

Secure Wireless Control and Remotely Triggered Actuation of a Gastric-Resident Ingestible Bioelectronic Capsule

by

Ian Ballinger

BS Electrical Engineering, University of Vermont, 2020

Submitted to the Department of Mechanical Engineering
in partial fulfillment of the requirements for the degree of

Master of Science in Mechanical Engineering

at the

Massachusetts Institute of Technology

September 2026

© 2026 Ian Ballinger. All rights reserved.

The author hereby grants to MIT a nonexclusive, worldwide, irrevocable, royalty-free license to exercise any and all rights under copyright, including to reproduce, preserve, distribute and publicly display copies of the thesis, or release the thesis under an open-access license.

Authored by: Ian Ballinger
Department of Mechanical Engineering
August 2026

Certified by: C. Giovanni Traverso
Professor of Mechanical Engineering, Thesis Supervisor

Accepted by: Nicolas G. Hadjiconstantinou
Professor of Mechanical Engineering
Graduate Officer, Department of Mechanical Engineering

Secure Wireless Control and Remotely Triggered Actuation of a Gastric-Resident Ingestible Bioelectronic Capsule

by

Ian Ballinger

Submitted to the Department of Mechanical Engineering on August 14, 2026,
in Partial Fulfillment of the Requirements for the Degree of

Master of Science in Mechanical Engineering

Abstract

A swallowed pill cannot be told when, where, or whether to release its payload. Its chemistry alone decides. This thesis develops an ingestible electronic capsule that removes that limit. The concrete case is commanded release of a bacterial therapeutic from a gastric-resident device, with on-demand treatment or prevention of Traveler's diarrhea as the motivating use case. The device holds the payload during gastric residence and opens it only on an authenticated wireless command.

That objective couples power, in-body communication, command authentication, gastric retention, and capsule-scale manufacturing into a single mechatronics problem in which battery current *magnitude* is the master constraint, because cell lifetime falls faster than linearly with it. A swallowable cell must hold an always-available, authenticated listening state across the entire window in which a command might arrive. Two generations reduced the design to practice. "iRES" established app-mediated Bluetooth Low Energy (BLE), a measured battery-lifetime budget, and a folded flexible-circuit build. "iTARGET-TD" integrated the successor control and payload-release hardware on an nRF5340 with stacked signaling, security, and power-delivery modules. A React Native smartphone application and Zephyr firmware share SPECK-128 challenge-response logic through common C++ sources.

Bench and in-body testing evaluated the platform against these requirements. On the bench, silver-oxide cells yielded a lifetime–current law, a single commanded actuation dissolved a payload membrane at a 5.06 mA peak and a 1.96 mA average over 47 minutes, and a staged in-body measurement placed control-link sensitivity near –102 dBm. Two individual devices (N = 1 each) each sustained an always-available link for roughly two to three months. In an integrated porcine study, a single gastric-resident device (N = 1) communicated continuously with its payload barrier endoscopically intact through a six-day window bounded by animal-facility scheduling rather than by the device, while separate retention studies established far longer residence. In a further intestinal-compartment study, a wireless command released a viable bacterial payload in vivo, recovered only in the triggered condition. These residence and release results come from complementary experiments, not one continuous trial. Together they show commanded in vivo release. They stop short of joining authentication and release in one synchronized trace, and of establishing therapeutic efficacy against Traveler's diarrhea.

Thesis Supervisor: C. Giovanni Traverso
Title: Professor of Mechanical Engineering

Contents

Acknowledgements

Chapter 1 · Introduction and Background

- 1.1 A Commandable Bacterial Therapeutic
- 1.2 The Engineering Problem
- 1.3 Literature Foundations
- 1.4 Thesis Objectives and Contributions

Chapter 2 · Requirements and Design Constraints

- 2.1 Power and Lifetime
- 2.2 Anatomical and Volumetric Compatibility
- 2.3 Wireless Link-Budget
- 2.4 Command Authentication
- 2.5 Retention, Reliability, and Manufacturability
- 2.6 Actuation and Release

Chapter 3 · Methods: Meeting the Requirements Across Two Generations

- 3.1 Power and Lifetime
- 3.2 Anatomical and Volumetric Integration
- 3.3 Wireless Communication and Link-Budget
- 3.4 Command Authentication
- 3.5 Retention, Reliability, and Manufacturability
- 3.6 Actuation and Release
- 3.7 Animal Study Methods

Chapter 4 · Results

- 4.1 Power and Lifetime
- 4.2 Command Authentication
- 4.3 iRES Design Reliability
- 4.4 Gastric Retention and Reliability
- 4.5 Remote In Vivo Communication and Bacterial Release

Chapter 5 · Conclusions, Discussion, and Future Work

- 5.1 Conclusions
 - 5.2 Discussion
 - 5.2.1 Limitations
 - 5.3 Future Work
 - 5.3.1 Capabilities the platform unblocks
- Supplemental Figures

Appendix A · Conservative RF Incident-Power Bound

Appendix B · Telemetry Collection-Pipeline Notes

List of Figures

- Figure 1.** Gastric-resident deployment and triggered-disassembly concept.
- Figure 2.** Platform lineage from iRES to iTARGET-TD.
- Figure 3.** A Schematic Overview of the iTARGET-TD Platform Device.
- Figure 4.** Measured sensor-mode current draw on iTARGET-TD.
- Figure 5.** Simulated deep-life brownout margin versus bypass capacitance.
- Figure 6.** The as-built iRES flexible circuit.
- Figure 7.** The as-built iTARGET-TD module stack.
- Figure 8.** Water-phantom bench setup.
- Figure 9.** Antenna impedance characterization in the water surrogate.
- Figure 10.** Fixed matching-network performance assessed by VSWR.
- Figure 11.** Tolerance error propagation through the fixed matching network.
- Figure 12.** Solved matching network and bench antenna measurement for iTARGET-TD.
- Figure 13.** Bench received-signal-strength gate for the matched device.
- Figure 14.** The iTARGET-TD antenna board.
- Figure 15.** Wireless signal path from radio to medium.
- Figure 16.** Authenticated actuation exchange (challenge → response → gate).
- Figure 17.** Designed security architecture, from provisioning to commanded dose.
- Figure 18.** Electrochemical actuation front-end.
- Figure 19.** Measured silver-oxide cell lifetime versus constant-current load.
- Figure 20.** Commanded dissolution-current trace with trigger and rupture frames.
- Figure 21.** A701 long-duration telemetry over 2.9 months.
- Figure 22.** A902 telemetry to battery death over 2.2 months.
- Figure 23.** SPECK-128 reference-vector agreement and challenge-response exchange.
- Figure 24.** In-medium bench comparison of MIT- and TRT-built devices.
- Figure 25.** Payload barrier intact at day 0 and day 6 in vivo (gastric-residence study, N = 1).
- Figure 26.** Bench triggering and retrieval of the retained device (gastric-residence study).
- Figure 27.** Wirelessly triggered release of a bacterial payload in vivo (intestinal-compartment study, iTARGET-TD).
- Figure 28.** Through-tissue control-link budget under gastric shielding.
- Supplementary Figure S1.** Power-delivery board.
- Supplementary Figure S2.** MCU board.
- Supplementary Figure S3.** Biometric module board.

Supplementary Figure S4. Antenna board.

Supplementary Figure S5. Instrument-level discharge capture.

List of Tables

Table 1. Chapter 2 requirements summary with measurable pass/fail criteria. The stated derating and any extrapolation are design assumptions, not measured outcomes.

Table 2. iRES and iTARGET-TD at a glance: the two design–build–test generations of the platform.

Table 3. Design power budget for the iTARGET-TD electronics: selected geometry, current limits, and time allocations (targets and allocations, not measured outcomes).

Table 4. As-built iRES flexible-circuit manufacturing constraints.

Table 5. Through-tissue control-link budget, measured end-to-end on a single device instance. Because every loss term is captured empirically on one device rather than stacked from independent estimates, this budget is authoritative.

Acknowledgements

I would first and foremost like to thank my advisor, Professor C. Giovanni Traverso of MIT's Department of Mechanical Engineering, for the opportunity to pursue this work, for his guidance and vision, for ensuring I always had the resources I needed, and for his unrelenting support over these last few years.

I would also like to deeply thank Dr. Ceara Byrne and Dr. Siheng You for their close guidance, mentorship, and day-to-day availability, which were vital to getting this work across the finish line.

I would like to thank the members of the Traverso lab for their critical efforts and support on this project. In particular, I thank William Markowitz, Injoo Moon, and Victoria Prepelook for their embedded systems expertise. I thank Dr. James McRae, Dr. Gary Liu, Dr. So-Yoon Yang, and Dr. Adam Gierlach for our shared time together working on this project. I learned a great deal from each of them. I am also grateful to the students and post-docs in the lab for their insight, camaraderie, and many fruitful conversations, and especially to all of my prior mentees, whom I have had the absolute privilege of watching grow: Anson Chan, Benjamin Lin, Hari Narayanam, Harrison Sun, Mark Yang, Rajani Tabor, Sophia You, and Teresa Ricapito.

I want to acknowledge and thank the animal team: Benedict Laidlaw, Maria Platero, Ashley Guevara, Andrew Pettinari, Alison Hayward, Niora Fabian, Kailyn Schmidt, and many others through the years, for their dedication and support with the in vivo work that was essential to both this thesis and many other projects.

I would also like to thank the team at Triple Ring Technologies (TRT) for their partnership in fabricating and assembling iRES devices, and for the manufacturing collaboration that made the build-to-build reliability comparisons in this work possible.

To my family and friends, thank you for your continued love and support throughout this program.

Generative AI tools were used to assist with phrasing and grammar in preparing this document. They were not used for any data analysis and I, the author, reviewed all material in this thesis.

Code availability. The original analysis and simulation source code supporting this thesis is published here [\[1\]](#). Figures and results derived from that code can be regenerated from the repository.

Chapter 1 • Introduction and Background

A swallowed pill cannot be told when, where, or whether to release its payload. Its chemistry alone decides. An ingestible *mechatronic* pharmacy lifts that limit. The device holds a payload through a residence period and opens only on command. That converts a one-shot chemical dose into an *electrically controlled actuation*: a mechanical release governed by a wireless signal. This thesis reduces one concrete instance of that solution to practice: commanded release of a bacterial therapeutic from a gastric-resident device, with on-demand treatment or prevention of Traveler's diarrhea as the motivating use case.

The therapeutic problem sets the design space. The technical foundations reviewed below fix the device within it, coupling power, in-body communication, command authentication, gastric retention, and capsule-scale manufacturing into a set of requirements that cannot be satisfied in isolation. Two design–build–test generations then reduce those requirements to practice, and the final device is evaluated against each.

1.1 A Commandable Bacterial Therapeutic

The project sponsor set the application that sets much of the design space: on-demand treatment of Traveler's diarrhea from a user-controlled, gastric-resident carrier that delivers an engineered live bacterial therapeutic, preserved through gastric residence, in response to an authenticated smartphone signal. The payload was an engineered *E. coli* Nissle strain, an established probiotic chassis [\[2\]](#) and one member of the broader class of engineered bacteria under development as diagnostic and therapeutic agents [\[3\]](#) whose oral delivery has motivated ingestible carriers [\[4\]](#), formulated for gastric survival by the species-specific methods of Jimenez et al. [\[5\]](#). That formulation would solve the therapeutic efficacy of the payload (stored at body temperature over months in deployment). The mechatronics had to solve the full-duplex communication, gastric retention, and payload release interfaces.

The payload sat behind a stainless-steel membrane. That membrane held the payload until an electrochemical current dissolved the barrier on command. The public patent record discloses that chamber-and-seal architecture, with separate bench and in vivo release examples [\[6\]](#).




FIGURE 1. Gastric-resident deployment and triggered-disassembly concept. This is the deployment concept the project sponsor specified at the outset. The shaped device stays in place through the commanded-release window and disassembles on trigger. The mechanical and electrochemical realization follows in the methods.

1.2 The Engineering Problem

The device had to wait, communicate through tissue, reject unauthorized actuation, and retain enough energy to open the chamber. Battery power and capacity sit at the center of these demands. To accomplish those goals, a swallowable cell must hold an always-available, authenticated listening-and-actuation state across the entire window in which a command might arrive. That window must further accommodate bursts of current during radio communication events, and high sustained currents for electrochemical operations. In particular, those performance details must remain satisfied at the *end* of the multi-month operational lifespan of the device, and the device must be aware of its remaining battery life to a sufficient degree so that it can self-disassemble and safely disengage from the gastric-retention configuration automatically at the end-of-deployment. Battery lifetime also falls faster than linearly with current draw, so every milliamp of amplitude spent on any of these capabilities degrades operating life super-linearly. Power-system characterization and the load current of every hardware or software change are therefore the first metrics to watch during development, ahead of any application-domain work.

