

## Annex J (informative)

### Interrupt Control Considerations

1

#### 2 **J.1 Introduction**

3 In the area of realtime systems there often exist devices with very simple inter-  
4 faces, that can or should be operated without a full-fledged driver within the  
5 operating system: what is needed is the ability to access the control and status  
6 registers of the interface, and the ability to capture the interrupts that are gen-  
7 erated and have the application program handle them.

8 The functions defined in this section allow a process or thread to capture an inter-  
9 rupt, to block awaiting the arrival of an interrupt, and to protect critical sections  
10 of code which are contended for by a user-written interrupt service routine. Cap-  
11 turing an interrupt involves registering a user-written interrupt service routine  
12 (ISR). The introduction of user-written ISRs does not make application programs  
13 completely portable, but at least establishes a reference model that allows pro-  
14 grams to be rehosted without completely subverting their logic, and confines  
15 non-portable code to specified modules.

16 A single threaded process, or a process in an implementation which does not sup-  
17 port threads, is considered to consist of a single thread of control; in this case, the  
18 term *thread* in the description of the interfaces in this section shall refer to this  
19 single thread of control.

#### 20 **J.2 Definitions**

21 ⇒ **2.2.2 General Terms** *Add the following definitions, in the right sorted order:*

##### 22 **J.2.0.1 interrupt:**

23 (1) The suspension of a process to handle an event external to the process. *Syn:*  
24 **interruption**. *See also:* **interrupt latency; interrupt mask; interrupt**  
25 **priority; interrupt service routine**. (2) To cause the suspension of a process.  
26 (3) Loosely, a hardware interrupt request.

27 **J.2.0.2 interruption:**

28 *See:* **interrupt.**

29 **J.2.0.3 interrupt latency:**

30 The delay between a computer system's receipt of a hardware interrupt request  
31 and its handling of the request. *See also:* **interrupt priority.**

32 **J.2.0.4 interrupt mask:**

33 A mask used to enable or disable interrupts by retaining or suppressing bits that  
34 represent interrupt requests.

35 **J.2.0.5 interrupt priority:**

36 The importance assigned to a given interrupt request. This importance deter-  
37 mines whether the request will cause suspension of the current instructions and,  
38 if there are several outstanding interrupt requests, which will be handled first.

39 **J.2.0.6 interrupt request:**

40 An external or other hardware input that requests the execution of the current  
41 instruction flow be suspended to permit execution of an ISR.

42 **J.2.0.7 interrupt service routine:**

43 A routine that responds to interrupt requests by storing the contents of critical  
44 registers, performing the processing required by the interrupt request, and then,  
45 if no higher priority process is eligible to run, restoring the register contents and  
46 restarting the interrupted process.

47 **J.2.0.8 ISR:**

48 Abbreviation for interrupt service routine.

49 **J.3 Concepts**

50 The following opaque data type is defined by the implementation in the header  
51 `<intr.h>`.

52 `intr_t`

53 Identifies the source of a hardware interrupt in an implementation-defined  
54 manner. The implementation shall also supply some means of obtaining  
55 legal values of type `intr_t`, that represent supported interrupt sources a pro-  
56 cess may connect to.

57           The header `<intr.h>` shall define the following symbols:

58           POSIX\_INTR\_PRI\_LOWEST

59                   The lowest available interrupt priority.

60           POSIX\_INTR\_PRI\_LOW

61                   A low available interrupt priority.

62           POSIX\_INTR\_PRI\_MED\_LOW

63                   A medium low available interrupt priority.

64           POSIX\_INTR\_PRI\_MEDIUM

65                   A medium available interrupt priority.

66           POSIX\_INTR\_PRI\_MED\_HIGH

67                   A medium high available interrupt priority.

68           POSIX\_INTR\_PRI\_HIGH

69                   A high available interrupt priority.

70           POSIX\_INTR\_PRI\_HIGHEST

71                   The highest available interrupt priority.

## 72   **J.4 Interrupt Control Functions**

### 73   **J.4.1 Associate a User-Written ISR with an Interrupt**

74   Functions: `posix_intr_associate()`, `posix_intr_disassociate()`, `posix_intr_lock()`,  
75   `posix_intr_unlock()`.

