

International Standard



7846

INTERNATIONAL ORGANIZATION FOR STANDARDIZATION • МЕЖДУНАРОДНАЯ ОРГАНИЗАЦИЯ ПО СТАНДАРТИЗАЦИИ • ORGANISATION INTERNATIONALE DE NORMALISATION

Industrial real-time FORTRAN — Application for the control of industrial processes

Langage FORTRAN en temps réel industriel — Application pour la commande des processus industriels

First edition — 1985-09-15

STANDARDSISO.COM : Click to view the full PDF of ISO 7846:1985

UDC 681.3.06 : 800.92

Ref. No. ISO 7846-1985 (E)

Descriptors : data processing, programming languages, fortran.

Price based on 32 pages

Foreword

ISO (the International Organization for Standardization) is a worldwide federation of national standards bodies (ISO member bodies). The work of preparing International Standards is normally carried out through ISO technical committees. Each member body interested in a subject for which a technical committee has been established has the right to be represented on that committee. International organizations, governmental and non-governmental, in liaison with ISO, also take part in the work.

Draft International Standards adopted by the technical committees are circulated to the member bodies for approval before their acceptance as International Standards by the ISO Council. They are approved in accordance with ISO procedures requiring at least 75 % approval by the member bodies voting.

International Standard ISO 7846 was prepared by Technical Committee ISO/TC 97, *Information processing systems*.

Users should note that all International Standards undergo revision from time to time and that any reference made herein to any other International Standard implies its latest edition, unless otherwise stated.

Contents

	Page
0 Introduction	1
1 Scope and field of application	1
2 Definitions	1
Section one: Multiprogramming and real-time features	3
3 Introduction	3
4 Date and time information	3
4.1 Obtain date and time	DATIM 3
4.2 Obtain clock counts	CLOCK 3
5 General aspects of tasking	3
5.1 States and transitions	3
5.2 Multiple activation calls	5
5.3 Synchronization concepts	5
5.3.1 Eventmarks	5
5.3.2 Resourcemarks	5
5.3.3 Semaphores	6
6 Procedure references	6
6.1 Terms and summary of procedure references	6
6.2 Creation of a new task	CREATE 8
6.3 Eliminating a task from the real-time system	KILL 8
6.4 Scheduling a task	SKED 8
6.5 Starting a task by simplified calls	9
6.5.1 Starting a task immediately	STRT 9
6.5.2 Starting a task after a specified time delay	STRTAF 9
6.5.3 Starting a task at a specified absolute time	STRTAT 10
6.5.4 Starting a task in repeated execution	10

6.5.5	Initial start immediately	CYCL	10
6.5.6	Initial start after a specified time delay	CYCLAF	10
6.5.7	Initial start at specified absolute time	CYCLAT	10
6.5.8	Connection of a task to an event	CON	10
6.6	Elimination of previous scheduling	DSKED	11
6.6.1	Elimination of event connections	DCON	11
6.6.2	Elimination of time connections	CANCEL	11
6.7	Delaying continuation of a task	SUSPND	11
6.8	Suspending until event or time		12
6.8.1	Delay until event has occurred	HOLD	12
6.8.2	Delay for a specified relative time	DELAY	12
6.9	Eventmark operations		12
6.9.1	Setting an eventmark to the ON condition	POST	12
6.9.2	Clearing an eventmark	CLEAR	12
6.9.3	Testing an eventmark condition	TESTEM	12
6.9.4	Masking an eventmark	MKEM	13
6.9.5	Unmasking an eventmark	UNMKEM	13
6.10	Resourcemark operations		13
6.10.1	Setting a resourcemark to the locked condition	LOCK	13
6.10.2	Setting a resourcemark to the unlocked condition	UNLOCK	13
6.10.3	Testing and setting a resourcemark to the locked condition	TLOCK	13
6.11	Semaphore operations		13
6.11.1	Initialization of semaphore	PRESEM	14
6.11.2	Wait on semaphore	WAITS	14
6.11.3	Release of semaphore	SIGNAL	14
6.11.4	Reading a semaphore value	IRDSEM	14
6.12	Normal termination of execution	EXIT	15
Section two: Binary pattern and bit processing			16
7	Introduction		16
8	Binary pattern processing		16
8.1	Boolean operations		16
8.1.1	Inclusive OR	IOR	16

8.1.2	Boolean AND	IAND	16
8.1.3	Boolean complement	NOT	16
8.1.4	Exclusive OR	IEOR	16
8.2	Shift operations		16
8.2.1	Logical shift	ISHL	16
8.2.2	Arithmetic shift	ISHA	17
8.2.3	Circular shift	ISHC	17
9	Bit processing		17
9.1	Bit testing	BTEST	17
9.2	Set bit	IBSET	17
9.3	Clear bit	IBCLR	17
9.4	Change bit	IBCHNG	17
Section three: Process input/output			18
10	Introduction		18
11	Scope of the process I/O and general structure of I/O routines		18
12	Input/output of analog values		18
12.1	Sequential analog data input	AISQW	18
12.2	Analog data input in random sequence	AIRDW	19
12.3	Analog data output	AOW	19
13	Input/output of digital values		19
13.1	Digital input	DIW	19
13.2	Digital output		19
13.2.1	Digital pulse output	DOMW	19
13.2.2	Latched digital output	DOLW	19
Section four: File handling			20
14	Introduction		20
15	Background information		20
16	File system environment		20
17	Procedures to control file access		21
17.1	Introduction		21
17.2	Creation of files	CFILW	21
17.3	Deletion of files	DFILW	21

17.4	Opening files	OPENW	21
17.5	Closing files	CLOSEW	22
17.6	Modify access mode	MODAPW	22
Annexes			
A	Historical background		23
B	Further descriptions pertaining to individual clauses		26
C	Abortion of tasks		30
D	File handling		31
	Bibliography		32

STANDARDSISO.COM : Click to view the full PDF of ISO 7846:1985

Industrial real-time FORTRAN — Application for the control of industrial processes

0 Introduction

This International Standard describes a task model and a set of related routines to allow control of multi-task systems using FORTRAN as the programming language. Information concerning the development of this International Standard is given in annex A, and suggestions and justifications for some features are given in annex B. Annex C deals with the problem of aborting tasks and annex D with file handling. These annexes do not form part of this International Standard.

1 Scope and field of application

This International Standard establishes external procedure references for use in industrial computer control systems. These external procedure references provide access to time and date information, permit interface of programs with executive systems, process input and output functions, allow manipulation of bit strings and provide a method for file handling.

These procedures are intended for use with programs written in FORTRAN in accordance with ISO 1539.¹⁾ These programs are expected to be executable both in a solitary and in a multiprogramming environment under the control of a real-time executive system.

It is the responsibility of the executive system to deal with exceptions or errors, such as division by zero or referencing $\text{SQRT}(-1.0)$, so that such errors do not cause serious effects, for example unwanted or uncontrolled abortion, or affect other tasks.

This International Standard is applicable to all FORTRAN systems which require multi-tasking features.

NOTE — No extensions or variations from this International Standard should be implemented except for features explicitly declared in this International Standard to be processor dependent.

2 Definitions

For the purpose of this International Standard, the following definitions apply.

NOTE — Terms printed in italics in the definitions are themselves defined in this clause. However, in each definition, such items are only printed in italic at their first occurrence.

2.1 access mode: The right or permission to access (read or write) a *file* granted by the *processor* following a request for such permission.

2.2 basic clock counts: Counts of the basic unit of the system real-time clock as available for the user's program.

2.3 computation: A set of *operations* carried out on a set of data, as available for the user's program.

2.4 critical region: A part of a *sequential order of operations* operating on shared data such that this part of the program must have exclusive access to the shared data during the *execution*.

2.5 DORMANT: One of the definite states of a *task*.

A dormant task is known to the *executive system* and is not in any of the states *PENDING*, *RUNNING*, or *SUSPENDED*.

2.6 event: A significant discrete occurrence or incident which is intended to affect the *execution* of some *task* in a planned manner.

An event itself occurs instantaneously and sets an *eventmark*¹⁾.

2.7 eventmark: Internal variable of the *executive system*, used to indicate that an *event* has occurred.

If the user's system contains *parallel tasks*, the eventmarks are shared data elements for the *tasks*.¹⁾

2.8 execution: The collection of actions performed by a computer *processor* carrying out instructions in a sequential manner.

2.9 executable program: A program including all of its functions and subroutines in a form suitable for *execution*.

1) See also 5.3.1.

2.10 executive routine; executive system: The part of the *processor* which supports the procedures described in this International Standard.

2.11 file: A collection of related records treated as a unit.

For the purpose of this International Standard, the records are considered to be of fixed length. Record storage and access are independent of the internal format of records.

2.12 initiation: The action taken by the *executive system* to start the *execution* of a *task* at its first executable statement.

2.13 multiprocessing¹⁾: A mode of *operation* that provides for parallel processing by two or more *processors* of a multiprocessor.

2.14 multiprogramming¹⁾: A mode of *operation* that provides for the interleaved *execution* of two or more computer programs by a single processor.

2.15 multitasking¹⁾: A multiple *operation* that provides for the concurrent performance or interleaved *execution* of two or more *tasks*.

2.16 NON-EXISTENT: One of the definite (formal) states of a *task*.

A non-existent task is unknown to the *executive system*.

2.17 object task; designated task; referenced task: The *task* that is wanted or expected to be started, halted, stopped, or otherwise affected as a consequence of a system subroutine call.

2.18 operation: A deterministic rule for the generation of a finite set of data from another finite set of data.

2.19 overrun: Occurs when the condition for *initiation* of a *task* becomes true while the task is still running because of a previous *initiation*.

2.20 parallel tasks; concurrent tasks: A set of *tasks* whose *operations* may overlap in *time*.

2.21 PENDING: One of the definite states of a *task*.

A pending task has been associated with an *event* or *time* condition such that when the condition occurs, the task will be transferred to state *RUNNING* and initiated.

2.22 processor: The combination of a data processing system and the mechanism by which programs are transformed for use on that data processing system.