1.3 Literature Foundations

This capsule sits within the established field of ingestible electronics for diagnostics and therapy [7]. Two prior lines of work define its lineage directly. The first is remotely triggered drug delivery: a wirelessly controlled implantable microchip released drug on command in a first-in-human study [8], building on the controlled-release microchip in which an electrochemically dissolved metal membrane gates a sealed reservoir [9]. That electrochemical-barrier principle is the one this device uses to seal and open its payload chamber. The second line is the integration of living engineered bacteria with ingestible wireless electronics, demonstrated by an ingestible bacterial-electronic capsule for gastrointestinal sensing [10]. That precedent, together with the broader maturation of engineered live biotherapeutic products for gastrointestinal delivery [11], frames the payload this platform is built to deliver. The device here descends from the same MIT lineage of secure ingestible systems that the following sections build on.

Retention, communication, and command authentication each rest on established engineering that the design adopts as a control. Keeping the device in the stomach uses a shaped, gastric-resident geometry that resists pyloric passage, an established approach to prolonged oral residence [12], within the size envelope characterized for ingestible capsules that transit the gut [13]. Here that geometry has one mechanical job. It holds the electronics and sealed payload in the stomach through the commanded-release window.

The same tissue that surrounds the antenna governs communication out of the body. Biological tissue presents frequency-dependent dielectric loading and attenuation to an in-body antenna [14], in-body biotelemetry is a mature design area [15], and impedance detuning varies with gastrointestinal location [16]. The device used 2.4 GHz Bluetooth Low Energy (BLE) for direct control from a commodity smartphone [17].

Authenticating the release command is a cryptography problem constrained by the power budget. The device gated actuation behind a challenge-response exchange using SPECK-128, a lightweight block cipher designed for constrained hardware [18] [19], addressing the class of threats surveyed for implantable and ingestible devices [20]. Jeon et al. report the predecessor BLE-and-SPECK implementation [17], and later Chandrakasan-group work extended the same secure-ingestible lineage toward post-quantum protocols [21]. The work presented here adds simple command authentication. It also adds the software and hardware interfaces that the device needs to support authenticated command handling given this background.

1.4 Thesis Objectives and Contributions

This thesis reduces to practice the electronic capsule described above (a gastric-resident, wirelessly commandable device that holds a bacterial payload through residence and releases it only on an authenticated command) and evaluates that platform against the coupled power, communication, authentication, retention, and manufacturing requirements it imposes. At swallowable scale the battery-current *magnitude* is the master design variable: cell lifetime falls faster than linearly with it, so every subsystem is priced in current, not in average power alone.

The specific contributions are:

- a capsule-scale wireless-control architecture carried across two design–build–test generations (iRES → iTARGET-TD);
- a measured silver-oxide battery lifetime–current law and a current-allocation power budget derived from it;
- an in-medium antenna-matching and acceptance methodology that specifies component-tolerance grades before the board is built;
- an authenticated command-handling path (a SPECK-128 challenge-response) shared as common C++ across the smartphone application and the device firmware;
- an integrated electrochemical payload-release subsystem; and
- complementary in vivo evidence of multi-day gastric residence, an intact payload barrier at day 6, sustained through-tissue communication, and commanded release of a viable bacterial payload.

Chapter 2 · Requirements and Design Constraints

2.1 Power and Lifetime

Ingestible devices that guarantee long-lived deployments must package a primary cell that sustains operation across the whole deployment window. Here that window is 60 days, with energy reserves for six therapy actuations plus a retention trigger, and the volumetric budget forces the cell to be small. Silver-oxide cell lifetime falls faster than linearly with the *magnitude* of current drawn, so the requirement prices sleep current and radio duty cycle as explicit costs against the 60-day target, and sets them against a 5 mA maximum prescribed electrochemical peak with a recommended operating point near 2 mA. Two levers therefore control the budget: one suppresses the dynamic range of the on-device power-draw signal, and the other lowers the DC effective operating current. The measured lifetime–current fit that quantifies this sensitivity appears in §4.1.

2.2 Anatomical and Volumetric Compatibility

Three device notions must stay distinct throughout this thesis. The *ingestible concept* is the long-run platform: a capsule small enough to be voluntarily swallowed. The *research prototype* actually built and tested here is administered endoscopically under sedation, and its delivered envelope is 16 mm in diameter by 44.8 mm in length. The *clinical device* is the future swallowable successor that would carry this work through human trials. Human translation was descoped during execution of this body of work: administration moved from a conscious swallow to endoscopic overtube passage, and overtube passage then set the governing geometric constraint on the delivered hardware. The unassisted-swallow envelope below therefore remains the standing target for any translational successor rather than a bound on the prototype.

Two geometric bounds apply to this device, and they are not the same bound. The long-run clinical bound is voluntary swallowing. Clinical capsule endoscopy fixes that one empirically rather than theoretically. Iddan et al. reported a wireless video-telemetry capsule endoscope small enough to be swallowed at 11 × 30 mm, carrying no external wires, fibre-optic bundles, or cables ^[13]. In ten human volunteers that capsule was swallowed without discomfort and passed, with a mouth-to-evacuation time of 24 h over

a range of 10 to 48 h. Later reviews treat pill-sized ingestibles as a standing device class alongside implantables and injectables [15]. An envelope near 11 mm by 30 mm is therefore an empirically supported ceiling for an unassisted swallow.

Swallowability was an active requirement at the start of this work. The program initially targeted that unassisted-swallow envelope so the result could translate through human trials. Human translation was descoped from the research effort during the period of execution, and the geometric requirement relaxed accordingly for later generations. Porcine gastric administration happens under sedation, so the device enters through an endoscopic overtube rather than down a conscious esophagus. Overtube passage therefore became the governing geometric constraint on the delivered hardware. The swallowable envelope above remains the standing target for any translational successor.

Gastric residence imposes a second constraint, and it is relational rather than absolute. A resident geometry has to resist pyloric passage once deployed [12]. That holds only while the deployed retention span exceeds the length of the body it anchors. Capsule length is therefore bounded above by the effective retention diameter rather than by an independent dimension. Retention performance and payload volume trade against each other through that one relation.

Inside whatever envelope results, the cells and the electronics compete for one cross-section. Cell chemistry fixes energy density, so cell volume converts directly into the operating lifetime of §2.1. The requirement is therefore an allocation rather than a dimension. Any cross-section granted to the electronics is lifetime removed from the power budget, and the allocation has to hold while preserving the electronics and payload seals. §3.2 gives the allocation each generation actually made.

2.3 Wireless Link-Budget

The application required control from a standard smartphone with no additional transceivers. That set the command path at 2.4 GHz BLE. It also allocated at least 20 dB between received signal strength and the applicable receiver floor as the link-margin requirement. The measured link quantities that this requirement is later checked against, a staged in-body sensitivity near -102 dBm (§4.5) and roughly 36 dB of margin in the pre-deployment bench check (§3.3), describe distinct configurations and must not be combined into one operating point. Transmit power was low and heavily duty-cycled, and formal SAR, tissue-heating, and regulatory electromagnetic-safety evaluation were out of scope for this work and are required before human translation. The radio-frequency human-exposure limits defined by IEEE [22] and ICNIRP [23] bound transmit-power choices for an in-body radio. Appendix A gives a conservative worst-case incident-power bound that supports the chosen transmit power.

§4.5 shows the primary RSSI trace behind this approximately -102 dBm in-body sensitivity measurement, alongside the in vivo triggering experiment. The full through-tissue link budget, measured end-to-end on one device instance, is tabulated there: a $+0$ dBm transmit reference, the benchtop received level, the in-body detuning and tissue path-loss terms, and the -102 dBm receiver floor. It fixes the pre-deployment bench acceptance criterion at -35 dBm received signal strength, with -55 dBm the strict minimum below which the received signal falls beneath the receiver floor once the device is shielded in the body.

2.4 Command Authentication

A wireless command that can open the payload chamber is a safety-critical attack surface. By construction, whoever can produce a valid command can actuate the device. The security and privacy exposure of connected implantable and ingestible medical devices is well documented [20], and the U.S. Food and Drug Administration now treats cybersecurity as a premarket requirement, expecting a documented authentication and authorization posture as a precondition for device submission [24]. The requirement here is therefore authorization rather than confidentiality. The capsule must wakeup into an actuation-capable state only in response to a legitimate command. An unauthenticated wakeup is therefore a safety hazard.

To meet that requirement within the device power budget, a pre-shared-key SPECK-128 challenge-response exchange gates actuation. That exchange builds on a lightweight block cipher suited to constrained microcontrollers [19]. The device issues a challenge. It enters the actuation-capable state only after the controller returns a response proving possession of the shared key. The mobile application and embedded firmware compute that exchange from common C++ sources, and shared reference vectors check both. The two endpoints therefore cannot silently diverge before integration. The successor platform additionally moved device communication onto a Bluetooth Mesh network, whose specification provides encryption and authentication by default. The SPECK gate then serves as a transport-independent authorization check for devices that do not support that standardized security (§3.4).

2.5 Retention, Reliability, and Manufacturability

Three constraints govern whether the device stays useful and buildable. **Retention** specified greater-than-60-day gastric residence with triggered disassembly while preserving the electronics and payload seals. The available release-platform evidence instead records shorter, unreconciled observations: at least six days of gastric communication and a separate day-6 payload-barrier image. Those observations do not

establish the 60-day target. **Reliability** entered as a failure-driven mitigation. In one identified study, communication degraded on an approximately 12-hour timescale, and retrieval implicated moisture and condensation, a recognized failure mode for sealed implantable electronics [25]. The response was desiccant, sealing, and material changes, which mitigate the hazard. **Manufacturability** constrained the electronics to a two-layer, 6/6 mil flexible-circuit process. That process also demanded a minimized bill of materials during an acute supply shortage. The design made the circuit foldable for packaging into a capsule, and that commensurate bending fatigue tolerance set the via and stiffener rules [26].

2.6 Actuation and Release

Payload actuation and release is the terminal requirement the whole device exists to serve. The device must hold the sealed payload intact throughout gastric residence and release it only on an authenticated command, delivering up to six therapy actuations plus a single retention-disassembly trigger within the electrochemical current and energy limits of §2.1 (a 5 mA prescribed peak against a recommended ≈ 2 mA operating point). The release mechanism must seal against the gastric environment through the commanded-release window and open on demand without compromising the surrounding electronics or the retention structure. This requirement anchors the actuation method (§3.6) and the actuation result (§4).

Table 1. Chapter 2 requirements summary with measurable pass/fail criteria. The stated derating and any extrapolation are design assumptions, not measured outcomes.

Requirement	Target	Rationale	Test method	Achieved?
Command latency	≤ 30 min	Application	Scan-interval test	Yes, bench (scan-interval design)
Operating lifetime	60 days	Deployment goal	Device telemetry	Partial: 2.2 and 2.9-month single-cell runs ($N = 1$ each, §4.1), with 60 d at the ≈ 18 μ A operating point an extrapolation of the fitted power law
Gastric residence	> 60 days	Therapeutic window	Porcine study	Integrated device to day 6 ($N = 1$, schedule-limited, §4.4). Prior whole-device retention 62–100 days ($N = 3$)
Payload seal integrity	Through residence	Prevent premature	Endoscopic inspection	Intact at day 6 by endoscopy ($N = 1$, §4.4)

Requirement	Target	Rationale	Test method	Achieved?
		release		
Actuation	Wireless command	User-triggered release	In vivo experiment	In vivo, CFU-confirmed single release (N = 1, §4.5) within the current envelope (5.06 mA peak, §4.1). Six-dose sequence not shown
Authentication	Reject invalid commands	Safety	Protocol test	Bench: SPECK-128 exchange verified, implementation-consistency only (§4.2)
Link margin	≥ 20 dB	Robust BLE	Bench / in vivo RSSI	Configuration-dependent: ~20 dB in bench telemetry, 0 dB at worst-case gastric shielding on the -102 dBm floor (§4.5)

Chapter 3 · Methods: Meeting the Requirements Across Two Generations

Two design–build–test generations built the platform. The first, iRES, answered the power, communication, and manufacturing questions that bounded the problem. The second, iTARGET-TD, tested whether those foundations could support authenticated control, modular electronics, and commanded bacterial release. The first generation supplied measured inputs to the second. Those inputs were a realizable folded flexible circuit, a battery-lifetime budget, and 2.4 GHz link measurements. The second generation integrated the release hardware on a stacked, modular architecture.

Evaluation proceeded from bench to in vivo testing. Bench measurements characterized the power, actuation, communication, and authentication subsystems under stated conditions. Separate porcine studies then showed a gastric-resident device sustaining wireless communication for several days. Endoscopic inspection showed an intact payload barrier, and a wireless command triggered release of a viable bacterial payload.




FIGURE 2. Platform lineage from iRES to iTARGET-TD. The first generation (iRES, left) placed the nRF52811, battery contacts, electrochemical circuitry, subsystem connectors, and programming interface on a folded monolithic flexible circuit. The second (iTARGET-TD, right) used an nRF5340 and a stacked rigid-PCB architecture of a microcontroller(MCU), sensing, and power-delivery modules. That stack integrated authenticated BLE control with the payload-release electronics.




