# Outcome-Defined Silicon

Specifying Front-End Chips for Cancer Imaging from the Clinical Decision Backward

Abhay Chakra Sadineni
University of California, Berkeley
abhaychakra@berkeley.edu

November 2026
DOI: 10.5281/zenodo.24210908

**Abstract**

The chips inside a medical imaging system are usually specified the way all chips are, by channel count, resolution, noise, bandwidth and power, chosen by engineers from what the previous generation offered and what the bill of materials allows. The clinical outcome the system exists to support enters at the end, as a validation study. This paper argues for running the specification the other way. Starting from a clinical decision stated as a sensitivity and specificity at an operating point, it derives the detectability the imaging task requires, then the signal-to-noise ratio each channel must deliver, and from that the front-end specification. A worked example for photoacoustic (optoacoustic) detection of a vascularised lesion carries the chain through optical transport, thermoelastic conversion, acoustic attenuation and thermal noise, showing where each link enters. It stops short of a feasibility verdict, because transducer sensitivity and geometric spreading are not modelled. Three results follow. The product of channel count and pulse averaging that a design needs grows as $e^{2\mu_{\mathrm{eff}}z}\,10^{\alpha f z/10}$ with lesion depth, about elevenfold per additional centimetre for the assumed tissue ($\mu_{\mathrm{eff}} = 1$ cm$^{-1}$ and 0.6 dB cm$^{-1}$ MHz$^{-1}$ at 3 MHz), so depth targets dominate cost. Converter resolution is set by the dynamic range between superficial signals and the noise floor, not by the lesion. And because superficial absorption varies with skin pigmentation, a specification written for a population-level sensitivity must carry roughly three extra bits of dynamic range that a specification written for a reference phantom would omit. The paper sets out what an outcome-defined specification contains and why the approach fits the design-input traceability that the U.S. Quality Management System Regulation, in force since 2 February 2026, already asks for.

**Keywords:** photoacoustic imaging; optoacoustics; medical device design; analog front end; ideal observer; design inputs; product strategy.

## 1 Introduction

An imperfect analogy helps. The transistor was not at first a better vacuum tube on the tube's own measures of gain, noise and life. It replaced the tube by performing the tube's function in a different way.

Medical imaging hardware is specified today much as tubes were. A photoacoustic or ultrasound system needs an analogue front end, meaning low-noise amplifiers, variable gain, filters and converters for each element of the transducer array. Those parts are chosen from catalogues and roadmaps described by channel count, bits, input-referred noise, bandwidth and milliwatts per channel. The clinical question, whether this system finds a cancer that matters in the patients it will see, is asked last, in a validation study, after the silicon has been fixed.

This paper proposes the reverse order. Write the clinical decision first, as a number. Derive the imaging performance it demands. Derive from that, through the physics, what each channel must do. The front-end specification becomes an output of the clinical requirement instead of an input to it. The idea is close to what regulators already expect of device design inputs, which must be derived from the intended use and user requirements [1]. Task-based design of photoacoustic systems already exists. Recent work ranks illumination designs by a Bayesian Cramér-Rao bound on reconstruction error [2], and model observers have long been used to evaluate imaging systems against detection tasks [3]. What this paper adds is narrower. It carries a clinical operating point down to the analogue front-end electronics, and it argues that the specification must hold across the patient population rather than for one phantom.

## 2 The chain

Figure 1 shows the five links.

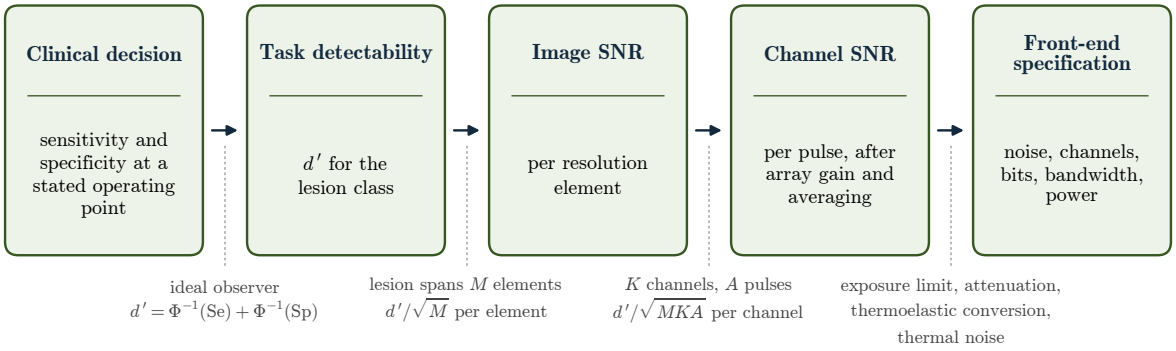



Figure 1. From clinical decision to front-end specification. Beneath each arrow is the relation or the physics that carries that link.

