

Ultrasonic Cochlear Excitation — Device Scoping

Where the proposal lands in the functional chain, what it would have to be, and what would have to be shown

Matthew "Bear" Schulz

Independent researcher

Belmont, California, United States

ORCID: [0009-0000-0956-5743](https://orcid.org/0009-0000-0956-5743)

Licence: Creative Commons Attribution 4.0 International (CC BY 4.0)

Preprint · 27 August 2026

Keywords: focused ultrasound · cochlear implant · acoustic output limits · IEC 60601-2-37 · implanted motes · artificial basilar membrane · vestibular schwannoma

Companion papers: *Closing the Loop Behind the Cochlea* and *The Auditory Pathway as a General-Purpose Data Channel*

0. What this document is, and what it is not

A falsification test says whether an idea is wrong. It does not say what would be built if the idea were right. A paper that proposes ultrasonic excitation without naming the device leaves that gap, and this document fills it: the functional placement, the physical envelope the regulations impose, the device that could exist inside that envelope, and the sequence of evidence required before anyone puts one on a person.

Nothing here is a design. No transducer has been modelled, no field simulated, no phantom built. This is the scoping stage that precedes design, and its main product is a list of the things that must be true.

Confidence, stated up front. The regulatory limits in §3 are quoted from primary sources and are reliable. The physical calculations follow from those limits and standard acoustics. The biological claims — whether spiral ganglion neurons respond to focused ultrasound at these intensities, and with what selectivity — are **unestablished** and are the reason this is scoping rather than specification.

1. Where ultrasonic excitation lands in the functional chain

The eight-stage decomposition places every hearing device on the same map. An ultrasonic excitation route substitutes for **stage 5, Delivery**, and possibly part of **stage 6, Transduction**. It does not touch stages 1, 2 or 4 directly.

But substituting a stage changes what the stages around it can do, and that is where the interesting consequences are.

Table 1.

stage	present practice (electrical)	with ultrasonic delivery	consequence
1 Acquisition	microphone	unchanged	none
2 Analysis	filterbank, 12-22 channels	filterbank, channel count set by achievable focal separation	more channels only if stage 5 delivers them
3 Mapping	channel → electrode	channel → focal coordinate	mapping becomes continuous rather than discrete
4 Encoding	envelope, some fine structure	unchanged in principle	none directly
5 Delivery	current from a fixed electrode	steered acoustic focus	the substitution
6 Transduction	direct depolarisation of ganglion cells	mechanism unestablished	the open question
7 Calibration	per-electrode threshold and comfort	per-focal-point, and the focus can be moved	calibration gains a spatial degree of freedom it did not have
8 Feedback	evoked response via the same electrodes	evoked response, plus pulse-echo from the same aperture	the aperture can image what it stimulates

Two of those consequences deserve to be stated as claims rather than table cells.

The channel count is the point, and it is conditional on a number nobody should assume. Electrical stimulus spreads widely through perilymph because the intracochlear fluids are electrolytes, and the resulting channel interaction constrains frequency selectivity and speech recognition. That much is established across independent sources. What is *not* a fixed quantity is the spread distance. It is characterised as a space constant Z in an exponential decay fitted per patient from evoked-response measurements, and device patents note directly that patient-specific values improve on assumed ones.

The advantage therefore has to be stated as a function rather than a figure:

Table 2.

electrical space constant	independent electrical sites	vs 1.5 mm acoustic focus
1.0 mm	10.0	1.3×
1.5 mm	6.7	2.0×
2.0 mm	5.0	2.7×
3.0 mm	3.3	4.0×
4.0 mm	2.5	5.3×

At 3-4 mm the case for a device programme is strong. At 1.5 mm it is a factor of two, which does not justify one. At 1.0 mm the argument is dead. Measuring the distribution of this parameter across real recipients therefore precedes everything else in §5 — it is cheaper than any other stage, it uses data implant manufacturers already collect at fitting, and it decides whether the rest of this document has a reason to exist.

The aperture is bidirectional, and that is not a minor detail. An electrode stimulates and can record its own evoked response — that is the closed loop this paper argues for. An acoustic aperture can do the same *and* form an image of the structure it is stimulating, from the same hardware, in the same session. Stage 8 stops being an electrical measurement of a neural response and becomes an electrical measurement plus a geometric one. For a population whose anatomy is being deformed by a growing lesion, a device that can see the structure it is driving is a different proposition from one that cannot.

And the problem inverts. A cochlear implant is a delivery device that happens to carry telemetry. An ultrasonic aperture is an imaging device that happens to be able to stimulate. Those are different products with different regulatory paths, different failure modes, and different first indications — and the imaging one may well be the one that reaches a patient first, because it makes no therapeutic claim.

2. The device that could exist

Specified only to the level the physics constrains. Every number is a starting point for design, not an output of it.

Carrier frequency: 0.5-1.5 MHz. Below 0.5 MHz the focal spot exceeds 3 mm and the channel-count argument evaporates. Above 1.5 MHz attenuation through bone rises steeply and focal depth of field shortens.

Aperture siting: mastoid or ear-canal, not the skull vault. The cochlea sits within the temporal bone. Transcranial delivery through the vault costs roughly 28 dB at 1 MHz through 7 mm of bone; a mastoid path avoids the vault entirely. This is the geometry a bone-anchored hearing device already uses, with an established surgical approach, established osseointegration and an existing regulatory precedent for the mounting.