2.23 processor dependent: Describes an action of the *processor* that is not specified in this International Standard.

2.24 repetitive execution: Occurs when a *task* is repeatedly initiated, whether at fixed intervals or by repetitive *events*.

2.25 resource mark: Internal variable of the *executive system*, used to indicate that a resource is exclusively reserved for a *task*.²⁾

2.26 RUNNING: One of the definite states of a *task*.

A running task is executing in its *virtual processor*.

2.27 semaphore: A variable of the *executive system*, used for the exchange of synchronizing information between interacting parallel *tasks*.

All semaphore *operations* in this International Standard imply *critical region* protection, provided by the *executive system*.

2.28 sequential order of operations: An order of *operations* the results of which are as if the operations are performed strictly one after another.

2.29 SUSPENDED: One of the definite states of a *task*.

A suspended task has temporarily halted the *execution* of its *virtual processor*, and is waiting for a specified condition to continue the execution of its virtual processor.

2.30 task: A *computation* which can be scheduled.

The *operations* of this computation are performed in a strict *sequential order of operations*.

(See also 2.17 and 2.28.)

2.31 time:

a) **absolute time:** Complete time and date specification.

b) **relative time:** A time increment or difference.

2.32 virtual processor: An environment in which a *task* can run from the time it is initiated until it terminates without consideration of availability of resources, managed by the *processor* and not by the user's program.

A particular implementation serves to map a set of virtual processors onto a set of real processor(s). This mapping is *processor dependent*.

1) Definition taken from ISO 2382/10.^[2]

2) See also 5.3.2.

Section one: Multiprogramming and real-time features

3 Introduction

This section describes several procedure references available for the user's program, and relating to multiprogramming and, in particular, real-time operation. For all calls given in this section, the operations are generally considered indivisible, i.e. their operation will behave as if they are not interrupted.

4 Date and time information

For programming in a real-time environment, the user must have access to the time variables of the executive system. These time variables are obtained by system calls as described below.

Unambiguous time specification requires unique designation of time including complete date and an acknowledged calendar, defining time zero. Execution of reference to subroutine DATIM provides this complete information. The date refers to the Gregorian calendar.

The calls are

CALL DATIM(t1) for obtaining current date and time,

CALL CLOCK(j,k1,k2) for obtaining the basic clock counts.

4.1 Obtain date and time

The form of the call is

CALL DATIM(t1)

where t1 designates an integer array, into whose first 8 elements will be placed the absolute time, as expressed by the executive system's real-time clock at the time when the call is executed. These elements are as follows:

- first element: Counts of the basic clock
- second element: Milliseconds (0 to 999)
- third element: Seconds (0 to 59)
- fourth element: Minutes (0 to 59)
- fifth element: Hours (0 to 23)
- sixth element: Day (1 to 31)
- seventh element: Month (1 to 12)
- eighth element: Year

4.2 Obtain clock counts¹⁾

Execution of a reference to this subroutine allows the user program to obtain the current value of the system real-time clock, expressed in basic clock counts.

The form of the call is

CALL CLOCK(j,k1,k2)

where

j designates an integer variable or integer array element into which the current value of the clock will be placed as a positive integer. j is counted up to a maximum value as given by k2, then set to zero and counted again;

k1 is the number of basic clock counts per second. It is an integer value returned by the system. This argument shall be an integer variable or an integer array element;

k2 is the maximum value j can attain. It is an integer value returned by the system. This argument shall be an integer variable or an integer array element.

The modulus size of j is given by argument k2 (modulus = k2 + 1). In practice, this will usually be equal to the modulus of the hardware counter realizing the clock.

5 General aspects of tasking

5.1 States and transitions

At any time, a task is in one and only one state. Actions executed by the executive system, other tasks, or the designated task itself, may cause transition from one state to another. These transitions are performed instantly, i.e. they are considered ideally to take no time.

This mathematical model of a task may be visualized by a "state diagram", such as that shown in the figure, in which the states are nodes, illustrated as circles, while transitions are drawn as pointed arrows from one node to another. [3]

1) See also annex B.

A multitasking system consists of several parallel tasks and can be considered as modelled by a number of disjointed but similar diagrams. It is feasible to apply a three-dimensional picture, with the similar diagrams sandwiched on top of each other and oriented so that the identical states of the individual diagrams cover each other.

State transitions are generally caused by subroutine calls, occurrence of events, or expiration of time limits. The name, form, and interpretation of the subroutine calls are standardized and described below.

In this International Standard, a task is described by a mathematical model illustrated by the state diagram shown in the figure. This model adheres to the following basic principles:

- transitions are non-ambiguous, i.e. for a given stimulus in a given state, the task can transit to only one possible new state;
- transitions are performed instantly, i.e. in zero time;
- a task exists in one state only at a time;
- the state model describes the behaviour of tasks as seen by the application programmer.




Figure — State model and transition diagram

The following symbolism is used for transitions in the state diagram:

- Capital letters in box : Effect on a task imposed by another task, i.e. a subroutine call in one task has the indicated effect on the designated task.
- Capital letters without box : Effect imposed by the task itself while in the state RUNNING.
- Small letters : Conditions under which the executive system performs the indicated state transitions.

With reference to principle d) above, no attempt is made to describe executive system actions transparent to the application programmer. Consequently, the state RUNNING is related to the task's virtual processor (see 2.32). It is immaterial for the state of a task whether a physical processor happens to be assigned to it or the execution is temporarily hampered by the executive system due to limited availability of physical processors and the task's low priority. Thus, the model is as well adapted to multiprocessors as it is to single processor computers.

5.2 Multiple activation calls¹⁾

Different tasks may issue apparently conflicting transition calls for the same object task. This will be the case during normal operation, as well as under error conditions, since the state of an object task is unknown at the time a transition call is made by another task.

In accordance with the principle of a task being in one state only at a time (see 5.1), a distinction is made between state transitions and calls for such transitions. Transition calls are received by the executive system, which will apply its own scheduling strategy in handling such calls.

Depending on available resources, such as internal table space, etc., of the executive system, a transition call will be accepted or rejected. An argument will be returned after the reference, with an appropriate value indicating whether the reference has been accepted normally or rejected.

5.3 Synchronization concepts

This International Standard provides three concepts for synchronization between tasks and the resolution of resource contention, as follows:

- eventmarks;
- resourcemarks;
- semaphores.

Eventmarks and semaphores are mainly used for synchronization purposes, whereas the resourcemarks are mainly used for the resolution of resource contention.

Eventmarks, resourcemarks, and semaphores are local variables of the executive system and they may not be accessed other than by the mechanisms described in this International Standard.

5.3.1 Eventmarks

In the management of concurrent tasks, it is necessary to associate certain tasks with certain events. These events may be either external or internal events. External events are of some physical nature, such as a contact closure, but the connection between an external event and its eventmark is beyond the scope of this International Standard. An internal event arises from specific program action (see 6.9.1 for a description of CALL POST).

Eventmarks are selected by reference to a numeric selector in the range 1 (one) to n , where n is processor dependent. Eventmarks have two states

ON
OFF

The association of events to tasks is done by the subroutine calls CALL SKED (see 6.4), CALL CON (see 6.5.8), CALL SUSPND (see 6.7), and CALL HOLD (see 6.8.1).

An eventmark is turned to ON by the occurrence of an internal or external event. If one or more tasks are associated with this event, the executive system will cause each associated task to begin or continue execution. An eventmark is turned to OFF by direct program control or by the executive system as it services the tasks associated with the eventmark. Eventmarks can be changed only by the procedure references defined in this International Standard and by the executive system that services the events.

Eventmarks may also be set to the OFF state by a specific program action, the CALL CLEAR (see 6.9.2), and they may be masked and unmasked (see 6.9.4 and 6.9.5). Additionally, the value of an eventmark may be tested by the logical function TESTEM (see 6.9.3).

5.3.2 Resourcemarks

A simple means to resolve contention for various resources is provided by the resourcemark concept, which permits one and only one task to use a resource at a time.

Resourcemarks are selected by reference to a numeric selector in the range 1 (one) to n , where n is processor dependent. Resourcemarks have two states

LOCKED
UNLOCKED

This International Standard does not define what can be considered to be a resource. It is the responsibility of the user to associate a resource with a resourcemark. If the user wants to reserve a resource exclusively for the running task, he chooses a resourcemark for the resource and performs the reference CALL LOCK (see 6.10.1). Should the resourcemark be in the state UNLOCKED at this moment, the resourcemark will change to state LOCKED, reserving the corresponding resource exclusively for this task. This reservation is later released by the task, by execution of a reference to EXIT (see 6.12) or UNLOCK (see 6.10.2).

1) See also annex B.

Should a task attempt to LOCK a resource mark which is already locked, the executive system will transfer this task into the state SUSPENDED, where it remains until the resource mark becomes unlocked by some other task.

If several tasks are waiting for a resource mark to be UNLOCKED, only one will be selected for execution by the executive system. (For details, see 6.10.2.) This is a main difference from the eventmark concept, where all tasks waiting for an eventmark are transferred to state RUNNING if the eventmark is set to ON.

A resource mark can also be set to LOCKED by reference to the logical function TLOCK (see 6.10.3).

5.3.3 Semaphores

Semaphores represent a synchronization concept useful for cases where a more advanced mechanism than those available by resource marks and eventmarks is desired. Semaphores are selected by reference to a numeric selector in the range 1 (one) to n , where n is processor dependent. Unlike eventmarks and resource marks, a semaphore has an integer value, s . The value s of a semaphore must be initialized and may later be set by the reference CALL PRESEM (see 6.11.1).

If the task needs to wait until the value of a semaphore is greater than or equal to a specific value, the task uses the reference CALL WAITS (see 6.11.2). By execution of CALL WAITS, the executive system tests if the specified decrement j is greater than s . In this case, the calling task is transferred to state SUSPENDED, where it remains until the semaphore value s is incremented by another task to a value greater than or equal to j . If j is not greater than s , the value s is decremented by j and the calling task continues its execution.

