

[illegible]

```
MM      MM      PPPPPPPP      SSSSSSSSS      CCCCCCCC      HH      HH      EEEEEEEEEEE      DDDDDDDDD
MM      MM      PPPPPPPP      SSSSSSSSS      CCCCCCCC      HH      HH      EEEEEEEEEEE      DDDDDDDDD
MMMM     MMMM     PP      PP      SS      CC      CC      HH      HH      EE      DD      DD
MMMM     MMMM     PP      PP      SS      CC      CC      HH      HH      EE      DD      DD
MM      MM      MM      PP      PP      SS      CC      CC      HH      HH      EE      DD      DD
MM      MM      MM      PP      PP      SS      CC      CC      HH      HH      EE      DD      DD
MM      MM      PPPPPPPP      SSSSSSS      SS      CC      CC      HHHHHHHHHH      EEEEEEEEE      DD      DD
MM      MM      PPPPPPPP      SSSSSSS      SS      CC      CC      HHHHHHHHHH      EEEEEEEEE      DD      DD
MM      MM      PP      SS      CC      CC      CC      HH      HH      EE      DD      DD
MM      MM      PP      SS      CC      CC      CC      HH      HH      EE      DD      DD
MM      MM      PP      SS      CC      CC      CC      HH      HH      EE      DD      DD
MM      MM      PP      SSSSSSSSS      CCCCCCCC      HH      HH      EEEEEEEEEEE      DDDDDDDDD
MM      MM      PP      SSSSSSSSS      CCCCCCCC      HH      HH      EEEEEEEEEEE      DDDDDDDDD

LL      IIIIIII      SSSSSSSSS
LL      IIIIIII      SSSSSSSSS
LL      II      SS
LL      II      SS
LL      II      SS
LL      II      SS
LL      II      SSSSSSS
LL      II      SSSSSSS
LL      II      SS
LL      II      SS
LL      II      SS
LL      II      SS
LLLLLLLLLLLL      IIIIIII      SSSSSSSSS
LLLLLLLLLLLL      IIIIIII      SSSSSSSSS
```


