

Smart sensing daemon for Miosix

Design of an energy efficient sensor daemon for Miosix

Rizzi Alessandro Maria

Matr. 783504, (alessandromaria.rizzi@mail.polimi.it)

Report for the master course of Real Time Operating Systems

Reviser: PhD. Patrick Bellasi (bellasi@elet.polimi.it)

Received: March, 05 2013

Abstract

The goal of the project is to design and implement a daemon for Miosix embedded OS performing the "smart-sensing", which is a new power-efficient way of performing a series of reads from a sensor spaced out by a fixed amount of time. The difference between this method and a normal data acquisition from the sensor lies in the fact that the board remains in suspension most of the time, waking up in a low-consumption mode only to read sensor data.

1 Introduction

In context of wireless sensor network great attention is given to energy-efficiency issues.

The reason for this is the fact that every node of the network must operate continuously for a long time period of time without any physical maintenance, which would be impractical due to the large number of nodes and their barely reachable location.

Besides, the main purpose of a WSN is the monitoring of some quantities acquired by some sensors.

From these two general properties of WSN it's clear that data acquisition from sensors is a key activity for such a network and should be optimized as much as possible on the energy-efficiency aspect.

We can imagine that on each node of the WSN will be executed different task. The major part of the task would be devoted to data acquisition and elaboration which means that will be busy for a certain amount of time performing some reads from a sensor and the spending some time in elaborating the data.

It would be reasonable suppose that the node would spend most of the time in waiting the data from the sensors. In the most simple implementation the board would remain always active, maybe just reducing the microcontroller consumption.

Otherwise it's possible to improve the energy consumption of the node first of all by making the board be suspended if there are no active task. The suspension mode is a special mode which is present in most of microcontroller platforms which has very little energy consumption: both the microcontroller and the memory are not powered.

Then it's possible to improve further the energy-efficiency by performing the requested reads in a special low-consumption mode.

The project aims to the development of the complete in-

frastructure for this activity, starting from an operating system which is able to run multiple task and to suspend/resume them.

The rest of the section is divided as follows: in subsection 1.1 is presented the Miosix Embedded OS; in subsection 1.2 the board used.

1.1 Miosix Embedded OS

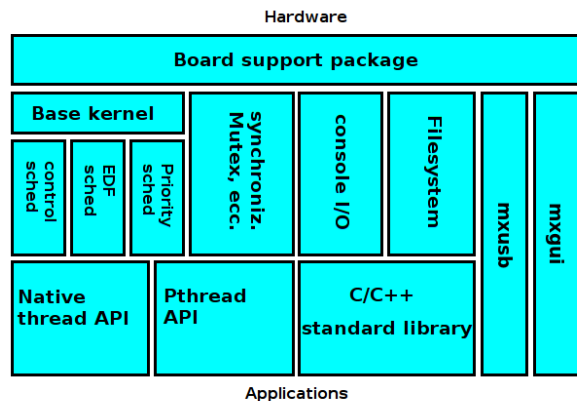


Figure 1: Miosix kernel architecture.

Miosix is a kernel for microcontrollers. Its goals are to provide an environment as much "standard compliant" as possible (in which developing application for an embedded system isn't much different than developing a standard desktop application) and avoiding performance or code-footprint penalties for the unused features. It provides:

- Multithreading with a pthread-like API
- C and C++ support including standard C and STL libraries.

- Device drivers for most complex devices
- Filesystem support with POSIX-like API
- Different scheduler to be chosen

It supports only 32bit microcontrollers and has an experimental support for multitasking which lacks some features (like dynamic-loading of executables from the filesystem). Its basic architecture is presented in **Figure 1**.

For this project has been particularly important some recent additions to it, which are multitasking with the possibility of suspending the system when there are no active tasks.

1.2 The stm3220g-eval board




Figure 2: Picture of the STM stm3220g-eval board.

The board used in the project is the ST Microelectronics stm3220g-eval (**Figure 2**).

It has a 120MHz Cortex-M3 microcontroller (STM32F207IG), which has 1MiB of FLASH, 128KiB of RAM and 2MiB of external SRAM.

The board is pretty rich in terms of hardware features, including an USB port, an RS-232 port, an audio DAC, a 320x240 3.2" LCD display, a MicroSD card slot, 4 LEDs, a 4-direction joystick, an RTC with backup battery and a potentiometer (which has been used as source of acquisition data).

In addition were added 128KiB of MRAM in order to save the processes during suspension.

In the development of the project the board has been flashed through the USB port, which has been used also as a serial console. Due to the lack of a JTAG in-circuit debugger, this feature has not been used in the development phase.

2 ADC driver

The first step in the development of the project has been the implementation of a module for Miosix capable of perform reads from a sensor.

This operation has been implemented as performing analog reads from a specific GPIO pin.

In particular the module has to be able to perform the following operations:

- Initialize a specific GPIO pin on a specific GPIO port to perform analog reads
- Initialize a specific ADC to perform an analog read
- Perform a single read on a given ADC channel

For the specific application context (WSN) characterized by a very low frequency of reads, high performances are not required.

So it's no need of using DMA. In fact also the initialization phase is repeated at every read.

One problem arised regarding the way of determinating the following informations of a given input: GPIO pin, GPIO port, ADC to be used, ADC channel connected to the GPIO port. This informations are not easily derivable from the other ones (it's necessary to look at the reference manual to obtain the correct mapping informations). Besides, while the ADC mechanism is quite general in all the ST Microelectronics' microcontrollers family (it has been succesfully tested in STM32F4 Discovery board), the mapping between GPIO pin, ADC and ADC channel are subject to change on different boards.

So the solution adopted has been to pack all the required informations into a 32bit integer as following:

- bits 0-4: ADC channel (0 the first, 1 the second, etc...)
- bits 5-7: ADC (0 the first, 1 the second, etc...)
- bits 8-15: GPIO pin
- bits 16-19: GPIO port (A=0, B=1, C=2, etc...)
- bits 20-31: Not used

In this way an input is identified by a single 32 bit integer. The input used in the project test has been the potentiometer: this because it has many possible values and it's easy to modify it and check the correctness.