#### 76   **J.4.1.1 Synopsis**

```
77   #include <intr.h>
78   int posix_intr_associate (intr_t intr,
79                            int (*intr_handler)(void *area),
80                            volatile void *area, size_t arealize);

81   int posix_intr_disassociate (intr_t intr,
82                                int (*intr_handler)(void *area));

83   int posix_intr_lock (intr_t intr);

84   int posix_intr_unlock (intr_t intr);
```

#### 85   **J.4.1.2 Description**

86   If the Interrupt Control option is supported:

B

87 If the number of ISRs currently connected to *intr* is less than `{_POSIX_-`  
 88 `INTR_CONNECT_MAX}` then the *posix\_intr\_associate()* function shall associ-  
 89 ate the given user-written ISR *intr\_handler* with a given interrupt *intr*. The  
 90 interrupt service routine shall conform to the following function prototype:

```
91 int int_handler (void *area);
```

92 After executing the *posix\_intr\_associate()* function, the issuing thread shall  
 93 become connected to the given interrupt. The system shall call  
 94 *intr\_handlers* in the reverse order that the ISRs were registered until one of  
 95 the *intr\_handlers* returns a code signifying that the interrupt has been han-  
 96 dled. The most recently registered ISR is thus called first.

97 Although an ISR is initially located in a process' address space, it executes  
 98 in an implementation-defined context, subject to a number of  
 99 implementation-defined restrictions. It is unspecified what restrictions  
 100 may be imposed by an implementation.

101 The execution context of an ISR may have an address space different from  
 102 the normal process' space, so any data areas accessed by the ISR must be  
 103 clearly identified as such. The argument *area* identifies a communication  
 104 region, whose size is *areasize*, where the ISR and the thread shall be able to  
 105 exchange data; when the ISR is called, it shall receive the address of the  
 106 communication region as its first argument. Implementations may have  
 107 additional arguments of implementation-defined types.

108 The return code which is returned by the ISR determines whether the inter-  
 109 rupt has been handled by this ISR, and whether the thread that registered  
 110 the ISR should be notified of the successful handling of this interrupt.

111 An interrupt handler wakes up a thread by posting one of the *notify* return  
 112 codes, which causes a thread waiting in a corresponding  
 113 *posix\_intr\_timedwait()* to unblock. No other ISR-to-thread notification  
 114 mechanism is specified.

115 Notification is described in clause C.3.2. Possible return codes (defined in  
 116 `<intr.h>`) include:

```
117 POSIX_INTR_HANDLED_NOTIFY
```

118 The ISR handled this interrupt, and the thread that registered the  
 119 ISR should be notified that the interrupt occurred.

```
120 POSIX_INTR_HANDLED_DO_NOT_NOTIFY
```

121 The ISR handled this interrupt, but the thread that registered the  
 122 ISR should not be notified that the interrupt occurred.

```
123 POSIX_INTR_NOT_HANDLED
```

124 The ISR did not handle this interrupt; if there are other ISRs con-  
 125 nected to this interrupt, then the next ISR should be called.

126 The *posix\_intr\_disassociate()* function shall cancel any existing association  
 127 between the interrupt *intr* and the ISR *interrupt\_handler*.

128 If a thread calls *posix\_intr\_disassociate()* and the thread does not have the  
 129 specified ISR registered for the specified interrupt, the

Copyright © 1998 IEEE. All rights reserved.

This is an unapproved IEEE Standards Draft, subject to change.

130 *posix\_intr\_disassociate()* function shall fail.