MPSCHED
Table of contents

- MULTIPROCESSOR SCHEDULER

L 15

16-SEP-1984 02:07:03 VAX/VMS Macro V04-00

Page 0

(1) 136
(1) 291
(1) 390

MPSSRESCHEDIPL5 - MF RESCHEDULING (MA780) INTERRUPT HANDLER
MPSSRESCHED - RESCHEDULING INTERRUPT HANDLER
MPSSMPSCHED - SECONDARY PROCESSOR SCHEDULING

```
0000 1  :
0000 2  : Version:      'V04-000'
0000 3  :
0000 4  :
0000 5  :      .MCALL MFPR
0000 6  :      .TITLE MPSCHED - MULTIPROCESSOR SCHEDULER
0000 7  :      .IDENT 'V04-000'
0000 8  :
0000 9  : *****
0000 10 : *
0000 11 : * COPYRIGHT (c) 1978, 1980, 1982, 1984 BY
0000 12 : * DIGITAL EQUIPMENT CORPORATION, MAYNARD, MASSACHUSETTS.
0000 13 : * ALL RIGHTS RESERVED.
0000 14 : *
0000 15 : * THIS SOFTWARE IS FURNISHED UNDER A LICENSE AND MAY BE USED AND COPIED
0000 16 : * ONLY IN ACCORDANCE WITH THE TERMS OF SUCH LICENSE AND WITH THE
0000 17 : * INCLUSION OF THE ABOVE COPYRIGHT NOTICE. THIS SOFTWARE OR ANY OTHER
0000 18 : * COPIES THEREOF MAY NOT BE PROVIDED OR OTHERWISE MADE AVAILABLE TO ANY
0000 19 : * OTHER PERSON. NO TITLE TO AND OWNERSHIP OF THE SOFTWARE IS HEREBY
0000 20 : * TRANSFERRED.
0000 21 : *
0000 22 : * THE INFORMATION IN THIS SOFTWARE IS SUBJECT TO CHANGE WITHOUT NOTICE
0000 23 : * AND SHOULD NOT BE CONSTRUED AS A COMMITMENT BY DIGITAL EQUIPMENT
0000 24 : * CORPORATION.
0000 25 : *
0000 26 : * DIGITAL ASSUMES NO RESPONSIBILITY FOR THE USE OR RELIABILITY OF ITS
0000 27 : * SOFTWARE ON EQUIPMENT WHICH IS NOT SUPPLIED BY DIGITAL.
0000 28 : *
0000 29 : *****
0000 30 :
0000 31 : ++
0000 32 :
0000 33 : Facility: Executive , Hardware fault handling
0000 34 :
0000 35 : Abstract: Multi-processing scheduler routines
0000 36 :
0000 37 : Environment: MODE=Kernel
0000 38 :
0000 39 : Author: RICHARD I. HUSTVEDT, Creation date: 15-MAY-1979
0000 40 :
0000 41 : Modified by:
0000 42 :
0000 43 : V03-007 KDM0032 Kathleen D. Morse 19-Nov-1982
0000 44 : Remove performance measurement for wait time between
0000 45 : reschedules on secondary.
0000 46 :
0000 47 : V03-006 KDM0018 Kathleen D. Morse 19-Nov-1982
0000 48 : Implement wait for event flag system services on
0000 49 : secondary.
0000 50 :
0000 51 : V03-005 KDM0014 Kathleen D. Morse 27-Sep-1982
0000 52 : Fix performance data collection of scheduling usage
0000 53 : of exec mode AST.
```



```
0000 53 : V03-004 KDM0008 Kathleen D. Morse 31-Aug-1982
0000 54 : Prevent race between STOP/CPU and rescheduling interrupts.
0000 55 : A SVPCTX could be done twice for the current process on
0000 56 : the secondary.
0000 57 :
0000 58 : V03-003 KDM0007 Kathleen D. Morse 31-Aug-1982
0000 59 : Remove SCH$IDLE from MPSS$RESCHED logic flow. Dropping
0000 60 : IPL from 7 to 3 allowed AST queuing and if an IPL 5
0000 61 : reschedule interrupt occurred, then the same PCB address
0000 62 : ended up in MPSS$GL_CURPCB and SCH$GL_CURPCB, causing problems.
0000 63 :
0000 64 : 01 -
0000 65 : --
0000 66 :
0000 67 :
0000 68 : INCLUDE FILES:
0000 69 :
0000 70 :
0000 71 :
0000 72 : MACROS:
0000 73 :
0000 74 :
0000 75 :
0000 76 : EQUATED SYMBOLS:
0000 77 :
0000 78 :
0000 79 : $DYNDEF ; Structure type code definitions
0000 80 : $IPLDEF ; Interrupt priority level definitions
0000 81 : $LCKDEF ; Interlock bit definitions
0000 82 : $MPSDEF ; Secondary processor state definitions
0000 83 : $PCBDEF ; PCB definitions
0000 84 : $PHDDEF ; PHD definitions
0000 85 : $PRDEF ; Processor register definitions
0000 86 : $PSLDEF ; Program status longword definitions
0000 87 : $STATEDEF ; State definitions
0000 88 :
0000 89 :
0000 90 : OWN STORAGE:
0000 91 :
0000 92 :
0000 93 : ++
0000 94 :
0000 95 : IPL Usages:
0000 96 :
0000 97 : Primary Processor Secondary Processor
0000 98 : -----
0000 99 :
0000 100 : 0 - Unused 0 - Unused
0000 101 : 1 - Unused 1 - Unused
0000 102 : 2 - AST delivery 2 - AST delivery
0000 103 : 3 - Rescheduling 3 - Rescheduling
0000 104 : 4 - IO posting 4 - Unused
0000 105 : 5 - Multi-processor interrupt 5 - Xdelta
0000 106 : 6 - Fork 6 - Unused
0000 107 : 7 - Software timer interrupt 7 - Quantum end software interrupt
0000 108 : 8 - Fork 8 - Unused
0000 109 : 9 - Fork 9 - Unused
```