The relevant part of implementation is shown in **Figure 3**. In particular the init method initialize both the GPIO and ADC as shown. There are two version of read the

first, which accepts no parameters, is meant for performing a read from an already initialized object. The second one, which requires the deviceId, is a static method which perform all the require operations (initialization, read) in a single step and is the only used outside the class.

```

16 |   AdcDriver::AdcDriver(uint32_t deviceId) {
17 |       decodeDeviceId(deviceId);
18 |   }
19 |
20 |   void AdcDriver::init() {
21 |       initGPIO();
22 |       initADC();
23 |   }
24 |
25 |   unsigned short AdcDriver::read() {
26 |       ADC_TypeDef *myADC = AdcInumADC();
27 |       myADC->SQR3 = ADCchannel;
28 |       myADC->CR2 |= ADC_CR2_SWSTART;
29 |       while ((myADC->SR & ADC_SR_EOC) == 0);
30 |       unsigned short data = myADC->DR;
31 |       return data;
32 |   }
33 |
34 |   unsigned short AdcDriver::read(uint32_t deviceId) {
35 |       AdcDriver adc(deviceId);
36 |       adc.init();
37 |       return adc.read();
38 |   }
39 |
40 |   void AdcDriver::decodeDeviceId(uint32_t deviceId) {
41 |       ADCchannel = (unsigned char) (deviceId & ADC_CHANNEL_MASK);
42 |       numADC = (unsigned char) ((deviceId & ADC_MASK) >> 5);
43 |       gpioPin = (unsigned char) ((deviceId & GPIO_PIN_MASK) >> 8);
44 |       gpioPort = (unsigned char) ((deviceId & GPIO_PORT_MASK) >> 16);
45 |   }
46 |
47 |   void AdcDriver::initGPIO() {
48 |       GpioPin gpio(gpioMapping[gpioPort], gpioPin);
49 |       gpio.mode(Mode::INPUT_ANALOG);
50 |   }
51 |
52 |   void AdcDriver::initADC() {
53 |       InterruptDisableLock_dLock; //Using the slow one so I don't care if kernel is started or not
54 |       RCC->APB3ENR |= ADC_REG(numADC);
55 |       ADC_TypeDef *myADC = AdcInumADC();
56 |       ADC->CCR |= ADC_CCR_ADCPRE_1; //ADC prescaler 80MHz/6=14MHz
57 |       myADC->CR1 = 0;
58 |       myADC->CR2 = ADC_CR2_ADON; //The first assignment sets the bit
59 |       myADC->SQR1 = 0; //Do only one conversion
60 |       myADC->SQR2 = 0;
61 |       myADC->SQR3 = 0;
62 |       myADC->SMPR1 = 7 << 18; //480 clock cycles of sample time for temp sensor
63 |   }

```

Figure 3: Extract from ADC driver implementation

3 SmartSensing module

This part is the core of the project. It has the task of keep track of the data requests and schedule them. It's main idea is quite simple: it keeps track of the request of an acquisition job and schedule the correct time to perform the required reads either by scheduling a new wake up or by the kernel daemon if they occur when the system is already running. The module can be divided in the following conceptual part:

- Data structures
- Request registration and query
- Kernel daemon

In the rest part of this section we analyze these different conceptual parts plus the integration with the rest of the kernel from the module point of view.

3.1 Data structures

In this subsection are described the data structures utilized and the methods which operate upon them. The basic idea is to subdivide the data requested in two sets: one with the data which are in common with all the acquisition jobs, and one with the group of data of each job. Since the data must survive the suspension process during which all the RAM content is lost, it is allocated at

the end of the backup SRAM, which is maintained during the whole process, as shown in **Figure 4**. The reason for this allocation is to minimize the compatibility issues with Miosix.




Figure 4: Structure of the end of backup SRAM

The data is disposed as follows: first the general data, which is formed by a canary, used to check if the other data structures haven't been overwritten by other data stored in the backup SRAM, and an integer which store the next system full boot time in milliseconds. Then there is a fixed-size array of the data of each job. This choice has been done due to the extremely low number of allowed jobs which are a consequence of the low number of tasks allowed (just 4). More precisely the job acquisition data is composed of a fixed-size array which holds the data acquired, the number of total samples requested at the job, the number of remaining acquisition, the time in milliseconds when the next acquisition has to be taken, the space in milliseconds between two acquisitions, the device identifier and the id of the process which has requested the job. Each array element is referred from now on in this report and in the code as "queue" because it behaves like a queue which maintains the data coming from the sensor.

```

25 |   struct SmartSensingStatus {
26 |   #ifndef CHECK_CANARY
27 |       /**
28 |        * @brief canary
29 |        * Canary to check if the backup ram area of the class has been corrupted
30 |        */
31 |       unsigned long int canary;
32 |   #endif
33 |       /**
34 |        * @brief nextSystemRestart
35 |        * Time in seconds of the next system restart
36 |        */
37 |       unsigned long long nextSystemRestart;
38 |   };
39 |
40 |
41 |   template <typename T, unsigned int N>
42 |   /**
43 |    * @brief The SSQueue struct
44 |    * Structure which holds a queue
45 |    */
46 |   struct SSQueue {
47 |       T data[N];
48 |       unsigned int size;
49 |       unsigned int remaining;
50 |   /**
51 |    * @brief nextTime next time in milliseconds
52 |    */
53 |       unsigned long long int nextTime;
54 |   /**
55 |    * @brief period period in milliseconds
56 |    */
57 |       unsigned int period;
58 |       uint32_t deviceId;
59 |       pid_t processId;
60 |   };

```

Figure 5: Data structures

Due to the extremely low size of the number of jobs, the

algorithm implemented for searching values into the array is just an exhaustive scan.

