Design and Performance Evaluation of an Implantable Ultrasonic Networking Platform for the Internet of Medical Things

G. Enrico Santagati, *Member, IEEE*, Neil Dave, *Student Member, IEEE*, and Tommaso Melodia, *Fellow, IEEE, Senior Member, ACM*

***Abstract*— Wireless networks of electronically controlled implantable medical sensors and actuators will be the basis of many innovative and potentially revolutionary therapies. The biggest obstacle in realizing this vision of networked implants is posed by the dielectric nature of the human body, which strongly attenuates radio-frequency (RF) electromagnetic waves. In this paper we present the first hardware and software architecture of an Internet of Medical Things (IoMT) platform with ultrasonic connectivity for intra-body communications. This platform can be used as a basis for building future IoT-ready medical implants and wearable devices. We demonstrate that ultrasonic waves can be efficiently generated and received with low-power mm-sized components, and that despite the conversion loss introduced by ultrasonic transducers, the attenuation of 2.4 GHz RF is substantially greater than ultrasound, e.g., ultrasound attenuates 70 dB less than RF at distances over 10 cm. We show that the proposed IoMT platform requires significantly less transmission power than 2.4 GHz RF with equal reliability, e.g., 35 dBm less at distances over 12 cm with 10^{-3} Bit Error Rate (BER) thus enabling a lower energy per bit and a longer device lifetime. Finally, we experimentally establish 2.4 GHz RF links do not function at distances greater than 12 cm, while ultrasonic links achieve a reliability of 10^{-6} BER up to a distance of 20 cm with less than 0 dBm transmission power.**

***Index Terms*— Body sensor networks, internet of things, acoustic communications.**

I. Introduction

WIRELESS networks comprised of implantable medical devices will be the basis of many innovative and possibly revolutionary therapies. Two major applications of these networks are *artificial pancreases*, i.e., implanted continuous glucose monitors wirelessly interconnected with adaptive insulin pumps which could improve the lives patients with type-1 diabetes and *ICU trauma monitoring sensors* that detect changes in pH or white blood cell concentration to prevent infections and monitor patient stability during recovery after intensive procedures. Other potential applications include functional electrical stimulation, a particular type of neurostimulation that attempts to restore motion in people with disabilities by injecting electrical currents into nerves to subsequently stimulate motion in extremities affected by paralysis. Future *neurostimulators* will be comprised of several miniaturized standalone stimulation devices that attach to different groups of neurons and wirelessly cooperate with each other to modulate electrical signals and restore healthy behaviors among targeted organs. Future *pacemakers* will consist of wirelessly connected sensing and pacing devices implanted in separate heart chambers which cooperatively enable advanced cardiac resynchronization therapy. Additionally in-vivo continuous multi-site monitoring of the heart and other organs may create a deeper understanding of how these organs react to different activities, external conditions, and drugs over extended periods of time. New and existing applications created by expanding the wireless capabilities of medical implants have fostered the growth of a rapidly expanding emergent market we refer to as *"The Internet of Medical Things"* (IoMT).

Limitations of Current Wireless Technology

Unfortunately, the dielectric nature of the human body generates major obstacles toward achieving this vision of *networked implantable devices*. Biological tissues are composed primarily of water (65%). Since radio-frequency (RF) waves are absorbed by aqueous mediums, higher transmission (Tx) power is needed to establish reliable links; this reduces the battery life (or, equivalently, increases the battery size) of the implantable device, in a domain where low power, miniaturization, and battery duration are major concerns. To put this in context, state-of-the-art pacemakers require power in the order of 50-100 μW for pacing [2], while commercial RF transceivers for implants operate in the order of tens of mW, making continuous telemetry impractical. Additionally, RF-based technologies are vulnerable to interference from other existing RF systems, and can be easily jammed or eavesdropped with cheap off-the-shelf devices. In-vivo RF use also pose safety concerns as microwaves are generally perceived as dangerous - the World Health Organization classifies RF waves as "possibly carcinogenic to humans".

Ultrasonic Wireless Communications

In this paper we present a superior alternative to RF communications in-vivo through the design and implementation of an ultrasonic wireless IoMT platform for implant connectivity

Manuscript received October 30, 2018; revised June 10, 2019; accepted September 24, 2019; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor H. Seferoglu. Date of publication January 20, 2020; date of current version February 14, 2020. This work was supported by the National Science Foundation under Grant CAREER CNS-1253309 and Grant CNS-1618731. A preliminary, shorter version of this article appeared in the Proceedings of IEEE Infocom 2017. *(Corresponding author: Tommaso Melodia.)*

The authors are with the Department of Electrical and Computer Engineering, Northeastern University, Boston, MA 02115 USA (e-mail: santagati@northeastern.edu; dave@northeastern.edu; melodia@northeastern.edu).

Digital Object Identifier 10.1109/TNET.2019.2949805

that serves as a basis for building medical implants that communicate safely, reliably and with low-power consumption. Ultrasounds are mechanical waves that propagate at frequencies above the upper limit of human hearing, i.e., $20\,\mathrm{kHz}$. Compared to RF-waves, ultrasounds are absorbed significantly less by human tissues (e.g., around $70\,\mathrm{dB}$ less attenuation for a $1\,\mathrm{MHz}$ ultrasonic link compared to a $2.45\,\mathrm{GHz}$ RF link over $10-20\,\mathrm{cm}$ [3]–[5]). Medical data obtained in recent decades has demonstrated that heat dissipated in tissues during ultrasound propagation is minimal compared to RF [6]. Due to this, the FDA allows a much higher transmission power for ultrasonic waves ($720\,\mathrm{mW/cm^2}$) in tissues compared to RF ($10\,\mathrm{mW/cm^2}$) [7]. By using ultrasound to connect medical implants, patients will benefit from devices that provide wireless real-time telemetry which minimally affect battery life. This will reduce the number of surgeries required to replace implant batteries. Real-time remote monitoring applications of the platform could reduce clinical visits and hospital length of stay, resulting in better use of medical resources. Low-power ultrasonic communications can enable advanced therapies that require reliable wireless data links through tissues between multiple implanted sensor and actuator devices. Additionally ultrasonic communications cannot be easily eavesdropped or jammed without physical contact making them significantly more secure than their RF counterparts and, eliminates the electromagnetic compatibility concern of adding to a already crowded RF spectrum. Ultrasound is transparent to the RF spectrum management procedures of healthcare facilities which in addition to its reliable biocompatible performance ultimately make it safer to use in RF dense environments.

Prior Work

The idea of using ultrasonic waves for intra-body communications has been previously investigated in [8]–[11]. In [8], the authors presented a channel model for ultrasonic intra-body communication, while in [9] the feasibility of nanoscale ultrasonic communications between nano-implants was studied. In [10], the authors investigated ultrasonic backscatter to enable communication between an external reader and an implantable device.

In [11], we proposed Ultrasonic WideBand (UsWB), an ultrasonic multipath-resilient physical (PHY) and medium access control (MAC) layer protocol, then experimentally demonstrated the feasibility of ultrasonic communications in gelatinous phantoms that mimic the properties of human tissue. In [12], we presented an experimental Mbit/s ultrasonic transmission scheme tested on ultrasonic phantoms and porcine meat, while in [13] ultrasonic wireless power transfer was demonstrated by remotely charging mm-sized implantable devices.

Ultrasonic waves have also been used in in-air or through bone short-range communications between smartphones and wearable devices [14]–[16]. In [14], the authors propose in-air communication systems that operate in the near-ultrasonic frequency range for low-data medium-range directional links, i.e., up to $20\,\mathrm{m}$.

In [15], we presented U-Wear, the first networking framework for wearable medical devices based on ultrasonic communication. In [16], the authors leverage the mechanisms of ultrasonic bone conduction to enable secure high-frequency acoustic communication between a smartwatch and a smartphone. Alternative intra-body communication technologies have been investigated using capacitive coupling [17] and galvanic coupling [18] where information is modulated either by an electrical potential or current flow. However, biocompatibility concerns limit transmission power and operational duty cycle to reduce thermal effects caused significant spectral absorption in tissues [19].

Paper Contribution

This paper presents the following core contributions:

- We present *the first hardware and software architecture of an IoMT platform with ultrasonic connectivity for in-vivo communication*. The IoMT platform consists of a modular, reconfigurable hardware and software architecture that can be flexibly adapted to different application and system specific scenarios to enable telemetry, remote control of medical implants, as well as implant-to-implant communications.
- We discuss the implementation of two size-, energy-, and resource-constrained prototypes based on the IoMT platform architecture, i.e., an implantable IoMT-mote and a wearable IoMT-patch, that implement and communicate with one another using state-of-the-art ultrasonic communication protocols. *The IoMT-mote is the first miniaturized software-defined implantable device with ultrasonic communication and networking capabilities. We also demonstrate, for the first time, the feasibility of ultrasonic communications using miniaturized, energy constrained embedded devices.*
- We extensively evaluate the performance of the ultrasonic connectivity offered by the IoMT prototypes in terms of energy consumption and communication reliability. Using ultrasonic phantoms and porcine meat as communication media, we *for the first time compare this ultrasonic IoMT platform against state-of-the-art low-power RF-based wireless technologies* operating in the industrial, scientific, and medical (ISM) $2.4\,\mathrm{GHz}$ band, e.g., Bluetooth Low Energy (BLE). The ISM band is currently the RF band of choice for wireless medical devices, because it easily interfaces medical devices with commercial devices such as smartphones through the existing commercial RF infrastructure, and offers larger bandwidth than other medical bands [20].
  Comparatively, we show that ultrasonic waves can be efficiently generated and received with low-power, mm-sized components, and that despite the conversion loss introduced by ultrasonic transducers the gap between $2.4\,\mathrm{GHz}$ RF waves and ultrasonic attenuation is still substantial, e.g., ultrasounds offer $70\,\mathrm{dB}$ less attenuation over $10\,\mathrm{cm}$. We show how the proposed IoMT platform requires significantly less Tx power compared to BLE with equal reliability, e.g., $35\,\mathrm{dBm}$ lower Tx power over $12\,\mathrm{cm}$ for 10^{-3} Bit Error Rate (BER) leading to lower energy per bit cost and a longer device lifetime. Finally, we show experimentally that BLE links do not function at $12\,\mathrm{cm}$, depths and greater while ultrasonic links achieve 10^{-6} BER up to $20\,\mathrm{cm}$ with less than $0\,\mathrm{dBm}$ Tx power.