Array: phased, steerable in two axes. A single fixed focus addresses one cochlear location and is useless. The channel-count argument requires electronic steering across the cochlear length. Element count, pitch and aperture follow from the steering range and grating-lobe constraints and are a design problem, not a scoping one.

Coupling: the hard practical problem. Ultrasound needs continuous acoustic contact. A worn device needs coupling that survives skin movement, temperature and perspiration across a full waking day. This is more likely to determine whether the concept is wearable than any of the acoustics.

Duty cycle: waking hours only, and the reason is instructive. Cochlear implant sound processors are removed at night. This is not merely convention: the manufacturer's own patient labelling for one system states "*Sleeping — Do not wear the sound processor while sleeping, as you may not become aware of your processor becoming unusually warm or hot,*" and manufacturer guidance elsewhere advises removal at bedtime for battery life and safety. A device dissipating tens of milliwatts already carries that warning.

The lesson transfers with force to anything thermally limited. **The user cannot be the thermal sensor**, because sleep removes their ability to act as one — which means an independent thermal cutout is not a refinement but the reason that labelling exists. Three consequences: a roughly 16-hour duty cycle, an overnight window that becomes mandatory thermal recovery rather than lost operating time, and a hard requirement for autonomous thermal shutdown.

2A. Power budget, reverse-engineered

The biological threshold is unknown, so the required intensity cannot be derived forward. It can be derived backward: what the regulatory ceiling *permits* sets an upper bound on power draw, and comparing that against a wearable budget produces a requirement on the biology.

At 1 MHz the wavelength is 1.54 mm, giving a focal spot area of 0.0186 cm². At the I_SPTA.3 ceiling of 720 mW/cm² that is **13.4 mW acoustic per focal point**, time- averaged. Each stimulated location carries its own limit, so an N-channel device delivers N times that. Taking an end-to-end electrical-to-focal efficiency of 0.096 — 0.20 path transmission, 0.60 transducer, 0.80 amplifier, all ASSUMED and all design-relevant rather than decision-relevant:

Table 3.

channels	acoustic at focus	electrical	plus DSP	hours on a 0.84 Wh cell
1	13.4 mW	140 mW	155 mW	5.4
4	53.6 mW	559 mW	574 mW	1.5
8	107.3 mW	1118 mW	1133 mW	0.7
13	174.3 mW	1816 mW	1831 mW	0.5

A cochlear implant system draws 20–40 mW and runs about a day on that cell. **An eight-channel ultrasonic aperture operating at the ceiling draws about 1.1 W — roughly 38 times a cochlear implant, and under an hour on the same battery. At the ceiling it is not wearable.**

The requirement that falls out. Inverting: inside a 40 mW cochlear-implant- comparable budget, an 8-channel device can deliver 16.1 mW/cm² per channel and a 13-channel device 9.9 mW/cm² — **2.2% and 1.4% of the regulatory ceiling respectively**. A wearable, battery-powered version therefore requires the biological excitation threshold to sit at a small fraction of a percent of the diagnostic time-average limit. A 100 mW body-worn unit with a larger cell relaxes that by roughly a factor of four. A mains-powered clinic instrument has no such constraint.

This converts an unanswerable question into a sharp one. Not *do neurons respond to focused ultrasound* — which Stage 4 would answer with a threshold — but **do they respond below about 1% of the diagnostic ceiling**, which is the number that decides whether a wearable device is possible at all.

And it selects the indication independently. The regulatory path favours a diagnostic device because it makes no therapeutic claim. The thermal argument favours short sessions over continuous wear. The power budget favours mains over battery. Three unrelated constraints point at the same product: **a clinic instrument used for minutes, not a worn device used all day**. That convergence is the most useful thing in this document.

2B. An alternative architecture: powered motes rather than a steered focus

A different arrangement deserves evaluation because it attacks the weakest point of everything above. Instead of steering a focus through aberrating bone to a millimetre target, place small transducing elements *at* the targets and use the external field only to power and address them.

This is not novel and that is a point in its favour. Ultrasonically powered, millimetre-scale implanted motes are an established research architecture for neural recording and stimulation, and published devices sit near 1 mm operating around 1 MHz — consistent with the constraint that a harvesting element be comparable to the wavelength it captures.

What it solves, and it is the thing Stage 0 puts at risk. Spatial selectivity comes from *where the elements are placed*, not from how tightly a beam can be focused through bone. If aberration through the temporal bone turns out to spoil the focus — the most likely way \$1's channel-count argument dies — a mote array is unaffected, because it never depended on the focus in the first place. It also lowers the required delivered intensity, since the element converts locally rather than the field having to do work at the target.

Where it fits anatomically, computed:

Table 4.

structure	dimension	1 mm mote fits?
scala tympani, basal turn	~1.5 mm	barely, with nothing to spare
scala tympani, apical turn	~0.5 mm	no
round window membrane	~2 mm	yes
internal auditory canal	~4 mm	yes, but see below
promontory surface	~25 mm accessible	yes

Placement route. The promontory and the round window niche are the sites to use. Both are reached through a standard tympanotomy, both sit outside the cochlear fluid space, and both avoid the two routes that are unavailable here: the internal auditory canal, which requires a retrosigmoid or middle fossa craniotomy, and the scala itself, where placing anything risks the cochlea. That last exclusion is specific to this population — in sensory loss an intracochlear site costs nothing because the cochlea is already gone, while here the working cochlea is the premise of the entire paper.