The methods for access and manipulate the queues can be subdivided in the following categories:

- Methods for retrieve a single queue which match a given parameter/condition

In this category there are `getFirstFreeQueue` (**Figure 6**) which retrieve the first unused queue and `getQueueFromProcessId` (**Figure 7**) which retrieve the index of the first queue which belongs to a process with the given id.

```

302 |     /**
303 |     * @brief getFirstFreeQueue
304 |     * Retrieve the first queue available
305 |     * It's supposed to be invoked only if the kernel is active
306 |     * @return index of the first queue, -1 if none available
307 |     */
308 |     int getFirstFreeQueue() const{
309 |         for (unsigned int i = 0; i < Q; i++) {
310 |             if (queue[i].size == 0) {
311 |                 return i;
312 |             }
313 |         }
314 |         return -1;
315 |     }

```

Figure 6: `getFirstFreeQueue` method

```

409 |     /**
410 |     * @brief getQueueFromProcessId
411 |     * Retrieve the queue associated to a given process
412 |     * Can be invoked either during boot phase or when the kernel is active
413 |     * @param processId id of the process which owns the queue
414 |     * @return index of the requested queue, -1 if no queue correspond the given process
415 |     */
416 |     int getQueueFromProcessId(pid_t processId) const{
417 |         for(unsigned int i=0;i<Q;i++){
418 |             if ((queue[i].size>0) && (queue[i].processId==processId)) {
419 |                 return (int)i;
420 |             }
421 |         }
422 |         return -1;
423 |     }

```

Figure 7: `getQueueFromProcessId` method

- Methods performing operations on a single queue

Here there are `initQueue` (**Figure 8**), `readQueue` (**Figure 9**) and `resetQueue` (**Figure 10**), which respectively initialize a queue with the given data, add to a queue a new read from its associated input; and set a queue not utilized.

```

388 |     /**
389 |     * @brief initQueue
390 |     * Initialize a queue
391 |     * It's supposed to be invoked only if the kernel is active
392 |     * @param i index of a free queue
393 |     * @param processId id of the process which owns the corresponding task
394 |     * @param threadId id of the thread which owns the corresponding task
395 |     * @param deviceId id which carries the information on the input source of the reads
396 |     * @param size number of reads to be performed
397 |     * @param period time in milliseconds between two reads
398 |     */
399 |     void initQueue(unsigned int i, pid_t processId, Thread* threadId, uint32_t deviceId, unsigned int size, unsigned int period) {
400 |         queue[i].deviceId = deviceId;
401 |         queue[i].size = size;
402 |         queue[i].remaining = size;
403 |         queue[i].nextTime = getTick() + period;
404 |         queue[i].period = period;
405 |         queue[i].processId=processId;
406 |         this->threadId[i]=threadId;
407 |     }

```

Figure 8: `initQueue` method

```

317 |     /**
318 |     * @brief readQueue
319 |     * Performs a read on the selected queue
320 |     * Can be invoked either during boot phase or when the kernel is active
321 |     * No checks are performed either on the index or if the queue is full
322 |     * @param i index of the selected queue
323 |     */
324 |     void readQueue(int i) {
325 |         unsigned short value = AdcDriver::read(queue[i].deviceId);
326 |         queue[i].data[queue[i].size - queue[i].remaining] = value;
327 |         char debug[200];
328 |         sprintf(debug, "T: %5li PID: %4i SS: %4x", getTick(), queue[i].processId, value);
329 |         if (isKernelRunning()){
330 |             puts(debug);
331 |         }
332 |         else{
333 |             IRQbootlog(debug);
334 |             IRQbootlog("\r\n");
335 |         }
336 |         queue[i].remaining--;
337 |     }

```

Figure 9: `readQueue` method

```

372 |     /**
373 |     * @brief resetQueue
374 |     * Set a given queue as available
375 |     * Can be invoked either during boot phase or when the kernel is active
376 |     * @param i index of the queue
377 |     * @return false if the given index is invalid, true otherwise
378 |     */
379 |     bool resetQueue(unsigned int i){
380 |         if((i<0)||i>=Q){
381 |             return false;
382 |         }
383 |         queue[i].size = 0;
384 |         queue[i].remaining = 0;
385 |         return true;
386 |     }

```

Figure 10: `resetQueue` method

- Methods for extract timing information from the whole set of queues

We consider the methods `getNextEvent` and `getNextSecond` (**Figure 11**) which given the current time and optionally an event time return the first occurrence in the future which is either a queue read event or the event given, respectively in millisecond and second. In particular the latter has been designed to cope with the different granularity of time, which is potentially problematic.

```

267 |     /**
268 |     * @brief getNextSecond
269 |     * Retrieve next time in seconds in which the next event will happen
270 |     * Can be invoked either during boot phase or when the kernel is active
271 |     * @param currentTime time in millisecond to be considered as current time
272 |     * @param minTime if different from 0 it counts as an additional event time in milliseconds
273 |     * @return next event time in seconds
274 |     */
275 |     unsigned int getNextSecond(unsigned long long currentTime, unsigned long long minTime) const{
276 |         unsigned int nextEvent=getNextEvent(currentTime,minTime);
277 |         if((currentTime%1000)>(nextEvent%1000)){
278 |             return nextEvent/1000;
279 |         }
280 |         return (nextEvent+999)/1000;
281 |     }
282 |
283 |     /**
284 |     * @brief getNextEvent
285 |     * Retrieve next time in milliseconds in which the next event will happen
286 |     * Can be invoked either during boot phase or when the kernel is active
287 |     * @param currentTime time in millisecond to be considered as current time
288 |     * @param minTime if different from 0 it counts as an additional event time in milliseconds
289 |     * @return next event time in milliseconds
290 |     */
291 |     unsigned long long getNextEvent(unsigned long long currentTime, unsigned long long minTime) const{
292 |         for (unsigned int i = 0; i < Q; i++) {
293 |             if (queue[i].size > 0 && queue[i].remaining > 0 && (queue[i].nextTime<currentTime) ) {
294 |                 if ((minTime == 0) || (minTime >= queue[i].nextTime)) {
295 |                     minTime = queue[i].nextTime;
296 |                 }
297 |             }
298 |         }
299 |         return minTime;
300 |     }

```

Figure 11: `getNextEvent` and `getNextSecond` methods

- Methods performing operations on the entire set of queues

Here there are methods which scans all the queue checking if they match a condition and, if so, perform an action on the selected queue.

The method `updateQueue` (**Figure 12**), given the current time in milliseconds, apply the `readQueue` method to all those queues which require a new read.

