used to do this. Second, you may have code in an editor buffer that you would like to compile. The `Zwei` editor has commands to read code into Lisp and compile it. Third, you may have a program (a group of function definitions and other forms) written in a file on the file system. The function `qc-file` can translate this file into a *QFASL* file that describes the compiled functions and associated data. The *QFASL* file format is capable of representing an arbitrary collection of Lisp objects, including shared structure and cycles of pointers. The name derives from "Q", a prefix once used to mean "for the Lisp machine, not for Maclisp", and "FASL", an abbreviation for "fast loading".

### 16.2 How to Invoke the Compiler

**compile** *function-spec* &optional *definition*

If *definition* is supplied, it should be a **lambda**-expression. Otherwise *function-spec* (this is usually a symbol, but see section 10.2, page 154 for details) should be defined as an interpreted function and its definition will be used as the **lambda**-expression to be compiled. The compiler converts the **lambda**-expression into a FEF, saves the **lambda**-expression as the `:previous-definition` property of *function-spec* if it is a symbol, and changes *function-spec*'s definition to be the FEF. (See `fdefine`, page 169.)

If *function-spec*'s definition is already a FEF, and that FEF's debugging info alist records the interpreted definition it was compiled from, that same definition is compiled again. The original definition is recorded in a FEF's debugging info alist whenever the function is compiled in core (such as by means of `compile`, but not if the function is loaded from a *QFASL* file, except for `defsubst`s).

**uncompile** *function-spec*

If *function-spec* is defined as a compiled function that records the original definition that was compiled, then *function-spec* is redefined with that original definition. This undoes the effect of calling `compile` on *function-spec*.

**compile-lambda** *lambda-exp function-spec*

Returns a compiled function object produced by compiling *lambda-exp*. The function name recorded by the compiled function object is *function-spec*, but that function spec is not defined by **compile-lambda**.

**compile-encapsulations** *function-spec*

Compiles all encapsulations that *function-spec* currently has. Encapsulations (see section 10.10, page 175) include tracing, breakons and advice. Compiling tracing or breakons makes it possible (or at least more possible) to trace or breakon certain functions that are used in the evaluator. Compiling advice makes it less costly to advise functions that are used frequently.

Any encapsulation that is changed will cease to be compiled; thus, if you add or remove advice, you must do **compile-encapsulations** again if you wish the advice to be compiled again.

**compile-encapsulations-flag***Variable*

If this is non-nil, all encapsulations that are created are compiled automatically.

**qc-file** *filename &optional output-file load-flag in-core-flag package file-local-declarations dont-set-default-p read-then-process-flag*

This function takes a formidable number of arguments, but normally only one argument is supplied. The file *filename* is given to the compiler, and the output of the compiler is written to a file whose name is *filename* except with a file type of "QFASL". The input format for files to the compiler is described on section 16.3, page 230. Macro definitions, **subst** definitions, and **special** declarations created during the compilation are undone when the compilation is finished.

The optional arguments allow certain modifications to the standard procedure. *output-file* lets you change where the output is written. *package* lets you specify in what package the source file is to be read. Normally the system knows, or asks interactively, and you need not supply this argument. *load-flag* and *in-core-flag* are incomprehensible; you don't want to use them. *file-local-declarations* is for compiling multiple files as if they were one. *dont-set-default-p* suppresses the changing of the default file name to *filename* that normally occurs.

Normally, a form is read from the file and processed and then another form is read and processed, and so on. But if *read-then-process-flag* is non-nil, the whole source file is read before any of it is processed. This is not done by default; it has the problem that compile-time reader-macros defined in the file will not work properly.

**qc-file-load** *filename &optional output-file load-flag in-core-flag package functions-defined file-local-declarations dont-set-default-p read-then-process-flag*

**qc-file-load** compiles a file and then loads in the resulting QFASL file.

**compiler:compiler-verbose***Variable*

If this variable is non-nil, the compiler prints the name of each function that it is about to compile.

**compiler:peep-enable***Variable*

The peephole optimizer is used if this variable is non-nil. The only reason to set it to nil is if there is a suspicion of a bug in the optimizer.

See also the `disassemble` function (page 641), which lists the instructions of a compiled function in symbolic form.

### 16.3 Input to the Compiler

The purpose of `qc-file` is to take a file and produce a translated version which does the same thing as the original except that the functions are compiled. `qc-file` reads through the input file, processing the forms in it one by one. For each form, suitable binary output is sent to the QFASL file so that when the QFASL file is loaded the effect of that source form will be reproduced. The differences between source files and QFASL files are that QFASL files are in a compressed binary form, which reads much faster but cannot be edited, and that function definitions in QFASL files have been translated from Lisp forms to FEFs.

So, if the source contains a `(defun ...)` form at top level, then when the QFASL file is loaded the function will be defined as a compiled function. If the source file contains a form that is not of a type known specially to the compiler, then that form (encoded in QFASL format) will be output "directly" into the QFASL file, so that when the QFASL file is loaded that form will be evaluated. Thus, if the source file contains `(setq x 3)`, then the compiler will put in the QFASL file instructions to set `x` to `3` at load time (that is, when the QFASL file is loaded into the Lisp environment). It happens that QFASL files have a specific way to `setq` a symbol. For a more general form, the QFASL file would contain instructions to recreate the list structure of a form and then call `eval` on it.

