CS 105

ε-CALCULUS

BY

PAUL L. RICHMAN

TECHNICAL REPORT NO. CS 105
AUGUST 16, 1968

COMPUTER SCIENCE DEPARTMENT
School of Humanities and Sciences
STANFORD UNIVERSITY

ε-CALCULUS*

By

Paul Lawrence Richman

ABSTRACT: We use recursive function theory to lay the basis for a partially constructive theory of calculus, which we call the ε-calculus. This theory differs from other theories that have grown out of recursive function theory in that

(1) it is directly related to the variable-precision computations used in scientific computation today, and

(2) it deals explicitly with intermediate results rather than ideal answers.

As $\varepsilon \to 0$, intermediate results in the ε-calculus approach their corresponding answers in the calculus. Thus we say "the ε-calculus approaches the calculus, as $\varepsilon \to 0$." It is hoped that investigations in the ε-calculus will lead to a better understanding of numerical analysis. Several new results in this direction are presented, concerning instability and also machine numbers. Discrete notions of limit, convergence, continuity, arithmetic, derivative and integral are also presented and analyzed.

---

*This research was supported by the National Science Foundation under contract NSF GP 5962, by the Office of Naval Research under contract N00014-67-A-0112-0029, and by the Air Force under contract AF1047-66.

Acknowledgment: Deepest thanks go to my advisor, George Forsythe, for suggesting the idea of an ε-calculus and for his encouragement and direction during the initial stages of my research. I owe a special debt of gratitude to Dana Scott for his many excellent conceptual and notational contributions to the ε-calculus.

Their fast and accurate typing has earned Gail D. Schwartz, Lynne E. Leontovich, and Phyllis A. Winkler special mention. Thanks also go to Dorothy T. McGrath for her artistic illustrations.

Above all, I thank my wife, Jackie, for her unfailing patience, encouragement and confidence. I dedicate this work to her.

This research was supported by the National Science Foundation under contract NSF GP 5962, by the Office of Naval Research under contract N00014-67-A-0112-0029, and Air Force AF1047-66.

Table of Contents

Table of Illustrations

Chapter 1: Introduction

1.1 Summary

By a "notion" we mean a property of or an operation defined on a function or functions. The calculus can be thought of as a collection of elementary notions such as limit, convergence, continuity, derivative and integral, together with certain proved relations between notions, such as the reciprocal relationship between integration and differentiation. Fundamental to all this are the concepts of a real number and a real function. In the usual textbook developments, these basic concepts are not presented constructively, the notions are not necessarily effective or computable in any sense and relations between notions are often proved unconstructively. This is in direct contrast to E. Bishop's Foundations of Constructive Analysis [B1] and in partial contrast to recursive or computable analysis (Turing, Mazur, Grzegorczyk, Goodstein, Specker, Klaua, Aberth and Kreisel, to mention a few researchers in the area). Bishop defines constructive concepts of real number and real function, develops constructive notions and proves relations between notions constructively. (He then goes on to constructive theories of sets, metric spaces, complex analysis, measure, integration, normed linear spaces, locally compact abelian groups and commutative algebras.) His work is based on Brouwer's intuitionistic mathematics. In their work, "constructive" is an undefined or primitive term. Recursive analysis also has constructive concepts of real number and real function (see [G2, pp. 61-2]) and deals with constructive notions, but it allows unconstructive proofs

(see Kreisel [K1, p. 101]) It is based on recursive function theory, initiated by Church. In recursive analysis, "constructive" is defined in terms of recursive functions.

Both of these constructive theories are presented in a way which makes them foreign to numerical computation as it is done on today's computers. Here, we use recursive function theory to develop a theory of not only constructive, but even finitely computable real functions and defined notions, which we call ε-functions and ε-notions; these represent the intermediate results which arise from numerical computation. We call the resulting theory ε-calculus. This theory is directly related to modern day numerical computation. ε-Functions are essentially defined over a finite set, $R(\varepsilon)$, of ε-precision machine numbers. $R(\varepsilon)$ approximates the real numbers and each ε-function and ε-notion approximates respectively a function and a notion from calculus. And, as $\varepsilon \to 0$, $R(\varepsilon)$ approaches (i.e. becomes dense in) the reals and each ε-function or ε-notion approaches (in a sense to be defined) its corresponding function or notion. Thus we say the ε-calculus is a discretization of the calculus such that, as $\varepsilon \to 0$, the ε-calculus approaches the calculus.

The value of the ε-calculus to numerical analysis is that it presents a model of variable-precision computations. The study of $R(\varepsilon)$, ε-functions and ε-notions within the context of this model should lead to a better understanding of numerical computation. Our principal results in this direction are

(1) a new and simple definition of numerical instability (the kind caused by propagation of roundoff-error) together with

a suggestive geometric characterization (ch. 3), and

(2) an algorithm for overcoming such instabilities (ch. 3 and ch. 4).

Other new results presented here include

(1) a characterization of the concept of variable-precision machine numbers (sec. 2.2), and

(2) two new definitions of computable real functions, one allowing functions with discontinuities (ch. 7).

Before we present the ε-calculus, we give a motivating example to point out some of the basic problems involved in forming such a theory (i.e., involved in going from ideal mathematics to actual numerical computation), and to develop some of our basic notation.

1.2 A Motivating Example

Let us use "precision of computation" in a general way to mean the accuracy of a given mathematical approximation together with the precision of the arithmetic used to evaluate this approximation. It is often said that "numerical analysis is not very interesting because all you have to do to get more accuracy in a numerical result is increase the precision of computation." As a broad and optimistic point of view, the above statement is quite reasonable. But, when applied to particular cases, it can be quite false. Increasing the precision of computation can drastically decrease the accuracy of the result! For example, consider an algorithm which uses

$$f(x,y) = (g(y) - g(x))/(y-x)$$

to approximate $\ell \equiv \frac{d}{dt} g(t)\big|_{t=x}$. Fix x. For simplicity, suppose $f(x,y) \to \ell$ monotonically as $|y-x| \to 0$, and that $f(x,y_1)$ is computed in a certain form of single-precision arithmetic to give a single-precision approximation, $F(\varepsilon_1; x, y_1)$, to ℓ (here, "ε_1" denotes "single-precision"). This would be the value of the ε_1-limit corresponding to $\lim_{y \to x} f(x,y)$. We can increase the precision of computation by

(1) replacing y_1 by y_2 with $0 < |y_2-x| < |y_1-x|$, yielding a more accurate mathematical approximation, $f(x,y_2)$ (more accurate because of the monotonicity assumption), and

(2) evaluating $f(x,y_2)$ in a certain form of double-precision arithmetic, yielding a double-precision approximation,

$F(\varepsilon_2; x, y_2)$, to ℓ (here, "ε_2" denotes "double-precision" This would be the value of the ε_2-limit. But $F(\varepsilon_2; x, y_2)$ is not necessarily closer to ℓ than $F(\varepsilon_1; x, y_1)$; in fact, if y_2 is too close to x , $F(\varepsilon_2; x, y_2)$ may be much worse than $F(\varepsilon_1; x, y_1)$ (e.g., see example 3.1-2, where $g(x)$ is taken to be $x + 1$) This is illustrated in figure 1.2-1 as three graphs (with x fixed)

(a) $f(x,y)$ versus $1/(y-x)$,

(b) $F(\varepsilon_1; x, y)$ versus $1/(y-x)$, and

(c) $F(\varepsilon_2; x, y)$ versus $1/(y-x)$,

where y varies in the interval $(x, x+1]$.

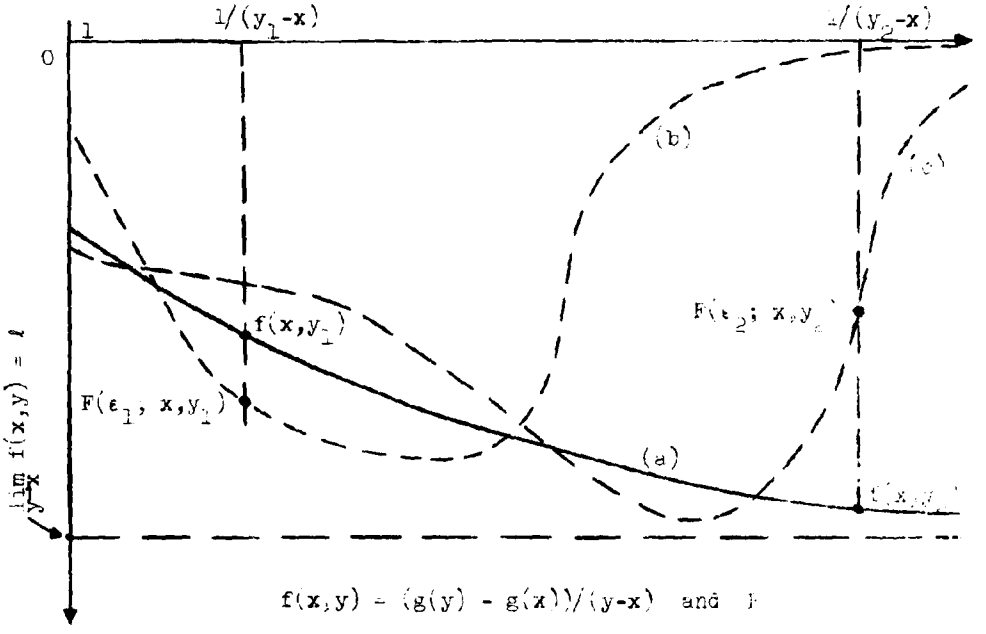

FIGURE 1.2-1

Notice that graph (b) stays close to f for awhile, but then falls off sharply to zero. Graph (c) stays close to f for awhile longer, but then it too falls off to zero. In general, F which exhibit such behavior are called unstable (this is discussed in detail in ch. 3). See Riesel [R1] for a similar example.

The tools normally used to deal with such instabilities are roundoff-error bounds, RF, and truncation-error bounds, TF. RF bounds the error incurred by using F in place of f ; TF bounds the error incurred by using $f(x,Y)$ in place of $\lim_{y \to x} f(x,y)$. And RF + TF bounds the error incurred by using F in place of $\lim_{y \to x} f(x,y)$. RF and TF are shown in figure 1.2-2, which is a redrawing of figure 1.2-1.

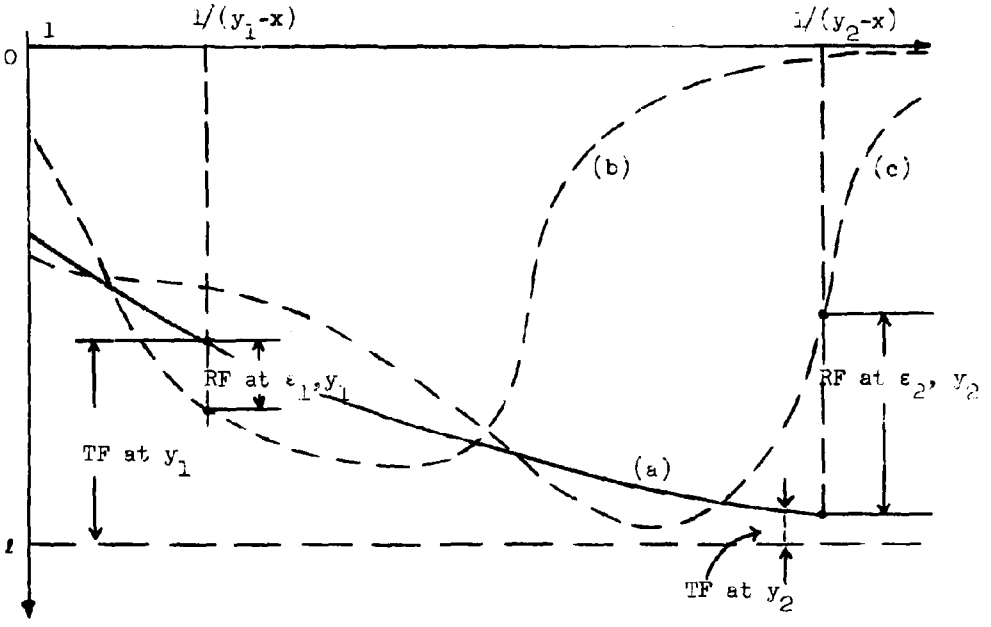

FIGURE 1.2-2
RF and TF

Of course, roundoff error is just a particular kind of truncation error; namely, it is the truncation error caused by using $F(\varepsilon; x, y)$ in place of $\lim_{\varepsilon \to 0} F(\varepsilon; x, y)$ (which is $= f(x,y)$). Both errors are caused by replacing infinite processes by finite processes. However, it is useful to distinguish roundoff error from truncation error so that they can be dealt with separately. The motivation for introducing these bounds comes from R. E. Moore's theory of interval analysis [M3, M4]. The key idea is that such bounds can be used to give precise information about a numerical result; i.e., an interval which contains the result.

1.3 An Outline

In the next two sections we give our basic notation and we discuss recursive natural functions and recursive natural operators. In chapter 2, we present the basic concept of a variable-precision computation, including the concepts of machine numbers, real inputs, subroutines, ε-functions and ε-operators. We give three examples of machine-numbers: floating-point, logarithmic and rational. And we give our main reason for introducing truncation-error bounds.

In chapter 3, we define and discuss numerical instability. A geometric characterization of instability is given which leads to the concept of "an ε-wave", to a proof that there is some desirable behavior even in the presence of instability (thm. 3.3-1), and finally to a (very inefficient) algorithm for overcoming instability (def. 3.3-4). This motivates our definition of ε-limit, given in the first section of chapter 4 (defs. 4.1-1, 4.1-2). We prove that, under certain conditions, the ε-limit of an ε-function approaches the limit of its corresponding ideal function as $\varepsilon \to 0$ (thm. 4.1-1). This ε-limit is shown to be a potentially efficient algorithm for overcoming instabilities of an approximation function by using a stably convergent truncation-error bound.

In section 4.2, we define ε-comparison relations, $<_\varepsilon$ and $=_\varepsilon$, and we prove that the truth-value of the ε-comparison of two real inputs approaches the truth-value of their comparison as $\varepsilon \to C$. These considerations are basic to what follows, and must be understood. In section 4.3, we use these ε-comparison relations to define ε-convergence and ε-continuity (pointwise). Again we prove that

these ε-notions approach (in a certain sense) their corresponding notions as $\varepsilon \to 0$. In section 4.5, we do the same for ε-convergence and ε-continuity over intervals. In preparation for this, we prove in section 4.4 some theorems about the kinds of discontinuities an ideal function can have while there exists an ε-function corresponding to it. These latter results are also made use of when we define ε-integrability in section 6.2.

In chapter 5, we define ε-operators for ε-arithmetic, ε-limit, ε-composition and ε-recursion. We also define two initial ε-functions, the identity and the constant ε-functions. The choice of these ε-operators and initial ε-functions was motivated by the operations and initial functions used in Mendelson [M1, pp. 120-1] to define the recursive natural functions. We illustrate the use of these ε-operators and ε-functions by using them to define an ε-function corresponding to e^x .

In chapter 6, we use the ε-operators and initial ε-functions of chapter 5 together with the ε-convergence of section 4.3 to define ε-differentiability and the ε-derivative, ε-integrability and ε-integral. In section 6.3, we prove the ε-calculus analog to the fundamental theorem of calculus.

In chapter 7, we define two notions of computable real function (based on ε-functions), and we prove that one of them is equivalent to one of the standard definitions from recursive analysis. We also prove that the operators and initial functions of chapter 5 are complete, in a certain sense.

The discussions of ε-convergence and ε-continuity in sections

4.3 and 4.5 and of ε-derivative and ε-integral in chapter 6 are only of definitional interest. Chapter 5 and the rest of chapter 4 are of more general interest; developments presented there should be useful in extending our theory.

1.4 Notation

Next our basic notation is presented. We begin with a list, using S and T to denote sets and m, an integer ≥ 0:

Symbol or Expression	Meaning
$=$	equal in numeric value
$\equiv$	equivalent
$\Rightarrow$	implies
$\Leftrightarrow$	if and only if
$\in$	set membership
$\cup$	union
$\cap$	intersection
$\subset$	inclusion: $S \subset T \Leftrightarrow (x \in S \Rightarrow x \in T)$
$S-T$	$S \cap$ (the complement of T)
"{" , "}"	used only in defining sets
$\{\ \}$	the null set
$\bar{x}_m$	if $m \geq 1$, the list $x_1, x_2, \ldots, x_m$; $\bar{x}_0$ is the empty list
$S^{(m)}$	if $m \geq 1$, $S \times \ldots \times S$ (m-fold); $S^{(0)}$ is $\{\bar{x}_0\}$
$f:S^{(m)} \to T$	f is a function from $S^{(m)}$ to T
$-$ (bar)	generally indicates repetition on a subscript, as in $\bar{x}_m$
$[x]$	greatest integer in x
n	$\{0, 1, 2, \ldots\}$
$R(\varepsilon)$	ε-precision machine numbers (sec. [illegible])

$\mathcal{M}$	the set of all machine numbers (sec. 2.2)
$\hat{I}(\cdot)$, $\check{I}(\cdot)$	conversion functions (sec. 2.8)
$\hat{+}$, $\check{+}$, $\hat{-}$, ...	roundup and rounddown machine arithmetic (sec. 2.8)
$N_\varepsilon(a)$	machine neighborhood of a (sec. 2.8)
$=_\varepsilon$, $<_\varepsilon$	ε-comparison relations (sec. 4.2)
$\mathcal{F} \approx f(P)$	$\mathcal{F}$ approximates f over P (sec. 2.5)
bool [statement]	= 1 if statement is true, = 0 otherwise (sec. 4.2)

Note that $S^{(0)} \neq \{\ \}$, since $\bar{x}_0 \in S^{(0)}$. Define R and $\bar{R}$ by

$$R \equiv \{x: \ x \text{ is a finite real number}\},$$

$$\bar{R} \equiv R \cup \{-\omega, \omega, \omega\};$$

ω stands for "undefined". Thus "$x = \omega$" means "x is undefined in terms of the members of $R \cup \{-\omega, \omega\}$". We will treat ω like any other point in $\bar{R}$, except that $\omega > x$ for all $x \in \bar{R} - \{\omega\}$, and ω is isolated from the rest of $\bar{R}$ (the null set is the only neighborhood of ω). For example, $\omega - \omega = \omega$, $0/0 = \omega$, $\omega + 3 = \omega$, $(-1) \times \omega = \omega$, $\lim_{i \to \omega} (-1)^i = \omega$, etc. All our constants, variables and functions will take values in $\bar{R}$.

We will use the usual neighborhood definition of limit for the doubly extended real line, with the addition that a limit which does not exist in the usual sense has the value ω.

To simplify inequalities, we let

$$|a-b| = 0 \qquad \text{if } a = b ,$$

even when $a = b \in \{-\infty, \infty, \omega\}$. We do this because we use $|a-b|$ to measure the distance between a and b. It is easy to show that this distance function satisfies the triangle inequality,

$$|a-b| \leq |a-c| + |c-b| ,$$

for $a, b, c \in \overline{\mathbb{R}}$. (In showing this, it is best to refer to the special rules for arithmetic involving $\pm\infty$ and ω given in sec. 2.8).

We will use notation of the form

$$B = \begin{cases} A_1 & \text{if } R_1 \\ \cdot & \cdot \\ \cdot & \cdot \\ \cdot & \cdot \\ A_{n-1} & \text{if } R_{n-1} \\ A_n & \text{otherwise,} \end{cases}$$

to indicate that $B = R_j$ $(1 \leq j \leq n)$ when j is the smallest integer such that R_j is true, and R_n is defined to be true always.

1.5 Recursive Functions and Operators

Inductive schemes for defining recursive natural functions can be found in Mendelson [M1, pp. 120-1] and elsewhere. Let $\mathfrak{n}$ be the set of nonnegative integers, 0, 1, ... The recursive functions from $\mathfrak{n}^{(m)} \to \mathfrak{n}$ $(m \geq 0)$ are those mappings from $\mathfrak{n}^{(m)} \to \mathfrak{n}$ which obtain the image point via constructive operations on the domain point. Recursive functions essentially characterize the input/output of Turing machines.

By recursive operator we mean a standard recursive functional with its integer arguments left unspecified. Thus a recursive operator φ of n function arguments constructively maps n functions, $\alpha_i : \mathfrak{n}^{(m_i)} \to \mathfrak{n}$ $(i = 1, \ldots, n)$, into a function, $\varphi(\alpha_1, \ldots, \alpha_n) : \mathfrak{n}^{(m)} \to \mathfrak{n}$ for some $m \geq 0$. Inductive schemes for defining recursive functionals may be found in Schoenfield [S2], Grzegorczyk [G1], Klaua [K1] and elsewhere.

The reader does not need to know any more about recursive functions and recursive operators than what we have just stated. We will not use their inductive definitions. As usual, we say a process is "effective" or "constructive" precisely when that process can be carried out by purely mechanical means (i.e., by a Turing machine).

REMARKS: In section 1.2 we saw that increasing the precision of computation may decrease the accuracy of a numerical result. In section 5.3 we show that this does not apply to purely arithmetic processes, i.e., rational function evaluation. There, increasing precision ultimately leads to convergence. The trouble arises when limits are involved: e.g., $\lim_{y \to x} f(x,y)$ in section 1.2, and $\lim_{n \to \infty} \sum_{i=1}^{n} a_i$ in Riesel's example [R2]. The root of this trouble is an interchange of limits which may not work. This merits further explanation. (Keep the example of section 1.2 in mind for the following.) In general, we will have

$$\lim_{\varepsilon \to 0} F(\varepsilon; x,y) = f(x,y)$$

for $y \neq x$. This implies

$$\lim_{y \to x} \lim_{\varepsilon \to 0} F(\varepsilon; x,y) = \lim_{y \to x} f(x,y) \quad .$$

But in order to compute successive approximations to this limit, we must define an ε-limit ε-operator, $\underset{y \to x}{\mathrm{LIM}}$, such that $\underset{y \to x}{\mathrm{LIM}}\, F(\varepsilon; x,y)$ is the finitely computable ε-approximation to $\lim_{y \to x} f(x,y)$ and such that

$$\lim_{\varepsilon \to 0} \underset{y \to x}{\mathrm{LIM}}\, F(\varepsilon; x,y) = \lim_{y \to x} f(x,y) \quad .$$

This interchange of limits is investigated in chapter 3.

BLANK PAGE

Chapter 2: Basic Concepts

2.1 Variable-Precision Computations

Before launching into a description of our model, we first loosely describe the kinds of finite computations which we are interested in modeling. These are characterized by having the following facilities:

(1) variable-precision machine numbers;

(2) the ability to make decisions based on a comparison of the values of machine numbers;

(3) variable-precision input routines for inputting the members of $\mathbb{R}$ into machine numbers (at a specified precision), and for giving roundoff-error bounds on these inputted values; and

(4) variable-precision arithmetic (ε-arithmetic) for machine numbers.

We formalize (1) - (3) in this chapter; (4) is formalized in section 5.3.

2.2 Machine Numbers and Their Comparison

Let $S_i (i = 1, 2, \ldots)$ be a finite subset of $\tilde{R}$ with $\{-\infty, \infty, \omega\} \subset S_i$. In order to formalize (1) and (2) above, we define the concept "S_i is effectively generable from i." Essentially we mean by this that there is an algorithm, with input parameter i, which produces as output all the elements of S_i. But this is not precise because we as yet have not said what "produces as output a member x of $\tilde{R}$" means: e.g., x may have a nonrepeating decimal expansion, so our algorithm cannot in general produce x by producing a decimal expansion for x. We define this concept precisely as follows. Let $\mathfrak{n}$ be as in section 1.5. Define the function, θ: $\mathfrak{n}^{(2)} \to \{-\infty, \ldots, -1, 0, 1, \ldots, \infty, \omega\}$, by

$$\theta(i, j) = \begin{cases} -\infty & \text{if } i = j = 1 \\ \infty & \text{if } i = j = 2 \\ \omega & \text{if } i = j = 3 \\ i-j & \text{otherwise.} \end{cases} \tag{2.2-1}$$

Let α_1 and α_2 be functions from $\mathfrak{n}$ to $\mathfrak{n}$ and define $\alpha(\cdot) \equiv \theta(\alpha_1(\cdot), \alpha_2(\cdot))$. We say <u>$\alpha_1, \alpha_2$ compute $a \in \tilde{R}$</u> (or <u>α computes a</u>) precisely when

$$|a - \alpha(n)/n| \leq 1/n \qquad \text{for } n = 1, 2, \ldots,$$

and we use a and $<\alpha>$ interchangeably. See Grzegorczyk [G2, p. 61] for a similar definition. When α_1 and α_2 are recursive and α_1, α_2 compute a, we say <u>a (or $<\alpha>$ is a</u>

computable number. For given α_1, α_2: $\mathbb{N}^{(m+1)} \to \mathbb{N}$ $(m \geq 1)$ and a fixed $\bar{x}_m \in \mathbb{N}^{(m)}$, we treat $\alpha(\bar{x}_m, \cdot) \equiv \theta(\alpha_1(\bar{x}_m, \cdot), \alpha_2(\bar{x}_m, \cdot))$ like a function of one variable, so that when $\alpha(\bar{x}_m, \cdot)$ computes $a \in \mathfrak{R}$ we use $< \alpha(\bar{x}_m, \cdot) >$ interchangeably with a.

Let $S_1, S_2, \ldots$ be as described above. Let $\ell(i)$ be the number of elements in S_i. We say S_i is effectively generable from i precisely when $\ell(\cdot)$ is recursive and there are recursive functions, α_1, α_2: $\mathbb{N}^{(3)} \to \mathbb{N}$, such that, with $\alpha(\cdot) \equiv \theta(\alpha_1(\cdot), \alpha_2(\cdot))$, we have

$$S_i = \{ < \alpha(i, 1, \cdot) >, < \alpha(i, 2, \cdot) >, \ldots, < \alpha(i, \ell(i), \cdot) > \},$$
$$< \alpha(i, 1, \cdot) > = -\infty, < \alpha(i, 2, \cdot) > = \infty, < \alpha(i, 3, \cdot) > = \omega .$$

We call the pair (α, ℓ) a generator of S_i.

The concept of variable-precision machine numbers is formalized as follows. We use the positive real constants, $\varepsilon_1, \varepsilon_2, \ldots$, to denote the possible levels of precision: ε_1 denotes single-precision, etc. We use ε to denote a variable which takes values in the set

$$\mathcal{E} \equiv \{\varepsilon_1, \varepsilon_2, \ldots\} . \tag{2.2-2}$$

DEFINITION 2.2-1: A machine number system, $(\mathcal{R}, \mathcal{E})$, is a set $\mathcal{E}$ of constants $\varepsilon_1, \varepsilon_2, \ldots$ together with a mapping, $\mathcal{R}$: $\mathcal{E} \to$ (set of subsets of $\mathfrak{R}$), where

(I) $\varepsilon_i \to 0$ strictly monotonically as $i \to \infty$,

(II) $\bigcup_{i \geq 1} \mathcal{R}(\varepsilon_i)$ is dense in R,

(III) $\mathcal{R}(\varepsilon_1) \subset \mathcal{R}(\varepsilon_2) \subset \ldots$,

(IV) $\mathcal{R}(\varepsilon)$ is finite, for each ε in $\mathcal{E}$,

(V) $\{0, 1, -\infty, \infty, \omega\} \subset \mathcal{R}(\varepsilon_1)$, and

(VI) $\mathcal{R}(\varepsilon_i)$ is effectively generable from i.

$\mathcal{R}(\varepsilon)^{(n)}$ represents a discretization of Euclidean n-space. Condition I reflects the statement "decreasing ε increases the precision"; no other use of the values of the ε_i will be made in this paper. Conditions II and VI are the only really essential restrictions; if $\mathcal{R}$ and $\mathcal{E}$ satisfy them, minor modifications will produce an $\mathcal{R}'$ and $\mathcal{E}'$ which satisfy I-VI. (If I is violated, replace $\mathcal{E}$ by any $\mathcal{E}'$ satisfying it; if III is violated, let $\mathcal{R}'(\varepsilon_i) \equiv \bigcup_{j=1}^{i} \mathcal{R}(\varepsilon_j)$, etc.) Condition II allows us to get at any number in R through the exclusive use of machine numbers. The nesting condition, III, says that we may reuse, at precision ε_{i+1}, any machine number that we used at precision ε_i; this will be used in dealing with instability and in the proof that our ε-limit approaches limit as $\varepsilon \to 0$. Condition IV will simplify our treatment of instability. Condition VI means that $\mathcal{R}$ could really be used as the basis of a variable-precision number system on a computer; it insures that the switching of precision can be done automatically. (We investigate other implications of this condition below.) That 0 and 1 are in $\mathcal{R}(\varepsilon)$ will prove convenient in many situations, but never will this be a necessity. However, having $-\infty, \infty, \omega \in \mathcal{R}(\varepsilon)$ greatly simplifies our model. We give three examples to clarify these ideas and to show the variety of machine number systems which satisfy I-VI.

Example 2.2-1: Let β be an integer ≥ 2. Let $0.a_1 a_2 \cdots a_i$

denote a base β fraction. Define $\varepsilon_i^{\odot} = 1/(2\beta^{i-1}+1)$.

A base β normalized floating-point number system is given by

$$\mathcal{E}^{\odot} \equiv \{\varepsilon_1^{\odot}, \varepsilon_2^{\odot}, \ldots\},$$

$$\mathcal{R}^{\odot}(\varepsilon_i^{\odot}) \equiv \{0, -\infty, \infty, \omega\} \cup \{x: |x| = 0.a_1 a_2 \cdots a_i \times \beta^e,$$

$a_1 > 0$, and e is an integer, $|e| < \beta^i\}$.

We chose these $\varepsilon_i^{\odot}$ because each real number x such that $\beta^{-\beta^i} \leq |x| \leq (\beta^i-1)\beta^{\beta^i-i-1}$ can (in principle) be inputted into $\mathcal{R}^{\odot}(\mathcal{E}_i^{\odot})$ with a relative error $\leq$ this $\varepsilon_i^{\odot}$.

EXAMPLE 2.2-2: Let β_1 be some finite number > 1. Define $\beta_{i+1} = \beta_i^{1/10}$ and $\varepsilon_i^* = 1-2/(\beta_i+1)$. A base β_1 logarithmic number system is given by

$$\mathcal{E}^* \equiv \{\varepsilon_1^*, \varepsilon_2^*, \ldots\}$$

$$\mathcal{R}^*(\varepsilon_i^*) \equiv \{0, -\infty, \infty, \omega\} \cup \{x: |x| = \beta_i^e, e \text{ an integer}, |e| < 10^{2i}\}$$

We chose these ε_i^* because each real number x, with $\beta_1^{-10^{i+1}+10^{1-i}} \leq |x| \leq \beta^{10^{i+1}-10^{1-i}}$, can (in principle) be inputted into $\mathcal{R}^*(\varepsilon_i^*)$ with a relative error $\leq$ this ε_i^*. We had to use different bases, β_i, approaching 1 as $i \to \infty$, so that condition II is met. The fact that some of the β_i will be irrational causes no difficulties.

EXAMPLE 2.2-3: Define $\varepsilon_i^{\#} = 1/(10^i+1)$ and

$$\mathcal{E}^{\#} \equiv \{\varepsilon_1^{\#}, \varepsilon_2^{\#}, \ldots\},$$

$$\mathcal{R}^{\#}(\varepsilon_i^{\#}) \equiv \{0, -\infty, \infty, \omega\} \cup \{x: x = p/q \text{ for integers } p \text{ and } q$$

$$\text{with } |p|, |q| < 10^i\}.$$

We chose these $\varepsilon_i^{\#}$ because each real number x, with $1/(10^i-1) \leq |x| \leq 10^i-1$ can (in principle) be inputted into $R^{\#}(\varepsilon_i^{\#})$ with a relative error $\leq$ this $\varepsilon_i^{\#}$. (This last statement is more difficult to prove than the corresponding statements of the other examples, so a proof is included in the appendix to this chapter. The other statements are also proved there as simple corollaries.)

Another important property which R must possess is that the members of $R(\varepsilon_i)$ must be representable in some simple form which varies with i in a simple way. This is necessary so that the members of $R(\varepsilon)$ can be represented simply in the computer. For example, any $x \in R^{\odot}(\varepsilon_i^{\odot}) - \{0, -\infty, \infty, \omega\}$ can be represented by a pair of integers (a, e) with $\beta^{i-1} \leq |a| < \beta^i$ and $|e| < \beta^i - i$ since x must equal $(a\,\beta^{-i})\,\beta^e$ for some such a and e. And any x in $R^*(\varepsilon_i^*) - \{0, -\infty, \infty, \omega\}$ can be represented by a pair of integers $(\pm i, e)$ with $|e| < 10^{2i}$, since x must equal $\pm \beta_i^e$ for some such e. A similar statement can be made about $R^{\#}$. In fact, any R which satisfies condition VI possesses this representability property. Suppose (α_R, ℓ_R) is a generator for $R(\varepsilon_i)$. Then any x in $R(\varepsilon_i)$ can be represented by a pair of integers (i, j), with $1 \leq j \leq \ell_R(i)$, since x must equal some $< \alpha_R(i, j, \cdot) >$. So much for representability. Next we consider comparison of machine numbers.