The method `wakeCompletedProcess` (**Figure 13**) wakes up all the processes whose queue requires no further reads.

```

244 | /**
245 |  * @brief updateQueue
246 |  * Update all the initialized queues
247 |  * Can be invoked either during boot phase or when the kernel is active
248 |  * @param time time in milliseconds. All the reads which are before it
249 |  * will be performed.
250 |  */
251 | void updateQueue(unsigned long long time) {
252 |     completedTask=false;
253 |     for (unsigned int i = 0; i < Q; i++) {
254 |         if ((queue[i].remaining > 0) && (queue[i].nextTime <= time)) {
255 |             readQueue(i);
256 |             queue[i].nextTime += queue[i].period;
257 |             if (queue[i].nextTime <= time){
258 |                 queue[i].nextTime = time + queue[i].period;
259 |             }
260 |             if (queue[i].remaining==0){
261 |                 completedTask=true;
262 |             }
263 |         }
264 |     }

```

Figure 12: updateQueue method

```

354 | /**
355 |  * @brief wakeCompletedProcess
356 |  * Wake up all the processes whose tasks have been terminated
357 |  * Can be invoked only when the kernel is active
358 |  */
359 | void wakeCompletedProcess(){
360 |     for(unsigned int i=0;i<Q;i++){
361 |         if((queue[i].size>0)&&(queue[i].remaining==0)){
362 |             if(threadId[i]!=NULL){
363 |                 threadId[i]->wakeup();
364 |             }
365 |             else{
366 |                 SuspendManager::wakeUpProcess(queue[i].processId);
367 |             }
368 |         }
369 |     }
370 | }

```

Figure 13: wakeCompletedProcess method

3.2 Request registration and query

In this subsection we discuss about the methods used by the rest of the operating system to interact with the module to allocate a new acquisition job or retrieve the data of a completed one (Figure 14).

The former (setQueue) just initializes a new jobs after checking that the provided data are correct.

The latter (readQueue) retrieves the reads from a completed job and frees the place in the array occupied by it.

```

89 | /**
90 |  * @brief setQueue
91 |  * Initialize a new queue
92 |  * @param processId id of the process which own the queue
93 |  * @param threadId id of the thread which invoke the method
94 |  * @param deviceId id of the selected source
95 |  * @param size number of sample to be captured
96 |  * @param sampling interval in milliseconds
97 |  * @return 0 if the work is scheduled, a negative integer if an error occurred
98 |  */
99 | int setQueue(pid_t processId, Thread* threadId, uint32_t deviceId, unsigned int size, unsigned int period) {
100 |     Lock<Mutex> lock(sharedData);
101 |
102 |     if ((size == 0) || (size > N)) {
103 |         return -1;
104 |     } else if (period < 1000) {
105 |         return -2;
106 |     } else if (getQueueFromProcessId(processId)-1) { //Check that no other queues exist fo the process
107 |         return -3;
108 |     }
109 |
110 |     int index = getFirstFreeQueue();
111 |     if (index < 0) {
112 |         return -3;
113 |     }
114 |     initQueue((unsigned int) index, processId, threadId, deviceId, size, period);
115 |     if (smartSensingThread==NULL){
116 |         smartSensingThread=wakeup();
117 |         smartSensingThread=NULL;
118 |     }
119 |     return 0;
120 | }
121 |
122 | /**
123 |  * @brief readQueue
124 |  * Retrieve data from a completed read
125 |  * @param processId id of the process which own the queue
126 |  * @param data pointer to the place where the reads will be stored
127 |  * @param n if an error occurred, number of sample written otherwise
128 |  * the error can occur if the process hasn't a queue or if the queue
129 |  * hasn't been completely filled yet.
130 |  * In every case the process queue is reseted.
131 |  */
132 | int readQueue(pid_t processId, unsigned short* data) {
133 |     Lock<Mutex> lock(sharedData);
134 |     int i=getQueueFromProcessId(processId);
135 |     if(i<0){
136 |         return -1;
137 |     }
138 |     if (queue[i].remaining==0){
139 |         resetQueue(i);
140 |         return -1;
141 |     }
142 |     unsigned int writingSize = queue[i].size;
143 |     for (unsigned int j = 0; j < writingSize; j++) {
144 |         data[j] = queue[i].data[j];
145 |     }
146 |     resetQueue(i);
147 |     return writingSize;
148 | }

```

Figure 14: setQueue and readQueue methods

3.3 Kernel daemon

It has the task of performing reads which can occur when the kernel is active.

It is simply composed of a thread (Figure 15) which delays itself on the next acquisition time, performs the read and so on. If there are no active jobs it just wait until a new one is created.

It's interesting to note the peculiar usage of C++ RAII design pattern in the lock and unlock mechanisms: a mutex lock (or unlock) is associated to an object. When it goes out of scope, the mutex is automatically unlocked (or locked); otherwise i.e. the lack of a mutex lock after the unlock at line 453 appears wrong. This usage is encouraged by the Miosix design.

```

425 | /**
426 |  * @brief getDaemonSleepTime
427 |  * Retrieve the time the kernel daemon must wait before performing a read
428 |  * It's supposed to be invoked only if the kernel is active
429 |  * @return time in milliseconds before the next read
430 |  */
431 | unsigned int getDaemonSleepTime() const{
432 |     long long nextRead = (long long)getNextEvent(getTick(),0);
433 |     long long currentTime = getTick();
434 |     if (nextRead && (nextRead-currentTime>0)){
435 |         return nextRead-currentTime;
436 |     }
437 |     return 0;
438 | }
439 |
440 | /**
441 |  * @brief innerThread
442 |  * This method implements the kernel daemon loop
443 |  * It's supposed to be invoked only if the kernel is active
444 |  */
445 | void innerThread(){
446 |     Lock<Mutex> lock(sharedData);
447 |     for(;;){
448 |         updateQueue(getTick());
449 |         wakeCompletedProcess();
450 |         if (getNextEvent(getTick(),0)==0){
451 |             smartSensingThread=Thread::getCurrentThread();
452 |             {
453 |                 Unlock<Mutex> unlock(sharedData);
454 |                 Thread::wait();
455 |             }
456 |         }
457 |         unsigned int sleepTime=getDaemonSleepTime();
458 |         if (sleepTime>0){
459 |             Unlock<Mutex> unlock(sharedData);
460 |             Thread::sleep(sleepTime);
461 |         }
462 |     }
463 | }
464 |

```

Figure 15: kernel daemon thread

3.4 System integration

In this subsection we discuss the system integration from the smart sensing module point of view. We defer to the next section the description of how the whole system interact with the module.

Basically it's composed by the methods onBoot and onSuspend (Figure 16).

The onBoot method is executed before that the system is fully started. It has the task of perform the required reads and either continue the boot process if a full start is required (or is the first boot of the board) or reschedule the next wakeup and suspend the board.