By execution of CALL SIGNAL (see 6.11.3), the executive system increments by j the value s of the specified semaphore. This incrementation may bring a task, waiting for this semaphore, from state SUSPENDED to state RUNNING, if s becomes greater than or equal to j of the suspended task. If multiple tasks are waiting for this semaphore and are candidates for running, after the incrementation of s by j , the order in which these tasks transit to state RUNNING is processor dependent.

By setting s to certain values and choosing different values for j , a variety of advanced synchronization concepts and solutions for resource contention are possible.

In addition to the calls described above, the value of a semaphore may be read using the integer function IRDSEM (see 6.11.4).

6 Procedure references

6.1 Terms and summary of procedure references¹⁾

This clause contains a summary of the subroutine calls and function references described in subsequent clauses.

The following designations for parameters apply to several of the calls. If the exact meaning of these parameter designations deviates from that described below, it will be marked specifically in the detailed description of the call. If the meaning is exactly as defined here, the description of the parameter will be omitted in the description of the call and a reference will be made to this subclause.

i specifies the task to be affected (object task). The argument shall be an integer array. The contents of i may be partly generated by CREATE. i is used as input parameter in all other calls.

t, t_1, t_2 designate integer arrays, whose first 8 elements contain a specification of absolute or relative time. Negative values of elements are not permitted. These elements are as follows:

- first element: Basic clock counts
- second element: Milliseconds
- third element: Seconds
- fourth element: Minutes
- fifth element: Hours
- sixth element: Day(s)
- seventh element: Month(s)
- eighth element: Year(s)

If, for absolute times, value 0 is used for one of the three date elements, this shall be interpreted as "current date", "current month", or "current year" by the executive system.

The interpretation of a relative time specification containing months or years different from zero is processor dependent.

m is set on return to the calling program, to indicate the disposition of the request as follows:

- 0 or less: undefined
- 1: request accepted
- 2 or greater: request rejected (error condition)

This argument shall be an integer variable or integer array element. The processor may define specific values greater than or equal to 2 to distinguish between certain reasons for rejection.

The list of function and subroutine calls, described in detail in subsequent subclauses, is given in table 1.

1) See also annex B.

Table 1

Full description in subclause	Call	Parameters
6.2	CALL CREATE(i,m)	i : identification of created task and associated program
6.3	CALL KILL(i,m) (opposite of CREATE)	
6.4	CALL SKED(i,s,e1,t1,t2,e2,m) (general scheduling)	s : mode selector e1 : eventmark reference t1 : absolute or relative time for first initiation t2 : time period for cyclic initiations e2 : reference to eventmark for overrun
6.5.1	CALL STRT(i,m) (start immediately)	
6.5.2	CALL STRTAF(i,t1,m) (start after time delay)	t1 : time delay before initiation
6.5.3	CALL STRTAT(i,t1,m) (start at absolute time)	t1 : absolute time for initiation
6.5.5	CALL CYCL(i,t2,m) (cyclic, with immediate first initiation)	t2 : length of time interval
6.5.6	CALL CYCLAF(i,t1,t2,m) (cyclic, with delayed first initiation)	t1 : time delay before first initiation t2 : length of time interval
6.5.7	CALL CYCLAT(i,t1,t2,m) (cyclic, with absolute time specification of first initiation)	t1 : absolute time for first initiation t2 : length of time interval
6.5.8	CALL CON(i,e,m) (establish event connection)	e : eventmark reference
6.6	CALL DSKED(i,s,e,m) (eliminate scheduling)	s : mode selector e : eventmark reference
6.6.1	CALL DCON(i,e,m) (eliminate event connection)	e : eventmark reference
6.6.2	CALL CANCEL(i,m) (eliminate time scheduling)	
6.7	CALL SUSPND(s,e,t,n,m) (suspend continuation of calling task for time period or until event)	s : mode selector e : reference to eventmark for end of delay t : time delay n : indicator for cause of end of delay
6.8.1	CALL HOLD(e,m) (suspend until event occurs)	e : reference to eventmark for end of delay
6.8.2	CALL DELAY(t,m) (suspend for a relative time)	t : time delay
6.9.1	CALL POST(e,m) (setting of an eventmark)	e : eventmark reference
6.9.2	CALL CLEAR(e,m) (resetting of an eventmark)	e : eventmark reference
6.9.3	TESTEM(e,m) (testing the state of an eventmark)	e : eventmark reference function value : condition of the eventmark
6.9.4	CALL MKEM(e,m) (setting the mask of an eventmark)	e : eventmark reference
6.9.5	CALL UNMKEM(e,m) (clearing the mask of an eventmark)	e : eventmark reference
6.10.1	CALL LOCK(r,m) (locking of a resourcemark)	r : resourcemark reference
6.10.2	CALL UNLOCK(r,m) (unlocking of a resourcemark)	r : resourcemark reference

Table 1 (concluded)

Full description in subclause	Call	Parameters
6.10.3	TLOCK(r,m) (testing and locking of a resourcemark)	r : resourcemark reference function value : condition of the resourcemark
6.11.1	CALL PRESEM(r,s,m) (initialization of semaphore)	r : semaphore reference s : initial value of semaphore
6.11.2	CALL WAITS(r,j,m) (wait on semaphore)	r : semaphore reference j : decrement
6.11.3	CALL SIGNAL(r,j,m) (release semaphore)	r : semaphore reference j : increment
6.11.4	IRDSEM(r,m) (read semaphore value)	r : semaphore reference function value : value of semaphore variable
6.12	CALL EXIT (termination of execution)	

6.2 Creation of a new task

A new task is introduced to the real-time system by reference to subroutine CREATE. The designated task will be associated with some specified program, considered a resource like other resources, necessary for the task to perform. The associated program is normally assumed to exist in an executable form. Formally, and in terms of the state model, the task is transferred from NON-EXISTENT to DORMANT as the effect of the reference (see the figure and 5.1).

A mechanism is assumed to exist outside the scope of this International Standard to create and initiate at least a first task, i.e. the parent, which in its turn may create other tasks.

The form of the call is

CALL CREATE(i,m)

where

i specifies an integer array which contains all the information necessary to specify the task and its associated program. The latter includes, among other items, its designation, where the program can be found (such as description of file), residency while existent (primary memory resident or swappable), etc. The array may also contain the task's processor priority.

This array will in general also contain output information, i.e. references distinguishing this task from other tasks created from an identical program code, reference to individual data sets, etc.

Such information may be used by other procedure references in clause 6. For a description of parameter i, see 6.1.

All details of this array are processor dependent.

m is as described in 6.1.

6.3 Eliminating a task from the real-time system¹⁾

A reference to subroutine KILL will eliminate a designated task from the real-time system, by transferring it to state NON-EXISTENT. If the designated task is in state DORMANT or PENDING, the effect shall be carried out immediately. If the designated task is in state RUNNING or SUSPENDED, the termination shall affect only future executions. Thus, the designated task will continue its present execution without any intervention by this call.

The form of the call is

CALL KILL(i,m)

where i and m are as described in 6.1.

6.4 Scheduling a task

Execution of a reference to the subroutine SKED, or its derivatives as listed in 6.5, shall schedule the initiation of a designated task, establishing the condition for subsequent transition to RUNNING. If the designated task is in state DORMANT when the call is made, the designated task shall transit to state PENDING. If it is already in state PENDING, the reference shall augment its conditions for subsequent transfer to RUNNING according to the arguments of the call, such that the designated task will transit to RUNNING when any of the conditions still valid becomes true. The transition to state RUNNING requires the task being in state PENDING, otherwise an overrun condition occurs. The augmentation of running conditions is subject to any resource limitation of the processor, and any violation of such limitation will result in an error return.

When a task transits to RUNNING, the condition that caused this transition is removed from the possible complex of combined conditions, such that the other conditions remain. When a running task subsequently exits, by transition from RUNNING to DORMANT, possible scheduling conditions, remaining from previous scheduling calls, shall cause immediate transition to state PENDING. (Note, however, the definition of the term "overrun" and the provision for its indication, as explained below.)

1) See also annex B.

Each normally accepted reference to the subroutine SKED causes the following effect.

After expiration of a specified time delay or at a desired absolute time or upon the occurrence of a specified event, the object task is transferred to state RUNNING and begins at the first executable statement of the program. The actual time resolution obtainable in a specific industrial computer system is subject to the resolution of that system's real-time clock. If the object task is initiated by an event occurrence, the executive system will set the eventmark OFF, as part of the initiation. If more than one task is waiting for a specified eventmark to change to the condition ON, when the change occurs all these tasks will transit into state RUNNING.

The form of this procedure reference is

CALL SKED(i,s,e1,t1,t2,e2,m)

where

i and m are as described in 6.1.

s is an integer expression, specifying three categories of task scheduling (see below), as follows:

- the values of the argument s between 10 and 15 cause the task to be initiated once;
- the values of the argument s between 20 and 25 cause the task to be initiated periodically by time;
- the values of the argument s between 30 and 35 cause the task to be initiated whenever a specified eventmark becomes ON.

e1 specifies the eventmark for scheduling; it is an integer expression.

t1 designates an integer array containing the absolute, or relative, time for the scheduling. (The term relative time for scheduling means the time delay from the time the reference is executed until the intended running.) (See 6.1.)

t2 is an integer array to designate the time period for cyclic runs. (See 6.1.)

e2 is an eventmark for overrun. It will be turned ON if overrun occurs. The action performed with the object task is processor dependent.

The values of the argument s between 10 and 15 define the first and only execution of the task.

- If s = 10: start immediately
- If s = 11: start at absolute time t1
- If s = 12: start after time t1
- If s = 13: start at event e1 (once)
- If s = 14: start at absolute time t1 or event e1 (once)
- If s = 15: start after time t1 or at event e1 (once)

The values of the argument s between 20 and 25 define the first execution of the task which is subsequently executed repetitively at the time period t2.