0000	110	:	10 - Fork	10 - Unused
0000	111	:	11 - Fork	11 - Unused
0000	112	:	12 - Unused	12 - Unused
0000	113	:	13 - Unused	13 - Unused
0000	114	:	14 - Unused	14 - Unused
0000	115	:	15 - Xdelta	15 - Unused
0000	116	:	16 - Unused	16 - Unused
0000	117	:	17 - Unused	17 - Unused
0000	118	:	18 - Unused	18 - Unused
0000	119	:	19 - Unused	19 - Unused
0000	120	:	20 - Device interrupt	20 - Unused
0000	121	:	21 - Device interrupt	21 - Unused
0000	122	:	22 - Device interrupt	22 - Unused
0000	123	:	23 - Device interrupt	23 - Unused
0000	124	:	24 - Timer interrupt	24 - Timer interrupt
0000	125	:	25 - Unused	25 - Unused
0000	126	:	26 - Unused	26 - Unused
0000	127	:	27 - Unused	27 - Unused
0000	128	:	28 - Unused	28 - Unused
0000	129	:	29 - Unused	29 - Unused
0000	130	:	30 - Unused	30 - Unused
0000	131	:	31 - Unused	31 - Unused
0000	132	:		
0000	133	--		
00000000	134		.PSECT AEXENONPAGED, LONG	; Nonpaged exec


```
0000 136 .SBTTL MPSS$RESCHEDIPL5 - MP RESCHEDULING (MA780) INTERRUPT HANDLER
0000 137 :++
0000 138 : MPSS$RESCHEDIPL5 - MULTI-PROCESSOR RESCHEDULING (MA780) INTERRUPT HANDLER
0000 139 :
0000 140 : This routine is entered via the IPL 5 rescheduling interrupt.
0000 141 : The vector for this interrupt is coded to cause execution
0000 142 : on the interrupt stack.
0000 143 :
0000 144 : ENVIRONMENT:
0000 145 :
0000 146 :     IPL=5 MODE=kernel IS=1
0000 147 :
0000 148 :     Executed by the primary processor.
0000 149 :
0000 150 : INPUT:
0000 151 :
0000 152 :     00(SP)=PC at reschedule interrupt
0000 153 :     04(SP)=PSL at interrupt
0000 154 :
0000 155 :--
0000 156 .ALIGN LONG ; Interrupt handlers must be
0000 157 ; longword aligned
0000 158 MPSS$RESCHEDIPL5:: ; Interrupt handler for
0000 159 ; multi-processor reschedule interrupt
3F BB 0000 160 PUSHR #^M<R0,R1,R2,R3,R4,R5> ; Save needed registers
0002 161
0002 162 .IF DF,MPPFMSWT
0002 163 BSBW MPSS$PFM_RESCHD ; Gather performance measurement data
0002 164 .ENDC
18 10 0002 165
0004 166 BSBB MPSS$SCHSCND ; Schedule secondary
0004 167
0004 168 .IF DF,MPPFMSWT
0004 169 BSBW MPSS$PFM_SCHDSUC ; Gather performance measurement data
0004 170 .ENDC
3F BA 0004 171
0006 172 POPR #^M<R0,R1,R2,R3,R4,R5> ; Restore registers
0007 173 REI ; Return from interrupt
0007 174
0007 175 .ENABL LSB
0007 176 PFAIL_EXIT: ; Exit when secondary is in powerfail
0007 177 MOVL #MPSS$K_INITSTATE,W^MPSS$GL_STATE ; Set state for restart
000C 178 BRB 1$ ; Done with rescheduling
000E 179 MPSS$EXIT:
000E 180 CMPL W^MPSS$GL_STATE,#MPSS$K_BUSYSTATE ; Check for powerfail window
0013 181 BNEQ 1$ ; where secondary was set busy after it
0015 182 TSTL W^MPSS$GL_PFAILTIM ; saw the powerfail occur. In this case,
0019 183 BNEQ 7$ ; go drop the process owned by secondary
001B 184 1$: RSB ; All done rescheduling secondary
001C 185 MPSS$SCHSCND:: ; Schedule secondary processor
001C 186 SETIPL #IPL$ SYNCH ; Synch scheduler with event reporting
001F 187 4$: BBSSI #LCK$Q INTERLOCK,W^MPSS$GL INTERLOCK,5$ ; Flush cache queue
0025 188 ASSUME MPSS$K_IDLESTATE LT MPSS$K_DROPSTATE
0025 189 ASSUME MPSS$K_BUSYSTATE GT MPSS$K_DROPSTATE
0025 190 ASSUME MPSS$K_INITSTATE GT MPSS$K_BUSYSTATE
0025 191 ASSUME MPSS$K_STOPSTATE GT MPSS$K_BUSYSTATE
02 0000'CF 00 E6 001F 187 4$:
0025 192 5$: CMPL W^MPSS$GL_STATE,#MPSS$K_DROPSTATE ; Is state already idle?
```



```
41 19 002A 193          BLSS      MPSS$SCHEDIPL5      ; Yes, just schedule a new one
E0 14 002C 194          BGTR      MPSS$EXIT           ; Br if busy
      002E 195          ;
      002E 196          ; The primary must check to see if the process was put into the DROP state
      002E 197          ; because it needed the primary to check for an event flag wait condition
      002E 198          ; or because it is ready for rescheduling.
      002E 199          ;
05 0000'CF 02 E7 002E 200 7$: BBCCI #MPSS$V_SECWAITCK,W^MPSS$GL_SECREQFLG,8$ ; Br if normal resched