131 If a thread has connected one or more user-written ISRs to a given inter-  
 132 rupt *intr*, then that thread calling the *posix\_intr\_lock()* function shall  
 133 prevent the system from calling those ISRs or notifying the connected  
 134 thread until delivery is re-enabled by means of the *posix\_intr\_unlock()* func-  
 135 tion, thus allowing the thread to perform operations in an atomic way with  
 136 respect to the ISR; if an ISR is executing when the thread that connected  
 137 that ISR to an interrupt calls *posix\_intr\_lock()*, then the *posix\_intr\_lock()*  
 138 function shall not return until that ISR has completed; the methods used by  
 139 an implementation to obtain these results are implementation-defined. To  
 140 allow implementation using a hardware disable-interrupts instruction,  
 141 *posix\_intr\_lock()* need not be a cancellation point. It is  
 142 implementation-defined whether locking an ISR causes other ISRs to be  
 143 locked. It is implementation-defined whether interrupts that arrive while  
 144 an interrupt is locked are queued or discarded. It is  
 145 implementation-defined whether registration under lock is supported.

146 It is implementation-defined whether these functions require an appropri-  
 147 ate privilege from the calling thread.

148 It is implementation-defined whether ISRs remain registered after the  
 149 registering thread terminates. It is implementation-defined which POSIX  
 150 operations, if any, may be executed from an interrupt handler.

151 Otherwise:

152 Either the implementation shall support the *posix\_intr\_associate()*,  
 153 *posix\_intr\_disassociate()*, *posix\_intr\_lock()*, and *posix\_intr\_unlock()* func-  
 154 tions as described above, or these functions shall not be provided. B

### 155 **J.4.1.3 Returns**

156 Upon successful completion, *posix\_intr\_associate()*, *posix\_intr\_disassociate()*,  
 157 *posix\_intr\_lock()*, and *posix\_intr\_unlock()* shall return zero. Otherwise an error  
 158 code shall be returned.

### 159 **J.4.1.4 Errors**

160 If any of the following conditions occur, the *posix\_intr\_associate()*,  
 161 *posix\_intr\_disassociate()*, *posix\_intr\_lock()*, and *posix\_intr\_unlock()* function shall  
 162 return the corresponding non-zero error code:

163 [EINVAL] The *intr* argument does not identify a supported interrupt that  
 164 can be connected to a user-specified ISR. B

165 B

166 [EPERM] The calling thread does not have an appropriate privilege to call  
 167 this function and the implementation requires such a privilege.

168 If the following condition occurs, it is implementation-defined whether the  
 169 *posix\_intr\_associate()* function shall return the corresponding non-zero error code:

Copyright © 1998 IEEE. All rights reserved.

This is an unapproved IEEE Standards Draft, subject to change.

170 [EAGAIN] The interrupt identified by the *intr* argument currently has the  
171 implementation-defined maximum number of ISRs connected.

172 If the following condition occurs, the *posix\_intr\_disassociate()*, *posix\_intr\_lock()*,  
173 and *posix\_intr\_unlock()* functions shall return the corresponding non-zero  
174 error code:

175 [ENOISR] The thread has not registered an ISR for the given interrupt.

176 If the following condition is detected, the *posix\_intr\_associate()* function shall  
177 return the corresponding non-zero error code:

178 [EINVAL] The arguments *area* and/or *areasize* and/or *intr\_priority* are  
179 invalid for the implementation.

#### 180 J.4.1.5 Cross-References

### 181 J.4.2 Await Interrupt Notification

182 Function: *posix\_intr\_timedwait()*.

#### 183 J.4.2.1 Synopsis

184 #include <intr.h>

185 int posix\_intr\_timedwait (int *flags*, const struct timespec \**timeout*);

#### 186 J.4.2.2 Description

187 If the Interrupt Control option is supported:

188 The *posix\_intr\_timedwait()* function causes the calling thread to block until  
189 notified that an interrupt has occurred. If an interrupt notification was  
190 delivered to the calling thread prior to the call to the  
191 *posix\_intr\_timedwait()* function, and this notification has not previously  
192 caused a call to the *posix\_intr\_timedwait()* function to be unblocked, then  
193 the calling thread is not blocked and instead the *posix\_intr\_timedwait()*  
194 function returns immediately.

195 The input argument *flags* contains only implementation-defined input  
196 values.

197 If the value of the *timeout* input argument is non-null, the wait for an inter-  
198 rupt to occur shall be terminated when the specified timeout period expires.