The fact that each member of $R(\varepsilon_i)$ has a unique representation in terms of α_R means that, given i, j and k, we can effectively decide whether $< \alpha_R(i, j, \cdot) >$ is $>$, $<$ or $= < \alpha_R(i, k, \cdot) >$;

when $j \neq k$ we have $< \alpha_{\mathcal{R}}(i, j, \cdot) > \neq < \alpha_{\mathcal{R}}(i, k, \cdot) >$ and we can determine which is greater by computing $\alpha_{\mathcal{R}}(i, j, n)$ and $\alpha_{\mathcal{R}}(i, k, n)$ for some finite number of values of n. This gives us (2) of section 2.1. Further, the following two conditions imply the existence of a generator for S_i:

(1) there are recursive functions α'_1, α'_2, ℓ' such that, with $\alpha'(\cdot) \equiv \theta(\alpha'_1(\cdot), \alpha'_2(\cdot))$, we have

$$S_i \equiv \bigcup_{j=1}^{\ell'(i)} \{ < \alpha'(i, j, \cdot) > \} \qquad \text{for } i = 1, 2, \ldots ,$$

(2) the relation $< \alpha'(i, j, \cdot) > = < \alpha'(i, k, \cdot) >$ is effectively decidable from j and k.

Thus condition VI on $\mathcal{R}$ is not too restrictive in (implicitly) requiring $(\alpha_{\mathcal{R}}, \ell_{\mathcal{R}})$ to be nonredundant.

Throughout the rest of this paper, we assume that $\mathcal{R}$, $\mathcal{E}$ and a corresponding generator $(\alpha_{\mathcal{R}}, \ell_{\mathcal{R}})$ are given and fixed. All of the following definitions are implicitly relative to these $\mathcal{R}$, $\mathcal{E}$ and $(\alpha_{\mathcal{R}}, \ell_{\mathcal{R}})$.

For later use, we define the <u>machine number set</u>, $\mathcal{M}$, by

$$\mathcal{M} = \bigcup_{i \geq 1} \mathcal{R}(\varepsilon_i) . \tag{2.2-3}$$

If $\mathcal{R}$ and $\mathcal{E}$ are the $\mathcal{R}^{\#}$ and $\mathcal{E}^{\#}$ of example 2.2-3, then $\mathcal{M}$ is $\mathcal{M}^{\#}$, the set of rationals together with $-\infty$, ∞ and ω.

2.3 Input Routines

We handle the inputting of members of $\mathfrak{R}$ essentially by assuming that members of $\mathfrak{R}$ are given by giving an input routine. To avoid confusing a number with its input routine, we introduce a new concept, that of a real input.

DEFINITION 2.3-1: A real input $x \equiv (X, RX)$ is a pair of mappings, $X, RX: \mathcal{E} \to \mathfrak{M}$, such that

(1) $X(\varepsilon)$, $RX(\varepsilon) \in \mathfrak{A}(\varepsilon)$ and $RX(\varepsilon) \geq 0$ for all $\varepsilon \in \mathcal{E}$,

(2) $|X(\varepsilon) - X(\eta)| \leq RX(\varepsilon) + RX(\eta)$ for all $\varepsilon, \eta \in \mathcal{E}$,

(3) $\lim_{\varepsilon \to 0} RX(\varepsilon) = 0$.

If $x \equiv (X, RX)$ satisfies (1) and (2), we call x a poor real input. We call (X, RX) an input routine for x.

It follows that the numeric value corresponding to the real input x is just $\lim_{\varepsilon \to 0} X(\varepsilon)$. When a real input x is used in a context that calls for a numeric value, we let that value be $\lim_{\varepsilon \to 0} X(\varepsilon)$. Thus for each ε we have

$$RX(\varepsilon) \geq |X(\varepsilon) - x| , \tag{2.3-1}$$

and so $RX(\varepsilon)$ is just a roundoff-error bound, bounding the error caused by using $X(\varepsilon)$ in place of (the numeric value of) x.

A real input can be thought of as a variable ranging not only over $\mathfrak{R}$ but also over input routines. We will find it unnecessary to distinguish notationally between a real variable (ranging over $\mathfrak{R}$) and a real input, or between a real constant and a fixed real input. When a real input is named x, its input routine will be

named (X, RX), etc. Note that for real inputs x and y, $x = y$ ("=" means "equal in numeric value" throughout this paper, see sec 1.4) precisely when

$$(2.3\text{-}2) \qquad |X(\varepsilon) - Y(\varepsilon)| \leq RX(\varepsilon) + RY(\varepsilon) \quad \text{for all } \varepsilon \in \mathcal{E} .$$

This relation will be useful in defining ε-comparison relations.

The following conventions will simplify notation later:

(1) when, in a given context, the value of the real input x is known to be in $\mathcal{R}(\varepsilon)$ and (X, RX) has not been explicitly specified, it will be assumed that $X(\delta) = x$ (in value) and $RX(\delta) = 0$ for all $\delta \leq \varepsilon$,

(2) when we state that $x \in \mathcal{M}$, for a real input x, we mean that $[RX(\varepsilon) = 0$ or $RX(\varepsilon) < \infty = |X(\varepsilon)|$ for some $\varepsilon]$ and $x \notin \mathcal{M}$ means $[|x| \neq \infty$ and $RX(\varepsilon) \neq 0$ for all $\varepsilon]$, and

(3) we will use the same notation for sets of numbers and sets of real inputs. If P is (in a given context) a set of numbers and it is not known that $P \subset \mathcal{M}$, then the set, P, of real inputs contains all real inputs x with value in the number set P . If the number set P is known to be a subset of $\mathcal{M}$, then the real input set P contains all real inputs x with value in the number set P such that $x \in \mathcal{M}$ (under convention (2) above).

By (3), $\tilde{R}$ may denote the set of all numbers or the set of all real inputs, depending on context. We use a rule analogous to (3) when I is a set of m-tuples of numbers.

2.4 Multiple-Precision Subroutines

Let $\bar{x}_0$ denote the empty list and for $m \geq 1$ let $\bar{x}_m$ denote $x_1, x_2, \ldots, x_m$. For the moment, let $\bar{x}_m$ be $m \geq 0$ variable poor real inputs and let $x_{m+1}, \ldots, x_{m+n}$ be $n \geq 0$ fixed real inputs. A multiple-precision subroutine of $m \geq 0$ variables and $n \geq 0$ constants is essentially a computer subroutine with input ε and with access to any finite number of values of the input routines for $\overline{x_{m+n}}$, say

$$\left.\begin{matrix} X_j(\varepsilon_1),\ X_j(\varepsilon_2), \ldots,\ X_j(\varepsilon_M) \\ RX_j(\varepsilon_1),\ RX_j(\varepsilon_2), \ldots,\ RX_j(\varepsilon_M) \end{matrix}\right\} \quad \text{for } j = 1, \ldots, m+n \ , \tag{2.4-1}$$

and with output in $\mathfrak{R}(\varepsilon)$, with the requirement that if any $X_j(\varepsilon)$ or $RX_j(\varepsilon)$ is ω then the output is ω. We call ε and $\bar{x}_m$ the inputs. With inputs ε and $\bar{x}_m$, we denote the output value of the subroutine F by $F(\varepsilon; \bar{x}_m)$. If for some $j \leq m$ we have $\lim_{\varepsilon \to 0} \sup RX_j(\varepsilon) > 0$, then we allow F to not halt when its inputs are $\bar{x}_m$ and (any) ε. When this happens we write $F(\varepsilon; \bar{x}_m) = \omega$ because $F(\varepsilon; \bar{x}_m)$ is undefined in terms of the members of $\tilde{R} - \{\omega\}$ (see sec. 1.4).

This can be formalized as follows. For $n = 1, 2, \ldots$ let p_n be the n^{th} prime: $p_1 = 2$, etc. Let x be a poor real input and suppose that for $i \geq 1$,

$$X(\varepsilon_i) = \langle \alpha_{\mathfrak{R}}(i, k_{2i-1}, \cdot) \rangle \quad RX(\varepsilon_i) = \langle \alpha_{\mathfrak{R}}(i, k_{2i}, \cdot) \rangle \ .$$

For $j \geq 1$ define

$$gn_j(x) = p_1^{k_1} \times p_2^{k_2} \times \ldots \times p_{2j}^{k_{2j}} \ .$$

Define the j^{th} Gödel number of $\bar{x}_m$ $(m \geq 0)$ by

$$(2.4\text{-}2) \qquad GN_j(\bar{x}_m) = p_1^{gn_j(x_1)} \times p_2^{gn_j(x_2)} \times \ldots \times p_m^{gn_j(x_m)} ,$$

where the empty product is 1 (i.e., $GN_j(\bar{x}_0) = 1$). $GN_M(\bar{x}_m)$ contains complete information about the numbers shown in (2.4-1). We say that γ is m-determining $(m \geq 1)$ precisely when γ is recursive and, for $i = 1, 2, \ldots$,

(1) $0 \leq \gamma(i, j) \leq \ell_R(i)$ for any j,

(2) for any poor real inputs $\bar{x}_m$, $\gamma(i, GN_k(\bar{x}_m)) = 0$ for $1 \leq k < i$, and if any $X_j(\varepsilon_i)$ or $RX_j(\varepsilon_i)$ is ω then $\gamma(i, GN_i(\bar{x}_m)) = 3$, and

(3) for any real inputs $\bar{x}_m$ there is an $M \geq i$ with $\gamma(i, GN_M(\bar{x}_m)) \neq 0$.

Thus γ waits until sufficient information about $\bar{x}_m$ has been collected (in $GN_M(\bar{x}_m)$), and then γ returns a nonzero value. When $\liminf_{\varepsilon \to 0} RX_j(\varepsilon) > 0$ for some j, $GN_k(\bar{x}_m)$ may never contain enough information about x_j for γ to return a nonzero value. (Of course, even when $\liminf_{\varepsilon \to 0} RX_j(\varepsilon) = 0 < \limsup_{\varepsilon \to 0} RX_j(\varepsilon)$ for some j we may have $\gamma(i, GN_k(\bar{x}_m)) = 0$ for some i and all k, but this will not be due to lack of information about x_j.) We say γ is 0-determining precisely when γ is recursive and $1 \leq \gamma(i, 1) \leq \ell_R(i)$. Let $\tilde{\mu}_n(\gamma(i, GN_n(\bar{x}_m)) \neq 0)$ be $\gamma(i, GN_M(\bar{x}_m))$ where M is the least value of n such that $\gamma(i, GN_n(\bar{x}_m)) \neq 0$, or let it be 3 when there is no such n.

DEFINITION 2.4-1: A multiple-precision subroutine of $m \geq 0$ variables and $n \geq 0$ constants, $x_{m+1}, \ldots, x_{m+n}$, is a mapping, $F : \mathcal{E} \times \{\text{poor real inputs}\}^{(m)} \to \mathcal{M}$, such that there is an $m+n$-determining γ which satisfies the following for any poor real inputs $\bar{x}_m$ and any $i \geq 1$:

$$(2.4\text{-}3) \qquad F(\varepsilon_i; \bar{x}_m) = < \alpha_R(i, \tilde{\mu}_n(\nu(i \;\; GN_n(\bar{x}_{m+n})) \neq 0), \cdot) > \; .$$

We say γ determines F relative to (α_R, ℓ_R). We stress that so long as the $\bar{x}_m$ are real inputs (not just poor), the computation of $F(\varepsilon; \bar{x}_m)$ via γ will always halt. Essentially the only subroutines whose computation may fail to halt are those which, with inputs ε and $\bar{x}_m$, try to find an $\eta \leq \varepsilon$ such that, say $RX_1(\eta) \leq t$, for some tolerance level $0 < t < \infty$; for example, when x_1 is a poor real input with $RX_1(\varepsilon) = \infty$ for all ε, such a subroutine will fail to halt. We will use "subroutine" as an abbreviation for "multiple-precision subroutine." An immediate consequence of the above definitions is

$$F(\varepsilon; \bar{x}_m) = \omega \quad \text{if any } X_j(\varepsilon) \text{ or } RX_j(\varepsilon) \text{ is } \omega \; .$$

This convention is taken from Scott [S1]. Note that these definitions have all been relative to (α_R, ℓ_R).

THEOREM 2.4-1: If F is a subroutine relative to (α_R, ℓ_R) and (α'_R, ℓ_R) is some other generator of $R(\varepsilon_i)$ then F is a subroutine relative to (α'_R, ℓ_R).

This means that the concept of subroutine is independent of which generator of $\mathcal{R}(\varepsilon_i)$ one uses.

Proof: Let $GN_j(\bar{x}_m)$ and $GN'_j(\bar{x}_m)$ denote the j^{th} Gödel numbers of $\bar{x}_m$ relative to $(\alpha_{\mathcal{R}}, \ell_{\mathcal{R}})$ and $(\alpha'_{\mathcal{R}}, \ell_{\mathcal{R}})$, respectively. The considerations of section 2.2 show that there is an effective procedure which, when given any generator for $\mathcal{R}(\varepsilon_i)$ and any $j \geq 1$, can order the members of $\mathcal{R}(\varepsilon_j)$. This means that there are recursive functions φ_1 and φ_2 such that $\varphi_1(i, 0) = 0$ and

$$< \alpha'_{\mathcal{R}}(i, \varphi_1(i, j), \cdot) > = < \alpha_{\mathcal{R}}(i, j, \cdot) > ,$$

$$\varphi_2(i, GN'_i(\bar{x}_m)) = GN_i(\bar{x}_m) ,$$

for any $i \geq 1$ and $1 \leq j \leq \ell_{\mathcal{R}}(i)$. Define γ' by

$$\gamma'(i, GN'_j(\bar{x}_{m+n})) = \varphi_1(i, \gamma(i, \varphi_2(j, GN'_j(\bar{x}_{m+n})))) ,$$

for $i \geq 1$, $j \geq 1$ and any $\bar{x}_m$. If γ determines F relative to $(\alpha_{\mathcal{R}}, \ell_{\mathcal{R}})$ then γ' determines F relative to $(\alpha'_{\mathcal{R}}, \ell_{\mathcal{R}})$. This completes the proof.

2.5 ε-Function

An ideal function of m variables ($m \geq 0$) is a mapping, $f: \tilde{R}^{(m)} \to \tilde{R}$, with the constraint that $f(\bar{x}_m) = \omega$ if any $x_i = \omega$.

DEFINITION 2.5-1: For $m \geq 0$, an ε-function, $\mathfrak{F}$, of m variables corresponding to an ideal function f (of m variables) over a set P of m-tuples of real inputs is a triple (F, RF, TF) of subroutines, where for any real inputs $\bar{x}_m$ we have

(1) for each ε, $RF(\varepsilon; \bar{x}_m) \geq |F(\varepsilon; \bar{x}_m) - f(\bar{x}_m)|$,

(2) $[\bar{x}_m \in P \text{ and } f(\bar{x}_m) \neq \omega] \Rightarrow \lim_{\varepsilon \to 0} RF(\varepsilon; \bar{x}_m) = 0$, and

(3) if $m = 0$, 1 then $TF \equiv \omega$; otherwise, for all ε,

$$TF(\varepsilon; \bar{x}_m) \geq |f(\bar{x}_m) - \lim_{y \to x_{m-1}} f(\bar{x}_{m-1}, y)| \quad . \tag{2.5-1}$$

We call P a domain set of $\mathfrak{F}$ and we write

$$\mathfrak{F} \approx f(P) \quad ,$$

to be read "$\mathfrak{F}$ corresponds to (or approximates) f over (or mod) P." This definition is illustrated in figure 2.5-1 for the case $m = 2$ and $x = \infty$.

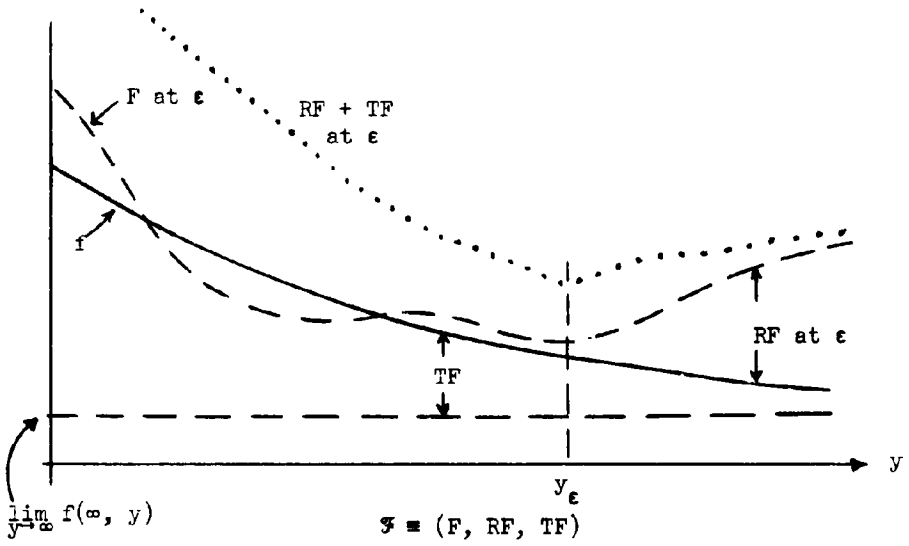

FIGURE 2.5-1

RF is a roundoff-error bound, bounding the error incurred by using F in place of f . TF is a truncation-error bound, bounding the error incurred by using $f(\bar{x}_m)$ in place of $\lim_{y\to\bar{x}_{m-1}} f(\bar{x}_{m-1}, y)$. For example, if $f(x, n) = \sum_{i=1}^{n} g(i)$, then $TF(\varepsilon; \infty, n)$ bounds the truncation error $|\sum_{i=n+1}^{\infty} g(i)|$. $RF(\varepsilon; \bar{x}_{m-1}, Y) + TF(\varepsilon; \bar{x}_{m-1}, Y)$ bounds the error incurred by using $F(\varepsilon; \bar{x}_{m-1}, Y)$ in place of $\lim_{y\to\bar{x}_{m-1}} f(\bar{x}_{m-1}, y)$. For the above graph, this bound is smallest when $Y = y_\varepsilon$.

Conditions (1) and (3) on $\mathcal{F}$ require that the bounds RF and TF work properly for <u>any</u> real inputs $\bar{x}_m$ and <u>any</u> ε . Condition (2) requires the convergence of F to f and RF to 0 at $\bar{x}_m \in P$ for

which $f(\bar{x}_m)$ is defined. (We must have $\exists F \neq \Omega$ in order to effectively compute, via F, an approximation to f that is correct to within some desired and arbitrarily small tolerance.) If instead of (2) we have

$$(2')[\bar{x}_m \in P \text{ and } f(\bar{x}_m) \neq \omega] \Rightarrow \lim_{\varepsilon \to 0} f(\varepsilon; \bar{x}_m) = f(\bar{x}_m) \quad ,$$

we say $\mathfrak{F}$ weakly corresponds to f over P and we write

$$\mathfrak{F} \sim f(P) \quad .$$

We call P a weak domain set of $\mathfrak{F}$. An immediate consequence of these definitions is

THEOREM 2.5-1: If $Q \subset P$ and $\mathfrak{F} \approx f(P)$ (or $\mathfrak{F} \sim f(P)$) then $\mathfrak{F} \approx f(Q)$ (or $\mathfrak{F} \sim f(Q)$).

We will use $\mathfrak{F}(\varepsilon; \bar{x}_m)$ to denote the triple of values,

$$(F(\varepsilon; \bar{x}_m), RF(\varepsilon; \bar{x}_m), TF(\varepsilon; \bar{x}_m)) \quad .$$

For any triple $\bar{a}_3$ of numbers, we will use $(\bar{a}_3)_i$ to denote $a_i (1 \leq i \leq 3)$. Thus we have

$$(\mathfrak{F}(\varepsilon; \bar{x}_m))_1 = F(\varepsilon; \bar{x}_m) \quad ,$$

etc. ε-Functions are finitely computable in the following sense: there is a Turing machine which, when given an object[1] for computing the $X_j(\varepsilon)$ and $RX_j(\varepsilon)$ for any given ε, can output the triple of values, $\mathfrak{F}(\varepsilon; \bar{x}_m)$ for any given ε.

Dealing with the instabilities alluded to in the introduction was not our main reason for introducing the truncation-error bounds (TF). We had to introduce them because any definition of "$\mathfrak{F}$ ε-converges at x," which is based only on the values of $F(\varepsilon; x, y)$ and $RF(\varepsilon; x, y)$ (ε and x are fixed here), cannot have much to do with "f converges at x) (which is true when $\lim_{y \to x} f(x,y) \neq \omega$); remember that F and RF can only take on a finite number of different values for each fixed ε. The truncation-error bound, TF, gives us the needed local information about f.

Suppose γ_1, γ_2, and γ_3 determine the subroutines F, RF and TF respectively. Then we say $(\gamma_1, \gamma_2, \gamma_3)$ determines $\mathfrak{F} \equiv (F, RF, TF)$ and (γ_1, γ_2) partially determines $\mathfrak{F}$. When we say "given $\mathfrak{F}$" we mean "given $(\gamma_1, \gamma_2, \gamma_3)$ determining $\mathfrak{F}$."

1/We call the thing which computes (X_i, RX_i) an "object" rather than a Turing machine because there may be no such Turing machine; there are only a countable number of Turing machines, but there are an uncountable number of values of $\overline{x}_m$. See Shoenfield [S1, p. 248] for similar considerations.

2.6 An Example: e^x

Following is an example of how one might go about defining an ε-function corresponding to e^x over $\tilde{R}$. This example is formalized in section 5.7.

Let $[y]$ denote the greatest integer in y and let $\mathrm{sgn}(x)$ denote the sign function at x (which is ω if x is ω, and otherwise is -1 if $x < 0$, is 0 if $x = 0$, and is 1 if $x > 0$).

Define

$$f(x,k,y) = \begin{cases} \omega & \text{if } k=\omega, \text{ or } y \geq \infty \\ 0 & \text{if } y < 1 \\ \sum_{n=1}^{[y]} (-|x|)^{n-1}/(n-1)! & \text{otherwise,} \end{cases}$$

$$tf(x,k,y) = \begin{cases} |x|^{[y]}/[y]! & \text{if } k = \infty \text{ and } |x| + 1 < y < \omega \\ \omega & \text{otherwise,} \end{cases}$$

$$f_{exp}(x) = \begin{cases} \infty & \text{if } x = \infty \\ 0 & \text{if } x = -\infty \\ (\lim_{y \to \infty} f(x,\infty,y))^{-\mathrm{sgn}(x)} & \text{otherwise.} \end{cases}$$

$tf(x,\infty,y)$ bounds the remainder term, $\left| \sum_{n=[y]}^{\infty} (-|x|)^n/n! \right|$. The point here is that tf can be computed using only arithmetic, $[\cdot]$, $|\cdot|$ and numerical comparisons. It should not surprise the reader that there is a subroutine, F, such that $F(\varepsilon; \bar{x}_3) \to f(\bar{x}_3)$ as $\varepsilon \to 0$ for most $\bar{x}_3$. We can use the methods of interval analysis, or an error analysis in the style of Wilkinson [W2], to obtain a subroutine F.

We can use interval analysis to obtain a subroutine, TF, which satisfies

$$TF(\varepsilon;\ \bar{x}_3) \geq tf(\bar{x}_3) \quad .$$

The result is $\mathcal{F} \equiv (F,\ RF,\ TF)$, an ε-function corresponding to f over most of $\widetilde{R}^{(3)}$. For each ε, let y_ε be some member of $R(\varepsilon)$ with $y_\varepsilon \neq \infty$ and $\lim_{\varepsilon \to 0} y_\varepsilon = \infty$. Let $RG(\varepsilon;\ x)$ be either of the two smallest numbers in $R(\varepsilon)$ which are $\geq RF(\varepsilon;x,\infty,y_\varepsilon) + TF(\varepsilon;x,\infty,y_\varepsilon)$. We can define an ε-function corresponding to

$$g(x) = \begin{cases} \omega & \text{if } |x| \geq \infty \\ e^x & \text{if } x \leq 0 \\ e^{-x} & \text{if } x \geq 0 \end{cases} ,$$

by

$$\mathcal{G}(\varepsilon;x) \equiv (F(\varepsilon;x,\infty,y_\varepsilon),\ RG(\varepsilon;x),\ \omega) \quad .$$

Thus $\mathcal{G}(\varepsilon_{i+1};x)$ is computed at a higher precision of computation (see sec. 1.2) than is $\mathcal{G}(\varepsilon_i;x)$ and we will have $\mathcal{G} \approx g(P)$ for some set $P \subset \widetilde{R}$. It is easy to get $\mathcal{F}_{exp}$ from $\mathcal{G}$.

This method of approximating e^x by an alternating series has the numerical disadvantage of involving cancellation, but it affords the use of the simple and rapidly convergent truncation-error bound, $|x|^n/n!$ (when $n \geq |x|$). A method based on $\Sigma|x|^n/n!$ would involve no cancellation (so lower precision arithmetic could be used) but we would have to use a more complicated and more slowly converging truncation-error bound of the form $|x|^n\, 2.74^{[|x|+1]}/n!$ (valid for any $n \geq 0$). We use the former method here because it simplifies the formalization in section 5.7.

2.7 Operators

Operators and ε-operators will be our principal vehicles for defining notions and ε-notions in chapters 4-6. Let S_f be the set of all ideal functions of 0, 1, 2... variables. An operator of $n \geq 0$ ideal functions over $S \subset S_f^{(n)}$ is a mapping, $\phi : S \to S_f$. Let $\overline{f}_n$ denote the list of ideal functions, $f_1, \ldots, f_n$, and likewise for $\overline{\mathcal{F}}_n$ and $\overline{P}_n$. We say $\overline{\mathcal{F}}_n$ corresponds to $\overline{f}_n$ over $\overline{P}_n$ precisely when $\mathcal{F}_i \approx f_i(P_i)$ $(1 \leq i \leq n)$, and we write

$$\overline{\mathcal{F}}_n \approx \overline{f}_n(\overline{P}_n) \ .$$

Let S be the set of all $\overline{\mathcal{F}}_n$ such that there is a $\overline{f}_n \in S$ and a $\overline{P}_n$ with $\overline{\mathcal{F}}_n \approx \overline{f}_n(\overline{P}_n)$. Let S_w be the set of all weak ε-functions. A weak ε-operator, (Φ, Q), corresponding to ϕ over $S' \subset S$ is a mapping, $\Phi : S' \to S_w$, together with a set function, Q, which depends on $\overline{\mathcal{F}}_n$, $\overline{f}_n$, $\overline{P}_n$, such that

(1) if $\overline{\mathcal{F}}_n \approx \overline{f}_n(\overline{P}_n)$, $\overline{\mathcal{F}}_n \in S'$ and $\overline{f}_n \in S$, then

$\Phi(\overline{\mathcal{F}}_n) \sim \phi(\overline{f}_n)(Q(\overline{\mathcal{F}}_n, \overline{f}_n, \overline{P}_n))$, and

(2) there exist recursive operators Ψ_1 and Ψ_2 such that,

if $\overline{\mathcal{F}}_n \approx \overline{f}_n(\overline{P}_n)$, $\overline{\mathcal{F}}_n \in S'$, $\overline{f}_n \in S$ and $(Y_{3i-2}, Y_{3i-1}, Y_{3i})$ determines $\mathcal{F}_i$ $(1 \leq i \leq n)$, then $(\Psi_1(\overline{Y_{3n}}), \Psi_2(\overline{Y_{3n}}))$ partially determines $\Phi(\overline{\mathcal{F}}_n)$.

When these conditons hold, we write

$$(\Phi, Q) \sim \phi(S') \ .$$

Condition (1) requires that Φ gives ε-approximations to ϕ and that $\Phi \to \phi$ as $\varepsilon \to 0$ in the sense that we at least have

$$\lim_{\varepsilon \to 0} (\Phi(\overline{\mathfrak{F}}_n)(\varepsilon;\ \overline{x}_p))_1 = \phi(\overline{f}_n)(\overline{x}_p)$$

for $\overline{x}_p \in Q(\overline{\mathfrak{F}}_n, \overline{f}_n, \overline{P}_n)$ at which $\phi(\overline{f}_n)(\overline{x}_p) \neq \omega$. Conditon (2) requires that Φ be finitely computable from its arguments; it requires Φ to constructively map the determiners of $\overline{\mathfrak{F}}_n$ into a partial determiner of $\Phi(\overline{\mathfrak{F}}_n)$. We have left truncation-error bounds out of (2) because, for the ε-operators which we will present later, we do not believe there is an automatic way to define a good truncation-error bound for $\Phi(\overline{\mathfrak{F}}_n)$ from the determiners of the $\overline{\mathfrak{F}}_n$ (see def. 4.1-1). In general, such bounds depend on certain analytic properties of $\phi(\overline{f}_n)$, properties which cannot be effectively recovered from the numeric information given by the determiners of $\overline{\mathfrak{F}}_n$. We avoid this problem in most cases by assuming such bounds to be given. If $\phi(\overline{f}_n)$ is a function of 0 or 1 variables, then the TF part of $\mathfrak{F} \equiv \phi(\overline{\mathfrak{F}}_n)$ is, by definition, identically ω, and this problem does not arise. We will have more to say about this in chapter 5.

If condition (1) above holds with "$\Phi(\overline{\mathfrak{F}}_n) \approx \phi(\overline{f}_n)$", we say <u>$(\Phi, Q)$ is an ε-operator corresponding to ϕ over S'</u> (and not just a weak ε-operator), and we write

$$(\Phi, Q) \approx \phi(S') \quad .$$

The "goodness" of (Φ, Q) depends on how nontrivial the relation between $\overline{\mathfrak{F}}_n, \overline{f}_n, \overline{P}_n$ and $Q(\overline{\mathfrak{F}}_n, \overline{f}_n, \overline{P}_n)$ is, how large S' is, and especially on how efficient Φ is, in terms of the number of evaluations

of the $\bar{\mathcal{F}}_n$ required to evaluate $\Phi(\bar{\mathcal{F}}_n)(\varepsilon; \bar{x}_m)$, and the accuracy achieved (i.e., the size of $(\Phi(\bar{\mathcal{F}}_n)(\varepsilon; \bar{x}_m))_2$). For example, let $\Phi_{bad}(\bar{\mathcal{F}}_n) \equiv (\omega, \omega, \omega)$ and $Q_{bad}(\bar{\mathcal{F}}_n, \bar{f}_n, \bar{P}_n) \equiv \{\}$, the null set. Then $(\Phi_{bad}, Q_{bad}) \approx \emptyset(S)$ for any operator $\emptyset$ over S and its corresponding S. Of course this is not a good ε-operator in any sense. The formalization of a measure of the goodness of (Φ, Q) is a worthwhile and as yet unsolved problem. When this is satisfactorily solved, the rules for ε-izing a notion will be complete.

For simplicity, when we present particular ε-operators, we will give a constructive analytic definition of Φ, rather than giving Ψ_1 and Ψ_2. It is a simple but tedious task to construct particular Ψ_1 and Ψ_2 from such a definition.

2.8 Rounding Subroutines

Roundup and rounddown ε-arithmetic subroutines, $A_{1,*}$ and $A_{2,*}$, for $*$ being $+$, $-$, $\times$, $\div$, are multiple-precision subroutines which essentially give upper and lower bounds on the ideal arithmetic operations, $+$, $-$, $\times$, $\div$. We make this precise as follows. For any $a \in R$, define an **ε-neighborhood of a**, $N_\varepsilon(a)$, by

$$N_\varepsilon(a) = \{a_1, a_2, a_3, a_4\} \quad , \tag{2.8-1}$$

where $[a_1$ and a_2 are the largest members of $R(\varepsilon)$ satisfying $a_1 \le a_2 \le a$, or, if this is not possible, then $a_1 = a_2 \le a]$, and $[a_3$ and a_4 are the smallest members of $R(\varepsilon) - \{\omega\}$ satisfying $a \le a_3 \le a_4$ or, if this is not possible, then $a \le a_3 = a_4]$. The fact that $-\infty$ and ∞ are in $R(\varepsilon)$ means that $N_\varepsilon(a)$ is well defined and nonempty for all $a \in R$. For $a \in \{-\infty, \infty, \omega\}$, define

$$N_\varepsilon(a) = \{a\} \quad .$$

For each ε, we require the $A_{n,*}$ to satisfy

(1) $RX_i(\varepsilon) \ne \omega\ (i=1,2) \Rightarrow A_{n,*}(\varepsilon; \overline{x}_2) = A_{n,*}(\varepsilon; \overline{x}_2(\varepsilon))$, and

(2) for any $a, b \in R(\varepsilon)$ we have $c_{n,*} = A_{n,*}(\varepsilon; a,b)$ for n

$N_\varepsilon(a*b)$ and $c_{1,*} \le a * b \le c_{2,*}$.