The onSuspend method instead deals with the definition of the next wake up time that must include the ones required for the scheduled reads.

```

150  /**
151  * @brief onBoot
152  * Hook of the boot process
153  */
154  void onBoot() {
155  if (firstBoot()) {
156  init();
157  } else {
158  #ifdef CHECK_CANARY
159  if (status->canary!=CANARY){
160  errorHandler(UNEXPECTED);
161  }
162  #endif
163  updateQueue(getTick());
164  debugZnt(getNextSecond(getTick(),status->nextSystemRestart*1000));
165  if ((completedTask)|| (status->nextSystemRestart*1000 <= (unsigned long long)getTick())) {
166  status->nextSystemRestart=0;
167  IRQbootLog("SS: Restart\r\n");
168  return;
169  }
170  } else if (getNextEvent(getTick(),status->nextSystemRestart*1000) == 0) {
171  errorHandler(UNEXPECTED);
172  }
173  } else {
174  updateQueue(getTick()+500); //No problem
175  if (completedTask){
176  return;
177  }
178  //Go to sleep!
179  SuspendManager::suspend(getNextSecond(getTick(),status->nextSystemRestart*1000));
180  }
181  }
182  }
183  }
184  /**
185  * @brief onSuspend
186  * Hook of the suspension process
187  * @param resumeTime time in seconds to be suspended
188  */
189  void onSuspend(unsigned long long resumeTime) {
190  Lock<Data> lock(sharedData);
191  status->nextSystemRestart = resumeTime;
192  unsigned long long currentTime=getTick();
193  updateQueue(currentTime+500); //No problem
194  if (completedTask){
195  return;
196  }
197  SuspendManager::suspend(getNextSecond(currentTime,status->nextSystemRestart*1000));
198  }

```

Figure 16: onBoot and onSuspend methods

```

40 int main()
41 {
42  Thread::create(ledThread_STACK_MIN);
43  SuspendManager::startHibernationDaemon();
44  iprintf("tick=%llu\n",getTick());
45  if (firstBoot())
46  {
47  puts("First boot");
48  SMART_SENSING::getSmartSensingInstance()->startKernelDaemon();
49  } else {
50  puts("RTC boot");
51  SuspendManager::resume();
52  SMART_SENSING::getSmartSensingInstance()->startKernelDaemon();
53  int ec;
54  Process::wait(6ec);
55  iprintf("Process terminated\n");
56  if (WIFEXITED(ec))
57  {
58  iprintf("Exit code is %d\n",WEXITSTATUS(ec));
59  } else if (WIFSIGNALED(ec)) {
60  if (WTERMSIG(ec)==SIGSEGV) iprintf("Process segfaulted\n");
61  }
62  }
63  }
64  ElfProgram prog1(reinterpret_cast<const unsigned int*>(test1_elf),test1_elf_len);
65  ElfProgram prog2(reinterpret_cast<const unsigned int*>(test2_elf),test2_elf_len);
66  for (int i=0;i++)
67  {
68  getchar();
69  pid_t child=Process::create(prog1);
70  Process::create(prog2);
71  int ec;
72  pid_t pid;
73  if (i%2==0) pid=Process::wait(6ec);
74  else pid=Process::waitpid(child,6ec,0);
75  iprintf("Process %d terminated\n",pid);
76  if (WIFEXITED(ec))
77  {
78  iprintf("Exit code is %d\n",WEXITSTATUS(ec));
79  } else if (WIFSIGNALED(ec)) {
80  if (WTERMSIG(ec)==SIGSEGV) iprintf("Process segfaulted\n");
81  }
82  }

```

Figure 18: main procedure

4 System integration

In this section is described how the new module is integrated into Miosix, in particular in three phases: the boot, the suspension and the process syscall.

4.1 Boot integration

The boot integration is performed by placing a hook in the boot process of Miosix (Figure 17) just before that the kernel is started with the call of onBoot, which has been previously presented, at line 134.

```

122  /**
123  * \internal
124  * Performs the part of initialization that must be done before the kernel is
125  * started, and starts the kernel.
126  * This function is called by the stage 1 boot which is architecture dependent.
127  */
128  extern "C" void _init()
129  {
130  using namespace miosix;
131  if (areInterruptsEnabled()) errorHandler(INTERRUPTS_ENABLED_AT_BOOT);
132  IRQbspInit();
133  //After IRQbspInit() serial port is initialized, so we can use BOOTLOG
134  SMART_SENSING::getSmartSensingInstance()->onBoot();
135  IRQbootLog(getMiosixVersion());
136  IRQbootLog("\r\nStarting Kernel... ");
137  //Create the first thread, and start the scheduler.
138  Thread::create(mainLoader,MAIN_STACK_SIZE,MAIN_PRIORITY,NULL);
139  startKernel();
140  //Never reach here
141  }

```

Figure 17: _init procedure

Moreover has been modified the main procedure (Figure 18), called after the completion of the kernel boot process. This has been done in order to start the daemon which performs the reads occurring when the system is active. The modifications are the addition of call to startKernelDaemon at lines 48 and 52.

4.2 Suspension integration

This part is particularly important because it coordinates the suspension manager of Miosix and our smart sensing module.

The suspension manager has the task of keeping track of which processes are active and if there are no active task it can decide to put the system into suspension.

It's important to specify that at the moment all the processes which are "suspension-ready" (the ones who are in a state where can be useful to suspend the system) have an associated resume time.

The only syscall that allows the system suspension is the sleep one, mainly because it's the only syscall during which the system can really do nothing; while in all other cases the process is waiting for something carried out by another process/the operating system itself, which requires the operating system being active (as we see later this project have caused further modifications regarding this aspect).

So if there are no working processes and the suspension manager decides to put the system into suspension (this happens in order to avoid the suspension process overhead when the time to be spent in suspension is too little) it saves the status of all the processes into the MRAM/backup SRAM and then the system is put into suspension.

This task is performed, in particular, by the hibernateDaemon procedure (Figure 19) which is run in a separate thread.

More precisely the last part of it (which sets the suspension time and actually performs the suspension) has been divided in the suspend method (Figure 20). In its stead has been put a call to the onSuspend method (line 217), which has been previously discussed.

This has been done in order to allow the smart sensing module to change the suspension time.

