

- (d) Not every interval encoding necessarily encodes an interval object, but when it does, that object is unique. Each interval object has at least one encoding and might have more than one.

[Note. Items (c) and (d) are standard and necessary properties of representations. By contrast, the properties (a) and (b) of the maps from Level 1 to Level 2, and back, are fundamental design decisions of the standard.]

6. Expressions and the functions they define

6.1. Definitions. An expression is some symbolic form used to define a function—in general not a static object within program code, but derived dynamically from a particular program execution. Expressions are central to interval computation, because the Fundamental Theorem of Interval Arithmetic (FTIA) is about interpreting an expression in different ways:

- as defining a mathematical real point function f ;
- as defining various (depending on the finite precision interval types used) interval functions that give proven enclosures for the range of f over an input box \mathbf{x} ;
- as defining corresponding decorated interval functions that can give the stronger conclusion that f is everywhere defined, or everywhere continuous, on \mathbf{x} —enabling, for example, an automatic check of the hypotheses of the Brouwer Fixed Point Theorem.

The standard specifies behavior, at the individual operation level, that enables such conclusions, whether or not the notion “expression” exists in a host programming language.

A *formal expression* defines a relation between certain mathematical variables—the *inputs*—and others—the *outputs*—via the application of named *operations*. It is by definition an acyclic (having an acyclic graph, see below) set of dependences between mathematical variables, defined by equations

$$v = \varphi(u_1, \dots, u_k), \quad (1)$$

where v and the u_i come from a nonempty finite set \mathcal{X} of *variable-symbols*; φ comes from a finite set \mathcal{F} of *formal library operations*; $k \geq 0$ is the *arity* of φ ; and distinct equations have distinct v ’s—the *single assignment* property.

Three essentially equivalent descriptions of an expression are as follows.

- (a) Drawing an edge from each u_i to v for each dependence-equation (1) defines the *computational graph* \mathcal{G} —Figure 6.1(a)—a directed graph over the node set \mathcal{X} . The dependences define an expression if and only if \mathcal{G} is *acyclic*. There is then a nonempty set of *output* nodes having no outgoing edge, and a possibly empty set of *input* nodes having no incoming edge.

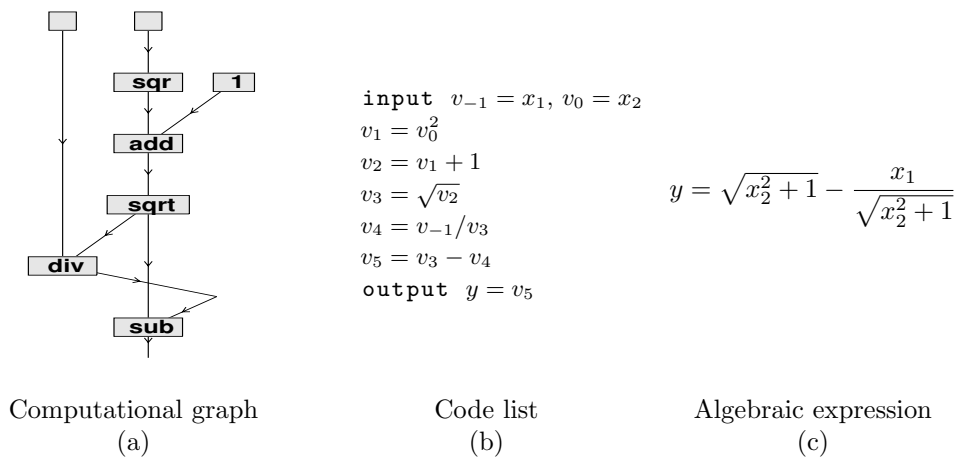


FIGURE 6.1. Essentially equivalent notations for an expression. In (a), the structure is shown by labeling nodes with operations only; the order of arguments is shown by reading incoming edges left to right, e.g., the inputs to **sub** are the results of the preceding **sqrt** and **div**, in that order. Similarly the input nodes are x_1 and x_2 left to right. Note form (c) has redundancy; (a) and (b) do not.