The external standing wave does not survive, and does not need to. Node positions in a skull depend on individual geometry and shift with jaw position and posture, so a standing wave cannot be relied on to place anything. But a mote array does not need a standing wave — it needs power delivery and addressing, which a plain beam supplies. Dropping that element costs the proposal nothing.

The constraint that decides it is imaging, and it is specific to this population. These patients are defined by needing repeated MRI of the internal auditory canal. Cochlear implantation and that surveillance can coexist: studies of implanted schwannoma patients report the canal and cerebellopontine angle sufficiently visible for follow-up in most cases at 1.5 T with artifact-reduction sequences. But the mitigation that makes it work is **positioning the**

artifact source away from the region of interest — published guidance recommends placing the receiver in an exaggerated superior-posterior position for exactly this reason, and a case report documents a compatible implant's shadow artifact concealing a large tumour elsewhere in the same patient.

A distributed array placed at the promontory and round window forfeits that mitigation, because the elements sit where the imaging must look.

This is a falsification point rather than a closure, and the distinction matters. The evidence here is one documented case of concealment against several series reporting adequate visualisation, and it is drawn from abstracts and extracts rather than full papers. That is not enough to close an architecture for an entire population. Alternative and complementary imaging exists, artifact-reduction sequences continue to improve, and element geometry and siting are design variables that have not been explored. **The correct statement is that artifact burden becomes a design constraint that must be traded against monitoring benefit, and that the trade belongs to clinicians and imaging physicists rather than to this document.** Where a hard physical or engineering limit closes something, this paper says so and names the wall; this is not one of those.

A safety question that applies to any repeated ultrasonic exposure of a tumour, and it must be acknowledged rather than assumed benign. Low-intensity ultrasound has documented effects on cell proliferation *in both directions*: inhibition and apoptosis are reported for osteosarcoma, leukaemia and hepatocellular lines, and an entire mechanotherapy and oncotripsy literature is built on that; but pulsed low-intensity ultrasound is also reported to *increase* proliferation and matrix production in human dermal fibroblasts. The effect is cell-type dependent. Review literature on diagnostic ultrasound notes that biological effects such as chromosome aberration and growth retardation have been recorded, but mostly at levels well above diagnostic intensity.

No data was found for vestibular schwannoma specifically, and its absence is the finding. A Schwann-cell tumour repeatedly insonated over years of surveillance is a scenario with no direct evidence either way. Any programme proposing repeated ultrasonic exposure of this population must answer that question before a first human session, and it should be answered in the direction of proliferation as well as inhibition — a diagnostic tool that accelerates the thing it is monitoring would be a catastrophic failure mode, and a tool that retards it would be a finding of a completely different kind.

Verdict, and the version that carries the least burden. As a *stimulation* array it faces the highest artifact burden, since more elements mean more shadow in the region that must be imaged — and that burden is population-specific rather than general, since in sensory loss surveillance imaging is not the point of care.

What survives is the passive half. The original application of ultrasonically powered motes was **recording**, not stimulation, and a single extracochlear element reporting local activity — powered and read out by external ultrasound, with no battery and no percutaneous lead — is a monitoring device rather than a therapy. Its artifact burden is one element rather than an array, and it lands on the same indication that the regulatory path, the thermal argument and the power budget have each independently selected. **That is now a fourth constraint pointing at the same product, arrived at from a completely different direction.**

How it fails. If a single extracochlear element degrades canal imaging enough that surveillance quality falls below what clinicians judge adequate, it fails for this population — because here the imaging *is* the care. That is a measurable question with a professional judgement attached, and both parts belong to the radiologists and surgeons who would have to live with the result.

2C. Passive tuned resonator arrays — the reed, not the pipe

A third architecture removes electronics from the implant entirely. Place a graded set of mechanical resonators, each tuned to its own frequency, and drive them with an external field. Each element responds only near its resonance, so **frequency decomposition happens in hardware with no controller, no chip and no addressing logic.** The power source sits outside the body — belt, wrist, behind the ear — and the implant is passive metal and ceramic.

One word in that description decides whether it works. A pipe or cavity resonator is set by an acoustic wavelength, and at audio frequencies there is no implantable version: a quarter-wave air column for 1 kHz is 86 mm, and in fluid it is 375 mm. The pipe *is* the wavelength. That is the same wall that closed the skull-as-substrate route.

A mechanical resonator escapes it, because its frequency comes from stiffness and mass rather than from a wavelength. The governing relation is the same for any prismatic beam,

$$f_0 = (\gamma^2/2\pi\sqrt{12}) \cdot (t/l^2) \cdot \sqrt{E/\rho}$$

with the mode constant γ set by the end conditions — 1.875 for a cantilever, giving a coefficient of 0.162, and 4.73 for a beam clamped at both ends, giving 1.028.