In addition this was necessary because the suspend method is called also by the onBoot method in smart sensing module, which has the necessity of performing a new suspension after that some reads were carried out. In this special case there is no need of saving the state (it hasn't been already changed) while it's useful to re-suspend the system.

```

156 #/
157 *The following thread function wait for all the processes to be suspended and
158 * then decide whether to hibernate or not..
159 #/
160 void SuspendManager::hibernateDaemon(void*)
161 {
162     Lock<Mutex> l(suspMutex);
163     for(;;)
164     {
165         hibernating.wait(l);
166         syscallReturnTime.sort(compareResumeTime);
167         list<syscallResumeTime>::iterator it;
168         it=syscallReturnTime.begin();
169         //NOTE: the following if, as well as the upper and lower bounds,
170         //will be replaced by the policy, once refined
171         if((it->resumeTime-getTick()/1000)<hibernationThreshold) continue;
172
173         list<ProcessStatus>* backupProc;
174         for(it=syscallReturnTime.begin();it!=syscallReturnTime.end();it++){
175             map<pid_t,Process*>::iterator proc=Process::processes.find(it->pid);
176             ProcessStatus* procStatus;
177             if((proc!=Process::processes.end())&&(proc->second->numActiveThreads==0)){
178                 procStatus=new ProcessStatus(*it->status);
179                 backupProc->push_back(procStatus);
180                 it->status=procStatus;
181             }
182         }
183
184         ProcessStatus* proc=getProcessesBackupAreaBase();
185         fprintf("Swapping %d processes\n",suspendedProcesses.size());
186         list<Process*>::iterator findProc;
187
188         for(findProc=suspendedProcesses.begin();
189             findProc!=suspendedProcesses.end();findProc++)
190         {
191             (*findProc)->serialize(proc);
192             if((*findProc)->toBeSwappedOut)
193             {
194                 //TODO: optimize
195                 Mram& mram=Mram::instance();
196                 mram.exitSleepMode();
197                 //Copy the process image from RAM to MRAM
198                 mram.write(reinterpret_cast<unsigned int*>
199                     (*findProc)->image.getProcessBasePointer(),
200                     ProcessPool::instance().getBaseAddress(),
201                     (*findProc)->image.getProcessBasePointer(),
202                     (*findProc)->image.getProcessImageSize());
203                 mram.enterSleepMode();
204                 //Now serialize the state of the SRAM allocator
205                 ProcessPool::instance().serialize(getBackupSramBase());
206                 //Now serialize the state of the backup SRAM allocator
207                 (*getBackupSramBase()+
208                     getAllocatorSramAreaSize()/sizeof(int))=
209                     suspendedProcesses.size();
210             }
211             proc++;
212         }
213         for(list<ProcessStatus*>::iterator it=backupProc.begin();it!=backupProc.end();it++){
214             delete *it;
215         }
216         SMART_SENSING::getSmartSensingInstance()->onSuspend(syscallReturnTime.begin()->resumeTime);
217     }
218 }
219 }

```

Figure 19: hibernateDaemon procedure

```

221 void SuspendManager::suspend(unsigned long long resumeTime) {
222     long long prev = getTick();
223     int sleepTime = resumeTime - prev / 1000;
224     getBackupSramBase()[1021] = prev & 0xfffffff; //FIXME: hack
225     getBackupSramBase()[1022] = prev >> 32;
226     getBackupSramBase()[1023] = sleepTime;
227     doSuspend(sleepTime);
228 }

```

Figure 20: suspend method

4.3 Syscall integration part 1

The project has the purpose of give an useful tool to the processes. However in order to be utilized by them there must be provided a way of communication between the processes and the smart sensing module. This has been first of all achieved by the usage of non-blocking syscalls: one for schedule the reads, and one for retrieving the data. In order to make them work they must be spaced by a sleep of the required time.

This method has the advantage of having a simple implementation (the behaviour of the suspension manager has not to be modified) and has been chosen to test the entire module.

A more advanced implementation is presented in the next subsection.

4.4 Syscall integration part 2

In this implementation all the work is performed by a single syscall which is a blocking read.

This has required an additional effort because first of all the smart sensing module has to wake up a process whose required job has been completed. Then also the suspension manager logic has to be modified because there is a new suspension-prone syscall which hasn't a defined resume time. In addition also the process manager has to be modified because the new particular syscall has to be resumed INSIDE the same syscall, while previously was resumed at the first instruction of the process AFTER the syscall. This last problem is caused by the peculiar way in which the suspension is performed: during a suspension ALL the kernel data is lost and only the processes ones data (i.e. stack and registers) are saved.

To cope with this several modification has been made.

First of all has been added in the SuspendManager class a new method, wakeUpProcess (Figure 21), which has the task of wake up a process whose smart sensing operation has finished.

```

100 void SuspendManager::wakeUpProcess(pid_t processId){
101     map<pid_t,Process*>::iterator findProc;
102     Lock<Mutex> l(suspMutex);
103     for(list<syscallResumeTime>::iterator i=syscallReturnTime.begin();i!=syscallReturnTime.end();i++){
104         if(i->pid==processId){
105             findProc=Process::processes.find(processId);
106             if(findProc!=Process::processes.end()){
107                 fprintf("** PID %i %i\n",processId,i->status->pid);
108                 fprintf("T: %S%i PID: %i RESTART SSIN: %i\n",getTick(),processId);
109                 Process::create(i->status,i->threadNum,true);
110             }
111             syscallReturnTime.erase(i);
112             return;
113         }
114     }
115 }
116 #/
117 *The following thread create the threads of processes at the time they must be
118 * resumed
119 #/
120 void SuspendManager::wakeUpDaemon(void*)
121 {
122     map<pid_t,Process*>::iterator findProc;
123     Lock<Mutex> l(suspMutex);
124     while(syscallReturnTime.empty()!=false)
125     {
126         syscallResumeTime ret=syscallReturnTime.front();
127         if(ret.resumeTime==getTick()/1000)
128         {
129             findProc=Process::processes.find(ret.pid);
130             //check if the process is already alive, it could happen that
131             //the main thread has already been spawned and is also terminated
132             //so other threads waiting to be resumed must be not be created.
133             //In any case, at the end of the cycle, the process must be
134             //erased from the syscallReturnTime list
135             if(findProc!=Process::processes.end()){
136                 //FIXME: This causes exception
137                 fprintf("T: %S%i PID: %i RESTART GENERIC\n",getTick(),ret.pid);
138                 Process::create(ret.status,ret.threadNum);
139             }
140             syscallReturnTime.pop_front();
141         }
142         else {
143             long long resumeTime=ret.resumeTime;
144             {
145                 Unlock<Mutex> u(l);
146                 int currentTime=resumeTime-getTick()/1000;
147                 if(currentTime>0){
148                     sleep(currentTime);
149                 }
150             }
151         }
152     }
153 }
154 }

```

Figure 21: wakeUpProcess and wakeUpDaemon methods

Then the process manager of Miosix has been modified in order to perform a two part syscall.