If s = 20: start immediately plus cyclic by t2

If s = 21: start at t1 plus cyclic by t2

If s = 22: start after t1 plus cyclic by t2

If s = 23: start at e1 plus cyclic by t2

If s = 24: start at t1 or e1 plus cyclic by t2

If s = 25: start after t1 or at e1 plus cyclic by t2

The values of the argument s between 30 and 35 define the first execution of the task which is subsequently executed whenever the eventmark specified by e1 becomes ON.

If s = 30: start immediately plus repeated by e1

If s = 31: start at t1 plus repeated by e1

If s = 32: start after t1 plus repeated by e1

If s = 33: start at e1 plus repeated by e1

If s = 34: start at t1 or e1 plus repeated by e1

If s = 35: start after t1 or at e1 plus repeated by e1

6.5 Starting a task by simplified calls

The following calls for scheduling represent subsets of the features of CALL SKED. They are established for ease of programming only.

CALL STRT(i,m)	start immediately
CALL STRTAF(i,t1,m)	start after time t1
CALL STRTAT(i,t1,m)	start at time t1
CALL CYCL(i,t2,m)	start immediately plus cyclic by t2
CALL CYCLAF(i,t1,t2,m)	start after time t1 plus cyclic by t2
CALL CYCLAT(i,t1,t2,m)	start at time t1 plus cyclic by t2
CALL CON(i,e1,m)	start at event e1 plus repeated by e1

6.5.1 Starting a task immediately

Execution of a reference to subroutine STRT establishes a condition for immediate transfer of the designated task to RUNNING via PENDING. Execution begins at the task's first executable statement. If in state DORMANT when the reference is made, the task will transit to state PENDING and continue immediately to state RUNNING.

The form of the call is

CALL STRT(i,m)

where i and m are as described in 6.1.

6.5.2 Starting a task after a specified time delay

Execution of a reference to subroutine STRTAF establishes a time delay as the condition for transfer of the designated task to RUNNING via PENDING.

After expiration of a specified time delay after the time when the reference is executed, the designated task is expected to transfer to state RUNNING and will do so if its state at that time is PENDING. If in state DORMANT when the reference is

made, the task will transit to state PENDING. In state RUNNING, the task will begin at its first executable statement.

The form of the call is

CALL STRTAF(i,t1,m)

where

i and m are as described in 6.1.

t1 designates an integer array, specifying the time delay after which the object task is to start its execution (see 6.1).

6.5.3 Starting a task at a specified absolute time

Execution of a reference to subroutine STRTAT establishes an absolute time as the condition for transfer of the designated task to RUNNING via PENDING.

At the specified absolute time, the object task is expected to transfer to state RUNNING and will do so, if its state at that time is PENDING. If in state DORMANT when the reference is made, the task will transit to state PENDING. In state RUNNING, the task will begin at its first executable statement. The task is started immediately if the specified absolute time is already passed when the reference to STRTAT is executed.

The form of the call is

CALL STRTAT(i,t1,m)

where

i and m are as described in 6.1.

t1 designates an integer array, specifying the absolute time at which the object task is to start its execution (see 6.1).

6.5.4 Starting a task in repeated execution

The calls for CYCL, CYCLAF, CYCLAT, and CON for repeated execution have the following common features :

The designated task will be transferred to state PENDING if its present state is DORMANT or when it becomes DORMANT. Further, the reference establishes the condition for subsequent transfer of the designated task from state PENDING to RUNNING for a first execution and additionally causes future repeated executions.

A reference to subroutine CYCL, CYCLAF, CYCLAT, or CON has the same immediate effect as a reference for single execution; additionally, after its termination (for example by EXIT), the object task will be transferred immediately from state DORMANT to state PENDING for the next repeated execution, as indicated in the figure by "repetitions".

The actual running may be delayed unintentionally while in state RUNNING, because of running of other programs. Such delays will not be accumulated for tasks with cyclic repetitions.

If the execution is not finished before the time for the next execution, an overrun situation exists and the action taken with the repeated task is processor dependent.

Re-scheduling under the specified conditions shall continue until actively terminated by a call of subroutine CANCEL (see 6.6.2) or DCON (see 6.6.1).

The statements above are similarly valid for the CALL SKED with parameter s = 20 to 25 and 30 to 35; SKED, however, has a special parameter e2, set to ON in the case of the overrun situation.

6.5.5 Initial start immediately

The form of the call is

CALL CYCL(i,t2,m)

where

i and m are as described in 6.1;

t2 designates an integer array, specifying the nominal length of the time interval (see 6.1).

6.5.6 Initial start after a specified time delay

The form of the call is :

CALL CYCLAF(i,t1,t2,m)

where

i and m are as described in 6.1;

t1 designates an integer array, specifying a time delay for the initial activation, as measured from the time the call was made (see 6.1);

t2 designates an integer array, specifying the nominal length of the time interval (see 6.1).

6.5.7 Initial start at specified absolute time

The form of the call is

CALL CYCLAT(i,t1,t2,m)

where

i and m are as described in 6.1;

t1 designates an integer array, specifying the absolute time at which the designated task is supposed to enter state RUNNING initially. This argument is exactly equivalent to parameter t1 of call for subroutine STRTAT (see 6.1);

t2 designates an integer array, specifying the nominal length of the time interval (see 6.1).

6.5.8 Connection of a task to an event

Execution of a reference to subroutine CON establishes a specified event as the condition for transition from state PEN-

DING to RUNNING. First, the task will transit to state PENDING if the state is DORMANT when the reference is executed, or when the state becomes DORMANT. Then, if the related eventmark is or becomes ON, the designated task will transit from state PENDING to RUNNING and will begin with its first executable statement. The association between the event and the designated task remains until actively cancelled by a reference to DSKED or DCON (see 6.6 and 6.6.1).

The form of the call is

CALL CON(i,e,m)

where

i and m are as described in 6.1.

e specifies the eventmark; it is an integer expression.

6.6 Elimination of previous scheduling

Execution of a reference to subroutine DSKED or its derivatives, DCON and CANCEL, shall cancel specified scheduling conditions for an object task. Thus, it eliminates further effects from previous calls to subroutine SKED or its simplified derived subroutines.

If the object task is in state RUNNING or SUSPENDED, the cancelling shall affect only future executions. Thus, the object task will conclude its present execution, without any intervention by this call.

If the object task is in state PENDING, the cancelling will cause a transition to state DORMANT if no scheduling conditions remain.

The form of the call is

CALL DSKED(i,s,e,m)

where

i and m are as described in 6.1.

s is an integer expression selecting one of the following conditions:

If s = 1: eliminate all time-based scheduling and scheduling by events as specified by argument e;

If s = 2: eliminate all time-based scheduling including repetition by time. Possible event connections remain unaltered;

If s = 3: eliminate event-based scheduling, including event-based repetition, as specified by argument e. Possible time connections remain unaltered.

e is an integer expression, which specifies the eventmark(s), whose connection is to be eliminated, as follows:

If e = -1: all eventmark connections for the designated task;

If e = 0: no event cancelling;

If e > 0: reference to one specific eventmark, as specified by the value of e.

6.6.1 Elimination of event connections

Execution of a reference to subroutine DCON shall cancel any connection between an object task and a specified event. Thus, it eliminates further effects from a previous call to subroutine SKED or CON. If the object task is in state RUNNING or SUSPENDED, the cancelling shall affect only future executions. Thus, the object task will conclude its present execution, without any intervention by this call.

If the object task is in state PENDING, the cancelling will cause a transition to state DORMANT if no scheduling conditions remain.

The form of the call is

CALL DCON(i,e,m)

where

i and m are as described in 6.1.

e specifies the eventmark; an integer expression (see 6.6).

6.6.2 Elimination of time connections

Execution of a reference to subroutine CANCEL shall cancel the future initiations of an object task due to time scheduling by previous calls of SKED or its simplified versions. The executive system shall ensure that no further move to the state PENDING shall take place due to previous cyclic scheduling.

If the object task is in any active state (RUNNING or SUSPENDED), the elimination shall affect only future executions. Thus, the object task will conclude its present execution without any intervention by this call.

If the object task is in state PENDING, the cancelling will cause a transition to state DORMANT if no scheduling conditions remain.

The form of the call is

CALL CANCEL(i,m)

where i and m are as described in 6.1.

6.7 Delaying continuation of a task

Execution of a reference to the subroutine SUSPND shall provide a means whereby a running task is suspended (i.e. transits to state SUSPENDED) for a specified length of time or until a specified event has occurred. The task shall then transit back to state RUNNING and shall resume execution with the statement immediately following the call of subroutine SUSPND.

If more than one task is suspended and waiting for a specified eventmark to change to the condition ON, when the change occurs all tasks will transit to state RUNNING.

The time delay is defined as the nominal duration from the time when the call was made until the program resumes execution in its virtual processor, by being transferred to state RUNNING. The actual instants for the entering and leaving states RUNNING and SUSPENDED are subject to the resolution of the system's real-time clock and to the interrogating and activating actions performed by the executive system.

The form of the call is

CALL SUSPND(s,e,t,n,m)

where

s is an integer expression selecting the condition on which the suspension shall end, as follows:

- If s = 1: at absolute time t;
- If s = 2: after time t;
- If s = 3: at event e;
- If s = 4: at absolute time t or event e;
- If s = 5: after time t or at event e.

e specifies the eventmark for ending suspension; it is an integer expression.

t designates an integer array, specifying the absolute, or relative time; for the suspension (see 6.1).

n is a return parameter to indicate the cause of the end of the delay in case of s = 4 or s = 5:

- If n = 1: end of delay by event e;
- If n = 2: end of delay by time t.

m is as described in 6.1.

6.8 Suspending until event or time

The following calls for suspension represent subsets of the features of CALL SUSPND. They are established for the ease of programming only.

CALL HOLD(e,m) suspend until the event e has occurred

CALL DELAY(t,m) suspend for a time delay as specified by t

6.8.1 Delay until event has occurred

Execution of a reference to subroutine HOLD shall suspend the calling task until a specified event has occurred. The task shall then transit to state RUNNING and thus resume execution with the statement immediately following the call of subroutine HOLD.