The design lever is a low stiffness-to-density ratio, which is the opposite of the intuition that reaches for silicon. Since frequency scales with $\sqrt{E/\rho}$, a material with a small ratio reaches low frequencies in a shorter beam. Gold's $\sqrt{E/\rho}$ is 2,023 m/s against silicon's 8,542 — a factor of 4.2 — so a gold beam is about half the length of a silicon one at the same frequency and thickness.

Recomputed for a doubly clamped gold beam 450 nm thick:

Table 5.

target frequency	required length
250 Hz	1.93 mm
1 kHz	0.97 mm
4 kHz	0.48 mm
8 kHz	0.34 mm

About a third of a millimetre to two millimetres covers the speech range. That is comfortably implantable, and it is not merely an arithmetic result: published finite-element and lumped-element modelling of exactly this structure — a graded array of doubly clamped beams designed as an artificial basilar membrane — reports gold spanning 23.4 kHz down to 252 Hz across lengths of 200 to 2000 μm , and identifies low E/ρ as the selection criterion for the same reason (Abdul, Shibly, Asary & Ruma, *Eng. Proc.* 2023, 48, 15; doi:10.3390/CSAC2023-14896, open access).

One dimensional check, offered because it affects anyone reproducing that design. Recomputing their reported resonances from their own stated relation returns values about a thousand times too high at the stated 450 μm thickness, and reproduces every reported figure to within a few percent at **450 nm**. The frequencies are internally consistent with each other and with the physics; the thickness as printed appears to be a unit slip. Anyone building from that geometry should confirm it with the authors rather than taking the printed dimension forward. The finding does not affect the paper's conclusions, which rest on the frequencies.

The tiling budget works out favourably. Covering 100 Hz to 8 kHz is 6.3 octaves. Twelve elements need $Q \approx 2.7$, twenty need $Q \approx 4.6$, thirty-two need $Q \approx 7.3$ — and real cochlear filters run Q between roughly 1 and 10. A 12-20 element array sits in the same regime the ear itself uses, which is expected rather than surprising: both are solving the same tiling problem.

And the environment supplies the damping the design needs. A cantilever in vacuum has Q in the thousands, which at 1 kHz rings for over 300 ms and would smear speech into unintelligibility. Immersed in tissue fluid, viscous damping drops Q to single digits — the range actually wanted. The implant medium is the damper, and a problem that would need engineering in air solves itself in the body.

Apodization, honestly placed. Grading coupling strength across the elements to shape the array's overall response is standard array design practice, and it maps onto the gain profile a cochlear filterbank already implements. It is a design technique available here, not a transfer of content from elsewhere, and it should be described that way.

Where it does not help, and it is the motivating case. A passive graded array performs mechanical frequency decomposition and delivers each band to a place. An intact cochlea already does exactly that, better, for nothing. In retrocochlear loss the cochlea is the working part, so this architecture duplicates the one organ that is not broken while still delivering into the pathway that is. **Closed for this population, for the same reason as everything else in this document.**

2D. Cross-population scope

Where these closures open doors. Four architectures have now been closed for slow-growth retrocochlear loss: the steered ultrasonic focus as a therapy, the powered mote array, the passive resonator array, and earlier the crosstalk-derived direction cue. **Every one of those closures is population-specific rather than general**, and the constraint that closes each is worth reading as a pointer to where it does not apply.

Table 6.

architecture	what closes it here	populations where that constraint is absent
powered mote array	artifact burden on the surveillance MRI that defines this population's care — a design constraint to be traded, not a closure	sensory loss, where imaging is not the point of care
passive resonator array	duplicates an intact cochlea	sensory and conductive loss, where the mechanical chain or the transducer really has failed
ultrasonic excitation as therapy	depends on the nerve that is compressed	any loss where the nerve is intact — which is most of them
bone-conduction direction cue	masks the good ear's own spectral cues	none found; this one closes generally

Three specific redirections follow.

Sensory loss, and the largest population. Graded piezoelectric cantilever arrays are already studied as self-powered artificial basilar membranes: the array replaces microphone, filterbank and much of the power budget by decomposing mechanically and generating charge at each element. That is prior art and an active field. What appears less explored is the same array as an externally *driven* actuator rather than a sensor. The eight-channel electrical ceiling discussed in §1 belongs to every implant recipient, so the spatial argument there is general if the biology cooperates.

Conductive and mixed loss, where a passive array has its clearest job. Ossicular replacement prostheses are single stiff struts whose transfer function is whatever the geometry happens to give. A graded resonator array is a prosthesis with a *designed* transfer function — gain shaped across frequency by element tuning rather than by luck. It does not need to beat a cochlea; it needs to replace a broken mechanical link with a better-shaped one, and in mixed loss with a dead middle ear and a live cochlea that is exactly the available job.

Auditory neuropathy and other retrocochlear disease. The measurement arguments of the main paper — emission phase gradient for robustness, efferent suppression as a loop assay, the pattern between paired hair-cell tests — do not depend on a tumour. They depend only on the failure sitting behind the cochlea, which is the defining feature of a larger class.

2D.1 The specification that survives

Stated per population. Naming where a constraint is absent is not the same as saying what would be built there. Each row below is the device that survives its population's gates.