First has been added a variable to the Process class which keep track of a pending syscall related to the process. The method create in the class Process (Figure 22) has been modified by adding the variable pendingOperation, which if it's true, it means that a smart sensing read syscall is pending.

```

120 pid_t Process::create(ProcessStatus* status, int threadId, bool pendingOperation)
121 {
122     Lock<Mutex> l(SuspendManager::suspMutex);
123     map<pid_t, Process*>::iterator findProc;
124     findProc=processes.find(status->pid);
125
126     if(findProc==processes.end())
127         throw runtime_error("unable to recreate the process after hibernation");
128     Process* proc=findProc->second;
129
130     proc->pendingOperation = pendingOperation;
131
132     proc->image.resume(status);
133
134     //TODO: look at it -- begin
135     #ifndef __CODE_IN_XRAM
136     //FIXME -- begin
137     //Till a flash file system that ensures proper alignment of the programs
138     //loaded in flash is implemented, make the whole flash visible as a big MPU
139     //region
140     extern unsigned int _etext asm("_etext");
141     unsigned int flashEnd=reinterpret_cast<unsigned int>(&_etext);
142     if(flashEnd & (flashEnd-1)) flashEnd-1<<fibs(flashEnd);
143     #ifdef MPU_MIOSIX_PRIVATE:MPUConfiguration(0, flashEnd,
144     proc->image.getProcessBasePointer(), proc->image.getProcessImageSize());
145     // MPU_MIOSIX_PRIVATE:MPUConfiguration(program.getElfBase(), roundedSize,
146     // image.getProcessBasePointer(), image.getProcessImageSize());
147     //FIXME -- end
148     #else // CODE_IN_XRAM
149     loadedProgram=ProcessPool::instance().allocate(roundedSize);
150     memcpy(loadedProgram, program.getElfBase(), elfSize);
151     MPU_MIOSIX_PRIVATE:MPUConfiguration(loadedProgram, roundedSize,
152     image.getProcessBasePointer(), image.getProcessImageSize());
153     #endif // __CODE_IN_XRAM
154     //TODO: look at it -- end
155
156
157     Thread *thr=Thread::createUserspace(Process::start,
158     status->interruptionPoints[threadId].registers,
159     Thread::DEFAULT, proc);
160     proc->toBeSwappedOut=true;
161
162     if (thr==0)
163     {
164         Lock<Mutex> l(procMutex);
165         processes.erase(proc->pid);
166         if(Thread::getCurrentThread()->proc!=0)
167         {
168             Thread::getCurrentThread()->proc->childs.remove(proc);
169             also kernelChilds.remove(proc);
170             delete proc;
171             throw runtime_error("Thread recreation failed");
172         }
173         //Cannot throw bad_alloc due to the reserve in Process's constructor.
174         //This ensures we will never be in the uncomfortable situation where a
175         //thread has already been created but there's no memory to list it
176         //among the threads of a process
177         proc->threads.push_back(thr);
178         proc->numActiveThreads++;
179         thr->wakeup(); //Actually start the thread, now that everything is set up
180         pid_t result=proc->pid;
181         return result;
182     }

```

Figure 22: create method

This variable is used in the start method of Process class (Figure 23), which is used to complete the syscall before continuing the execution of the process, with the method completeSmartSensingOperation that simply copy the reads into the memory of the process and set as return value the number of bytes read (as the standard POSIX read requires).

```

478 void Process::completeSmartSensingOperation(miosix_private::SyscallParameters &sp){
479     sp.setReturnValue(miosix::SMART_SENSING::getSmartSensingInstance()->readQueue(pid,
480     reinterpret_cast<short unsigned int*>(sp.getSecondParameter()));
481 }
482
483 void *Process::start(void *argv)
484 {
485     Process *proc=Thread::getCurrentThread()->proc;
486     if(proc==0) errorHandler(UNEXPECTED);
487     unsigned int entry=proc->program->getEntryPoint();
488     #ifndef __CODE_IN_XRAM
489     entry=reinterpret_cast<unsigned int>(proc->program->getElfBase())+
490     reinterpret_cast<unsigned int>(proc->loadedProgram);
491     #endif // __CODE_IN_XRAM
492     if(proc->suspended==false)
493         Thread::setupUserspaceContext(entry, proc->image.getProcessBasePointer(),
494         proc->image.getProcessImageSize());
495     else
496     {
497         if(argv){
498             if(proc->pendingOperation){
499                 iprintf("RESTART\n");
500                 miosix_private::SyscallParameters sp(reinterpret_cast<unsigned int*>(argv));
501                 proc->completeSmartSensingOperation(sp);
502             }
503             Thread::resumeUserspaceContext(reinterpret_cast<unsigned int*>(argv));
504         }
505         else
506             errorHandler(UNEXPECTED);
507         //in the following block the process is removed from the list of the
508         //suspended processes, after hibernation
509         {
510             Lock<Mutex> l(SuspendManager::suspMutex);
511             proc->suspended=false;
512             SuspendManager::suspendedProcesses.remove(proc);
513         }
514     }
515     bool running=true;
516     do {

```

Figure 23: completeSmartSensingOperation and part of start methods

The core of the syscall is presented at Figure 24. Here we can notice a switch over the possible syscall, where the interesting part is the read one (with id equal to 4).

The code first of all check the value of the file description. If it's equal to 4 it's interpreted as 5 reads from the potentiometer spaced by 1 second each.

This trick has been used because the filesystem management is beyond the scope of this project and this has been a rapid way for testing the smart sensing mechanism. We can imagine that the information of device id, number of reads and period can be memorized in some data structures by a modified open syscall and then utilized here.

So the smart sensing operation is set and at line 567 is signaled to the suspension manager that the process can enter to suspension. Here is used another trick to maintain as much as possible the compatibility with the suspension manager: the smart sensing read is treated like a sleep of a very long interval.

The rest of the syscall (line 568-569) wait and complete the syscall in the case that all the smart sensing operation is performed without putting the system in suspension. In this case the thread of the process is woken up by the smart sensing class when the job is finished. The whole mechanism works because in case of suspension all the kernel data is lost and so both the process thread status and the thread id inside the SmartSensing class.