The form of the call is

CALL HOLD(e,m)

where

e specifies the eventmark for ending suspension; it is an integer expression.

m is as described in 6.1.

6.8.2 Delay for a specified relative time

Execution of a reference to subroutine DELAY shall transit the calling task to state SUSPENDED for a specified duration. The task shall then transit back to state RUNNING, resuming execution with the statement immediately following the call of subroutine DELAY.

The form of the call is

CALL DELAY(t,m)

where

t designates an integer array, specifying the relative time for the suspension (see 6.1).

m is as described in 6.1.

6.9 Eventmark operations

6.9.1 Setting an eventmark to the ON condition

Execution of a reference to subroutine POST shall set a designated eventmark to the ON condition. If the eventmark was already ON, there shall be no action. Because of an earlier transition call in some task, the ON condition may cause a task to be transferred from state PENDING or SUSPENDED to state RUNNING.

The form of the call is

CALL POST(e,m)

where

e specifies the eventmark; it is an integer expression.

m is as described in 6.1.

6.9.2 Clearing an eventmark

Execution of a reference to subroutine CLEAR shall cause the designated eventmark to become OFF. If the eventmark was already OFF, there shall be no action.

The form of the call is

CALL CLEAR(e,m)

where

e specifies the eventmark; it is an integer expression.

m is as described in 6.1.

6.9.3 Testing an eventmark condition

Execution of a reference to function TESTEM shall return a logical value TRUE if the specified eventmark was ON and a logical value FALSE if the eventmark was OFF. If the eventmark is unknown to the processor, a logical FALSE value will be returned, and the error parameter will indicate an error condition.

The form of this function reference is

TESTEM(e,m)

where

e specifies the eventmark; it is an integer expression.

m is as described in 6.1.

6.9.4 Masking an eventmark

Execution of a reference to subroutine MKEM does not change the state of the designated eventmark but causes it to be masked. The masking effect is that the eventmark may freely change its state without any effect on tasks that might be PENDING or SUSPENDED waiting for this eventmark to be set.

The form of the call is

CALL MKEM(e,m)

where

e specifies the eventmark whose corresponding event is to be masked; it is an integer expression.

m is as described in 6.1.

6.9.5 Unmasking an eventmark

Execution of a reference to the subroutine UNMKEM shall allow actions associated with the specified eventmark to be executed. If the eventmark is in the ON condition, then all actions associated with the specified eventmark will be executed.

The form of the call is

CALL UNMKEM(e,m)

where

e specifies the eventmark whose corresponding event is to be unmasked; it is an integer expression.

m is as described in 6.1.

6.10 Resourcemark operations

6.10.1 Setting a resourcemark to the locked condition

Execution of a reference to the subroutine LOCK shall cause the specified resourcemark to be locked. If the specified resourcemark was locked already, the executive system shall suspend the execution of the task.

The form of the call is

CALL LOCK(r,m)

where

r specifies the resourcemark; it is an integer expression.

m is as described in 6.1.

6.10.2 Setting a resourcemark to the unlocked condition

Execution of a reference to the subroutine UNLOCK shall cause one of the following actions:

a) if the resourcemark is unlocked, there shall be no action;

b) if the resourcemark is locked and there are no tasks suspended as a result of a previously unsuccessful attempt to lock the associated resource, the resourcemark shall be unlocked;

c) if the resourcemark is locked and there are one or more tasks suspended as a result of previous attempts to lock the associated resource, one and only one task shall transit to state RUNNING. The associated resourcemark remains locked. The criteria used to select the task to transit to state RUNNING are processor dependent.

The form of the call is

CALL UNLOCK(r,m)

where

r specifies the resourcemark; it is an integer expression.

m is as described in 6.1.

6.10.3 Testing and setting a resourcemark to the locked condition

Execution of the function TLOCK shall first test the specified resourcemark. The function shall return the value TRUE if the resourcemark is unlocked; it shall return the value FALSE, however, if the resourcemark is locked. After this test, the resourcemark will be locked.

The form of this function reference is

TLOCK(r,m)

where

r specifies the resourcemark; it is an integer expression.

m is as described in 6.1.

The reason for the use of TLOCK as compared to CALL LOCK is to allow a task to either reserve a resource if it is unlocked, or to continue execution if it is locked. This would not be possible with CALL LOCK.

6.11 Semaphore operations¹⁾

All semaphore variables have the form of local variables of the executive system, and the only means of access is through an argument, r, which refers to one particular semaphore. The value of a particular semaphore variable thus referred to is termed s in the following.

The effects of subroutines SIGNAL and WAITS are, respectively, the increasing and decreasing of the semaphore value by an amount j, a positive integer conveyed as the second argument of the calls. For WAITS, the decrease will only take place if the result is not less than zero; otherwise the calling task is suspended before the decrementing takes place, and continuation will not occur until after $s \geq j$, i.e. the decrementation linked to the continuation will yield a non-negative value.

1) See also annex B.

The subroutine calls for the synchronizing mechanisms are described in more detail below.

6.11.1 Initialization of semaphore

Execution of a reference to the subroutine PRESEM has two purposes. Firstly, it declares the intention of use of a particular semaphore, permitting the system to give diagnostic warnings at run time if another semaphore operation is issued referring to a semaphore that is not initialized. Secondly, PRESEM establishes the initial value for the semaphore. Normally, PRESEM is referenced only in the initialization phase of the execution of a real-time program. The error parameter *m* will indicate an error condition if no semaphore exists with the indicated designation.

The form of the call is

CALL PRESEM(*r,s,m*)

where

r specifies a semaphore; it is an integer expression.

s is the initial value given to the semaphore. Until the CALL PRESEM is executed for a particular semaphore, and the internal variable is assigned value *s*, the internal value is undefined, and another system call referring to this semaphore shall give an error return. This parameter is an integer expression.

A negative value is permissible. This is the only way a semaphore may attain a negative value. The effect of a negative initial value is that a correspondingly greater increase by virtue of CALL SIGNAL is required before the releasing action can take place.

m is as described in 6.1.

6.11.2 Wait on semaphore

Execution of a reference to subroutine WAITS will involve a possible suspension of the calling task, as controlled by the referenced semaphore.

By the end of the call, the semaphore value *s* will be reduced by the amount *j*. The reduction and subsequent continuation will only take place when this can be done giving *s* a non-negative value. Otherwise, the calling task is suspended until this decrementation can take place.

The form of the call is

CALL WAITS(*r,j,m*)

where

r is as described in 6.11.1.

j is an integer expression, specifying the amount by which the semaphore variable is to be decremented, if applicable. The value of *j* shall be positive (one or greater), and *j* = 1 corresponds to the simple semaphore most commonly used.

m is as described in 6.1.

6.11.3 Release of semaphore

Execution of a reference to the subroutine SIGNAL shall increment an integer semaphore variable, designated *s* in the expression below, and referred to by one of the arguments of the reference. The subroutine shall be granted exclusive access to this semaphore during its operation and will execute the following modification of its value :

$$s = s + j$$

where *j* is an argument (see below).

A change to a positive value of sufficient magnitude, caused by this operation, may provide a possible releasing transfer to state RUNNING for a task waiting in state SUSPENDED for this event to occur.

Other tasks, suspended in their execution of CALL WAITS relating to the same semaphore *r*, shall have their suspension condition re-evaluated after the present SIGNAL call is terminated. This will provide an opportunity for suspended tasks to be released and resume operation, as described above. This continuation is subject to all common restrictions pertaining to exclusive operation on a same semaphore. Thus, the effect will be that only one of the suspended tasks will be examined at a time. This examination may involve reduction of the semaphore value again, as a consequence of releasing operation of WAITS for the examined task. This examination continues as long as a possibility remains that the semaphore value is greater than or equal to the *j*-value of some suspended task. The order in which suspended tasks are checked otherwise is processor dependent.

The form of the call is

CALL SIGNAL(*r,j,m*)

where

r is as described in 6.11.1.

j is an integer expression, specifying the amount by which the semaphore variable is to be incremented, if applicable. The value of *j* shall be positive (one or greater), and *j* = 1 corresponds to the simple semaphore most commonly used.

m is as described in 6.1.

6.11.4 Reading a semaphore value

A semaphore value may be interrogated by execution of a reference to the integer function IRDSEM. The purpose of this function is not that it be used as a synchronization operation, but only to provide a means to supervise synchronization in a system. Thus, a reference to function IRDSEM can, for example, provide information about how far a buffer or another shared resource is from saturation. When accepted (honoured), the reference will be granted exclusive access to the semaphore. If the semaphore is already being accessed by another system call when the reference to IRDSEM is made, the reference will be subject to the same deferred response and contention mechanisms as the other semaphore operations. On

return, the function designator will have a value equal to the internal semaphore value when the reference was accepted.

The form of the reference is

IRDSEM(r,m)

where

r is as described in 6.11.1.

m is as described in 6.1.

6.12 Normal termination of execution¹⁾

Execution of a reference to subroutine EXIT shall terminate the execution of a task and return the task to state DORMANT.

Eventmarks and semaphores shall not be affected. Resource marks previously locked by this task shall be unlocked (see 5.3.2), and files released.

The form of this call is

CALL EXIT

The common FORTRAN operations STOP and END provide alternative means to terminate execution of a task. However, the effect on connected units, such as files, is as described in ISO 1539^[1]; the effect on other resources described in this International Standard is processor dependent.

1) See also annex B.

Section two: Binary pattern and bit processing

7 Introduction

The function references described in this section provide mechanisms for operations on bit patterns as well as individual bits of the internal representation of integer variables.

The operations presuppose that integer numbers are considered as if they are *unsigned integers*. The arithmetic shift operation (see 8.2.2) is an exception.

With respect to arguments referencing individual bits, a bit numbering rule in which bit number 0, the rightmost bit, is the least significant bit is assumed.

8 Binary pattern processing

8.1 Boolean operations

Boolean operations provided are

OR, AND, EOR and NOT.