199 If the *timeout* input argument is null, the wait is terminated only by the  
200 interrupt.

201 The timeout expires after the interval specified by *timeout* has elapsed since  
202 the wait began. If the Timers option is supported, the timeout is based on  
203 the CLOCK\_REALTIME clock; if the Timers option is not supported, the  
204 timeout is based on the system clock as returned by the POSIX.1 function  
205 *time()*. The resolution of the timeout is determined by the resolution of the

B

B

B

206 clock that it uses.

207 Under no circumstance will the function fail with a timeout if the interrupt  
 208 notification occurred prior to the *posix\_intr\_timedwait()* call. The validity  
 209 of the *timeout* argument need not be checked if the interrupt notification  
 210 occurred prior to the *posix\_intr\_timedwait()* call. Invocation of  
 211 *posix\_intr\_timedwait()* shall implicitly release an *posix\_intr\_lock()*.

212 This function shall fail if no ISR is currently registered by the calling  
 213 thread.

214 Otherwise:

215 Either the implementation shall support the *posix\_intr\_timedwait()* func-  
 216 tion as described above or this function shall not be provided. B

### 217 J.4.2.3 Returns

218 Upon successful completion, *posix\_intr\_timedwait()* shall return zero. Otherwise  
 219 an error code shall be returned.

### 220 J.4.2.4 Errors

221 If any of the following conditions occur, the *posix\_intr\_timedwait()* function shall  
 222 return the corresponding non-zero error code.

223 [EINVAL] The thread would have blocked, but the timeout argument  
 224 specified a nanoseconds value less than zero or greater than or  
 225 equal to 1000 million.

226 [EINTR] A signal interrupted this function.

227 [ENOISR] The thread has not registered an ISR for the given interrupt.

228

229 [ETIMEDOUT] The interrupt notification was not received before the specified  
 230 timeout expired. B

### 231 J.4.2.5 Cross-References

## 232 J.5 Rationale for Interrupt Control

### 233 J.5.1 The Interrupt Model

#### 234 J.5.1.1 Background

235 The purpose of this interface is to allow connection of non-standard  
 236 interrupt-generating hardware in a standard way.

237 Such hardware, when enabled, may generate a continuous stream of interrupts,  
 238 not following the request-response model typical of common I/O devices such as

Copyright © 1998 IEEE. All rights reserved.  
 This is an unapproved IEEE Standards Draft, subject to change.

239 disks and tapes. A typical example would be a radar antenna generating a stream  
240 of azimuth-change and north-crossing interrupts as the antenna rotates. Another  
241 example would be the stream of angle resolver pulses from the joints of a robot.

242 Many hardware architectures have fewer interrupts than peripheral devices,  
243 requiring interrupts to be shared. In such cases, more than one ISR is invoked by  
244 a given interrupt, and the identity of the device or devices generating an interrupt  
245 must be established by polling the devices connected to that interrupt.

246 Some kinds of non-standard hardware generates multiple and related streams of  
247 interrupts which should all be handled by a single thread. Again, a radar  
248 antenna, with its two interrupt streams, provides a classic example.

#### 249 **J.5.1.2 The Model**

250 To each suitable interrupt one may attach zero or more interrupt service routines  
251 (ISRs). When the interrupt is activated, these ISRs are executed in reverse order  
252 of registration; that is, last registered, first executed.

253 Each ISR must first poll its device to determine if that device is asserting the  
254 interrupt. If not, the ISR returns immediately with a return value signifying that  
255 the interrupt was not handled. If the ISR's device is asserting the interrupt, the  
256 ISR does whatever is needed (a matter of local design), and returns with a return  
257 value signifying that the interrupt has been handled, and further that the regis-  
258 tering thread should or should not be notified (awakened). The decision to notify  
259 or not is made by the user-written ISR code, but notification is performed by the  
260 vendor supplied code which invokes these ISRs.

261 Note that there are no direct error returns from these ISRs to the invoking kernel;  
262 device error handling is performed within the ISRs and reported as needed in the  
263 communications area.

264 The first ISR to handle an interrupt consumes it, preventing execution of ISRs in  
265 the remainder of the list. If there are two devices simultaneously asserting the  
266 interrupt, the second device will continue to assert the interrupt, forcing re-  
267 traversal of the ISR chain, from the top.