The Lisp machine editor ZWEI assumes that source files are formatted so that an open parenthesis at the left margin (that is, in column zero) indicates the beginning of a function definition or other top level list (with a few standard exceptions). The compiler assumes that you follow this indentation convention, enabling it to tell when a close-parenthesis is missing from one function as soon as the beginning of the next function is reached.

If the compiler finds an open parenthesis in column zero in the middle of a list, it invents enough close parentheses to close off the list that is in progress. A compiler warning is produced instead of an error. After that list has been processed, the open parenthesis is read again. The compilation of list that was forcefully closed off is probably useless, but the compilation of the rest of the file is usually correct. You can read the file into the editor and fix and recompile just the function that was unbalanced.

A similar thing happens on end of file in the middle of a list, so that you get to see any warnings for the function that was unbalanced.

Certain special forms including `eval-when`, `progn`, `local-declare`, `declare-flavor-instance-variables`, and `comment` are customarily used around lists that start in column zero. These symbols have a non-nil `si:may-surround-defun` property that makes the compiler permit this. You can add such properties to other symbols if you want.

**compiler:qc-file-indentation**

*Variable*

The compiler checks for open-parentheses in column zero if this variable is non-nil.

Sometimes we want to put things in the file that are not merely meant to be translated into QFASL form. One such occasion is top level macro definitions; the macros must actually get defined within the compiler in order for the compiler to be able to expand them at compile time. So when a macro form is seen, usually it should be evaluated at compile time as well as put into the QFASL file.

Another thing we sometimes want to put in a file is compiler declarations. These are forms which should be evaluated at compile time to tell the compiler something. They should not be put into the QFASL file, unless they are useful for working incrementally on the functions in the file, compiling them one by one from the editor.

Therefore, a facility exists to allow the user to tell the compiler just what to do with a form. One might want a form to be:

Put into the QFASL file (compiled, of course), or not.

Evaluated within the compiler, or not.

Evaluated if the file is read directly into Lisp, or not.

The `eval-when` special form is used to control this. An `eval-when` form looks like

```
(eval-when times-list
  form1
  form2
  ...)
```

The *times-list* may contain one or more of the symbols `load`, `compile`, or `eval`. If `load` is present, the *forms* are written into the QFASL file to be evaluated when the QFASL file is loaded (except that `defun` forms will put the compiled definition into the QFASL file instead). If `compile` is present, the *forms* are evaluated in the compiler. If `eval` is present, the *forms* are evaluated when read into Lisp; this is because `eval-when` is defined as a special form in Lisp. (The compiler ignores `eval` in the *times-list*.) For example,

```
(eval-when (compile eval) (macro foo (x) (cadr x)))
```

would define `foo` as a macro in the compiler and when the file is read in interpreted, but not when the QFASL file is fasloaded.

**eval-when (*time...*) *body...***

*Special Form*

When seen by the interpreter, if one of the *times* is the symbol `eval` then the *body* forms are evaluated; otherwise `eval-when` does nothing.

But when seen by the compiler, this special form does the special things described above.

For the rest of this section, we will use lists such as are given to `eval-when`, e.g. `(load eval)`, `(load compile)`, etc., to describe when forms are evaluated.

If a form is not enclosed in an `eval-when`, then the times at which it will be evaluated depend on the form. The following table summarizes at what times evaluation will take place for any given form seen at top level by the compiler.

`(eval-when times-list form ...)`

*times-list* specifies when the *form*... should be performed.

`(declare (special ...))` or `(declare (unspecial ...))`

The `special` or `unspecial` is performed at `(load compile)` time.

`(declare anything-else)`

*anything-else* is performed only at `(compile)` time.

`(special ...)` or `(unspecial ...)`

`(load compile eval)`

`(macro ...)` or `(defmacro ...)` or `(defsubst ...)`

or `(defflavor ...)` or `(defstruct ...)`

`(load compile eval)`. However, during file to file compilation, the definition is kept in effect only for the one file. It is not done "for real" until the file is loaded.

`(comment ...)` Ignored at all times.

`(compiler-let ((var val) ...) body...)`

Processes the *body* in its normal fashion, but at `(compile eval)` time, the indicated variable bindings are in effect. These variables will typically affect the operation of the compiler or of macros. See section 17.4.6, page 265.

`(local-declare (decl decl ...) body...)`

Processes the *body* in its normal fashion, with the indicated declarations added to the front of the list which is the value of `local-declarations`.

`(defun ...)` or `(defmethod ...)` or `(defselect ...)`

`(load eval)`, but at load time what is processed is not this form itself, but the result of compiling it.

*anything-else* `(load eval)`

Sometimes a macro wants to return more than one form for the compiler top level to see (and to be evaluated). The following facility is provided for such macros. If a form

`(progn (quote compile) form1 form2 ...)`

is seen at the compiler top level, all of the *forms* are processed as if they had been at compiler top level. (Of course, in the interpreter they will all be evaluated, and the `(quote compile)` will harmlessly evaluate to the symbol `compile` and be ignored.) See section 17.4.3, page 260, for additional discussion of this.

To prevent an expression from being optimized by the compiler, surround it with a call to `dont-optimize`.

**dont-optimize** *form*

*Special Form*

In execution, this is equivalent to simply *form*. However, any source-level optimizations that the compiler would normally perform on the top level of *form* are not done.