FIGURE 3. A Schematic Overview of the iTARGET-TD Platform Device.

That first generation established what was achievable under the tightest constraints. Its most consequential output was a lesson about where the design had been over-constrained, more than any single subsystem. Chief among those lessons was volume. Keeping the device in the stomach of a porcine model required lengthening the gastric-retention arms repeatedly before it reliably stayed put. Each extension relaxed the cross-section the electronics had to fit. That relaxation is the origin of the defining choice in iTARGET-TD: a stacked, modular architecture. That architecture spends the freed volume on larger cells and dedicated security and sensing subsystems instead of one larger board. The sections below give the design reasoning behind how each generation met each requirement in turn. Supplemental figures S1–S4 collect the subsystem board schematics that realize these choices on iTARGET-TD.

iRES initially demonstrated a microcontroller and BLE radio with an nRF52811 [27]. iTARGET-TD moved to an nRF5340 [28], with Zephyr [29] carrying application logic, drivers, and BLE integration, and common C++ sources sharing the command protocol with the smartphone application [30], and that shared protocol reaches the React Native application as a C++ TurboModule over the JavaScript Interface (JSI) [31]. iRES and iTARGET-TD are two generations of one project rather than a survey of the field: iRES established a capsule-scale flexible circuit, battery characterization, app-mediated BLE, and an electrochemical payload interface, and iTARGET-TD carried those into a modular successor that added authenticated command handling and bacterial release [32] [6].

Table 2. iRES and iTARGET-TD at a glance: the two design–build–test generations of the platform.

Attribute	iRES	iTARGET-TD
-----------	------	------------

Attribute	iRES	iTARGET-TD
Primary purpose	Low-power communication + manufacturability	Integrated secure control + payload release
Microcontroller	nRF52811	nRF5340
PCB architecture	Folded flexible circuit	Stacked rigid PCBs
Silver-oxide cell diameter	6.8 mm	9.6 mm
Target lifetime	60 days	$\approx$ 1 week (faster data rate config.)
Retention configuration	Nitinol arms + rigid chassis	Nitinol arms + rigid chassis
Payload release	Stainless-steel electrochemical barrier	Stainless-steel electrochemical barrier
Long-duration evidence	Multi-month telemetry (A701, A902; N = 1 each, largely bench)	Day-6 barrier photo (N = 1) + commanded release

3.1 Power and Lifetime

The electronics exist only to deliver release and disassembly, so power minimization was paramount. The hard problem was sixty days of always-available command reception and actuation on a swallowable cell. The first generation sized an ultra-low-power sleep architecture around that constraint. It selected the nRF52811 specifically for its low active and sleep current. It also characterized its silver-oxide cells by bench discharge across five constant-current loads. That discharge yielded the current-budget law that converts the 60-day target into a per-subsystem current allocation. That law penalizes current *magnitude* worse than linearly, so iTARGET-TD treated dynamic current as a primary design variable. The power-delivery module added a ripple-rejection capacitor bank to smooth the battery-observed current spikes from radio and actuation bursts. The budget also set the radio duty cycle. The second generation then deliberately used the additional volume headroom inherited from iRES. Larger cells bought energy headroom for the added security and sensor loads that the first generation could not afford.

Inverting that law gives iRES a single hard ceiling on average current. Solving the measured lifetime–current fit of §4.1 at the 60-day target returns approximately 29 μ A as the largest DC average the 6.8 mm cells can sustain. That ceiling is a property of the cells rather than of the architecture, so it is parametrized by cell diameter and does not carry over unchanged to a different cell. It is also optimistic in three ways. It describes a

fresh single cell under a constant load, it ignores shelf-life self-discharge, and it extrapolates roughly an order of magnitude below the 500 μA lowest measured point of the fit. iRES therefore budgeted to 0.6 of the ceiling, approximately 18 μA , holding the remainder as margin against self-discharge, unit-to-unit spread, and the extrapolation itself. That 0.6 is an engineering margin rather than a measured derating factor. No discharge in the characterized set constrains the factor at this current, so it is stated as a design choice and not as a result. Every subsystem allocation in this chapter is drawn against the 18 μA figure rather than against the 29 μA arithmetic limit.

iTARGET-TD is not bound by that number, for two independent reasons. It carries larger 9.6 mm cells, which moves the underlying capacity and therefore the ceiling. Its scope also changed. Automatic disassembly after months of residence was descoped from the research effort, so continuous multi-month availability stopped setting the power budget. The design prioritized sensor-data fidelity instead, and sized its operating point to roughly one-week lifetime measurements. That is a deliberate exchange of endurance for data density, not a regression against the iRES budget.

Table 3. Design power budget for the iTARGET-TD electronics: selected geometry, current limits, and time allocations (targets and allocations, not measured outcomes).

Quantity	Value	Kind
Silver-oxide cell diameter	9.6 mm	Selected geometry
Shared capsule diameter	12 mm	Design constraint
Maximum prescribed electrochemical current	5 mA	Design constraint
Continuous operation at 5 mA	75 min	Budget calculation
Recommended electrochemical current	2 mA	Design point
Therapy allocation	6 doses $\times$ 30 min each	Design allocation
Retention-trigger allocation	7 min	Design allocation

Sizing the cells for a 60-day life is as much a current-*dynamic-range* problem as an average-energy one. As a primary cell depletes, its internal resistance rises nonlinearly with elapsed time and depleted capacity as its chemistry changes. That rise goes well beyond the constant-resistance assumption typically presented within a datasheet capacity rating. At end of life of a cell, this elevated source resistance can still deliver small, steady currents. Low-current operations such as ~ 2 mA foil dissolution therefore continue to succeed. Device performance *appears* nominal when testing covers only that scenario. However, the same cell can no longer supply a brief (several microseconds) radio-transmit burst. Its terminal voltage collapses below the brownout point of the regulator. Radio performance degrades at that point. High-current operations therefore fail first, and silently, even when low-current self-checks pass.

Usable capacity is intrinsically tied to the *magnitude* of the current drawn, worse-than-linearly. That is consistent with the measured $t \propto I^{-1.37}$ law and with prior reports. The point at which the cell counts as dead is therefore set by the magnitude of the current the load draws, so the same cell dies sooner under high-current radio bursts than under a steady low-current load.




FIGURE 4. Measured sensor-mode current draw on iTARGET-TD. The window spans two minutes. The IMU ran at 100 Hz with 1 s on-time per 15 s period. Temperature and PPG sampled each second. The trace shows a **587 µA average against a 60 mA peak**, roughly a 100× peak-to-average ratio. Lifetime penalizes current *magnitude* worse than linearly. Suppressing these spikes with the power-delivery ripple-rejection capacitor bank and lowering the DC average are therefore both first-order design levers.

Power delivery in iTARGET-TD preserves that dynamic range. It sizes a bypass capacitor bank to low-pass filter current bursts as the battery sees them. The cell supplies the slow average current through its rising internal resistance. A local capacitor bank delivers the fast burst charge and holds the terminal voltage above brownout. The simulation below drives the measured radio-burst current into a battery modeled at a 300 Ω deep-life internal resistance. The brownout threshold is 1.75 V. That resistance is estimated from Power Profiler Kit 2 discharge measurements, whose nonlinear resistance-versus-capacity rise motivates the value. A 1 µF bypass collapses to 0.95 V on every burst. A 22.6 µF bank (six 4.7 µF ceramics, derated for DC-bias and temperature loss of effective capacitance [33]) holds a 2.77 V minimum. Bypass capacitance is thus a compact, tunable lever. That lever extends useful operating life beyond the "dead-battery" state a constant-resistance datasheet metric would predict. It trades board area, rather than scarce chemistry volume, for burst headroom late in life.

The simulation is fully reproducible from the thesis code repository [1] (`src/brownout_bypass_sim.jl`). It integrates the single-node charge balance $C dv_t/dt = (V_{oc} - v_t)/R_{int} - i(t)$ as an ACME discrete-time state-space model at a fixed 10 µs step. The load current $i(t)$ is the measured Power Profiler Kit 2 radio-burst capture (10 µs

sample period, ≈ 15.6 mA peak). The cell is a Thévenin source at $V_{oc} = 3.1$ V with a deep-life $R_{int} = 300\ \Omega$ read from the high-depletion tail of the constant-current discharge runs. The regulator is modeled as a fixed 1.75 V brownout comparator on the terminal voltage. The 22.6 μ F bank is six 4.7 μ F X5R ceramics with a 0.8 DC-bias derating factor ($6 \times 4.7 \times 0.8\ \mu$ F). Capacitor ESR (a few m Ω) is neglected as roughly five orders of magnitude below the 300 Ω source resistance. A built-in internal-resistance sweep (100–500 Ω) reports the brownout margin's sensitivity to that 300 Ω estimate.




FIGURE 5. Simulated deep-life brownout margin versus bypass capacitance. The simulation drives battery terminal voltage with the measured radio-burst current at a deep-life internal resistance ($R = 300\ \Omega$, $OCV = 3.1$ V). The brownout threshold is 1.75 V. A 1 μ F bypass browns out on every burst (minimum 0.95 V). A 22.6 μ F bank holds 2.77 V. Internal resistance rises nonlinearly toward end of life. High-current radio bursts therefore brown out first, even while low-current operations such as foil dissolution still succeed. Bypass capacitance is therefore sized to preserve the current dynamic range of the supply across the operating window.

The power-delivery module realizes these levers in hardware. Its bypass bank places six 4.7 μ F X5R ceramics in parallel across the cell rail. That bank is the one sized above against the transmit-burst load. A TLV5518 (used as a load switch) gates a switched rail to the adapter connector under a dedicated enable line. DF40 board-to-board connectors join it to the MCU stack. The BLE transmit burst that this bank buffers is roughly 20 mA. That burst should be the first capability to degrade as internal cell resistance rises late in life. This is consistent with the brownout model above. A Keysight DC power analyzer in battery-profiler mode recorded each discharge load as an

instrument-level capture. Supplemental figure S5 reproduces one such capture. Supplemental figures S1–S4 carry the full power-delivery, MCU, biometric-sensor, and antenna board schematics.

3.2 Anatomical and Volumetric Integration

iRES fit the electronics into an approximately 9 mm capsule cross-section shared with 6.8 mm silver-oxide cells, as a narrow two-layer flexible circuit initially targeting 000-capsule compatibility. The circuit folded beside the cells, and the volume assigned to the cells was not available to the electronics. That board carried the microcontroller and radio, battery contacts, electrochemical front-end, connectors, and programming interface. The largest available cross-section dictated the form. That cross-section, however, was relaxed for iTARGET-TD. Reliable gastric retention in a porcine model forced repeated lengthening of the retention arms. Each extension relaxed the volume the electronics had to respect. Descoping human translation relaxed it again: the unassisted-swallow envelope no longer bound the design, replaced by endoscopic overtube passage for swine under sedation. The second generation used that freed volume for a new architecture rather than a single larger board. It split the electronics into signaling, security, and power-delivery modules on stacked rigid PCBs joined by board-to-board connectors, and allocated the cylindrical envelope along its axis instead of folding one sheet as iRES did. The modular stack is therefore a direct fallout of the iRES retention experience. Engineers could revise, retune, and validate each subsystem independently. That stacked method is also more volumetrically efficient, and the freed volume accommodated the larger cells of approximately 9.6 mm diameter that the power budget of §3.1 relies on.

The quantitative account of that retention-arm lengthening is part of the iRES retention work, reported in the iRES manuscript now in preparation, and is not reproduced here.

The delivered iTARGET-TD assembly resolves those constraints to specific numbers. The retention system carried over from iRES gives an effective retention diameter of 47.5 mm. The capsule measures 44.8 mm in length, which keeps it strictly below that retention span, as the relation of §2.2 requires. In its folded administration configuration the capsule measures 16 mm in diameter, which is set by endoscopic overtube passage rather than by unassisted swallowing. These dimensions are taken from the CAD assembly.




FIGURE 6. The as-built iRES flexible circuit. It carries the microcontroller and radio, battery contacts, electrochemical front-end, subsystem connectors, and programming interface on one two-layer sheet. The board folds beside the 6.8 mm cells within the shared capsule cross-section. That packaging constraint drove the bend-line and stiffener rules of §3.5.




FIGURE 7. The as-built iTARGET-TD module stack. The photograph shows the stacked, modular electronics as physically realized. Signaling, security, and power-delivery boards join through board-to-board connectors within the capsule cross-section. The photograph documents the as-built architecture, not final dimensions or measured performance.

3.3 Wireless Communication and Link-Budget

Authenticated communication and control through tissue is hard because biological media are lossy and high-permittivity [14], attenuating the signal and detuning the antenna at once, so a link that operates comfortably on the bench can fail in vivo [15]. That loading is not something the program specified in advance. It fell out of the tuning work, and the methods below had to characterize it before they could correct for it.