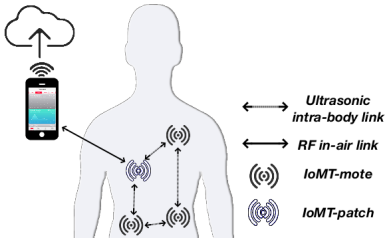

Fig. 1. IoMT-based System.

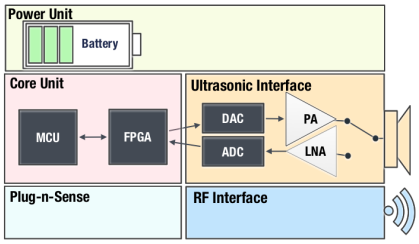

Fig. 2. IoMT platform hardware architecture.

The remainder of the paper is organized as follows.

In Section II, we describe the hardware and software architectures of the IoMT platform. In Section III and IV, we present the implementation of an IoMT-mote and an IoMT-patch. In Section V, we evaluate the performance of the IoMT ultrasonic connectivity, and we compare this with a state-of-the-art BLE chipset. Finally, in Section VI, we conclude the paper.

II. IoMT Platform Architecture

The IoMT platform is a modular software and hardware architecture to be used as a basis for future low-power IoMT-ready wearable and implantable devices that communicate wirelessly in-vivo using ultrasound. The IoMT platform enables (i) remote measurement, and cloud storage of patient physiology data measured by implantable sensors (telemetry); (ii) wireless control of actuators deployed in the body of the patient, e.g., stimulators, drug pumps, and pacing devices; (iii) treatments designed around closed-loop feedback through *implant-to-implant communications*, where actuators trigger responses to physiological data captured by sensors implanted throughout the body. For example, a smart coronary stent could detect clogs, allow doctors to remotely monitor the patient condition and automatically trigger drug delivery to prevent arterial re-occlusion [21]. Similarly a smart neurostimulator can be triggered by a heart rate sensor to anticipate an epileptic attack [22].

Figure 1 shows an application of the IoMT platform. A set of sensors and actuators (IoMT-motes) are deployed inside the patient and communicate with each other, or with wearable devices (IoMT-patches) through intra-body ultrasonic links (dotted lines). The IoMT-patches enable communication from the intra-body network to an access point connected to the Internet, e.g., a smartphone, through an RF link (continuous line). The platform architecture proposed was designed to easily and flexibly facilitate a network of various implanted sensors and actuators to enable multi-modal patient data acquisition and use multi-modal data to enable novel and innovative therapies. Compared to existing architectures which are optimized for one to one links, extra care was placed in the hardware and software design to provide users with reconfigurability in the physical, link and network layers to create custom intra-body networking solutions. Additional efforts were also placed into enabling the easy integration of this platform into existing medical devices.

A. Hardware Architecture

The IoMT-mote and the IoMT-patch are based on the modular ultrasonic IoMT modular hardware architecture shown in Figure 2. The *core unit* includes (i) a mm-size low-power field programmable gate array (FPGA) and (ii) a micro controller unit (MCU). This combination offers hardware and software reconfigurability in an ultra small form factor that in addition to its low energy consumption, provides a multitude of power management options to further optimize energy use in application specific scenarios. The miniaturized FPGA hosts the physical (PHY) layer communication functionalities. Reconfigurability at the physical layer was an intentional design choice done to maximize patient quality of life, even though it adds significant hardware complexity to the platform. The implantable nature of the platform creates unique challenges to ensuring the devices remain up to date as long as possible that are unlike the challenges faced by other wireless platforms.

For instance most wireless platforms use only a MCU in practice and update their firmware by pushing code through an "over air download" process. Firmware updates are sufficient for the lifetimes of most standard wireless chipsets however implants are intended to have a significantly higher device lifetime. Implant replacement is an invasive and undesirable procedure and so not only was the platform designed to be ultra-low power to facilitate an implant lifetime longer than the standard 5-10 years but also effort was put into reconfigurability at the physical layer through an FPGA so wireless updates can be pushed to the platform to reconfigure entire communication buses and update the platform to interface with future buses and technologies to remain up to date and functional in a rapidly changing wireless market.

The platform uses a unique combination of both an MCU and an FPGA to enable maximum reconfigurability across network layers but the functionality of the two additionally complement one another to create versatile nodes to implement a custom network with. For instance the MCU is responsible for data processing and executing software-defined functionalities to implement flexible and reconfigurable upper-layer protocols, e.g., non-time critical MAC functionalities as well as network, transport, and application protocols, while the *ultrasonic interface* enables ultrasonic wireless connectivity for both the IoMT-mote and the IoMT-patch. It consists of a receiver (Rx) and a transmitter (Tx) chain. The Rx chain includes a low-noise amplifier (LNA) which also functions as a bandpass filter as well as an analog-to-digital converter (ADC) to amplify and sample received signals, while the Tx chain is comprised of a digital-to-analog converter (DAC) and a power amplifier (PA) to create and amplify an analog signal from the digital waveform, generated by the core unit, prior to transmission.

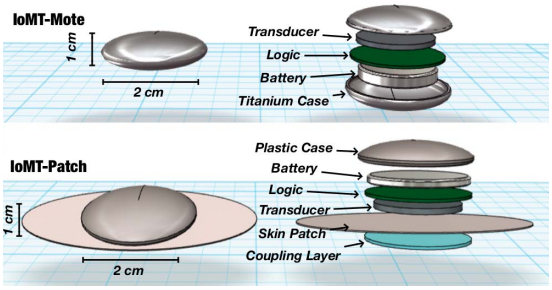

Fig. 3. Mock-ups of the IoMT-mote (top) and IoMT-patch (bottom).

The hardware architecture also contains an *RF interface* with an antenna to enable in-air RF wireless functionalities as well as a *Plug-n-Sense (PnS) module* comprised of standard interfaces that allow the IoMT-mote to network different sensors, e.g., pressure and glucose sensors, to the ultrasonic IoMT platform according to case specific application and therapy requirements. The PnS module offers analog, SPI, I2C and additional reconfigurable digital interfaces to support the integration of a multitude of peripheral sensors. The *power unit* includes a battery for powering the device and a voltage regulation system that filters out noise, EMI spikes, and functions as a watchdog system to alert the core unit of undervoltage or overvoltage conditions so the core unit can store critical data then proceed to safely shut down appropriate components by running the system in safe mode until a technician wirelessly diagnoses and potentially repairs the device either by removal or wirelessly updating device firmware to work around system failures.

Figure 3 shows the mock-ups of the IoMT-mote and the IoMT-patch including the logic, the battery, the ultrasonic transducer, and casing with desired target dimensions. The IoMT-mote will be enclosed in a titanium biocompatible casing, while the IoMT-patch will be enclosed in a plastic casing and attached to a disposable adhesive patch.

B. Software Architecture

The IoMT platform includes a unique MCU and FPGA software-defined architecture designed to network IoMT-devices and encloses a set of PHY, data link and network layer functionalities that can flexibly adapt to application and system requirements. The IoMT software framework offers real-time reconfigurability at the application layer, enabling application-specific data processing so that implants can be updated for new treatments wirelessly and don't need to be removed and re-implanted just to change the modality or application of a patients treatment. In particular, processed sensor data obtained from applications running in the nodes are decomposed into *primitive building blocks* that can rearranged to create new sensing applications that meet application specific requirements.

Figure 4 presents a high level description of the IoMT software architecture. The FPGA implements the PHY layer communication functionalities, as well as the interfaces, e.g., SPI and I2C, to connect the FPGA with the MCU and the peripherals (DAC, ADC). The MCU software architecture is based on a real-time operating system (RTOS) and executes

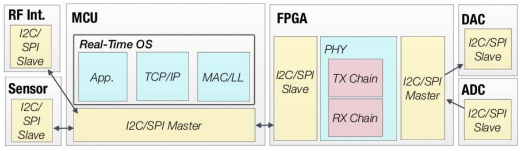

Fig. 4. IoMT platform software architecture.

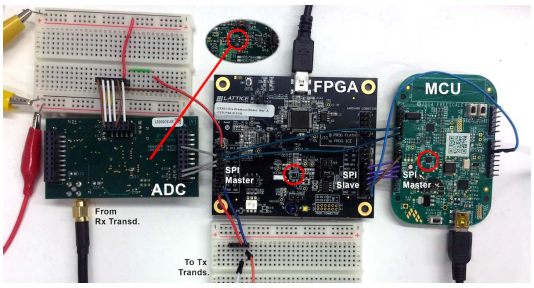

Fig. 5. Ultrasonic alpha-prototype node.

the upper layer communication protocols, e.g., link layer (LL), MAC, Network and Application layers. The MCU software also defines SPI and I2C interfaces to enable data exchange between the MCU and its peripherals (FPGA, RF interface and sensors).

III. IoMT-Mote Prototype

We now present the design of the IoMT-mote prototype based on the IoMT platform architecture discussed in Section II. The IoMT-mote is the first miniaturized software-defined implantable device with ultrasonic communication and networking capabilities.

A. Hardware Implementation

Figure 5 shows the hardware implementation in its alpha-prototype stage, i.e., using evaluation modules for each IoMT subsystem prior to committing to a single integrated design. The red circles in the figure indicate the component of interest in each development board i.e., the ADC, the MCU and the FPGA. In the alpha-prototype the ADC receives data directly from an Rx ultrasonic transducer conversely, the FPGA outputs digital waveforms to Tx ultrasonic transducers. The FPGA is connected to the MCU and ADC evaluation boards through SPI interfaces as slave and master, respectively.

1) Core Unit: The *core unit* of the node includes (i) a $2x2\ mm$ low-power field programmable gate array (FPGA) and (ii) a microcontroller unit (MCU).