      FFC9' 30 0034 201      BSBW      W^MPSS$WAITCK      ; Go perform event flag wait check for
      E6 11 0037 202      BRB        4$                  ; secondary processor and recheck for
      0039 203 8$:          ; DROP in case process was returned
      0039 204          ; to the secondary processor
      0039 205          ;
      0039 206          ; The primary processor takes the process that the secondary can
      0039 207          ; no longer execute and puts it back on the appropriate scheduler queue.
      0039 208          ;
      0039 209          ;
51 0000'CF D0 0039 209      MOVL      W^MPSS$GL_CURPCB,R1      ; Get address of current PCB
      52 0B A1 9A 003E 210      MOVZBL PCBS$B_PRI(R1),R2      ; Current priority
00 00000000'GF 52 E2 0042 211      BBSS      R2,G^SCH$GL_COMQS,10$ ; Mark queue non-empty
      2C A1 0C B0 004A 212 10$: MOVW      #SCH$C_COM,PCBS$W_STATE(R1) ; Set state to ready and computable
53 00000000'GF 42 7E 004E 213      MOVAQ     G^SCH$AQ_COM[R2],R3 ; Compute address of queue
      93 61 0E 0056 214      INSQHE     (R1),a(R3)+          ; Insert at tail of queue
      00 0000'CF 00 E6 0059 215      BBSSI     #LCK$V_INTERLOCK,W^MPSS$GL_INTERLOCK,15$ ; Flush cache queue
      0000'CF D5 005F 216 15$: TSTL      W^MPSS$GL_PFAILTIM    ; Is secondary in power fail?
      A2 12 0063 217      BNEQ      PFAIL_EXIT              ; Br if yes, power fail in progress
      0000'CF 01 D0 0065 218      MOVL      #MPSS$R_IDLESTATE,W^MPSS$GL_STATE ; Set secondary's state to idle
      006A 219      SOFTINT #IPL$_SCHED ; Reschedule primary processor
      006D 220      .DSABL LSB
      006D 221      ;
      006D 222      ;+
      006D 223      MPSS$SCHEDIPL5 - SCHEDULE NEW PROCESS FOR EXECUTION
      006D 224      ;
      006D 225      ; In this routine, the primary processor selects the highest priority
      006D 226      ; executable process and places it in the hands of the secondary
      006D 227      ; processor. It then tells the secondary processor to exit its idle
      006D 228      ; loop and start executing, by altering MPSS$GL_STATE.
      006D 229      ;
      006D 230      ENVIRONMENT:
      006D 231      ;
      006D 232      Executed by the primary processor.
      006D 233      ;
      006D 234      ;-
      006D 235      MPSS$SCHEDIPL5::
      006D 236      TSTL      W^MPSS$GL_PFAILTIM          ; Schedule for execution
      0071 237      BNEQ      PFAIL_EXIT                  ; Has secondary taken powerfail?
      0073 238          ; Br on yes, must set state to
      0073 239          ; INIT so that he can restart.
      007C 240          ;
      007E 241      NEXTQ: FFS      #0,#32,G^SCH$GL_COMQS,R2 ; Find first full state
      0086 242      BEQL      MPSS$IDLE                  ; No executable process?
      0089 243      QLOOP: MOVAQ     G^SCH$AQ_COM[R2],R3 ; Compute queue head address
      008C 244      MOV      R3,R4                        ; Get head of queue
      008F 245      CMPL      (R4),R4                    ; Get next in queue
      0091 246      BEQL      R3,R4                      ; Check for end of queue
      0095 247      BEQL      QEMPTY                     ; Br if yes
      009C 248      MOVL      PCB$B_PHD(R4),R5           ; Get process header address
      009E 249      EXTZV     #PSL$V_CURMOD,#PSL$S_CURMOD,PHD$B_PSL(R5),R0 ; Get current mode
      009E 249      BEQL      QLOOP                      ; Br if yes, can't run on secondary
      009E 249      CMPB      R0,PHD$B_ASTLVL(R5)         ; Is there an AST pending for process?
```



```

      0104  E4 18 00A3 250      BGEQ  QLOOP      ; Br on yes, don't sched for secondary
      DE  D5 00A5 251      TSTL  PHDSL_MPINHIBIT(R5) ; Any reason not to run on secondary?
      54  64 0F 00A9 252      BNEQ  QLOOP      ; Br on yes, don't sched for secondary
      08  12 00AB 253      REMQUE (R4),R4      ; Remove from queue
      00 00000000'GF 52  E5 00AE 254      BNEQ  20$      ; Queue not empty