Direct control from a commodity smartphone with no additional transceivers was a program requirement rather than a free design choice. That fixed the radio at 2.4 GHz BLE. The design problem was therefore not to select a band. It was to make that mandated band operate reliably inside the body. Because surrounding tissue detunes an in-body antenna, antenna impedance characterization covered free space and a water surrogate, and a single fixed matching network targeted the resulting bounded impedance locus, with match quality assessed by voltage standing-wave ratio (VSWR) against a project acceptance band [34]. A porcine model evaluated in-body performance, with the device placed via endoscopic overtube and an external receiver logging telemetry. Two regimes ran. The first was long-duration ambulatory telemetry over the full device-to-gateway-to-backend path. The second was a sedated high-gain sensitivity measurement with the receiver repositioned to minimize distance. A pre-deployment bench received-signal-strength check served as a manufacturing go/no-go gate.

The bench surrogate was water. A capsule was suspended on a thread at mid-depth in a 1000 mL beaker, and the separation to a fixed receiver was set with a tape measure so that runs could be repeated at a known distance. Water is a first-order stand-in. Its permittivity at 2.4 GHz sits above that of muscle, so it detunes the antenna somewhat harder than tissue does, and a formulated tissue-equivalent phantom would track in-body loading more closely. That refinement was not pursued, and the reason is programmatic rather than technical. Large-animal access was routine on this program, so any question that a better phantom would have answered could instead be settled directly in the animal. Water was therefore sufficient for its actual job, which was to provide a repeatable lossy-medium bench check between animal sessions.



FIGURE 8. Water-phantom bench setup. A capsule hangs on a thread at mid-depth in a 1000 mL beaker of water, the surrogate medium for these runs. A tape measure sets the separation to the receiver held in the bench vise, which makes the geometry repeatable between runs. Water loads the antenna somewhat harder than muscle at 2.4 GHz, so this arrangement is a repeatable bench check rather than a tissue-equivalent reproduction of in-body loading. Routine large-animal access meant in-body behavior could be confirmed directly, so a formulated phantom was not required.




FIGURE 9. Antenna impedance characterization in the water surrogate. These measurements designed the single fixed matching network. Left: the surrogate measurement setup. Right: the matched-antenna response in surrogate.




FIGURE 10. Fixed matching-network performance assessed by VSWR. The assessment ran against the project acceptance band across the bounded impedance locus. These are bench design characterizations of the matching network, taken during device manufacture, not in vivo results.

A single fixed network only holds the match if the parts that build it are repeatable, so component tolerance was treated as a design variable in its own right. Setting that model up first required both of its endpoints. The source endpoint is the nRF port impedance, which the vendor does not publish directly, so it was back-calculated from the datasheet-recommended y-network and its 50 Ω design target, giving an inferred port impedance of $72 - j26 \Omega$. The load endpoint is the antenna itself, measured in the water phantom with the matching network shorted so that the bare element was seen alone. That measured in-phantom impedance, 17.75Ω , is the matching target and sets the normalization at the center of the chart. Between those endpoints sit six lumped elements around a 6 mm 50 Ω line: series 2.2 nH, shunt 0.7 pF, series 2.2 nH, the line, then shunt 1.78 pF, series 2.48 nH, and series 4.8 pF. Each element was then perturbed across its as-built tolerance band, ± 0.1 pF on the capacitors and ± 0.1 nH on the inductors, and the resulting spread of input impedances plotted against acceptance circles at VSWR 1.5 and 3. The reachable region stays bounded near the acceptance band, so a single fixed network holds the match under the tolerances of the parts actually stocked. This analysis is what set the purchasing requirement and specified those ± 0.1 pF and ± 0.1 nH component grades, before any board was built.




FIGURE 11. Tolerance error propagation through the fixed matching network. The model runs between an inferred nRF port impedance of $72 - j26 \Omega$, back-calculated from the datasheet-recommended y-network and its 50Ω target, and a matching target of 17.75Ω , the antenna element measured in the water phantom with the matching network shorted. That target sets the chart normalization, so a perfect match lands at the center. Each element is perturbed across its as-built tolerance band, $\pm 0.1 \text{ pF}$ on the capacitors and $\pm 0.1 \text{ nH}$ on the inductors, and the reachable input impedance is plotted against acceptance circles at VSWR 1.5 and 3. The traced path steps the impedance element by element from the source to the antenna, and the shaded region is the reachable input-impedance spread under those tolerances, which stays bounded near the acceptance band. This is a theoretical sensitivity model of the network, not a measurement of built boards, and it is the basis on which the $\pm 0.1 \text{ pF}$ and $\pm 0.1 \text{ nH}$ component grades were specified. The network and its tolerances are reproducible in an online Smith-chart tool [35].

That tolerance requirement then had to be satisfiable with the discrete components actually available in stock. The load for the final solve came from a direct measurement of the antenna element on its own, taken on a vector network analyzer in air rather than in the phantom. Measuring in air is far easier to set up, and it is defensible at one frequency. The in-air and in-water measured responses appear to converge near 2440 MHz, so to a first approximation an in-air reading there also describes the loaded element. This was treated as a practical working assumption rather than a precisely characterized crossing. The antenna element targets BLE operation, so an approximate convergence inside the band is not entirely a coincidence, though its exact placement is convenient. That measurement gives $17 - j14.45 \Omega$ at 2440 MHz. This is the in-air antenna element figure, distinct from the 17.75Ω in-phantom value that normalizes the tolerance model above. Solving it against a 50Ω source with an impedance-matching calculator [36] returns a shunt capacitor of 1.818 pF and a series inductor of 2.487 nH .

The pi footprint was populated with the closest nominal values on hand, 1.8 pF (± 0.1 pF) and 2.4 nH (± 0.1 nH), with the third position left open. Substituting stocked nominals for the solved values is exactly the kind of perturbation the tolerance study bounded, so the build tests that analysis rather than departing from it. The assembled device then went through the bench received-signal-strength gate. Adjacent to the receiver it returned roughly -20 dBm, and at 1 m it held roughly -60 dBm, which clears the go/no-go threshold with margin.




FIGURE 12. Solved matching network and bench antenna measurement for iTARGET-TD. Left: the network solved for the antenna load measured in air at right, $17 - j14.45 \Omega$ at 2440 MHz against a 50Ω source, returning a shunt capacitor of 1.818 pF and a series inductor of 2.487 nH. The pi footprint was populated with the nearest stocked nominals, 1.8 pF (± 0.1 pF) and 2.4 nH (± 0.1 nH), with the third position left open. Right: the antenna element measured on its own on a vector network analyzer over a 500 MHz span. This reading was taken in air rather than in the phantom, for ease of setup. The marker sits at 2440 MHz. Air and water responses approximately converge near that frequency, so to a first approximation an in-air reading there also describes the loaded element. This in-air element figure is distinct from the 17.75Ω in-phantom value that normalizes the tolerance model in the preceding figure.




FIGURE 13. Bench received-signal-strength gate for the matched device. A scan of the assembled device advertising at a 103 ms interval. Signal strength sits near -20 dBm with the receiver adjacent to the device and near -60 dBm at 1 m, clearing the manufacturing go/no-go threshold. These are bench readings from an uncalibrated receiver in an ordinary lab environment. They are separate from the in-body quantities discussed in §2.3 and must not be combined with them.




FIGURE 14. The iTARGET-TD antenna board. This is the board that carried the matched design into the animal. The $50\ \Omega$ RF port of the MCU reaches a 2450AT07A0100 chip antenna through a board-to-board connector. The discrete matching footprint (R1, C2, C3) carries a do-not-populate mark. The delivered configuration therefore relies on the chip antenna and the matching network solved above, not on board-level tuning.




FIGURE 15. Wireless signal path from radio to medium. The 50 Ω RF port of the nRF5340 reaches the surrounding medium through a board-to-board connector, the fixed π matching network solved in §3.3, and the chip antenna. Each interface between these stages is a place the in-body match can detune, which is why the matching network was specified against a bounded impedance locus rather than a single nominal.

BLE test parameters. The control link runs on the LE 1M PHY (`CONFIG_BT_PHY_UPDATE=n`; no long-range coded PHY) as a connectable peripheral advertising under the device name *ADAPTER Remote Activation*. Advertising uses Zephyr's

`BT_GAP_ADV_FAST_INT_MIN_2 – MAX_2` window, a 100–150 ms interval (the bench go/no-go capture above was taken at ≈ 103 ms). Radio transmit power is +0 dBm in the default firmware configuration. The reported link measurements were taken at +4 dBm, the nRF radio maximum. During development the device woke on an approximately 30 s scan cycle with a few seconds of advertising (§3.4). The budgeted deployment configuration relaxes that to the ≈ 10 min cadence set by the 30 min command-latency requirement. The receiver was a smartphone BLE scanner reading advertisement RSSI (the history view reproduced above), with bench readings of ≈ -20 dBm adjacent and ≈ -60 dBm at 1 m. The exact receiver model, antenna orientation, and per-point repetition counts were not recorded uniformly across the bench captures.

3.4 Command Authentication

The team selected SPECK-128 for its low cycle and energy cost on the constrained microcontroller, its small code size, and its implementability as a single shared C++ routine across the React Native application and the Zephyr firmware. It gates actuation with a lightweight authorization check inside the power budget. It is a constrained implementation demonstrating feasibility, not a recommended clinical security architecture. An early concept targeted dedicated wake-up silicon. The challenge-response instead proved cheap enough to ride the existing BLE scan cycle of the radio. It cost a few extra clock cycles each time a scan already ran. That removed the need for both a separate high-power listening mode and any dedicated hardware. One state machine, written once and shared as common C++, served both the React Native application and the Zephyr firmware. A C++ TurboModule exposes it to the application. Both endpoints therefore computed the identical exchange, and shared reference vectors checked the two implementations before integration. Latency rather than throughput set the duty cycle. A command had to arrive within roughly thirty minutes. The budgeted configuration therefore woke to scan on the order of every ten minutes.

The research builds scanned far more often (about every 30 s, with a few seconds of advertising) during development. Meeting a 30-minute reception latency against a 10-minute scan interval was therefore comfortable. The authenticated wakeup cost only the incremental cipher work atop scans the radio already ran. A full wireless-security threat evaluation (replay, relay, key management, denial of service, and downgrade among others) was out of scope for this work, because the pre-human, animal-facility stage did not require pushing wireless security beyond an authenticated command gate. The successor iTARGET-TD generation subsequently moved device communication onto a Bluetooth Mesh network. Bluetooth Mesh mandates network- and application-layer encryption and authentication by specification, so on that transport a standardized security layer is present by default. The foil-dissolution actuation firmware builds used for the release experiments deliberately kept plain, unprovisioned BLE rather than mesh, so that any smartphone could issue the command instead of only a phone provisioned into the mesh network. That isolated the release result from the provisioning step as an experimental variable. Adding mesh to the actuation build is straightforward and was demonstrated on a parallel mesh-server build. The SPECK challenge-response is therefore most valuable not as a competitor to that standard but as a portable authorization gate for links and devices that do not implement Bluetooth Mesh's now-standard security: a self-contained check that rides an ordinary BLE scan cycle without depending on a mesh stack.

The exchange is a single SPECK-128/128 block operation: a 128-bit block and a 128-bit pre-shared key, two 64-bit words each, run over 32 rounds with rotation constants $\alpha = 8$ and $\beta = 3$. The device draws a fresh 128-bit challenge from its hardware entropy source on each exchange. The application returns the SPECK-128 encryption of that challenge under the pre-shared key. The device recomputes the same transform and gates entry to the actuation-capable state on a match. The pre-shared key is loaded once through the key-set routine and is never transmitted or read back across the native bridge (a Diffie–Hellman key-agreement path is scaffolded in the cipher module but the deployed gate uses the static pre-shared key). Because the block transform is identical on both endpoints and shares one C++ source, the production core is verified against the public SPECK-128/128 reference vector before integration. On an invalid or malformed response the shared state machine takes its reject path into a hard-fault state, emitting a remote-debug signal rather than re-prompting, so a failed attempt does not simply loop for retries. Replay across sessions, challenge-cache expiration, and role-scoped re-authorization of actuation are not hardened in this implemented subset. They belong instead to the fuller designed architecture described next (a Signal-style double-ratchet with time-limited, role-scoped permissions) and, at the transport layer, to the Bluetooth Mesh standard the successor platform adopts (§2.4). The guiding choice in both is to reuse well-vetted, maintained constructions rather than to extend a bespoke gate.




FIGURE 16. Authenticated actuation exchange (challenge → response → gate). The device issues a challenge on a scan it already runs. The application returns a SPECK-128 response computed from the pre-shared key. The device verifies it and enters the actuation-capable state only on a valid response. Both endpoints compute the exchange from common C++ sources. The challenge and key are each 128 bits (SPECK-128/128, 32 rounds, $\alpha = 8$, $\beta = 3$). The challenge is a fresh hardware-random value per exchange, and an invalid response drives the shared state machine to a hard-fault rather than a retry. Broader threat hardening (replay, relay, downgrade) belongs to the designed architecture below and to the Bluetooth Mesh transport layer (§2.4), not to this implemented gate.