FPGA: We use the Lattice Semiconductor iCE40 Ultra, which currently offers the greatest logic density for its form factor (4k look-up tables in a $2.08 \times 2.08\,\mathrm{mm}$ package). It is highly power effective with a very low static current drain of ($71\ \mu A$) which is ideal for maximizing implant battery life. it can also be reconfigured in real time wirelessly by processing received data and running it through its hardened internal SPI bus. The internal FPGA oscillator drifts substantially over time and so a high precision external oscillator is used to ensure the clocks are synchronized between multiple nodes. Additionally the external oscillator is more robust to changes in ambient temperature and is more likely to maintain a similar clock rate implanted, as when tested at room temperature. We chose to use a 12 MHz crystal to support running the

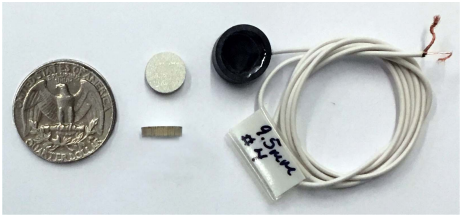

Fig. 6. Ultrasonic transducers with and without casing.

nodes between $700\,\mathrm{kHz}$ - $5\,\mathrm{MHz}$ (thus enabling the platform to be reconfigured for different transducers and enabling extra flexibility in the TX software design) and to conserve on power compared to using a $48\,\mathrm{MHz}$ crystal.

MCU: The MCU is responsible for implementing upper-layer protocols and coordinating the Tx/Rx operation of its respective node. Since every node is a transceiver where the Tx/Rx antennae chains share a transducer, switching between modes is handled by the MCU based on hardware interrupts generated by incoming messages. We use the Freescale KL03, a ultra-low-power ARM Cortex-M0+ MCU due to the relatively dense suite of embedded peripherals the Kl03 offers in its highly constrained form factor ($1.6 \times 2.0\,\mathrm{mm}$) such as a 12-bit ADC, used to interface the MCU with external analog sensors, as well as three communication interfaces make the chip a low power solution that significantly increases the versatility and modularity of the IoMT platform.

2) Ultrasonic Interface: The *ultrasonic interface* enables ultrasonic wireless connectivity through the use of data converters, low-noise amplifiers (LNA), and custom ultrasonic transducers.

Ultrasonic transducers: The IoMT-mote prototype uses a custom-made, miniaturized *ultrasonic transducer* to generate and receive ultrasonic waves [23].

The transducer is built around a $9.5\,\mathrm{mm}$ diameter thin-disk piezoelectric element that operates around $700\,\mathrm{kHz}$, and is custom fabricated to be highly damped which enables it to support a relatively large bandwidth of $200\mathrm{kHz}$ with respect to the center frequency. The $700\,\mathrm{kHz}$ central frequency offers an excellent tradeoff between in-vivo signal attenuation (which increases with frequency),transducer thickness (element size decreases with frequency), available bandwidth (increases with frequency), and radiation directivity (increases with frequency) [3]. A diameter of $9.5\,\mathrm{mm}$ is a good compromise between size, conversion loss (increases with smaller disks), and directionality (decreases with smaller disks).

For prototyping, the disk is epoxied in a waterproof casing, alongside a coupling layer, electrodes, and a micro-coaxial cable. The final IoMT-mote will embed the raw piezoelectric disk, logic and battery in a titanium casing. Figure 6 presents the ultrasonic transducer in its waterproof casing and the unshielded piezoelectric disk next to a quarter to illustrate the size of transciever element.

ADC: The ultrasonic interface operates at $700\mathrm{kHz}$ central frequency, therefore, the ADC sampling frequency for the IoMT platform is currently set to $2\,\mathrm{MHz}$. The ADC supports up to a $3\,\mathrm{MHz}$ sampling rate and was intentionally chosen to have a common hardware footprint to support our modular design since it can be easily swapped out for ADCs that support higher sampling rates in application specific scenarios where transducers that support greater frequencies are desired. The current platform design uses a small low-power serial ADC, i.e., TI ADS7883, and samples the received signal at $2\,\mathrm{MHz}$. The ADS7883 is connected as a slave peripheral to the FPGA through a SPI interface and is clocked by the FPGA since the FPGA is configured as SPI master. The SPI connection requires only three pins to output data to the FPGA, i.e., clock, enable, and data out. Compared to parallel ADCs where samples are loaded in parallel, to minimize delays for applications such as mass storage and high frequency DSP, serial ADCs can reduce design complexity and improve the form factor of a product. Parallel ADCs require at least an equivalent amount of pins as the ADC resolution i.e, 8 pins for an 8 bit ADC, 16 pins for a 16 bit, serial ADCs have fewer pins and so are often found in smaller packages than their parallel counterparts. Using a serial ADC doesn't impose its typical design restrictions on our platform since the platform operates at significantly lower frequencies than RF and so using a serial ADC ultimately decreases the size of our platform and uses less of the limited real estate of the FPGA i.e, only uses 3 of the 36 total pins, allowing the FPGA to interface with a greater number of peripheral devices without sacrificing receiver performance since the slower sampling rate reconstructs the transducer frequency range without aliasing.

The serial ADC has a 12-bit sample resolution and for design simplicity we add a 4-bit padding so the serial interface can operate at a standard 16 bit word size. To reliably run at a sample rate of $2\,\mathrm{MHz}$ the SPI link must operate at $32\,\mathrm{MHz}$, as clocked by the SPI Master.

To avoid an unnecessarily high dynamic power consumption, which is proportional to the circuit clock frequency, and to avoid extra static power consumption created by running an external crystal of $32\,\mathrm{MHz}$ or greater we generate a ultra-low power PLL using the iCE40 Ultra FPGA, to internally synthesize a $32\,\mathrm{MHz}$ clock from the external $12\,\mathrm{MHz}$ crystal to drive the SPI Master block.

LNA: To amplify and filter unwanted frequencies out of our received signal we implement a signal conditioning circuit using a low-noise, low power operational amplifier (TI OPA835). The circuit has tunable gain and adds a DC offset of $1.6\,\mathrm{V}$ to the incoming signal so that a single supply $3.3\,\mathrm{V}$ ADC can be used for the platform. This eliminates extra clutter in the alpha-prototype phase and the need for additional circuitry in the beta prototype phase needed to generate a dual supply from a single power source. To increase the receiver sensitivity and therefore operate at lower Tx powers, we use a low-noise, variable gain amplifier (VGA), TI AD8338, prior to the signal conditioning circuit. The AD8338 offers low current consumption, i.e., $3\mathrm{mA}$, and a voltage controlled gain between $0 - 80\,\mathrm{dB}$. By reducing the Tx power we save energy at the transmitter, but we increase power consumption at the receiver to power the preamplifier. Therefore, the use of the VGA is application dependent.

Tx chain: Due to the impulse based transmission scheme implemented in this prototype (see Section III-C1), the system transmits square pulses generated by the FPGA, with no need for analog conversion. The digital pulses are fed into

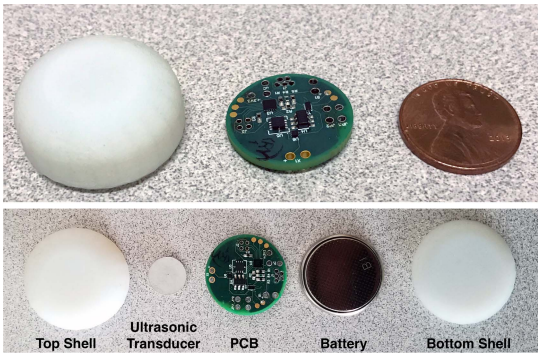

Fig. 7. IoMT-Mote prototype enclosed in a plastic shell and the custom PCB compared to a penny (top). Breakdown of the IoMT-Mote prototype including plastic shell, ultrasonic transducer, battery and custom PCB (bottom).

the transducer which filters out the out-of-band frequency components, and transforms *the square wave into a narrow-band pulse centered at* 700 kHz. Removing the need for an electronic DAC reduces the design size, energy consumption, complexity, and cost of the device. Additionally the lack of a DAC allows for an overall more cost effective system for intra-body beamforming using the IoMT-motes.

3) Power Unit: For the alpha-prototype the *power unit* consists of a commercial power supply with adjustable voltage levels to facilitate prototyping in addition to USB power supplies for both the MCU and FPGA. In the final beta-prototype, the power unit will be comprised of a small implantable-grade 3.3V battery with a dual-channel low-dropout (LDO) regulator which will provide separate voltage rails for the FPGA Core and the rest of the system i.e, the MCU, FPGA IO bank and peripheral devices.

B. Integration and Miniaturization

Figure 7 (top) shows the hardware implementation of the IoMT-mote beta-prototype enclosed in a biocompatible plastic (PEEK) shell, together with the custom IoMT platform printed circuit board (PCB) compared to a penny. PEEK is a cheaper but less durable biocompatible material used as a substitute for titanium in this stage of prototyping. The shell encases the ultrasonic transducer, the battery, and the custom PCB as shown in Fig. 7 (bottom).

The custom PCB serves as the miniaturized integrated version of all the electronics shown in the alpha prototype. The board was designed to match the diameter of the smallest commercially available coin cell battery (1cm). Since currently the widest component on our platform is the power source, it dictates the diameter of the case. Imposing this size constraint on the platform was integral to making the IoMT-mote a feasible solution for applications such as multi-chamber heart pacing and neurostimulation. The first iteration of the miniaturized design has a PCB diameter of 2 cm and 4 mm maximum thickness (with the components placed on it).

To scale down the system this drastically, special attention was given to part selection and interfaces used on the MCU and FPGA. To match the 3.3V of the coin cell battery, all parts were selected that run on a 3.3 V supply. Modern implants tend to use digital components that support supply voltages ranging from 3.3 V to 5 V. However, due to the growth of the IoT space and ultra-low power systems, most modern sensors and digital components offer a 3.3 V alternative and so this module was designed to support peripherals with voltages of 3.3 V and under.