      2C A4 0E B0 00B0 255      BBCC  R2,G^SCH$GL_COMQS,20$ ; Set queue empty
      0000'CF 54  D0 00B8 256 20$:      MOVW  #SCH$C_CUR,PCB$W_STATE(R4) ; Set state to current
      0000'CF 54  D0 00BC 257      MOVL  R4,W^MPSS$GL_CURPCB ; Note current PCB loc
      00C1 259
      00C1 260      .IF  DF,MPPFMSWT
      00C1 261      BSBW  W^MPSS$PFM_EXCHG      ; Collect perf meas data
      00C1 262
      0000'CF 03  D0 00C1 263      .ENDC
      00C1 264      MOVL  #MPSS$K_BUYSSTATE,W^MPSS$GL_STATE ; Set state to busy
      00C6 265
      00C6 266      .IF  DF,MPPFMSWT
      00C6 267      BSBW  MPSS$PFM_CTXSW      ; Gather performance measurement data
      00C6 268      .ENDC
      00C6 269
      05 00C6 270      RSB      ; Normal return
      00C7 271
      51 20 52  D6 00C7 272 QEMPTY: INCL  R2      ; Next queue
      52  C3 00C9 273      SUBL3  R2,#32,R1      ; Compute number of remaining queues
      0B 15 00CD 274      BLEQ  MPSS$IDLE      ; Br if none left
      52  EA 00CF 275      FFS  R2,R1,G^SCH$GL_COMQS,R2 ; Find next active state
      A4 12 00D8 276      BNEQ  NEXTQ      ; Continue if another non-empty state
      00DA 277 MPSS$IDLE:      ; No active, executable process
      0000'CF 00000000'GF 9E 00DA 278      MOVAB G^SCH$GL_NULLPCB,W^MPSS$GL_CURPCB ; Set to null PCB
      05 00E3 279      RSB      ; And return
      00E4 280
      00E4 281
      00E4 282 QEMPTY:      BUG_CHECK      QUEUEEMPTY,FATAL
      00E4 283
      00E8 284
      00E8 285 SCH$IDLE:
      00E8 286      SETIPL #IPL$ SCHED
      00000000'GF 20 90 00EB 287      MOVB  #32,G^SCH$GB_PRI
      0029 31 00F2 288      BRW  MPSS$SCHED      ; And try again
      00F5 289
```



```
00F5 291 .SBTTL MPSS$RESCHED - RESCHEDULING INTERRUPT HANDLER
00F5 292 :++
00F5 293 : MPSS$RESCHED - RESCHEDULING INTERRUPT HANDLER
00F5 294 :
00F5 295 : This routine is entered via the IPL 3 rescheduling interrupt.
00F5 296 : the vector for this interrupt is coded to cause execution on
00F5 297 : the kernel stack.
00F5 298 :
00F5 299 : This routine replaces SCH$RESCHED in a vanilla VMS system.
00F5 300 :
00F5 301 : ENVIRONMENT:
00F5 302 :
00F5 303 : IPL=3 MODE=kernel IS=0
00F5 304 :
00F5 305 : Executed by the primary processor.
00F5 306 :
00F5 307 : INPUT:
00F5 308 :
00F5 309 : 00(SP)=PC at reschedule interrupt
00F5 310 : 04(SP)=PSL at interrupt
00F5 311 :
00F5 312 :--
00F5 313 :.ALIGN LONG ; Align interrupt handler
00F8 314 MPSS$RESCHED:: ; Reschedule interrupt handler
00F8 315 SETIPL #IPL$_SYNCH ; Synch scheduler with event reporting
00FB 316 SVPCTX ; Save context of process
00FC 317 MOVL G^SCH$GL_CURPCB,R1 ; Get address of current PCB
0103 318 MOVZBL PCBSB_PRT(R1),R2 ; Current priority
0107 319 BBSS R2,G^SCH$GL_COMQS,10$ ; Mark queue non-empty
010F 320 10$: MOVW #SCH$C_COM,PCBSW_STATE(R1) ; Set state to resident computable
0113 321 MOVAQ G^SCH$AQ_COMH[R2],R3 ; Compute address of queue
011B 322 INSQUE (R1),a(R3)+ ; Insert at tail of queue
011E 323 :
011E 324 :+
011E 325 : MPSS$SCHED - SCHEDULE NEW PROCESS FOR EXECUTION
011E 326 :
011E 327 : This routine selects the highest priority executable process
011E 328 : and places it in execution.
011E 329 :
011E 330 : This routine replaces SCH$SCHED in a vanilla VMS system.
011E 331 :
011E 332 : ENVIRONMENT:
011E 333 :
011E 334 : Executed by the primary processor.
011E 335 :
011E 336 :--
011E 337 :
011E 338 MPSS$SCHED:: ; Schedule for execution
011E 339 SETIPL #IPL$_SYNCH ; Synch scheduler with event reporting
0121 340 BSBW MPSS$SCHSCND ; Schedule secondary
0124 341 FFS #0,#32,G^SCH$GL_COMQS,R2 ; Find first full state
012D 342 BEQL SCH$IDLE ; No executable process?
012F 343 MOVAQ G^SCH$AQ_COMH[R2],R3 ; Compute queue head address
0137 344 REMQUE a(R3)+,R4 ; Get head of queue
013A 345 BVS QEMPTY ; Br if queue was empty (bugcheck)
013C 346 BNEG 20$ ; Queue not empty
013E 347 BBCC R2,G^SCH$GL_COMQS,20$ ; Set queue empty
```