The implemented gate is the demonstrated core of a larger security architecture the program specified and reviewed with cryptography-expert collaborators but did not carry into the pre-human builds. That designed protocol has four key properties. First, firmware is updated at the point of administration: the device is flashed with the current signed binary immediately before ingestion, so no compiled image sits on the device in its resting state, which raises the cost of reverse engineering, and proactively blocks a set of time-of-validation to time-of-use attacks. Second, a two-way integrity check lets the device run only firmware that is both signed by the manufacturer (the *provisioner*) and decryptable by a hardware-backed private key generated at manufacture and never

extracted, so a software or hardware supply-chain compromise is detectable with the other half as a co-root of trust. A failed software root triggers a pre-installed automatic-disassembly routine that prevents erroneous permanent residence, and a hardware compromise stays observable to the provisioner. The provisioner that signs firmware and sets device policy need not be the device manufacturer. A companion organization can self-host that role instead. Third, end-to-end encryption follows a Signal-style double-ratchet established at setup among patient, provisioner, and healthcare-provider roles, so patient health data is never identifying and the provider-patient relationship is enforced by the protocol itself, a deliberate reuse of a well-vetted construction rather than a bespoke cipher suite. Because that channel runs end to end from the healthcare provider to the device, the provider sets the scope of allowed commands directly, and the patient's phone only relays sealed traffic. The end-user hardware therefore need not be trusted, and compromise of a patient-supplied phone does not widen what the link can command. Fourth, communication is minimized for the hostile environments the program's setting anticipates: the device assumes a hostile channel from administration until a permissions-update step, and steady-state operation needs only two packets (a wakeup pattern and a command, which may be merged into one event) with an optional status reply. The hostile framing is a default posture rather than a response to a named adversary. At the scale of large embedded deployments, every device should be treated as low-hanging fruit, so each one is optimized for a minimal attack surface by default. Merging the wakeup and command into a single sealed packet is therefore the minimal steady-state configuration. Role-scoped, time-limited permissions (a bounded number of authorized wakeups and a permissions timeout) bound what a captured link can do and for how long.




FIGURE 17. Designed security architecture, from provisioning to commanded dose.

Redrawn from the program's designed protocol. In the secured hardware-access phase the provisioner locks and attests the device and prints an encrypted serial on the packaging. Over the internet/Bluetooth setup phase the patient and healthcare provider establish double-ratchet sessions, claim their roles from the packaged serial, update firmware at the point of administration, and receive a time-limited permission grant (a bounded number of authorized wakeups and a permissions timeout). In the Bluetooth-only phase the device treats the channel as hostile and accepts only a wakeup pattern followed by a command, optionally replying with status. The SPECK challenge-response above is the bench-implemented subset of this command gate. The double-ratchet sessions, attestation, and role-scoped permissions were designed but not implemented in the pre-human builds.

3.5 Retention, Reliability, and Manufacturability

Retention used a shaped gastric-resident geometry, comprising releasable Nitinol retention arms and a rigid plastic main-chassis, to hold the electronics and sealed payload in the stomach through the commanded-release window ^[12], with triggered electrochemical disassembly to release the device by detaching the retention arms from the main chassis. Reliability entered as a failure-driven response. Resident testing implicated moisture and condensation in all three of communication-, retention-, and viability-degradation failures. The successor iTARGET-TD design therefore integrated a desiccant and revised its sealing, adhesive, and material choices. Telemetry analysis

separated backend and infrastructure outages from real device failures.

Manufacturability then favored the modular design of §3.2. The team could iterate and tune individual subsystems independently at build-time, most consequentially the antenna. A single detuned or failed subsystem therefore no longer forced a full-device rebuild.

The detailed iRES retention and long-duration reliability results are reported in a manuscript in preparation (McRae et al.).

Manufacturing yield and per-build assembly statistics are deliberately not reported. The prototypes were assembled by hand on a lab-scale, single-operator line, which is not a representative proxy for a production manufacturing flow, so yields measured there would not generalize and could mislead. The as-built process constraints that do carry design information are summarized below.

Table 4. As-built iRES flexible-circuit manufacturing constraints.

Constraint	Value / rule
Layers	2
Trace / space	6/6 mil
Core	5 mil polyimide
Copper	1 oz
Package style	QFN
Substrate	Flex PCB with stiffener
Vias	Enlarged diameter, reduced count
Bend lines	7 required
Sourcing	Buy constrained parts early; minimize BOM length

3.6 Actuation and Release

The membrane material was a deliberate trade. Gold foil was a candidate for its inertness, but its nanofabrication seemed unnecessary pitched against stainless steel. Stainless steel is durable and readily purchasable. A 2 mm-diameter, 25 μ m-thick stainless-steel foil seals each chamber, one foil per chamber across the six-chamber design. The composition and biological fate of the electrochemical dissolution products were not characterized in this work. Stainless steel can release chromium, nickel, and other alloy constituents, so a dedicated toxicological evaluation of the dissolution products is required before human translation. The therapeutic event was then an electrochemical release. A controlled current dissolves the stainless-steel membrane

sealing the payload chamber in the acidic stomach environment. That dissolution still proceeds up to neutral pH, though at higher energy consumption. Any measured dissolution power trace therefore comes from a neutral-pH environment as a worst case. The dissolution-current trace characterized bench actuation. A split protocol tested commanded in vivo release, separating two burdens a single experiment would have coupled. One gastric study covered residence, payload viability, and post-extraction operation. A separate intestinal-compartment experiment made wireless-triggered release directly observable. The experimental detail and outcomes are reported with the results. The electrochemical foil-dissolution release mechanism described in this section was developed by Dr. So-Yoon Yang. The author's contribution on this front was integrative, providing the electrical actuation subsystem, its interfaces, and the command-availability and power budget that drive the dissolution current.




FIGURE 18. Electrochemical actuation front-end. The circuit drives the controlled dissolution current across the stainless-steel membrane sealing the payload chamber. Circuit design only. The measured actuation trace is reported with the results.

3.7 Animal Study Methods

The in vivo results reported in §4.4 and §4.5 come from two separate porcine studies: a gastric-residence and communication study, and an intestinal-compartment commanded-release study. This subsection states the methods common to both and the parameters specific to each. Both were conducted under the MIT Committee on Animal Care protocol described below.

Animals and allocation. The animals were female Yorkshire swine (*Sus scrofa domestica*). The approved protocol covers a porcine working-weight range of roughly 35–120 kg. The animals used in these two experiments weighed approximately 90 kg. The commanded-release study used a single instrumented device in a clamped intestinal compartment together with one paired, untriggered control. The gastric-residence and

communication study used a single instrumented device. The residence and release experiments were carried out in separate terminal sessions rather than in the same animal.

Anesthesia and administration. Animals were anesthetized with tiletamine/zolazepam (~4–6 mg/kg IM) and xylazine (~2–4 mg/kg IM). For endoscopic work the plane of anesthesia was maintained with isoflurane (2–3% in O₂) delivered by endotracheal intubation, with thermal support and sterile ophthalmic ointment. Devices were introduced endoscopically through an overtube into the stomach. Anatomical placement was gastric for the residence study, accessed endoscopically in a survival experiment context, and a clamped intestinal compartment for the commanded-release study, accessed laparoscopically in a terminal experiment context.

Monitoring and duration. The integrated gastric-residence study for iTARGET observed the device to day 6, a window set by animal-facility scheduling rather than a device limit. Preceding whole-device retention studies for iRES reached far longer. During procedures animals were monitored continuously under anesthesia. In survival and retention settings they were observed at least twice daily for fecal output, abdominal distension, lethargy, inappetence, and weight loss, with radiographic surveillance as needed.

Retrieval. Devices were retrieved endoscopically using the overtube and endoscope, or triggered to disassemble for passage, or retrieved at the time of animal euthanasia.

Success and failure criteria. Outcomes are reported descriptively against qualitative criteria, because each experiment used a single instrumented animal. Outcomes include retention through the observation window, maintenance of the wireless link, payload-barrier integrity, and commanded release confirmed by payload appearance in the target compartment.

Bacterial viability. The payload was *Escherichia coli* Nissle 1917. Released and control samples were serially diluted and plated, and viability was reported as colony-forming units per millilitre (CFU/mL), comparing the triggered device against the untriggered control at the sampled time points (0 and 2 h).

Ethical approval. All procedures were performed under MIT Committee on Animal Care (CAC) protocol 0821-066-24, "Ingestible Transceiver-Actuable Resident Gastrointestinal bioElectronic Therapeutic for Travelers Diarrhea (iTARGET-TD)," with Prof. C. Giovanni Traverso as principal investigator, under DARPA sponsorship with the required Department of Defense (ACURO) animal-use review. The work reported here was carried out at the MIT Division of Comparative Medicine.

Chapter 4 • Results

4.1 Power and Lifetime

Bench discharge of the silver-oxide cells across five constant-current loads produced the lifetime–current fit $t = 11.456I^{-1.371}$ ($R^2 = 0.9977$) shown below. This is a budgeting input over the tested loads, not a measured greater-than-60-day device run. The five discharge loads all sit at or above 500 μA , so the fit is measured only in that band. The intermediate $\approx 29 \mu\text{A}$ point and the $\approx 18 \mu\text{A}$ design operating point lie below the measured range, and any lifetime quoted there is an extrapolation of the power law rather than a measured value, as measuring those operating points directly was infeasible. A single dissolution-current trace characterized commanded actuation in a neutral pH medium. Acidic media performed the operation more quickly, so this trace was selected as the worst-case scenario and used as budgetary input. It reached a peak of 5.06 mA and averaged 1.96 mA over 47 min 16 s. It first showed evidence of release at 9 min 12 s. The trace continued stress-testing the system after release began. The trace measures one instance of a release event.




FIGURE 19. Measured silver-oxide cell lifetime versus constant-current load. Left: the primary discharge curves, cell terminal voltage against time, for the five constant-current loads spanning 500 μA to 20 mA. Right: the resulting lifetimes with the fitted relation $t = 11.456I^{-1.371}$ ($R^2 = 0.9977$), where t is in hours and I in milliamps. The fit is a project-specific budgeting input over the tested loads. It characterizes best-case single-cell performance across the discharge currents in the range of interest. The main takeaway: the lifetime law scales worse than linearly with the *magnitude* of current draw from the cells. Two levers therefore become vital. One suppresses the dynamic range of on-device power-draw signals. The other lowers the DC effective operating current.




FIGURE 20. Commanded dissolution-current trace with trigger and rupture frames.

Nordic Power Profiler Kit 2 current against time for the single worst-case neutral-pH actuation, re-zeroed at the trigger event and shown over the first 30 minutes. The current rises to a 5.06 mA peak, and the foil-rupture spike appears near 9 min. The two frames mark the trigger event ($t = 0:00$) and foil rupture ($t = 9:05$). The video and the current capture were recorded as separate runs, so the frame timing is an illustrative alignment rather than a literal co-recording.

Beyond the constant-current fit, two iRES devices ran through the full device-to-gateway-to-backend telemetry path. Each showed multi-month operation on a single primary-cell stack ($N=1$ each). The logs held received signal strength (RSSI) against relative time, and temperature where the on-board temperature stream reported. The gateway wall-clock carries no experimental meaning here. Every axis is therefore a delta from first contact ($T+0$). Both devices held their link well above the nRF52811 sensitivity floor for the entire run. That floor is -96 dBm, with a ~ 20 dB practical margin at -80 dBm. Early data gaps in both runs were collection-pipeline artifacts, not link losses (see Appendix B).

Serial number A701 (firmware rev 7) operated for roughly 2.9 months with continuous link coverage, and remained operating in a benchtop incubator at the end of the run. A brief endoscopic placement ($\sim 40^\circ\text{C}$) verified it in vivo first. It then moved immediately into a benchtop incubator for the remainder of the run. Its early RSSI and temperature gaps are collection-pipeline artifacts detailed in Appendix B, not device faults. Once the temperature telemetry resumed, the incubator held the temperature stream near 37°C for the rest of the run. The small ($\sim 1^\circ\text{C}$) step changes are shared-incubator setpoint adjustments made by other users and did not affect this experiment.

Serial number A902 (firmware rev 9) operated for roughly 2.2 months across a hybrid in/out-of-animal sequence. It stayed in radio contact throughout. It started as manufacturing-time checks at room temperature. It then went to the animal. The elevated ~42 °C core temperature there appears to be genuine physiological data per the in vivo team. It is noted here but is not a focus of this report. It returned to the bench. It finally sat in an agitated incubator (~39 °C) for the remainder of the run to battery death. Temperature reporting was noticeably less reliable than RSSI. The BLE Health Thermometer Service requires a connected exchange with acknowledgements. RSSI is measurable from advertisement alone, with no connection event. This informed an advertisement-only telemetry posture that removes unnecessary bidirectional traffic for biometric reporting. It retains a connected path only where a bidirectional link is actually wanted. Here, temperature reporting efficacy dropouts follow from non-symmetric link performance. Uplink and downlink directions do not materialize as an identical RSSI metric even with nominally identical hardware on each node. Return loss and antenna tuning on the ingested-device side differ from those of the standard receiving hardware. The residual variance in the RSSI floor is an open question, attributed provisionally to humidity swings (~60–100% RH) rather than to the link itself.