Table 7.

population	the surviving specification	first gate
Retrocochlear, this document's case	no therapeutic device survives. A clinic-sited, mains-powered acoustic aperture used for minutes as a pathway probe and imaging aperture — bidirectional, diagnostic indication, no continuous wear	does safe-intensity excitation localise finely enough to distinguish adjacent cochlear regions?
Sensory (cochlear) loss	a steered acoustic aperture as an alternative to electrical delivery, or a distributed powered-mote array; both are attacking the same conduction-limited channel ceiling by different routes	the electrical space-constant distribution — if it is small, both lose their advantage
Conductive and mixed loss	a graded passive resonator array as an ossicular replacement with a designed transfer function: 12-20 elements, 1-7 mm, Q 3-9, no electronics, no power in the body	does a designed transfer function outperform a stiff strut in the same ear?
Normal hearing, elective	out of scope for this document; treated in the companion paper on the general-purpose channel	not an engineering gate

Read down the "first gate" column. Three of the four are answerable before any hardware is built, and two of them are the same question asked of different populations. That is the practical output of the cross-population exercise: not four programmes, but one measurement that gates two of them.

The general lesson, recorded because it is the reusable part. A narrow population imposes constraints that a broad one does not: here, surveillance imaging, an intact cochlea, and a compressed nerve. Each of those closed an architecture. Testing an idea against the *narrowest* case first is therefore efficient — it fails fast, and each failure names the constraint responsible, which is a map of where the idea does belong. **Four architectures closed against one hard case have produced three redirections into larger populations, which is a better return than four untested proposals.**

3. The regulatory envelope, and the constraint that binds

Quoted from FDA Track 3 guidance and IEC 60601-2-37, non-ophthalmic:

Table 8.

quantity	limit
derated spatial-peak temporal-average intensity, $I_{\text{SPTA.3}}$	$\leq 720 \text{ mW/cm}^2$
derated spatial-peak pulse-average intensity, $I_{\text{SPPA.3}}$	$\leq 190 \text{ W/cm}^2$
mechanical index, $MI = P_{\text{neg}}/\sqrt{f}$	≤ 1.9
derating factor	0.3 dB/cm/MHz

Peak negative pressure permitted by MI: 0.95 MPa at 0.25 MHz, 1.34 at 0.5, 1.90 at 1.0, 2.33 at 1.5.

The binding constraint is not MI. A hearing device delivers sound continuously, and continuous operation is governed by the time-average limit. At 720 mW/cm^2 the equivalent continuous pressure amplitude is about **149 kPa** — an order of magnitude below the peak the mechanical index would permit.

Stated as a duty-cycle budget, which is the form that matters: at the pulse-average ceiling, duty cycle is limited to **$I_{\text{SPTA.3}}/I_{\text{SPPA.3}} = 0.38\%$** . A device that must deliver audio continuously therefore cannot operate anywhere near the pulse limit, and its entire operating envelope sits under the time-average constraint.

Thermal. Published guidance treats an in-situ rise of $\leq 1.5^\circ\text{C}$ as usable without thermal reservation and $\geq 41^\circ\text{C}$ for five minutes or more as potentially hazardous; computed maxima at the $I_{\text{SPTA.3}}$ ceiling for focused geometries reach about 2°C . A continuously operating device therefore sits in the region where the thermal index must be actively

managed rather than assumed adequate — and the target is adjacent to bone, which heats faster than soft tissue, so the bone thermal index rather than the soft-tissue index governs.

ALARA applies throughout. Guidance explicitly states that defaulting to maximum output is not good practice.

4. Standards that would apply

Named so a reader can check them, not because compliance has been assessed.

Table 9.

standard	scope
IEC 60601-1	general safety, medical electrical equipment
IEC 60601-2-37	ultrasonic diagnostic and monitoring equipment; output display
IEC 62304	medical device software lifecycle
ISO 14971	risk management
ISO 10993 series	biocompatibility of contacting materials
IEC 60601-1-2	electromagnetic compatibility
IEC 60601-1-11	home-use environment, if worn outside a clinic
ISO 14708 series	active implantable devices, if any part is implanted
IEC 60118 / ANSI-ASA S3.22	hearing aid characteristics and measurement
IEC 62359	thermal and mechanical index determination

Two regulatory observations that materially affect strategy. A **diagnostic** device making no therapeutic claim follows a different and considerably shorter path than a therapeutic one — which is an argument for the imaging-and-probe indication of §1 reaching a patient first. And a device intended for an orphan indication may qualify for **Humanitarian Device Exemption**-type pathways in some jurisdictions; whether this population qualifies has not been assessed and is a question for a regulatory professional rather than for this document.

5. What would have to be shown, in order

Each stage gates the next. Failure at any stage stops the programme, and the earlier stages are cheap.

Stage 0 — settle the space constant. Obtain the distribution of the electrical space constant across real implant recipients from evoked-response data already collected at fitting. *Kills the programme if:* the median sits near or below 1.5 mm, because the acoustic advantage then falls to a factor of two or less. This costs nothing beyond data access and it gates everything below it.