These operations are implemented as integer functions. The implicit type for OR, AND and EOR is indicated by the use of I as the first letter of the function name. Their parameters, *j* and *k*, are integer expressions.

After execution of the functions, the parameters remain unchanged. The operations are performed on corresponding (equally numbered) bits of the two operands, giving the corresponding bit value of the result. The values indicated in the following truth tables represent the individual and corresponding bits of the arguments and of the function value.

8.1.1 Inclusive OR

The form of this function is

IOR(*j*,*k*)

The value of the function is computed from the values of the parameters *j* and *k* according to the following truth table:

<i>j</i>	0	1	0	1
<i>k</i>	0	0	1	1
Function value	0	1	1	1

8.1.2 Boolean AND

The form of this function is

IAND(*j*,*k*)

The value of the function is computed from the values of the parameters *j* and *k* according to the following truth table:

<i>j</i>	0	1	0	1
<i>k</i>	0	0	1	1
Function value	0	0	0	1

8.1.3 Boolean complement

The form of this function is

NOT(*j*)

The value of the function is the logical complement of the parameter value, *j*, according to the following truth table:

<i>j</i>	0	1
Function value	1	0

8.1.4 Exclusive OR

The form of this function is

IEOR(*j*,*k*)

The value of the function is computed from the values of the parameters *j* and *k* according to the following truth table:

<i>j</i>	0	1	0	1
<i>k</i>	0	0	1	1
Function value	0	1	1	0

8.2 Shift operations

The shift operations provided are logical, arithmetic and circular. The shift operations are implemented as integer functions. The functions have two parameters, *j* and *n*, considered as integer expressions, as follows:

j specifies the value (binary pattern) to be shifted;

n specifies the shift count:

n > 0 indicates a left shift,

n = 0 indicates no shift,

n < 0 indicates a right shift.

If the absolute value of the shift count is greater than the number of bits in a numeric storage unit, then the result is *undefined*.

The parameter values are not changed by the shift operations.

8.2.1 Logical shift

The form of this function is

ISHL(*j*,*n*)

All bits representing the parameter *j* are shifted *n* places. Bits shifted out from the left end or the right end, as the case may be, are lost. Zeros are shifted in from the opposite end.

8.2.2 Arithmetic shift

The form of this function is

ISHA(j,n)

Argument j and the function value are considered as signed integers. All bits representing the parameter j are shifted n places. In the case of a right shift ($n < 0$), zeros are shifted into the left end if j is positive, and ones are shifted in if j is negative. The bits shifted out of the right end are lost. In case of a left shift ($n > 0$), zeros are shifted into the right end while the bits shifted out of the left end are lost. In a left shift an arithmetic overflow may occur.

8.2.3 Circular shift

The form of this function is

ISHC(j,n)

All bits representing the parameter j are shifted circularly n places, i.e., the bits shifted out of one end are shifted into the opposite end. No bits are lost.

NOTE — The number of bits representing j is processor dependent.

9 Bit processing

Individual bits of an integer can be tested with the functions for bit processing. The functions have two parameters j and k , which are integer expressions, as follows:

j specifies the binary pattern;

k specifies the selected bit, numbered as in clause 7.

If k is negative or greater than the number of bits that are used to represent an integer value, the result of the function is undefined.

The parameter values are not changed by these functions.

9.1 Bit testing

The form of this function is

BTEST(j,k)

This function is of type LOGICAL. The k^{th} bit of parameter j is tested. If it is 1, the value of the function is TRUE; if it is 0, the value of the function is FALSE.

9.2 Set bit

The form of this integer function is

IBSET(j,k)

The value of the function is equal to the value of parameter j with the k^{th} bit set to 1.

9.3 Clear bit

The form of this integer function is

IBCLR(j,k)

The value of the function is equal to the value of parameter j with the k^{th} bit set to 0.

9.4 Change bit

The form of this integer function is

IBCHNG(j,k)

The value of the function is equal to the value of parameter j with the k^{th} bit complemented.

Section three: Process input/output

10 Introduction

The user of industrial real-time FORTRAN must be able to address specific process devices for his application. As the majority of input/output (I/O) systems are computer dependent, this can only be standardized in a universal way by standardized calls to driver routines which are especially written for each I/O system. Computer independent I/O systems are either standardized, or standards are under consideration (CAMAC, GPIB, MEDIA, etc.). Standardized I/O calls designed for these systems are equally valid to the calls presented here but are outside the scope of this International Standard.

11 Scope of the process I/O and general structure of I/O routines

The process peripheral is the link between the process, or its terminal devices, and the central processing unit of the computer. Data describing the space and time behaviour of the process are received by a processing unit and prepared so that they can be transferred to the central processing unit through an I/O interface. The variety of tasks required by the process has resulted in a large number of peripheral devices from different manufacturers. However, in the course of many years of hardware development, largely compatible and generally accepted lines of development have become established. They may be characterized by the following statements.

I/O ports are distinguished by their individual addresses. The I/O port designations (addresses) used in the procedure references will probably, in most systems, be identical to the individual hardware addresses, but this is not mandatory. This relation is considered a processor dependent feature beyond the scope of this International Standard.

ISO 1539 specifies that one statement must be completed before processing of the next statement begins. The standard process I/O described herein adheres to this rule. The calling task will wait for completion and this operating mode is indicated by the last letter W (waiting) of all subroutine names.

The following designations for parameters apply to several of the references. If the exact meaning of these parameter designations deviates from that described below, it will be marked specifically in the detailed description of the reference. If the meaning is exactly as defined here, the description of the parameter will be omitted in the description of the reference and a reference will be made to this clause.

The procedures for process I/O normally have four (in a special case, five) parameters which in the following will be designated in the general form with i , j , k , m (and n if needed). Such an I/O call has the general form

CALL procio(i , j , k , m)

where

procio indicates one of the subroutines subsequently described;

i specifies the number of values to be transferred; it is an integer expression;

j specifies the name of an integer array or array element that contains necessary information for description of the I/O ports, i.e. address and data conversion information. The orderly representation of information is processor dependent;

k specifies the name of an integer array or array element that contains the input or output values;

m is a status indicator. Its value characterizes the "success" of a call, as follows

$m < 0$: undefined,

$m = 1$: all data have been transferred,

$m > 2$: error conditions.

This argument shall be an integer variable or an integer array element.

12 Input/output of analog values

For input, a distinction is made between hardware implemented sequential and random input. In the first case, for sequential input, the input parameter j contains the address of the first analog input; further addresses will be generated automatically. In the second case, the full sequence of addresses must be given in the array j . For output, the form is always random, i.e. all addresses are given in array j .

Generally, the format of j is system dependent. See annex B, clause B.10.

The relationship between the range of an input or output port and its corresponding element in k is processor dependent. Some suggested design guidelines may be found in annex B, clause B.10.

12.1 Sequential analog data input

Execution of a reference to the subroutine AISQW reads a sequence of analog input ports of sequential addresses.

The form of the call is

CALL AISQW(i , j , k , m)

where

i specifies the number of analog input ports to be read. The parameter is an integer expression;

j is a description of either hardware or software information for the acquisition and for the conversion of the first and the following analog ports. It is the name of an integer array or of an element. See annex B, clause B.10;

k is an array for recording the converted analog values. It is the name of an integer array or of an element. See annex B, clause B.10.

m is as described in clause 11.

12.2 Analog data input in random sequence

The subroutine AIRDW reads a sequence of analog input ports in a specified order.

The form of the call is

CALL AIRDW(i,j,k,m)

where

i specifies the number of input ports to be read. The parameter is an integer expression;

j is a description of either hardware or software information for the acquisition and for the conversion of each analog value. It is the name of an integer array. See annex B, clause B.10;

k is an array for recording the converted analog values. See annex B, clause B.10;

m is as described in clause 11.

12.3 Analog data output

The subroutine AOW outputs a sequence of analog values to a collection of analog output ports in a specified order.

The form of the call is

CALL AOW(i,j,k,m)

where

i specifies the number of analog output ports. The parameter is an integer expression;

j contains information for the data conversion and transfer. It is the name of an integer array. See annex B, clause B.10;

k is an array from which the analog values are output. It is the name of an integer array. See annex B, clause B.10;

m is as described in clause 11.

13 Input/output of digital values

For this type of input/output, it is assumed that, while the effective information may be represented at times by a single bit, it will, nevertheless, be necessary to transfer digital values (considered as entities using whole numeric storage units or words) into or out of an integer array (for example 16 bits for each numeric storage unit).

13.1 Digital input

The form of the call is

CALL DIW(i,j,k,m)

where

i specifies the number of digital values input; it is an integer expression;

j contains hardware and in some case software information for conversion and transfer. It is the name of an integer array. A possible reset specification can also be contained in j;

k is an array in which the digital values will be stored. It is the name of an integer array;

m is as described in clause 11.

13.2 Digital output

For the output, pulse output (Digital Output Momentary) is distinguished from digital output with a permanently held value (Digital Output Latching).

13.2.1 Digital pulse output

Execution of a reference to the subroutine DOMW performs pulsed output to a collection of I/O ports. A pulse will occur on those bits selected by a binary 1 of the corresponding bit and element of an array k. No pulse appears on bits selected by binary 0. The pulse duration is indicated in a suitable form by parameter n.

The pulse will begin either when the DOMW is executed or at the next clock tick. This, as well as the representation of the pulse, for example polarity, voltage, etc., depends on the hardware used.

The form of the call is

CALL DOMW(i,j,k,n,m)

where

i specifies the number of digital values output; it is an integer expression;

j contains hardware information for transfer of each digital value. This parameter is the name of an integer array;

k is an array representing the digital values to be output. It is the name of an integer array. The first outputted value will be taken from the first array element (i.e. the element whose index is 1);

n is the number of time units of a clock depending on the hardware used. If the processor does not allow selection of duration, this argument is ignored but must be present. This argument shall be an integer expression;

m is as described in clause 11.

13.2.2 Latched digital output

For latched digital output (DOLW), in addition to the output field, a mask field is also required to indicate which bits are to be changed in the output. The parameter k is therefore subdivided into k1 and k2.