Taken together, this development cycle ended in a stable, in vivo compatible BLE telemetry system. That system ran for multiple months. The conditions reported at the end of each run are directly comparable. They yield instructive inputs for the iTARGET-TD generation design stage.




FIGURE 21. A701 long-duration telemetry over 2.9 months. The device ran iRES, MIT firmware rev 7. The plot uses relative time (T+0 = first contact). A brief in vivo endoscopic verification (~40 °C) came first. The device then ran in a benchtop incubator for the rest of the trace. RSSI (blue) stays well above the nRF52811 sensitivity (–96 dBm) and the practical link floor (–80 dBm) throughout. The early RSSI gap is the collection-pipeline buildout month, not a link loss. Temperature (red) is absent longer for a collection-side reason, not a device fault. The frontend initially parsed only the 9xx temperature-reporting convention, a choice that prioritized the in-animal diagnosis of A902. A per-packet fix to the frontend then recovered the 7xx temperature of A701. Once reporting resumes it holds near 37 °C, with the small shared-incubator setpoint steps annotated. Hourly medians overlay the raw samples. N=1.




FIGURE 22. A902 telemetry to battery death over 2.2 months. The device ran iRES, MIT firmware rev 9. The plot uses relative time (T+0 = first contact, no calendar dates). The device stayed in radio contact across the bench-prototyping, in-animal (~42 °C), post-retrieval-bench, and agitated-incubator (~39 °C) phases. Temperature reporting is visibly spottier than RSSI. The BLE Health Thermometer Service needs a connected, acknowledged exchange, whereas RSSI is recoverable from advertisement alone. That motivates advertisement-only telemetry in the next generation. The residual RSSI-floor variance remains open (suspected humidity, ~60–100% RH). End-to-end RSSI is also not symmetric even with identical hardware, owing to return-loss and tuning differences on the ingested-device side. Hourly medians overlay the raw samples. N=1.

4.2 Command Authentication

The SPECK-128 challenge-response implementation matched its reference vectors, and the shared mobile-and-embedded exchange completed successfully in the integrated application-plus-firmware test. This establishes that a correct, authenticated command is required before the device enters an actuation-capable state. It does not establish full medical-device security or regulatory compliance. In particular, this implementation-consistency test does not evaluate resistance to replay, challenge reuse, key extraction, denial of service, unauthorized pairing, relay, or downgrade attacks, nor the consequences of a compromised smartphone, modified firmware, or random-number failure. Those evaluations were out of scope for the pre-human stage of this work. They are also, by design, not addressed by expanding this bespoke gate. Hardening against replay, relay, downgrade, and the rest is deferred to the Bluetooth Mesh security layer the successor platform adopts (§3.4): a maintained, standardized transport security is the appropriate home for those guarantees, because wireless security is an ongoing maintenance commitment rather than a one-time property, and a self-rolled scheme would have to be maintained indefinitely to stay safe. The SPECK gate is therefore retained only as a portable authorization check for links that lack that standard.




FIGURE 23. SPECK-128 reference-vector agreement and challenge-response exchange.
The test covered a successful application-to-firmware exchange. Verifies protocol correctness across the two endpoints, not a full security evaluation.

4.3 iRES Design Reliability

Assembled devices also ran directly against the media they would have to transmit through. In one bench session, four devices hung in a water-filled beaker on a common support. A smartphone running nRF Connect and a bench nRF development board scanned them concurrently. The bench board logged RSSI and temperature over a serial link. The session compared one MIT-built device (A902) against three TRT-built devices (B106, B107, B128). The two builds differed only in the 3D-printed spacer surrounding the antenna and were otherwise identical, so the session suggests a build-to-build effect from that one variable. Immersion and scan geometry were the same for all four. That comparison made build-to-build differences in link performance visible without requiring a calibrated range. Readings also went through air and through glass at varying distance and orientation. This check is qualitative. The smartphone reports uncalibrated RSSI, and transmitter orientation and beaker position moved by hand between readings. The setup was not an anechoic or otherwise controlled RF environment. It stands as manufacturing go/no-go evidence, not as an antenna-gain measurement.



FIGURE 24. In-medium bench comparison of MIT- and TRT-built devices. Left: the physical arrangement. Four devices hang in a water-filled beaker on a common support. A bench nRF development board logs RSSI and temperature per device over serial. The smartphone that took the photograph also ran the concurrent scan, and its case sits empty on the bench. Right: the smartphone nRF Connect scan of the same four devices from that position. The comparison sets one MIT-built device (A902) against three TRT-built devices (B106, B107, B128) under identical immersion and scan geometry. The two builds differed only in the 3D-printed spacer surrounding the antenna and were otherwise identical. It therefore reads as a qualitative build-to-build comparison rather than a calibrated link measurement. The smartphone reports uncalibrated RSSI, and orientation moved by hand between readings.

4.4 Gastric Retention and Reliability

The capstone retention result was a single fully-integrated iTARGET-TD device demonstrated in one porcine gastric-residence study (N = 1 animal), in which the device communicated continuously and held its payload barrier intact through day 6, confirmed by endoscopy. The six-day observation window was set by animal-facility logistics rather than by any device limit: the terminal endpoint was scheduled around the animal-handling calendar, in this case ahead of a seasonal facility-staffing reduction, and the device was still communicating with an intact barrier when the study ended. Component-level studies preceding this integrated demonstration bound the residence ceiling well above six days. Prior whole-device gastric-retention studies retained the resident geometry for 62 to 100 days across three animals [37], and sealed-payload viability persisted on the order of ten days in the gastric environment. The earlier iRES generation additionally demonstrated in-stomach commanded disassembly of the retention system in vivo, exercising the same electrochemically triggered actuation

mechanism later carried into the iTARGET-TD payload-release chambers. Those iRES retention and disassembly counts and their supporting data belong to the iRES retention work, yet to be reported by a manuscript in preparation (McRae et al.). The author's contribution there was integrative (the electrical subsystem, its interfaces, command availability, data transfer, study co-execution, and the power budget), rather than the retention mechanics. The integrated six-day run therefore composites retention, sustained communication, and barrier integrity that the earlier studies had each shown separately, and its duration reflects scheduling rather than the platform's demonstrated residence capability. This gastric-residence study is one of the two complementary in vivo experiments of §5.1. The commanded in vivo release is the separate intestinal-compartment study of §4.5.

Reliability testing also surfaced an environmental failure absent from short acute placements. In a separate placement, communication degraded on an approximately 12-hour timescale, with retrieval implicating moisture and condensation. That observation drove the desiccant and sealing mitigations described in the methods in order to counteract diffusive water vapor ingress through the polycarbonate chassis of the iTARGET-TD device.




FIGURE 25. Payload barrier intact at day 0 and day 6 in vivo (gastric-residence study, N = 1). Endoscopic images from the gastric-residence study show day 0 (left) and day 6 (right). The barrier is intact at day 6. Radio communication showed the integrated resident device operational at day 6. Endoscopic extraction from the animal then checked seal integrity. It also checked maintained bacterial viability and triggerability of a sealed device that had resided in vivo. This gastric-residence study is distinct from the commanded in vivo release of §4.5, which was a separate intestinal-compartment experiment.




FIGURE 26. Bench triggering and retrieval of the retained device (gastric-residence study). The image shows a viable payload deployed on command from a device retained in vivo for 6 days and then extracted. This triggering and viability check ran on the benchtop after extraction, because a large-animal model does not allow reliable endoscopic *catch* of an ejected payload for viability testing. Commanded release observed directly in vivo is reported separately in §4.5.

4.5 Remote In Vivo Communication and Bacterial Release

Commanded release was made directly observable in vivo with a fully-integrated iTARGET-TD device in an intestinal-compartment study. Retrofitting a basket-style release into the gastric retention geometry proved difficult, so wireless release was not demonstrated in the stomach. The experiment instead placed the integrated device in an isolated intestinal compartment, where an ejected payload could be captured and assayed. A wireless command released a viable bacterial payload on the same 2.4 GHz BLE control path used throughout. Against a control device under identical placement, viable colonies were recovered only in the test condition two hours after the commanded release, and were absent at the 0 h baseline and in the control. That is commanded in vivo payload deployment with a viability-confirmed output. The release was attempted once and succeeded once, in a single animal (N = 1). Devices reached the animal only after clearing pre-deployment quality control. Earlier fabrication builds that failed QC were caught and rejected on the bench and were never implanted, because the relevant defects were detectable beforehand.

The communication link held through tissue only because the antenna matched the in-body environment. In the matched-versus-unmatched comparison, the matched design recovered a usable link where the unmatched antenna lost the signal. A staged in-body

sensitivity measurement reached approximately -102 dBm, quantifying the control-link budget under tissue RF shielding and detuning. The available artifacts do not join authentication and release in one synchronized trace. The experiment does not establish therapeutic efficacy against Traveler's diarrhea, and it was separate from the long-duration gastric-residence study of §4.4.




FIGURE 27. Wirelessly triggered release of a bacterial payload in vivo (intestinal-compartment study, iTARGET-TD). An integrated device (top left) was placed in an isolated intestinal compartment of a pig alongside a control device (without a trigger system, otherwise identical) in its own compartment. The integrated device maintained active BLE communication while the intestine was returned to the abdominal cavity (top right). After the commanded release, viable colonies were recovered only in the test condition at 2 h. The control device and the 0 h baseline yielded none (bottom: recovered plates and CFU/mL). This study is complementary to, and separate from, the gastric-residence study of §4.4.




FIGURE 28. Through-tissue control-link budget under gastric shielding. The same device instance measured on the bench at roughly -55 dBm degrades to about -75 dBm after in-body frequency shifting and to approximately -102 dBm once tissue and surgical-clamp shielding are added. This staged measurement quantifies the in-body sensitivity the matched antenna design was built to meet. The table below records the full cascade and the pre-flight acceptance criteria derived from it.

Table 5. Through-tissue control-link budget, measured end-to-end on a single device instance. Because every loss term is captured empirically on one device rather than stacked from independent estimates, this budget is authoritative.

Stage	RSSI (dBm)	Δ (dB)
Transmit reference (device)	+ 0	—
Benchtop received, same device instance (non-tuned matching network device-side)	- 55	- 55
After in-body frequency-shift detuning	- 75	- 20
After tissue + surgical-clamp path loss (gastric shielding)	- 102	- 27
Receiver sensitivity floor (absolute)	- 102	—

The most-shielded gastric configuration lands on the -102 dBm receiver floor, zero margin at the worst case. That floor is absolute, so the measured in-body degradation sets the pre-deployment acceptance policy directly: a benchtop RSSI of at least -55 dBm is strictly required (any weaker device falls beneath the floor once shielded), and the acceptance criterion is tightened to -35 dBm so that each accepted device retains positive working headroom after the full in-body loss cascade. A -35 dBm bench reading predicts -82 dBm after the same 47 dB of in-body loss, approximately 20 dB above the -102 dBm floor.

Chapter 5 · Conclusions, Discussion, and Future Work

5.1 Conclusions

This thesis set out a single control problem: commanded release of a bacterial therapeutic from a gastric-resident device. It treated power, in-body communication, command authentication, retention, and manufacturing as one coupled electronics problem rather than five independent requirements. Two design–build–test generations reduced that problem to practice. The first, iRES, established the power budget, the app-mediated 2.4 GHz BLE link, and a manufacturable capsule-scale flexible circuit. The second, iTARGET-TD, carried those foundations into modular electronics with authenticated command handling and integrated payload-release hardware.

One constraint organized the rest. At swallowable scale the cell's lifetime falls faster than linearly with the *magnitude* of current drawn from it, so every subsystem was priced in current rather than in average power alone. That single law propagated into the decisions that define the platform. Telemetry reads received signal strength from advertisements alone, rather than from a connected exchange the lifetime law charges for super-linearly. An authenticated wakeup rides the radio's existing scan cycle instead of a separate high-power listening mode. A local capacitor bank is sized against a deep-life internal-resistance estimate, so that late-life radio bursts do not brown out while low-current membrane dissolution still succeeds. The subsystem innovations are therefore corollaries of one power argument, not independent features.