Examples:

```
(dont-optimize (apply 'foo (list 'a 'b)))
```

actually makes a list and calls `apply`, rather than doing

```
(foo 'a 'b)
```

```
(dont-optimize (si:flavor-method-table flav))
```

actually calls `si:flavor-method-table` as a function, rather than substituting the definition of that `defsubst`.

`dont-optimize` can even be used around a `defsubst` inside of `setf` or `locf`, to prevent open-coding of the `defsubst`. In this case, a function will be created at load time to do the setting or return the location.

```
(setf (dont-optimize (zwei:buffer-package buffer))
      (pkg-find-package "foo"))
```

Subforms of *form*, such as arguments, are still optimized or open coded, unless additional `dont-optimize`'s appear around them.

## 16.4 Compiler Declarations

Declarations provide auxiliary information on how to execute a function or expression properly, in addition to "what expression to compute". Many declarations are relevant to techniques of compilation and are irrelevant when a function is interpreted. Some do not affect execution at all and only provide information about the function, for the sake of `arglist`, for example.

Declarations may apply to an entire function or to any expression within it. Declarations can be made around any subexpression by writing a `local-declare` around the subexpression. Declarations can be made on an entire function by writing a `declare` at the front of the function's body.

**local-declare** (*declaration...*) *body...*

*Special Form*

A `local-declare` form looks like

```
(local-declare (decl1 decl2 ...)
  form1
  form2
  ...)
```

Each *decl* is consed onto the list local-declarations while the *forms* are being evaluated (in the interpreter) or compiled (in the compiler).

**declare** *declaration...**Special Form*

The special form **declare** is used for writing declarations that apply to an entire function definition.

A *declare* inside a function definition, just after the argument list, is equivalent to putting a *local-declare* around the function definition. More specifically,

```
(defun foo (a b)
  (declare (special a b))
  (bar))
```

is equivalent to

```
(local-declare ((special a b))
  (defun foo (a b)
    (bar)))
```

Note that

```
(defun foo (a b)
  (local-declare ((special a b))
    (bar)))
```

will not do the job, because the declaration is not in effect for the binding of the arguments of **foo**.

**declare** is preferable to *local-declare* in this sort of situation, because it allows the *defuns* themselves to be the top-level lists in the file. While *local-declare* might appear to have an advantage in that one *local-declare* may go around several *defuns*, it tends to cause trouble to use *local-declare* in that fashion.

**declare** has a similar meaning at the front of the body of a *progn*, *prog*, *let*, *prog\**, *let\**, or internal *lambda*. For example,

```
(prog (x)
  (declare (special x))
  ...)
```

is equivalent to

```
(local-declare ((special x))
  (prog (x)
    ...))
```

At top level in the file, (**declare** *declarations...*) is equivalent to (*eval-when* (**compile**) *declarations...*). This use of **declare** is nearly obsolete, and should be avoided.

Elsewhere, and in the interpreter, **declare**'s are ignored.

Here is a list of declarations that have system-defined meanings:

(**special** *var1 var2 ...*)

The variables *var1*, *var2*, etc. will be treated as special variables during the compilation of the *forms*.

(**unspecial** *var1 var2 ...*)

The variables *var1*, *var2*, etc. will be treated as local variables during the

compilation of the *forms*.

**(:def name . definition)**

*name* will be defined for the compiler during the compilation of the *forms*. The compiler uses this to keep track of macros and open-codable functions (defsubst) defined in the file being compiled. Note that the caddr of this item is a function.

**(propname symbol value)**

Within *forms*, (getdecl *symbol propname*) will return *value*. This is how the compiler keeps track of defdecls.

These declarations are significant only when they apply to an entire defun

**(:arglist . arglist)**

Records *arglist* as the argument list of the function, to be used instead of its lambda-list if anyone asks what its arguments are. This is purely documentation.

**(:values . values) or (:return-list . values)**

Records *values* as the return values list of the function, to be used if anyone asks what values it returns. This is purely documentation.

**(:si:function-parent parent-function-spec)**

Records *parent-function-spec* as the parent of this function. If, in the editor, you ask to see the source of this function, and the editor doesn't know where it is, the editor will show you the source code for the parent function instead.

**(:self-flavor flavorname)**

Instance variables of the flavor *flavorname*, in *self*, will be accessible in the function.

### **local-declarations**

*Variable*

The value of this variable is a list of all declarations that are temporarily in effect. During compilation, it pertains to the code being compiled. During interpretation, it pertains to the code being interpreted. (It is not used during execution of compiled code, since all processing of the declarations was done at compile time.)

As a result, at any time a macro is expanded, the value of **local-declarations** pertains to the code being expanded.

### **sys:file-local-declarations**

*Variable*

During file-to-file compilation, the value of this variable is a list of all declarations that are in effect for the rest of the file. Macro definitions, defdecls, and special declarations that come from defvars are all recorded on this list.

### **special variable...**

*Special Form*

Declares each *variable* to be "special" for the compiler. Usually it is better to use defvar or defconst. special is used to make it possible to compile one file that refers to a variable without first having to load another file that defines the variable.



**unspecial** *variable...**Special Form*

Removes any "special" declarations of the *variables* for the compiler.