Stage 1 — field simulation. Model the acoustic field from a mastoid-sited aperture through temporal bone into the cochlea. Establish achievable focal spot size and steering range, and confirm that in-situ intensity at the target stays within §3 while remaining above any plausible biological threshold. *Kills the programme if:* the focal spot in real temporal bone is no smaller than the 3–4 mm electrical spread constant, because then the entire channel-count argument fails and there is no reason to proceed. **This stage requires no hardware and should be done first.**

Stage 2 — phantom validation. Build the field in a tissue-mimicking phantom with embedded temporal bone, and measure with a calibrated hydrophone. *Kills it if:* measured fields diverge from simulation by more than the design margin, or aberration through bone is uncorrectable at the available element count.

Stage 3 — cadaveric temporal bone. Verify coupling, siting and field geometry in real anatomy with real variability. *Kills it if:* inter-specimen variation in achievable focus exceeds the spatial resolution the concept depends on.

Stage 4 — animal, with the auditory confound controlled. This is the stage that answers the actual open question: does focused ultrasound excite spiral ganglion neurons directly, and with what spatial selectivity? The methodology here is established and must be followed — the ultrasound neuromodulation literature has learned, expensively, that focused ultrasound produces an indirect auditory response that masquerades as direct neural activation, and deafened animal models are now standard specifically to separate the two. *Kills it if:* apparent responses are indirect, or selectivity is no better than electrical stimulation.

Stage 5 — human, diagnostic indication only. A probe that maps where the pathway conducts, making no therapeutic claim. *Kills it if:* thresholds in conscious humans exceed the \$3 envelope, or the map adds nothing to existing tests.

Stage 6 — therapeutic indication. Only reachable if all of the above pass, and only in a population where an alternative is inadequate.

Realistic scale. Stage 0 is a data request. Stages 1 and 2 are a graduate project. Stages 3 to 5 are a funded multi-year programme with a clinical partner. Stating that plainly is part of scoping: an idea that requires stage 4 to be interesting is not an idea a paper can settle.

6. What is not known, and who would have to answer it

- **What is the real distribution of the electrical space constant?** Unknown to this document, and it decides whether the central argument holds. Stage 0.
- **Do spiral ganglion neurons respond to focused ultrasound at intensities inside the \$3 envelope?** Unknown. Requires stage 4. This is the question the whole concept rests on and nothing in this document establishes it.
- **What is the mechanism if they do?** Radiation force, mechanosensitive channel gating, cavitation and thermal effects have all been proposed for ultrasonic neuromodulation generally, and the mechanism determines the safe operating envelope. Unresolved in the wider field.
- **Is the focal spot in real temporal bone anywhere near the free-field figure?** Bone aberrates severely. Requires stage 1 and stage 3.
- **Can continuous operation stay within thermal limits at a bone-adjacent target?** Requires stage 1 with a thermal model, not an acoustic one alone.
- **Does the 30 dB dynamic range of bone-conducted ultrasound perception apply to a direct-excitation route?** If it does, the delivery channel is narrower than speech requires and the concept needs compression it may not tolerate.

7. Where this leaves the concept

The document opened by promising a device and it should close by saying what the device is, what was given up to get there, and what would settle it.

What survived, in one paragraph. For slow-growth retrocochlear loss there is no surviving therapeutic device, and four separate constraints agree on that: the nerve dependency of any acoustic route, a power budget that puts an eight-channel aperture at roughly 1.1 W against a cochlear implant's 30 mW, a thermal limit that forbids continuous wear where the user cannot act as the sensor, and a regulatory path that favours a device making no therapeutic claim. What survives is a **clinic-sited, mains-powered, bidirectional acoustic aperture used for minutes at a time as a pathway probe** — an instrument, not a prosthesis. Four independent lines of argument selecting the same product is the strongest structural result here, and none of them was aimed at that answer.

What was given up. A steered focus as therapy, closed by nerve dependency. A distributed powered-mote array, constrained by artifact burden on the surveillance imaging that defines this population's care — a trade for clinicians rather than a wall. A passive resonator array, which duplicates a cochlea that here is intact. And the standing wave, the needle route, and the cavity resonator, each closed by a wavelength comparison that cost one calculation apiece.

What would settle it, in order and by cost. The electrical space-constant distribution, obtainable from evoked-response data already collected at fitting, gates the entire spatial argument and requires only data access. Field simulation through real temporal bone gates the focus and requires no hardware. The schwannoma-and-ultrasound proliferation question gates any repeated exposure of this population and has no direct evidence in either direction — it is the gate that must not be skipped. Only then does anything get built.

What this document is not. No transducer has been modelled, no field simulated, no phantom built, and the biology on which the whole concept rests is unestablished. The regulatory envelope is real and quoted; the specification above is what fits inside it, not evidence that anything should be placed there.

References

Retrieved and verified during preparation. Entries marked **[partial]** were identified by title and source but their complete bibliographic record was not obtained; they are listed with enough detail to be found. No author list, volume or year appearing below without that marker was inferred.