The work addressed each requirement, with measured evidence for power, communication, authentication, retention, and actuation. A measured battery lifetime–current law makes current draw budgetable, complemented by two single-device demonstrations of continuous multi-month (~2–3 month) telemetry operation on one primary-cell stack. A matched in-body BLE control link is established where an unmatched antenna loses the signal. A SPECK-128 challenge-response admits actuation only on an authenticated command. Multi-day gastric residence held an intact payload barrier observed at day 6. A commanded in vivo release deployed a viable bacterial payload. Together these establish the principal subsystems and show commanded in vivo release, authenticated command handling (on the bench), and multi-day gastric residence in complementary experiments. They do not yet demonstrate all of these functions in a single integrated in vivo study.

These demonstrations are complementary rather than one continuous experiment. No single synchronized trace captures authentication and release alongside therapeutic efficacy against Traveler's diarrhea. The lifetime law with its single actuation trace does not establish a 60-day device run or six in vivo doses. The contribution is therefore the electronics platform, its design specifications, and its so-far measured in vivo behaviors.

5.2 Discussion

The long-lived telemetry runs of §4.1 are the empirical complement to the bench lifetime–current law reported there. A701 and A902 each sustained an always-available BLE link on a single primary-cell stack for roughly two-to-three months (N=1). That is the behavior the current-budget law was meant to secure. Read together, they support the central power claim of this thesis. For a swallowable primary cell, how much charge the cell can usefully deliver depends on how the load draws it, so the current profile governs the achievable lifetime. They also place the iTARGET-TD power-management strategy against a substantial prior literature.

The measured lifetime exponent, $t \propto I^{-1.37}$, is a device-specific case of a long-established super-linear relationship. That relationship links discharge current to deliverable capacity, and Peukert's law first captured it. Modern critical reviews confirm that the Peukert exponent is chemistry- and rate-dependent and caution explicitly against extrapolating a single datasheet capacity across load profiles [38] [39]. The iRES budget is therefore expressed as a current-*allocation* law rather than a milliamp-hour figure. That framing follows the conservative reading of the Peukert literature.

This thesis emphasizes one end-of-life failure mode: high-current radio bursts brown out first while low-current foil dissolution still succeeds. That mode corresponds to the battery power-capability literature, in which the cell's rising internal impedance sets the maximum deliverable pulse, and an impedance model rather than residual charge predicts it [40]. The iTARGET-TD brownout model drives the measured transmit-burst current into a deep-life internal resistance, against the regulator brownout threshold. That model applies exactly that power-capability framing at capsule scale. It is specialized to the silver-oxide chemistry and to the burst profile of a BLE transmit event.

The mitigation is likewise a constrained corner of an established design space. A local capacitor buffers a high-rate pulsed load. The source supplies only the slow average current through its internal resistance. The capacitor delivers the fast burst charge. That is the operating principle of battery–supercapacitor hybrid architectures for high-rate pulsed loads [41] and for extending the service life of battery-powered wireless sensor nodes [42]. What distinguishes the iTARGET-TD realization is scale and constraint. Those systems add a supercapacitor and an active bidirectional converter, prohibitive in both

volume and bill-of-materials inside a capsule. In iTARGET-TD a small ceramic bypass bank is instead sized against a deep-life internal-resistance estimate. It trades board area rather than scarce chemistry volume for burst headroom late in life. The contribution is thus not the buffering principle but its reduction to a capsule-manufacturable form.

Finally, the A902 telemetry observation feeds back into the same budget. Lifetime penalizes current *magnitude* super-linearly, so how often a telemetry scheme forces high-current connected radio events dominates its cost. RSSI is recoverable from advertisements alone, while the Health Thermometer Service requires a connected, acknowledged exchange. That finding makes advertisement-only telemetry a power-management decision as much as a protocol one. It removes bidirectional traffic that the current-budget law charges for super-linearly.

The literature above supplies the mechanisms and the prior art for interpreting these device-specific measurements. It also frames the future work below.

5.2.1 Limitations

These comparisons position the strategy rather than validate it beyond the present evidence. The lifetime law is a five-load bench fit. The deep-life internal resistance is an estimate from Power Profiler Kit 2 discharge data. The bypass-bank result is a simulation against that estimate. The multi-month runs are single devices. The residence, barrier-integrity, and commanded-release results come from complementary experiments rather than one synchronized integrated trace: no single synchronized trace captures authentication and release alongside therapeutic efficacy against Traveler's diarrhea, and the lifetime law with its single actuation trace does not establish a 60-day device run or six in vivo doses. The SPECK-128 challenge-response was verified for implementation consistency only. A full wireless-security threat evaluation (replay, challenge reuse, key extraction, denial of service, unauthorized pairing, relay, and downgrade) was out of scope for the pre-human stage of this work.

5.3 Future Work

The open work follows directly from those boundaries. Each item below resolves a specific dependency the present evidence leaves unmet. Several are necessarily speculative, and are offered as the natural next steps a follow-on program would take rather than as completed designs.

- **Integrated authenticated gastric-resident release.** The single largest gap is the one the limitations name: no experiment yet joins authentication, retention, and release in one synchronized in vivo trace. The immediate next study is a single gastric-resident device that receives an authenticated command and releases its payload in place, with the authentication exchange and the actuation current captured on one timeline. This is the experiment that would convert the present complementary evidence into an integrated demonstration.
- **Extended residence and actuation.** Tie the multi-day communication and the day-6 barrier observation to one identified device and protocol, and extend testing toward the greater-than-60-day residence target. This has so far been prohibitively expensive to pursue on the program timeline, but it is a scheduling and resourcing problem rather than an unsolved design one.
- **Replicated reliability.** The multi-month telemetry rests on two single-device runs ($N = 1$ each). A powered cohort (multiple devices per configuration, reported with success rates and failure modes) would move the lifetime-law validation from an illustrative pair to a statistical result, and would let the build-to-build reliability comparison of §4.2 isolate its variable properly.
- **Cybersecurity threat modeling.** The SPECK-128 exchange was verified for implementation consistency only. A formal threat model (replay, challenge reuse, key extraction, denial of service, unauthorized pairing, relay, downgrade, compromised-phone, firmware-modification, and random-number-failure cases), with matched negative-case testing, is required before any security claim beyond authenticated command handling. This work would also fix the protocol parameters (challenge and key length, nonce generation, replay handling, key provisioning, command expiration) at a level a reviewer could audit.
- **Dissolution-product toxicology.** The stainless-steel membrane releases uncharacterized electrochemical products that may include chromium and nickel. A dedicated toxicological evaluation of those products, at the masses actually released per deployment, is a prerequisite for human translation.
- **Sealing and sterility.** The payload must remain sealed and sterile through gastric residence and then release on command. A characterized sealing process, with sterility validation and shelf-life data at body temperature, is needed to carry the device from a research prototype toward a manufacturable product.
- **Release-dose reproducibility.** Commanded release was shown as a viability-preserving event, not as a metered dose. Quantifying the released bacterial dose, and its reproducibility across devices and commands, is required before the payload can be tied to a therapeutic outcome.

- **Therapeutic efficacy.** Pair commanded release with outcome studies for Traveler's diarrhea. Biosensor chemistry and large-animal efficacy work lie outside this thesis, but they are the studies that would establish clinical benefit rather than device function.
- **Dead-battery and failed-release behavior.** The device must be safe when it does not work as intended. Defined behavior for a battery that dies before release, and for an actuation command that fails to open the chamber, together with the retrieval or passage path in those cases, is a safety requirement rather than an optional study.
- **Retrieval and passage strategy.** The end-of-deployment path (self-disassembly, retrieval, or safe passage) needs an explicit, tested strategy, including the case where disassembly does not complete.
- **Manufacturing quality controls.** Carrying the build from hand-assembled prototypes to a controlled process requires stated yields, acceptance testing, and defined failure modes across build sites.

5.3.1 Capabilities the platform unblocks

The items above close the gaps in what has been demonstrated. A second class of future work instead builds *outward* from what the platform now makes possible. These directions are more speculative than the validation studies above (they are sketches of what a follow-on program could pursue rather than designs this thesis evaluates), but each follows directly from the architecture already reduced to practice: six independently addressable chambers, an authenticated per-command actuation path, a duty-cycled biometric sensing front-end, and a power budget that can afford all of them.


- **Programmable multi-agent dosing.** The demonstrated device carried a single bacterial payload, but the six-chamber architecture and per-chamber authenticated actuation together support a very different mode of use: a single resident device delivering a *scheduled sequence* of distinct agents, one chamber released at a time across a deployment. That reframes the capsule from a single-shot actuator into a programmable, in-body pharmacy. The natural payloads are small-molecule drugs relevant to a field-deployed setting: a hemostatic agent for internal bleeding, a radioprotective agent for radiation exposure, and stimulants or other agents for alertness and sleep management.
- **Closed-loop sense-and-release.** The biometric front-end (temperature and photoplethysmography) and an added gastric-chemistry sensor turn commanded release into *closed-loop* release: a measured physiological trigger drives actuation without an external command. Two demonstrations follow immediately from the payload set above: an optical bleeding sensor paired with a hemostatic chamber, and

a gastric pH sensor paired with a pH-conditioned release schedule, each closing the loop between the sensing hardware already on the board and the actuation hardware already validated. The optical-sensing half of that loop is not hypothetical: an ingestible reflectance capsule from the Traverso Lab already detects a gastrointestinal blood-flow pathology in vivo by pulsing light-emitting diodes and reading reflected light onboard [43], which is the sensing front end a closed-loop hemostatic device would pair with an authenticated release.


- **Survival studies with in vivo pharmacokinetic confirmation.** Moving from terminal to survival experiments, with serial blood sampling over days, would let measured drug levels in blood confirm and quantify each commanded release directly. That is the most efficient route to the metered-dose and release-reproducibility evidence the validation list calls for, and it is the bridge from device function to a therapeutic-outcome study. The same survival model supports repeated dosing from successive chambers in one animal over a multi-day residence.
- **Additional actuation modalities.** The authenticated command path and the late-life power budget are not specific to chemical release. The same resident platform could carry a non-pharmacological intervention (gastric electrostimulation or mechanostimulation being the immediate examples), actuated over the same authenticated link and priced in the same current budget, extending the device from a drug-delivery vehicle toward a more general commanded gastrointestinal therapeutic.

Power, communication, authentication, retention, and actuation have each been reduced to practice and demonstrated in complementary experiments. Showing these requirements are met together at once, in one integrated in vivo study, is the central open goal the platform is built to support. Beyond it, the multi-agent, closed-loop, and survival-study directions above are the capabilities that integrated platform was designed to reach.


Supplemental Figures




SUPPLEMENTARY FIGURE S1. Power-delivery board. Six 4.7 μF X5R ceramics sit in parallel across the cell rail. They form the bypass bank of §2.1 and §3.1, sized against the transmit-burst load. A TLV5518 regulator (used as a load switch) gates a switched rail out to the adapter connector under a dedicated enable line. DF40C-10DP board-to-board connectors carry both rails into the stack. The electrode and indicator LED nets are also brought out here.




SUPPLEMENTARY FIGURE S2. MCU board. Left: the net and connector map, with J1 to the module board and J2 to the antenna board. Those connectors carry the SWD, QSPI, trace, and RF nets. Right: the nRF5340 with its DC/DC regulator inductors, decoupling, and 32 MHz crystal. The antenna matching network feeds the 50 Ω RF port that §3.3 traces to the chip antenna.



SUPPLEMENTARY FIGURE S3. Biometric module board. An LMT70 temperature sensor, with a TPS61099 boost converter supplying the LED rail. TLV55xx load switches gate the sensor rails. The board also carries a MAX30101 photoplethysmography and heart-rate front end. These are the sensing loads whose duty-cycled current the §2.1 budget has to absorb.



SUPPLEMENTARY FIGURE S4. Antenna board. The 50 Ω RF port from the MCU board reaches a 2450AT07A0100 chip antenna through the board-to-board connector. The discrete matching footprint (R1, C2, C3) carries a do-not-populate mark. Reproduced in §3.3 alongside the matching network it serves.



SUPPLEMENTARY FIGURE S5. Instrument-level discharge capture. A silver-oxide cell discharged on a Keysight N6705C with an N6781A module. The mode was battery-profiler, primary-battery configuration. The load was a constant $-500 \mu\text{A}$ with a 1.260 s pulse and 5.000 ms rest. The capacity rating was 28.000 mAh. The run reached 14.185 mAh depleted at 1.594 V over roughly 1 d 06 h. This is one of the five constant-current loads behind the $t = 11.456I^{-1.371}$ fit of §4.1. Its agreement with that law is a constituent measurement of the fit rather than an independent validation of it.

Appendix A · Conservative RF Incident-Power Bound

A worst-case incident-power check supports the chosen transmit power. The check treats the ≈ 9 mm capsule–tissue interface (area ≈ 0.64 cm²) as intercepting the entire radiated output. That bounds the peak incident power density at ≈ 3.9 mW/cm² (≈ 39 W/m²) at 4 dBm. The bound is an overestimate, since a sub-centimeter in-body antenna couples far less than its full output through that face. The radio transmits only briefly per scan. The six-minute averaging windows of the IEEE and ICNIRP limits therefore see this peak scaled by the transmit duty cycle. Even at the research-mode ~ 30 s scan cadence that duty cycle is of order 10^{-5} . That places the averaged incident power near 10^{-3} W/m² or below. These are order-of-magnitude design bounds, not a measured SAR or a compliance determination.