When symbol properties are referred to during macro expansion, it is desirable for properties defined in a file to be "in effect" for the compilation of the rest of the file. This will not happen if `get` and `defprop` are used, because the `defprop` will not be executed until the file is loaded. Instead, you can use `getdecl` and `defdecl`. These are normally the same as `get` and `defprop`, but during file-to-file compilation they also refer to and create declarations.

**getdecl** *symbol property*

This is a version of `get` that allows the properties of the *symbol* to be overridden by declarations.

If `local-declarations` or `sys:file-local-declarations` contains a declaration of the form (*property symbol value*), `getdecl` returns *value*. Otherwise, `getdecl` returns the result of (`get symbol property`).

`getdecl` is typically used in macro definitions. For example, the `setf` macro uses `getdecl` to get the `setf` property of the function in the expression for the field to be set.

**putdecl** *symbol property value*

Causes (`getdecl symbol property`) to return *value*.

`putdecl` usually simply does a `putprop`. But if executed at compile time during file-to-file compilation, it instead makes an entry on `file-local-declarations` of the form (*property symbol value*).

In either case, this stores *value* where `getdecl` will find it; but if `putdecl` is done during compilation, it affects only the rest of that compilation.

**defdecl** *symbol property value**Special Form*

When executed, this is like `putdecl` except that the arguments are not evaluated. It is usually the same as `defprop` except for the order of the arguments.

Unlike `defprop`, when `defdecl` is encountered during file-to-file compilation, it is executed, creating a declaration which remains in effect for the rest of the compilation. (The `defdecl` form also goes into the QFASL file to be executed when the file is loaded). `defprop` would have no effect whatever at compile time.

`defdecl` is often useful as a part of the expansion of a macro. It is also useful as a top-level expression in a source file.

```
(defdecl foo setf ((foo x) . (set-foo x si:value)))
```

in a source file would allow (`setf (foo arg) value`) to be used in functions in that source file; and, once the file was loaded, by anyone.

The next three functions are primarily for Maclisp compatibility. In Maclisp, they are declarations, used within a `declare` at top level in the file.

**\*expr** *symbol...**Special Form*

Declares each *symbol* to be the name of a function. In addition it prevents these functions from appearing in the list of functions referenced but not defined, printed at the end of the compilation.

**\*lexpr** *symbol...**Special Form*

Declares each *symbol* to be the name of a function. In addition it prevents these functions from appearing in the list of functions referenced but not defined, printed at the end of the compilation.

**\*fexpr** *symbol...**Special Form*

Declares each *symbol* to be the name of a special form. In addition it prevents these names from appearing in the list of functions referenced but not defined, printed at the end of the compilation.

There are some advertised variables whose compile-time values affect the operation of the compiler. The user may set these variables by including in his file forms such as

```
(declare (setq open-code-map-switch t))
```

However, these variables seem not to be needed very often.

**run-in-maclisp-switch***Variable*

If this variable is non-nil, the compiler will try to warn the user about any constructs that will not work in Maclisp. By no means will all Lisp Machine system functions not built in to Maclisp be cause for warnings; only those that could not be written by the user in Maclisp (for example, `make-array`, `value-cell-location`, etc.). Also, lambda-list keywords such as `&optional` and initialized `prog` variables will be mentioned. This switch also inhibits the warnings for obsolete Maclisp functions. The default value of this variable is `nil`.

**obsolete-function-warning-switch***Variable*

If this variable is non-nil, the compiler will try to warn the user whenever an "obsolete" Maclisp-compatibility function such as `maknam` or `samepnamep` is used. The default value is `t`.

**allow-variables-in-function-position-switch***Variable*

If this variable is non-nil, the compiler allows the use of the name of a variable in function position to mean that the variable's value should be `funcall'd`. This is for compatibility with old Maclisp programs. The default value of this variable is `nil`.

**open-code-map-switch***Variable*

If this variable is non-nil, the compiler will attempt to produce inline code for the mapping functions (`mapc`, `mapcar`, etc., but not `mapatoms`) if the function being mapped is an anonymous lambda-expression. This allows that function to reference the local variables of the enclosing function without the need for special declarations. The generated code is also more efficient. The default value is `t`.

**all-special-switch***Variable*

If this variable is non-nil, the compiler regards all variables as special, regardless of how they were declared. This provides compatibility with the interpreter at the cost of efficiency. The default is nil.

**inhibit-style-warnings-switch***Variable*

If this variable is non-nil, all compiler style-checking is turned off. Style checking is used to issue obsolete function warnings, won't-run-in-Maclisp warnings, and other sorts of warnings. The default value is nil. See also the `inhibit-style-warnings` macro, which acts on one level only of an expression.

**compiler-let** ((*variable value*)...) *body*...*Macro*

Syntactically identical to `let`, `compiler-let` allows compiler switches to be bound locally at compile time, during the processing of the *body* forms.

Example:

```
(compiler-let ((open-code-map-switch nil))
  (map (function (lambda (x) ...)) foo))
```

will prevent the compiler from open-coding the `map`. When interpreted, `compiler-let` is equivalent to `let`. This is so that global switches which affect the behavior of macro expanders can be bound locally.

## 16.5. Using Compiler Warnings

When the compiler prints warnings, it also records them in a data base, organized by file and by function within file. Old warnings for previous compilations of the same function are thrown away, so the data base contains only warnings that are still applicable. This data base can be used to visit, in the editor, the functions that got warnings. You can also save the data base and restore it later.