The form of the call is

CALL DOLW(i,j,k1,k2,m)

where

i specifies the number of digital words; it is an integer expression;

j contains hardware information for every digital value that is output. This parameter is the name of an integer array;

k1 is an array representing the digital values to be output. The parameter is the name of an integer array;

k2 designates an array whose values define digital outputs which can be changed by the subroutine. A bit set in the k2 array indicates that the digital output will be changed to the state defined by the corresponding bit position in the corresponding integer array element in k1. The order of the elements in k1 and k2 will correspond to the order in j. This argument shall be an integer array name, or an integer array element;

m is as described in clause 11.

Section four: File handling

14 Introduction

The external procedure references described in this section provide means for controlling the access of files, and also provide means for resolving problems of file access contention in a multitasking/multiprocessing environment. In such an environment, it is expected that concurrent tasks will attempt to access the same file at the same time; therefore, the external procedure references defined herein provide the information necessary for the processor to resolve such simultaneous access in an orderly manner. The method for resolution of access control is left to the processor.

The procedure references in this section are intended to provide the methods by which the task can inform the processor of the manner in which it intends to use the file, but the references are not intended to require specific properties or attributes to be associated with the referenced files. They provide the means to avoid contention problems when used in conjunction with sound program design but the implementation of this International Standard is no assurance that such problems will not arise.

15 Background information

Files exist in most computing systems and can have various attributes and features, such as

- a file can contain data, programs, or catalogue information;
- there can be a variety of ways for file access such as sequential, direct, and stream;
- a file can be created or deleted by a task, by a system utility, or at system generation time;

- a file can have security attributes associated with the file for the purpose of ensuring file privacy;
- when a file is associated with a task, this association can be restricted by the processor for reasons of privacy;
- a file can be associated with a set of related concurrent tasks and this association can be restricted to assure orderly resolution of contention problems among the concurrent tasks;
- a file can be internal or external to a task;
- a file can reside on fixed or removable media;
- a file can reside on main storage or backing storage;
- restrictions for reasons of privacy or contention may apply to a file or a component of a file such as records and data items.

16 File system environment

In industrial real-time computer systems, concurrent task operation with shared resources such as files is a common occurrence. This International Standard does not address all the areas of file management but is concerned with the problems that most commonly arise in industrial real-time computer systems.

Table 2 shows those features covered by this International Standard and those that are excluded; however, the excluded features may affect the result of a request for association of a concurrent task to a file. Such restrictions on association are processor dependent and are outside the scope of this International Standard.

Table 2 — Features and attributes of files

Included	Excluded
Files whose contents are considered to be data	Files whose contents are not considered to be data by the accessing concurrent task
Files which exist on fixed media only or on removable media that are not removed	Files that exist on removable media which are removed
Files that reside in main storage or in backing storage	Files that are internal to a concurrent task
Files that are external to a concurrent task	Creation and deletion of files by a system utility or at system generation
Creation and deletion of files by a concurrent task	Methods of file access
The association of a file to a concurrent task for both system-created and for concurrent task-created files	Restrictions on file access as applied to a component of a file
Restrictions on file access as applied to the file	Attributes of a file for the purpose of ensuring file privacy
The association of a file to a concurrent task irrespective of the method of access (for example direct, sequential, or stream)	

17 Procedures to control file access

17.1 Introduction

The procedure references defined in this section are non-interruptible, that is the processor will only execute one such procedure reference at a time. This requirement ensures that the features described are executed in an orderly manner.

The argument *m*, shown below, shall be set equal to or greater than two (2) in value when the request is not accepted by the executive system. Individual implementation may specify unique values of *m* within the allowable range to designate the specific reason for which the request was rejected.

17.2 Creation of files

Execution of a reference to the subroutine CFILW shall establish, but not open, a named file. Files established by CFILW do not have any privacy attribute to restrict a concurrent task from accessing the files. The contents of a newly created file are not defined by this International Standard.

The form of the call is

CALL CFILW(*j*,*n1*,*n2*,*m*)

where

j specifies the file. The argument is either

- a) an integer expression, or
- b) an integer array name, or
- c) a procedure name, or
- d) a character expression.

The processor shall define which of the above four forms are acceptable;

n1 specifies the number of character storage units per record in this file. This argument shall be an integer expression;

n2 specifies the maximum number of records in this file. This argument shall be an integer expression;

m is set on return to the calling task to indicate the disposition of the request. The value shall be 1 or greater:

- 1 — File successfully created
- 2 or greater — File not created

This argument shall be an integer variable name or integer array element name.

17.3 Deletion of files

Execution of a reference to the subroutine DFILW shall remove a file from the file system. Any file created by the mechanism described in 17.2 can be deleted by the execution of a reference to DFILW, but deletion will not be effected if the file is currently open to any task.

The form of the call is

CALL DFILW(*j*,*m*)

where

j specifies the file. The argument is either

- a) an integer expression, or
- b) an integer array name, or
- c) a procedure name, or
- d) a character expression.

The processor shall define which of the above four forms are acceptable;

m is set on return to the calling task to indicate the disposition of the request. The value shall be 1 or greater:

- 1 — File successfully deleted
- 2 or greater — File not deleted

This argument shall be an integer variable name or integer array element name.

17.4 Opening files

Execution of a reference to the subroutine OPENW shall associate the unit specified by the task with the named file, and shall define the desired access mode of that task to the file.

The form of the call is

CALL OPENW(*i*,*j*,*k*,*m*)

where

i specifies the unit by which the file, named by the argument *j*, is referenced in the task. This argument shall be an integer expression;

j specifies the file. The argument is either

- a) an integer expression, or
- b) an integer array name, or
- c) a procedure name, or
- d) a character expression.

The processor shall define which of the above four forms are acceptable;

k specifies the access mode desired by the task. It is a declaration of the task's intended use of the file. This argument shall be an integer expression. The following values are defined:

- 1 Unlocked — Read/write or write-access is requested by the calling task; other tasks are also allowed the same access,
- 2 Protected read — Read access is requested by the calling task and allowed to other tasks,
- 3 Locked — Read/write or write-access is requested by the calling task. The calling task excludes any file access by other tasks;

m is set on return to the calling task to indicate the disposition of the request. The value shall be 1 or greater:

- 1 — File successfully opened to the calling task
- 2 or greater — File not opened to the calling task

This argument shall be an integer variable name or integer array element name.

The following limitations apply:

If the file is currently open to another task, the disposition of a possible request for a particular access mode will be as follows:

- Unlocked — Fails if another task currently has the file open in the Locked or Protected modes; otherwise succeeds.
- Protected read — Fails if another task currently has the file open in the Locked or Unlocked modes.
- Locked — Fails.

Any attempt to open a file will be successful only if the file exists. If the file was created by a mechanism outside the scope of this International Standard, the attributes given to the file at its creation may restrict the granting of an access mode to the task.

17.5 Closing files

Execution of a reference to the subroutine CLOSEW shall end the task association of the specified logical unit with a named file.

The form of the call is

CALL CLOSEW(*i*,*m*)

where

i specifies the unit. The argument shall be an integer expression;

m is set on return to the calling task to indicate the disposition of the request. The value shall be 1 or greater:

- 1 — File successfully closed to the calling task
- 2 or greater — Non-performance

This argument shall be an integer variable name or integer array element name.

17.6 Modify access mode

Execution of a reference to the subroutine MODAPW shall change the calling task's access mode of a file previously opened by the calling task without closing and reopening the file.

If the calling task does not have access mode to the file, the request fails.

If the request for change cannot be granted, the previous access mode remains in force.

The form of the call is

CALL MODAPW(*i*,*k*,*m*)

where

i specifies the unit. This argument shall be an integer expression;

k specifies the new access mode desired. This argument shall be an integer expression. The following values are defined:

- 1 Unlocked — Read/write or write-access is requested by the calling task; other tasks are also allowed the same access,
- 2 Protected read — Read access is requested by the calling task and allowed to other tasks,
- 3 Locked — Read/write or write-access is requested by the calling task. The calling task excludes any file access by other tasks;

m is set on return to the calling task to indicate the disposition of the request. The value shall be 1 or greater:

- 1 — Access mode requested is granted to the task
- 2 or greater — Access mode before the request remains in force

This argument shall be an integer variable name or integer array element name.

The following limitations apply:

If the file is currently open to another task, the request for a change of mode will fail.

If the file was created by a mechanism outside the scope of this International Standard, the attributes given to the file at its creation may restrict the granting of an access mode to the task.

Annex A

Historical background

(This annex does not form part of the standard.)

A.1 Introduction

The FORTRAN language was originally developed by IBM in 1955. Since that time, FORTRAN has become the most widely used high level language for scientific applications, and large powerful user libraries are available in FORTRAN.

FORTRAN has been standardized internationally in ISO 1539.^[1]

A.2 Special requirements

The wide use and the proved excellence of FORTRAN for scientific applications soon led to its use for industrial real-time systems. These applications require special operations, i.e. real-time operations, bit-string manipulation and facilities for process I/O. These operations can only be accomplished in one of the following two ways:

- a) FORTRAN remains the basic language for arithmetic calculations and for reading and writing of data by standard peripherals. The special operations are realized by elements outside the syntax of FORTRAN;
- b) the complete real-time language remains within the syntax of FORTRAN including the special operations. This means that these operations have to be FORTRAN-subroutines or FORTRAN-functions.

The first approach leads to typical real-time languages which offer the user simple but powerful programming facilities for the special operations. In developing such extensions, the designers try to apply all the features of the real-time operating system used; thus these extensions become dependent on the actual system. This was especially true for the early developments of this kind (see, for example, bibliographic references [4], [5] and [6]). PROCOL, as a later development, avoids this disadvantage; this language was created in France for a series of French computers. It offers very advanced features for real-time programming (see, for example, bibliographic reference [7]).

All languages with extensions outside the syntax of FORTRAN need special compilers for their translation.