To apply the FTIA it suffices to consider expressions that are *scalar*, with a single output. (All the individual library arithmetic operations of the standard are scalar.)

- (b) Since \mathcal{G} is acyclic the equations can be ordered so that each one only depends on inputs or previously computed values, thus representing the expression as a *code list*—Figure 6.1(b). In the notation of A. Griewank [2], the inputs are written v_{1-n}, \dots, v_0 where $n \geq 0$, conventionally given the aliases x_1, \dots, x_n , so x_i is the same as v_{i-n} . The operations are

$$v_r = \varphi_r(u_{r,1}, \dots, u_{r,k_r}), \quad (r = 1, \dots, m),$$

where $\varphi_r \in \mathcal{F}$ with arity k_r , and each $u_{r,i}$ is a known v_j , that is $j = j(r, i) < r$. (Constants, which are operations of arity 0, may be referred to directly instead of assigned to a v_j .) Assuming a single output, it is v_m , given the alias y , so that the expression defines a formal function $y = f(x_1, \dots, x_n)$.

Either m or n but not both can be zero. The case $n = 0$ and $m \geq 1$ gives a *constant expression*. If $m = 0$ and $n = 1$, there are no operations and y is the same as x_1 , defining the *identity function* $y = f(x_1) = x_1$. In general for $m = 0$ and $n \geq 1$ there are n possibilities, the *coordinate projections* $\pi_j(x_1, \dots, x_n) = x_j$ ($j = 1, \dots, n$).

- (c) By allowing redundancy, an expression always can be converted to a normal *algebraic expression*—Figure 6.1(c)—over the variable-set \mathcal{X} and library \mathcal{F} , defined recursively as follows:
- if $x \in \mathcal{X}$ is a variable symbol, then x is an expression;
 - if $f \in \mathcal{F}$ is a function symbol of arity k and if e_i is an expression for $i = 1, \dots, k$ then the *function symbol application* $f(e_1, \dots, e_k)$ is an expression.

The redundancy is because this form has no way of referring to intermediate values by name, so that if such a value is used several times, the subexpression that computes it must be repeated each time it occurs, see Figure 6.1. If the algebraic expression is evaluated naively, such a subexpression is evaluated more than once, which affects efficiency but not the numerics of what is computed.

Because of its simplicity, this is the definition of expression used in the FTIA proof in Clause D.4.

6.2. Mapping to a library. In IA applications, \mathcal{F} and the arity of each of its functions are defined from the interface of an IA library. Formal expressions typically map to expressions in languages using IA libraries, or to sections of run-time data-flow in executions of programs using IA libraries. The statement of the FTIA transfers to applications of IA libraries via appropriate mapping of a formal expression to a language expression or a section of data-flow. In the absence of such a mapping, the conclusions of the FTIA cannot be drawn from execution of programs written in a host language.

Namely, each formal operation in \mathcal{F} must have a primary *point version* with real-number input(s) and output; one or more *interval versions* with interval input and output; and for each of these a corresponding *decorated interval version*. These produce *point evaluation*, *interval evaluation* and *decorated interval evaluation* of the expression, also termed evaluation in *point mode*, *interval mode* or *decorated interval mode*. Floating-point versions are not relevant to this standard.

The point version is a theoretical (Level 1) function, of which each interval version—there is at least one for each interval type provided by the implementation—is a finite-precision (Level 2) *interval extension*, and each decorated interval version is a *decorated interval extension*.

In this standard an implementation's library by definition comprises all its Level 2 versions of operations that it provides for any of its supported interval types. For the set-based flavor these are specified in §10.6, §10.7 and in Clause 12. Different interval evaluations of f come from using library operations of different Level 2 types, as the implementation may provide.

The set operations **intersection** and **convexHull** are not point-operations and cannot appear directly in an arithmetic expression. However they are useful for efficiently *implementing* interval extensions of functions defined piecewise, see Example (ii) in §11.8.