There are three editor commands that you can use to begin visiting the sites of the recorded warnings. They differ only in how they decide which files to look through:

**Meta-X Edit Warnings**

For each file that has any warnings, asks whether to edit the warnings for that file.

**Meta-X Edit File Warnings**

Reads the name of a file and then edits the warnings for that file.

**Meta-X Edit System Warnings**

Reads the name of a system and then edits the warnings for all files in that system (see `defsystem`, page 520).

While the warnings are being edited, the warnings themselves appear in a small window at the top of the editor frame, and the code appears in a large window which occupies the rest of the editor frame.

As soon as you have finished specifying the file(s) or system to process, the editor will proceed to visit the code for the first warning. From then on, to move to the next warning, use the command `Control-Shift-W`. To move to the previous warning, use `Meta-Shift-W`. You can also switch to the warnings window with `Control-XO` or with the mouse, and move around in

that buffer. When you use Control-Shift-W and there are no more warnings after the cursor, you return to single-window mode.

You can also insert the text of the warnings into any editor buffer:

#### Meta-X Insert File Warnings

Reads the name of a file and inserts into the buffer after point the text for that file's warnings. The mark is left after the warnings, but the region is not turned on.

#### Meta-X Insert Warnings

Inserts into the buffer after point the text for the warnings of all files that have warnings. The mark is left after the warnings, but the region is not turned on.

You can also dump the warnings data base into a file and reload it later. Then you can do Meta-X Edit Warnings again in the later session. You dump the warnings with `si:dump-warnings` and load the file again with `load`. In addition, `make-system` with the `:batch` option writes all the warnings into a file in this way.

**si:dump-warnings** *output-file-pathname* &rest *warnings-file-pathnames*

Writes the warnings for the files named in *warnings-file-pathnames* (a list of pathnames or strings) into a file named *output-file-pathname*.

**compiler:warn-on-errors**

*Variable*

If this variable is non-nil, errors in reading code to be compiled, and errors in macro expansion within the compiler, produce only warnings; they do not enter the debugger. The variable is normally `t`.

The default setting is useful when you do not anticipate errors during compilation, because it allows the compilation to proceed past such errors. If you have walked away from the machine, you do not come back to find that your compilation stopped in the first file and did not finish.

If you find an inexplicable error in reading or macroexpansion, and wish to use the debugger to localize it, set `compiler:warn-on-errors` to nil and recompile.

## 16.5.1 Controlling Compiler Warnings

By controlling the compile-time values of the variables `run-in-maclisp-switch`, `obsolete-function-warning-switch`, and `inhibit-style-warning-switch` (explained above), you can enable or disable some of the warning messages of the compiler. The following special form is also useful:

**inhibit-style-warnings** *form*

*Macro*

Prevents the compiler from performing style-checking on the top level of *form*. Style-checking will still be done on the arguments of *form*. Both obsolete function warnings and won't-run-in-Maclisp warnings are done by means of the style-checking mechanism, so, for example,

```
(setq bar (inhibit-style-warnings (value-cell-location foo)))
```

will not warn that `value-cell-location` will not work in Maclisp, but

(`inhibit-style-warnings (setq bar (value-cell-location foo))`) will warn, since `inhibit-style-warnings` applies only to the top level of the form inside it (in this case, to the `setq`).

Sometimes functions take argument that they deliberately do not use. Normally the compiler warns you if your program binds a variable that it never references. In order to disable this warning for variables that you know you are not going to use, there are two things you can do. The first thing is to name the variables `ignore` or `ignored`. The compiler will not complain if a variable by one of these names is not used. Furthermore, by special dispensation, it is all right to have more than one variable in a lambda-list that has one of these names. The other thing you can do is simply use the variable for effect (ignoring its value) at the front of the function. Example:

```
(defun the-function (list fraz-name fraz-size)
  fraz-size      ; This argument is not used.
  ...)
```

This has the advantage that `arglist` (see page 172) will return a more meaningful argument list for the function, rather than returning something with `ignores` in it.

The following function is useful for requesting compiler warnings in certain esoteric cases. Normally, the compiler notices whenever any function  $x$  uses (calls) any other function  $y$ ; it makes notes of all these uses, and then warns you at the end of the compilation if the function  $y$  got called but was neither defined nor declared (by `*expr`, see page 237). This usually does what you want, but sometimes there is no way the compiler can tell that a certain function is being used. Suppose that instead of  $x$ 's containing any forms that call  $y$ ,  $x$  simply stores  $y$  away in a data structure somewhere, and someplace else in the program that data structure is accessed and `funcall` is done on it. There is no way that the compiler can see that this is going to happen, and so it can't notice the function usage, and so it can't create a warning message. In order to make such warnings happen, you can explicitly call the following function at compile-time.

**compiler:function-referenced** *what by*

*what* is a symbol that is being used as a function. *by* may be any function spec. `compiler:function-referenced` must be called at compile-time while a compilation is in progress. It tells the compiler that the function *what* is referenced by *by*. When the compilation is finished, if the function *what* has not been defined, the compiler will issue a warning to the effect that *by* referred to the function *what*, which was never defined.

You can also tell the compiler about any function it should consider "defined":

**compiler:compilation-define** *function-spec*

*function-spec* is marked as "defined" for the sake of the compiler; calls to this function will not produce warnings.