In choosing approach b), where the language remains within the syntactical frame of FORTRAN, is gained the advantage that a first compilation and check of the user program can be done on any computer for which a FORTRAN compiler exists. On the other hand, this approach, using subroutines and functions, leads to somewhat clumsy handling of the added CALLs and FUNCTIONs and their associated parameters. Yet this may be considered as a minor inconvenience as the many users of FORTRAN are well accustomed to this kind of programming.

Industrial real-time FORTRAN has also to be compared with three other families of real-time languages

- industrial real-time BASIC;
- languages specifically designed for real-time applications, such as PEARL, RTL/2, etc;
- Ada^[8].

The industrial real-time BASIC languages are very easy to learn and to apply. They are well suited for simple and small problems. They can be implemented relatively easily in large as well as in small computer systems.

The languages of the specific real-time type deliver very powerful programming features to the user. These languages are therefore well suited to large and complex problems. Their implementation (compiler and real-time system) is relatively expensive.

Since 1980, Ada has been a possible alternative. Although Ada is more general than the specific real-time languages, its size and complexity puts it in the same category as far as implementation is concerned. Ada is considered to be very advanced with respect to synchronization (rendezvous); on the other hand, the possibilities for time-based scheduling appear less developed.

The industrial real-time FORTRAN languages implemented by approach a) above are often similar to the specific real-time languages. On the other hand, the industrial real-time FORTRAN languages implemented by approach b) are in many respects between BASIC and the specific real-time language type. Thus, this language offers the user an alternative to these two language facilities. Consequently, a language according to approach b) has to be relative simple; this means that the number and complexity of the additional operations have to remain restricted because of the ease of learning and programming.

A.3 Source of this International Standard

In order to prevent the development of many incompatible real-time languages, the "Workshop on Standardization of Industrial Computer Languages" was founded in 1970 at Purdue University. After a union with another organization in 1973, the workshop was named "International Purdue Workshop on Industrial Computer Systems". The "FORTRAN Committee" of the Purdue Workshop was very active and successful from the beginning. For various and good reasons, this committee chose approach b) above with all special operations being kept within the syntactical frame of the FORTRAN language. Recognizing that it could not develop a single comprehensive standard in a timely manner, the committee decided to develop a series of standards. A program of work was proposed to the Instrument Society of America (ISA) and accepted. ISA formed Technical Committee SP61. The committee developed in a relatively short time a first proposal for real-time FORTRAN which was approved and published by the ISA as ISA Standard S61.1 (1972)^[9]. The standard contains the following groups and numbers of special operations :

- 3 CALLs and the FORTRAN statement STOP for controlling the state of concurrently activated "programs"¹⁾.
- 6 CALLs for process I/O.
- 5 INTEGER FUNCTIONs for bit-string manipulation applied on an INTEGER used as a bit-string.

A second supplementary standard ISA S61.2 (1973)^[10] was published one year later. This standard contains the following groups and numbers of special operations:

- 7 CALLs for handling random unformatted files.
- 1 LOGICAL FUNCTION for testing an individual bit in an INTEGER.
- 2 CALLs for setting and clearing of an individual bit in an INTEGER.

These ISA standards have gained great acknowledgement and, by worldwide use, great importance (the proposal for file handling excepted). The early development and the publication of these standards by the American "FORTRAN Committee" of the Purdue Workshop and the ISA has been an important step for industrial process control.

In the meantime, the ISA standards were developed further and were finally approved as ANSI/ISA standards. Thus, ANSI/ISA S61.1^[11] contains all sections of the old standards except the file handling; the file handling is the subject of ANSI/ISA S61.2 (1978)^[12]. As a significant difference from earlier standards, ANSI/ISA S61.1 (1976) does not allow subroutine calls without wait on the completion of the procedure. The change was made to maintain compatibility with the external procedure reference operation of the host language, FORTRAN^[1].

As ISA S61.1 contains some procedure references which have gained wide acceptance, this standard can be considered as a "basic industrial real time FORTRAN" of the Purdue Workshop. It contains only 3 CALLs and the statement STOP for the management of parallel tasks. ¹⁾ Since more extensive mechanisms are needed, the American "FORTRAN Committee" of the Purdue Workshop has worked intensively on a supplementary proposal concerning the management of parallel tasks, organized as ISA S61.3^[13]. Parallel to this work, the working group "Prozess-FORTRAN" of the German VDI/VDE has developed a complete proposal "Prozess-FORTRAN 75", published as a draft standard of the VDI/VDE^[14].

"Prozess-FORTRAN 75" comprises 11 CALLs for the management of parallel tasks and thus offers a restricted but powerful tool for programming real-time operations. The binary pattern and bit processing is similar to ANSI/ISA S61.1 (1976); merely the description is different. There are additional INTEGER FUNCTIONs for arithmetic and circular shift and a CALL for a bit change. The process I/O is practically identical to ISA.

In ANSI/ISA S61.1 (1976), only the standardization of the analog I/O is performed in the direction of a restrictive standard. The file handling is different from ANSI/ISA S61.2 (1978).

In 1976, the Purdue Europe²⁾ Technical Committee 1 on industrial real-time FORTRAN was founded by members of the VDI/VDE working group "Prozess-FORTRAN" and other specialists in Europe. The committee decided to formulate a complete proposal for industrial real-time FORTRAN.

Up to the end of 1977, the cooperation between the American and European committees was rather loose, with only one joint meeting per year at the International Purdue Workshop meeting. Thus, the American and European committees worked rather independently of each other, with the result that the proposals produced were quite different in several respects. Therefore, ISO/TC97/SC5/WG1, the Purdue Europe Workshop, ECMA TC8, and the International Purdue Workshop all requested the American and European committees to develop proposals that were largely, or preferably completely, identical. In order to reach this objective, cooperation had to be intensified considerably.

1) In this International Standard, the term "task" is used instead of the term "program" used in bibliographic references [9], [10], [11] and [12]. See 2.30.

2) Now EWICS.

From the middle of 1978 onward, most meetings of each organization were attended by at least one member of each of the other two organizations concerned. Through this level of intense international cooperation significant technical progress was made. This resulted in rapid development of this joint standard. Concurrently, to complete its program of work, ISA SP61 developed a draft standard S61.3^[14] that is fully compatible with this International Standard.

To further the common goal of a single international standard for the application of FORTRAN to the control of industrial processes, the ISA took several actions:

- a) It recommended to ANSI that ANS/ISA S61.1^[11] be withdrawn as an item of work for ISO/TC97/SC5/WG1. This action was taken.
- b) It recommended to ANSI that ANS/ISA S61.2^[12] not be proposed as an item of work for WG1.
- c) It decided that the draft standard S61.3^[13], if approved, would not be forwarded to ANSI for processing as a national standard.

This International Standard will need to be adapted from time to time to allow the progress of FORTRAN and also the progress of real-time operating systems. Thus supplements must be expected.

STANDARDSISO.COM : Click to view the full PDF of ISO 7846-1985

Annex B

Further descriptions pertaining to individual clauses

(This annex does not form part of the standard.)

B.1 Definitions and naming rules

Some of the definitions given in clause 2 are inspired by or borrowed from such literature as bibliographic references [15], [16], and other sources. The definitions correspond largely to those commonly used in the International Purdue Workshop and are identical or nearly identical to definitions of ANSI/ISA S61.2^[12] and proposed standard ISA S61.3.^[13] Definitions of the states refer to the figure of section one and its description in clause 5 and may be clearer if the figure and its description are consulted.

The names of the calls in this International Standard have been chosen according to a few basic rules, as follows:

- a) The names should be descriptive — giving an indication of their use.
- b) They should be different from calls in any of the ISA standards S61.1, S61.2 and S61.3, if such a call is not completely identical to a call in the other standards. In the latter case, an identical name is desirable, and only then.
- c) The names should not be too common, otherwise they might easily coincide with names naturally selected by users for certain application programs.

Rule b) has the desired effect that an implementor may choose to keep calls already implemented according to an older standard, and existing user programs will not become obsolete just because of the new implementation.

B.2 Use of CALL CLOCK (see 4.2)

The call CALL CLOCK is useful, first of all for applications requiring a somewhat lower level of programming. All three parameters are output arguments, *j* (the first in line) supplying the instantaneous value of the system clock. This value directly represents the contents of the counting register of the clock.

The next two arguments supply system constants necessary and sufficient to achieve system independence in the use of this call, to allow the user to preserve portability and system independence for his programs.

The first of these arguments, *k*₁, supplies the basic clock frequency, whereas the last argument informs about the modulo of the clock, i.e. it reveals the number of bits of the clock register. The modulo is equal to the value of *k*₂ + 1. One clock tick after the value of *j* is equal to *k*₂, *j* will be reset to zero and start counting upwards again.

B.3 Choice of task model (see clause 5)

Commonly, in many implementations, a call to the executive system (by calls of SKED, STRTAF, STRTAT, etc.) for state transition to PENDING causes listing in a table, appropriately called "pending-table". This is a queue of transition calls waiting to be effected, i.e. intended to cause state transitions some time later. All entries in this table involve some conditions for future activations, and these conditions are either time-dependent (STRTAF, STRTAT) or event-dependent (CON). Entries with time conditions will generally be sorted according to the activation time. Often, executive system implementors will find it convenient to split this queue into two: a time queue and an event queue. With an ordered time queue, the executive system's task of monitoring when time is due for any task will be simple: it has only to check the first element, since this always represents the next task to be activated. The other part of the "pending-table" contains entries made by reference to the subroutine CON. All these elements refer to events as their conditions. The executive system will check this queue when external interrupts are received, as well as at some other instances, in order to be able to react to possible programmed events, for example eventmarks set to ON by reference to subroutine POST. Similar tables may exist to direct the supervision of suspended tasks; in fact, often these two sets of tables will be combined.

Transition calls such as those mentioned will cause entry in the "pending-table" up to the point where the maximum table space is exhausted. When this occurs, the transition call will be rejected with the error parameter (termed *m* in clause 6) indicating this fact by being set to a value greater than one.