To adhere to this design decision, a single-supply ADC was used. Therefore, incoming signals needed to be DC level shifted by half of the ADC supply voltage (1.65 V) to use the full resolution of the ADC and potentially support sampling an incoming signal with a peak to peak voltage that matches the maximum supported voltage of the IoMT platform (3.3 V). Dual supply ADCs exist in the same package size as the current ADC (which would have cleared up significant real estate on the PCB caused by generating and routing a DC bias on an already densely packed PCB). Although using a dual supply ADC would have resulted in a smaller PCB, ultimately the increased power consumption of generating an inverted voltage rail would significantly impact battery life and ultimately result in a larger IoMT module to accommodate a larger battery. To additionally save space, instead of shifting the signal by the the ideal 1.65 V we simply used the FPGA internal reference of 1.2 V instead. This reduces the incoming voltage range that the ADC can handle without sampling errors by .45 V. However, it also removes the need to include a voltage divider or step down regulation circuitry to generate 1.65 V from 3.3 V. Through testing the alpha prototype we had additionally found that the incoming signal is typically an order of magnitude smaller than its transmitted peak to peak 3.3 V due to attenuation introduced by the transducer and channel losses so this design decision does not impact our overall performance.

The MCU and FPGA support both SPI and I2C serial and programming protocols however to miniaturize the boards I2C was sparingly used to interface devices on the module. I2C components require two fewer pins and subsequently require routing two less lines than their SPI counterparts. However, each of these lines requires a pull up resistor. The smallest available commercial resistors still must be placed on the top and bottom layers of the PCB. It became apparent that the PCB for the module could not properly route the densely packed pins on the MCU and FPGA out to power sources and peripheral components with a two layer design. A typical two layer PCB design is where conductive material is routed on the top and bottom of a substrate, on the same surfaces components are placed on. To enable designs with higher densities typically PCBs can range from 4 layers upwards at the tradeoff of cost since each layer adds complexities to the fabrication process. Since the PCB would already need at least a 4 layer design to route the dense core unit components, it ultimately used up less real estate on the IoMT PCB to route several SPI interfaces through the internal layers of the PCB. The SPI interfaces facilitate communication between the MCU, FPGA, ADC, and up to an additional six peripheral components (shared between the MCU and FPGA).

To route every pin of the 36 pin FPGA and 20 pin MCU, which respectively have a.35 mm and.4 mm pitch between pins, standards set by the Association Connecting Electronics Industries (IPC) suggest implementing a six layer board, a dense four layer solution, or a less dense four layer solution

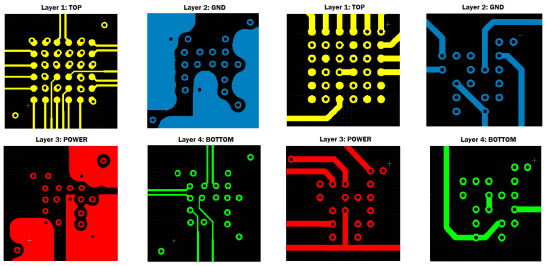

Fig. 8. IPC recommended 4-layer routing shown on left compared to IoMT 4-layer routing with thicker traces (greater safety margin).

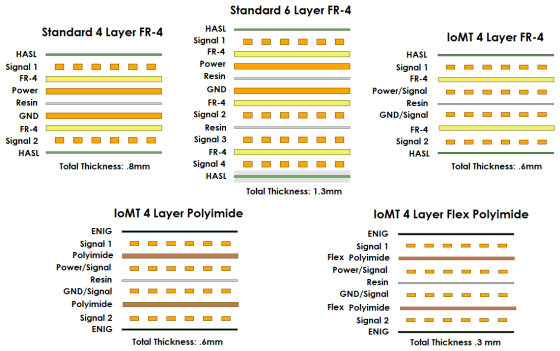

Fig. 9. Standard FR-4 vs. IoMT prototype, standard and flex PCB Layer Stackups.

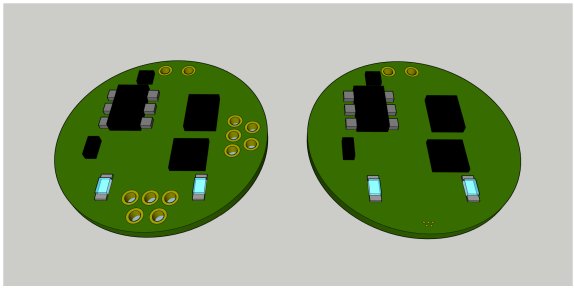

Fig. 10. (Left) Current custom programming ports (Right) future denser programmer footprint, identical on top and bottom of PCB.

with buried vias and a maximum conductive material width of .063 mm. The routing of the IoMT PCB is compared to the recommended dense four layer solution in Fig. 9 (left).

The routing of the IoMT module uses wider traces (.15 mm) than the standard trace routing width (.063 mm) to improve the overall cost and safety margin of the device. Although the system draws limited power during transmission and the voltage regulator can dissipate significant power and alert the rest of the system to shut down and avoid damage, a greater trace width reduces chances of the trace overheating and melting. Although our PCB use conductive traces thinner than the typical 70 um to allow for easier migration to flexible PCB materials in the future, the routing pattern used on the IoMT module doubles the safety margin from the recommended solution.

For this initial design standard FR-4 (class four flame retardant) double sided woven fiberglass was used as the laminate that separates copper layers. Similarly for the current iteration of the PCB Hot Air Solder Leveling (HASL) was used as the surface finish for electrical component placing. HASL was chosen for its cost effectiveness at this stage of design. The process immerses PCBs in a tin lead alloy and then uses hot air to blow off excess solder (the solder present on the laminate and not exposed copper). Due to this process it typically isn't recommended for high density components since the surface tension of solder between fine pitch pads may not hold components in place, as they drift after immersion during cooling. This is predominantly the case in components with rectangular pads with high density in one dimension (2 columns on the sides of an IC with densely placed pads) however due to the fact that the high density components of the IoMT platform are packaged in a ball grid array, the solder balls offer a great enough surface tension to adhere to the pads during the entire process regardless of their fine pitch. HASL boards currently meet FDA implant requirements as long as the PCB is encased within a hermetically sealed capsule, however due to increasingly strict RoHS standards it is highly appealing to use a more expensive leadless finish for the final product. Electroless Nickel Immersion Gold (ENIG) is a leadless finish that uses nickel as the PCB solder mask and immerses the nickel in gold to prevent oxidation on the edges of each pad that remain exposed after component placement. The process is very difficult to rework and so it is an ideal surface finish to use on a final product.

Finally, to miniaturize the entire PCB, a custom programming port needed to be made to boot firmware onto the MCU and FPGA using a PC. Typically a micro-USB connector (7.5mm $\times$ 5mm $\times$ 2mm) is used as a low profile port but due to the minuscule nature of the IoMT module this form factor is highly intrusive. The MCU is programmed using JTAG which is an industry standard for booting firmware onto ARM core systems, the FPGA uses an SPI interface with an additional line to reset the FPGA prior to programming it. Both systems require a 5 wire programmer and so the same custom port can be used for both. JTAG programmers range from 10-20 pins and can vary in size from 32.6 mm $\times$ 3 mm to 57.6 mm $\times$ by 3 mm respectively, both of which are far too large for these PCBs. These sizes tend to be because of redundant ground pins to sink extra current in higher power embedded applications such as automotive PCBs and due to extra pins to trigger interrupts to debug the chips in real time. Real time debugging was sacrificed to remove extra pins from our header and ultimately create a smaller form factor. A mini JTAG connector exists for embedded applications with a 1.27 mm pitch however for 10 pin connections this pitch still results in a 6.35 mm $\times$ 2.54 area. For this first iteration we created a smaller 5 pin connector as shown in figure 10 (left) with a similar pitch as the JTAG mini and connect it to a PC using custom cabling. This footprint takes up the smaller area of 3.8 $\times$ 2.4 mm.

This maintains the same horizontal pitch of the JTAG mini connector so if there is a need to solder pins for a research trial where the unit will be mounted at strange angles and needs to be frequently reprogrammed, this option will result in a firm secure connection. For finalized iterations of the board to further condense the design we intend to place pads on the top and bottom of the PCB (instead of using holes) as shown in figure 10 (right) and create a programming fixture. This fixture will have a similar radius as the PCBs and cut in a shape where the PCB can only be placed in the proper mechanical orientation for programming. A series of densely

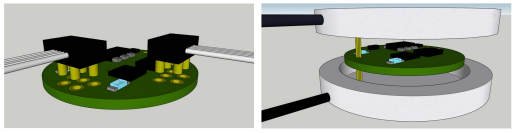

Fig. 11. (Left) Current programming connectors for core module (Right) programming fixture for denser programming port.

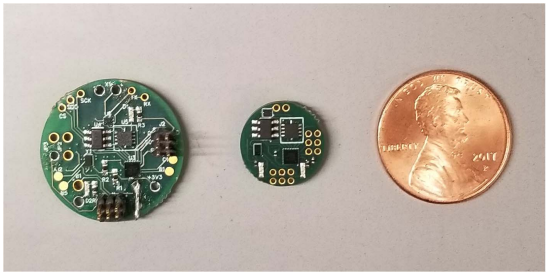

Fig. 12. (From left to right) 2 cm Diameter PCB, 1 cm Diameter PCB, US Penny.

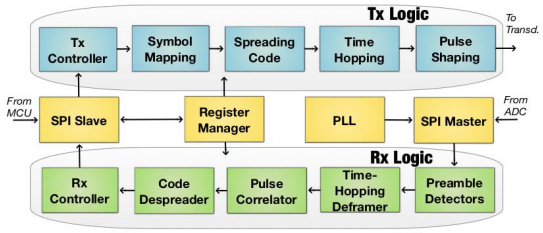

Fig. 13. Block scheme of the FPGA design.

TABLE I
RESOURCE OCCUPATION OF THE LOGIC IMPLEMENTED ON THE FPGA

Module	PLBs	%
SPI Slave	61/440	14%
Tx Logic	44/440	10%
SPI Master	82/440	18%
Rx Logic	252/440	57%
Tot	439/440	99%

packed spring loaded pins (.4 mm pitch) will be embedded onto two custom connectors in the top and bottom fixture covers respectively. These will make contact with the PCB and begin to compress as the fixture is closed. When the fixture is closed the spring will fully compress and lock giving the user feedback that a solid connection has been made and the user can proceed to program all systems on the PCB. This final design will reduce the area used for both connectors from 7.6 mm $\times$ 4.2 mm to 1.75mm $\times$.7 mm.

