Detection of Cough Events

By Tyler Coleman, ECE '17

Introduction

Multi-drug resistant (MDR) tuberculosis is an emerging danger to world health that threatens our ability to treat what is already one of those most deadly infectious diseases on earth. Cases of MDR tuberculosis are primarily encountered in developing areas of India, China, Eastern Europe, South America, and southern Africa. The infections in these areas are often unfortunately worsened by a lack of access to labs, equipment, and skilled technicians, making diagnosis and treatment all the more difficult.

Team Periwinkle's mission was to develop a reliable, hands-free device that effectively monitors a patient over a 24-hour period, with the sole purpose of identifying coughs and logging the times at which they occur. Preliminary research suggests a link between a positive response to treatment and a drop in patients' cough rates within the first two weeks of treatment. The desire for our device stems from the urgent need for technology supporting the advancement of this research in developing countries. It is our hope that our device will present a low-cost, reliable solution to aid these communities.

Cough detection sounds like a simple enough endeavor. After all, we have phones that can interpret human speech in a variety of languages right in our pockets. However, multiple key considerations must be taken into account during the development of a device whose sole purpose is to detect and record cough events.

Logging an event consists of two stages: the detection of the cough, and then the classification of the detected event. Once recorded, the cough's audio sample is stored to be later uploaded to a server for classification, where processing power is – relatively speaking – just about unlimited. Before that can happen, however, our device must be capable of detecting potential cough events to be recorded. These paper explores the challenges presented by hardware limitations when attempting to accomplish this task.

The Design Goals

Previous research (Tracey, Comina, & Larson, 2011) in this area has led to algorithms that are capable of analyzing long (24+ hour) audio samples, and identifying individual cough events within that extended file. One of our main goals for this project is to develop hardware capable of deciding when to actually record – essentially removing a large amount of unnecessary data from these long audio samples before they're even operated on to classify cough events. This accomplishes a few things:

1. The space requirement for stored data is significantly reduced
2. Power consumption is significantly reduced – the process no longer has to deal with the constant sampling and writing to memory of audio data
3. Privacy concerns are alleviated

We feel the last point is especially pertinent. The entire aim of this device is to aid in the treatment of a serious illness. The last thing we want is to induce unnecessary stress, and asking a patient to wear a device that records their every word for 24 hours is hardly a way to avoid doing so.

A key feature of the device is that it monitors accelerometer input and microphone input simultaneously to get a better idea of whether or not an event might be a cough. By looking at microphone input alone, we have no way of really knowing the difference between a cough and any other ambient noise that might carry with it similar energy until we upload the sample for further analysis. For our application of 24-hour wear that isn't meant to impact the life of the patient, this could mean recording every dog bark and car horn the patient passes by on their way to work. Not ideal. Similarly, quick movements could fool an accelerometer on its own into thinking a cough may have taken place, even though no sound was generated. By using the two in conjunction, we are able to observe and wait for a moment in which we identify sufficient volume associated with the sound of a cough, as well as sufficient movement detected by accelerometers placed at points related to the biomechanics of a cough. This all results in fewer recordings being stored for processing, which helps with all three design goals mentioned above.

The Hardware Limitations

Compared to the smart phones in our pockets, the device with which we intend to monitor our patient is packed with significantly less processing power and far less memory available for raw sample analysis. Each sensor – one microphone and up to four accelerometers – are continually sampled and stored temporarily in SRAM within our processor. To visualize this, imagine a container that you can partition and fill up with contents however you like, like the one shown to the right.

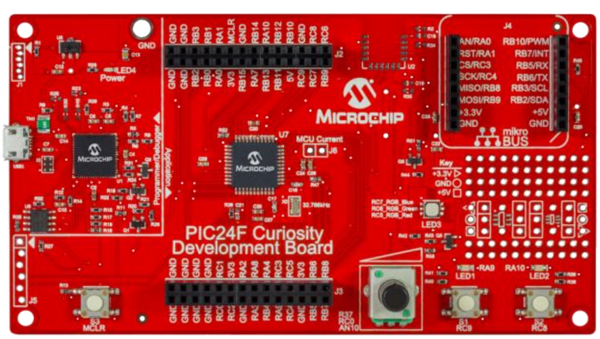

The development board on which our project is implemented. Contains about 1 millionth the memory and 1 thousandth the processing speed of an ordinary desktop.