268 If no ISR in an interrupt chain claims an interrupt, the behavior is unspecified.  
269 Typically, unclaimed interrupts are simply ignored.

270 Multiprocessors are implicitly handled, depending on the underlying hardware.  
271 In many systems, one associates each device with a processor that will handle its  
272 interrupts. In other systems, such as recent ones from Sun Microsystems, the ISR  
273 (a kernel thread) runs on any available processor.

274 Interrupt handling has no effect on scheduling queues, except that an interrupt  
275 can result in the unblocking of a process whose priority exceeds that of any that  
276 were executing when that interrupt arrived.



### 277 **J.5.1.3 Registration**

278 ISRs become connected to interrupts by registration, and disconnected by  
279 de-registration. A thread registering an ISR provides four pieces of information:  
280 the address of the ISR code, the address and size of the communication region of  
281 memory (used for data shared by ISR and the registering thread), the interrupt ID,  
282 and (implicitly) the thread ID of the registering thread.

283 A thread may have multiple ISRs registered, and each ISR may generate  
284 notification requests. The thread waits for any and all such notifications in one  
285 place, using the *posix\_intr\_timedwait()* function. In such cases, the ISRs must  
286 place whatever information is needed for the application to tell one device from  
287 another in their respective communication regions. It is necessary to have pre-  
288 cisely one wait-point to prevent deadlocks due to interrupts arriving in an unex-  
289 pected order.

290 To implement smooth and leakproof transfer of interrupt traffic from one ISR to  
291 another, it is sufficient to register the new ISR before deregistering the old ISR.  
292 For the short period of time that there are two ISRs for one device, the most  
293 recently registered ISR will consume all the traffic, allowing the old ISR to be dere-  
294 gistered at leisure.

295 The communication region is an area of memory that is visible both to the ISR  
296 during an interrupt and to the registering thread at all (other) times. It is the  
297 user's responsibility to allocate a suitable region, perhaps by the use of the Typed  
298 Memory facilities provided by the implementation. Shared access to the commun-  
299 ications region by both thread and ISR is mediated using *posix\_intr\_lock()* and  
300 *posix\_intr\_unlock()*.

### 301 **J.5.2 Portability**

302 Although interrupt handling isn't entirely portable, there is still profit in standar-  
303 dizing the interrupt control interface. First is the implicit standardization of core  
304 functionality. Second is programmer portability. Third is that interrupt handling  
305 code can follow the hardware device for which it was written. All of this is sup-  
306 ported by a great deal of embedded and/or realtime (often non-UNIX) system prac-  
307 tice. The resulting modularization and isolation of non-portable code also aids  
308 portability.

309 The model which repeatedly emerges from existing practice involves two facilities.  
310 First, the application should be able to arrange for an interrupt to notify a process  
311 or thread when the interrupt occurs. Second, the application should also be able  
312 to provide an interrupt handler or interrupt service routine (ISR) to immediately  
313 service each interrupt occurrence without a time-consuming process context  
314 switch. Users of this model may require one or the other or both of these facili-  
315 ties. Clause C.4.9 shows an example using the interfaces in this section to imple-  
316 ment an application which conforms to this model.

### 317 **J.5.3 Existing Practice**

318 Most operating systems designed specifically to support realtime applications pro-  
319 vide similar services for application interrupt handling and control. Ada language  
320 runtime environments which support interrupt handling also provide similar ser-  
321 vices. The fact that UNIX systems have not typically provided such services is  
322 indicative of the non-realtime heritage of UNIX. See also **Portability** above.

#### 323 **J.5.3.1 Interrupt Specification**

324 There is no portable way to specify an interrupt. Therefore, these interfaces must  
325 rely on an opaque type, `intr_t`, to identify a specific interrupt. Each implementa-  
326 tion which supports the Interrupt Control option must provide a mechanism for  
327 obtaining objects of this type which identify all user accessible interrupts sup-  
328 ported by the implementation. Such mechanisms may include constant objects of  
329 type `intr_t`, and functions, macros, typecasts which involve implementation specific  
330 interrupt identification procedures, returning objects of type `intr_t`. Once such an  
331 object has been associated with an interrupt of interest, this interrupt may be  
332 accessed via the portable interfaces in this section. The implementation-specific  
333 mechanisms cannot and will not be standardized herein.