The most recent implementation of the PCB has been miniaturized to a $1\,\mathrm{cm}$ diameter as shown in figure 12. This was achieved by removing the test points and breakouts to assess the electrical performance and functionality of the $2\,\mathrm{cm}$ diameter device. Further chemical and mechanical testing must be performed to determine changes to the make of the design for future iterations. However, the electrical performance of the device has been determined to be up to par. Future improvements with respect to electrical performance will come when newer top of the line FPGAs, other signal chain ICs as well as powering technologies become readily available.

C. Software Implementation

As discussed in Section II-B, the FPGA primarily implements the PHY layer, while the MCU runs upper-layer protocols, as well as application layer routines, such as sensor data acquisition and reconfigurable data processing operations. In the current prototype, we implemented the state-of-the-art UsWB transmission scheme and protocol [11]. UsWB is an impulse-based ultrasonic transmission and multiple access scheme. It uses short information-bearing carrierless ultrasonic pulses, following a pseudo-random adaptive time-hopping pattern with a superimposed adaptive spreading code. An Ultra-Wideband communication system was used compared to other modern networking solutions due to the robustness of impulse based communications to multipath which is a natural consequence of communication in-vivo due to the distribution of organic channel types within a user.

1) FPGA Design: The FPGA top-level module instantiates Tx and Rx chain blocks which implement the UsWB transmitter and receiver, respectively, as well as the SPI, PLL and register manager modules.

The UsWB PHY layer assumes time is divided in slots of duration T_c, with slots organized in frames of duration $T_f = NT_c$, where N is the number of slots per frame. Each user transmits one pulse per frame in a slot determined by a pseudo-random time-hopping sequence. Bits are mapped into a pseudo-orthogonal code of variable length, M, and code chips are mapped into pulses through pulse position modulation (PPM). The pair code and frame length (M, N) can be adjusted to satisfy reliability constraints.

Figure 13 shows a block diagram of the FPGA implementation. In Table I, we report the FPGA resource utilization for the modules discussed above in terms of programmable logic blocks (PLBs), and resource percentage.

After several optimization cycles, the final implementation occupies around 99% of the available logic cells. As expected, the receiver logic occupies more than 50% of the available resources on the FPGA, most of which are dedicated to the synchronization process. This was achieved after several optimization cycles and through implementing creative solutions to reduce processing complexity and therefore resource usage. To reduce receiver complexity, the correlator templates used for synchronization are square waveforms of amplitude '-1' and '1', implemented using 2-bit coefficients instead of 12-bit coefficients (same size as the input).

Tx chain design: The Tx chain receives a stream of bytes coming from the MCU through the SPI Slave module, and outputs the PHY digital waveforms which represent the modulated bits. These digital waveforms are then radiated by the transducer as an ultrasonic signal in the communication channel. The Tx controller receives data from the MCU and coordinates the PHY layer operations of the Tx chain. In the Symbol Mapping block the information bits are mapped into $\{-1,1\}$ binary symbols. These binary symbols are then spread in chips by the Spreading Code module. For each symbol, this block outputs M chips in $\{-1,1\}$. Chips are then forwarded to the Time-Hopping module that spreads them in time according to the selected time-hopping pattern, generated using a Linear Feedback Shifter Register (LFSR) module. Finally, the Pulse Shaping module maps the incoming chips to position-modulated pulses. The output is a train of position-modulated pulses following a predefined

time-hopping pattern. Each pulse consists of three cycles of a 700 kHz square wave. A longer electrical excitation provides a higher output pressure because of the resonant operation of the transducers. However, longer pulses lower the data rate. We found that three cycles optimize a tradeoff between data rate and ultrasound generation efficiency. Finally, packets are preceded by two preambles: (i) a 64 cycle square wave that is used at the receiver for packet detection and (ii) a train of three pulses properly spaced in time used at the receiver for achieving time-hopping synchronization.

Rx chain design: The custom receiver chain implements the receiver UsWB PHY layer functionalities. The received ultrasonic signal is converted to an electrical signal by the Rx transducer. The signal is amplified by the LNA, and converted to an analog signal by the ADC. Then, the custom receiver chain in the FPGA processes the digital waveform acquired through an SPI master interface. Finally, the receiver chain outputs a binary stream representing the received decoded data, which is delivered to the MCU through the SPI Slave interface.

The *Rx controller* triggers the start of the PHY layer processing when synchronization is achieved, and makes the decision on whether the received bits are an incoming signal or high energy noise based on the output of the PHY layer processing. The *preamble detectors* consist of a *packet* detector for coarse synchronization and a *time-hopping* synchronization block for fine synchronization.

After synchronization is achieved, the *time-hopping deframer*, the *code despreader*, and *pulse correlator* invert the operation done at the transmitter, and the Rx controller makes a decision on the state of the received bits.

Register manager: The *register manager* is responsible for storing and routing the design configuration parameters (written by the LL module running on the MCU) on a pool of setting registers in the FPGA. Through these setting registers, one can reconfigure key parameters of the PHY layer transmission scheme,

and enable real-time reconfiguration of the transceiver. The communication system is designed to allow real-time reconfiguration of several parameters, e.g., spreading code and time-hopping frame length, among others.

SPI module: The iCE40 Ultra has two hardened, i.e., already fabricated in the FPGA, SPI IP buses. The SPI interfaces enable communication with the external peripherals and the MCU. We configured one hardened SPI bus as slave, and one as master. The *SPI Slave* block is driven by the SPI Master module of the MCU. This SPI link is used to exchange data between the MCU and the FPGA, such as data queued for transmission, received data or PHY configuration parameters. The data rate on this link is $1\mathrm{Mbit/s}$, which is greater than the PHY layer data rate, which results in a perpetually backlogged PHY Tx chain. The *SPI Master* handles the communication with the ADC. Specifically, it triggers the sampling operations on the ADC, and reads back the sampled digital waveform.

PLL module: The iCE40 Ultra includes an ultra-low power Phase Locked Loop (PLL) that provides a variety of user-synthesizable clock frequencies. We use the PLL to internally synthesize a 32 MHz clock signal to individually drive the SPI Master as discussed in Section III-A2.

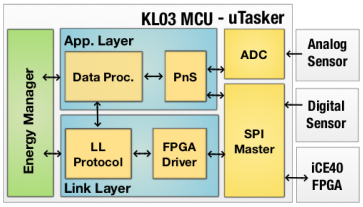

Fig. 14. KL03 MCU Firmware Architecture.

The PLL is shut down during transmission when the ADC is unused, to minimize energy consumption.

2) Core MCU Firmware: The MCU software architecture is implemented on a RTOS vs. a sequential loop so that application specific task scheduling can be easily integrated without interfering the timing of the core tasks, additionally since RTOS schedules tasks using an interrupt driven paradigm at the sacrifice of RAM, it frees up processing resources since the MCU doesn't need to consistently poll events to drive tasks and can instead wait for interrupts.

We selected μTasker as our RTOS, because it runs in resource constrained environments such as the KL03 MCU and provides a reconfigurable task scheduler based on an infrastructure which minimizes power consumption. This eliminates the need to design a scheduler from the ground up to obtain similarly low power consumption. Additionally μTasker abstracts most of its task processes from the hardware, and thus is not a hardware specific RTOS so the firmware can be easily migrated, with little or no change required, to cutting edge MCU solutions as they're released for future iterations of the IoMT platform.

Figure 14 shows the firmware architecture implemented on the KL03 MCU. In a typical application, the application layer would trigger a reading through the PnS interface from a digital or analog sensor. The sensor reading is then processed by the application specific data processing firmware and passed to the LL protocol module for transmission through the SPI Master interface.

Link layer: The UsWB LL protocol manages the data transmission over the UsWB PHY layer interface. The connection is established through an advertising process initiated by a slave node, which periodically transmits advertising packets.

A master node scans the channel, and upon receiving an advertising packet, connects to the slave. Both nodes agree on a connection interval for exchanging data periodically.

The connection uses a stop-and-wait flow control mechanism based on cumulative acknowledgements.

The link layer also implements driver functionalities on the FPGA for the UsWB transceiver such as initializing the transceiver, configuring PHY layer parameters and triggering transmit and receive operations.

Application layer: The application layer implements the PnS module to connect the IoMT-mote with sensors. The PnS module consists of a digital I2C/SPI Master interface that connects to digital sensors, and an analog interface based on the internal MCU ADC that reads analog sensor outputs.

Sensor data is processed by the reconfigurable data processing module implemented on the MCU and encrypted

end-to-end using a streamlined implementation of the Advanced Encryption Standard (AES) based on a 128bit key exchanged during paring between two devices.

The data processing module is based on the idea of decomposing data processing applications running in the IoMT-mote into primitive blocks, and offering real-time reconfigurability at the application layer. The data processing consists of a sequence of basic operations that extract desired medical parameters from sensor data. Real-time modular reconfiguration enables new processing functions to be wirelessly transmitted and installed on the IoMT-platform at runtime, so applications can be updated on the same motes without being removed. This maximizes code reusability and allows for new features to be extracted from the same sensors based on which medical parameters users desire to acquire.

Based on this modular approach, applications can be represented by chains of binary sequences, i.e., keys. Each primitive function is mapped to a binary key. A concatenation of keys represents a concatenation of operations, and creates an application. The IoMT-mote feeds these keys into an finite-state-machine (FSM) where each state represents a primitive block function. By parsing consecutive keys, the FSM transitions from state to state to process inputs and produce outputs.

Energy manager: We leverage μTasker primitives to access the KL03 power states, and fine tuned the μTasker framework to minimize the system energy consumption. Specifically, the energy management module is able to (i) adjust the core clock frequency at runtime according to the processing power required, (ii) select low-power mode according to application requirements at runtime, and (iii) automatically wake up different cores of the MCU based on which system functions need to be active during the IoMT-motes current state. The MCU current consumption can go from $1.8\,\mathrm{mA}$ when all systems are active down to $0.6\,\mu\mathrm{A}$ in its most conservative state. The MCU supports other intermediate states that trade current consumption for device response time and input-output interface functionality.

IV. IoMT-Patch Prototype

Here we describe the IoMT-patch implementation, focusing only on the modules that differ from the IoMT-mote.