Acoustic output limits and safety

1. U.S. Food and Drug Administration. *Marketing Clearance of Diagnostic Ultrasound Systems and Transducers* — guidance for industry and Food and Drug Administration staff.
2. IEC 60601-2-37. *Medical electrical equipment — Part 2-37: Particular requirements for the basic safety and essential performance of ultrasonic medical diagnostic and monitoring equipment.*
3. Health Canada. *Guidelines for the Safe Use of Diagnostic Ultrasound.*
4. Transient temperature rise under derated diagnostic intensity limits. *Ultrasound in Medicine & Biology.* **[partial]**

Ultrasound and cell proliferation

5. Low-intensity pulsed ultrasound and proliferation: inhibition and apoptosis reported in osteosarcoma, leukaemia and hepatocellular lines; increased proliferation and matrix production reported in human dermal fibroblasts. **[partial — several primary sources identified by description only. The direction of effect is cell-type dependent, and no data was found for vestibular schwannoma; see §2B.]**

Resonator arrays and artificial basilar membranes

6. Abdul, B.; Shibly, M.A.H.; Asary, A.R.; Ruma, N.J. Design and modelling of MEMS resonators for an artificial basilar membrane. *Engineering Proceedings* **2023**, 48(1), 15. doi:10.3390/CSAC2023-14896 (open access, CC BY). *Note: the reported beam thickness of 450 µm is inconsistent with the paper's own reported resonance frequencies by approximately three orders of magnitude and resolves at 450 nm; see §2C.*

Ultrasonically powered implants

7. Seo, D.; Neely, R.M.; Shen, K.; Singhal, U.; Alon, E.; Rabaey, J.M.; Carmena, J.M.; Maharbiz, M.M. Wireless recording in the peripheral nervous system with ultrasonic neural dust. *Neuron* **2016**, 91(3), 529–539. doi:10.1016/j.neuron.2016.06.034
8. Piech, D.K.; et al. A wireless millimetre-scale implantable neural stimulator with ultrasonically powered bidirectional communication. *Nature Biomedical Engineering* **2020**, 4(2), 207–222. doi:10.1038/s41551-020-0518-9

Cochlear implants: electrical spread, imaging, channel limits

9. Shannon, R.V.; Zeng, F.-G.; Kamath, V.; Wygonski, J.; Ekelid, M. Speech recognition with primarily temporal cues. *Science* **1995**, 270(5234), 303–304. doi:10.1126/science.270.5234.303
10. Transimpedance and spread-of-excitation measurement in cochlear implant recipients; impedance spectroscopy of human cochleae. **[partial — two primary sources identified by description only. The numeric space constant is a per-patient fitted parameter, not a constant; see §1.]**
11. MRI surveillance of the internal auditory canal in cochlear implant recipients, and a case report of magnet shadow artifact concealing pathology in neurofibromatosis type 2. **[partial]**

Bone-conducted ultrasound and transcranial transmission

12. Stenfelt, S. Transcranial attenuation of bone-conducted sound when stimulation is at the mastoid and at the bone conduction hearing aid position. *Otology & Neurotology* **2012**, 33, 105–114.
13. Rigato, C.; Reinfeldt, S.; Håkansson, B.; Fredén Jansson, K.-J.; Renvall, E.; Eeg-Olofsson, M. Effect of transducer attachment on vibration transmission and transcranial attenuation for direct drive bone conduction stimulation. *Hearing Research* **2019**, 381, 107763. doi:10.1016/j.heares.2019.06.006
14. Perception of bone-conducted ultrasound, dynamic range and carrier-independent pitch; and skull vibratory patterns from 250 Hz to 64 kHz reporting increasing contralateral isolation with frequency. **[partial — the ultrasonic extension of transcranial attenuation is not established by references 12 and 13 and is marked UNVERIFIED in the provenance table.]**

Focused ultrasound neuromodulation

15. Indirect auditory activation as a confound in focused ultrasound neuromodulation, with contralateral bias and deafened animal models used as controls. **[partial — several primary sources identified by description]**

Declaration

This page is a standing declaration attached to every preprint in this series. It states what was and was not done, so a reader can judge the work rather than infer its rigour from its tone.

Publication status

This is a preprint. It has not been peer reviewed. It is a checkpoint on a road, published so the reasoning is available and checkable while the work is still in progress, not a finished result presented as settled.

The literature pass is limited and incomplete. Searches were targeted and natural-language, run against general web indices. They did **not** include structured database searching with controlled vocabulary, systematic patent searching by classification code, clinical trial registries, or non-English literature. Patent literature in particular is where ideas of this kind predominantly sit and is where this method is weakest.

Prior art was read at abstract and extract level. Except where a full text was opened and said to have been, findings come from abstracts, extracts and directed keyword pulls rather than from complete papers. Statements of the form *the literature indicates X* should be read as *prior art at abstract level indicates X, to be confirmed against the full text*. This matters most where a single study is used to close a direction, and this document avoids doing that.

A screening pass is not a clearance. Where a claim is described as unclaimed or not found, the correct reading is *not found by these searches*, not *absent from the art*. Anything intended for filing, commercial use, or a novelty claim requires a professional prior-art search first.

Both successes and failures are reported. Approaches evaluated and not pursued are recorded in an appendix with the quantity that closed each one, on the view that a documented dead end is worth more to a reader than an additional proposal.

Authorship and assistance

Prepared by a part-time, unaffiliated research effort, published free to the world. An AI language model (Anthropic Claude) assisted with structural drafting, literature recall, arithmetic and unit checking, and targeted literature search. No automated tool generated any observation from a patient, a device, or an instrument. **The author is responsible for the correctness of all content and citations.**

Provenance of quantities

Every load-bearing number is tagged by how it was obtained. This exists because the characteristic failure of speculative engineering is the inherited constant: a figure taken from a reference propagates through documents unchecked, because each reader assumes an earlier reader verified it.