*From decision to detectability.* A detection task with sensitivity Se and specificity Sp at a chosen operating point corresponds, for an ideal observer facing equal-variance Gaussian statistics, to a detectability index

$$d' = \Phi^{-1}(\mathrm{Se}) + \Phi^{-1}(\mathrm{Sp}), \tag{1}$$

where $\Phi^{-1}$ is the inverse standard normal distribution [3, 4]. $\mathrm{Se} = \mathrm{Sp} = 0.90$ gives $d' = 2.56$.

*From detectability to image SNR.* If the lesion covers $M$ statistically independent resolution elements and the observer integrates over them, the signal-to-noise ratio needed in each element is $d'/\sqrt{M}$.

*From image SNR to channel SNR.* Delay-and-sum reconstruction over $K$ channels with uncorrelated noise, and averaging over $A$ laser pulses, raise SNR by $\sqrt{KA}$. Each channel, on each pulse, must therefore deliver

$$\mathrm{SNR}_{\mathrm{ch}} = \frac{d'}{\sqrt{MKA}}. \tag{2}$$

*From channel SNR to the front end.* The channel's signal comes from physics the designer does not control. Its noise comes partly from physics and partly from the chip. The next section works through both.

## 3 Photoacoustic detection at depth, worked through

Photoacoustic imaging sends a short laser pulse into tissue. Absorbed light heats the tissue faster than it can expand or conduct heat away, and the thermoelastic pressure rise launches an ultrasound wave that an array detects [5, 6]. The signal crosses four physical domains before it reaches a converter, which makes it a clean test of the chain.

*Optical.* Fluence at the skin is capped by laser safety limits. For pulsed exposure of skin between 700 and 1,050 nm, the maximum permissible exposure for a single pulse of 1 to 100 ns is $20 \times 10^{2(\lambda-0.7)}$ mJ cm$^{-2}$ with $\lambda$ in micrometres [7], or 31.7 mJ cm$^{-2}$ at 800 nm. For pulse trains, the limit on average irradiance, $0.2 \times 10^{2(\lambda-0.7)}$ W cm$^{-2}$ or about 0.32 W cm$^{-2}$ at 800 nm, allows this per-pulse fluence only up to a repetition rate of about 10 Hz, so averaging more pulses costs frame time as well as fluence. These values come from the 2014 edition of the standard, since superseded by the 2022 edition [8], against which they should be confirmed before design use. In the diffusive regime, fluence falls approximately as $F(z) = F_0\, e^{-\mu_{\text{eff}} z}$, with $\mu_{\text{eff}}$ set by the tissue's absorption and reduced scattering [9].

*Thermal to mechanical.* Under thermal and stress confinement the initial pressure is

$$p_0 = \Gamma\, \mu_a\, F, \tag{3}$$

where $\mu_a$ is the optical absorption coefficient and $\Gamma = \beta c^2 / C_p$, the Grüneisen parameter, combines the thermal expansion coefficient, sound speed and specific heat [10]. It is a thermodynamic property of the tissue, and it sets how much of the absorbed energy becomes sound.

*Acoustic.* The pressure wave loses amplitude with distance at roughly $\alpha$ dB per cm per MHz, with $\alpha$ of order 0.5 to 0.75 in soft tissue [11], and spreads geometrically by a factor that depends on lesion size and aperture.

*Electrical.* The front end cannot do better than the thermal noise of its source impedance, $v_n = \sqrt{4 k_B T R B}$ [12, 13]. For 50 Ω at body temperature over 10 MHz, $v_n = 2.9$ μV rms. Amplifier noise adds to this; a good front end keeps the excess small.

Table 1 carries illustrative values through the chain. They are chosen to be plausible, not measured, and the point of the table is the structure.

| Quantity | Illustrative value | Basis |
|---|---|---|
| Clinical operating point | Se = Sp = 0.90 | Requirement |
| Detectability $d'$ | 2.56 | Equation (1) |
| Wavelength, skin exposure limit | 800 nm, 31.7 mJ cm$^{-2}$ per pulse | ANSI Z136.1 |
| Effective attenuation $\mu_{\text{eff}}$ | 1.0 cm$^{-1}$ | Assumed |
| Lesion depth $z$ | 3 cm | Requirement |
| Fluence at lesion | 1.58 mJ cm$^{-2}$ | $F_0 e^{-\mu_{\text{eff}} z}$ |
| Grüneisen parameter $\Gamma$ | 0.2 | Assumed |
| Lesion absorption contrast $\Delta\mu_a$ | 0.1 cm$^{-1}$ | Assumed |
| Initial pressure contrast at lesion | 32 Pa | Equation (3) |
| Acoustic loss to surface, 3 MHz, $\alpha = 0.6$ | 5.4 dB (amplitude × 0.54) | $\alpha f z$ |
| Thermal noise, 50 Ω, 310 K, 10 MHz | 2.9 μV rms | Johnson-Nyquist |
| Resolution elements in lesion $M$ | 16 | Assumed |
| Channels $K$, pulses averaged $A$ | 128, 10 | Design choice |
| Required SNR per channel per pulse | 0.018 | Equation (2) |

Table 1. An illustrative budget. Values marked assumed are placeholders for measurements a real design would use.