Condition (1) states that the $A_{n,*}$ do not use the inputted roundoff-error bounds. Condition (2) requires $A_{1,*}$ and $A_{2,*}$ to give exact upper and lower approximations to $a * b$. It is easy to verify that such subroutines exist. For example, subroutines for $A_{1,+}, \ldots, A_{1,-}$

and $A_{2,-}$ can be defined as follows. (The $\hat{I}(\cdot)$ and $\check{I}(\cdot)$ defined here will be used later.)

EXAMPLE 2.8-1: Define ε-precision roundup and rounddown converted values, $\hat{I}(\varepsilon,\beta)$ and $\check{I}(\varepsilon,\beta)$, for the number $< \beta >$ as follows. Let n_0 be the least integer such that the interval $[(\beta(n_0) - 1)/n_0, (\beta(n_0) + 1)/n_0]$ overlaps at most one of the intervals $[(\alpha_R(i, l, n_0) - 1)/n_0, (\alpha_R(i, l, n_0) + 1)/n_0] (1 \le l \le l_R(i))$. If $\beta(n_0)$ is $-\infty$, ∞ or ω then let $\hat{l}$ and $\check{l}$ be 1, 2 or 3, respectively. Otherwise let the intervals about $\alpha_R(i, \hat{l}, n)$ and $\alpha_R(i, \check{l}, n_0)$ be the first ones lying completely to the right and left, respectively, of the one about $\beta(n_0)$. Define

(2.8-2) $\overset{*}{I}(\varepsilon_i, \beta) = < \alpha_R(i, \overset{*}{l}, \cdot) >$ for $*$ being $\frown$ and $\smile$.

Then $\overset{*}{I}(\varepsilon, \beta) \in N_\varepsilon(< \beta >)$ and $\check{I}(\varepsilon, \beta) \le (< \beta >) \le \hat{I}(\varepsilon, \beta)$.

Define $\beta^{\pm}_{i,j,k}$ by

$$\beta^{\pm}_{i,j,k}(n) = \alpha_R(i, j, 2n) \pm \alpha_R(i, k, 2n) \quad \text{for } n \ge 1 .$$

Then $\beta^{\pm}_{i,j,k}$ computes $< \alpha_R(i, j, \cdot) > \pm < \alpha_R(i, k, \cdot) >$. (See Bishop [B1, pp. 16, 21] for similar definitions of +, -, × and ÷.) Let $a = < \alpha_R(i, j, \cdot) >$ and $b = < \alpha_R(i, k, \cdot) >$. We can define the $A_{n,\pm}$ by

(2.8-3) $A_{1,\pm}(\varepsilon_i; a, b) = \hat{I}(\varepsilon_i, \beta^{\pm}_{i,j,k})$, $A_{2,\pm}(\varepsilon_i; a, b) = \check{I}(\varepsilon_i, \beta^{\pm}_{i,j,k})$.

For the rest of the paper, we assume particular $A_{n,*}$ to be given. For $\varepsilon, b \in R(\varepsilon)$ we will use $a *_\varepsilon b$ and $a *_\varepsilon b$ to denote $A_{1,*}(\varepsilon; a,b)$

and $A_{2,*}(\varepsilon; a,b)$, respectively, omitting the subscript ε whenever no confusion can arise. For subroutines F and G, we will abbreviate $F(\varepsilon; \bar{x}_m) \hat{+} G(\varepsilon; \bar{x}_m)$ by $(F \hat{+} G)(\varepsilon; \bar{x}_m)$, etc. In general, we will factor out arguments as much as possible, calling the resulting form argument factored form.

In order to prove that the $A_{n,*}$ converge to ideal arithmetic as $\varepsilon \to 0$, we must first state explicitly the special rules for arithmetic involving $\pm\infty$ and ω. Let $x,y,z \in \tilde{R}$ satisfy $-\infty < y < \infty$ and $0 < z \leq \infty$. The special rules are

General: $x * \omega = \omega * x = \omega$.

Addition: $\infty + (-\infty) = \omega$; $\infty + \infty = \infty + y = \infty$.

Multiplication: $\infty \times 0 = \omega$; $\infty \times z = \infty$.

Division: $x \div 0 = \infty \div \infty = \omega$; $y \div \infty = 0$.

These, combined with the usual definition of real arithmetic (see, for example, Bishop [B1, pp. 16, 21]) and the usual associative, commutative and distributive laws, completely define the arithmetic of $\tilde{R}$. For example, $\infty - \infty = \infty + (-\infty) = \omega$ and $\infty \times (-10) = (-1) \times (\infty \times 10) = -\infty$.

THEOREM 2.8-1: Suppose $\lim_{\varepsilon \to 0} X_i(\varepsilon) = x_i (i = 1,2)$ and $x_1 * x_2 \neq \omega$. Then $X_1(\varepsilon) \hat{*} X_2(\varepsilon)$ and $X_1(\varepsilon) \check{*} X_2(\varepsilon)$ approach $x_1 * x_2$ as $\varepsilon \to 0$.

Proof: If $x_1, x_2 \in R$ then $x_1 * x_2 \neq \omega$ implies that $x_1 * x_2$ is finite. We have that $X_1(\varepsilon) \hat{*} X_2(\varepsilon)$ and $X_1(\varepsilon) \check{*} X_2(\varepsilon)$ are in $N_\varepsilon(X_1(\varepsilon) * X_2(\varepsilon))$ and property II of $\tilde{R}$ together with the continuity of arithmetic give us convergence. If x_1 or x_2 is infinite then

$x_1 * x_2 \neq \omega$ and the above special rules for arithmetic involving $\pm\infty$ yield that $|x_1 * x_2|$ is one of

$$\infty + \infty, \infty - y, \infty \times z, y \div \infty \quad .$$

Convergence is clear in these cases. This completes the proof.

For $A,B \in \mathfrak{R}(\varepsilon)$, let

(2.8-4)
$$|A \ominus_\varepsilon B| = \begin{cases} 0 & \text{if } A = B \\ A \ominus_\varepsilon B & \text{if } A \geq B \\ B \ominus_\varepsilon A & \text{otherwise,} \end{cases}$$

(2.8-5)
$$|A \oplus_\varepsilon B| = \begin{cases} 0 & \text{if } A = B \\ \max(0, A \oplus_\varepsilon B) & \text{if } A \geq B \\ \max(0, B \oplus_\varepsilon A) & \text{otherwise} \quad . \end{cases}$$

This simplifies inequalities, because $|A \ominus_\varepsilon B|$ and $|A \oplus_\varepsilon B|$ are effective upper and lower bounds on the distance, $|A - B|$, between A and B (see sec. 1.4).

REMARKS: In our model, we have assumed that it is possible to use arbitrary levels of precision (arbitrarily small ε). But in practice, we almost always use single- or double-precision, and there is a finite upper limit on how high the precision can be (precisions higher than double-precision being provided via software). However, our model does not preclude an emphasis on single- and double-precision computations. We feel it is conceptually correct to keep arbitrary precision in mind in the design and analysis of algorithms; doing so helps keep algorithms machine independent and is kind to the occasional user who requires high accuracy.

In our definition of roundoff-error bounds and truncation-error bounds (sec. 2.3 and 2.5), we have taken the stand that numerical analysis should concern itself with rigorous approximation rather than just estimation. However, it should be possible to form an "ε-calculus of estimation" by defining these bounds to be statistical quantities. In fact, it should be possible to form an "ε-calculus of stable ε-functions" which involves no error bounds, as we indicate in the remarks after chapters 4 and 5. (Such an ε-calculus would not have very interesting ε-notions of ε-comparison, ε-convergence and ε-continuity.) The last two ε-calculi should be interesting to explore. The last one will probably resemble current scientific computation more closely than the ε-calculus developed in this thesis. In the "ε-calculus of stable ε-functions", a "poor real input" x would be a mapping $X: \mathcal{E} \to \mathcal{M}$ such that $X(\varepsilon) \in \mathcal{R}(\varepsilon)$. A "real input" x with value $c \in \tilde{R}$ would be a "poor real input" such that for every $\delta > 0$ there is an ε with $|X(\varepsilon) - c| < \delta$. A "subroutine" would essentially remain as before. And "an $\mathcal{F} \approx f(P)$" would be a "subroutine" F with $\lim_{\varepsilon \to 0} F(\varepsilon; \bar{x}_m) = f(\bar{x}_m)$ for all $\bar{x}_m \in P$ at which $f(\bar{x}_m) \neq \omega$.

2.A Appendix: Maximum Relative Error

Here we prove statements made about the examples of section 2.1

THEOREM 2.A-1: Assume that, for each ε, $R(\varepsilon) - \{\omega\}$ is symmetric about 0. Let $\Theta(\varepsilon)$ and $\tau(\varepsilon)$ be the smallest and largest finite positive numbers in $R(\varepsilon)$ and define

$$(2.A\text{-}1) \qquad E(\varepsilon) = \sup_{\Theta(\varepsilon) \le |x| \le \tau(\varepsilon)} \min_{y \in R(\varepsilon)} \left|\frac{x-y}{x}\right| .$$

Letting $a < a'$ range over positive finite neighbors in $R(\varepsilon)$ yields

$$(2.A\text{-}2) \qquad E(\varepsilon) = \max_a (a'-a)/(a'+a) .$$

Proof: By symmetry it suffices to consider only x with $\Theta(\varepsilon) \le x \le \tau(\varepsilon)$. For $a \in R(\varepsilon)$ with $\Theta(\varepsilon) \le a < \tau(\varepsilon)$, let a' denote the successor of a in $R(\varepsilon)$. We have

$$E(\varepsilon) = \max_a \left(\sup_{a \le x \le a'} \min_{y \in R(\varepsilon)} \left|\frac{x-y}{x}\right| \right)$$

$$= \max_a \sup_{a \le x \le a'} \min \left(\frac{x-a}{x}, \frac{a'-x}{x} \right) .$$

The facts that $(x-a)/x$ is monotone increasing and $(a'-x)/x$ is monotone decreasing for $x \neq 0$, and that these functions intersect at $(a + a')/2$ yield (2.A-2). This completes the proof.

This theorem gives immediate results for examples 2.2-1 and 2.2-2. For $a \in R^{\odot}(\varepsilon_i^{\odot})$ with $a = (b\beta^{-i})\beta^e$ and $\beta^{i-1} \le b < \beta^i$ we have

$a'-a = \beta^{e-i}$ and $\beta^{e-i}/(a'+a)$ is smallest for such a when $a = \beta^{-1}\beta^{e}$, yielding $a'+a = \beta^{-1}\beta^{e} + (\beta^{-1} + \beta^{-i})\beta^{e} = 2\beta^{e-1} + \beta^{e-i}$ and

$$E^{\odot}(\varepsilon_i^{\odot}) = \max_{|e| < \beta^i} \beta^{e-i}/(2\beta^{e-1} + \beta^{e-i}) = \varepsilon_i^{\odot} .$$

For $a \in \mathfrak{R}^{*}(\varepsilon_i^{*})$ with $a = \beta_i^{e}$ we have $a' = \beta_i^{e+1}$ and

$$E^{*}(\varepsilon_i^{*}) = \max_{|e| < 10^{2i}} (\beta_i^{e+1} - \beta_i^{e})/(\beta_i^{e+1} + \beta_i^{e}) = \varepsilon_i^{*} .$$

To prove this for $\mathfrak{R}^{\#}$ we need more machinery. First we note that it suffices to take the maximum in (2.A-2) over $\Theta(\varepsilon) \le a < 1$ (rather than $\Theta(\varepsilon) \le a < \tau(\varepsilon)$) for $\mathfrak{R}^{\#}$: suppose $a = p/q \ge 1$ and $a' = p'/q'$; then the successor of $b = q/p$ is $b' = q'/p'$ and we have

$$\frac{b'-b}{b'+b} = \frac{1/a' - 1/a}{1/a' + 1/a} = \frac{a'-a}{a'+a} .$$

For $n = 1, 2, \ldots$ define the Farey series of order n, F_n, to be the sequence of rationals, p/q, with $0 \le p \le q \le n$ and G.C.D.$(p,q) = 1$, written in increasing order. We shall require the following two well known lemmas (see Niven and Zuckerman [NZ, pp. 128-130]).

LEMMA 2.A-1: If p/q and p'/q' are consecutive fractions in F_n then $p'q - pq' = 1$.

LEMMA 2.A-2: If p/q and p'/q' are consecutive in F_n, then among all rationals with value between these two, $(p+p')/(q+q')$, is the unique one with smallest denominator.

We will also need the following two new results.

LEMMA 2.A-3: For $n \geq 2$, if $p/q < p'/q'$ are consecutive in F_n and $p > 0$, then $pq' \geq [\frac{1}{2}(n+1)]$.

Proof: F_2 is $< 0, \frac{1}{2}, \frac{1}{1} >$, so the theorem is true for F_2 by inspection. Suppose it is true for F_{n-1} . Any consecutive fractions in F_n will be either $\frac{p}{q}, \frac{p'}{q'}$ or $\frac{p}{q}, \frac{p+p'}{q+q'}$ or $\frac{p+p'}{q+q'}, \frac{p'}{q'}$ where $\frac{p}{q}, \frac{p'}{q'}$ are consecutive fractions in F_{n-1} (by lemma 2.A-2). We have $pq' \geq [n/2]$. In the last two cases this implies

$$p(q+q') = pq + pq' \geq 1 + [n/2] \geq [\tfrac{1}{2}(n+1)] \quad ,$$

$$(p+p')q' = pq' + p'q' \geq [n/2] + 1 \geq [\tfrac{1}{2}(n+1)] \quad ,$$

and the induction step follows. In the first case, if n is even then $[n/2] = [\frac{1}{2}(n+1)]$ so $pq' \geq [\frac{1}{2}(n+1)]$. Or, if $pq' > [n/2]$ then $pq' \geq [\frac{1}{2}(n+1)]$. Suppose n is odd and $pq' = [n/2]$. Then we have

$$q' \leq pq' = (n-1)/2 \quad ,$$

$$q \leq p'q = pq' + 1 = (n-1)/2 + 1 \quad ,$$

$$q + q' \leq n \quad .$$

It follows that $(p+p')/(q+q')$ is in F_n and so p/q, p'/q' could not have been consecutive in F_n . This completes the proof.

LEMMA 2.A-4: Let n be ≥ 2 and let $a < a'$ run through all consecutive fractions in F_n with $a \geq 0$. Then we have

$$\max_{a} \frac{a'-a}{a'+a} = \frac{1}{2[\frac{1}{2}(n+1)] + 1} \quad .$$

Proof: Let $a = p/q$ and $a' = p'/q'$. We have

$$\frac{p'/q' - p/q}{p'/q' + p/q} = \frac{p'q - pq'}{p'q + pq'} = \frac{1}{2pq' + 1} \leq \frac{1}{2[\frac{1}{2}(n+1)] + 1} \quad .$$

Further, the fractions $1/([\frac{1}{2}(n+1)] + 1)$, $1/[\frac{1}{2}(n+1)]$ are consecutive in F_n because $([\frac{1}{2}(n+1)] + 1) + [\frac{1}{2}(n+1)] > n$, and for these we have

$$\frac{1}{2pq' + 1} = \frac{1}{2[\frac{1}{2}(n+1)] + 1} \quad .$$

This completes the proof.

Taking $n = 10^i - 1$ yields

$$E^{\#}(\varepsilon_i^{\#}) = \max_{0 \neq a \in F_n} (a'-a)/(a'+a) = \varepsilon_1^{\#} \quad ,$$

completing our task.

BLANK PAGE

Chapter 3. Numerical Instability

3.1 A Definition

To simplify notation, we restrict this discussion to ε-functions and functions of two variables. Let x be a real input and suppose that $\lim_{y \to x} f(x, y)$ exists. Let F be a subroutine satisfying

$$\lim_{\varepsilon \to 0} F(\varepsilon; x, y) = f(x, y) , \tag{3.1-1}$$

for $y \in \mathcal{M} - \{x\}$ at which $f(x, y) \neq \omega$. This means that, with $\mathcal{F} \equiv (F, \omega, \omega)$, we have $\mathcal{F} \sim f(\{x\} \times (\mathcal{M} - \{x\}))$. We do not require (3.1-1) for $y = x$ because (1) we do not need to, and (2) this could cause problems when $f(x, y)$ is discontinuous at $y = x$ (x fixed).

We are interested here in a computation of $\lim_{y \to x} f(x, y)$ which proceeds at precision ε by selecting a $y_\varepsilon \in \mathcal{R}(\varepsilon)$ and then using $F(\varepsilon; x, y_\varepsilon)$ to approximate $\lim_{y \to x} f(x, y)$. We call the rule used to select these y_ε's (as a function of F, x and possibly other things) a stopping criterion because it tells us where to stop at precision ε. We say that this stopping criterion works at x precisely when

$$\lim_{\varepsilon \to 0} F(\varepsilon; x, y_\varepsilon) = \lim_{y \to x} f(x, y) \tag{3.1-2}$$

This framework is quite general. $f(\infty, n)$ might be the n^{th} iterate of an iterative procedure for evaluating $\lim_{n \to \infty} f(\infty, n)$, as in Newton's or Bernoulli's method for finding zeros of a polynomial, or as in numerical integration methods for ordinary differential equations. The assumption that $\lim_{n \to \infty} f(\infty, n)$ exists means that the discrete method converges in exact arithmetic, with exact starting values. (This is weaker than "convergence"

as defined in Ralston [R1, p. 171].

We are now ready to discuss numerical stability. As used in numerical analysis, stability deals with the way local rounding errors of some iterative procedure propogate and effect the total accumulated error. (See Henrici [H1, pp. 11, 302, 309] and Ralston [R1, p. 175 (under 1)]. Let us consider an example.

EXAMPLE 3.1-1: Let $q_1, q_2, \cdots$ be defined by

$$q_n = Q_n(q_0, q_1, \cdots, q_{n-1}) .$$

At precision ε, let $\tilde{q}_0$ approximate q_0, $\tilde{Q}_i$ approximate the i^{th} recurrence relation, Q_i, and define $\tilde{q}_1, \tilde{q}_2, \cdots$, by

$$\tilde{q}_n = \tilde{Q}_n(\tilde{q}_0, \tilde{q}_1, \cdots, \tilde{q}_{n-1}) .$$

The n^{th} local rounding error of this iterative procedure is

$$\tilde{q}_n - Q_n(\tilde{q}_0, \tilde{q}_1, \cdots, \tilde{q}_{n-1}) ,$$

and the total accumulated error is $\tilde{q}_n - q_n$. Let $f(x, y) = q_n$ when $n = [y] \geq 0$ and $x \neq \omega$; otherwise let $f(x, y) = \omega$. Suppose we are interested in the finite limit, $\lim_{y \to \infty} f(\infty, y) = \lim_{n \to \infty} q_n$. Let $F(\varepsilon; x, y)$ be defined in terms of the $\tilde{q}_n$ so that (3.1-1) is satisfied for $x = \infty$. (This is easy, but tedious, to do; F will be effective so long as the $\tilde{Q}_n$ are.) Then $F(\varepsilon; \infty, y) - f(\infty, y)$ is the total accumulated error. If, as $y \to \infty$ through finite values in $R(\varepsilon)$, $|F(\varepsilon; \infty, y) - f(\infty, y)|$ becomes large, it would be said that "numerical instability has set in at precision ε."

If this happens for infinitely many values of ε, it would be said that F is unstable at ∞. Suppose this happens, and let y_ε be some value in $R(\varepsilon)$ where the total error has become large. Then we would have $\lim_{\varepsilon \to 0} F(\varepsilon; \infty, y_\varepsilon) \neq \lim_{y \to \infty} f(\infty, y)$, even though we may have $\lim_{\varepsilon \to 0} y_\varepsilon = \infty$ and $y_\varepsilon \neq \infty$. This is the tragedy of numerical instability; <u>when F is unstable at ∞, there will be seemingly quite reasonable stopping criteria that do not work at ∞.</u> On the other hand, if F is stable at ∞ (in the sense usual to numerical analysis), then any such reasonable stopping criteria should work at ∞.

This idea of stability generalizes easily to any F, x and f satisfying the assumptions at the beginning of this section, whether or not they involve iterative methods and local rounding errors. In this generalization, it is important that "reasonable" stopping criteria choose y_ε's that satisfy

$$|X(\varepsilon) - y_\varepsilon| > RX(\varepsilon)$$

so that y_ε is effectively distinct from x at precision ε. This is necessary so that F is not unstable just because $f(x, y)$ is discontinuous at $y = x$, e.g., when $f(x, y)$ involves divisions by $(x-y)$. Define <u>the set, $\rho(\varepsilon; x)$, of members of $R(\varepsilon)$ that are effectively distinct from x at precision ε</u> by

$$(3.1\text{-}3) \qquad \rho(\varepsilon; x) \equiv \{Y.\ Y \in R(\varepsilon) \text{ and } |X(\varepsilon) - Y| > RX(\varepsilon)\}$$

<u>DEFINITION 3.1-1. We say a stopping criterion is reasonable at x precisely when its y_ε's satisfy</u>

(3.1-4) $\qquad y_\varepsilon \in \rho(\varepsilon;\ x) \cup \{\omega\}\ ,\ \lim_{\varepsilon \to 0} y_\varepsilon = x\ .$

If $RX(\varepsilon) = \omega$ then (3.1-4) forces the choice $y_\varepsilon = \omega$. This cannot happen when ε is sufficiently small.

DEFINITION 3.1-2: Suppose $\lim_{y \to x} f(x, y)$ is finite and $\mathcal{F} \sim f(\{x\} \times (\mathcal{M} - \{x\}))$. We say F is stable at x precisely when any stopping criterion which is reasonable at x, works at x. Otherwise, we say F is unstable at x.

Following is an example of an F unstable at 0.

EXAMPLE 3.1-2: Let $\Theta(\varepsilon)$ be the smallest positive number in $R(\varepsilon)$. Suppose a certain form of ε-arithmetic is to be used and that in this ε-arithmetic we have $0 + 1 = 1$, $\Theta(\varepsilon) + 1 = 1$, $1 - 1 = 0$ and $0/(0 - \Theta(\varepsilon)) = 0$ (see section 5.3 for a detailed discussion of ε-arithmetic). Suppose a subroutine F, evaluated at $(\varepsilon;\ x,\ y)$, approximates $f(x, y) = (x + 1 - (y + 1))/(x - y)$ by replacing x and y by $X(\varepsilon)$ and $Y(\varepsilon)$ and by replacing arithmetic by this ε-arithmetic. (That F satisfies (3.1-1) follows from corollary 5.3-1 in chapter 5.) Define $y_\varepsilon = \Theta(\varepsilon)$. Then these y_ε's satisfy (3.1-4) for x being 0, but $F(\varepsilon;\ 0,\ y_\varepsilon) = 0$ for all ε, while $\lim_{y \to 0} f(0, y) = \frac{d}{dt}(t+1)\big|_{t=0} = 1$. Hence F is unstable at 0.

3.2 A Geometric Characterization

F is unstable at x if and only if there is a reasonable at x stopping criterion whose y_ε's satisfy

$$\lim_{\varepsilon \to 0} |F(\varepsilon; x, y_\varepsilon) - f(x, y_\varepsilon)| > 0 . \tag{3.2-1}$$

Interpreting this geometrically, we find that the graph of $F(\varepsilon; x, Y)$ versus finite $Y \in R(\varepsilon)$ acts like an <u>ε-wave</u>. This is pictured in figure 3.2-1 for $x = \infty$ and $\varepsilon = \varepsilon_1, \varepsilon_2$.

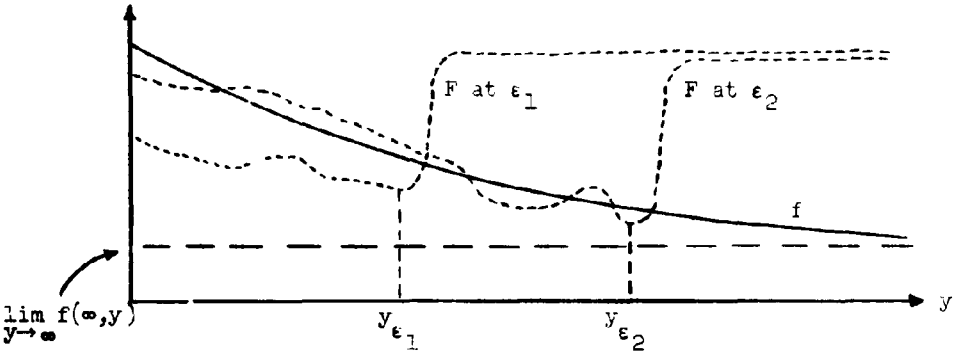

Figure 3.2-1 Instability

As $\varepsilon \to 0$, the ε-wave moves towards x. The <u>crest of the ε-wave</u> stays uniformly away from $\lim_{y \to x} f(x, y)$. (See example 3.1-2).

Two usual stopping criteria are

(1) choose y_ε to be the first value of Y for which $TF(\varepsilon; x, Y) \leq RF(\varepsilon; x, Y)$, as $Y \to x$ via some fixed approach, and

(2) choose y_ε to be the first value of Y (as $Y \to x$ via some fixed approach) such that $F(\varepsilon; x, Y)$ and the previous four values of F are equal, within some tolerance.

The trouble with such stopping criteria is that they can make $F(\varepsilon; x, y_\varepsilon)$ ride the crest of the ε-wave out to x, thereby destroying convergence.

3.3 A Stopping Criterion That Works

The question arises, is there any stopping criterion (effectively computable or not) which yields convergence even when F is unstable at x? The answer is given in the affirmative by the following definition and theorem.

DEFINITION 3.3-1: Stopping criterion S. C.$^{\#}$ selects y_ε to be that value of $Y \in \mathcal{R}(\varepsilon)$ closest to x at which $|F(\varepsilon; x, Y) - \lim_{y \to x} f(x, y)|$ assumes its minimum over all $Y \in \mathcal{R}(\varepsilon)$ (taking the smaller value for y_ε in case of a tie).

We call $(y_\varepsilon, F(\varepsilon; x, y_\varepsilon))$ the base of the ε-wave of F at x. Of course y_ε cannot be effectively computed from F and x.

THEOREM 3.3-1: If $\mathcal{F} \sim f(\{x\} \times (\mathcal{M} - \{x\}))$ and $\lim_{y \to x} f(x, y)$ is finite, then S. C.$^{\#}$ works at x.

Note that S. C.$^{\#}$ works whether F is stable or not. Thus there is some desirable behavior even in the presence of instability.

Proof: Let $\ell = \lim_{y \to x} f(x, y)$. For any real input y, we have

$$|F(\varepsilon; x, y) - \ell| \leq |F(\varepsilon; x, y) - f(x, y)| + |f(x, y) - \ell| . \tag{3.3-1}$$

Let an $\eta > 0$ be given. By choosing y sufficiently close to x, keeping $y \in \mathcal{M} - \{x\}$, the second term on the right side of (3.3-1) becomes $< \eta/2$. Then by making ε sufficiently small, this value or y is in $\mathcal{R}(\varepsilon)$ and the first term on the right side of (3.3-1) becomes $< \eta/2$ (the nesting property III of $\mathcal{R}(\varepsilon)$ allows this). For such y and ε, the left side of (3.3-1) is $< \eta$. The left side

bounds the distance of the $F(\varepsilon; x, y_\varepsilon)$ of S. C.$^{\#}$ from ℓ, and so S. C.$^{\#}$ works. This completes the proof.

Thus the height of the base of the ε-wave of F an x approaches $\lim_{y \to x} f(x, y)$ as $\varepsilon \to 0$.

3.4 An Algorithm for Overcoming Instability

We derive an effective analog to S. C.[#] as follows. Suppose $\mathcal{F} \approx f(\{x\} \times (\mathcal{M} - \{x\}))$ and $\lim_{y \to x} f(x, y)$ exists.

DEFINITION 3.4-1: S. C.[##] selects y_ε to be the smallest value of Y in $R(\varepsilon)$ for which $(RF \hat{+} TF)(\varepsilon; x, Y)$ assumes its minimum over all Y in $R(\varepsilon)$.

This is finitely computable because $R(\varepsilon)$ is a finite set. This stopping criterion keeps us close enough to the base of the ε-wave of F at x that we get convergence even in the unstable case, provided only that

there is a sequence, $y_1, y_2, \ldots$, with each $y_i \in \mathcal{M} - \{x\}$,

(3.4-1) such that

$$\lim_{i \to \infty} \lim_{\varepsilon \to 0} \sup TF(\varepsilon; x, y_i) = 0 .$$

THEOREM 3.4-1: Suppose that

(1) $\lim_{y \to x} f(x, y) \neq \omega$,

(2) $\mathcal{F} \approx f(\{x\} \times \mathcal{M} - \{x\}))$, and

(3) TF satisfies (3.4-1).

Then S. C.[##] works at x.

Proof: This proof in essentially the same as that of theorem 3.3-1. We will prove that for every $\eta > 0$ there is a $\delta > 0$ with

$$(RF \hat{+} TF)(\varepsilon; x, y_\varepsilon) < \eta \qquad \text{for all } \varepsilon \leq \delta ,$$

where y_ε is chosen by S. C.[##]

Let an $\eta > 0$ be given. By assumption, there is a Y in some $R(\delta_1) - \{x\}$ with

$$\lim_{\varepsilon \to 0} \sup TF(\varepsilon; x, Y) < \eta/4$$

For this Y there are $\delta_2, \delta_3 \in \mathcal{E}$ with

$$TF(\varepsilon; x, Y) < \eta/2 \qquad \text{for all } \varepsilon \leq \delta_2 ,$$
$$RF(\varepsilon; x, Y) < \eta/4 \qquad \text{for all } \varepsilon \leq \delta_3 .$$

There is a δ_4 such that $\mathcal{R}(\delta_4)$ contains an η_1 and an η_2 with $3\eta/4 \leq \eta_1 < \eta_2 < \eta$. Let $\delta = \min(\delta_1, \delta_2, \delta_3, \delta_4)$. We have

$$(RF \hat{+} TF)(\varepsilon; x, y_\varepsilon) \leq (RF \hat{+} TF)(\varepsilon; x, Y) < \eta \quad \text{for all } \varepsilon \leq \delta$$

The first inequality makes use of the nesting property III of $\mathcal{R}$. The second inequality uses the facts that

(1) $(RF \hat{+} TF)(\varepsilon; x, Y) \in N_\varepsilon((RF + TF)(\varepsilon; x, Y))$,

(2) $(RF + TF)(\varepsilon; x, Y) < 3\eta/4$, and

(3) $\eta_1, \eta_2 \in \mathcal{R}(\varepsilon)$ (by property III of $\mathcal{R}$).

This completes the proof.

Thus S. C.## is a (totally inefficient) algorithm for overcoming instability. It should be possible to find a more efficient algorithm for which theorem 3.4-1 holds because

(1) we do not need to find the exact minimum of $(RF \hat{+} TF)(\varepsilon; x, Y)$ over $Y \in \mathcal{R}(\varepsilon)$; we only need to stay "sufficiently close" to it, and

(2) in particular cases, it should be possible to localize the search for y_ε.

On the other hand, it should be possible to show that any stopping criterion, for which theorem 3.4-1 holds, must require so many

evaluations of EF ? Tr [illegible] choose y_g that it cannot be very efficient.

In chapter 4, we present an efficient algorithm that almost satisfies this theorem (it has a somewhat more stringent third hypothesis, regarding TF).

3.5 Applications

We have proven that numerical instability is not an essential limitation of finite computations. Convergent roundoff-error and truncation-error bounds can be combined with an unstable subroutine to form a convergent algorithm for computing the associated limit. Here, we consider applications of this result to the initial-value problem for ordinary differential equations. Instabilities can generally be classified as

(I) those due to the particular method of solution used, and

(II) those due to the problem being solved.

We will give an example of each and we will show that the instabilities in both of these examples can be overcome by S. C.[##]. Of course, the best way to overcome instabilities of type I is to find a stable method of solution.