**compiler:make-obsolete** *function reason*

*Special Form*

This special form declares a function to be obsolete; code that calls it will get a compiler warning, under the control of `obsolete-function-warning-switch`. This is used by the compiler to mark as obsolete some Maclisp functions which exist in Zetalisp but should not be used in new programs. It can also be useful when maintaining a large system, as a reminder that a function has become obsolete and usage of it should be phased out. An example of an obsolete-function declaration is:

```
(compiler:make-obsolete create-mumblefrotz
 "use MUMBLIFY with the :FROTZ option instead")
```

## 16.5.2 Recording Warnings

The warnings data base is not just for compilation. It can record operations for any number of different operations on files or parts of files. Compilation is merely the only operation in the system that uses it.

Each operation about which warnings can be recorded should have a name, preferably in the keyword package. This symbol should have four properties that tell the system how to print out the operation name as various parts of speech. For compilation, the operation name is `:compile` and the properties are defined as follows:

```
(defprop :compile "compilation" name-as-action)
(defprop :compile "compiling" name-as-present-participle)
(defprop :compile "compiled" name-as-past-participle)
(defprop :compile "compiler" name-as-agent)
```

The warnings system considers that these operations are normally performed on files that are composed of named objects. Each warning is associated with a filename and then with an object within the file. It is also possible to record warnings about objects that are not within any file.

To tell the warnings system that you are starting to process all or part of a file, use the macro `si:file-operation-with-warnings`.

**sys:file-operation-with-warnings** (*generic-pathname* *Special Form*  
*operation-name whole-file-p*) *body...*

*body* is executed within a context set up so that warnings can be recorded for operation *operation-name* about the file specified by *generic-pathname* (see page 469).

In the case of compilation, this is done at the level of `qc-file` (actually, it is done in `compiler:compile-stream`).

*whole-file-p* should be non-nil if the entire contents of the file will be processed inside the *body* if it finishes; this implies that any warnings left over from previous iterations of this operation on this file should be thrown away on exit. This is only relevant to objects that are not found in the file this time; the assumption is that the objects must have been deleted from the file and their warnings are no longer appropriate.

All three of the special arguments are specified as expressions that are evaluated.

Within the processing of a file, you must also announce when you are beginning to process an object:

**sys:object-operation-with-warnings** (*object-name* *location-function*) *body...* *Special Form*

*body* is execute in a context set up so that warnings are recorded for the object named *object-name*, which can be a symbol or a list. Object names are compared with **equal**.

In the case of compilation, this macro goes around the processing of a single function.

*location-function* is either nil or a function that the editor uses to find the text of the object. Refer to the file SYS: ZWEI; POSS LISP for more details on this.

*object-name* and *location-function* are specified with expressions that are evaluated.

You can enter this macro recursively. If the inner invocation is for the same object as the outer one, it has no effect. Otherwise, warnings recorded in the inner invocation apply to the object specified therein.

Finally, when you detect exceptions, you must make the actual warnings:

**sys:record-warning** *type severity location-info format-string &rest args*

Records one warning for the object and file currently being processed. The text of the warning is specified by *format-string* and *args*, which are suitable arguments for **format**, but the warning is *not* printed when you call this function. Those arguments will be used to reprint the warning later.

**sys:record-and-print-warning** *type severity location-info format-string &rest args*

Records a warning and also prints it.

*type* is a symbol that identifies the specific cause of the warning. Types have meaning only as defined by a particular operation, and at present nothing makes much use of them. The system defines one type: **si:premature-warnings-marker**.

*severity* measures how important a warning this is, and the general causal classification. It should be a symbol in the keyword package. Several severities are defined, and should be used when appropriate, but nothing looks at them:

**:implausible** This warning is about something that is not intrinsically wrong but is probably due to a mistake of some sort.

**:impossible** This warning is about something that cannot have a meaning even if circumstances outside the text being processed are changed.

**:probable-error**

This is used to indicate something that is certainly an error but can be made correct by a change somewhere else; for example, calling a function with the wrong number of arguments.

**:missing-declaration**

This is used for warnings about free variables not declared special, and such. It means that the text was not actually incorrect, but something else that is supposed to accompany it was missing.

**:obsolete**

This warning is about something that you shouldn't use any more, but which still does work.

<code>:very-obsolete</code>	This is about something that doesn't even work any more.
<code>:maclisp</code>	This is for something that doesn't work in Maclisp.
<code>:fatal</code>	This indicates a problem so severe that no sense can be made of the object at all. It indicates that the presence or absence of other warnings is not significant.
<code>:error</code>	There was a Lisp error in processing the object.

*location-info* is intended to be used to inform the editor of the precise location in the text of the cause of this warning. It is not defined as yet, and you should use `nil`.

If a warning is encountered while processing data that doesn't really have a name (such as forms in a source file that are not function definitions), you can record a warning even though you are not inside an invocation of `sys:object-operation-with-warnings`. This warning is known as a *premature warning* and it will be recorded with the next object that is processed; a message will be added so that the user can tell which warnings were premature.

Refer to the file `SYS: SYS; QNEW LISP` for more information on the warnings data base.