A. MCU and RF Interface

The IoMT-patch prototype replaces the KL03 MCU with a TI CC2650 BLE wireless MCU that coordinates transmissions over the ultrasonic interface, as discussed in Section III-A1; and transmissions over the RF interface to connect the system with the access point. Specifically, we consider two different access point solutions: (i) a multi-platform smartphone app that communicates via BLE with the IoMT-patch and gives the user direct access to implanted sensor data, and (ii) a 6LOWPAN edge router that enables IPv6 connectivity and direct data delivery to the cloud. The TI CC2650 is the smallest, lowest-power $2.4\,\mathrm{GHz}$ wireless MCU currently available on the market, and is designed to operate in energy constrained systems powered by small coin cell batteries. The CC2650 device contains an ARM Cortex-M3 that implements upper layers of the BLE protocol stack and user defined functionalities. A secondary low-power ARM Cortex-M0 processor is in charge of lower-level BLE functionalities.

B. Software Implementation

We implemented two different firmwares for the TI CC2650 wireless MCU: (i) a BLE-enabled firmware based on TI-RTOS for connecting the IoMT patch to a smartphone, and (ii) an IPv6-enabled firmware based on Contiki that offers 6LOWPAN capabilities for IPv6 support [24]. Finally, we implemented a BLE-enabled access point through a smartphone app that delivers sensor data to the user and a 6LOWPAN edge router that collects and publishes data on the cloud.

1) BLE-Enabled Implementation: The BLE-enabled firmware establishes the connection between the IoMT-patch and the BLE access point. The connection is established through an advertising process initiated by the IoMT-patch, which transmits advertising packets every $300\,\mathrm{ms}$ on three control channels. The advertising time was set to $300\,\mathrm{ms}$ as it reduces energy consumption while minimally impacting user experience. When the access point receives an advertising packet, it sends a connection request to the slave and they start exchanging data every $2\,\mathrm{s}$.

We implemented the system pipeline for a BLE Heart Rate Profile. In the background, the IoMT-patch initializes the ultrasonic intra-body link to retrieve heart rate data from the IoMT-mote. Data security in the over-the-air link is established through the frequency-hopping scheme adopted by the BLE physical layer, as well as end-to-end encryption. We developed an Android and iOS smartphone app based on the Qt framework, which implements a simple BLE listener that scans for BLE devices, connects to the Heart Rate service running on the IoMT-patch, and delivers the heart rate to the user. Through the GUI the user can trigger a reading from the IoMT-patch and visualize the acquired data.

2) IPv6-Enabled Implementation: The IPv6-enabled firmware offers 6LoWPAN [24] encapsulation and header compression to support IPv6 over 802.15.4 wireless networks. It also allows the IoMT-patch to connect to the Internet using open standards. The IoMT-patch is configured as a MQ Telemetry Transport (MQTT) client that periodically reads data from the IoMT-mote through the ultrasonic interface, and publishes sensor readings to a MQTT server. An edge router, i.e., a gateway between the 6LoWPAN mesh and Internet, provides conversion between 6LoWPAN and IPv6 header. We implement the edge router using the 6LBR 6LoWPAN Border Router solution [25] running on a Raspberry Pi.

V. Performance Evaluation

This section presents the performance evaluation of the ultrasonic wireless interface implemented on the alpha-prototypes and miniaturized PCBs in terms of communication reliability and energy consumption. The ultrasonic wireless interface is responsible for enabling communications between the IoMT-patch and the IoMT-mote, and therefore its performance directly affects the lifetime of the implantable device, the most energy constrained device in the system. We also compare the performances between the ultrasonic wireless interface and an intra-body BLE link.

A. Hardware Current Consumption.

This section details the energy consumption of the IoMT-mote prototype.

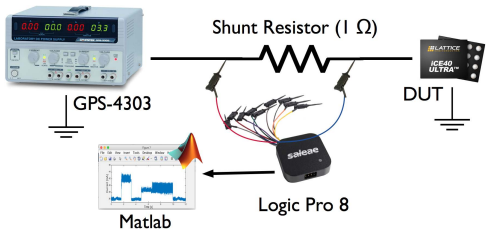

Fig. 15. Diagram of the current measurement setup (not in scale).

TABLE II
CURRENT AND POWER CONSUMPTION OF THE IOMT-MOTE

Component	Current [mA]		Power [mW]	
	Tx	Rx	Tx	Rx
MCU	1.8	1.8	6	6
FPGA	1.6	2.3	4	4.4
ADC	-	2	-	6.6
Preamp.	-	3	-	9
Tot.	3.4	9.1	10	26

We measured the current consumption of the IoMT-mote prototype using a custom current sensing system that uses a shunt resistor topology. A shunt resistor measures the voltage drop across a small resistor connected in series between the power supply and the load. The current drawn by the system is proportional to the voltage drop across the series resistor ($I = V/R$).

Figure 15 shows a diagram of the measurement setup. In these experiments, the voltage drop is measured using two analog inputs of the Saleae Logic Pro 8 logic analyzer to capture voltages at the two ends of a $1\,\Omega$ shunt resistor.

In Table II, we report the current and power consumption of the IoMT-mote. The IoMT-mote consumes 9.1 mA in Rx mode, and as low as 3.4mA in Tx mode. These results suggest that ultrasonic waves can be efficiently generated and received using low-energy and miniaturized components, which is a fundamental step in proving the proposed IoMT platform can be feasibly miniaturized.

Further optimization of the proposed hardware design could drastically reduce power consumption and substantially outperform RF-based devices. For example, active power consumption during receiving and transmitting can be reduced by replacing the FPGA with an application-specific integrated circuit (ASIC). While it is hard to estimate the energy gain from replacing an FPGA with an ASIC, studies suggest that the power reduction can be tenfold [26].

In Table III, we compare the IoMT-mote prototype current and power consumption with the consumption of the TI CC2650 BLE MCU and the Microsemi ZL70103 transceiver that operates in the MICS band. The MCU of the IoMT-mote prototype is set to 3 MHz, and we show the Rx current consumption with and without preamplifier. Finally, we assume the three devices operate at $5\,$dBm transmission power.

The IoMT prototype performs comparatively well against commercial wireless deviceswith respect to power consumption.

B. Propagation Loss.

The attenuation of $2.4\,$GHz RF waves in the ISM band and $700\,$kHz ultrasounds through porcine meat was measured at distances between 4-12 cm. This data is compared to simulated

TABLE III
CURRENT AND POWER CONSUMPTION OF THE IOMT-MOTE PROTOTYPE HARDWARE COMPARED TO THE TI CC2650 AND ZL70103 CONSUMPTION IN TX AND RX MODE

	I_{Tx} [mA]	I_{Rx} [mA]	P_{Tx} [mW]	P_{Rx} [mW]
IoMT-mote	3.4	6.1/9.1	10	17/26
CC2650	10.47	6.47	34.5	22.3
ZL70103	7.6	6.8	23.3	20.9

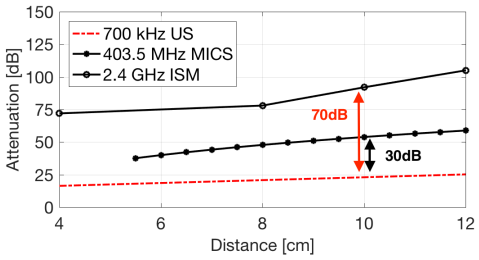

Fig. 16. Attenuation in porcine meat for RF $2.4\,$GHz ISM, RF $403.5\,$MHz MICS and for $700\,$kHz ultrasounds as a function of propagation distance.

attenuation of the 403.5 MHz MICS band in tissues [5] in Figure 16. The attenuation considers absorption by tissue, conversion losses and spread losses. Measurements for RF and ultrasound were obtained by transmitting signals between antennas and transducers respectively.

To avoid RF leakages the two CC2650 boards were enclosed in Faraday shielding bags which attenuate up to 82dB RF leakage and reduce the undesired effect of in-air RF propagation.

We confirm results reported in [3], [4] and observe that for $10\,$cm propagation distance, ultrasound attenuates $70\,$dB and $30\,$dB less than RF $2.4\,$GHz IMS and $403.5\,$MHz MICS, respectively.

C. Bit Error Rate Evaluation

The performance of the UsWB transmission scheme implementation on the IoMT-mote was assessed with respect to BER as a function of Tx power. We varied the Tx power from $5\,$dBm ($3\,$mW) to $-25\,$dBm ($3\,\mu$W) by connecting attenuators between the FPGA output and the transducer. The following results were obtained using ultrasonic phantoms that match the acoustic properties of human tissues. Specifically, an upper arm phantom that emulates muscle tissue containing veins and blood, as well as a thoracic phantom that emulates a thoracic spinal segment surrounded by muscle and skin [27]. Figure 17 shows the channel impulse response (CIR) of the two scenarios. The point at time zero indicates the instant of transmission, the blue points represent the time of arrival of the signal paths. In the upper arm phantom almost no multipath effect is observed, except for a secondary path caused by the reflection of the transmitted signal between the surfaces, which require exactly 3 times the propagation time to arrive at the receiver. Due to the soft/hard tissue interface in the thoracic phantom, the multipath effect is stronger.

l-pad attenuators were implemented using purely resistive components. Through attenuator tuning, the transmission power was varied from $5\,$dBm ($3\,$mW) to $-25\,$dBm ($3\,\mu$W). For each BER measurement up to 2500 packets of 48 bytes, i.e., approximately 768 kilobits, containing pseudorandom-generated raw data were transmitted. This

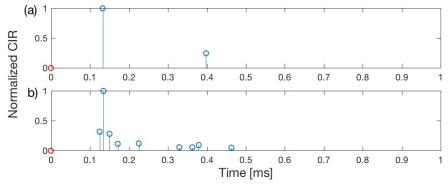

Fig. 17. Ultrasonic channel impulse response for (a) upper arm phantom and (b) thoracic phantom.

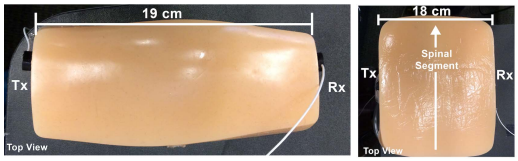

Fig. 18. Experiment setup for the upper arm phantom (left) and thoracic phantom (right).