Three results follow from the structure, whatever the exact values.

### 3.1 Depth sets the cost

Signal from the lesion scales as $e^{-\mu_{\text{eff}} z}$ optically and as $10^{-\alpha f z/20}$ acoustically. For a fixed $d'$, the required $KA$ scales as the inverse square of signal amplitude,

$$KA \propto e^{2\mu_{\text{eff}} z}\, 10^{\alpha f z/10}. \tag{4}$$

With $\mu_{\text{eff}} = 1$ cm$^{-1}$ and 0.6 dB cm$^{-1}$ MHz$^{-1}$ at 3 MHz, each additional centimetre of required depth multiplies the channel-pulse product by about 11 (Figure 2). Channels cost silicon area, power and cable; pulses cost frame rate and laser life. A clinical requirement that says "lesions to 4 cm" rather than "to 3 cm" is a requirement for roughly an order of magnitude more front-end resource. That trade belongs in front of the clinical team when the requirement is written, not discovered by the hardware team afterwards.

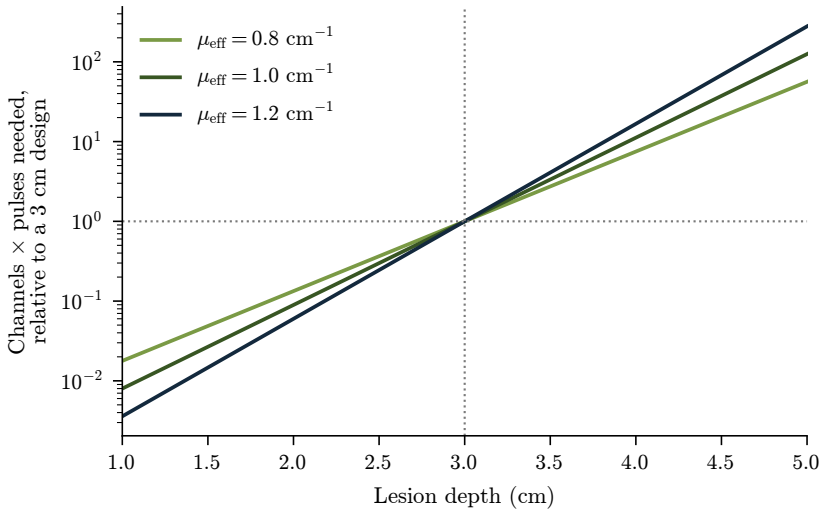



Figure 2. Channel-pulse product required to hold the same detectability as lesion depth changes, relative to a design for 3 cm, for three values of effective optical attenuation (acoustic loss 0.6 dB cm$^{-1}$ MHz$^{-1}$ at 3 MHz).

### 3.2 The lesion does not set the bits

At the design point, the lesion's signal on each channel sits below the noise. It emerges only after summing across channels and pulses. Converter resolution therefore has nothing to do with resolving the lesion on one channel. What sets it is the strongest signal the converter must hold without clipping, usually from superficial absorbers such as the skin, set against the noise floor that must not be swamped by quantisation. If the superficial signal is $R$ times the lesion's per-channel amplitude, and quantisation noise is held at one third of the analogue noise, the minimum resolution is

$$N \geq \log_2\!\left(\frac{6}{\sqrt{12}}\, R \cdot \text{SNR}_{\text{ch}}\right). \tag{5}$$

The specification for bits is thus a statement about the ratio of the brightest thing in the field to the faintest thing that matters, which is a clinical statement once the faintest thing is defined by the decision.

### 3.3 The population sets the margin

Melanin absorption falls with wavelength but remains significant in the near infrared [9]. The superficial signal, and with it $R$, therefore rises with skin pigmentation, while the light reaching

the lesion falls, which lowers the channel signal-to-noise ratio as well. If superficial absorption varies by a factor of ten across the intended population, an assumption made here for illustration, equation (5) adds $\log_2 10 \approx 3.3$ bits. A specification derived from a reference phantom with one skin tone will omit those bits. A specification derived from a sensitivity requirement that must hold across the population cannot. The same logic applies to every quantity in Table 1 that varies between patients. Breast density changes $\mu_{\text{eff}}$, and body habitus changes depth. An outcome-defined specification is written for the distribution of patients, not the median.