## 16.6 Compiler Source-Level Optimizers

The compiler stores optimizers for source code on property lists so as to make it easy for the user to add them. An optimizer can be used to transform code into an equivalent but more efficient form (for example, `(eq obj nil)` is transformed into `(null obj)`, which can be compiled better). An optimizer can also be used to tell the compiler how to compile a special form. For example, in the interpreter `do` is a special form, implemented by a function which takes quoted arguments and calls `eval`. In the compiler, `do` is expanded in a macro-like way by an optimizer into equivalent Lisp code using `prog`, `cond`, and `go`, which the compiler understands.

The compiler finds the optimizers to apply to a form by looking for the `compiler:optimizers` property of the symbol that is the car of the form. The value of this property should be a list of optimizers, each of which must be a function of one argument. The compiler tries each optimizer in turn, passing the form to be optimized as the argument. An optimizer that returns the original form unchanged (`eq` to the argument) has "done nothing", and the next optimizer is tried. If the optimizer returns anything else, it has "done something", and the whole process starts over again. Only after all the optimizers have been tried and have done nothing is an ordinary macro definition processed. This is so that the macro definitions, which will be seen by the interpreter, can be overridden for the compiler by optimizers.

Optimizers should not be used to define new language features, because they only take effect in the compiler; the interpreter (that is, the evaluator) doesn't know about optimizers. So an optimizer should not change the effect of a form; it should produce another form that does the same thing, possibly faster or with less memory or something. That is why they are called optimizers. If you want to actually change the form to do something else, you should be using macros.



**compiler:add-optimizer** *function optimizer optimized-into...* *Special Form*

Puts *optimizer* on *function*'s optimizers list if it isn't there already. *optimizer* is the name of an optimization function, and *function* is the name of the function calls which are to be processed. Neither is evaluated.

(**compiler:add-optimizer** *function optimizer optimize-into-1 optimize-into-2...*) also remembers *optimize-into-1*, etc., as names of functions which may be called in place of *function* as a result of the optimization. Then **who-calls** of *function* will also mention calls of *optimize-into-1*, etc.

## 16.7 Files that Maclisp Must Compile

Certain programs are intended to be run both in Maclisp and in Zetalisp. Their source files need some special conventions. For example, all **special** declarations must be enclosed in **declares**, so that the Maclisp compiler will see them. The main issue is that many functions and special forms of Zetalisp do not exist in Maclisp. It is suggested that you turn on **run-in-maclisp-switch** in such files, which will warn you about a lot of problems that your program may have if you try to run it in Maclisp.

The macro-character combination **#Q** causes the object that follows it to be visible only when compiling for Zetalisp. The combination **#M** causes the following object to be visible only when compiling for Maclisp. These work both on subexpressions of the objects in the file and at top level in the file. To conditionalize top-level objects, however, it is better to put the macros **if-for-lispm** and **if-for-maclisp** around them. The **if-for-lispm** macro turns off **run-in-maclisp-switch** within its object, preventing spurious warnings from the compiler. The **#Q** macro-character cannot do this, since it can be used to conditionalize any S-expression, not just a top-level form.

To allow a file to detect what environment it is being compiled in, the following macros are provided:

**if-for-lispm** *form* *Macro*

If (**if-for-lispm** *form*) is seen at the top level of the compiler, *form* is passed to the compiler top level if the output of the compiler is a QFASL file intended for Zetalisp. If the Zetalisp interpreter sees this it will evaluate *form* (the macro expands into *form*).

**if-for-maclisp** *form* *Macro*

If (**if-for-maclisp** *form*) is seen at the top level of the compiler, *form* is passed to the compiler top level if the output of the compiler is a FASL file intended for Maclisp (e.g. if the compiler is COMPLR). If the Zetalisp interpreter sees this it will ignore it (the macro expands into nil).

**if-for-maclisp-else-lispm** *maclisp-form lispm-form* *Macro*

If (**if-for-maclisp-else-lispm** *form1 form2*) is seen at the top level of the compiler, *form1* is passed to the compiler top level if the output of the compiler is a FASL file intended for Maclisp; otherwise *form2* is passed to the compiler top level.

**if-in-lispm** *form**Macro*

In Zetalisp, (*if-in-lispm form*) causes *form* to be evaluated; in Maclisp, *form* is ignored.

**if-in-maclisp** *form**Macro*

In Maclisp, (*if-in-maclisp form*) causes *form* to be evaluated; in Zetalisp, *form* is ignored.

In order to make sure that those macros are defined when reading the file into the Maclisp compiler, you must make the file start with a prelude, which should look like:

```
(declare (cond ((not (status feature lispm))
                (load '|AI: LISPM2; CONDIR|))))
;; Or other suitable filename
```

This will do nothing when you compile the program on the Lisp Machine. If you compile it with the Maclisp compiler, it will load in definitions of the above macros, so that they will be available to your program. The form (*status feature lispm*) is generally useful in other ways; it evaluates to *t* when evaluated on the Lisp machine and to *nil* when evaluated in Maclisp.

## 16.8 Putting Data in QFASL Files

It is possible to make a QFASL file containing data, rather than a compiled program. This can be useful to speed up loading of a data structure into the machine, as compared with reading in printed representations. Also, certain data structures such as arrays do not have a convenient printed representation as text, but can be saved in QFASL files. For example, the system stores fonts this way. Each font is in a QFASL file (on the *SYS: FONTS;* directory) that contains the data structures for that font. When the file is loaded, the symbol that is the name of the font gets set to the array that represents the font. Putting data into a QFASL file is often referred to as "*fasdumping* the data".