The last addition is in line 598, which calls the method cleanUp in the smart sensing module where the process is being destroyed. This is necessary to eventually free the queues of the process inside the SmartSensing class.

```

548 switch(sp.getSyscallId())
549 {
550     case 2:
551         running=false;
552         proc->waitCode=(sp.getFirstParameter() & 0xff)<<8;
553         break;
554     case 3:
555         sp.setReturnValue(write(sp.getFirstParameter(),
556         reinterpret_cast<const char*>(sp.getSecondParameter()),
557         sp.getThirdParameter()));
558         break;
559     case 4:
560         //STUB: ONLY FOR TESTING
561         if(sp.getFirstParameter()==4){
562             unsigned int size = sp.getThirdParameter()/sizeof(unsigned short int);
563             if(size>SMART_SENSING::getSmartSensingInstance()->setQueue(proc->pid, Thread::getCurrentThread(),
564             POTENTIOMETER_ID, size, 1000)<0){
565                 sp.setReturnValue(-1);
566                 break;
567             }
568             SuspendManager::enterInteruptionPoint(proc, threadId, BIG_TIME, 5, -1);
569             Thread::getCurrentThread()->wait();
570             proc->completeSmartSensingOperation(sp);
571             break;
572         }
573         //FIXME: check that the pointer belongs to the process
574         sp.setReturnValue(read(sp.getFirstParameter(),
575         reinterpret_cast<char*>(sp.getSecondParameter()),
576         sp.getThirdParameter()));
577         break;
578     case 5:
579         if(sp.getFirstParameter()==1000000)
580             SuspendManager::enterInteruptionPoint(proc, threadId,
581             sp.getFirstParameter()/1000000, 5, -1);
582         sp.setReturnValue(sp.getThirdParameter());
583         break;
584     default:
585         running=false;
586         proc->waitCode=SIGSYS; //Bad syscall
587         #ifdef NETX_ERPLOG
588         iprintf("Unexpected syscall number %d\n", sp.getSyscallId());
589         #endif //NETX_ERPLOG
590         break;
591 }
592
593 }
594
595 while(running){
596     Lock<Mutex> l(procMutex);
597     SMART_SENSING::getSmartSensingInstance()->cleanUp(proc->pid);
598     proc->numBis+=fibs;
599     list<Process*>::iterator it;
600     for(it=proc->childs.begin(); it!=proc->childs.end(); ++it) (*it)->ppid=0;
601     for(it=proc->zombies.begin(); it!=proc->zombies.end(); ++it) (*it)->ppid=0;
602     kernelChilds.splice(kernelChilds.begin(), proc->childs);
603     kernelZombies.splice(kernelZombies.begin(), proc->zombies);

```

Figure 24: extract from start method

In summary we recap the whole syscall mechanism which is quite complex due to the various suspension call. This can be summed up with the following steps:

1. The process invoke a smart sensing read syscall.
2. The syscall manager inside the kernel sets up the smart sensing operation, signals to the suspension module the possibility of a suspension and waits.

3. If the operation is completed without any suspension the process wakes up and the syscall returns normally, otherwise a suspension is performed.
4. The system can be started and suspended several times. The last one the smart sensing module calls the `wakeUpProcess` in the `SuspendManager` class directly if the read occurs when the kernel is already started or when the smart sensing daemon is started if the last read happens before the boot process.
5. The `wakeUpProcess` wake up the corresponding process with the flag of the pending syscall set.
6. The start method in the `Process` class detects the pending syscall, completes it and return the control to the process.
7. The process continues its execution as nothing were happened between.

5 Testing

For the testing of the project were used two programs. The first (**Figure 25**) performs a series of a smart sensing acquisition composed of 5 reads spaced by 1 second and 3 seconds of active wait (it's like a sleep but it doesn't allow the system going into suspension).

```

40 void printData(unsigned short* data,unsigned int len){
41     unsigned int i;
42     char text[]{"Read 1:"};
43     for(i=0;i<len;i++){
44         write(1,text,mystrlen(text));
45         debugHex(data[i]);
46     }
47 }
48
49 int main()
50 {
51     static const char str[]{"Test 1\n"};
52     static const char str2[]{"Unexpected command\n"};
53     for(;;)
54     {
55         char result[100];
56         unsigned short data[100];
57         write(1,str,mystrlen(str));
58
59         write(1,str,mystrlen(str));
60         int l=read(4,data,10);
61
62         char text[]{"Test 1\n"};
63         write(1,text,mystrlen(text));
64         printData(data,l);
65         int i;
66         for(i=0;i<3000;i++){
67             waitl();
68         }
69         usleep(4000000);
70     }
71 }
72

```

Figure 25: code of the first testing program

The second one (**Figure 26**) instead performs just a series of smart sensing acquisition again of 5 reads spaced by 1 second.

```

24
25 void printData(unsigned short* data,unsigned int len){
26     unsigned int i;
27     char text[]{"Read 2:"};
28     for(i=0;i<len;i++){
29         write(1,text,mystrlen(text));
30         debugHex(data[i]);
31     }
32 }
33 }
34
35 int main()
36 {
37     static const char str[]{"Test 2\n"};
38     static const char str2[]{"Unexpected command\n"};
39     for(;;)
40     {
41         char result[100];
42         unsigned short data[100];
43         write(1,str,mystrlen(str));
44
45         write(1,str,mystrlen(str));
46         int l=read(4,data,10);
47         char text[]{"Test 2\n"};
48         write(1,text,mystrlen(text));
49         printData(data,l);
50     }
51 }
52
53

```

Figure 26: code of the second testing program

The resulting log correctly outlines the alternation of smart sensing reads and system activations; also the kernel daemon part behaves correctly and the values read from the potentiometer were correct.

References

- [1] Wolfgang Wieser. *Programming STM32 F2, F4 ARMs under Linux: A Tutorial from Scratch*, 2012.
- [2] Trevor Martin. *The Insider's Guide To The STM32 ARM Based Microcontroller - 2nd Edition*. 2009.
- [3] ST Microelectronics. *STM3220G-EVAL evaluation board user manual (UM1057)*, 2012.
- [4] ST Microelectronics. *STM32F2xx reference manual (RM0033)*, 2011.
- [5] ST Microelectronics. *STM32F2xx programming manual (PM0056)*, 2011.
- [6] Federico Terraneo. *Control based design of OS components*, 2011.