51 00000000'GF 07 00FB 316 SVPCTX
52 0B A1 9A 0103 318 MOVZBL
00 000000'GF 52 E2 0107 319 BBSS
2C A1 0C B0 010F 320 10\$: MOVW
53 00000000'GF 42 7E 0113 321 MOVAQ
93 61 0E 011B 322 INSQUE
011E 323
011E 324 :+
011E 325 : MPSS\$SCHED - SCHEDULE NEW PROCESS FOR EXECUTION
011E 326 :
011E 327 : This routine selects the highest priority executable process
011E 328 : and places it in execution.
011E 329 :
011E 330 : This routine replaces SCH\$SCHED in a vanilla VMS system.
011E 331 :
011E 332 : ENVIRONMENT:
011E 333 :
011E 334 : Executed by the primary processor.
011E 335 :
011E 336 :--
011E 337 :
011E 338 MPSS\$SCHED::
011E 339 SETIPL
52 00000000'GF 20 FEF8 30 0121 340 BSBW
B9 13 0124 341 FFS
53 00000000'GF 42 7E 012D 342 BEQL
54 93 0F 0137 343 MOVAQ
A8 1D 013A 344 REMQUE
08 12 013C 345 BVS
00 00000000'GF 52 E5 013E 346 BNEG
BBCC 347


```
0A A4 0C 91 0146 348 20$:
00000000'GF 10 18 A4 54 D0 014C 349
54 00000000'GF 00 0000'CF 00 E6 014A 350
01 0000'CF 3A 12 014C 351
00000000'GF 7FFFFFFF 8F D3 0150 352
10 A4 8E D5 0157 353
54 6C A4 D0 015B 354
0104 C4 D5 015C 355
OC 12 015E 356
0199 357
0199 358 40$:
0199 359
0199 360
0199 361
0199 362
0199 363
0199 364
0199 365
0199 366
0199 367
0199 368
0199 369
0199 370
0199 371
0199 372
0199 373
0199 374
0199 375
019C 376
019C 377
01A3 378
01A5 379
01AC 380 45$:
01B1 381 50$:
01B3 382
01B8 383
01BB 384
01C3 385 60$:
01C6 386
01C7 387
01C7 388

CMPB #DYN$C_PCB,PCBSB_TYPE(R4) ; Must be a process control block
BNEQ QEMPTY ; Otherwise fatal error
MOVW #SCH$C_CUR,PCBSW_STATE(R4) ; Set state to current
MOVL R4,G^SCH$GGL_CURPCB ; Note current PCB location
MTPR PCBSL_PHYPCB(R4),#PR$PCBB ; Set PCB base physical address
LDPCTX ; Restore context
PUSHL R4 ; Save register
MOVL G^SCH$GGL_CURPCB,R4 ; Get address of current PCB
BBSSI #LCK$V_INTERLOCK,W^MPSS$GGL_INTERLOCK,40$ ; Flush cache queue
CMPL W^MPSS$GGL_STATE,#MPSS$K_IDLESTATE ; Is secondary idle?
BNEQ 50$ ; Br on no, already running a process
MFPR #PR$ASTLVL,-(SP) ; Get PR ASTLVL
TSTL (SP)+ ; Is ASTLVL = KERNEL?
BEQL 50$ ; Br on yes, dont cause reschedule
BITL #X7FFFFFFF,G^SCH$GGL_COMQS ; Is any process other than null COM?
BEQL 50$ ; Br on no, nothing else can run
PUSHAB PCBSL_ASTQFL(R4) ; Are there any AST's outstanding
CMPL (SP)+,PCBSL_ASTQFL(R4) ; for this process? If so, don't
BNEQ 50$ ; cause a reschedule AST.
MOVL PCBSL_PHD(R4),R4 ; Get address of PHD
TSTL PHD$M_PINHIBIT(R4) ; Any reason not to run on secondary?
BNEQ 45$ ; Br on yes, don't cause reschedule AST

IF DF,MPPFMSWT
BSBW W^MPSS$PFM_ASTSC
ENDC

MTPR #PSL$C_EXEC,#PR$ASTLVL ; Cause reschedule of process when it
; exits from kernel mode on primary
MOVL G^SCH$GGL_CURPCB,R4 ; Get address of current PCB
BRB 60$ ; Don't decrement priority
MOVL G^SCH$GGL_CURPCB,R4 ; Get address of current PCB
CMPB PCBSB_PRI(R4),PCBSB_PRI(R4) ; Check for base
BEQL 60$ ; Br if Priority=current
BBC #4,PCBSB_PRI(R4),60$ ; Don't float real time priority
INCB PCBSB_PRI(R4) ; Move toward base priority
MOVB PCBSB_PRI(R4),G^SCH$GB_PRI ; Set global priority
POPL R4 ; Restore register
REI ; Normal return
```



```
01C7 390 .SBTTL MPSSMPSCHED - SECONDARY PROCESSOR SCHEDULING
01C7 391 :++
01C7 392 : FUNCTIONAL DESCRIPTION:
01C7 393 :
01C7 394 : MPSSMPSCHED is entered to drop the current process running
01C7 395 : in the secondary processor and select a new one.
01C7 396 :
01C7 397 : ENVIRONMENT:
01C7 398 :
01C7 399 : Executed by the secondary processor.
01C7 400 :
01C7 401 :--
01C7 402 .ALIGN LONG
01C8 403 MPSSMPSCHED::
01C8 404 SETIPL #IPL$ SYNCH ; Synchronize with other events
01CB 405 BITL #PSL$M_CURMOD,4(SP) ; Check for kernel
01D3 406 BEQL GETOUT ; If so, exit already in progress
01D5 407
01D5 408 MPSSMPSCHED2::
01D5 409
01D5 410 .IF DF,MPPFMSWT
01D5 411 BSBW MP$PFM_SVPCTX ; Gather performance measurement data
01D5 412 .ENDC
01D5 413
01D5 414 SETIPL #IPL$ POWER ; Prevent duplicate SVPCTX from a MA780
01D8 415 ; interrupt for STOP/CPU
01D8 416 SVPCTX ; Save current process context
01D9 417 MOVL #MP$K_DROPSTATE,W*MP$SGL_STATE ; Indicate process being dropped
01DE 418 SETIPL #IPL$ SYNCH ; Now allow MA780 interrupts again
01E1 419
01E1 420 MPSSMPSCHED1::
01E1 421 ; Init entry point for startup
01E1 422
01E1 423 .IF DF,MPPFMSWT
01E1 424 BSBW MP$PFM_INT ; Gather performance measurement data
01E1 425 .ENDC
01E1 426 BSBW MP$INTPRIM ; Interrupt primary processor
01E4 427
01E4 428 : Now invalidate the system half of the translation buffer because
01E4 429 : the secondary is going into a busy-waiting loop until the primary
01E4 430 : gives it something to do. The primary may request that the secondary
01E4 431 : invalidate certain system addresses from its translation buffer while
01E4 432 : it is waiting. Since the secondary won't have these addresses in its
01E4 433 : translation buffer, it just keeps telling the primary that it has done
01E4 434 : the request, i.e., by clearing MP$SGL_INVALID. (If this clear was not
01E4 435 : in the loop, the primary would wait indefinitely for the secondary, who
01E4 436 : in turn was waiting indefinitely for the primary.)
01E4 437 :
01E4 438 INVALID ; Invalidate translation buffer
01E7 439 10$: CLRL W*MP$SGL_INVALID ; Acknowledge TB invalidate request
01EB 440 BBSSI #LCK$V_INTERLOCK,W*MP$SGL_INTERLOCK,20$ ; Flush cache queue
01F1 441 20$: CMPL #MP$K_BUSTYSTATE,W*MP$SGL_STATE ; Is state set to busy?
01F6 442 BNEQ 10$ ; No, loop
01F8 443
01F8 444 : The primary has scheduled a process for the secondary to execute.
01F8 445 : The secondary need only load up the new PCB and being executing.
01F8 446 :
```