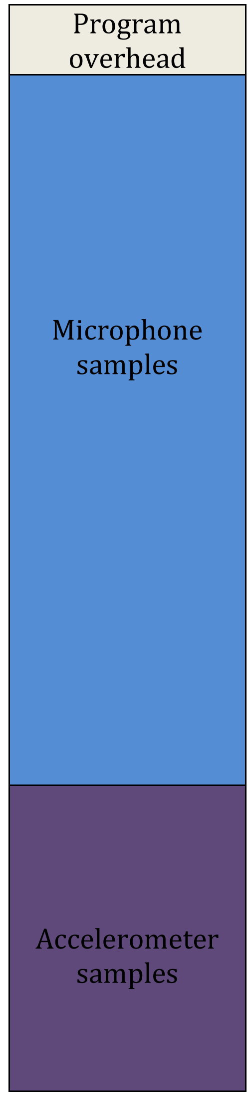

The Work-Around

In our case, we're filling up this container with samples – thousands of them per second. The samples themselves aren't of any use to us until they're organized like this, and all readily available for processing in their respective locations within memory. As new samples are read in, old samples are thrown out to make room for the new ones.

Periodically, a relatively simple routine is performed that calculates the energy within a series of samples. We compare this calculated energy to a pre-determined threshold to determine whether or not a series of entries in memory *might* be a cough. If we suspect an event has occurred, we begin writing audio data to an SD for more permanent storage. These are the files that are later analyzed off-board to characterize cough events within.

The Approach

Of course, how we partition memory is entirely up to our judgment and can have a profound impact on the performance of our device. Even when efficiently allocating space for each sensor, the most we can realistically store in SRAM is less than a quarter of a second of raw data at a time. This may not sound like much, but it is enough for our application. In fact, the great majority of the energy in a cough appears as a spike right at its beginning, during what is commonly referred to as the "explosive" phase of the cough (Smith, Earis, & Woodcock, 2006.). This spike is all that we need to begin our process of determining a cough.

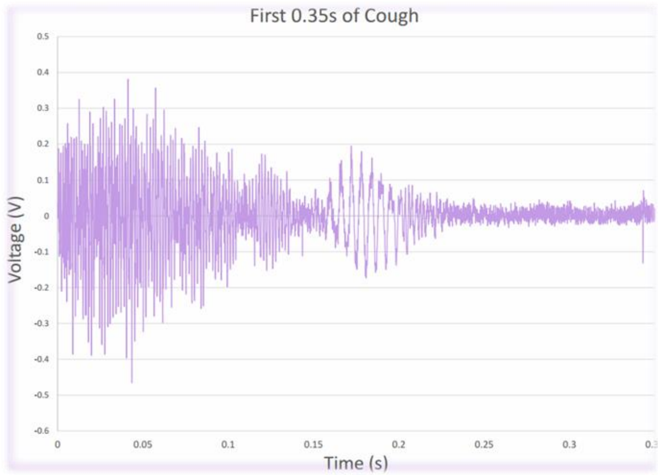

Visualization of cough audio. Note the spike in energy at the cough's beginning.

From this point, there are multiple approaches we can take. Our current idea is to trigger an automatic two second period in which all audio data is recorded, which should be sufficient time to capture the entirety of a cough in audio. If another spike is detected within this period (signaling the possibility that another cough has occurred), the timer is reset and at least another two seconds of audio is recorded. At the end of 24 hours, we would then expect to pull our SD card and obtain many audio files of at least two second length that can then be individually characterized.

Of course, it is highly unlikely that our method will prove the most efficient and best-performing in the long run. There are many ways to organize samples in memory. Another such approach could be to store more lengthy audio samples directly on the SD card, and analyze these even less often. By doing this, we could record all audio, and instead "trim away" any data that we later determine not to be a cough. There are countless possibilities on the hardware side of things to implement a cough detection algorithm given limited resources, each with their own advantages.

In Conclusion

In migrating to a more robust and localized detection algorithm, we hope to address the three design goals mentioned earlier. Our work this year is intended to provide a basis upon which future teams can expand and improve. The developing of a medical device is no easy feat, and comes with unique design goals that have influenced our work greatly. Finally, it is our hope that our work will have brought this device one step closer to its ultimate goal of helping those in need of better tuberculosis care.

References

1. Tracey, B., Comina, G., & Larson, S. (2011). Cough detection algorithm for monitoring patient recovery from pulmonary tuberculosis. IEEE Xplore.
2. Smith, J., Earis, J., & Woodcock, A. (2006). Establishing a gold standard for manual cough counting: video versus digital audio recordings. *Cough*. *Volume 2*.