6.3. The FTIA. Each library point-operation has a defined domain, the set of inputs where it can be evaluated. This leads to the idea of *natural domain* $\text{Dom}(f)$ of the point function $f(x) = f(x_1, \dots, x_n)$ defined by an expression: the set of points x where f is *defined* in the sense that the whole expression can be successfully evaluated.

[Example. From the domains of $/$ and $\sqrt{\cdot}$, one finds the natural domain of $\sqrt{1+1/x}$ is the union of the two intervals $-\infty < x \leq -1$ and $0 < x < +\infty$.]

In the set-based flavor, Moore’s basic theorem for a scalar function is as follows, with the above notation.

Theorem 6.1 (Fundamental Theorem of Interval Arithmetic). *Let $\mathbf{y} = f(\mathbf{x})$ be the result of interval-evaluation of f over a box $\mathbf{x} = (x_1, \dots, x_n)$ using any interval versions of its component library functions. Then*

- (i) (“Basic” form of FTIA.) *In all cases, \mathbf{y} contains the range of f over \mathbf{x} , that is, the set of $f(x)$ at points of \mathbf{x} where it is defined:*

$$\mathbf{y} \supseteq \text{Rge}(f | \mathbf{x}) = \{ f(x) \mid x \in \mathbf{x} \cap \text{Dom}(f) \}. \quad (2)$$

- (ii) (“Defined” form of FTIA.) *If also each library operation in f is everywhere defined on its inputs, while evaluating \mathbf{y} , then f is everywhere defined on \mathbf{x} , that is $\text{Dom}(f) \supseteq \mathbf{x}$.*
- (iii) (“Continuous” form of FTIA.) *If in addition to (ii), each library operation in f is everywhere continuous on its inputs, while evaluating \mathbf{y} , then f is everywhere continuous on \mathbf{x} .*

It is important to note that the theorem holds in finite precision, not just at Level 1. The decoration system gives basic tools for checking the conditions for the “defined” and “continuous” forms, during evaluation of a function.

6.4. Related issues. When program code contains conditionals (including loops), the run time data flow and hence the computed expression generally depends on the input data—for instance the example in §11.8 where a function is defined piecewise. The user is responsible for checking that a property such as global continuity holds as intended in such cases. The standard provides no way to check this automatically.

The standard requires that at Level 2, for all interval types, operations and inputs, the interval part of a decorated interval operation equal the corresponding bare interval operation. This ensures that converting bare interval program code to use decorated intervals leaves the data flow entirely unchanged (provided no conditionals depend on decoration values)—hence the computed expression and the interval part of its result are unchanged. If this were not so, there might in principle be an arbitrarily large discrepancy between the bare and the decorated versions of a computation that contains conditionals.

7. Flavors

7.1. Flavors overview. The standard permits different interval *flavors*, which embody different foundational (Level 1) approaches to intervals. An implementation shall provide at least one flavor. For brevity, phrases such as “A flavor shall provide, or document, a feature” mean that the implementation of that flavor shall provide the feature, or its documentation describe it.

Flavor is a property of program execution context, not of an individual interval, therefore just one flavor shall be in force at any point of execution. It is recommended that at the language level, the flavor should be constant at the level of a procedure/function, or of a compilation unit.

A flavor is identified by a unique name. Certain flavors, termed **included**, are specified in this standard. The (*list to be confirmed*) flavors are the currently included flavors. The procedure for submitting a new flavor for inclusion is described in Annex B. An implementation that has both included and non-included flavors is not conforming as a whole, but the part (§3.1) comprising the included flavors might be conforming.

The flavor concept enforces a common core of behavior that different kinds of interval arithmetic shall share:

- (i) There is a set of *common intervals* whose members are—in a sense made precise in §7.4—intervals of any flavor.
- (ii) There is a set of library operations, identified by their names, that are required in all flavors; see Clause 9.
- (iii) There is a set of *common evaluations* of library operations, with common intervals as input, that give—in a sense made precise in §7.2—the same result in any flavor.

In item (iii) the result means the tightest mathematical (Level 1) result, ignoring any interval widening due to finite precision (Level 2).