04 AE 03000000 8F D3 01CB 405
3E 13 01D3 406

0000'CF 02 D0 01D9 417
01DE 418

FE1C' 30 01E1 426

00 0000'CF 00 D4 01E7 439
0000'CF 03 E6 01EB 440
EF 12 01F1 441
01F6 442
01F8 443
01F8 444
01F8 445
01F8 446


```
00 0000'CF 00 E6 01F8 447 SETIPL #IPL$ POWER ; Disable powerfail interrupts
54 0000'CF D0 01FB 448 BBSSI #LCK$? INTERLOCK,W^MPSS$GL_INTERLOCK,30$ ; Flush cache queue
52 18 A4 D0 0201 449 30$: MOVL W^MPSS$GL_CURPCB,R4 ; Get current PCB address
10 52 DA 0206 450 MOVL PCB$L_PHYPCB(R4),R2 ; Get physical PCB address
020A 451 MTPR R2,#PR$_PCBB ; Set PCB base address
020D 452
020D 453 .IF DF,MPPFMSWT
020D 454 BSBW MPSS$PFM_LDPCTX ; Gather performance measurement data
020D 455 .ENDC
020D 456
0000'CF 04 06 020D 457 LDPCTX ; Load process context
D0 020E 458 MOVL #MPSS$K_EXECSTATE,W^MPSS$GL_STATE ; Indicate LDPCTX is done
02 0213 459 GETOUT: REI ; And execute
0214 460 .END
```