In compiled programs, the constants are saved in the QFASL file in this way. The compiler optimizes by making constants that are *equal* become *eq* when the file is loaded. This does not happen when you make a data file yourself; identity of objects is preserved. Note that when a QFASL file is loaded, objects that were *eq* when the file was written are still *eq*; this does not normally happen with text files.

The following types of objects can be represented in QFASL files: Symbols (but uninterned symbols will be interned when the file is loaded), numbers of all kinds, lists, strings, arrays of all kinds, instances, and FEFs.

When an instance is *fasdumped* (put into a QFASL file), it is sent a *:fasd-form* message, which must return a Lisp form that, when evaluated, will recreate the equivalent of that instance. This is because instances are often part of a large data structure, and simply *fasdumping* all of the instance variables and making a new instance with those same values is unlikely to work. Instances remain *eq*; the *:fasd-form* message is only sent the first time a particular instance is encountered during writing of a QFASL file. If the instance does not accept the *:fasd-form* message, it cannot be *fasdumped*.

**dump-forms-to-file** *filename forms-list &optional attribute-list*

Writes a QFASL file named *filename* which contains, in effect, the forms in *forms-list*. That is to say, when the file is loaded, its effect will be the same as evaluating those forms.

Example:

```
(dump-forms-to-file "foo" '((setq x 1) (setq y 2)))
(load "foo")
x => 1
y => 2
```

*attribute-list* is the file attribute list to store in the QFASL file. It is a list of alternating keywords and values, and corresponds to the *-\*-* line of a source file. The most useful keyword in this context is `:package`, whose value in the attribute list specifies the package to be used both in dumping the forms and in loading the file. If no `:package` keyword is present, the file will be loaded in whatever package is current at the time.

**compiler:fasd-symbol-value** *filename symbol*

Writes a QFASL file named *filename* which contains the value of *symbol*. When the file is loaded, *symbol* will be `setq`'ed to the same value. *filename* is parsed with the same defaults that `load` and `qc-file` use. The file type defaults to "QFASL".

**compiler:fasd-font** *name*

Writes the font named *name* into a QFASL file with the appropriate name (on the `SYS: FONTS;` directory).

**compiler:fasd-file-symbols-properties** *filename symbols properties dump-values-p dump-functions-p new-symbol-function*

This is a way to dump a complex data structure into a QFASL file. The values, the function definitions, and some of the properties of certain symbols are put into the QFASL file in such a way that when the file is loaded the symbols will be `setq`ed, `fdefined`, and `putpropped` appropriately. The user can control what happens to symbols discovered in the data structures being `fasdump`ed.

*filename* is the name of the file to be written. It is parsed with the same defaults that `load` and `qc-file` use. The file type defaults to "QFASL".

*symbols* is a list of symbols to be processed. *properties* is a list of properties which are to be `fasdump`ed if they are found on the symbols. *dump-values-p* and *dump-functions-p* control whether the values and function definitions are also dumped.

*new-symbol-function* is called whenever a new symbol is found in the structure being dumped. It can do nothing, or it can add the symbol to the list to be processed by calling `compiler:fasd-symbol-push`. The value returned by *new-symbol-function* is ignored.

## 16.9 Analyzing QFASL Files

QFASL files are composed of 16-bit nibbles. The first two nibbles in the file contain fixed values, which are there so the system can tell a proper QFASL file. The next nibble is the beginning of the first *group*. A group starts with a nibble that specifies an operation. It may be followed by other nibbles that are arguments.

Most of the groups in a QFASL file are there to construct objects when the file is loaded. These objects are recorded in the *fasl-table*. Each time an object is constructed, it is assigned the next sequential index in the *fasl-table*. The indices are used by other groups later in the file, to refer back to objects already constructed.

To prevent the *fasl-table* from becoming too large, the QFASL file can be divided into *whacks*. The *fasl-table* is cleared out at the beginning of each whack.

The other groups in the QFASL file will perform operations such as evaluating a list previously constructed or storing an object into a symbol's function cell or value cell.

If you are having trouble with a QFASL file and want to find out exactly what it does when it is loaded, you can use UNFASL to find out.

**si:unfasl-print** *input-file-name*

Prints on standard-output a description of the contents of the QFASL file *input-file-name*.

**si:unfasl-file** *input-file-name* &optional *output-file-name*

Writes a description of the contents of the QFASL file *input-file-name* into the output file. The output file type defaults to UNFASL and the rest of the pathname defaults from *input-file-name*.

## 16.10 Compiler Interlocking

Because the compiler uses a temporary area for its internal data (so as to avoid the need for frequent garbage collection), it is not reentrant. Only one process can use the compiler at any given time.

Calling the compiler when it is in use in another process will wait until the other compilation is finished, and also produce a notification: "Compiler waiting for resources in process LOSSAGE-5". If the other compilation is proceeding, eventually it will finish and this one will begin. If the other compilation is hung for some reason, such as because it is waiting to type out on an unexposed window, or is having trouble with file servers, you should go to that process and either cause the compilation to finish or abort it.

**compiler:locking-resources** *body...*

*Special Form*

Executes the body, having locked the compiler lock.