#### 334 **J.5.4 Interrupt Latency**

335 Connecting an interrupt to a POSIX realtime signal, while performing the neces-  
336 sary function of initiating application processing, often cannot alone guarantee  
337 adequate and timely response to rapid and/or time critical interrupts. There are  
338 two reasons for this. First, the sequence of execution of POSIX processes and/or  
339 threads is a function of the CPU scheduling policy, not the relative urgency of vari-  
340 ous interrupts; it is possible to work within the constraints of the scheduling pol-  
341 icy, but interrupt response and handling time is still likely to be  
342 non-deterministic. Second, even if the notified process becomes the running pro-  
343 cess immediately upon interrupt occurrence, the overhead necessary for process  
344 context switching is unlikely to support interrupt latency requirements in the ten  
345 to hundred microseconds range, or for interrupts occurring at a rapid rate. Also,  
346 many interrupting devices require execution of special code to respond to and  
347 deassert each and every interrupt.

348 The purpose of `posix_intr_associate()`, therefore, is to provide a path to first level  
349 interrupt servicing code whose latency is a function only of other interrupts occur-  
350 ring at or near the same time. Such code is intended to deal with the high speed  
351 portion of the interrupt handling. It should perform only functions which cannot  
352 be postponed until they are handled by a normally scheduled process; it typically  
353 executes with at least the interrupt of interest locked out, and therefore need not  
354 be reentrant. To achieve this low latency, it is expected that an implementation  
355 will bind ISR code as closely as possible to the hardware interrupt mechanism.  
356 The issues of response time (timeliness) and mutual exclusion (ISR,  
357 non-reentrant) are independent.

### 358 **J.5.5 Relationship To Realtime Profiles**

359 Handling of interrupts by user written code is typical in applications conforming  
360 to the two smaller realtime profiles from IEEE 1003.13, the *Minimal* and *Control*  
361 profiles. Systems which require these profiles typically utilize neither virtual  
362 memory nor an architecture supporting separate user and kernel modes. In such  
363 systems, interrupt servicing via the ISR model is easily implemented as simple  
364 procedure call in the context of the single executing process.

365 For the two more complex realtime profiles in IEEE 1003.13, the *Dedicated* and  
366 *Multi-Purpose* profiles, process and kernel separation is of concern to most imple-  
367 mentations, and the ISR model becomes somewhat more difficult to implement.  
368 Although, systems requiring these profiles are far more likely to utilize full scale  
369 device drivers integrated or loaded into the kernel in an implementation specified  
370 manner, some may require the interfaces of this section, especially  
371 *posix\_intr\_associate()*.

### 372 **J.5.6 Limitations Imposed on ISR Code**

373 An ISR needs to be able to execute without incurring unpredictable delays; it must  
374 complete in a timely manner. For this reason an ISR is normally restricted to ker-  
375 nel level calls since the boundary of kernel level calls' behavior is strongest. Invo- B  
376 cation of an ISR has the implicit effect of a call to *posix\_intr\_lock()* to block further  
377 invocations of that same ISR. Consequently, ISRs need not be reentrant. There is,  
378 therefore, no free choice on the interfaces that can be used and caution needs to be  
379 exercised in selecting interfaces. Invocation of POSIX interfaces within an ISR  
380 cannot be considered consistent with the timeliness requirement of an ISR.

381 The notion of the "current process" is erroneous within an ISR since it may exe-  
382 cute in the context of the kernel or of an arbitrary process. Since most POSIX  
383 interfaces may query or alter the state of the "current process", their use within  
384 ISR code is not only untimely, but erroneous.

385 It is expected that anyone wanting to write a user-written ISR is familiar with how  
386 to write I/O drivers for their particular system. Mistakes in an ISR can and will  
387 crash the system. Caution is advised.