EXAMPLE 3.5-1: Consider solving the initial-value problem $y' = -y$, $y(0) = 1$, by the corrector formula,

$$y_{n+1} = y_{n-1} - \frac{h}{3}(y_{n+1} + 4y_n + y_{n-1}) , \tag{3.5-1}$$

from Milne's method (see [R1, p. 182]). This is a well-known unstable formula. Since y_n depends on h, let us write $y_n(h)$. Taking $y_0(h) = 1$ and $y_1(h) = e^{-h}$, we find the solution of the above difference equation to be

$$y_n(h) = A(h)\, r_+(h)^n + B(h)\, r_-(h)^n ,$$

where

$$r_{\pm}(h) = \frac{-2h + \sqrt{3(h^2 + 3)}}{h + 3}$$

$$A(h) = \frac{e^{-h}(h + 3) + 2h}{2\sqrt{3(h^2 + 3)}} + \frac{1}{2} = 1 - B(h) .$$

For fixed $h > 0$, we have $|r_+(h)| < 1 < |r_-(h)|$ and $B(h) \neq 0$, so that $|y_n(h)| \to \infty$ as $n \to \infty$, whereas $y(nh) \to 0$. However, for any finite x we have

(3.5-2) $$\lim_{n \to \infty} y_n(\tfrac{x}{n}) = e^{-x} .$$

Let $f_x(\infty, y) = y_n(\frac{x}{n})$ when $n = [y] \geq 1$, and let it be ω otherwise. Let $F_x(\varepsilon; \infty, y)$ approximate $f_x(\infty, y)$ by evaluating (3.5-1) in some form of ε-arithmetic (see sec. 5.3), where the approximations used for the initial values converge to the correct values, $y_0(\frac{x}{n}) = 1$ and $y_1(\frac{x}{n}) = e^{-x/n}$, as $\varepsilon \to 0$. It follows from corollary 5.3-1 that F_x satisfies (3.1-1). RF_x can be defined as in sections 5.3 and 5.6, and TF_x can be defined so that

$$TF_x(\varepsilon; \infty, n) \geq |y_n(\tfrac{x}{n}) - e^{-x}| .$$

It follows that S. C.## works when it is applied to F_x.

Because of its extreme lack of efficiency, S. C.## could not be used in practice. But the ε-limit defined in section 4.1 could be applied to this F_x with reasonable efficiency.

EXAMPLE 3.5-2: Consider solving the system,

(3.5-3) $$y' = z \quad , \quad y(0) = 1,$$
$$z' = y \quad , \quad z(0) = -1 ,$$

by the Newton-Cotes closed formula,

$$(3.5\text{-}4) \qquad \begin{pmatrix} y_{n+1} \\ z_{n+1} \end{pmatrix} \equiv \begin{pmatrix} y_n \\ z_n \end{pmatrix} + \frac{h}{2} \begin{pmatrix} z_{n+1} + z_n \\ y_{n+1} + h_n \end{pmatrix} .$$

Again, we will write $y_n(h)$ and $z_n(h)$. The general solution (3.5-3) is $y = Ae^{x} + Be^{-x}$, and the initial values give $y = e^{-x}$. The general solution to (3.5-4) for y_n is

$$y_n(h) = A(h)\left(\frac{2-h}{2+h}\right)^n + B(h)\left(\frac{2+h}{2-h}\right)^n .$$

Taking $y_0 = -z_0 = 1$ yields $A(h) = 1$ and $B(h) = 0$. But rounding errors in computing (3.5-4) in ε-arithmetic will build up so that $|\tilde{y}_n(h)|$ becomes large as n does, even though $B(h)$ should have been zero. We computed $\tilde{y}_n(.05)$ and $\tilde{y}_n(\frac{5}{n})$ for $n = 20, 30, \ldots, 1000$, on an IBM360/65 computer in short and long-precision; $\tilde{y}_n(05)$ is graphed in figure 3.5-1 and all the data is given in tables 3.5-1 and 3.5-2.

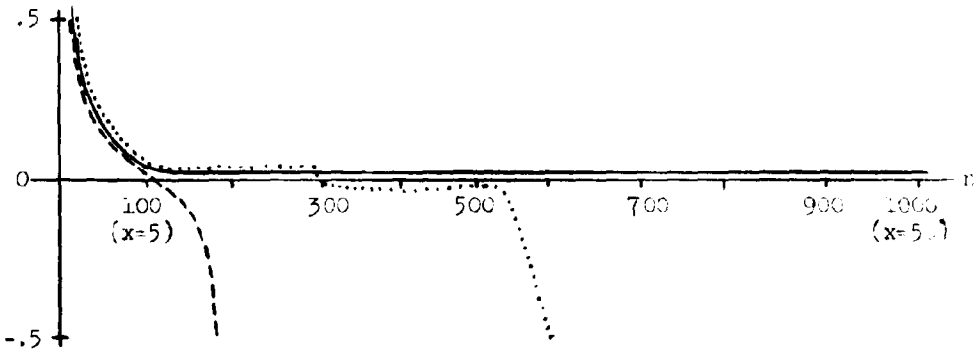

—— e^{-x}

--- $\tilde{y}_n(.05)$ short-precision

····· $\tilde{y}_n(.05)$ long-precision

Figure 3.5-1. $\tilde{y}_n(.05)$

We again have (3.5-2). The rest follows as in the last example, except the initial values for F_x are to be $y_0(h) = -z_0(h) = 1$.

It is possible to construct examples where $\lim_{n\to\infty} y_n(\frac{x}{n}) = y(x)$ does not hold, even when the initial values are assumed exact. The methods of this chapter cannot be used to overcome such instabilities.

TABLE 3.5-1

Data from example 3.5-2

n	double-precision $\tilde{y}_n(.05)$	single-precision	$y_n(.05)$
20	3.678027788566281-01	3.677294*-01	3.678027788554713*-01
30	2.230604167155121-01	2.229016*-01	2.230604167157031-01
40	1.352788841343381-01	1.349900*-01	1.352788841347261-01
50	9.204224112028951-02	8.154595*-02	8.204224112057361-02
60	4.975594950423881-02	4.892212*-02	4.975594950538521-02
70	3.017536426595[illegible]1-02	2.878845*-02	3.017536426792651-02
80	1.830037648945941-02	1.600527*-02	1.830037649273491-02
90	1.109858282533901-02	7.308926*-03	1.109858283075691-02
100	6.730929319207751-03	4.787731*-04	6.730929328151961-03
110	4.082089591768351-03	-6.229091*-03	4.082089606523791-03
120	2.475654486847361-03	-1.452738*-02	2.475654511182441-03
130	1.501403860652331-03	-2.653506*-02	1.501403900821561-03
140	9.105525425312201-04	-4.531748*-02	9.105526878620571-04
150	5.522204177985061-04	-7.586941*-02	5.522205269084761-04
160	3.349035596642871-04	-1.253403*-01	3.349037797758451-04
170	2.031[illegible]7947683137*-04	-2.070141*-01	2.031082443386551-04
180	1.231780516975561-04	-3.415465*-01	1.231785408511721-04
190	7.470257011420811-05	-5.632893*-01	7.470377667666561-05
200	4.530407968549431-05	-9.284276*-01	4.530540962057621-05
210	2.747406372183671-05	-1.531482*+00	2.747625665269531-05
220	1.665983066358511-05	-2.525161*+00	1.666345555568961-05
230	1.009988132991051-05	-4.163611*+00	1.010584356944161-05
240	6.119034196151341-06	-6.865200*+00	6.128865258737951-06
250	3.700746969034021-06	-1.131960*+01	3.716957340175791-06
260	2.227484504496541-06	-1.866377*+01	2.254213673402181-06
270	1.324[illegible]336968.251*-06	-3.077290*+01	1.367[illegible]723862851-06
280	7.564334948564051-07	-5.073912*+01	8.291060532141211-07
290	3.829965623433301-07	-8.366090*+01	5.028258413553371-07
300	1.073618406887591-07	-1.379449*+02	3.049475103907951-07
310	-1.408569028263001-07	-2.274450*+02	1.849407417314871-07
320	-4.25[illegible]4941[illegible]1-07	-3.750098*+02	1.121605417081481-07
330	-8.177726887238091-07	-6.183253*+02	6.802171873265161-08
340	-1.419327315050551-06	-1.019509*+03	4.125295991833101-08
350	-2.383322137951701-06	-1.681017*+03	2.501857717247901-08
360	-3.955923400692161-06	-2.771725*+03	1.517295292622391-08
370	-6.538711418648861-06	-4.570008*+03	9.201902207078971-09
380	-1.079122834747001-05	-7.535[illegible]27*+03	5.580654249727251-09
390	-1.779939922733111-05	-1.242391*+04	3.384485202531331-09
400	-2.935283439262221-05	-2.048504*+04	2.052580140888201-09
410	-4.840182534369191-05	-3.377685*+04	1.244823062490441-09
420	-7.981073278878271-05	-5.569184*+04	7.549446796447841-10
430	-1.316001524483491-04	-9.182394*+04	4.578493815689091-10
440	-2.169950587471461-04	-1.514001*+05	2.774707510564421-10
450	-3.578[illegible]1941491[illegible]1-04	-2.496332*+05	1.683962749388721-10
460	-5.899773229810091-04	-4.116275*+05	1.021280738457901-10
470	-9.728096351519491-04	-6.786766*+05	6.193735344041341-11
480	-1.604059145847331-03	-1.118995*+06	3.756298935976501-11
490	-2.644922074061471-03	-1.844997*+06	2.278073071041441-11
500	-4.361193744415671-03	-3.042077*+06	1.381577186868671-11
510	-7.191142257642811-03	-5.[illegible]15910*+06	8.378816059676981-12
520	-1.185742470435681-02	-8.270530*+06	5.081479285353361-12
530	-1.955162561675321-02	-1.363651*+07	3.081751830278411-12
540	-3.223854030044721-02	-2.248355*+07	1.868982201855791-12
550	-5.315790621231251-02	-3.707096*+07	1.133476886921161-12
560	-8.765170399818361-02	-6.112374*+07	6.874168474762961-13
570	-1.445282886891621-01	-1.007839*+08	4.168959487801931-13
580	-2.383116959344631-01	-1.661778*+08	2.528338267347031-13
590	-3.929505077116141-01	-2.739960*+08	1.533354985554611-13
600	-6.479333752601[illegible]41-01	-4.517663*+08	9.299298406260021-14

TABLE 3.5-1 (con't)

Data from example 3.5-2

n	$\tilde{y}_n(.05)$ double-precision	$\tilde{y}_n(.05)$ single-precision	$y_n(.05)$	$e^{-.05n}$
610	-1.068372862575608*+00	-7.448829*+08	5.639721875126985-14	5.675685232632745-14
620	-1.761632626251515*+00	-1.228193*+09	3.426307795218165-14	3.442477108469995-14
630	-2.904743857291725+00	-2.025111*+09	2.074305377650695-14	2.087967911665945-14
640	-4.789612062548945+00	-3.339055*+09	1.257998711626525-14	1.266416554909425-14
650	-7.897558213998185+00	-5.505450*+09	7.629352820973475-15	7.681204685202135-15
660	-1.302222913441915+01	-9.077445*+09	4.626954219343975-15	4.658886145103425-15
670	-2.147226358252625+01	-1.496717*+10	2.806597168432345-15	2.825757287115635-15
680	-3.540546695948175+01	-2.467860*+10	1.701806619517505-15	1.713908431542025-15
690	-5.837982967194045+01	-4.069136*+10	1.032090336291375-15	1.039538011670235-15
700	-9.626209759016799*+01	-6.709189*+10	6.259292037352855-16	6.305116760147025-16
710	-1.587259C7002114*+02	-1.106208*+11	3.796056937191245-16	3.824246628097155-16
720	-2.617220506062765+02	-1.823934*+11	2.302185005014075-16	2.319522830243585-16
730	-4.315516796677545+02	-3.007370*+11	1.396200290197245-16	1.406861712446155-16
740	-7.115825808052575+02	-4.958709*+11	8.467500422864305-17	8.533047625744105-17
750	-1.173323597524435+03	-8.176138*+11	5.135263465750925-17	5.175559C05801895-17
760	-1.934685167461215+03	-1.348101*+12	3.114370185499985-17	3.139132792048045-17
770	-3.190088995986955+03	-2.222778*+12	1.888764172864545-17	1.903980283286465-17
780	-5.260115688833865+03	-3.664986*+12	1.145474003615355-17	1.154822417301585-17
790	-8.673369390261565+03	-6.043004*+12	6.946027113845795-18	7.004352026168675-18
800	-1.430145717[0165*+04	-9.964057*+12	4.213^8523476864*-18	4.248354255291605-18
810	-2.358157384568575+04	-1.642879*+13	2.555099096987525-18	2.576757109154995-18
820	-3.988349406839405+04	-2.708767*+13	1.549584456908095-18	1.562882189335005-18
830	-6.411472452434325+04	-4.466251*+13	9.397725481262885-19	9.479359653504795-19
840	-1.057183259714695+05	-7.364124*+13	5.699416693239655-19	5.749522264293595-19
850	-1.743182167454555+05	-1.214235*+14	3.456509546941025-19	3.487261531994465-19
860	-2.874321023350015+05	-2.002489*+14	2.096260562030325-19	2.115131037591095-19
870	-4.739448062009565+05	-3.301066*+14	1.271313816509665-19	1.282891823608795-19
880	-7.814843140348625+05	-5.442843*+14	7.710104599224865-20	7.781132241133835-20
890	-1.288584C82138C15+06	-8.974361*+14	4.675927545111875-20	4.719495271526145-20
900	-2.124737383616005+06	-1.479734*+15	2.835797896870825-20	2.862518580549415-20
910	-3.503464781161095+06	-2.439871*+15	1.719819144824795-20	1.736205283100305-20
920	-5.776838854290495+06	-4.022938*+15	1.043014346745115-20	1.053061735755395-20
930	-9.525389645C99C85+06	-6.633027*+15	6.325542605975335-21	6.387142293058445-21
940	-1.570634912614435+07	-1.093664*+16	3.836235751202705-21	3.873997628687205-21
950	-2.589809047856135+07	-1.803272*+16	2.326552148254265-21	2.349698337452825-21
960	-4.270318232766595+07	-2.973332*+16	1.410978169641825-21	1.425164082740945-21
970	-7.041299208528055+07	-4.902548*+16	8.557123452829705-22	8.644057113036125-22
980	-1.161334793168125+08	-8.083218*+16	5.189616917094195-22	5.242885663363495-22
990	-1.914422242924365+08	-1.332748*+17	3.147333784970715-22	3.179970900197765-22
1000	-3.156677599818405+08	-2.197452*+17	1.908755523275955-22	1.928749847963935-22

TABLE 3.5-2

Data from example 3.5-2

(Note: $y(5) = e^{-5} = 6.7379469990855'-03$)

n	single-precision $\tilde{y}_n(\frac{5}{n})$	double-precision $y_n(\frac{5}{n})$	
20	6.259691'-03	6.56312402779648'-03	6.56312402790868'-03
30	5.460650'-03	6.66008827465367'-03	6.66008827578716'-03
40	5.009264'-03	6.69412021149203'-03	6.69412021215067'-03
50	5.949687'-03	6.70988861449936'-03	6.70988861592713'-03
60	1.785394'-03	6.71845852303765'-03	6.71845852888167'-03
70	-1.430578'-03	6.72362739157639'-03	6.72362739775014'-03
80	2.511602'-03	6.72698278674873'-03	6.72698278853473'-03
90	-9.466864'-03	6.72928349316458'-03	6.72928350441466'-03
100	4.787731'-04	6.73092931920775'-03	6.73092932815196'-03
110	-1.192247'-02	6.73214710541409'-03	6.73214712415119'-03
120	-9.010211'-03	6.73307338681753'-03	6.73307339919532'-03
130	-1.851527'-02	6.73379425406474'-03	6.73379428368321'-03
140	-3.144952'-02	6.73436628389642'-03	6.73436629883609'-03
150	-2.433643'-02	6.73482774281800'-03	6.73482778138931'-03
160	-2.318940'-02	6.73520547117676'-03	6.73520547763337'-03
170	-2.624966'-02	6.73551848763167'-03	6.73551850753462'-03
180	-2.789574'-02	6.73578078464965'-03	6.73578083277416'-03
190	-3.652655'-02	6.73600280280934'-03	6.73600284084923'-03
200	-4.467228'-02	6.73619235639365'-03	6.73619238936501'-03
210	-6.243775'-02	6.73635547974731'-03	6.73635551095046'-03
220	-7.417661'-02	6.73649684680107'-03	6.73649689923037'-03
230	-7.697296'-02	6.73662016382665'-03	6.73662025025945'-03
240	-5.590025'-02	6.73672842720946'-03	6.73672850653685'-03
250	-7.729518'-02	6.73682396844148'-03	6.73682403412359'-03
260	-1.126830'-01	6.73690867314930'-03	6.73690875302489'-03
270	-1.092784'-01	6.73698417848589'-03	6.73698423442461'-03
280	-8.893842'-02	6.73705172709532'-03	6.73705177405624'-03
290	-1.420971'-01	6.73711236526989'-03	6.73711244818440'-03
300	-9.052908'-02	6.73716704114330'-03	6.73716715674576'-03
310	-1.554755'-01	6.73721657511228'-03	6.73721665691298'-03
320	-6.847292'-02	6.73726155046883'-03	6.73726158543395'-03
330	-1.843618'-01	6.73730241728272'-03	6.73730249947136'-03
340	-2.111134'-01	6.73733975117423'-03	6.73733985321659'-03
350	-8.065218'-02	6.73737395472379'-03	6.73737405122982'-03
360	-1.433600'-01	6.73740522053359'-03	6.73740543922362'-03
370	-1.938782'-01	6.73743410804914'-03	6.73743431683434'-03
380	-2.166646'-01	6.73746083027843'-03	6.73746094479951'-03
390	-1.720492'-01	6.73748531402137'-03	6.73748555186108'-03
400	-3.021251'-01	6.73750823211830'-03	6.73750833465946'-03
410	-3.281339'-01	6.73752922285242'-03	6.73752947178680'-03
420	-2.661970'-01	6.73754901743251'-03	6.73754911718278'-03

430	-3.216861'-01	6.737567173922 45'-03	6.737567407986 35'-03
440	-2.298895'-01	6.737584223544 65'-03	6.737584465928 76'-03
450	-2.621343'-01	6.737600178477 28'-03	6.737600399369 77'-03
460	-4.247934'-01	6.737614939862 16'-03	6.737615305020 86'-03
470	-3.659283'-01	6.737628893763 21'-03	6.737629269413 90'-03
480	-4.369898'-01	6.737642028399 05'-03	6.737642370163 92'-03
490	-2.354691'-01	6.737654264357 03'-03	6.737654677046 61'-03
500	-4.692361'-01	6.737666125949 13'-03	6.737666252932 18'-03
510	-4.564425'-01	6.737676725326 80'-03	6.737677154585 83'-03
520	-5.697100'-01	6.737687167865 13'-03	6.737687433370 44'-03
530	-3.301094'-01	6.737696941777 62'-03	6.737697135847 93'-03
540	-5.685695'-01	6.737705940705 81'-03	6.737706304309 39'-03
550	-4.983457'-01	6.737714461659 73'-03	6.737714977235 88'-03
560	-4.773043'-01	6.737722877179 93'-03	6.737723189705 71'-03
570	-5.245698'-01	6.737730792937 70'-03	6.737730973746 47'-03
580	-6.700377'-01	6.737737921191 16'-03	6.737738358646 71'-03
590	-3.889994'-01	6.737745100408 16'-03	6.737745371237 60'-03
600	-6.671242'-01	6.737751662554 39'-03	6.737752036131 12'-03
610	-4.950620'-01	6.737757824878 57'-03	6.737758375939 28'-03
620	-3.840306'-01	6.737764932695 23'-03	6.737764411464 10'-03
630	-3.033572'-01	6.737769670839 37'-03	6.737770161872 05'-03
640	-2.756428'-01	6.737775521331 74'-03	6.737775644842 25'-03
650	-3.408707'-01	6.737780426092 31'-03	6.737780976705 28'-03
660	-3.852482'-01	6.737785498495 70'-03	6.737785872563 56'-03
670	-5.888407'-01	6.737790062835 62'-03	6.737790646401 75'-03
680	-7.872203'-01	6.737794711058 27'-03	6.737795211182 06'-03
690	-3.430081'-01	6.737798524985 58'-03	6.737799578937 28'-03
700	-6.937112'-01	6.737803509975 23'-03	6.737803760842 65'-03
710	-3.669302'-01	6.737807129597 80'-03	6.737807767296 41'-03
720	-7.862510'-01	6.737810815129 93'-03	6.737811607977 38'-03
730	-5.162380'-01	6.737814320977 63'-03	6.737815291907 11'-03
740	-1.016203'+00	6.737818416433 91'-03	6.737818827499 37'-03
750	-8.151451'-01	6.737821264757 86'-03	6.737822222612 99'-03
760	-7.357829'-01	6.737824675372 95'-03	6.737825484593 33'-03
770	-6.082242'-01	6.737827533365 80'-03	6.737828623111 99'-03
780	-5.016716'-01	6.737830716502 18'-03	6.737831636199 79'-03
790	-4.737896'-01	6.737833547803 33'-03	6.737834538288 59'-03
800	-1.252605'+00	6.737836529953 67'-03	6.737837332229 87'-03
810	-1.300713'+00	6.737839440880 58'-03	6.737840023332 90'-03
820	-1.331237'+00	6.737841400136 34'-03	6.737842616582 59'-03
830	-5.401559'-01	6.737844317445 90'-03	6.737845116665 91'-03
840	-6.525357'-01	6.737847154512 64'-03	6.737847527993 05'-03
850	-8.088494'-01	6.737848773792 56'-03	6.737849854717 04'-03
860	-9.591714'-01	6.737850999267 20'-03	6.737852100750 35'-03
870	-1.185460'+00	6.737853659706 21'-03	6.737854269779 10'-03
880	-1.424537'+00	6.737855409270 53'-03	6.737856365284 41'-03
890	-6.967568'-01	6.737857621854 79'-03	6.737858390553 33'-03
900	-9.581641'-01	6.737859520382 43'-03	6.737860348688 53'-03
910	-1.313843'+00	6.737861265694 31'-03	6.737862242625 77'-03
920	-1.630921'+00	6.737863337966 71'-03	6.737864075141 00'-03
930	-9.874165'-01	6.737864840511 56'-03	6.737865848860 44'-03
940	-1.385051'+00	6.737865804783 79'-03	6.737867566274 56'-03
950	-7.232992'-01	6.737868559010 75'-03	6.737869229740 96'-03
960	-1.191267'+00	6.737869991409 30'-03	6.737870841493 77'-03
970	-1.731914'+00	6.737871327267 19'-03	6.737872403656 46'-03
980	-1.096796'+00	6.737873226739 37'-03	6.737873918242 94'-03
990	-1.628775'+00	6.737874573948 60'-03	6.737875387165 15'-03
1000	-9.935147'-01	6.737875474873 56'-03	6.737876812240 56'-03

REMARKS: Our definitions of subroutine and stability (def. 2.4-1 and 3.1-2) depend on the machine number system $(R, \mathcal{E})$ being considered. We can eliminate this dependence by defining an algorithm $\mathcal{A}$ for an ideal function f to be a constructive mapping from {the set of all $(R, \mathcal{E})$} into {the set of all subroutines}, such that $F \equiv \mathcal{A}(R, \mathcal{E})$ is a subroutine relative to $(R, \mathcal{E})$ and F and f satisfy (3.1-1). Thus $\mathcal{A}$ is a recursive operator (see sec. 1.5), mapping any determiner (α_R, ι_R) of $(R, \mathcal{E})$ into a determiner γ of such an $F \equiv \mathcal{A}(R, \mathcal{E})$. Roughly speaking, Algol procedures and Fortran subroutines are examples of such algorithms. We would then say $\mathcal{A}$ is stable relative to $(R, \mathcal{E})$ at x if $\mathcal{A}(R, \mathcal{E})$ is stable at x. Note that x is a real input, and therefore x depends on $(R, \mathcal{E})$. We would say $\mathcal{A}$ is stable at c (a numeric constant) if, for any $(R, \mathcal{E})$ and any real input $x = c$, $\mathcal{A}(R, \mathcal{E})$ is stable at x. If we allow $\mathcal{A}$ to take more arguments, say a list of algorithms of the above type (as well as $(R, \mathcal{E})$), then we get stronger and more general concepts of stability, analogous to those found in the literature on the numerical solution of ordinary differential equations.

BLANK PAGE

Chapter 4: ε-Limit, ε-Comparison, ε-Convergence, and ε-Continuity

4.1 ε-Limit and Truncation-Error Bounds

Define an operator, ϕ_{lim}, over the set S_{lim} of ideal functions of two variables, by

$$(4.1\text{-}1) \qquad \phi_{lim}(f)(x) = \lim_{y \to x} f(x,y) \quad .$$

Thus ϕ_{lim} maps an ideal function of two variables into an ideal function of one variable. This operator represents a notion of limit. (We consider limits of the forms $\lim_{\overline{y}_m \to \overline{x}_m} g(\overline{x}_m, \overline{y}_m)$ and $\lim_{y \to x_m} h(\overline{x}_m, y)$ at the end of section 4.3 and in section 5.4.) Having x in the argument list of f considerably simplifies notation because the TF part of $\mathfrak{F}$ depends upon where y is going; if we are interested in $\lim_{y \to x} g(y)$, we will simply form f, with $f(k,y) = g(y)$ when $k \neq \omega$, and then consider $\lim_{y \to x} f(x,y)$.

To formulate an ε-notion of ε-limit corresponding to ϕ_{lim}, we must define a set S_{lim} of ε-functions and an ε-operator (Φ_{lim}, Q_{lim}) over S_{lim} such that if $\mathfrak{F} \in S_{lim}$ and $\mathfrak{F} \approx f(P)$ then $\Phi_{lim}(\mathfrak{F}) \approx \phi_{lim}(f)(Q_{lim}(\mathfrak{F}, f, P))$. One way to define Φ_{lim} is to select some effective stopping criterion (see ch. 3) and define

$$(4.1\text{-}2) \qquad \Phi_{lim}(\mathfrak{F})(\varepsilon;x) \equiv (F(\varepsilon; x, y_\varepsilon), (RF \hat{\dagger} TF)(\varepsilon; x, y_\varepsilon), \omega) \quad ,$$

where y_ε is the value chosen by the stopping criterion when it is given ε, x and $\mathfrak{F}$. If we selected the S. C.## of definition 3.4-1,

we would have a totally inefficient Φ_{lim} with good accuracy, a large S_{lim} and a good Q_{lim}. But its total lack of efficiency rules out this Φ_{lim}.

Another method is suggested by the proofs of theorems 3.3-1 and 3.4-1. Roughly, this method proceeds at precision ε by

(1) finding a $\delta \leq \varepsilon$ and a $Y \in R(\delta)$ such that the truncation-error bound, $TF(\delta; x, Y)$, is $\leq \varepsilon$,

(2) finding an $\eta \leq \delta$ such that $RF(\eta; x, Y) \leq \varepsilon$, and

(3) defining $\Phi_{lim}(\mathcal{F})(\varepsilon; x)$ to be (approximately) $(F(\eta; x, Y), 2\varepsilon, \omega)$.

Of course, these steps will have to be modified and S_{lim} will have to be defined so that this process halts for each $\mathcal{F} \in S_{lim}$, any ε and any real input x. As we shall see, the only stability requirements needed to insure that this method converges concern TF (and not F or F $\hat{+}$ RF).

DEFINITION 4.1-1: Suppose $\mathcal{F} \approx f(P)$ for some P. We say TF is stably convergent at x relative to f precisely when

$$\left[\lim_{y \to x} f(x,y) \neq \omega\right] \Rightarrow \lim_{\varepsilon \to 0} TF(\varepsilon; x, y_\varepsilon) = 0 \quad ,$$

as long as the y_ε's are chosen by any reasonable at x stopping criterion.

This is a stability requirement on TF because, if we assume that $\lim_{\varepsilon \to 0} TF(\varepsilon; x, Y)$ always exists, then it is equivalent to requiring that $\lim_{y \to x} f(x,y) \neq \omega$ should imply the existence of a tf such that

(i) $\lim_{\varepsilon \to 0} TF(\varepsilon; x, Y) = tf(x,Y)$ for all $Y \in \mathcal{M}$ at which $tf(x,Y) \neq \omega$,

(ii) TF is stable at x (under def. 3.1-2), and

(iii) $\lim_{y \to x} tf(x,y) = 0$.

Let S_{lim} be the set of all ε-functions $\mathcal{F}$ of two variables such that for each $x \in \tilde{R}$ and each ε,

(1) [$Y \in \mathcal{R}(\varepsilon)$ and $RF(\varepsilon; x, Y) \neq \omega$] implies $\lim_{\varepsilon \to 0} RF(\varepsilon; x, Y) = 0$, and

(2) [$TF(\varepsilon; x, Y) \neq \omega$ for some $Y \in \mathcal{R}(\varepsilon)$] implies [$\lim_{\varepsilon \to 0} TF(\varepsilon; x, y_\varepsilon) = 0$ as long as the y_ε's are chosen by any reasonable at x stopping criterion].

For the following, we assume that an effective, reasonable at any $x \neq \omega$, stopping criterion, S.C., is given. We also assume that a $\lambda: \mathcal{E} \to \mathcal{M}$ is given which satisfies

(i) $\lambda(\varepsilon) \in \mathcal{R}(\varepsilon)$ and $\lambda(\varepsilon) > 0$ for all ε,

(ii) $\lim_{\varepsilon \to 0} \lambda(\varepsilon) = 0$, and

(iii) $\lambda(\varepsilon_i) = \langle \gamma(i, \cdot) \rangle$ for some recursive function, γ .

We define Φ_{lim} in terms of S.C., λ and the $\hat{I}$ and $\check{I}$ of section 2.8 by

DEFINITION 4.1-2: Let $\mathcal{F} \in S_{lim}$, x and ε be given. Let $y_{\varepsilon_1}, y_{\varepsilon_2}, \ldots$ be the values selected by S.C. for TF and x .

If $TF(\varepsilon; x, y_\varepsilon) = \omega$ then define $\Phi_{lim}(\mathcal{F})(\varepsilon; x) \equiv (\omega, \omega, \omega)$.

Otherwise, let δ be the largest member of $\mathcal{E}$ such that $\delta \leq \varepsilon$ and $TF(\delta; x, y_\delta) \leq \lambda(\varepsilon)$. If $RF(\delta; x, y_\delta) = \omega$ then define

$\Phi_{\lim}(\mathfrak{F})(\varepsilon;\, x) \equiv (\omega, \omega, \omega)$. Otherwise let j be the smallest integer such that $\varepsilon_j \le \delta$ and $RF(\varepsilon_j;\, x, y_\delta) \le \lambda(\varepsilon)$. Suppose $F(\varepsilon_j;\, x, y_\delta)$ is $< \alpha_R(j, k, \cdot) >$ and let $\beta_\varepsilon(\cdot)$ be $\alpha_R(j, k, \cdot)$. Define

(4.1-3)
$$\Phi_{\lim}(\mathfrak{F})(\varepsilon;\, x) \equiv (\hat{I}(\varepsilon, \beta_\varepsilon),\ 2 \hat{\times} \lambda(\varepsilon) \hat{+} |\hat{I}(\varepsilon, B_\varepsilon) \hat{-} \check{I}(\varepsilon, \beta_\varepsilon)|,\ \omega) \quad .$$

For $\mathfrak{F} \approx f(P)$, define

(4.1-4)
$$Q_{\lim}(\mathfrak{F},\, f,\, P) \equiv \{x:\ \{x\} \times (\mathcal{M} \cap \{\text{some neighborhood of } x\}) \subset P \text{ and } TF \text{ is stably convergent at } x \text{ relative to } f\} .$$

THEOREM 4.1-1: We have

$$(\Phi_{\lim},\, Q_{\lim}) \approx \beta_{\lim}(\mathfrak{S}_{\lim}) \quad .$$

Proof: Suppose $\mathfrak{F} \approx f(P)$ and $\mathfrak{F} \in \mathfrak{S}_{\lim}$. Let $x \in Q_{\lim}(\mathfrak{F},\, f,\, P)$ be such that $\ell \equiv \lim_{y \to x} f(x,y) \neq \omega$. Then for sufficiently small ε, $TF(\varepsilon;\, x, y_\varepsilon) \neq \omega$ and we can find a δ with $TF(\delta;\, x, y_\delta) \le \lambda(\varepsilon)$. Let δ denote the largest such value $\le \varepsilon$. This means that $|f(x, y_\delta) - \ell| \le \lambda(\varepsilon)$. If δ is sufficiently small (it will be, if ε was) then $RF(\delta;\, x, y_\delta) \neq \omega$ and we can find an $\eta \le \delta$ with $RF(\eta;\, x, y_\delta) \le \lambda(\varepsilon)$. Let ε_j be the largest such η . This means that

$$|F(\varepsilon_j;\, x, y_\delta) - f(x, y_\delta)| \le \lambda(\varepsilon) \quad .$$

Thus we have

(4.1-5)
$$|F(\varepsilon_j;\, x, y_\delta) - \ell| \le 2\lambda(\varepsilon) \quad .$$

For each ε, the corresponding $F(\varepsilon_{j(\varepsilon)}; x, y_{\delta(\varepsilon)})$ equals some $< \alpha_R(j(\varepsilon), k(\varepsilon), \cdot) >$. Let $\beta_\varepsilon(\cdot)$ be $\alpha_R(j(\varepsilon), k(\varepsilon), \cdot)$. From (4.1-5) we know that

$$\lim_{\varepsilon \to 0} < \beta_\varepsilon > = \ell \quad , \tag{4.1-6}$$

if $\ell = \pm \infty$ then $< \beta_\varepsilon > = \ell$ for all sufficiently small ε .