## 4 What an outcome-defined specification contains

1. The clinical decision, the operating point, and the consequence of each error type.
2. The intended population, stated as distributions of the physical quantities that enter the budget, namely depth, lesion size, tissue attenuation and superficial absorption.
3. The physics envelope, made up of exposure limits, acoustic and optical properties and thermal noise.
4. The derived budget, with each link's assumption and margin, from detectability down to channel count, converter resolution, bandwidth and noise.
5. The verification evidence the deployed system must keep producing to show it still meets the decision, on the instrument actually in use. Related work by the author shows how quickly such a guarantee can fail silently when the instrument changes [14].

## 5 Why now

Three conditions make this practical in 2026. Front ends for array imaging are now built as integrated multi-channel parts, so a family of chips can be parameterised by channel count, gain range and resolution rather than fixed. Simulation of light, heat and sound in tissue is routine enough to run the chain for a population rather than a phantom. And the U.S. device quality regulation now incorporates ISO 13485:2016 by reference, effective 2 February 2026 [15]. An outcome-defined specification is that trace made quantitative.

For a semiconductor company, the commercial reading is direct. A chip specified against a clinical decision is easier to qualify into a device, because part of the device's design history is already written. For a device company, it makes the cost of each clinical claim visible before the hardware is frozen.

## 6 Limits

The ideal-observer model assumes Gaussian noise and a known signal. In practice, reconstruction errors and clutter, not thermal noise, often limit photoacoustic detection at depth, and they do not average away as $\sqrt{KA}$. The values in Table 1 are illustrative; a real budget would use the measured noise-equivalent pressure of the transducer, measured tissue properties and reader-study data. Geometric spreading and transducer sensitivity are folded into the channel SNR rather than modelled. None of this changes the order of the argument. It changes the numbers that go into it.

## Funding and interests

No external funding was received. The author declares no conflicts of interest.

## References

[1] International Organization for Standardization (2016). *ISO 13485:2016, Medical devices: Quality management systems, Requirements for regulatory purposes*. Geneva: ISO.

[2] Scope Crafts, E., Anastasio, M. A., and Villa, U. (2024). Optimizing quantitative photoacoustic imaging systems: The Bayesian Cramér-Rao bound approach. *Inverse Problems*, 40(12), 125012. `https://doi.org/10.1088/1361-6420/ad910a`

[3] Barrett, H. H., and Myers, K. J. (2004). *Foundations of Image Science*. Hoboken, NJ: Wiley.

[4] Green, D. M., and Swets, J. A. (1966). *Signal Detection Theory and Psychophysics*. New York: Wiley.

[5] Beard, P. (2011). Biomedical photoacoustic imaging. *Interface Focus*, 1(4), 602-631. doi:10.1098/rsfs.2011.0028.

[6] Xia, J., Yao, J., and Wang, L. V. (2014). Photoacoustic tomography: Principles and advances. *Progress in Electromagnetics Research*, 147, 1-22.

[7] American National Standards Institute (2014). *ANSI Z136.1-2014: American National Standard for Safe Use of Lasers*. Orlando, FL: Laser Institute of America.

[8] American National Standards Institute (2022). *ANSI Z136.1-2022: American National Standard for Safe Use of Lasers*. Orlando, FL: Laser Institute of America.

[9] Jacques, S. L. (2013). Optical properties of biological tissues: A review. *Physics in Medicine and Biology*, 58(11), R37-R61.

[10] Cox, B., Laufer, J. G., Arridge, S. R., and Beard, P. C. (2012). Quantitative spectroscopic photoacoustic imaging: A review. *Journal of Biomedical Optics*, 17(6), 061202.

[11] Szabo, T. L. (2014). *Diagnostic Ultrasound Imaging: Inside Out*, 2nd edition. Oxford: Academic Press.

[12] Johnson, J. B. (1928). Thermal agitation of electricity in conductors. *Physical Review*, 32(1), 97-109.

[13] Nyquist, H. (1928). Thermal agitation of electric charge in conductors. *Physical Review*, 32(1), 110-113.

[14] Sadineni, A. C. (2026). Spectral Passport: When does a Raman model's uncertainty guarantee survive a change of instrument? Preprint, November 2026. `https://doi.org/10.5281/zenodo.24210927`

[15] U.S. Food and Drug Administration (2024). Medical Devices; Quality System Regulation Amendments. Final rule, 21 CFR Part 820. 89 Fed. Reg. 7496 (2 February 2024), effective 2 February 2026.