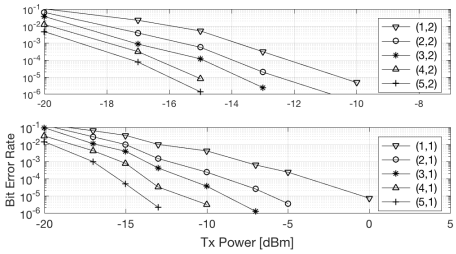

Fig. 19. BER for the no-amp scenario in the upper arm phantom for code length in {1, 5} and frame length 2 (top) and 1 (bottom).

allows us to detect a bit error rate resolution in the order of 10^{-6}.

Upper Arm Phantom: Two transducers were set up facing each other on opposite sides of the upper arm phantom (a distance of 19 cm), as shown in Fig. 18 (left).

Figure 19 shows the observed BER as a function of the Tx power, for code length varying in {1, 5} and frame length 1 (center) and 2 (top), when no preamplifier is used. As shown the spreading code scheme mitigates signal distortion, thus reducing channel errors.

The prototype achieves a data rate of $90\,\mathrm{kbit/s}$, with code length 1, frame length 2, i.e., pair (1,2), a 10^{-6} BER and an input power at the Tx transducer of about $-10\mathrm{dBm}$ ($0.1\mathrm{mW}$). A data rate up to about $180\,\mathrm{kbit/s}$ can be achieved (also with 10^{-6} BER) with pair (1,1) increasing the input power to $0\mathrm{dBm}$ ($1\ \mathrm{mW}$). Lower-power transmissions are also possible by compensating with longer spreading code. For example, in the current implementation, for a Tx power of $-15\,\mathrm{dBm}$ ($30\,\mu\mathrm{W}$), with a code length of 5 and frame length of 2, we obtain a data rate of $18\,\mathrm{kbit/s}$ and observe a BER lower than 10^{-6}.

Thoracic Phantom: The two transducers were placed facing each other, 18 cm apart, as shown in Fig. 18 (right). The thoracic phantom enables testing the communication performance through heterogeneous soft/hard tissues.

Figure 20 shows BER as a function of Tx power for frame length 2, code length 1 and 2, with different amplification gains at the receiver, compared to the 0-gain scenario when no preamplifier is used. Higher path loss and multipath effect caused by the soft-hard tissue interface, necessitate an increase of $15\mathrm{dBm}$ in Tx power to achieve the same BER performance of the upper arm scenario when no preamplifier is used. Introducing a gain at the receiver, increases receiver sensitivity

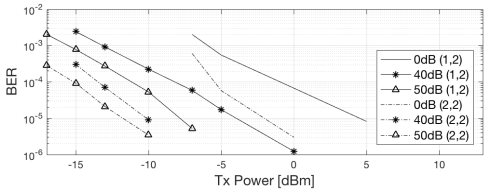

Fig. 20. BER in the thoracic phantom for code length in{1, 2}, frame length 2 and different amplification gains.

Fig. 21. Experiment setup for porcine meat scenario along 12 cm.

so the system can operate at lower Tx powers. By using $40\,\mathrm{dB}$ gain and $50\,\mathrm{dB}$ gain, we can get $90\,\mathrm{kbit/s}$ with 10^{-6} BER with $0\,\mathrm{dBm}$ and $-8\,\mathrm{dBm}$ Tx power, respectively. Because of the impedance mismatch between the transducer and the preamplifier, the Tx power does not decrease linearly with the receiver gain. Therefore, in future iterations of the IoMT platform an impedance matching circuit will be implemented to compensate for this loss.

Preamplifier use depends on the application scenario, it offers an increased Rx sensitivity for the IoMT mote at the cost of increased size and design complexity.

D. RF 2.4 GHz *Vs. Ultrasounds*

Packet Error Rate (PER): Here we compare the UsWB transmission scheme implemented in the IoMT-mote with a BLE PHY layer based on a $1\mathrm{Mbit/s}$ Gaussian frequency shift key (GFSK) implemented on the TI CC2650 in terms of PER through porcine meat. Porcine meat closely emulates human muscle tissues [28], [29], and provides us with a medium to evaluate ultrasonic and RF communications side by side with respect to reliability and energy consumption. To assess the feasibility of mote to mote communication the system was tested with two ultrasonic transducers facing each other $12\,\mathrm{cm}$ apart, as shown in Figure 21.

The PER is obtained as the ratio between the number of packets received with errors, and the total number of packets transmitted. Figure 22 (top) shows the IoMT-mote PER as a function of the Tx power, with code and frame length {1, 2}. For 10^{-6} BER, the prototype achieves $180\,\mathrm{kbit/s}$, with code and frame length 1, with a Tx power of about $-20\ \mathrm{dBm}$ ($10\,\mu\mathrm{W}$). By using frame length 2 to get rid of the ISI effect, we achieve $90\ \mathrm{kbit/s}$ data rate, with the same 10^{-6} BER, and Tx power of $-27\ \mathrm{dBm}$ ($2\ \mu\mathrm{W}$). Figure 22 (center) shows the PER performance of BLE in porcine meat as a function of the Tx power at different communication distances. We observe that at distances greater than $10\ \mathrm{cm}$, reliability drops dramatically. Specifically, at $12\,\mathrm{cm}$, the PER becomes as high as 80% with the maximum Tx power available, i.e., $5\ \mathrm{dBm}$, making communication almost unfeasible. At greater distances, the communication is completely disrupted. In Fig. 22 (bottom), PER performance of BLE over a $12\,\mathrm{cm}$ distance is compared with the IoMT-mote performance

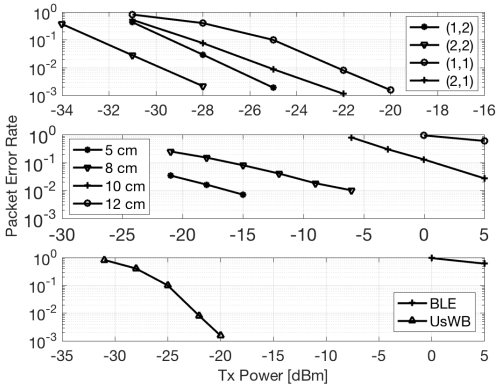

Fig. 22. IoMT-mote PER in porcine meat for code and frame length in {1,2} (top). BLE PER in porcine meat for different communication distances (center). BLE and IoMT-mote PER in 12 cm porcine meat (bottom).

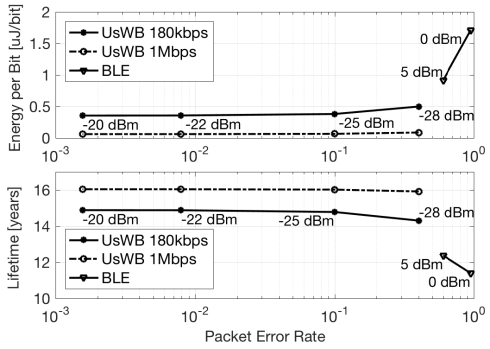

Fig. 23. IoMT-mote and BLE E_b (top) and lifetime (bottom) in porcine meat.

(code and frame length 1). BLE was found to require around 35 dBm higher Tx power to achieve the same reliability as the IoMT-mote. This gap can be further increased by implementing stronger synchronization and decoding operations at the PHY layer of the mote.

Energy per Bit and Device Lifetime: In a remote monitoring application in which 20 bytes of data are sent every minute between two devices the energy consumption of the IoMT-mote is compared to that of the CC2650 BLE devices, assuming current consumption values reported in Section V-A. We define energy per bit, E_b, as the ratio between the total energy spent by the two devices exchanging information data over the amount of successfully exchanged information data [J/bit]. Network lifetime is the minimum between the Tx and Rx devices battery lifetime [years]. For the BLE devices, the master node connects, followed by a connection event every 32 s, i.e., the maximum connection interval available. In this scenario, we measure E_b equal to $0.77\,\mu\mathrm{J/bit}$ for BLE against $0.37\,\mu\mathrm{J/bit}$ for the IoMT-mote.

We consider transmitting, receiving, processing and idle states only. The processing state occurs before and after a packet is transmitted and received, and we assume a consumption of 3 mA. We also assume $2\,\mu$A idle current consumption, and a 300 mAh battery. Under these conditions, the BLE network lifetime is 12.5, against 14.8 years achieved by the IoMT-mote.

In Fig. 23, we show the E_b (top) and the network lifetime (bottom) using the PER measurements discussed above. We observe that at over 12 cm the IoMT-mote outperforms

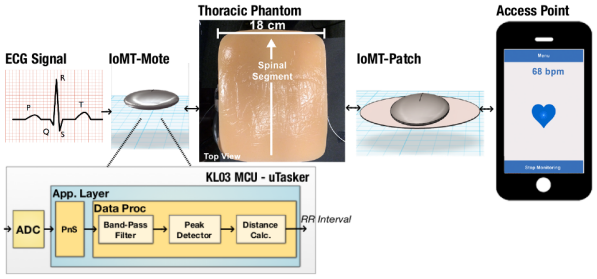

Fig. 24. Setup for the data processing performance evaluation.

BLE in terms of lifetime and E_b. In fact, the IoMT-mote can achieve much lower PER with lower Tx power, and therefore keep the E_b and network lifetime close to the ideal values of 14.8 years achieved when PER is ideally zero. On the other hand, BLE can only operate over 12 cm at the maximum Tx power, and still underperforms in terms of PER compared to the IoMT-mote. This further reduces the network lifetime and increases the E_b. Specifically, the IoMT-mote has a greater device lifetime of two years while achieving much higher reliability than BLE, i.e., about three orders of magnitude lower PER. In Fig. 23 we also show how the the IoMT-mote E_b and lifetime performance would scale if we increased the data rate from $180\,\mathrm{kbit/s}$ to $1\,\mathrm{Mbit/s}$ to match the BLE data rate. This can be achieved using wider-bandwidth ultrasonic transducers and/or using higher order modulation schemes. Results show how E_b can become as low as $0.07\,\mu\mathrm{J/bit}$ and network lifetime increase up to 16 years.

E. Data Processing

As a proof of concept, we implemented the IoMT-mote as a heart-rate monitor by leveraging the reconfigurable data processing module and we evaluated the processing accuracy in terms of displacement between the heart-rate reading and the expected heart-rate. Figure 24 shows the system setup for this experiment. The application measures the heart rate in beats-per-minute from a single-electrode ECG signal, based on the *RR interval* duration, i.e., distance between two consecutive R waveforms. The ECG signal is generated by a waveform generator and fed to the ADC of the IoMT-mote MCU.