It follows that

$$\lim_{\varepsilon \to 0} \hat{I}(\varepsilon, \beta_\varepsilon) = \lim_{\varepsilon \to 0} \check{I}(\varepsilon, \beta_\varepsilon) = \ell \quad , \tag{4.1-7}$$

because $\hat{I}(\varepsilon, \beta_\varepsilon)$ and $\check{I}(\varepsilon, \beta_\varepsilon)$ are both in $N_\varepsilon(< \beta_\varepsilon >)$. Further, by the triangle inequality and (4.1-5) we have

$$|\hat{I}(\varepsilon, \beta_\varepsilon) - \ell| \leq 2 \hat{\times} \lambda(\varepsilon) \hat{+} |\hat{I}(\varepsilon, \beta_\varepsilon) \hat{-} \check{I}(\varepsilon, \beta_\varepsilon)| \quad . \tag{4.1-8}$$

By (4.1-7) and theorem 2.8-1, the right side of (4.1-8) approaches 0 as $\varepsilon \to 0$, which means that

$$\lim_{\varepsilon \to 0} \Phi_{\lim}(\mathfrak{F})(\varepsilon; x) \equiv (\lim_{y \to x} f(x,y), 0, \omega) \quad .$$

This completes the proof.

$\Phi_{\lim}$ overcomes instabilities in F and/or $F \hat{+} RF$ by using a stably convergent TF . First it picks a place $(y_{\delta(\varepsilon)})$ at which to evaluate the ε-limit, and then it increases the precision until the crest of the $\varepsilon_{j(\varepsilon)}$-wave has moved past $y_{\delta(\varepsilon)}$. The efficiency of $\Phi_{\lim}$ will depend on

(1) how closely TF and RF approximate the errors that they bound, how difficult they are to evaluate, and

(2) how judiciously S.C. chooses its y_ε's; if they are unnecessarily close to x then j may have to be made very large (an expensive enterprise) before $RF(\varepsilon_j; x, y_\delta) \leq \lambda(\varepsilon)$, especially when F or $F \hat{+} RF$ is unstable at x.

Thus we say that Φ_{lim} offers a potentially efficient algorithm for overcoming instability.

We will use $\underset{y \to x}{LIM}\, \mathfrak{F}(\varepsilon; x, y)$ to denote $\Phi_{lim}(\mathfrak{F})(\varepsilon; x)$ and we call this the ε-limit of $\mathfrak{F}$ at x.

4.2 ε-Comparison Relations, $<_\varepsilon$ and $=_\varepsilon$

In the following sections, we will need an ε-less-than relation, $<_\varepsilon$, and an ε-equality relation, $=_\varepsilon$. We will define these relations so that

(1) $\mathcal{F}(\varepsilon; \bar{x}_m) <_\varepsilon \mathcal{G}(\varepsilon; \bar{y}_n)$ is true when, based only on the information given by $\mathcal{F}(\varepsilon; \bar{x}_m)$ and $\mathcal{G}(\varepsilon; \bar{y}_n)$, $f(\bar{x}_m)$ must be less than $g(\bar{y}_n)$, and

(2) $\mathcal{F}(\varepsilon; \bar{x}_m) =_\varepsilon \mathcal{G}(\varepsilon; \bar{y}_n)$ is true when, based only on the information given by $\mathcal{F}(\varepsilon; \bar{x}_m)$ and $\mathcal{G}(\varepsilon; \bar{y}_n)$, $f(\bar{x}_m)$ might be equal to $g(\bar{y}_n)$.

Essentially, the ε-less-than relationship holds when the interval $[F - RF, F + RF]$ lies entirely to the left of the interval $[G - RG, G + RG]$, and ε-equality holds when these intervals overlap (see figure 4.2-1). Of course $=_\varepsilon$ will not be an equivalence relation because it will not be transitive.

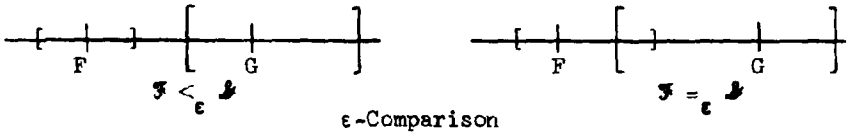

(a) FIGURE 4.2-1 ε-Comparison (b)

DEFINITION 4.2-1: Let x and y be poor real inputs. For $\varepsilon \in \mathcal{E}$, define $x =_\varepsilon y$ to be true (and $x \neq_\varepsilon y$ to be false) precisely when

$$|X(\varepsilon) \ominus Y(\varepsilon)| \leq RX(\varepsilon) \oplus RY(\varepsilon) \quad . \tag{4.2-1}$$

Define $x <_\varepsilon y$ to be true (and $x \not<_\varepsilon y$ to be false) precisely when $x \neq_\varepsilon y$ and $X(\varepsilon) < Y(\varepsilon)$.

For $a_1, a_2 \geq 0$, $a_3 \in R(\varepsilon)$, when the triple $(\bar{a}_3)$ appears in an ε-comparison (ε fixed) this triple is to be understood to denote the poor real input $a \equiv (A, RA)$ defined by

$$(A(\delta), RA(\delta)) \equiv \begin{cases} (\omega, \omega) & \delta > \varepsilon \\ (a_1, a_2) & \delta \leq \varepsilon \end{cases} .$$

This convention allows us to ε-compare ε-function values directly. At all times precisely one of $\mathcal{F}(\varepsilon; \bar{x}_m) =_\varepsilon \mathcal{A}(\varepsilon; \bar{y}_n)$, $\mathcal{F}(\varepsilon; \bar{x}_m) <_\varepsilon \mathcal{A}(\varepsilon; \bar{y}_n)$ and $\mathcal{A}(\varepsilon; \bar{y}_n) <_\varepsilon \mathcal{F}(\varepsilon; \bar{x}_m)$ holds.

THEOREM 4.2-1: Let x and y be real inputs. Then

$$x = y \Leftrightarrow [x =_\varepsilon y \text{ for all } \varepsilon]$$
$$x < y \Rightarrow [x <_\varepsilon y \text{ for all sufficiently small } \varepsilon]$$
$$x < y \not\Leftarrow [x <_\varepsilon y \text{ for some } \varepsilon] .$$

Proof: If $x = y$ then we have, for all ε,

$$|X(\varepsilon) \dot{-} Y(\varepsilon)| \leq |X(\varepsilon) - Y(\varepsilon)| \leq RY(\varepsilon) + RY(\varepsilon) \leq RX(\varepsilon) \hat{+} RY(\varepsilon) ,$$

and so $x =_\varepsilon y$ for all ε. If $x =_\varepsilon y$ for all ε, then applying theorem 2.8-1 and taking the limit of (4.2-1) as $\varepsilon \to 0$ yields $|x - y| \leq 0$, so $x = y$.

If $x < y$ then, for all sufficiently small ε,

(4.2-2) $$X(\varepsilon) < Y(\varepsilon), \qquad |X(\varepsilon) \dot{-} Y(\varepsilon)| > RX(\varepsilon) \hat{+} RY(\varepsilon) ,$$

the second inequality holding because, by theorem 2.8-1, $|X(\varepsilon) \dot{-} Y(\varepsilon)| \to |x-y| > 0$ whereas $RX(\varepsilon) \dot{+} RY(\varepsilon) \to 0$. Thus $x <_\varepsilon y$ for all sufficiently small ε. If $x <_\varepsilon y$ (some ε) then (4.2-2) holds and so

$$|X(\varepsilon) - Y(\varepsilon)| > RX(\varepsilon) + RY(\varepsilon) \quad ,$$

implying $x < y$. This completes the proof.

Let bool[statement] be 1 if the statement is true and 0 if it is false. The notions of comparison given by

$$\not{D}_*(\bar{f}_2)(\bar{x}_m) = \begin{cases} \omega & \text{if } f_i(\bar{x}_m) = \omega \ (i = 1 \text{ or } 2) \\ \text{bool}\,[f_1(\bar{x}_m) * f_2(\bar{x}_m)] & \text{otherwise,} \end{cases}$$

for $*$ being $=$ and $<$, can be ε-ized easily, to yield <u>weak</u> ε-operators, $(\mathfrak{P}_*, Q_*)$. The weakness of these ε-operators is due to the fact that the information given by $\mathfrak{F}_i(\varepsilon; \bar{x}_m)(i = 1, 2)$ may never (for any ε) be sufficient to determine that $f_1(\bar{x}_m)$ <u>must</u> equal $f_2(\bar{x}_m)$. See Bishop [B1, p. 24] and Aberth [A1, pp. 287-8] for similar considerations.

Techniques from interval analysis can be formalized in the ε-calculus to yield a weak ε-operator corresponding to the operator,

$$\not{D}_{pos}(f)(a, b) = \begin{cases} \omega & \text{if } a \text{ or } b \text{ is in } \{-\infty, \infty, \omega\} \\ \text{bool}\,[f(x) > 0 \text{ for } a \leq x \leq b] & \text{otherwise} \quad . \end{cases}$$

We leave this to the reader.

4.3 ε-Convergence and ε-Continuity: Pointwise

In this context, we say f converges at x precisely when $\lim_{y \to x} f(x, y) \neq \omega$; i.e., precisely when the limit exists in the usual sense. Otherwise, we say f diverges at x .

DEFINITION 4.3-1: Fix x and ε . We say $\mathfrak{F}$ ε-converges at x precisely when

(4.3-1) $$\underset{y \to x}{LIM}\, \mathfrak{F}(\varepsilon; x, y) \neq_{\varepsilon} \omega \ .$$

Otherwise we say $\mathfrak{F}$ ε-diverges at x .

We say f is continuous at x precisely when

(4.3-2) $$f(x, x) = \lim_{y \to x} f(x, y) \neq \omega \ .$$

Otherwise we say f is discontinuous at x . Note that (4.3-2) uses the transitivity relation $a = b \neq c \Rightarrow a \neq c$ to insure that $f(x, x) \neq \omega$. Since $=_{\varepsilon}$ and $\neq_{\varepsilon}$ do not satisfy such a transitivity relation, we must explicitly insure this in

DEFINITION 4.3-2: Fix x and ε . We say $\mathfrak{F}$ is ε-continuous at x precisely when

$$\mathfrak{F}(\varepsilon; x, x) =_{\varepsilon} \underset{y \to x}{LIM}\, \mathfrak{F}(\varepsilon; x, y) \neq_{\varepsilon} \omega \text{ and } \mathfrak{F}(\varepsilon; x, x) \neq_{\varepsilon} \omega \ .$$

Otherwise we say $\mathfrak{F}$ is ε-discontinuous at x . We say $\mathfrak{F}$ is strongly ε-discontinuous at x precisely when

$$\mathfrak{F}(\varepsilon;\ x,\ x) \neq_{\varepsilon} \underset{y\to x}{\mathrm{LIM}}\,\mathfrak{F}(\varepsilon;\ x,\ y) \quad .$$

If we are interested in the continuity of $g(y)$ at $y = x$, then we simply form f, with $f(k, y) = g(y)$ when $k \neq \omega$, and investigate the ε-continuity at x of some $\mathfrak{F}$ corresponding to f . Let ε_j be as in definition 4.1-2, when a value for ε_j is found (i.e., when $TF(\varepsilon;\ x,\ y_{\varepsilon}) \neq \omega$); otherwise let ε_j be ε . Let $\mathfrak{F}_{\varepsilon}$ denote the finite subset of $\mathfrak{M}^{(3)}$ given by

$$\mathfrak{F}_{\varepsilon} \equiv \{\mathfrak{F}(\eta;\ x,\ Y):\ Y \in \mathcal{R}(\eta) \text{ and } \varepsilon_j \leq \eta \leq \varepsilon\} \quad .$$

From the definitions of ε-function, of ε-limit and of ε-equality, it follows that

(1) $\mathfrak{F}$ ε-converges at x when, based only on information contained in $\mathfrak{F}_{\varepsilon}$, f <u>must</u> converge at x, and

(2) $\mathfrak{F}$ ε-diverges at x when, based only on information contained in $\mathfrak{F}_{\varepsilon}$, f <u>might</u> diverge at x .

Let $\mathfrak{F}'_{\varepsilon}$ denote the finite subset of $\mathfrak{M}^{(3)}$ given by

$$\mathfrak{F}'_{\varepsilon} \equiv \{\mathfrak{F}(\eta;\ x,\ x):\ \varepsilon_j \leq \eta \leq \varepsilon\} \quad .$$

As above, we have

(1) $\mathfrak{F}$ is ε-continuous at x when, based only on $\mathfrak{F}_{\varepsilon} \cup \mathfrak{F}'_{\varepsilon}$, f <u>might</u> be continuous at x, and

(2) $\mathfrak{F}$ is strongly ε-discontinuous at x when, based only on $\mathfrak{F}_{\varepsilon} \cup \mathfrak{F}'_{\varepsilon}$, f <u>must</u> be discontinuous at x .

These definitions can be expressed in operator, ε-operator form as follows. Let $S_{\lim}$ be as in section 4.1. Define operators

ϕ_{conv}, ϕ_{cont}, each over S_{lim}, by

$$\phi_{conv}(f)(x) = \text{bool}\,[f \text{ converges at } x] \ ,$$
$$\phi_{cont}(f)(x) = \text{bool}\,[f \text{ is continuous at } x] \ ,$$

so long as $x \neq \omega$. Of course, $\phi_{conv}(f)(\omega) = \phi_{cont}(f)(\omega) = \omega$. Define ε-operators corresponding weakly to the above by equating, for $x \neq_{\varepsilon} \omega$,

$$\Phi_{conv}(\mathfrak{F})(\varepsilon;\, x) \equiv (\text{bool}\,[\mathfrak{F} \text{ } \varepsilon\text{-converges at } x] \ , \quad \text{bool}\,[\mathfrak{F} \text{ } \varepsilon\text{-diverges at } x], \omega) \ , \tag{4.3-3}$$

$$\Phi_{cont}(\mathfrak{F})(\varepsilon;\, x) \equiv (\text{bool}\,[\mathfrak{F} \text{ is } \varepsilon\text{-continuous at } x] \ , \quad \text{1-bool}\,[\mathfrak{F} \text{ is strongly } \varepsilon\text{-discontinuous at } x], \omega) \ , \tag{4.3-4}$$

$$Q_{cont}(\mathfrak{F}, f, P) \equiv Q_{lim}(\mathfrak{F}, f, P) \cap \{x: \ (x, x) \in P\} \ . \tag{4.3-5}$$

THEOREM 4.3-1: We have

$$(\Phi_{conv}, Q_{lim}) \sim \phi_{conv}(S_{lim}) \ , \tag{4.3-6}$$

$$(\Phi_{cont}, Q_{cont}) \sim \phi_{cont}(S_{lim}) \ . \tag{4.3-7}$$

Further, if f converges at x for all $x \in Q_{lim}(\mathfrak{F}, f, P)$ then $\Phi_{conv}(\mathfrak{F})$ is not weak. If f is discontinuous at x and $f(x, x) \neq \omega$ and $\lim_{y \to x} f(x, y) \neq \omega$ for all $x \in Q_{cont}(\mathfrak{F}, f, P)$, then $\Phi_{cont}(\mathfrak{F})$ is not weak.

Proof: Consider ε-convergence first. Suppose $\mathfrak{F} \approx f(P)$ and x is in $Q_{lim}(\mathfrak{F}, f, P)$. If f converges at x then by theorem 4.1-1,

$\lim_{y \to x} \mathfrak{F}(\varepsilon; x, y) \to (\lim_{y \to x} f(x, y), 0, \omega)$ as $\varepsilon \to 0$ and so $\lim_{y \to x} \mathfrak{F}(\varepsilon, x, y)$ must be $\neq_\varepsilon \omega$ for all sufficiently small ε. If f diverges at x, then $\lim_{y \to x} f(x, y) = \omega$ and so $\lim_{y \to x} \mathfrak{F}(\varepsilon; x, y) =_\varepsilon \omega$ must hold for all ε. This and (4.3-3) yield (4.3-6) and the first remark after (4.3-7).

Consider ε-continuity. Suppose $\mathfrak{F} \sim f(P)$ and x is in $Q_{cont}(\mathfrak{F}, f, P)$. If f is continuous at x then theorems 4.1-1, 4.1-2 and the fact that $(x, x) \in P$ give us ε-continuity for all sufficiently small ε. If f is discontinuous at x then $f(x, x) = \omega$ or $\lim_{y \to x} f(x, y) = \omega$ or else $f(x, x) \neq \lim_{y \to x} f(x, y)$. In the first two cases we have $\mathfrak{F}(\varepsilon; x, x) =_\varepsilon \omega$ or $\underset{y \to x}{LIM}\, \mathfrak{F}(\varepsilon; x, y) =_\varepsilon \omega$ for all ε, and so $\mathfrak{F}$ is always ε-discontinuous at x. In the third case theorems 4.1-1 and 4.2-1 imply that $\mathfrak{F}(\varepsilon; x, y) \neq_\varepsilon \underset{y \to x}{LIM}\, \mathfrak{F}(\varepsilon; x, y)$ for all sufficiently small ε, so $\mathfrak{F}$ is ε-discontinuous at x for all sufficiently small ε. This and (4.3-4) yield (4.3-7) and the last remark. This completes the proof.

Let f be an ideal function of $m + 1$ variables and g an ideal function of $2m$ variables $(m \geq 1)$. We can easily discretize "f converges at $\bar{x}_m$" (true if $\lim_{y \to \bar{x}_m} f(\bar{x}_m, y) \neq \omega$) and "$f$ is continuous at $\bar{x}_m$" (true if $f(\bar{x}_m, x_m) = \lim_{y \to \bar{x}_m} f(\bar{x}_m, y) \neq \omega$). However, with our present setup we cannot discretize "g converges at $\bar{x}_m$" (true if $\lim_{\bar{y}_m \to \bar{x}_m} g(\bar{x}_m, \bar{y}_m) \neq \omega$) and "$g$ is continuous at $\bar{x}_m$" (true if $g(\bar{x}_m, \bar{x}_m) = \lim_{\bar{y}_m \to \bar{x}_m} g(\bar{x}_m, \bar{y}_m) \neq \omega$) because our truncation-error bounds do not give the necessary local information about all the possible

approaches of $\bar{y}_m$ to $\bar{x}_m$. (The $\bar{x}_m$ appearing in $g(\bar{x}_m, \bar{y}_m)$ may just be dummy variables telling where $\bar{y}_m$ is to go.) We could have done this latter discretization if we had assumed truncation-error bounds, $TG(\varepsilon; \bar{x}_m, \bar{y}_m)$, for limits of the form $\underline{\lim}_{\bar{y}_m \to \bar{x}_m} g(\bar{x}_m, \bar{y}_m)$.

We would then have defined an ε-limit of the form $\underline{LIM}_{\bar{y}_m \to \bar{x}_m} \mathcal{L}(\varepsilon; \bar{x}_m, \bar{y}_m)$.

Of course this ε-limit would not be more powerful, computationally, than $\underline{LIM}_{\bar{y} \to \bar{x}_m} \mathcal{F}(\varepsilon; \bar{x}_m, y)$, because m successive applications of the latter ε-limit are essentially as good as one application of the former; the difference between these two ε-limits is that the latter one will approach $\bar{x}_m$ along the m-dimensional axes whereas the former one may take any approach. This follows from the relation,

$$\left[\underline{\lim}_{\bar{y}_m \to \bar{x}_m} g(\bar{x}_m, \bar{y}_m) \neq \omega \right] \Rightarrow \underline{\lim}_{\bar{y}_m \to \bar{x}_m} g(\bar{x}_m, \bar{y}_m) = \underline{\lim}_{y_1 \to x_1} \cdots \underline{\lim}_{y_m \to x_m} g(\bar{x}_m, \bar{y}_m) .$$

Thus, if the limit exists, the domain set of $\mathcal{L}$ is large enough and the truncation-error bounds involved are stably convergent, then both these ε-limits will work. We have avoided the more complicated form of ε-limit in order to simplify notation.

4.4 Discontinuities

In order to discretize convergence and continuity over intervals, we must know more about the kinds of discontinuities f can have in P while there still exists an ε-function corresponding to f over P. Consider the ideal function f defined for finite x and y by

$$f(x, y) = \text{bool}\,[x < y] \quad .$$

Define an ε-function $\mathfrak{F}$ by

$$\mathfrak{F}(\varepsilon; x, y) \equiv (\text{bool}\,[x <_{\eta} y \text{ for some } \eta \geq \varepsilon],\ \text{bool}\,[x \neq_{\eta} y \text{ or } RX(\eta) = RY(\eta) = 0 \text{ for some } \eta \geq \varepsilon],\ \omega),$$

so long as x and y are $\neq_{\varepsilon} -\infty, \infty, \omega$. In this case we have

$$\mathfrak{F} \approx f(\{(x, y):\ x \neq y \text{ or } x = y \in \mathcal{M}\}) \quad .$$

However, we only have

$$\mathfrak{F} \sim f(\tilde{R}^{(2)}) \quad ,$$

the correspondence being weak because $RX(\eta) \neq 0$ for $\eta \geq \varepsilon$ implies $\mathfrak{F}(\varepsilon; x, x) \equiv (1, 1, \omega)$ and so $\lim_{\varepsilon \to 0} RF(\varepsilon; x, x) \neq 0$ for any $x \notin \mathcal{M}$. Let $\mathfrak{F}'$ be any ε-function weakly corresponding to f over $\tilde{R}^{(2)}$. Then for $x = y \notin \mathcal{M}$, $RF'(\varepsilon; x, y)$ cannot go to 0 with ε because the inputted values of x and y will always be inexact, and so $\mathfrak{F}'$ will never have enough information to decide for sure that $x = y$.

Many variations on this basic theme are possible. The underlying principle is given by

THEOREM 4.4-1: Suppose $\mathfrak{F} \approx f(P)$ and $\bar{x}_m \in P$ is a point of discontinuity of f (i.e., $f(\bar{x}_m) = \omega$ or $\lim_{\bar{y}_m \to \bar{x}_m} f(\bar{y}_m) \neq f(\bar{x}_m)$) and $f(\bar{x}_m) \neq \omega$. Then at least one $x_j \in \mathcal{M}$.

Proof: Suppose $\mathfrak{F}$, f, P and $\bar{x}_m$ satisfy the hypotheses, but $|x_i| \neq \infty$ and $RX_i(\varepsilon) \neq 0$ for all ε and $i = 1,2,\ldots, m$. We will prove that this implies $\lim_{\varepsilon \to 0} RF(\varepsilon; \bar{x}_m) \neq 0$, a contradiction. Suppose γ_1 and γ_2 are the given determiners of F and RF and that F and RF involve respectively r_1 and r_2 subroutine constants. For $k = 1,2$ let $n_k(\varepsilon_i)$ be the least value of n such that $\gamma_k(i, GN_n(\bar{x}_{m+r_k})) \neq 0$, and define

$$n(\varepsilon_i) = \max(n_1(\varepsilon_i), n_2(\varepsilon_i)) \quad .$$

Let $\sigma(\varepsilon)$ be the m-dimensional rectangle of real inputs,

$$\sigma(\varepsilon) \equiv \{\bar{y}_m: \ GN_{n(\varepsilon)}(\bar{y}_m) = GN_{n(\varepsilon)}(\bar{x}_m)\} \quad .$$

All the sides of $\sigma(\varepsilon)$ have positive length. Let ℓ_+ and ℓ_- be the limit superior and limit inferior of $f(\bar{y}_m)$ as $\bar{y}_m \to \bar{x}_m$, and define

$$K = \begin{cases} \frac{1}{2}(\ell_+ - \ell_-) & \text{if } \ell_+ \neq \ell_- \\ \frac{1}{2}|\ell_+ - f(\bar{x}_m)| & \text{otherwise} \end{cases} ,$$

$$h(\varepsilon, c) = \sup_{\bar{y}_m \in \sigma(\varepsilon)} |c - f(\bar{y}_m)| \quad \text{for } c \in \bar{R} \quad .$$

$K > 0$ because f is discontinuous at $\bar{x}_m$. There are $\bar{y}_m \in \sigma(\varepsilon)$ which make $f(\bar{y}_m)$ arbitrarily close to the one of ℓ_+, ℓ_-, $f(\bar{x}_m)$

which is furthest from c. Thus we have

$$h(\varepsilon, c) \geq K \quad \text{for any } c \in \tilde{R} \text{ and any } \varepsilon .$$

For any real inputs $\bar{y}_m$ we have

$$RF(\varepsilon; \bar{y}_m) \geq |F(\varepsilon; \bar{y}_m) - f(\bar{y}_m)| .$$

For $\bar{y}_m \in \sigma(\varepsilon)$ we have $F(\varepsilon; \bar{y}_m) = F(\varepsilon; \bar{x}_m)$ and $RF(\varepsilon; \bar{y}_m) = RF(\varepsilon; \bar{x}_m)$, yielding

$$RF(\varepsilon; \bar{x}_m) \geq \sup_{\bar{y}_m \in \sigma(\varepsilon)} |F(\varepsilon; \bar{x}_m) - f(\bar{y}_m)| = h(\varepsilon, F(\varepsilon; \bar{x}_m)) .$$

Thus $RF(\varepsilon; \bar{x}_m) \geq K > 0$ for any ε, the desired contradiction. This completes the proof.

COROLLARY 4.4-1: Suppose $\mathcal{F}$, f, P and $\bar{x}_m$ satisfy the hypotheses of the above theorem. Then for each $j = 1,2,\ldots, m$, either $x_j \in \mathcal{M}$ or the function $g(y) = f(x_1, \ldots, x_{j-1}, y, x_{j+1}, \ldots, x_m)$ is discontinuous at $y = x_j$.

Proof: Define $\mathcal{F}$ by setting $TG \equiv \omega$ and

$$G(\varepsilon; y) = F(\varepsilon; x_1,\ldots, x_{j-1}, y, x_{j+1},\ldots, x_m) ,$$

$$RG(\varepsilon; y) = RF(\varepsilon; x_1,\ldots, x_{j-1}, y, x_{j+1}, \ldots, x_m) .$$

Then $\mathcal{F} \approx g(\lceil x_j \rceil)$ and $g(x_j) = f(\bar{x}_m) \neq \omega$. If g is discontinuous at x_j then, by the above theorem, $x_j \in \mathcal{M}$. This completes the proof.

COROLLARY 4.4-2: If $\mathfrak{F} \approx f(P)$, $(\bar{x}_m, x_m) \in P$, $f(\bar{x}_m, x_m) \neq \omega$ and f is discontinuous at $\bar{x}_m$, then $x_m \in \mathcal{M}$.

Proof: Under the given assumptions, $(\bar{x}_m, x_m)$ is a point of discontinuity of f, so, by corollary 4.4-2, either $x_m \in \mathcal{M}$ or $g(y) = f(\bar{x}_m, y)$ is continuous in y at $y = x_m$. The latter alternative is ruled out by assumption, so we have $x_m \in \mathcal{M}$. This completes the proof.

COROLLARY 4.4-3: Let P be a set of m-tuples of numbers and suppose $\mathfrak{F} \approx f(P)$. Then f is continuous at every $\bar{x}_m \in P$ with $f(\bar{x}_m) \neq \omega$ and $|x_i| \neq \infty$ $(i = 1, \ldots, m)$.

The second use of P is as a set of real inputs (see sec. 2.3). For example, P might be $\mathcal{R}^{(m)}$.

Proof: Suppose $\mathfrak{F}$, f and P satisfy the hypotheses. Assume that $\bar{x}_m \in P$ with $f(\bar{x}_m) \neq \omega$ (and hence $x_i \neq \omega$ for all i) and $|x_i| \neq \infty$ for $i = 1, \ldots, m$. Define $\bar{y}_m$ by

$$Y_i(\varepsilon) = X_i(\varepsilon) \quad ,$$

$$RY_i(\varepsilon) = \max(\Theta(\varepsilon), RX_i(\varepsilon)) \quad ,$$

for $i = 1, \ldots, m$ and all ε, where $\Theta(\varepsilon)$ is the smallest positive number in $\mathcal{R}(\varepsilon)$. Then each y_i is a real input and $\bar{y}_m \in P$. From theorem 4.4-1, we know that f cannot be discontinuous at $\bar{y}_m$, and hence not at $\bar{x}_m$. This completes the proof.

4.5 ε-Convergence and ε-Continuity: Over Intervals

For simplicity, we consider open intervals. Let $o(a, b)$ denote the open interval between a and b ($a, b \in \tilde{R}$). Define the open ε-interval between a and b, $\mathcal{O}(\varepsilon; a, b)$, by

$$\mathcal{O}(\varepsilon; a, b) \equiv \{Y: Y \in R(\varepsilon) \cap o(a, b), Y \neq_{\varepsilon} a, Y \neq_{\varepsilon} b\} .$$

For $Y \in R(\varepsilon)$, the decision $Y \in \mathcal{O}(\varepsilon; a, b)$ is effective, given real inputs a and b, and we have

$$\mathcal{O}(\varepsilon; a, b) \subset o(a, b) \quad \text{for all } \varepsilon ,$$

$$\bigcup_{i \geq 1} \mathcal{O}(\varepsilon_i; a, b) \equiv o(a, b) \cap \mathcal{M} . \tag{4.5-1}$$

We say f converges over $o(a, b)$ precisely when f converges at all $x \in o(a, b)$. Otherwise, we say f diverges in $o(a, b)$.

DEFINITION 4.5-1: We say $\mathcal{F}$ ε-converges over $\mathcal{O}(\varepsilon; a, b)$ precisely when $\mathcal{F}$ ε-converges at all $x \in \mathcal{O}(\varepsilon; a, b)$. Otherwise we say $\mathcal{F}$ ε-diverges in $\mathcal{O}(\varepsilon; a, b)$.

We say f is continuous over $o(a, b)$ precisely when f is continuous at all $x \in o(a, b)$. Otherwise we say f is discontinuous in $o(a, b)$.

DEFINITION 4.5-2: We say $\mathcal{F}$ is ε-continuous over $\mathcal{O}(\varepsilon; a, b)$ precisely when $\mathcal{F}$ is ε-continuous at all $x \in \mathcal{O}(\varepsilon; a, b)$. Otherwise we say $\mathcal{F}$ is ε-discontinuous in $\mathcal{O}(\varepsilon; a, b)$. We say $\mathcal{F}$ is strongly ε-discontinuous in $\mathcal{O}(\varepsilon; a, b)$ precisely when there

is an $x \in \mathfrak{G}(\varepsilon; a, b)$ such that $\mathfrak{F}$ is strongly ε-discontinuous at x.

We express this in operator, ε-operator form by defining, for $a \neq \omega$ and $b \neq \omega$ and $f \in S_{lim}$,

$$\phi_{convo}(f)(a, b) = \text{bool [f converges over } o(a, b)]$$
$$\phi_{conto}(f)(a, b) = \text{bool [f is continuous over } o(a, b)] ,$$

and defining, for $a \neq_\varepsilon \omega$ and $b \neq_\varepsilon \omega$,

$$\Phi_{convo}(\mathfrak{F})(\varepsilon; a, b) \equiv (\text{bool [}\mathfrak{F} \text{ } \varepsilon\text{-converges over } \mathfrak{G}(\varepsilon; a, b)], \omega, \omega) ,$$

$$\Phi_{conto}(\mathfrak{F})(\varepsilon; a, b) \equiv (\text{bool [}\mathfrak{F} \text{ } \varepsilon\text{-continuous over } \mathfrak{G}(\varepsilon; a, b)] ,$$

$$\text{1-bool [}\mathfrak{F} \text{ is strongly } \varepsilon\text{-discontinuous in } \mathfrak{G}(\varepsilon; a, b)], \omega) .$$

Since the evaluation of $\Phi_{conto}(\mathfrak{F})(\varepsilon; a, b)$ and $\Phi_{convo}(\mathfrak{F})(\varepsilon; a, b)$ involves the evaluation of $\Phi_{lim}(\mathfrak{F})$ only at ε-points $(\varepsilon; x)$ for which $x \in \mathfrak{R}(\varepsilon)$, the set S_o of ε-functions to which Φ_{convo} and Φ_{conto} may be applied is defined as follows. Let S_o be the set of all ε-functions $\mathfrak{F}$ of two variables such that for each ε and each $x \in \mathfrak{R}(\varepsilon)$,

(1) $[Y \in \mathfrak{R}(\varepsilon)$ and $RF(\varepsilon; x, Y) \neq \omega]$ implies $\lim_{\varepsilon \to 0} RF(\varepsilon; x, Y) = 0$, and

(2) $[TF(\varepsilon; x, Y) \neq \omega$ for some $Y \in \mathfrak{R}(\varepsilon)]$ implies $[\lim_{\varepsilon \to 0} TF(\varepsilon; x, y_\varepsilon) = 0$ as long as the y_ε's are chosen by any reasonable at x stopping criterion].