Two classes are separated, and they warrant different levels of concern. A number is **decision-relevant** if changing it by a factor of two changes whether the work should proceed. It is **design-relevant** if it changes how something is built but not whether. An unverified design-relevant number is untidy; an unverified decision-relevant number can invalidate the work.

Table 10.

tag	meaning
SOURCED	traced to a named primary or authoritative source; count of independent sources given
COMPUTED	derived in this work from SOURCED inputs; the calculation is reproducible from the stated relation
UNVERIFIED	recalled or widely repeated, not checked against a source in preparing this document
ASSUMED	chosen as representative for the purpose of an estimate; not a measurement

Where a quantity is SOURCED to a single source and carries safety consequence, that is noted, because one source is not corroboration.

Table 11. Provenance table — this document.

quantity	tag	source / note
I_SPTA.3 ≤ 720 mW/cm ²	SOURCED $\times 3$	FDA Marketing Clearance of Diagnostic Ultrasound Systems and Transducers; IEC 60601-2-37 commentary; Health Canada Guidelines for the Safe Use of Diagnostic Ultrasound
I_SPPA.3 ≤ 190 W/cm ²	SOURCED $\times 2$	FDA guidance; <i>Ultrasound in Med. & Biol.</i> review of transient temperature rise
MI ≤ 1.9	SOURCED $\times 4$	FDA guidance; Health Canada; <i>J. Ultrasound Med.</i> safety overview; clinical protocol documents
derating 0.3 dB/cm/MHz	SOURCED $\times 3$	as above
thermal criteria 1.5 °C / 41 °C for 5 min	SOURCED $\times 1$	Health Canada guidance. Single source, safety-relevant — corroboration needed
CI processor removal at night	SOURCED $\times 2$	Cochlear Nucleus CI422 patient labelling ("you may not become aware of your processor becoming unusually warm or hot"); MED-EL patient guidance
electrical stimulus spreads; channel interaction limits selectivity	SOURCED $\times 3$	<i>Hearing Research</i> transimpedance/spread-of-excitation study (n=38); <i>APL Materials</i> impedance spectroscopy of human cochleas; multiple device patents
electrical space constant, numeric value	UNVERIFIED — DECISION-RELEVANT	a per-patient fitted parameter, not a constant. Device patents state patient-specific values improve on assumed ones. Stage 0 exists to settle it
~8 effective channels	UNVERIFIED	widely repeated; primary study not checked
skull attenuation vs frequency	ASSUMED — design-relevant	representative figures. Conclusion (mastoid over vault) survives a factor-of-two error
skull thickness 7 mm	ASSUMED	varies roughly 4–10 mm
ultrasonically powered motes ~1 mm at ~1 MHz	SOURCED $\times 1$	neural-mote literature
graded beam arrays as artificial basilar membranes	SOURCED $\times 2$	Abdul et al., <i>Eng. Proc.</i> 2023, 48, 15 (doi:10.3390/CSAC2023-14896), FE and LE modelling; plus the wider MEMS artificial-basilar-membrane literature
beam resonance relation and mode constants ($\gamma = 1.875 / 4.73$)	SOURCED $\times 1$, COMPUTED	standard beam theory; both end conditions reproduced independently here
gold spans 23.4 kHz–252 Hz over 200–2000 μm lengths	SOURCED $\times 1$	Abdul et al., as above — reported thickness inconsistent with reported frequencies by $\sim 10^3$; resolves at 450 nm. Flagged in §2C
low E/ ρ as the material selection criterion	SOURCED $\times 1$	Abdul et al., stated explicitly
material constants for Au, Pt, Cr, Mo, Si	SOURCED $\times 1$ / ASSUMED	Au, Pt, Cr, Mo from Abdul et al. Table 2; silicon standard values
cochlear filter Q ≈ 1 –10	UNVERIFIED	widely repeated; primary source not checked
cantilever, pipe and Q-tiling relations	COMPUTED	from the stated formulae
CI/MRI surveillance of the canal feasible with positioning and MARS	SOURCED $\times 3$	15-exam MED-EL series; 8-scan retained-magnet series; 1.5 T artifact-positioning study
CI magnet artifact can conceal pathology	SOURCED $\times 1$	NF2 case report

cochlear and IAC dimensions	ASSUMED	standard anatomical ranges
BCU dynamic range ~30 dB; carrier-independent pitch	SOURCED ×1	<i>Neuroscience Letters</i> study of bone-conducted ultrasound under cisplatin
ultrasonic contralateral isolation rises with frequency	SOURCED ×1	Dunlap, Lenhardt & Clarke skull vibratory patterns, via patent citation. Second source needed
tissue $c = 1540$ m/s, $\rho = 1050$ kg/m ³	ASSUMED	standard soft-tissue values
focal spot area, per-channel acoustic power, electrical draw, % of ceiling	COMPUTED	from the SOURCED limits above and the ASSUMED efficiencies stated in §2A
transducer / path / amplifier efficiencies	ASSUMED — design-relevant	0.20 / 0.60 / 0.80; plausible, unsourced