The sampled ECG signal is then processed as shown in the simplified primitive block sequences. The smartphone app connects through BLE to the IoMT-patch, which retrieves the data from the IoMT-mote through the ultrasonic intra-body link. The thoracic phantom was used to emulate in-tissue ultrasonic propagation.

In these experiments, we vary the heart-rate by changing the output frequency of the wave generator. A processing accuracy of 98.7% was observed.

VI. Conclusions

The feasibility of the first hardware and software architecture of an intra-body IoMT platform with ultrasonic connectivity for intra-body communications was presented, and compared to state-of-the-art low-power RF-based wireless technology. It was demonstrated that ultrasonic waves can be efficiently generated and received with low-power, mm-sized components, and that ultrasonic communications require much

lower Tx power compared to BLE with equal reliability, thus leading to lower energy per bit cost and longer device lifetime. We also show experimentally that BLE links do not function at distances greater than $12\,\mathrm{cm}$, while ultrasonic links achieve a reliability of 10^{-6} up to $20\mathrm{cm}$ with less than $0\mathrm{dBm}$ Tx power. By using wider-band transducers and further optimizing the hardware consumption of the prototypes E_b can decrease to an order of magnitude less than BLE, and achieve even longer device lifetime.

REFERENCES

[1] G. E. Santagati and T. Melodia, "An implantable low-power ultrasonic platform for the Internet of medical things," in *Proc. IEEE Conf. Comput. Commun. (INFOCOM)*, Atlanta, GA, USA, May 2017, pp. 1–9.

[2] A. Al-Ahmad, K. A. Ellenbogen, A. Natale, and P. J. Wang, *Pacemakers and Implantable Cardioverter Defibrillators: An Expert's Manual*. Minneapolis, MN, USA: Cardiotext Publishing, 2010.

[3] L. Galluccio, T. Melodia, S. Palazzo, and G. E. Santagati, "Challenges and implications of using ultrasonic communications in intra-body area networks," in *Proc. IEEE Int. Conf. Wireless Demand Netw. Syst. (WONS)*, Courmayeur, Italy, Jan. 2012, pp. 182–189.

[4] D. Kurup, W. Joseph, G. Vermeeren, and L. Martens, "Path loss model for in-body communication in homogeneous human muscle tissue," *Electron. Lett.*, vol. 45, no. 9, pp. 453–454, Apr. 2009.

[5] K. Sayrafian-Pour, W.-B. Yang, J. Hagedorn, J. Terrill, and K. Yazdandoost, "A statistical path loss model for medical implant communication channels," in *Proc. IEEE 20th Int. Symp. Pers., Indoor Mobile Radio Commun.*, Sep. 2009, pp. 2995–2999.

[6] A. Y. Cheung and A. Neyzari, "Deep local hyperthermia for cancer therapy: External electromagnetic and ultrasound techniques," *Cancer Res.*, vol. 44, no. 9, pp. 4736s–4744s, Oct. 1984.

[7] U.S. Food Drug Administration (FDA). (2008). *Information for Manufacturers Seeking Marketing Clearance of Diagnostic Ultrasound Systems and Transducers*. [Online]. Available: http://goo.gl/ErZl9g

[8] Y. Davilis, A. Kalis, and A. Ifantis, "On the use of ultrasonic waves as a communications medium in biosensor networks," *IEEE Trans. Inf. Technol. Biomed.*, vol. 14, no. 3, pp. 650–656, May 2010.

[9] T. Hogg and R. A. Freitas, Jr., "Acoustic communication for medical nanorobots," *Nano Commun. Netw.*, vol. 3, no. 2, pp. 83–102, Feb. 2012.

[10] M. Peisino and P. Ryser, *Deeply Implanted Medical Device Based on a Novel Ultrasonic Telemetry Technology*. Lausanne, Switzerland: EPFL, 2013.

[11] G. E. Santagati, T. Melodia, L. Galluccio, and S. Palazzo, "Medium access control and rate adaptation for ultrasonic intrabody sensor networks," *IEEE/ACM Trans. Netw.*, vol. 24, no. 4, pp. 1121–1134, Aug. 2015.

[12] E. Demirors, G. Alba, G. E. Santagati, and T. Melodia, "High data rate ultrasonic communications for wireless intra-body networks," in *Proc. IEEE Symp. Local Metrop. Area Netw. (LANMAN)*, Rome, Italy, Jun. 2016, pp. 1–6.

[13] J. Charthad, M. J. Weber, T. C. Chang, M. Saadat, and A. Arbabian, "A mm-sized implantable device with ultrasonic energy transfer and RF data uplink for high-power applications," in *Proc. IEEE Custom Integr. Circuits Conf.*, Sep. 2014, pp. 1–4.

[14] H. Lee, T. H. Kim, J. W. Choi, and S. Choi, "Chirp signal-based aerial acoustic communication for smart devices," in *Proc. IEEE Conf. Comput. Commun. (INFOCOM)*, Hong Kong, Apr./May 2015, pp. 2407–2415.

[15] G. E. Santagati and T. Melodia, "U-wear: Software-defined ultrasonic networking for wearable devices," in *Proc. ACM Conf. Mobile Syst., Services Appl. (MOBISYS)*, Florence, Italy, May 2015, pp. 241–256.

[16] S.-C. Kim and S.-C. Lim, "Transferring data from smartwatch to smartphone through mechanical wave propagation," *Sensors*, vol. 15, no. 9, 2015, Art. no. 21394.

[17] K. Zhang *et al.*, "Modeling and characterization of the implant intra-body communication based on capacitive coupling using a transfer function method," *Sensors*, vol. 14, no. 1, p. 1740, 2014.

[18] M. Swaminathan *et al.*, "Multi-path model and sensitivity analysis for galvanic coupled intra-body communication through layered tissue," *IEEE Trans. Biomed. Circuits Syst.*, vol. 10, no. 2, pp. 339–351, Apr. 2015.

[19] M. Swaminathan, U. Muncuk, and K. R. Chowdhury, "Tissue safety analysis and duty cycle planning for galvanic coupled intra-body communication," in *Proc. IEEE Int. Conf. Commun. (ICC)*, May 2016, pp. 1–6.

[20] *Medical Device Radiocommunications Service (MedRadio)*. Accessed: Oct. 2018. [Online]. Available: https://goo.gl/pjKY4C

[21] D. Son *et al.*, "Bioresorbable electronic stent integrated with therapeutic nanoparticles for endovascular diseases," *ACS Nano*, vol. 9, no. 6, pp. 5937–5946, 2015.

[22] F. T. Sun and M. J. Morrell, "Closed-loop neurostimulation: The clinical experience," *Neurotherapeutics*, vol. 11, no. 3, pp. 553–563, 2014.

[23] APC International Ltd. *First Steps Towards Piezoaction*. Accessed: Oct. 2018. [Online]. Available: https://www.americanpiezo.com

[24] N. Kushalnagar, J. Hui, and D. Culler, *Transmission of IPv6 Packets Over IEEE 802.15.4 Networks*, document RFC 4944, 2007.

[25] L. Deru, S. Dawans, M. Ocaña, B. Quoitin, and O. Bonaventure, "Redundant border routers for mission-critical 6LoWPAN networks," in *Real-World Wireless Sensor Networks*. Cham, Switzerland: Springer, 2014, pp. 195–203.

[26] I. Kuon and J. Rose, "Measuring the gap between FPGAs and ASICs," in *Proc. ACM/SIGDA 14th Int. Symp. Field Program. Gate Arrays (FPGA)*, 2006, pp. 21–30.

[27] *Blue Phantom Ultrasound Training Models*. Accessed: Oct. 2018. [Online]. Available: http://www.bluephantom.com/

[28] C. R. Mol and P. A. Breddels, "Ultrasound velocity in muscle," *J. Acoust. Soc. Amer.*, vol. 71, no. 2, pp. 455–461, 1982.

[29] T. Koch *et al.*, "Ultrasound velocity and attenuation of porcine soft tissues with respect to structure and composition: I. Muscle," *Meat Sci.*, vol. 88, no. 1, pp. 51–58, 2011.

G. Enrico Santagati (Member, IEEE) received the B.S. and M.S. degrees in telecommunication engineering from the University of Catania, Italy, in 2010 and 2012, respectively, and the Ph.D. degree in computer engineering from Northeastern University in 2017. At the same time, he was working at the Wireless Networks and Embedded Systems Laboratory, Northeastern University, under the guidance of Prof. T. Melodia. His current research interests are in ultrasonic intra-body sensor networks and software defined radios. He is specialized in the design and development of energy-constrained embedded systems with wireless software-defined connectivity and the Internet-of-Things capabilities.

Neil Dave (Student Member, IEEE) received the bachelor's and master's degrees from Northeastern University in 2016 and 2019, respectively. He is also responsible for hardware design and systems integration of devices with the Wireless Networks and Embedded Systems Laboratory under the guidance of Prof. T. Melodia. His primary research interests are in ultrasonic intra-body sensor networks, biomedical devices, and underwater communications.

Tommaso Melodia (Fellow, IEEE) received the Ph.D. degree in electrical and computer engineering from the Georgia Institute of Technology in 2007. He is currently the William Lincoln Smith Professor with the Department of Electrical and Computer Engineering, Northeastern University. He is also the Director of the Institute for the Wireless Internet of Things, and the Director of Research for the PAWR Project Office, a public-private partnership that is developing four city-scale platforms for advanced wireless research in the United States. His research focuses on modeling, optimization, and experimental evaluation of wireless networked systems, with applications to 5G networks and Internet of Things, software-defined networking, and body area networks. His research is supported mostly by the U.S. federal agencies, including the National Science Foundation, the Air Force Research Laboratory, the Office of Naval Research, the Army Research Laboratory, and DARPA. He is a Senior Member of the ACM. He is the Editor-in-Chief of *Computer Networks*, and a former Associate Editor of the IEEE TRANSACTIONS ON WIRELESS COMMUNICATIONS, the IEEE TRANSACTIONS ON MOBILE COMPUTING, the IEEE TRANSACTIONS ON MULTIMEDIA, among others.