In order for Φ_{convo} and Φ_{combo} to work well on $\mathcal{F} \in \mathbf{S}_{\rho}$, $\mathcal{F}$ will have to approximate its f uniformly over $o(a, b) \cap \mathcal{M}$ in a sense to be defined. Otherwise, for example, f may converge over $o(a, b)$, but for each i, $\mathcal{F}$ may ε_i-diverge at x for each $x \in \mathcal{O}(\varepsilon_i; a, b) - \mathcal{O}(\varepsilon_{i-1}; a, b)$. We would then have

(1) $x \in \mathcal{M} \cap o(a, b) \Rightarrow$ [$\mathcal{F}$ ε-converges at x for all sufficiently small ε], and

(2) $\mathcal{F}$ ε-diverges in $\mathcal{O}(\varepsilon; a, b)$ for all ε.

DEFINITION 4.5-3: We say $\mathcal{F}$ approximates f uniformly at a, b, written $\mathcal{F} \overset{u}{\approx} f[a, b]$, precisely when there is a $\delta > 0$ such that for each $\varepsilon \leq \delta$ and every $(x, y) \in \mathcal{O}(\varepsilon; a, b)^{(2)}$ we have

(1) $f(x, y) \neq \omega \Rightarrow \mathcal{F}(\varepsilon; x, y) \neq_{\varepsilon} \omega$, and

(2) $\lim_{z \to x} f(x, z) \neq \omega \Rightarrow TF(\varepsilon; x, y) \neq \omega$.

Let P be a set of pairs of real inputs. We say P covers $o(a, b)^{(2)}$ precisely when, for each pair of numbers $(c, d) \in o(a, b)^{(2)}$ there is a pair $(x, y) \in P$ with $x = c$ and $y = d$. Let Q be a set of real inputs. We say Q covers $o(a, b) \cap \mathcal{M}$ precisely when, for each ε and each number $c \in o(a, b) \cap \mathcal{R}(\varepsilon)$ there is a real input $x \in Q$ with $X(\delta) = c$ and $RX(\delta) = 0$ for $\delta \leq \varepsilon$. Define

$$Q_o(\mathcal{F}, f, P) \equiv \{(a, b): a \neq b, P \text{ covers } o(a, b)^{(2)}, \mathcal{F} \overset{u}{\approx} f[a, b],$$

$f(x, y) \neq \omega$ for $(x, y) \in o(a, b)^{(2)}$, $Q_{lim}(\mathcal{F}, f, P)$ covers $o(a, b) \cap \mathcal{M}\}$.

THEOREM 4.5-1: We have

$$(\Phi_{convo}, Q_o) \sim \phi_{convo}(S_o) \quad ,$$

$$(\Phi_{conto}, Q_o) \sim \phi_{conto}(S_o) \quad .$$

Proof: Suppose $\mathcal{F} \approx f(P)$, $\mathcal{F} \in S_o$ and $(a, b) \in Q_o(\mathcal{F}, f, P)$. Then corollary 4.4-2 and the fact that $f(x, y) \neq \omega$ for $(x, y) \in o(a, b)^{(2)}$ imply that f converges over $o(a, b) - \mathcal{M}$. Suppose f converges over the rest of $o(a, b)$. Then there is a $\delta > 0$ such that, for each $\varepsilon \leq \delta$ and every $(x, y) \in \mathcal{O}(\varepsilon; a, b)^{(2)}$, none of F, RF and TF equals ω at $(\varepsilon; x, y)$. This means that $\mathcal{F}$ ε-converges over $\mathcal{O}(\varepsilon; a, b)$ for all $\varepsilon \leq \delta$ and so

$$(4.5\text{-}2) \quad \lim_{\varepsilon \to 0} \Phi_{convo}(\mathcal{F})(\varepsilon; a, b) = (\phi_{convo}(f)(a, b), \omega, \omega) \quad .$$

On the other hand, suppose there is an $x_o \in \mathcal{O}(a, b) \cap \mathcal{M}$ such that f diverges at x_o. For all sufficiently small ε, $x_o \in \mathcal{O}(\varepsilon; a, b)$ and $TF(\varepsilon; x_o, y) = \omega$ for all y. This means that $\mathcal{F}$ ε-diverges in $\mathcal{O}(\varepsilon; a, b)$ for all sufficiently small ε, again implying (4.5-2).

Consider continuity. Corollary 4.4-2 implies that f is continuous over $o(a, b) - \mathcal{M}$. Suppose f is continuous over $o(a, b) \cap \mathcal{M}$ also. By the uniformity assumption, for all $\varepsilon \leq \delta$ we have $\mathcal{F}(\varepsilon; x, x) \neq_\varepsilon \omega$, $\underset{y \to x}{LIM}\, \mathcal{F}(\varepsilon; x, y) \neq_\varepsilon \omega$ and $\mathcal{F}(\varepsilon; x, x) =_\varepsilon \underset{y \to x}{LIM}\, \mathcal{F}(\varepsilon; x, y)$ for all $x \in \mathcal{O}(\varepsilon; a, b)$, i.e., that $\mathcal{F}$ is ε-continuous over $\mathcal{O}(\varepsilon; a, b)$. This implies

$$(4.5\text{-}3) \quad \lim_{\varepsilon \to 0} (\Phi_{conto}(\mathcal{F})(\varepsilon; a, b))_1 = \phi_{conto}(f)(a, b) \quad .$$

On the other hand, suppose there is an $x_o \in o(a, b) \cap \mathcal{M}$ such that f is discontinuous at x_o. This x_o is in $\mathcal{O}(\varepsilon; a, b)$ for all sufficiently small ε. By theorem 4.3-1, $\mathfrak{F}$ is ε-discontinuous at x_o for all sufficiently small ε. Thus $\mathfrak{F}$ is ε-discontinuous in $\mathcal{O}(\varepsilon; a, b)$ for all sufficiently small ε, again yielding (4.5-3). This completes the proof.

It is not difficult to generalize this to half closed and closed intervals, and to a definition of "$\mathfrak{F}$ ε-converges for $x_m \in \mathcal{O}(\varepsilon; a, b)$ at $\bar{x}_{m-1}$" and "$\mathfrak{F}$ is ε-continuous for $x_m \in \mathcal{O}(\varepsilon; a, b)$ at $\bar{x}_{m-1}$" for $\mathfrak{F}$ of $m + 1$ variables.

REMARKS: In the "ε-calculus of stable ε-functions" mentioned in the remarks at the end of chapter 2, an ε-limit ε-operator could be defined by using a particular, reasonable at any $x \neq \omega$, stopping criterion to define

$$\Phi'_{\lim}(F)(\varepsilon;\ x) = F(\varepsilon;\ x,\ y_\varepsilon) \quad ,$$

$$Q'_{\lim}(F,\ f,\ P) \equiv \{x:\ \lim_{y \to x} f(x,\ y) \neq \omega \Rightarrow F \text{ is stable at } x\} \quad .$$

However, it would not be possible to define "ε-comparison" relations, $<_\varepsilon$ and $=_\varepsilon$, satisfying theorem 4.2-1. It would not be possible to define "ε-convergence" for reasons mentioned in section 2.5. Due to the lack of "ε-comparison" and "ε-convergence", it would not be possible to define "ε-continuity" either. A better name for this "ε-calculus" would be "a model of scientific computation" because the model would still be strong enough to do basic computation, but the reliability of results would have to be checked outside of the model, by physical tests or by an error analysis.

It is interesting that theorem 4.4-1 would no longer hold in this model. When "$\mathfrak{F} \approx f(P)$", any $\bar{x}_m \in P$ could be a point of discontinuity of f, provided P does not contain all "real inputs" $\bar{y}_m$ equal in value to $\bar{x}_m$ or P does not contain a neighborhood of $\bar{x}_m$. However, we would have

THEOREM: Suppose $I_1, \ldots, I_m$ $(m \geq 1)$ are intervals contained in R. Let $I \equiv I_1 \times \ldots \times I_m$. Suppose "$\mathfrak{F} \approx f(I)$" and $f(\bar{x}_m)$ is finite for $\bar{x}_m \in I$. Then f is continuous in I.

The second use of I is as a set of m-tuples of "real inputs", under the convention in section 2.3. By continuous in I we mean continuous with respect to limits taken from the interior of I.

Proof: For simplicity, we consider only the case $m = 1$. Let $y_1, y_2, \ldots$ be arbitrary numbers in I approaching $x \in I$. Let z_1 be a "real input" with $z_1 = y_1$. Let η_1 be the largest value of η such that $|Z_1(\varepsilon) - y_1| \leq 1$ for $\varepsilon \leq \eta$ and $|F(\eta; z_1) - f(y_1)| \leq 1$. Suppose F uses only $Z_1(\varepsilon_1), Z_1(\varepsilon_2), \ldots, Z_1(\gamma_1)$ in evaluating $F(\eta_1; z_1)$ (see sec. 2.4). For $i = 2, 3, \ldots$ define z_i, η_i and γ_i as follows. Let z_i be a "real input" with $z_i = y_i$ and $Z_i(\varepsilon) = Z_{i-1}(\varepsilon)$ for $\gamma_{i-1} \leq \varepsilon$. Let η_i be the largest value of η such that $\eta < \gamma_{i-1}$, $|Z_i(\varepsilon) - y_i| \leq i/i$ for $\varepsilon \leq \eta$, and $|F(\eta; z_i) - f(y_i)| \leq 1/i$. Suppose F uses only $Z_i(\varepsilon_1), Z_i(\varepsilon_2), \ldots, Z_i(\gamma_i)$ in evaluating $F(\eta_i; z_i)$. Note that $\gamma_i \leq \eta_i < \gamma_{i-1}$. Define $W: \mathcal{E} \to \mathcal{M}$ by $W(\varepsilon) = Z_i(\varepsilon)$ where i is such that $\eta_{i+1} < \varepsilon \leq \eta_i$ (or, if $\eta_1 \leq \varepsilon$ then $i = 1$). For $\eta_{i+1} < \varepsilon \leq \eta_i$ we have

$$|W(\varepsilon) - x| = |Z_i(\varepsilon) - x| \leq |Z_i(\varepsilon) - y_i| + |y_i - x|$$

$$\leq 1/i + |y_i - x| .$$

As $\varepsilon \to 0$ we have $|y_i - x| \to 0$ and so $|W(\varepsilon) - x| \to 0$. Hence $w \equiv W$ is a "real input" with value x. We have

$$|f(x) - f(y_i)| \leq |f(x) - F(\eta_i; z_i)| + |F(\eta_i; z_i) - f(y_i)|$$

$$\leq |f(x) - F(\eta_i; z_i)| + 1/i .$$

But $Z_i(\varepsilon) = W(\varepsilon)$ for $\varepsilon \geq \gamma_i$, so $F(\eta_i; z_i) = F(\eta_i; w)$, yielding

$$|f(x) - f(y_i)| \leq |f(x) - F(\eta_i; w)| + 1/i \quad .$$

As $i \to \infty$ we have $\eta_i \to 0$, yielding $f(y_i) \to f(x)$. This completes the proof.

Chapter 5: More ε-Operators

Our final tasks are to define ε-derivative, $\frac{D}{Dx}(\mathcal{F})$, ε-integral, $\int Dt(\mathcal{F})$, and to prove the fundamental theorem of the ε-calculus, essentially that

$$\int Dt\left(\frac{D}{Dt}(\mathcal{F})\right)(\varepsilon;\ a,\ b) =_{\varepsilon} \mathcal{F}(\varepsilon;\ b) - \mathcal{F}(\varepsilon;\ a) \quad ,$$

where "−" here denotes ε-subtraction of ε-functions, and is defined below. Our definitions will be based on

$$\frac{d}{dx}f(x) = \lim_{y \to x} (f(x) - f(y))/(x-y) \quad ,$$

$$\int_a^b f(t)dt = \lim_{n \to \infty} \frac{b-a}{n} \sum_{j=1}^{n} f\left(a + j\frac{(b-a)}{n}\right) \quad .$$

For this, we will need ε-operators for <u>ε-arithmetic</u> (Φ_+, Φ_-, ...), <u>ε-limit</u> (Φ_{lim}), <u>ε-composition</u> (Φ^n_{comp}), and <u>ε-recursion</u> (Φ_{rec}). It is interesting to note that all these ε-operators except Φ_{lim} will work in a fixed precision; i.e., when $\Phi(\mathcal{F})(\varepsilon;\ \bar{x}_m)$ is being evaluated, only values of $\mathcal{F}$ at points $(\varepsilon;\ \bar{y}_m)$ are required by Φ. If we were to define Φ_{lim} in terms of the S.C.## of section 3.4, then Φ_{lim} would have this property also. We will also need two input ε-functions, $\mathcal{J}^m_j$ and $\mathcal{C}^m_k$, the identity and the constant ε-functions. As mentioned earlier (see sec. 2.7), we will only be able to give partial definitions of these ε-operators because we have no automatic procedure for generating stably convergent truncation-error bounds; probably no such procedure exists. (However, it may be possible to generate such bounds from a definition of the ideal function expressed

in terms of the operators and initial functions of this chapter; this is a worthwhile research project.) We will assume truncation-error bounds to be given; for completeness, unspecified truncation-error bounds may be taken to be identically ω. The roundoff-error bounds which we generate will be of the per step variety (as seen in interval analysis). Such bounds are notoriously inefficient in real situations. If better bounds are available, they can be used in place of our automatically generated bounds. (It may be possible to automatically improve such automatically generated bounds if we are given a definition of the ideal function under consideration in terms of the operators and initial functions of this chapter.)

For the following sections, we need definition 4.1-1 for TF of $m + 1 \geq 2$ variables (it was stated for TF of 2 variables).

DEFINITION 5-1: Suppose $\mathcal{F} \approx f(P)$. We say TF is stably convergent at $\bar{x}_m$ relative to f $(m \geq 1)$ precisely when

$$\lim_{y \to \bar{x}_m} f(\bar{x}_m, y) \neq \omega \Rightarrow \lim_{\varepsilon \to 0} TF(\varepsilon; \bar{x}_m, y_\varepsilon) = 0 \quad ,$$

as long as the y_ε's are chosen by a reasonable at x_m stopping criterion. We always say TF of one variable are stably convergent at $\bar{x}_0$ relative to f.

We will use the notation,

$$TF \downarrow \bar{x}_m(f) \quad \text{or} \quad TF \downarrow P(f) \quad ,$$

to denote that TF is stably convergent at $\bar{x}_m$, or at all $\bar{x}_m \in P$, relative to f .

In the following, we will need the mapping, $V: \mathcal{P}^{(3)} \to \mathcal{P}^{(2)}$, given by

$$V[(a, b, c)] \equiv (a, b) \quad .$$

5.1 Identity ε-Functions

For $1 \leq j \leq m$ and $x_i \neq \omega$ $(i = 1,\ldots, m)$, define the identity (ideal) functions of m variables by

$$i_j^m(\bar{x}_m) = x_j \quad .$$

For $1 \leq j \leq m$, define

$$\mathcal{I}_j^m(\varepsilon; \bar{x}_m) \equiv (X_j(\varepsilon), RX_j(\varepsilon), TI_j^m(\varepsilon; \bar{x}_m)) \quad ,$$

as long as no $x_i =_\varepsilon \omega$, where TI_j^m is to be defined. Of course TI_1^1 is identically ω. For $1 \leq j < m$, $TI_j^m(\varepsilon; \bar{x}_m) = 0$ so long as no $x_i =_\varepsilon \omega$. For $m \geq 2$, define TI_m^m by

$$TI_m^m(\varepsilon; \bar{x}_m) = (RI_{m-1}^m \hat{+} RI_m^m \hat{+} |I_{m-1}^m \hat{\supset} I_m^m|)(\varepsilon; \bar{x}_m) \quad .$$

THEOREM 5.1-1: $\mathcal{I}_j^m \approx i_j^m(\tilde{R}^{(m)})$ and $TI_j^m \downarrow \tilde{R}^{(m-1)}(i_j^m)$.

Proof: The only thing requiring proof is that $TI_m^m (m \geq 2)$ is stably convergent. But this follows immediately from theorem 2.8-1, completing the proof.

5.2 Constant ε-Functions

For any real input k define the constant k ideal function of m variables ($m \geq 1$) by

$$c_k^m(\bar{x}_m) = k \quad ,$$

so long as no $x_i = \omega$. Define ε-functions C_k^m by

$$C_k^1(\varepsilon;\, x_1) \equiv (I_1^2(\varepsilon;\, k,\, x_1),\, RI_1^2(\varepsilon;\, k,\, x_1),\, \omega) \quad ,$$

$$C_k^m(\varepsilon;\, \bar{x}_m) \equiv \mathcal{J}_1^{m+1}(\varepsilon;\, k,\, \bar{x}_m) \quad \text{for} \quad m \geq 2 \quad .$$

THEOREM 5.2-1: $C_k^m \approx c_k^m(\tilde{R}^{(m)})$ and $TC_k^m \doteq \tilde{R}^{(m-1)}(c_k^m)$.

This follows immediately from theorem 5.1-1.

5.3 ε-Arithmetic

For * being +, -, ×, ÷, define operators $\hat{b}_*$ by

$$\hat{b}_*(f, g)(\bar{x}_m) = f(\bar{x}_m) * g(\bar{x}_m) \qquad (m \geq 1) \ .$$

We define ε-operators corresponding to these by first defining ε-arithmetic for machine numbers and corresponding ε-arithmetic roundoff-error bounds.

ε-Arithmetic subroutines, FL_* for * being +, -, ×, ÷, are subroutines of two variables which approximate ideal arithmetic. Let $N_\varepsilon(a)$ be the ε-neighborhood of $a \in \tilde{R}$, as defined in section 2.8. The FL_* must satisfy

(1) for $x \neq_\varepsilon \omega$ and $y \neq_\varepsilon \omega$, $FL_*(\varepsilon; x, y) = FL_*(\varepsilon; X(\varepsilon), Y(\varepsilon))$, and

(2) for any $a, b \in R(\varepsilon)$, $FL_*(\varepsilon; a, b)$ is in $N_\varepsilon(a * b)$.

For example, the rounding subroutines $A_{n,*}$ of section 2.8 satisfy these. Condition (1) states that the FL_* do not use the inputted error bounds. Condition (2) requires that, when FL_* operates at ε-precision on members of $R(\varepsilon)$, it must get an answer within two machine numbers from the correct answer, unless the correct answer is in $\{-\omega, \infty, \omega\}$, in which case FL_* must get the correct answer.

Define a function, $w: \{(\varepsilon, X): \varepsilon \in \mathcal{E}, X \text{ finite and in } R(\varepsilon)\} \to \mathcal{M}$, by

$$w(\varepsilon, X) = \max(|X - Y_1|, |X - Y_2|) \ ,$$

where Y_1 is the second member of $R(\varepsilon)$ below X (or the first below X if there is only one) and Y_2 is the second member of $R(\varepsilon) - \{\omega\}$

above X (or the first above X if there is only one). For $X \in \mathfrak{R}(\varepsilon)$, let $\widehat{|X|}$ denote $|X \hat{\geq} 0|$ and $\underbrace{|X|}$ denote $|X \check{\geq} 0|$. Define error bounds, RFL_*, for the FL_*, by

$$(5.3\text{-}1) \qquad R_{\pm}(\varepsilon; x, y) = (RX \hat{+} RY)(\varepsilon) ,$$

$$(5.3\text{-}2) \qquad R_{\times}(\varepsilon; x, y) = (\widehat{|X|} \hat{\times} RY \hat{+} \widehat{|Y|} \hat{\times} RX \hat{+} RX \hat{\times} RY)(\varepsilon) ,$$

$$(5.3\text{-}3) \qquad R_{\div}(\varepsilon; x, y) = \begin{cases} \omega & \text{if } (\underbrace{|Y|} \check{\geq} RY)(\varepsilon) \leq 0 , \\ ((RY \hat{\times} \widehat{|X|} \hat{\div} \underbrace{|Y|} \hat{+} RX) \hat{\div} (\underbrace{|Y|} \check{\geq} RY))(\varepsilon) & \text{otherwise,} \end{cases}$$

$$(5.3\text{-}4) \qquad RFL_*(\varepsilon; x, y) = \begin{cases} \omega & \text{if } FL_*(\varepsilon; x, y) = \omega \text{ or } RX(\varepsilon) = \omega \\ & \text{or } RY(\varepsilon) = \omega \text{ or } [FL_*(\varepsilon; x, y) = \pm\omega \\ & \text{and } -\omega <_{\varepsilon}^{x} y <_{\varepsilon} \omega] \\ 0 & \text{if } FL_*(\varepsilon; x, y) = \pm\omega \\ w(\varepsilon, FL_*(\varepsilon, x, y)) \hat{+} R_*(\varepsilon; x, y) & \text{otherwise .} \end{cases}$$

We define ε-arithmetic ε-operators from the FL_* and RFL_* as follows. Let S_{arith} be the set of pairs of ε-functions both of $m \geq 1$ variables. When $\mathfrak{F} \approx f(\mathfrak{P})$, let $f(\bar{x}_m)$ denote (as well as its numeric value) the poor real input, $(F(\cdot; \bar{x}_m), RF(\cdot; \bar{x}_m))$. Suppose $\bar{\mathfrak{F}}_2 \in S_{arith}$ and $\bar{\mathfrak{F}}_2 \approx \bar{f}_2(\bar{\mathfrak{P}}_2)$. Define

$$V[\Phi_*(\bar{\mathfrak{F}}_2)(\varepsilon; \bar{x}_m)] = (FL_*, RFL_*)(\varepsilon; f_1(\bar{x}_m), f_2(\bar{x}_m)) .$$

For $m \geq 2$, we assume the third part of $\mathfrak{P}_*(\overline{\mathfrak{F}}_2)$ to be given. We will abbreviate $\rho_*(f, g)$ by $f*g$ and $\mathfrak{P}_*(\mathfrak{F}, \mathfrak{G})$ by $\mathfrak{F}*\mathfrak{G}$. Also, we let $-\mathfrak{F}$ denote $C_0^m - \mathfrak{F}$. Define

$$Q_{arith}(\overline{\mathfrak{F}}_2, \overline{f}_2, \overline{P}_2) \equiv P_1 \cap P_2 \quad .$$

THEOREM 5.3-1: For * being +, -, ×, ÷, we have

$$(\mathfrak{P}_*, Q_{arith}) \approx \rho_*(S_{arith}) \quad .$$

Proof: It suffices to prove that, for each ε and any real inputs x and y,

$$RFL_*(\varepsilon; x, y) \geq |FL_*(\varepsilon; x, y) - (x * y)| \quad , \tag{5.3-5}$$

and that

$$x * y \neq \omega \Rightarrow \lim_{\varepsilon \to 0} RFL_*(\varepsilon; x, y) = 0 \quad . \tag{5.3-6}$$

Let ε, x and y be given. If $RX(\varepsilon)$ or $RY(\varepsilon)$ is ∞ or $FL_*(\varepsilon; x, y) = \omega$ or $[FL_*(\varepsilon; x, y) = \pm\infty$ and $-\infty <_\varepsilon x <_\varepsilon \infty$ and $-\infty <_\varepsilon y <_\varepsilon \infty]$ then (5.3-5) holds because $RFL_*(\varepsilon; x, y) = \omega$. If $RX(\varepsilon)$ and $RY(\varepsilon)$ are finite, $|x| = \infty$ or $|y| = \infty$, and $FL_*(\varepsilon; x, y) = \pm\infty$ then $FL_*(\varepsilon; x, y) = X(\varepsilon) * Y(\varepsilon) = x * y$ by conditions (1) and (2) on the FL_*; in this case (5.3-5) reduces to $0 \geq 0$, which is true. Suppose $FL_*(\varepsilon; x, y)$, $RX(\varepsilon)$ and $RY(\varepsilon)$ are finite. Then either x, y and $x * y$ are finite or $x * y = x \div \infty = 0$. By the triangle inequality, we have

$$|FL_*(\varepsilon; x, y) - (x * y)| \leq |FL_*(\varepsilon; x, y) - (X(\varepsilon) * Y(\varepsilon))| + |X(\varepsilon) * Y(\varepsilon) - (x * y)| \quad . \tag{5.3-7}$$

By conditions (1) and (2) on FL_* , we have

$$|FL_*(\varepsilon; x, y) - (X(\varepsilon) * Y(\varepsilon))| \leq w(\varepsilon, FL_*(\varepsilon; x, y)) \quad .$$

For the second term on the right side of (5.3-7), we have

$$|X(\varepsilon) \pm Y(\varepsilon) - (x \pm y)| \leq RX(\varepsilon) + RY(\varepsilon) \leq R_{\pm}(\varepsilon; x, y) \quad ,$$

$$|X(\varepsilon) \times Y(\varepsilon) - x \times y| \leq |X(\varepsilon) \times Y(\varepsilon) - X(\varepsilon) \times y| + |X(\varepsilon) \times y - x \times y|$$

$$\leq |X(\varepsilon)| \times RY(\varepsilon) + RX(\varepsilon) \times |y| \leq R_{\times}(\varepsilon; x, y) \quad .$$

If $*$ is $\div$ and $|y| = \infty$ then the above assumptions imply that x is finite and $x \div y = 0$; in this case, $R_{\div}(\varepsilon; x, y)$ is either $\infty \div \infty = \omega$ or it is (some finite number) $\div \infty = 0$, taking the latter value for all sufficiently small ε , so (5.3-5) holds in this case. Suppose y is finite. Then

$$|X(\varepsilon) \div Y(\varepsilon) - x \div y| = |(y \times X(\varepsilon) \div Y(\varepsilon) - x) \div y|$$

$$= |(y - Y(\varepsilon)) \times X(\varepsilon) \div Y(\varepsilon) + X(\varepsilon) - x| \div |y|$$

$$\leq (RY(\varepsilon) \times |X(\varepsilon) \div Y(\varepsilon)| + RX(\varepsilon)) \div |y| \leq R_{\div}(\varepsilon; x, y) \quad .$$

Thus (5.3-5) holds in all cases.

As in theorem 2.8-1, it follows that, for $x * y \neq \omega$,

$$\lim_{\varepsilon \to 0} FL_*(\varepsilon; x, y) = x * y \quad ,$$

(5.3-8)

$$|x * y| = \infty \Rightarrow [FL_*(\varepsilon; x, y) = x * y \text{ for all sufficiently small } \varepsilon] \quad .$$

This implies that, for $x * y$ being finite,

$$\lim_{\varepsilon \to 0} w(\varepsilon, FL_*(\varepsilon; x, y)) = 0 \quad .$$

For such $x * y$, it follows from theorem 2.8-1 that

$$\lim_{\varepsilon \to 0} R_*(\varepsilon; x, y) = 0 \quad ,$$

and (5.3-6) follows. When $|x * y| = \infty$, (5.3-6) follows from (5.3-4) and (5.3-8). This completes the proof.

We say that <u>f is rational</u> if it can be defined from the i_j^m and c_k^m by a finite number of arithmetic operations. We say that <u>$\mathfrak{F}$ is rational</u> if it can be defined from the $\mathcal{I}_j^m$ and $\mathbf{C}_k^m$ by a finite number of ε-arithmetic ε-operations.

<u>COROLLARY 5.3-1</u>: <u>Let $\mathfrak{F}$ be the rational ε-function whose definition corresponds to that of the rational function f . Then</u>

$$\mathfrak{F} \approx f(\tilde{R}^{(m)}) \quad .$$

This follows from theorems 5.1-1, 5.2-1 and 5.3-1 by a simple induction argument, which we omit. For example, this means that $\mathfrak{F} \equiv (\mathcal{I}_1^2 + \mathbf{C}_1^2 - (\mathcal{I}_2^2 + \mathbf{C}_1^2)) \div (\mathcal{I}_1^2 - \mathcal{I}_2^2)$ corresponds to the f of example 3.1-2 over $\tilde{R}^{(2)}$.

5.4 ε-Limit

Here we generalize the definitions of section 4.1. Let S_{lim} be the set of all ideal functions of 2,3,... variables. Define an operator ϕ_{lim} over S_{lim} by

$$\phi_{lim}(f)(\bar{x}_m) = \lim_{y \to x_m} f(\bar{x}_m, y) \quad .$$

Assume a $\lambda(\cdot)$ as in section 4.1 and an effective, reasonable at any $x \neq \omega$, stopping criterion, S.C., have been given. Let S_{lim} be the set of all ε-functions of 2,3,... variables such that if $\mathfrak{F} \in S_{lim}$ then for each $\bar{x}_m \in R^{(m)}$ and each ε

(1) [$Y \in R(\varepsilon)$ and $RF(\varepsilon; \bar{x}_m, Y) \neq \omega$] implies $\lim_{\varepsilon \to 0} RF(\varepsilon; \bar{x}_m, Y) = 0$, and

(2) [$TF(\varepsilon; \bar{x}_m, Y) \neq \omega$ for some $Y \in R(\varepsilon)$] implies

[$\lim_{\varepsilon \to 0} TF(\varepsilon; \bar{x}_m, y_\varepsilon) = 0$, as long as the y_ε's are chosen by a reasonable at x_m stopping criterion].

Define Φ_{lim} by

DEFINITION 5.4-1: Let $\mathfrak{F} \in S_{lim}$, $\bar{x}_m$ and ε be given. Let $y_{\varepsilon_1}, y_{\varepsilon_2}, \ldots$ denote the values selected by S.C. for TF and $\bar{x}_m$. If $TF(\varepsilon; \bar{x}_m, y_\varepsilon) = \omega$, define

(5.4-1) $$V[\Phi_{lim}(\mathfrak{F})(\varepsilon; \bar{x}_m)] \equiv (\omega, \omega) \quad .$$

Otherwise, let δ be the largest member of $\mathcal{E}$ such that $\delta \leq \varepsilon$ and $TF(\delta; \bar{x}_m, y_\delta) \leq \lambda(\varepsilon)$. If $RF(\delta; \bar{x}_m, y_\delta) = \omega$, apply (5.4-1).

Otherwise let j be the smallest integer such that $\varepsilon_j \leq \delta$ and $RF(\varepsilon_j; \bar{x}_m, y_\delta) \leq \lambda(\varepsilon)$. Suppose $F(\varepsilon_j; \bar{x}_m, y_\delta)$ is $\langle \alpha_R(j, k, \cdot) \rangle$. Let $\beta_\varepsilon(\cdot)$ be $\alpha_R(j, k, \cdot)$. Define

$$V[\Phi_{lim}(\mathfrak{F})(\varepsilon;\bar{x}_m)] \equiv (\hat{I}(\varepsilon,\beta_\varepsilon),\ 2\ \hat{x}\ \lambda(\varepsilon)\ \hat{+}\ |\hat{I}(\varepsilon,\beta_\varepsilon)\ \hat{-}\ \check{I}(\varepsilon,\beta_\varepsilon)|) \quad .$$

For $m \geq 2$ we assume the third part of $\Phi_{lim}(\mathfrak{F})$ to be given. Define Q_{lim} by

$$Q_{lim}(\mathfrak{F},f,P) \equiv \{\bar{x}_m : \{\bar{x}_m\} \times (\mathfrak{M} \cap \{\text{some neighborhood of } x_m\} \subset P \text{ and } TF \downarrow \bar{x}_m(f)\} \quad .$$

Theorem 4.1-1 generalizes immediately to

THEOREM 5.4-1: We have

$$(\Phi_{lim}, Q_{lim}) \approx \beta_{lim}(S_{lim}) \quad .$$

5.5 ε-Composition

For $n = 1,2,\ldots,$ let S^n_{comp} be the set of all $(n+1)$ - tuples $(f, \bar{g}_n)$, of ideal functions, where f takes n variables and each g_i takes m variables (some $m \geq 1$). Define a composition operator ϕ^n_{comp} over S^n_{comp} by

$$\phi^n_{comp}(f, \bar{g}_n)(\bar{x}_m) = \overline{f(g_n(\bar{x}_m))} \quad .$$

We abbreviate $\phi^n_{comp}(f, \bar{g}_n)$ by $f(\overline{g_n})$. (In context, it will be clear whether the g_i are functions or variables.) We will also use $\phi_+(f, g)(\bar{g}_n)$ or $(f + g)(\bar{g}_n)$ interchangeably with $\phi^n_{comp}(f + g, \bar{g}_n)$, etc.