### 388 **J.5.7 Handler Specification**

389 The working group is undecided on how an interrupt handler (not executing in  
390 process context) may be specified to the *posix\_intr\_associate()* function. Since the  
391 handler code may need to ultimately run in a kernel address mapping different  
392 from that of the process doing the registration, specifying the virtual address of  
393 non-relocatable code would appear problematic. The extent of the handler (i.e.  
394 what functions it can call or data it can access) is also unclear in the virtual  
395 memory case. Finally, at least one example of existing practice requires that a  
396 handler be dynamically loaded into the kernel from a file; in this case, a pathname  
397 would be required as the handler specification.

398 The working group opposes converting the handler specification argument of  
 399 *posix\_intr\_associate()* to an opaque type. It has been suggested that a pathname  
 400 be used for this argument, and this pathname could identify a memory resident  
 401 code segment on systems such as those conforming to the minimal realtime profile  
 402 (like the pathname interpretation for *exec()* or *posix\_spawn()* on such systems).  
 403 Unless the working group can agree on a workable method of handler  
 404 specification, the original method; specifying a pointer to a function, will be  
 405 retained.

#### 406 **J.5.8 *posix\_intr\_timedwait()* versus *Sigwait()***

407 Why doesn't *posix\_intr\_timedwait()* precisely follow the existing *sigwait()* model?  
 408 This was discussed in the working group, which for simplicity decided not to  
 409 attempt to map hardware interrupts onto signals, as the underlying models are  
 410 only in general similar, differing greatly in the details. The great mass of existing  
 411 code renders the signals model and *sigwait()* API essentially immutable, so it was  
 412 decided to make *posix\_intr\_timedwait()* independent of *sigwait()*, and as simple as  
 413 possible, to reduce the burden on implementors, and for performance and predic-  
 414 tability.

B  
B

#### 415 **J.5.9 Interrupt Specification**

416 Several examples may serve to clarify the use of *intr\_t*:

#### 417 **Figure J-1 – *intr\_t* Examples**

418

##### 419 **An implementation supplied constant:**

```
420 posix_intr_associate (IVEC_240, &handler, &data, sizeof(data) );
```

##### 421 **Asking a device how it will interrupt:**

```
422 posix_devctl (fd, GET_INTERRUPT_ID, &interrupt, sizeof(intr_t), NULL);  
423 posix_intr_associate (interrupt, &handler, &data, sizeof(data) );
```

##### 424 **Simple cast of known interrupt vector:**

```
425 posix_intr_associate ((intr_t)240, &handler, &data, sizeof(data) );
```

426

427 **J.5.10 Application Example**

428 The following C-Language program fragment demonstrates usage of the interfaces  
429 in this section:

430 **Figure J-2 – An Interrupt Control Application Fragment**

```

431  /* Collect digitized data to a file - The A to D converter runs */
432  /* at 30khz sampling rate, has a 256-sample circular buffer, */
433  /* and interrupts after each 128 samples.                               */
434  queue my_queue; /* statically allocated */
435  main()
436      {
437      int my_handler(queue *my_queue);
438      posix_intr_associate(INTR_240, &my_handler, &my_queue, sizeof(my_queue));
439
440      start_A_to_D();
441      while (TRUE)
442          {
443              int localbuffer[128];
444              if (posix_intr_timedwait(0, A_to_D_timeout()) == 0)
445                  {
446                      posix_intr_lock(INTR_240);
447                      while (dequeue(&my_queue, localbuffer))
448                          {
449                              posix_intr_unlock(INTR_240);
450                              fwrite(localbuffer, 1, sizeof(localbuffer),
451                                  stdout);
452                              posix_intr_lock(INTR_240);
453                          }
454                      posix_intr_unlock(INTR_240);
455                  }
456              else
457                  /*handle errors, including timeout*/
458          }
459
460  int my_handler(queue *my_queue)
461      {
462      int localbuffer[128];
463      read_A_to_D(localbuffer);
464      enqueue(my_queue, localbuffer);
465      return POSIX_INTR_HANDLED_NOTIFY;
466      }

```

Copyright © 1998 IEEE. All rights reserved.  
This is an unapproved IEEE Standards Draft, subject to change.

Copyright © 1998 IEEE. All rights reserved.  
This is an unapproved IEEE Standards Draft, subject to change.