Appendix B · Telemetry Collection-Pipeline

Notes

The sparse coverage common to the first month of both runs is not a device fault. It spans the buildout of the collection pipeline itself. The telemetry backend moved from a local single-device prototype to a shared database and transport pipeline. That pipeline supported multiple devices. During that month a single shared laptop served as the collection front-end. It traveled between the animal facility and the benchtop incubator area several times until the distributed pipeline came online. The outages were Windows updates and OneDrive filling local laptop storage, not lost radio links.

The prolonged temperature data gap in the A701 run has a separate cause on the collection side rather than on the device. The 7xx and 9xx firmware series encode temperature with different reporting conventions. The collection frontend parsed only the 9xx convention. That choice dates to the delivery of A902 to the animal. Diagnosing the elevated core temperature of A902 and the composite system anomalies took priority. A701 meanwhile sat safely in the incubator, where its missing temperature stream was not critical. A per-packet fix to the frontend then recovered the A701 temperature telemetry.

References

- [1] Ballinger, Ian. *thesis_dist: Analysis and Simulation Source Code Supporting the Thesis "Secure Wireless Control and Remotely Triggered Actuation of a Gastric-Resident Ingestible Bioelectronic Capsule."* GitHub, 2026. Accessed 13 Aug. 2026. https://github.com/IanBallinger/thesis_dist
- [2] Sonnenborn, U. "Escherichia coli strain Nissle 1917 – from bench to bedside and back: history of a special Escherichia coli strain with probiotic properties." *FEMS Microbiology Letters*, vol. 363, no. 19, 2016, fnw212. <https://doi.org/10.1093/femsle/fnw212>
- [3] Riglar, D. T., and P. A. Silver. "Engineering bacteria for diagnostic and therapeutic applications." *Nature Reviews Microbiology*, vol. 16, no. 4, 2018, pp. 214–225. <https://doi.org/10.1038/nrmicro.2017.172>
- [4] Jimenez, M., et al. "Microbial therapeutics: New opportunities for drug delivery." *Journal of Experimental Medicine*, vol. 216, no. 5, 2019, pp. 1005–1009. <https://doi.org/10.1084/jem.20190609>
- [5] Jimenez, M., et al. "Synthetic extremophiles via species-specific formulations improve microbial therapeutics." *Nature Materials*, vol. 23, 2024, pp. 1436–1443. <https://doi.org/10.1038/s41563-024-01937-6>
- [6] Traverso, C. G., et al. "Oral Pill Gastric Retention and Disassembly Designs." International Patent Application WO 2025/071738 A2, PCT/US2024/038763, 3 Apr. 2025. <https://patents.google.com/patent/WO2025071738A2/en>
- [7] Steiger, C., et al. "Ingestible electronics for diagnostics and therapy." *Nature Reviews Materials*, vol. 4, no. 2, 2019, pp. 83–98. <https://doi.org/10.1038/s41578-018-0070-3>
- [8] Farra, R., et al. "First-in-human testing of a wirelessly controlled drug delivery microchip." *Science Translational Medicine*, vol. 4, no. 122, 2012, 122ra21. <https://doi.org/10.1126/scitranslmed.3003276>
- [9] Santini, J. T., Jr., et al. "A controlled-release microchip." *Nature*, vol. 397, no. 6717, 1999, pp. 335–338. <https://doi.org/10.1038/16898>
- [10] Mimee, M., et al. "An ingestible bacterial-electronic system to monitor gastrointestinal health." *Science*, vol. 360, no. 6391, 2018, pp. 915–918. <https://doi.org/10.1126/science.aas9315>
- [11] Charbonneau, M. R., et al. "Developing a new class of engineered live bacterial therapeutics to treat human diseases." *Nature Communications*, vol. 11, no. 1, 2020, 1738. <https://doi.org/10.1038/s41467-020-15508-1>
- [12] Bellinger, A. M., et al. "Oral, ultra-long-lasting drug delivery: Application toward malaria elimination goals." *Science Translational Medicine*, vol. 8, no. 365, 2016. <https://doi.org/10.1126/scitranslmed.aag2374>
- [13] Iddan, G., et al. "Wireless capsule endoscopy." *Nature*, vol. 405, no. 6785, 2000, p. 417. <https://doi.org/10.1038/35013140>
- [14] Gabriel, S., et al. "The dielectric properties of biological tissues: III. Parametric models for the dielectric spectrum of tissues." *Physics in Medicine and Biology*, vol. 41, no. 11, 1996, pp. 2271–2293. <https://doi.org/10.1088/0031-9155/41/11/003>
- [15] Kiourti, A., and K. S. Nikita. "A Review of In-Body Biotelemetry Devices: Implantables, Ingestibles, and Injectables." *IEEE Transactions on Biomedical Engineering*, vol. 64, no. 7, 2017, pp. 1422–1430. <https://doi.org/10.1109/TBME.2017.2668612>
- [16] Cil, E., et al. "On the Use of Impedance Detuning for Gastrointestinal Segment Tracking of Ingestible Capsules." *IEEE Transactions on Antennas and Propagation*, vol. 71, no. 2, 2023, pp. 1977–1981. <https://doi.org/10.1109/TAP.2023.3234374>

- [17] Jeon, Y., et al. "Secure and Stable Wireless Communication for an Ingestible Device." *2023 45th Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*, 2023, pp. 1–6. <https://doi.org/10.1109/EMBC40787.2023.10340731>
- [18] Beaulieu, R., et al. "The SIMON and SPECK Families of Lightweight Block Ciphers." *Cryptology ePrint Archive*, no. 2013/404, 2013. <https://eprint.iacr.org/2013/404>
- [19] Beaulieu, R., et al. "The SIMON and SPECK lightweight block ciphers." *Proceedings of the 52nd Annual Design Automation Conference (DAC '15)*, 2015, pp. 1–6. <https://doi.org/10.1145/2744769.2747946>
- [20] Camara, C., et al. "Security and privacy issues in implantable medical devices: A comprehensive survey." *Journal of Biomedical Informatics*, vol. 55, 2015, pp. 272–289. <https://doi.org/10.1016/j.jbi.2015.04.007>
- [21] Jang, S., et al. "Post Quantum Secure Communication Protocol with Ultra Low-Power Hardware Solution for Ingestible Medical Device." *2025 47th Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*, 2025, pp. 1–7. <https://doi.org/10.1109/EMBC58623.2025.11254204>
- [22] IEEE. "IEEE Standard for Safety Levels with Respect to Human Exposure to Electric, Magnetic, and Electromagnetic Fields, 0 Hz to 300 GHz." IEEE Std C95.1-2019, IEEE, 2019. <https://doi.org/10.1109/IEEESTD.2019.8859679>
- [23] International Commission on Non-Ionizing Radiation Protection (ICNIRP). "Guidelines for Limiting Exposure to Electromagnetic Fields (100 kHz to 300 GHz)." *Health Physics*, vol. 118, no. 5, 2020, pp. 483–524. <https://doi.org/10.1097/HP.0000000000001210>
- [24] U.S. Food and Drug Administration. "Cybersecurity in Medical Devices: Quality Management System Considerations and Content of Premarket Submissions." Guidance for Industry and FDA Staff, Feb. 2026. <https://www.fda.gov/regulatory-information/search-fda-guidance-documents/cybersecurity-medical-devices-quality-management-system-considerations-and-content-premarket>
- [25] Vanhoestenbergh, A., and N. Donaldson. "Corrosion of silicon integrated circuits and lifetime predictions in implantable electronic devices." *Journal of Neural Engineering*, vol. 10, no. 3, 2013, 031002. <https://doi.org/10.1088/1741-2560/10/3/031002>
- [26] Girard, G., et al. "Experimental characterization of rolled annealed copper film used in flexible printed circuit boards: Identification of the elastic-plastic and low-cycle fatigue behaviors." *Microelectronics Reliability*, vol. 115, 2020, 113976. <https://doi.org/10.1016/j.microrel.2020.113976>
- [27] Nordic Semiconductor. "nRF52811 System on Chip." *Nordic Semiconductor*, product specification summary. Accessed 13 Aug. 2026. <https://www.nordicsemi.com/Products/nRF52811>
- [28] Nordic Semiconductor. "nRF5340 System-on-Chip." *Nordic Semiconductor*, product page. Accessed 13 Aug. 2026. <https://www.nordicsemi.com/Products/nRF5340>
- [29] Zephyr Project. *Zephyr RTOS Documentation*. Linux Foundation. Accessed 13 Aug. 2026. <https://docs.zephyrproject.org/>
- [30] Bluetooth SIG. *Bluetooth Core Specification*. Version 6.1, vol. 6, parts A–B, Bluetooth SIG, 2025. <https://www.bluetooth.com/wp-content/uploads/Files/Specification/HTML/Core-61/out/en/low-energy-controller/physical-layer-specification.html>
- [31] React Native. "About the New Architecture." *React Native*, official documentation (TurboModules and JSI); the New Architecture has been available for opt-in since React Native 0.68 (2022) and enabled by default since 0.76 (2024). Accessed 13 Aug. 2026. <https://reactnative.dev/architecture/landing-page>
- [32] Traverso, C. G., et al. "Low Power Battery and Signaling Systems." International Patent Application WO 2025/024274 A2, PCT/US2024/038741, 30 Jan. 2025. <https://patents.google.com/patent/WO2025024274A2/en>

- [33] Fortunato, M. "Temperature and Voltage Variation of Ceramic Capacitors, or Why Your 4.7 μF Capacitor Becomes a 0.33 μF Capacitor." *Analog Devices*, technical article. Accessed 13 Aug. 2026. <https://www.analog.com/en/resources/technical-articles/temperature-and-voltage-variation-ceramic-capacitor.html>
- [34] Pozar, D. M. *Microwave Engineering*. 4th ed., John Wiley & Sons, 2011. <https://openlibrary.org/isbn/9780470631553>
- [35] *Online Smith Chart Tool*. onlinesmithchart.com, interactive web application. Accessed 13 Aug. 2026. https://onlinesmithchart.com/?circuit=blackBox_72_-26_20__seriesInd_2.2_nH_5_0__shortedCap_0.7_pF_15_0_0__seriesInd_2.2_nH_5_0__transmissionLine_6_mm_50_50_1__shortedCap_1.78_pF_6_0_0__seriesInd_2.48_nH_5_0__seriesCap_4.8_pF_2_0_0&fSpan=100&zo=17.75&vswrCircles=1.5_3
- [36] Analog Devices. "RF Impedance Matching Calculator." *Analog Devices*, interactive design tools. Accessed 13 Aug. 2026. <https://www.analog.com/en/resources/interactive-design-tools/rf-impedance-matching-calculator.html>
- [37] McRae, James C. *Ingestible Devices for Continuous Health Monitoring and Autonomous Drug Delivery*. 2025. Massachusetts Institute of Technology, PhD dissertation. *Dspace@MIT*, not yet publicly available. <https://dspace.mit.edu/search?spc.page=1&query=&configuration=site&f.author=McRae,%20James,query>
- [38] Doerffel, D., and S. Abu Sharkh. "A critical review of using the Peukert equation for determining the remaining capacity of lead-acid and lithium-ion batteries." *Journal of Power Sources*, vol. 155, no. 2, 2006, pp. 395–400. <https://doi.org/10.1016/j.jpowsour.2005.04.030>
- [39] Galushkin, N. E., et al. "A Critical Study of Using the Peukert Equation and Its Generalizations for Determining the Remaining Capacity of Lithium-Ion Batteries." *Applied Sciences*, vol. 10, no. 16, 2020, 5518. <https://doi.org/10.3390/app10165518>
- [40] Lee, S., et al. "Maximum pulse current estimation for high accuracy power capability prediction of a Li-Ion battery." *Microelectronics Reliability*, vol. 55, no. 3–4, 2015, pp. 572–581. <https://doi.org/10.1016/j.microrel.2014.12.016>
- [41] Shin, D., et al. "Battery-supercapacitor hybrid system for high-rate pulsed load applications." *Design, Automation & Test in Europe (DATE)*, 2011. <https://doi.org/10.1109/DATE.2011.5763295>
- [42] Ongaro, F., et al. "Li-Ion Battery-Supercapacitor Hybrid Storage System for a Long Lifetime, Photovoltaic-Based Wireless Sensor Network." *IEEE Transactions on Power Electronics*, vol. 27, no. 9, 2012, pp. 3944–3952. <https://doi.org/10.1109/TPEL.2012.2189022>
- [43] Chen, J., et al. "An Ingestible Capsule for Luminance-Based Diagnosis of Mesenteric Ischemia." *Science Robotics*, vol. 10, no. 107, 2025, article eadx1367. <https://doi.org/10.1126/scirobotics.adx1367>