For the present, let $g_i(\bar{x}_m)$ denote (together with its numeric value) the poor real input, $(G_i(\cdot;\bar{x}_m)$, $RG_i(\cdot;\bar{x}_m))$. Define

$$V[\Phi^n_{comp}(\mathfrak{F}, \bar{\mathfrak{f}}_n)(\varepsilon;\bar{x}_m)] \equiv V[\mathfrak{F}(\varepsilon; \overline{g_n(\bar{x}_m)})] \quad .$$

Let S^n_{comp} be the set of all $(n + 1)$ - tuples $(\mathfrak{F}, \bar{\mathfrak{f}}_n)$, of ε-functions such that $(\mathfrak{F}, \bar{\mathfrak{f}}_n) \approx (f, \bar{g}_n)(P, \bar{P}_n)$ for some $(f, \bar{g}_n) \in S^n_{comp}$ and some $P, \bar{P}_n$, and such that the computation of $\mathfrak{F}(\varepsilon;\overline{g_n(\bar{x}_m)})$ via the determiners of $\mathfrak{F}$ and $\bar{\mathfrak{f}}_n$ halts for any real inputs $\bar{x}_m$ (see section 2.4). We assume the third part of $\Phi^n_{comp}(\mathfrak{F}, \bar{\mathfrak{f}}_n)$ to be given. We will abbreviate $\Phi^n_{comp}(\mathfrak{F}, \bar{\mathfrak{f}}_n)$ by $\mathfrak{F}(\bar{\mathfrak{f}}_n)$. Define Q^n_{comp} by

$$Q^n_{comp}(\mathfrak{F}, \bar{\mathfrak{f}}_n, f, \bar{g}_n, P, \bar{P}_n) \equiv \{\bar{x}_m : \bar{x}_m \in \bigcap_{i=1}^{n} P_i \text{ and } \overline{g_n(\bar{x}_m)} \in P\}$$

(Note that $\overline{g_n(\bar{x}_m)}$ again denotes poor real inputs, as explained above. And $\overline{g_n(\bar{x}_m)}$ cannot be in P unless each $g_i(\bar{x}_m)$ is a real input.)

THEOREM 5.5-1: We have

$$(\phi^n_{comp}, Q^n_{comp}) \approx \rho^n_{comp}(S^n_{comp}) \quad .$$

5.6 ε-Recursion

We are interested in the following form of recursion. Let g_∞, g_0 and h be given ideal functions of m, m and $m+1$ variables ($m \geq 1$). We define f by recursion as follows:

$$(5.6\text{-}1) \quad f(\bar{x}_m) = \begin{cases} g_\infty(\bar{x}_m) & \text{if } x_m = \infty \\ g_0(\bar{x}_m) & \text{if } x_m < 1 \\ h(\bar{x}_m, f(\bar{x}_{m-1}, x_m - 1)) & \text{otherwise.} \end{cases}$$

We put this in operator form by defining

$$\phi_{rec}(g_\infty, g_0, h) = f \quad ,$$

where f is as in (5.6-1), and defining S_{rec} to be the corresponding set of (g_∞, g_0, h). The following example illustrates the use of this recursion. The operators of this example will be used later.

EXAMPLE 5.6-1: Define the operators σ_+ and $\sigma_\times$ over the S_{lim} of section 5.4 by

$$\sigma_+(g)(\bar{x}_m) = \sum_{i=0}^{[x_m - 1]} g(\bar{x}_m, x_m - i) \quad ,$$

$$\sigma_\times(g)(\bar{x}_m) = \prod_{i=0}^{[x_m - 1]} g(\bar{x}_m, x_m - i) \quad ,$$

where the empty sum is defined to be $c_0^m(\bar{x}_m)$, the empty product is $c_1^m(\bar{x}_m)$, and $\sum_{i=0}^{\infty}$ and $\prod_{i=0}^{\infty}$ are $\equiv \omega$. Let g_+ be c_0^{m+1} and $g_\times$ be c_1^{m+1}. For $*$ being $+$ and $\times$, we have

$$e_*(g) \equiv \Phi_{rec}(\mathcal{C}_\omega^{m+2}, g_*, g(\overline{\mathcal{I}_{m+2}^{m+2}}) * \mathcal{I}_{m+2}^{m+2})(\overline{\mathcal{I}_m^m}, \mathcal{I}_m^m) \quad .$$

Both e_+ and $e_\times$ will be used in section 5.7, where we define an ε-function which corresponds to e^x over $\tilde{R}$.

Let $\mathcal{A}_\infty, \mathcal{A}_0, \mathcal{N}$ be given, and define the F part $\mathcal{F} \equiv \Phi_{rec}(\mathcal{A}_\infty, \mathcal{A}_0, \mathcal{N})$ by letting $\mathcal{A}_{-1}$ denote $\mathcal{N}(\overline{\mathcal{I}_m^m}, \mathcal{F}(\overline{\mathcal{I}_{m-1}^m}, \mathcal{I}_m^m - C_1^m))$ and equating

$$(5.6\text{-}2) \quad F(\varepsilon;\bar{x}_m) = \begin{cases} G_\infty(\varepsilon;\bar{x}_m) & \text{if} \quad \mathcal{I}_m^m(\varepsilon;\bar{x}_m) =_\varepsilon \infty \\ G_0(\varepsilon;\bar{x}_m) & \text{if} \quad \mathcal{I}_m^m(\varepsilon,\bar{x}_m) <_\varepsilon 1 \\ G_1(\varepsilon;\bar{x}_m) & \text{if} \quad (\mathcal{I}_m^m - C_1^m)(\varepsilon;\bar{x}_m) <_\varepsilon \mathcal{I}_m^m(\varepsilon;\bar{x}_m) \\ \omega & \text{otherwise} \quad . \end{cases}$$

The third test is needed because it can happen that $\mathcal{I}_m^m(\varepsilon,\bar{x}_m) <_\varepsilon \infty$ but $(\mathcal{I}_m^m - C_1^m - C_1^m - \ldots - C_1^m)(\varepsilon;\bar{x}_m) \not<_\varepsilon 1$ no matter how many C_1^m's are ε-subtracted from $\mathcal{I}_m^m$. Thus the evaluation of $F(\varepsilon,\bar{x}_m)$ via (5.6-2) with the third test replaced by "otherwise" (and the fourth alternative removed) would not halt for certain ε and $\bar{x}_m$. However, evaluation of $F(\varepsilon;\bar{x}_m)$ via (5.6-2) will always halt.

For the following definition of RF, we will need an ε-comparison operator, $\leq_\varepsilon$. We want $S(\varepsilon;\bar{x}_m) \leq T(\varepsilon;\bar{y}_n)$ to hold when based only on information given by $S(\varepsilon;\bar{x}_m)$ and $T(\varepsilon;\bar{y}_n)$, $s(\bar{x}_m)$ must be $\leq t(\bar{y}_n)$.

DEFINITION 5.6-1. Let x and y be poor real inputs. Define $x \leq_\varepsilon y$ to be true (and $x \not\leq_\varepsilon y$ to be false) precisely when $X(\varepsilon) \leq Y(\varepsilon)$, $RX(\varepsilon)$ and $RY(\varepsilon)$ are finite, and

$$|X(\varepsilon) \dot{-} Y(\varepsilon)| \geq RX(\varepsilon) \hat{+} RY(\varepsilon) \quad .$$

We adopt the convention in section 4.2 concerning the use of triples, $\bar{a}_j$, with $\leq_\varepsilon$.

Define the subroutine ER by

$$(5.6\text{-}3) \qquad ER(\varepsilon;\bar{x}_m) = (RG_0 \hat{+} RG_1 \hat{+} |G_0 \hat{-} G_1|)(\varepsilon;\bar{x}_m) \ .$$

ER bounds the error caused by using G_j in place of g_k for $(j,k) \in \{(0,0), (0,1), (1,0), (1,1)\}$. RF is defined by

$$(5.6\text{-}4) \quad RF(\varepsilon;\bar{x}_m) = \begin{cases} RG_\infty(\varepsilon;\bar{x}_m) & \text{if } \infty \leq_\varepsilon \mathcal{J}_m^m(\varepsilon;\bar{x}_m) \\ RG_0(\varepsilon;\bar{x}_m) & \text{if } \mathcal{J}_m^m(\varepsilon;\bar{x}_m) <_\varepsilon 1 \\ RG_1(\varepsilon;\bar{x}_m) & \text{if } 1 \leq_\varepsilon \mathcal{J}_m^m(\varepsilon;\bar{x}_m) \text{ and} \\ & \quad (\mathcal{J}_m^m - C_1^m)(\varepsilon;\bar{x}_m) <_\varepsilon \mathcal{J}_m^m(\varepsilon,\bar{x}_m) \\ ER(\varepsilon;\bar{x}_m) & \text{if } (\mathcal{J}_m^m - C_1^m)(\varepsilon;\bar{x}_m) <_\varepsilon \mathcal{J}_m^m(\varepsilon;\bar{x}_m) \\ \omega & \text{otherwise .} \end{cases}$$

As usual, we assume TF to be given. Let N_m be given by

$$N_m \equiv \{\bar{x}_m : \ x_m \text{ is a positive integer and no } x_j = \omega\} \ .$$

If $m = 1$, let $S \equiv \{\bar{x}_0\}$; otherwise let S be some subset of $\bar{R}^{(m-1)}$. Let $T \equiv S \times (\bar{R} - \{\infty\})$, where $\{\bar{x}_0\} \times P$ is defined to be P , for any set P .

For $P_1 \equiv S \times (\omega)$ and $P_2 \equiv \Gamma$ define Q_{rec} by

$$Q_{rec}(\mathfrak{F}_3, \bar{g}_3, \bar{P}_3) \equiv \begin{cases} S \times \tilde{R} - N_m & \text{if } P_3 \equiv (T - N_m) \times \tilde{R} \\ S \times \tilde{R} & \text{if } P_3 \equiv T \times \tilde{R} \text{ and } g_0(\bar{x}_{m-1}, 1) = \\ & \quad h(\bar{x}_{m-1}, 1, y) \text{ for any } \bar{x}_{m-1} \in S \text{ and} \\ & \quad \text{any } y \neq \omega \\ \{\,\} & \text{otherwise .} \end{cases}$$

For P_1 and P_2 not of this form, let Q_{rec} be $\{\,\}$.

Let S_{rec} be the set of all $\mathfrak{F}_3 \approx \bar{f}_3(\bar{P}_3)$, for some $\bar{f}_3 \in S_{rec}$ and some $\bar{P}_3$, such that the computation of $\Phi_{rec}(\mathfrak{F}_3)(\varepsilon; \bar{x}_m)$ via the determiners of $\mathfrak{F}_3$ halts for any real inputs $\bar{x}_m$ (see section 5.3).

THEOREM 5.6-1: We have

$$(\Phi_{rec}, Q_{rec}) \approx \beta_{rec}(S_{rec})$$

Proof: Suppose $T \equiv S \times (\tilde{R} - \{\omega\})$, $\mathfrak{F}_\omega \approx g_\omega(S \times \{\omega\})$, $\mathfrak{F}_0 \approx g_0(1)$, $\mathfrak{N} \approx h((T - N_m) \times \tilde{R})$, $(\mathfrak{F}_\omega, \mathfrak{F}_0, \mathfrak{N}) \in S_{rec}$ and $(g_\omega, g_0, h) \in S_{rec}$. Let $f \equiv \beta_{rec}(g_\omega, g_0, h)$ and $\mathfrak{F} \equiv \Phi_{rec}(\mathfrak{F}_\omega, \mathfrak{F}_0, \mathfrak{N})$. Suppose $\bar{x}_m \in S \times \bar{R} - N_m$ and $f(\bar{x}_m) \neq \omega$. If $x_m < 1$ then $f(\bar{x}_m) = g_0(\bar{x}_m)$ and $V[\mathfrak{F}(\varepsilon; \bar{x}_m)] \equiv V[\mathfrak{F}_0(\varepsilon; \bar{x}_m)]$ for all sufficiently small ε, yielding convergence. If $x_m = \omega$ then $f(\bar{x}_m) = g_\omega(\bar{x}_m)$ and $V[\mathfrak{F}(\varepsilon; \bar{x}_m)] \equiv V[\mathfrak{F}_\omega(\varepsilon; \bar{x}_m)]$ for all sufficiently small ε, again yielding convergence. Suppose $1 < x_m < \omega$. Let $\mathfrak{F}_{x_m}$ and $\mathfrak{F}_{x_m}$ be given by

$$\mathcal{S}_{x_m} \equiv \mathcal{S}^m_m - \underbrace{C^m_1 - \ldots - C^m_1}_{[x_m - 1]} \quad ,$$

$$\mathcal{B}_{x_m} \equiv \mathcal{K}(\overline{\mathcal{S}^m_{m-1}},\ \mathcal{S}_{x_m},\ \mathcal{B}_C(\overline{\mathcal{S}^m_{m-1}},\ \mathcal{S}_{x_m} - C^m_1)) \quad .$$

For all sufficiently small ε we have

$$\mathcal{S}_{x_m}(\varepsilon;\bar{x}_m) \not<_\varepsilon 1 \ , \quad (\mathcal{S}_{x_m} - C^m_1)(\varepsilon;\bar{x}_m) <_\varepsilon 1 \ ,$$

and so

$$(5.6\text{-}5) \quad V[\mathcal{F}(\varepsilon;\bar{x}_m)] \equiv V[\mathcal{K}(\mathcal{S}^m_m,\ \mathcal{K}(\mathcal{S}^m_{m-1},\ \mathcal{S}^m_m - C^m_1,\ \ldots\ \mathcal{B}_{x_m}\ \ldots))(\varepsilon;\bar{x}_m)] \quad .$$

We also have

$$(5.6\text{-}6) \quad f(\bar{x}_m) = h(\bar{x}_m,\ h(\bar{x}_{m-1},\ x_m - 1,\ \ldots\ h(\bar{x}_{m-1},\ x_m - [x_m - 1],\ g_0(\bar{x}_{m-1},\ x_m - [x_m])) \ldots)) \quad .$$

Successive applications of corollary 5.3-1 and theorem 5.5-1, working from the inside to the outside on (5.6-5), give us convergence.

In addition to the above hypotheses, suppose $\mathcal{K} \approx h(T \times \tilde{R})$ and $g_0(\bar{x}_{m-1},1) = h(\bar{x}_{m-1},1,y)$ for any $\bar{x}_{m-1} \in S$ and any $y \neq \omega$. From the above, we have $\mathcal{F} \approx f(S \times \tilde{R} - N_m)$. Suppose $\bar{x}_m \in S \times \tilde{R} \cap N_m$ and $f(\bar{x}_m) \neq \omega$. Define $\mathcal{B}'_{x_m}$ by

$$\mathcal{B}'_{x_m}(\varepsilon;\bar{y}_m) \equiv (G_{x_m}(\varepsilon;\bar{y}_m),\ ER(\varepsilon;\bar{y}_{m-1},z),\omega) \quad \text{for all } \varepsilon \text{ and } \bar{y}_m \ ,$$

where z is the real input $(I_{x_m}(\cdot;\bar{y}_m), RI_{x_m}(\cdot;\bar{y}_m))$. (When $\bar{y}_m$ is $\bar{x}_m$, the value of z is 1.) Equations (5.6-5) and (5.6-6) become

(5.6-7) $V[\mathfrak{F}(\varepsilon;\bar{x}_m)] \equiv V[\mathfrak{K}(\mathfrak{J}_m^{\bar{m}}, \mathfrak{K}(\mathfrak{J}_{m-1}^{\bar{m}}, \mathfrak{J}_m^m - C_1^m, \ldots \mathfrak{L}'_{x_m} \ldots))(\varepsilon;\bar{x}_m)]$.

(5.6-8) $f(\bar{x}_m) = h(\bar{x}_m, h(\bar{x}_{m-1}, x_m\text{-}1, \ldots h(\bar{x}_{m-1}, 1, g_0(\bar{x}_{m-1}, 0))\ldots))$.

If $x_m = 1$, these last equations are

$$V[\mathfrak{F}(\varepsilon;\bar{x}_m)] \equiv V[\mathfrak{L}'_{x_m}(\varepsilon;\bar{x}_m)] \quad ,$$

$$f(\bar{x}_{m-1}, 1) = h(\bar{x}_{m-1}, 1, g_0(\bar{x}_{m-1}, 0)) \quad ,$$

and z is just x_m . In general, we have

$$ER(\varepsilon;\bar{x}_{m-1}, z) = (RG_0 \hat{+} RG_1 \hat{+} |G_0 \hat{\supset} G_1|)(\varepsilon;\bar{x}_{m-1},z) \quad ,$$

$$G_0(\varepsilon;\bar{x}_{m-1}, z) = (\mathfrak{L}_0(\mathfrak{J}_{m-1}^{\bar{m}}, \mathfrak{J}_{x_m})(\varepsilon;\bar{x}_m))_1 \quad ,$$

$$G_1(\varepsilon;\bar{x}_{m-1}, z) = G_{x_m}(\varepsilon;\bar{x}_m) \quad ,$$

$$RG_1(\varepsilon;\bar{x}_{m-1}, z) = RG_{x_m}(\varepsilon;\bar{x}_m) \quad ,$$

$$g_0(\bar{x}_{m-1}, z) = g_0(\bar{x}_{m-1}, 1) \quad ,$$

$$g_1(\bar{x}_{m-1}, z) = h(\bar{x}_{m-1}, 1, g_0(\bar{x}_{m-1}, 0)) = g_0(\bar{x}_{m-1}, 1) \quad ,$$

the last equality following from our additional hypotheses. Further, by theorem 5.5-1 and corollary 5.3-1, we have

$$\lim_{\varepsilon\to 0} G_0(\varepsilon;\bar{x}_{m-1}, z) = g_0(\bar{x}_{m-1}, 1) \quad ,$$

$$\lim_{\varepsilon\to 0} G_1(\varepsilon;\bar{x}_{m-1}, z) = g_0(\bar{x}_{m-1}, 1) \quad ,$$

$$\lim_{\varepsilon\to 0} RG_0(\varepsilon;\bar{x}_{m-1},z) = \lim_{\varepsilon\to 0} RG_1(\varepsilon;\bar{x}_{m-1},z) = 0 \quad .$$

By this and theorem 2.8-1, we have

$$\lim_{\varepsilon\to 0} ER(\varepsilon;\bar{x}_{m-1},z) = 0 \quad .$$

From this and (5.6-7) and (5.6-8), it follows that

$$\lim_{\varepsilon\to 0} V[\mathfrak{F}(\varepsilon;\bar{x}_m)] \equiv (f(\bar{x}_m), 0) \quad .$$

This completes the proof.

Let $\mathscr{J}_+$ be C_0^{m+1} and $\mathscr{J}_\times$ be C_1^{m+1}. Let $\mathscr{J}$ be an ε-function of $m+1$ variables. For $*$ being $+$ and $\times$, define

$$\Sigma_*(\mathscr{J}) \equiv \Phi_{rec}(C_\omega^{m+1}, \mathscr{J}_*, \mathscr{J}(\mathscr{J}_{m+1}^{\overline{m+2}}) * \mathscr{J}_{m+2}^{m+2})(\bar{\mathscr{J}}_m^m, \mathscr{J}_m^m) \quad ,$$

$$Q_{sum}(\mathscr{J},g,P) \equiv \begin{cases} S \times \tilde{R} - N_m & \text{if } P \equiv S \times (\tilde{R} - N_1 - \{\omega\})^{(2)} \\ \{\,\} & \text{otherwise .} \end{cases}$$

Let S_{sum}^* be the set of all $\mathscr{J}$ of $m+1 = 2,3,\ldots$ variables such that the computation of $\Sigma_*(\mathscr{J})(\varepsilon;\bar{x}_m)$ via the determiner of $\mathscr{J}$ halts for any real inputs $\bar{x}_m$.

COROLLARY 5.6-1: For * being + and x , we have

$$(\Sigma_*, Q_{sum}) \approx \sigma_*(S^*_{sum}) \quad .$$

Proof: Suppose $\mathcal{L} \approx g(S \times (\tilde{R} - N_1 - \{\omega\})^{(2)})$. Then

$$\mathcal{L}(\overline{\mathcal{J}^{m+2}_{m+1}}) * \mathcal{J}^{m+2}_{m+2} \approx g(\overline{i^{m+2}_{m+1}}) * i^{m+2}_{m+2}(S \times (\tilde{R} - N_1 - \{\omega\})^{(2)} \times \tilde{R})$$ and

theorems 5.6-1 and 5.5-1 yield $\Sigma_*(\mathcal{L}) \approx \sigma_*(g)(S \times \tilde{R} - N_m)$. This completes the proof.

5.7 An ε-Function Corresponding to e^x Over $\mathfrak{R}$

First we mention that $e^{\omega} = \omega$, $e^{-\infty} = 0$ and $e^{\infty} = \infty$. Let $f_{exp}(x) = e^x$ as in section 2.6. Define two basic ε-functions, $\mathcal{P}$ and $\mathcal{Q}$, by

$$\mathcal{P} \equiv \Sigma_+(\mathcal{C}_1^2) \quad , \tag{5.7-1}$$

$$\mathcal{Q} \equiv \mathcal{E}_{rec}(\mathcal{C}_\infty^1, \mathcal{C}_1^1 - \mathcal{I}_1^1, \mathcal{I}_1^2 - \mathcal{C}_1^2)(\mathcal{I}_1^1 + \mathcal{C}_1^1) \quad . \tag{5.7-2}$$

By corollary 5.6-1, $\mathcal{P} \approx \sigma_+(c_1^2)(\mathfrak{R} - N_1)$; $p \equiv \sigma_+(c_1^2)$ is essentially an entier for positive reals, because

$$p(x) = \begin{cases} \omega & \text{if } x = \omega, \infty \\ 0 & \text{if } x < 1 \\ [x] & \text{otherwise} \end{cases} \quad . \tag{5.7-3}$$

By theorem 5.6-1, $\mathcal{Q} \approx a(\mathfrak{R})$, where $a(x) = |x|$. We use $\mathcal{P}$ and $\mathcal{Q}$ to define $\mathcal{F}_{exp}$ as follows:

$$\mathcal{T} \equiv \Sigma_x(\mathcal{I}_1^4 + \mathcal{P}(\mathcal{I}_4^4)) \quad , \tag{5.7-4}$$

$$\mathcal{F} \equiv \Sigma_+(\mathcal{T}(\mathcal{I}_1^3, \mathcal{I}_2^3, \mathcal{I}_3^3 - \mathcal{C}_1^3))(-\mathcal{Q}(\mathcal{I}_1^3), \mathcal{I}_3^3) \quad , \tag{5.7-5}$$

$$\mathcal{L} \equiv \mathcal{E}_{rec}(\mathcal{C}_0^1, \mathcal{E}_{lim}(\mathcal{F})(\mathcal{I}_1^1, \mathcal{C}_\infty^1), \mathcal{E}_{lim}(\mathcal{F})(\mathcal{I}_1^2, \mathcal{C}_\infty^2)) \quad , \tag{5.7-6}$$

$$\mathcal{F}_{exp} \equiv \mathcal{E}_{rec}(\mathcal{C}_\infty^1, \mathcal{L}(\mathcal{C}_1^1 - \mathcal{I}_1^1), \mathcal{C}_1^2 + \mathcal{L}(\mathcal{I}_1^2 - \mathcal{C}_1^2))(\mathcal{I}_1^1 + \mathcal{C}_1^1) \quad . \tag{5.7-7}$$

$\mathcal{T}$ forms terms, $x^{n-1}/(n-1)!$. $\mathcal{F}$ forms sums, $\sum_{n=1}^{N} (-|x|)^{n-1}/(n-1)!$.

$\mathcal{J}$ checks for $x = -\infty$, to insure that $\mathcal{F}_{exp}(\varepsilon;-\infty) = 0$, and otherwise it approximates $\sum_{n=0}^{\infty}(-|x|)^n/n!$. $\mathcal{F}_{exp}$ checks for $x = \infty$, to insure that $\mathcal{F}_{exp}(\varepsilon;\infty) = \infty$, and otherwise it computes $\mathcal{J}(\varepsilon;x)$ and reciprocates this value when $x \geq 0$. The only essential part of this definition that is missing is a definition of the TF part of $\mathcal{F}$, because TF is the only truncation-error bound used here. In the following discussion, we will define TF and use our previous theorems to prove that $\mathcal{F}_{exp} \approx f_{exp}(\tilde{R})$. The only nontrivial part of this development is the definition of a stably convergent TF.

Theorems 5.1-1, 5.5-1 and 5.3-1 yield that $\mathcal{I}_1^4 \div \rho(\mathcal{I}_4^4) \approx i_1^4 \div p(i_4^4)(\tilde{R}^{(4)} - N_4)$. This and corollary 5.6-1 imply that $\mathcal{T} \approx t(\tilde{R}^{(3)} - N_3)$, where

$$t(x,m,n) = \begin{cases} \omega & \text{if } x = \omega, m = \omega, \text{ or } n \geq \infty \\ 1 & \text{if } n < 1 \\ x^{[n]}/[n]! & \text{otherwise.} \end{cases}$$

Thus $\mathcal{F} \approx f(\tilde{R}^{(3)} - N_3)$, where f is defined in section 2.6. For the next step, we need the TF part of $\mathcal{F}$. We could define TF from $\mathcal{TF}' \equiv \mathcal{T}(-\alpha(\mathcal{I}_1^3), \mathcal{I}_2^3, \mathcal{I}_3^3)$ without making use of any special properties of $(R, \mathcal{E})$, but a considerable derivation is required to insure that the resulting TF is stably convergent at all (x, ∞) with x finite. For simplicity, we instead sketch a definition of TF for the case where $(R, \mathcal{E})$ is a floating-point number system and ε-arithmetic satisfies the usual relations needed for an error analysis in the style of Wilkinson [W2]. We work from the tf of section 2.6. Let a finite x be given and let

$$Y \equiv Y(\varepsilon) = (|X(\varepsilon)| \hat{+} RX(\varepsilon)) \hat{\div} ((1 \hat{-} \varepsilon) \hat{\times} (1 \hat{-} \varepsilon) \hat{\times} (1 \hat{-} \varepsilon)) \quad .$$

For integers $n > Y$ let $fl_\varepsilon(Y^n/n!)$ denote the product $Y \times (Y/2) \times (Y/3) \times \ldots \times (Y/m)$ evaluated in floating-point ε-arithmetic, where m is the largest integer with $Y < m \leq n$ such that overflow and underflow do not occur (or m is ω, if there is no such integer). Using this m instead of n is two-thirds of the trick needed to define TF for an arbitrary (R, ε). Assume that the value I used in place of i in (Y/i) satisfies $I = i \times (1+\eta_i(\varepsilon))$, where $|\eta_i(\varepsilon)| \leq \varepsilon$. Then for some $|\eta'_j(\varepsilon)| \leq \varepsilon$ $(j=1,\ldots, 3m)$, we have

$$\begin{aligned} fl_\varepsilon(Y^n/n!) &= (Y^m/m!) \times \prod_{j=1}^{3m} (1 + \eta'_j(\varepsilon)) \\ &\geq ((|X(\varepsilon)| + RX(\varepsilon))^m/m!) \times \prod_{j=1}^{3m} \frac{(1+\eta'_j(\varepsilon))}{(1-\varepsilon)} \\ &\geq |x|^m/m! \quad . \end{aligned}$$

Further, we have

$$\lim_{\varepsilon \to 0} fl_\varepsilon(Y^n/n!) = |x|^n/n! \quad ,$$

so we define

$$TF(\varepsilon;x,k,y) = \begin{cases} fl_\varepsilon(Y^n/n!) & \text{if } k = \infty \text{ and } |x| + 1 \leq_\varepsilon y \leq_\varepsilon \infty \\ \omega & \text{otherwise} \quad , \end{cases}$$

where n is the largest integer $\leq_\varepsilon y$. We prove that TF is stably convergent at (x, ∞) for finite x as follows. Let $y_{\varepsilon_1}, y_{\varepsilon_2}, \ldots$

satisfy $y_\varepsilon \in R(\varepsilon) - \{\infty\}$ and $\lim_{\varepsilon\to 0} y_\varepsilon = \infty$. For some $|\eta_j''(\varepsilon)| \le 2\varepsilon$ we have

$$Y(\varepsilon) = (|X(\varepsilon)| + RX(\varepsilon))(1-\varepsilon)^{-3} \times \prod_{j=1}^{7} (1 + \eta''(\varepsilon)) \quad ,$$

$$TF(\varepsilon;x,\infty,y_\varepsilon) \le ((|X(\varepsilon)| + RX(\varepsilon))(\tfrac{1+\varepsilon}{1-\varepsilon})^3(1 + 2\varepsilon)^7)^m/m! \quad ,$$

for all sufficiently small ε. As $\varepsilon \to 0$, $m \to \infty$ and the right side above goes to 0. Thus TF is stably convergent at all (x, ∞) with x finite.

Thus $\mathfrak{F} \equiv (F, RF, TF)$ is in S_{lim} and $\Phi_{lim}(\mathfrak{F})(\mathcal{I}_1^m, \mathcal{I}_\infty^m) \approx \Phi_{lim}(f)(i_1^m, c_\infty^m)(R \times \tilde{R}^{(m-1)})$. This and theorem 5.6-1 yield $\mathcal{I} \approx g(\tilde{R})$, where g is defined in section 2.6. This implies that $\mathfrak{F}_{exp} \approx f_{exp}(\tilde{R})$.

REMARKS: It is easy to define initial ε-functions and ε-operators analogous to those of this chapter for the "ε-calculus of stable ε-functions" discussed at the end of chapters 2 and 4. However, our example in section 5.7 would have to be changed, because the subtractions in $\Sigma(-|x|)^n/n!$ makes F unstable at (x,∞). This can be remedied by defining F in terms of $\Sigma|x|^n/n!$.

Chapter 6. ε-Derivative and ε-Integral

6.1 ε-Differentiability and ε-Derivative

Define a difference operator, d, over the set S_d of ideal functions of one variable by

$$d(f) \equiv (f(i_1^2) - f(i_2^2)) \div (i_1^2 - i_2^2) .$$

We say f is differentiable at x precisely when d(f) converges at x . Otherwise, we say f is nondifferentiable at x . Define a difference ε-operator, (D, Q_D), over the set S_D of ε-functions of one variable by

$$D(\mathfrak{F}) \equiv (\mathfrak{F}(\mathfrak{i}_1^2) - \mathfrak{F}(\mathfrak{i}_2^2)) \div (\mathfrak{i}_1^2 - \mathfrak{i}_2^2) ,$$
$$Q_D(\mathfrak{F}, f, P) \equiv P \times P .$$

By theorems 5.1-1, 5.5-1 and corollary 5.3-1, we have

$$(D, Q_D) \approx d(S_D) .$$

DEFINITION 6.1-1: We say $\mathfrak{F}$ is ε-differentiable at x precisely when $D(\mathfrak{F})$ ε-converges at x . Otherwise, we say $\mathfrak{F}$ is ε-nondifferentiable at x .

Our previous analysis of ε-convergence at x carries over immediately to ε-differentiability at x , so we will not bother to express this in operator, ε-operator form.

Define a derivative operator, $\frac{d}{dt}$, over S_d by

$$\frac{d}{dt}(f) \equiv \phi_{lim}(d(f)) .$$

Let $S_{d/dt}$ be the set of all ε-functions, $\mathcal{F}$, of one variable such that $D(\mathcal{F}) \in S_{lim}$. Define an ε-derivative ε-operator, $(\frac{D}{Dt}, Q_{d/dt})$, by

$$\frac{D}{Dt}(\mathcal{F}) \equiv \Phi_{lim}(D(\mathcal{F})),$$

$$Q_{d/dt}(\mathcal{F}, f, P) \equiv Q_{lim}(D(\mathcal{F}), d(f), P \times P).$$

We call $\frac{D}{Dt}(\mathcal{F})(\varepsilon; x)$ <u>the ε-derivative of $\mathcal{F}$ at x</u>. By theorem 4.1-1 we have

$$(\frac{D}{Dt}, Q_{d/dt}) \approx (S_{d/dt}).$$

Thus, under the usual conditions, the ε-derivative at x of an ε-function approaches the derivative at x of its corresponding ideal function as $\varepsilon \to 0$.