MPSCHED
Symbol table

- MULTIPROCESSOR SCHEDULER

J 16

16-SEP-1984 02:07:03 VAX/VMS Macro V04-00
5-SEP-1984 02:07:20 [MP.SRC]MPSCHED.MAR;1

Page 11
(1)

```

BUG$_QUEUEEMPTY      ***** X 02
DYN$C_PCB             = 0000000C
GETOUT               = 00000213 R 02
IPL$POWER            = 0000001F
IPL$SCHED             = 00000003
IPL$SYNCH            = 00000008
LCK$V_INTERLOCK      = 00000000
MP$EXIT              = 0000000E R 02
MP$SGL_CURPCB        ***** X 02
MP$SGL_INTERLOCK      ***** X 02
MP$SGL_INVALID        ***** X 02
MP$SGL_PFAILTIM       ***** X 02
MP$SGL_SECREQFLG      ***** X 02
MP$SGL_STATE          ***** X 02
MP$IDCE              = 000000DA R 02
MP$INTPRIM           ***** X 02
MP$K_BUSYSTATE        = 00000003
MP$K_DROPSTATE        = 00000002
MP$K_EXECSTATE        = 00000004
MP$K_IDLESTATE        = 00000001
MP$K_INITSTATE        = 00000005
MP$K_STOPSTATE        = 00000006
MP$MPSCHED            = 000001C8 RG 02
MP$MPSCHED1           = 000001E1 RG 02
MP$MPSCHED2           = 000001D5 RG 02
MP$RESCHED            = 000000F8 RG 02
MP$RESCHEDIPL5        = 00000000 RG 02
MP$SCHED              = 0000011E RG 02
MP$SCHEDIPL5          = 0000006D RG 02
MP$SCHSCND            = 0000001C RG 02
MP$V_SECWAITCK        = 00000002
MP$WAITCK             ***** X 02
NEXTQ                 = 0000007E R 02
PCB$B_PRI             = 0000000B
PCB$B_PRI8            = 0000002F
PCB$B_TYPE            = 0000000A
PCB$L_ASTQFL          = 00000010
PCB$L_PHD             = 0000006C
PCB$L_PHYPCB          = 00000018
PCB$W_STATE           = 0000002C
PFAIL_EXIT            = 00000007 R 02
PHD$B_ASTLVL          = 000000CF
PHD$L_MPINHIBIT       = 00000104
PHD$L_PSL             = 000000C4
PR$ASTLVL             = 00000013
PR$IPL                = 00000012
PR$PCBB               = 00000010
PR$SIRR               = 00000014
PR$TBIA               = 00000039
PSL$C_EXEC            = 00000001
PSL$M_CURMOD          = 03000000
PSL$S_CURMOD          = 00000002
PSL$V_CURMOD          = 00000018
QEMPTY                = 000000C7 R 02
QEMPTY               = 000000E4 R 02
QLOOP                 = 00000089 R 02
SCH$AQ_COMH           ***** X 02

```

```

SCH$AQ_COMT           ***** X 02
SCH$C_COM             = 0000000C
SCH$C_CUR             = 0000000E
SCH$GB_PRI            ***** X 02
SCH$GL_COMQS          ***** X 02
SCH$GL_CURPCB         ***** X 02
SCH$GL_NULLPCB        ***** X 02
SCH$IDCE              = 000000E8 R 02

```


+-----+
! Psect synopsis !
+-----+

PSECT name	Allocation	PSECT No.	Attributes
ABS	00000000 (0.)	00 (0.)	NOPIC USR CON ABS LCL NOSHR NOEXE NORD NOWRT NOVEC BYTE
\$ABSS	00000000 (0.)	01 (1.)	NOPIC USR CON ABS LCL NOSHR EXE RD WRT NOVEC BYTE
AEXENONPAGED	00000214 (532.)	02 (2.)	NOPIC USR CON REL LCL NOSHR EXE RD WRT NOVEC LONG

+-----+
! Performance indicators !
+-----+

Phase	Page faults	CPU Time	Elapsed Time
Initialization	30	00:00:00.06	00:00:01.02
Command processing	108	00:00:00.86	00:00:03.85
Pass 1	252	00:00:06.96	00:00:20.38
Symbol table sort	0	00:00:00.84	00:00:01.80
Pass 2	97	00:00:01.76	00:00:06.14
Symbol table output	9	00:00:00.09	00:00:00.26
Psect synopsis output	2	00:00:00.02	00:00:00.02
Cross-reference output	0	00:00:00.00	00:00:00.00
Assembler run totals	500	00:00:10.59	00:00:33.47

The working set limit was 1350 pages.
36516 bytes (72 pages) of virtual memory were used to buffer the intermediate code.
There were 40 pages of symbol table space allocated to hold 597 non-local and 17 local symbols.
465 source lines were read in Pass 1, producing 15 object records in Pass 2.
22 pages of virtual memory were used to define 21 macros.

+-----+
! Macro library statistics !
+-----+

Macro library name	Macros defined
_\$255\$DUA28:[MP.OBJ]MP.MLB;1	4
_\$255\$DUA28:[SYS.OBJ]LIB.MLB;1	9
_\$255\$DUA28:[SYSLIB]STARLET.MLB;2	6
TOTALS (all libraries)	19

745 GETS were required to define 19 macros.

There were no errors, warnings or information messages.

MACRO/LIS=LIS\$:MPSCHED/OBJ=OBJ\$:MPSCHED MSRC\$:MPPREFIX/UPDATE=(ENH\$:MPPREFIX)+MSRC\$:MPSCHED/UPDATE=(ENH\$:MPSCHED)+EXECMLS/LIB+LIB\$:M

0248 AH-BT13A-SE
VAX/VMS V4.0

DIGITAL EQUIPMENT CORPORATION
CONFIDENTIAL AND PROPRIETARY