It deserves mention here that there is a function, f, such that

(1) $f(x)$ is finite for all $x \in R$,

(2) there is a $\mathcal{F} \approx f(\tilde{R})$, and

(3) f is nondifferentiable at every point in R.

See Grzegorczyk [G1, pp. 199-201].

BLANK PAGE

6-2 ε-Integrability and ε-Integral

Let C denote a finite closed interval of numbers. For P being a set of real inputs, we say P covers C almost everywhere precisely when C has a subset C' of Lebesgue measure zero such that for any $c \in C - C'$ there is an $x \in P$ with $x = c$. As is usual, we write "a.e." for "almost everywhere." Let "over C" be implicit in the statements "f is continuous a.e.", bounded or integrable. From analysis, we know that the bounded, Rieman integrable functions are precisely those that are bounded and continuous a.e. (See Royden [R3, p. 70].)

Suppose f is a bounded ideal function and that $\mathcal{F} \approx f(P)$. In order for the information contained in the set $\{\mathcal{F}(\varepsilon; x): \varepsilon \in \mathcal{E}, x \in C \cap P\}$ to determine whether f is continuous a.e., P will have to cover C a.e. (Remember that $\mathcal{F}$ gives no information because all of $\mathcal{F}$'s of one variable are $= \omega$.) Similarly, the set,

$$\{f: \mathcal{F} \approx f(P) \text{ and } f \text{ is bounded}\},$$

will contain both integrable and nonintegrable ideal functions unless P covers C a.e.

Now, suppose we have a definition of "$\mathcal{F}$ is ε-integrable (over C)" which is based only on the values of $\mathcal{F}$. Then, the weakest requirement on the size of P which might make the conditions,

(1) $\mathcal{F} \approx f(P)$,

(2) f is bounded, and

(3) $\mathcal{F}$ is ε-integrable for all sufficiently small ε,

equivalent ot [f is integrable] is

(4) P covers C a.e.

However, by theorem 4.4-1, we know that conditions 1, 2 and 4 by themselves inply that f is integrable, i.e., that f is continuous everywhere in C except possibly at points in $\mathcal{M} \cap C$, a set of measure zero. Hence essentially the only definition of "$\mathcal{F}$ is ε-integrable", which uses only the values of $\mathcal{F}$ and for which we have

$$(1) - (4) \text{ hold } \Leftrightarrow f \text{ is integrable,}$$

is "$\mathcal{F}$ is ε-integrable precisely when $1 = 1$."

Consider basing our definition on the values of $\mathcal{F}' \equiv \mathcal{F}(\delta_2^2)$, an ε-function which may have a worthwhile truncation-error bound. (This is reasonable because it is only by a fluke of notation that $\mathcal{F}$ of one variable have no worthwhile truncation-error bound.) Then we have the additional information given by TF', which satisfies

$$TF'(\varepsilon; x, Y) \geq \left| f(Y) - \lim_{y \to x} f(y) \right| .$$

Again, let us assume that f is bounded. Let z be a poor real input such that, for each ε, $|Z(\varepsilon) - c| \leq RZ(\varepsilon)$ for all $c \in C$. If $TF'(\varepsilon; z, z) < \omega$ then we will know that $\lim_{y \to c} f(y)$ exists for all $c \in C$. And this implies that f is integrable, by

THEOREM 6.2-1: Suppose f is bounded over C and $\lim_{y \to c} f(y)$ exists for all $c \in C$. Then f has at most a countable number of discontinuities in C.

So far as we know, this is a new result.

Proof: (This was proved independently by Bill Glassmire and Paul Rosenthal.) Define

$$g(c) = \lim_{y \to c} f(y) \qquad \text{for } c \in C .$$

First we prove that g is continuous. Suppose $y_i \to c$ as $i \to \infty$. For each i there is an $x_i \neq c$ with

$$|g(y_i) - f(x_i)| < 1/i, \quad |y_i - x_i| < 1/i ,$$

because $g(y_i) = \lim_{x \to y_i} f(x)$. Thus

$$\lim_{i \to \infty} x_i = c, \quad \lim_{i \to \infty} g(y_i) = g(c) .$$

Since y_i was an arbitrary approach, this means that g is continuous in C. Next we prove that $b \neq f$ only on a countable set. Suppose not. We have

$$\{x: \ g(x) > f(x)\} \equiv \bigcup_{n=1}^{\infty} \{x: \ g(x) > f(x) + 1/n\} .$$

If this set is uncountable then at least one of the sets on the right is uncountable; suppose $E_n \equiv \{x: \ g(x) > f(x) + 1/n\}$ is. Then its members have a cluster point, x_0, and there is a sequence, $x_1, x_2, \ldots$, from E_n and approaching x_0 such that

$$g(x_m) > f(x_m) + 1/n \qquad \text{for } m = 1, 2, \ldots,$$

$$\lim_{m \to \infty} g(x_m) = g(x_0) \geq g(x_0) + 1/n ,$$

a contradiction. Similarly, the set $\{x: \ g(x) < f(x)\}$ is countable. Therefore, the set $\{x: \ g(x) \neq f(x)\}$ is countable. This completes the proof.

Thus we can define ε-integrability as follows. For finite a and b, let c[a, b] denote the closed interval between a and b. For given, finite real inputs a and b, let $z \equiv z(a, b)$ be a poor real input such that

(1) for each ε, $|Z(\varepsilon) - x| \leq RZ(\varepsilon)$ for all $x \in c[a, b]$,

(2) $\bigcap_{i \geq 1} \{x: |Z(\varepsilon_i) - x| \leq RZ(\varepsilon_i)\} \equiv c[a, b]$, and

(3) $Z(\cdot)$ and $RZ(\cdot)$ are effectively computable from a and b.

It is easy to verify that such z exist.

DEFINITION 6.2-1: Suppose $\mathfrak{F} \approx f(P)$ for some P. Let $\mathfrak{F}' \equiv \mathfrak{F}(\mathcal{S}_2^2)$. Let $z \equiv z(a, b)$ be as above. For finite a and b, we say TF' is good relative to f, a and b precisely when [f is integrable over c[a, b]] $\Rightarrow$ [$TF'(\varepsilon; z, z) < \omega$ for all sufficiently small ε].

Let TF' and $z \equiv z(a, b)$ be as above. We put these results in operator, ε-operator form by defining

$$S_{int} \equiv \{\text{ideal functions of one variable, bounded over } o(-\infty, \infty)\},$$

$$\phi_{int}(f)(a, b) = \text{bool}[f \text{ is integrable over } c[a, b]],$$

provided a and b are finite, and

$$S_{int}\{\mathfrak{F}: \mathfrak{F} \approx f(P) \text{ for some } f \in S_{int} \text{ and some } P,$$
$$\text{and computation of } TF'(\varepsilon; x, x) \text{ via the determiner}$$
$$\text{of TF' halts for any poor real input x (see sec. 2.4)}\},$$

$\Phi_{int}(\mathcal{F})(\varepsilon;\ a,\ b) \equiv (\text{bool}[TF'(\varepsilon;\ x,\ z) < \omega],\ \text{bool}[TF'(\varepsilon;\ z,\ z) = \omega], \omega)$,

provided $-\infty <_{\varepsilon}^{a} b <_{\varepsilon} \infty$, and

$$Q_{int}(\mathcal{F},\ f,\ P) \equiv \{(a,\ b):\ TF' \text{ is good relative to } f,\ a \text{ and } b\} .$$

An immediate consequence of the above analysis is

THEOREM 6.2-2: We have

$$(\Phi_{int},\ Q_{int}) \sim \phi_{int}(S_{int}) .$$

Further, if f is integrable over $c[a,\ b]$ for all $(a,\ b) \in Q_{int}(\mathcal{F},\ f,\ P)$, then $\Phi_{int}(\mathcal{F})$ is not weak, and vice versa.

Now for the integral. Using the notation of chapter 5, we define a partial sum and an integral operator over S_{int} by

$$h^m \equiv (i_2^m - i_1^m) \div p(i_4^m) \qquad (m = 4,\ 5) ,$$

$$t \equiv i_1^5 + p(i_5^5) \times h^5 ,$$

$$\phi_{psum}(f) \equiv h^4 \times \sigma_+(f(t)) ,$$

$$\int dt(f)(a,\ b) = \begin{cases} \phi_{lim}(\phi_{psum}(f))(i_1^2,\ i_2^2,\ c_\infty^2)(a,\ b) \\ \qquad \text{if } \phi_{int}(f)(a,\ b) = 1 \\ \omega \qquad \text{otherwise} . \end{cases}$$

We define ε-operators for these by

$$\aleph^m \equiv (\mathcal{S}_2^m - \mathcal{S}_1^m) + \mathcal{P}(\mathcal{S}_4^m) \qquad (m = 4, 5) \ ,$$

$$\mathcal{T} \equiv \mathcal{S}_1^5 + \mathcal{P}(\mathcal{S}_5^5) \times \aleph^5 \ ,$$

$$\Phi_{psum}(\mathfrak{F}) \equiv \aleph^4 \times \Sigma_+(\mathfrak{F}(\mathcal{T})) \ ,$$

$$Q_{psum}(\mathfrak{F}, f, P) \equiv Q_{sum}(\mathfrak{F}(\mathcal{T}), f(t), Q^1_{comp}(\mathfrak{F}, \mathcal{T}, f, t, P,$$

$$\widetilde{R}^{(3)} \times (\widetilde{R} - N_1)^{(2)})) - N_4 \ ,$$

$$\int Dt(\mathfrak{F})(\varepsilon; a, b) = \begin{cases} \Phi_{lim}(\Phi_{psum}(\mathfrak{F}))(\mathcal{S}_1^2, \mathcal{S}_2^2, C_\infty^2)(\varepsilon; a, b) \\ \qquad \text{if} \quad \Phi_{int}(\mathfrak{F})(\varepsilon; a, b) \equiv (1, 0, \omega) \\ (\omega, \omega, \omega) \qquad \text{otherwise} \end{cases} ,$$

$$Q_{integral}(\mathfrak{F}, f, P) \equiv Q_{int}(\mathfrak{F}, f, P) \cap$$

$$Q^3_{comp}(\Phi_{lim}(\Phi_{psum}(\mathfrak{F})), \mathcal{S}_1^2, \mathcal{S}_2^2, C_\infty^2, \ldots) \ ,$$

$$S_{integral} \equiv \{\mathfrak{F}: \ \mathfrak{F} \in S_{int} \ \text{ and } \ \Phi_{psum}(\mathfrak{F}) \in S_{lim}\} \ .$$

The theorems of chapter 5 and theorem 6.2-2 yield

THEOREM 6.2-3: We have

$$(\int Dt, Q_{integral}) \sim \int dt(S_{integral}) \ .$$

Further, $\int Dt(\mathfrak{F})$ is weak if and only if $\Phi_{int}(\mathfrak{F})$ is weak.

6.3 The Fundamental Theorem of the ε-Calculus

Following is the ε-calculus analog to the fundamental theorem of the calculus.

THEOREM 6.3-1: Fix $a, b \in R$ and assume that

(1) $\frac{d}{dt}(f)$ is bounded and integrable over $c[a, b]$, and

(2) $\mathfrak{F}_{\approx} f(\{\})$.

Then, for any ε we have

$$\int Dt\left(\frac{D}{Dt}(\mathfrak{F})\right)(\varepsilon; a, b) =_{\varepsilon} (\mathfrak{F}(\mathscr{S}_2^2) - \mathfrak{F}(\mathscr{S}_1^2))(\varepsilon; a, b) \quad . \tag{6.3-1}$$

Proof: For any $a_1, a_2 \geq 0$, $a_3 \in R(\varepsilon)$ let $\rho(\bar{a}_3)$ be

$$\rho(\bar{a}_3) \equiv \{x: \ |a_1 - x| \leq a_2\} \quad .$$

For any ε-function $\mathscr{A}$ of m variables and any $(\varepsilon; \bar{x}_m)$ we have $[\mathscr{A} \approx g(\{\})$ implies $g(\bar{x}_m) \in \rho(\mathscr{A}(\varepsilon; \bar{x}_m))]$. This means that

$$A = \int dt\left(\frac{d}{dt}(f)\right)(a, b) \in \rho\left(\int Dt\left(\frac{D}{Dt}(\mathfrak{F})\right)(\varepsilon; a, b)\right) \quad ,$$

$$B = f(b) - f(a) \in \rho((\mathfrak{F}(\mathscr{S}_2^2) - \mathfrak{F}(\mathscr{S}_1^2))(\varepsilon; a, b)) \quad .$$

The fundamental theorem of the calculus tells us that $A = B$, yielding (6.3-1). This completes the proof.

BLANK PAGE

Chapter 7: Computable Real Functions and Completeness

7.1 Computable Real Functions

We say that a set of real inputs P covers $\tilde{R}^{(m)}$ $(m \geq 1)$ precisely when each value in $\tilde{R}^{(m)}$ is taken on by some member of P. We say P covers $\tilde{R}^{(0)}$ precisely when $P \equiv \{\bar{x}_0\}$. Let K_1 be the class of all ideal functions, f, such that there is a P covering $\tilde{R}^{(m)}$ $(m \geq 0)$ and an $\mathfrak{F}$ with $\mathfrak{F} \approx f(P)$. We say f is computable$_1$ precisely when $f \in K_1$. K_1 depends on $(R, \mathcal{C})$ and it contains many functions with discontinuities. We will not consider K_1 further.

By specialization of an ideal function f of $m \geq 1$ variables, let us mean the replacement of a variable by a numeric constant, yielding an ideal function of m-1 variables. Let K_2' be the class of all ideal functions f such that there is an $\mathfrak{F} \approx f(\tilde{R}^{(m)})$ and F and RF are subroutines of m variables and no constants. Let K_2 be the smallest class of ideal functions containing K_2' and closed under specialization. We say f is computable$_2$ precisely when $f \in K_2$. If $f \in K_2$ then there is an $\mathfrak{F} \approx f(\tilde{R}^{(m)})$ such that F and RF are subroutines of m variables and $n \geq 0$ constants. (We do not know whether the reverse is true.) As we shall see, K_2 is independent of $(R, \mathcal{C})$. By theorem 4.4-1, we know that any $f \in K_2$ is continuous at all $\bar{x}_m \in R^{(m)}$ with $f(\bar{x}_m) \neq \omega$.

Let θ be as in section 2.2 and η as in 1.5. Let us say $\alpha_1, \alpha_2 : \eta \to \eta$ give $a \in \tilde{R}$ precisely when

(1) for each $n \geq 1$, either $\alpha_1(n) = \alpha_2(n) = 3$ or

$$|a - \theta(\alpha_1(n), \alpha_2(n))/n| \leq 1/n \quad ,$$

(2) if $a \neq \omega$ then, for all sufficiently large n, either

$$\alpha_1(n) \neq 3 \quad \text{or} \quad \alpha_2(n) \neq 3 \ .$$

For $m \geq 1$, we say $\overline{\alpha_{2m}}$ give $\bar{x}_m$ precisely when α_{2i-1}, α_{2i} give x_i $(i = 1,2,\ldots, m)$. Also, we say $\bar{\alpha}_0$ give $\bar{x}_0$. We say recursive operators ψ_1, ψ_2 give f precisely when for any $\overline{\alpha_{2m}}$ and $\bar{x}_m \in \tilde{R}^{(m)}$,

$$[\overline{\alpha_{2m}} \text{ give } \bar{x}_m] \Rightarrow [\psi_1(\overline{\alpha_{2m}}), \psi_2(\overline{\alpha_{2m}}) \text{ give } f(\bar{x}_m)] \ .$$

Let K_3' be the class of all ideal functions f for which there exist recursive operators ψ_1, ψ_2 which give f. Let K_3 be the smallest class of ideal functions which contains K_3' and which is closed under specialization. We say f is computable$_3$ precisely when $f \in K_3$. This is analogous to Grzegorczyk's definition of computable continuous real functions of one variable [G2]. K_3 obviously does not depend on $(\mathcal{R}, \mathcal{E})$ and we have

THEOREM 7.1-1: $K_2 \equiv K_3$.

Proof: We prove this by proving $K_2' \equiv K_3'$. Our proof is based on two transformation functions, t_1 and t_2. Suppose α_1, α_2 give $a \in \tilde{R}$ and define the poor real input $t_1(\alpha_1, \alpha_2) \equiv x$ by

$$\alpha(\cdot) \equiv \theta(\alpha_1(\cdot), \alpha_2(\cdot)) \ ,$$

$$\chi(\varepsilon) = \hat{I}(\varepsilon, \alpha) \ ,$$

$$RX(\varepsilon) = \begin{cases} \omega & \text{if } X(\varepsilon) = \omega \ , \\ |\hat{I}(\varepsilon, \alpha) \ominus \check{I}(\varepsilon, \alpha)| & \text{otherwise} \ . \end{cases}$$

If $a \neq \omega$ then this x is a real input. Let y be a poor real input and define $t_2(y) \equiv (\beta_1, \beta_2)$ as follows. Let an $n \geq 1$ be given. If $RY(\varepsilon_n) = \omega$ or $\limsup_{\varepsilon \to 0} RY(\varepsilon) > 0$ then define $\beta_1(n) = \beta_2(n) = 3$. Otherwise let $j(n)$ be the smallest value of j such that $RY(\varepsilon_j) \leq \check{I}(\varepsilon_j, \gamma_{3n})$, where $\gamma_{3n}(k) = [k/3n]$ so that $<\gamma_{3n}> = 1/3n$. Then

$$|Y(\varepsilon_{j(n)}) - y| \leq 1/3n \ .$$

Suppose $Y(\varepsilon_{j(n)})$ is $<\alpha_R(j(n), k(n), \cdot)>$ and let $\delta_n(\cdot)$ be $\alpha_R(j(n), k(n), \cdot)$. If $<\delta_n> = -\infty, \infty$ or ω then $<\delta_n> = y$; in this case, define $\beta_1(n) = \beta_2(n) = 1$, 2 or 3 respectively. Suppose $<\delta_n>$ is finite. For any integer i, define

$$r(i) = \begin{cases} i/3 & \text{if } i = 3[i/3] \\ (i-1)/3 & \text{if } i = 3[i/3] + 1 \\ (i+1)/3 & \text{otherwise} \ . \end{cases}$$

Note that $r(i)$ is always an integer. Define

$$\beta_1(n) = |r(\delta_n(3n))| \ ,$$

$$\beta_2(n) = |r(\delta_n(3n))| - r(\delta_n(3n)) \ .$$

In this case we have

$$\beta_1(n) - \beta_2(n) = r(\delta_n(3n)) \quad ,$$

$$|y - r(\delta_n(3n))/n|$$

$$\leq |y - Y(\varepsilon_{j(n)})| + |Y(\varepsilon_{j(n)}) - \delta_n(3n)/3n| + 1/3n \leq 1/n \quad .$$

Thus if y is a real input then β_1, β_2 give y. Further, if y is a real input or if $RY(\cdot) = \omega$ then β_1 and β_2 are computable from (Y, RY).

We prove $K_2' \subset K_3'$ as follows. Suppose $f \in K_2'$. Then there is an $\mathfrak{F} \approx f(\widetilde{R}^{(m)})$ such that F and RF are subroutines of no constants (see sec. 2.4). We will construct recursive operators ψ_1, ψ_2 which give f. Suppose $\overline{\alpha_{2m}}$ give $\bar{x}_m$. Define poor real inputs $\bar{y}_m$ by

$$y_i = t_1(\alpha_{2i-1}, \alpha_{2i}) \quad .$$

Let $f(\bar{y}_m)$ denote the poor real input $(F(\cdot; \bar{y}_m), RF(\cdot; \bar{y}_m))$ and define ψ_1 and ψ_2 by

$$(\psi_1(\overline{\alpha_{2m}}), \psi_2(\overline{\alpha_{2m}})) = t_2(f(\bar{y}_m)) \quad .$$

Then ψ_1 and ψ_2 are recursive operators and they give f. Thus $f \in K_3'$.

We prove $K_3' \subset K_2'$ as follows. Suppose $f \in K_3'$. Then there are recursive operators ψ_1, ψ_2 giving f. We will define an $\mathfrak{F} \approx f(\widetilde{R}^{(m)})$ such that F and RF are subroutines of no constants. Suppose $\bar{x}_m$ are poor real inputs. Define $\overline{\alpha_{2m}}$ by

$$(\alpha_{2i-1}, \alpha_{2i}) = t_2(x_i) \quad .$$

Define F and RF by

$$(F(\cdot\,;\,\overline{x}_m),\ RF(\cdot\,;\,\overline{x}_m)) \equiv t_1(\psi_1(\overline{\alpha_{2m}}),\ \psi_2(\overline{\alpha_{2m}}))\ ,$$

and set $TF \equiv \omega$. Then $\mathfrak{F} \equiv (F, RF, TF)$ is the desired $\mathfrak{E}$-function. This completes the proof.

Let K be the class of all $f: R \to R$ such that there is an $f' \in K_3'$ (or K_2') with $f(x) = f'(x)$ for all $x \in R$. K is precisely Grzegorczyk's class of computable continuous real functions [G2]. In [G2] Grzegorczyk proves K to be equivalent to several other classes of computable real functions which have appeared in the literature. In [G2, p. 192] he proves that the $f \in K$ are computably uniformly continuous in any segment. He also constructs an $f \in K$ which is not differentiable at any point [G2, p. 199].

7.2 Completeness

Let K^* be the class of all $f: R^{(m)} \to R$ $(m \geq 0)$ such that there is an $f' \in K_3$ with $f'(\bar{x}_m) = f(\bar{x}_m)$ for all $\bar{x}_m \in R^{(m)}$. Let K^{**} be the class of all $f: R^{(m)} \to R$ $(m \geq 0)$ which can be defined exclusively in terms of the c^n_k, i^n_j, ϕ_+, $\phi_\times$, $\phi_\div$, ϕ_{lim}, ϕ^n_{comp}, ϕ_{rec} from chapter 5.

THEOREM 7.2-1: $K^* \subset K^{**}$.

In this sense, the initial functions and operators of chapter 5 are complete.

Proof: We only sketch the proof. The Stone-Wierstrauss theorem (see Bishop [B1, pp. 97, 100]) shows that we can construct an arbitrarily close (in the sup norm) polynomial approximation to a continuous function over a compact set if we are given

(1) access to any finite number of (arbitrarily close approximations to) values of the function over the compact set, and

(2) the modulus of continuity of f.

Grzegorczyk [G1, p. 192] has shown that every $f \in K^*$ of one variable has a computable modulus of continuity, and his proof generalizes to f of any number of variables. Thus any $f \in K^*$ can be written as a polynomial in m variables:

$$f(\bar{x}_m) = \lim_{n \to \infty} \sum_{\bar{j}_m} c(n, p_1^{j_1} \times \ldots \times p_m^{j_m}) \times x_1^{j_1} \times \ldots \times x_m^{j_m} ,$$

where p_i denotes the i^{th} prime number, $c(n, k)$ is the k^{th} coefficient of the n^{th} polynomial, the sum is taken over all $\bar{j}_m$

such that $0 \leq j_i \leq n$ $(i = 1,2,\ldots, m)$, and where the n^{th} polynomial approximates f over the m-dimensional square, $[-n, n]^{(m)}$, with a maximum error less than $1/n$. Further, we can assume that each $c(n, k)$ is rational. Thus, in order to show that $K^* \subset K^{**}$ we need only show that K^{**} includes all computable rational functions, $c(n, k)$. But, since division is one of the closure operations of K^{**}, we need only show that K^{**} contains all recursive rational functions, $b(n, k)$. It is obvious that the initial functions and operations, except the effective minimum, used to define the recursive functions in [M1, p. 120-1] can be simulated by the operators and initial functions of K^{**}. That the effective minimum operator can also be simulated in this way follow from the following:

$$t^n \equiv i_m^n - i_{m+1}^n - c_1^n + f(i_{m-1}^n, i_m^n - c_1^n) \quad \text{for} \quad n = m+1, m+2 \quad ,$$

$$h \equiv \phi_{rec}(c_\omega^{m+1}, i_m^{m+1}, t^{m+2})(\overline{i_m^{m+1}}, t^{m+1}) \quad ,$$

$$g \equiv \phi_{rec}(c_\omega^m, c_0^m, h)(\overline{i_{m-1}^m}, i_m^m + c_1^m) \quad ,$$

$$\mu y(f(\overline{i_{m-1}^m}, y) \neq 0) \equiv \phi_{lim}(g(\overline{i_{m-1}^{m+1}}, i_{m+1}^{m+1}))(\overline{i_{m-1}^m}, c_\omega^m) \quad .$$

This deserves some explanation. For $i > j \geq 0$, if $i = j+1$ and $f(\bar{x}_{m-1}, j) = 0$ then $h(\bar{x}_{m-1}, i, j) = i$, or otherwise $h(\bar{x}_{m-1}, i, j) = j$. If $f(\bar{x}_{m-1}, k) = 0$ for $k = 0, 1, \ldots, \ell$ then $g(\bar{x}_{m-1}, \ell) = \ell+1$; otherwise, $g(\bar{x}_{m-1}, \ell)$ is the least value of n such that $f(\bar{x}_{m-1}, n) \neq 0$. This completes the proof.

BLANK PAGE

Summary and Conclusions

We have developed a theory of numerical computation based on recursive function theory, with a flavor of interval analysis. This theory concerns itself with a general class of variable-precision computations and the finite-precision (or intermediate) results arising in such computations. For example, the floating-point computations of modern digital computers are in this class. Our main goal was to form a realistic model of such computations. This was done by developing the concepts of

(1) a machine number system $(\mathcal{R}, \mathcal{E})$ (sec. 2.2),

(2) a real input $x \equiv (X, RX)$ (sec. 2.3),

(3) a subroutine F (sec. 2.4),

(4) an ε-function $\mathcal{F} \equiv (F, RF, TF)$ (sec. 2.5), and

(5) ε-arithmetic (sec. 2.8 and 5.3).

If this model had been our only goal, we would probably have dispensed with roundoff-error and truncation-error bounds (the RX, RF and TF indicated above) because such bounds are usually not computed on the computer. (We discuss the removal of these bounds in the remarks at the end of chapters 2, 4 and 5.) However, our secondary goal necessitated the incorporation of these bounds. This secondary goal was to find out how concepts from the calculus such as convergence, continuity, differentiability and integrability apply, at each fixed level of precision, to numerically computed functions which, after all, can be viewed at a fixed precision as a discrete set of points on a graph. This secondary goal was achieved by associating the numerically computed function, F, with its underlying mathematical (or ideal) function, f, through the

use of roundoff-error bounds, RF, and truncation-error bounds, TF. Thus we defined an ε-function $\mathfrak{F}$ to be a triple (F, RF, TF).

In trying to apply convergence and continuity to ε-functions, we were lead to an investigation of stopping criteria and stability (ch. 3). Out of this came a new and simple definition of stability, the concept of an ε-wave, and a proof that instability can be overcome, given the requisite error bounds.

As presented in chapters 2 and 3, the concepts of subroutine ε-function and stability are machine dependent because they are defined in terms of a fixed machine number system. In the remarks at the end of chapter 3, we show how these concepts can be made machine independent.

The part of the ε-calculus dealing with notions from the calculus is of definitional interest only. For example, one may have wondered whether there is a definition of ε-continuity which satisfies the following: for each fixed precision ε, many numerically computed functions which look possibly continuous at a point x, but whose corresponding ideal function is discontinuous at x, may be accepted as ε-continuous at x; but, as $\varepsilon \to 0$ these functions should be weeded out as ε-discontinuous at x. We found (in sec. 4.3) that it is possible to form such a definition by making use of computable information about (i.e., bounds on) truncation and roundoff errors. We do not expect such definitions to be of practical importance.

On the other hand, the part of the ε-calculus which models scientific computation should have practical implications. Our work on stopping criteria and stability tends in this direction. But we as yet have no concrete applications.

REFERENCES

[A1] Aberth, Oliver. Analysis in the computable number field. J. Assoc. Comput. Mach. V 15, No. 2 (April 1968) 275-299.

[B1] Bishop, Errett. Foundations of Constructive Analysis. McGraw-Hill. New York. 1967.

[G1] Grzegorczyk, Andrzej. Computable functionals. Fund. Math. 42 (1955) 168-202.

[G2] ———————. On definitions of computable real continuous functions. Fund. Math. 44 (1957) 61-71.

[G3] ———————. Some approaches to constructive analysis. Constructivity in Mathematics. North-Holland. Amsterdam (1959) 43-61.

[K1] Klaua, Dieter. Konstructive Analysis. (German) Deutscher Verlag der Wissenschaften. Berlin. 1961.

[K2] Kreisel, Georg. Interpretations of analysis by means of functionals of finite types. Constructivity in Mathematics. North-Holland. Amsterdam (1959) 101-128.

[M1] Mazur, S. Computable Analysis. Panstwowe Wydawnictwo Naukowe. Warszawa. 1963.

[M2] Mendelson, Elliott. Introduction to Mathematical Logic. Van Nostrand. Princeton, N. J. 1964.

[M3] Moore, Ramon E. Interval Analysis. Printice-Hall. Englewood Cliffs, N. J. 1966.

[M4] ———————. The automatic analysis and control of errors in digital computing based on the use of interval numbers. Error in Digital Computation. John Wiley and Sons. New York. V 1 (1964) 61-130.

[N-Z] Niven, Ivan M., Zuckerman, Herbert S. An Introduction to the Theory of Numbers. John Wiley and Sons. New York. 1966.

[R1] Ralston, Anthony. A First Course in Numerical Analysis. McGraw-Hill. New York. 1965.

[R2] Riesel, Hans. A case of numerical divergence. BIT (1961) 130-1.

[R3] Royden, H. L. Real Analysis. Macmillan. New York. 1964.

[S1] Scott, Dana. Some definitional suggestions for automata theory. Journal of Computer and System Sciences. V 1, No. 2 (August 1967) 187-212.

[S2] Shoenfield, Joseph R. Mathematical Logic. Addison-Wesley. Reading, Mass. 1967.

[W1] van Wijngaarden, A. Numerical analysis as an independent science BIT 6 (1966) 66-81.

[W2] Wilkinson, J. H. Rounding Errors in Algebraic Processes. Prentice-Hall. Englewood Cliffs, N. J. 1963.

Unclassified

Security Classification

DOCUMENT CONTROL DATA - R & D

(Security classification of title, body of abstract and indexing annotation must be entered when the overall report is classified)

1. ORIGINATING ACTIVITY (Corporate author) Computer Science Department Stanford University Stanford, California 94305	2a. REPORT SECURITY CLASSIFICATION Unclassified 2b. GROUP -L-

3. REPORT TITLE

ε-CALCULUS

4. DESCRIPTIVE NOTES (Type of report and inclusive dates)

Manuscript for Publication (Technical Report)

5. AUTHOR(S) (First name, middle initial, last name)

Richman, Paul L.

6. REPORT DATE August 16, 1968	7a. TOTAL NO. OF PAGES 138	7b. NO. OF REFS 19
8a. CONTRACT OR GRANT NO. N00014-67-A-0112-0029 b. PROJECT NO. c. d.	9a. ORIGINATOR'S REPORT NUMBER(S) CS 105 9b. OTHER REPORT NO(S) (Any other numbers that may be assigned this report) none	

10. DISTRIBUTION STATEMENT

Releasable without limitations on dissemination.

11. SUPPLEMENTARY NOTES ---	12. SPONSORING MILITARY ACTIVITY Office of Naval Research

13. ABSTRACT

We use recursive function theory to lay the basis for a partially constructive theory of calculus, which we call the ε-calculus. This theory differs from other theories that have grown out of recursive function theory in that

(1) it is directly related to the variable-precision computations used in scientific computation today, and

(2) it deals explicitly with intermediate results rather than ideal answers.

As $\varepsilon \to 0$, intermediate results in the ε-calculus approach their corresponding answers in the calculus. Thus we say "the ε-calculus approaches the calculus, as $\varepsilon \to 0$." It is hoped that investigations in the ε-calculus will lead to a better understanding of numerical analysis. Several new results in this direction are presented, concerning instability and also machine numbers. Discrete notions of limit, convergence, continuity, arithmetic, derivative and integral are also presented and analyzed.

DD FORM 1 NOV 68 1473 (PAGE 1)

S/N 0101-807-6801

Unclassified

Security Classification

Unclassified

Security Classification

14 KEY WORDS	LINK A		LINK B		LINK C	
	ROLE	WT	ROLE	WT	ROLE	WT
recursive						
computable						
numerical analysis						
automatic analysis						

DD FORM NOV 1473 (BACK)

(PAGE 2)

Unclassified

Security Classification