Risk Assessment of the ICHNOS Mechanism

ICHNOS · Patras Medicine iGEM 2026

Part of our Safety and Security work

October 2026

Scope: The ICHNOS genetic circuit (sensing modules, peptide transmission, the TetR feedback node and the tandem fluorescent timer), assessed for the ways the mechanism could fail, mislead or create risk.

Companion document: Hazard Analysis of Wet-Lab Materials, covering reagents, consumables, chemical hazards and waste. This document does not repeat that material.

Contents

How to read this document **3**
1. The mechanism in brief **3**
2. Design decisions that reduce risk **3**
3. Biological hazard: what the assessment excludes **4**
4. Construct identity and host physiology **5**
5. What the sensing modules report **6**
5.1 Adaptation is a property of the pathways, not a defect of the circuit 7
5.2 Each promoter reports its pathway, not a chemical 7
5.3 The copper control: what it establishes 8
6. Signal transmission: TIP and TetR **8**
6.1 The fusion is the unit that matters 8
6.2 TetR and the loop 9
6.3 The carbon source is an input to the circuit 9
7. The timer and measurement integrity **9**
7.1 What the ratio reports 9
7.2 Growth rate couples the two channels 10
7.3 Interpretation and artefacts 10
7.4 Viability and the readable range 11
8. Genetic stability **12**
9. Exposure and containment **12**
10. Biosecurity and data integrity **12**
11. Risk register **13**
12. Limits of use **14**

How to read this document

A risk assessment of a biosensor has to answer two quite different questions and we have organised the document so that neither is allowed to hide the other.

The first is the conventional one: **does the mechanism create a biological hazard?** Sections 3 and 4 address it and the answer is largely favourable. ICHNOS is a recording architecture in a Risk Group 1 organism whose output is fluorescence. Several of our design choices reduce risk deliberately rather than incidentally and we set those out before anything else, because they are the part of this assessment that is already settled.

The second question is the one specific to instruments: could the mechanism produce a result that is wrong in a way nobody notices? Sections 5 to 8 address it. Within contained research this is a question about data quality. It would become a safety question only in a future context where a reading supported a claim that something is safe, a context ICHNOS is not in today and one we treat throughout as a boundary to be respected rather than a present danger.

We have written the uncertainties out in detail because that is the honest state of a proof-of-concept system and because an uncertainty we have identified and planned for is a managed risk, while one we have not is a liability. Where a limitation is intrinsic to a published method rather than to our implementation, we say so.

1. The mechanism in brief

Three strains, each designed to carry a single construct integrated at the HO locus of *S. cerevisiae* CEN.PK2-1C, selected by *HIS3* complementation.

Each follows the same architecture. An input promoter responds to a condition: a synthetic Yap1-responsive promoter for oxidative stress, the native KAR2 promoter for ER stress via Ire1–Hac1 and the native CUP1 promoter as a copper-inducible functional control. The promoter drives TIP, a short peptide expressed as a fusion with a GB1 solubility domain, a flexible linker and an SV40 nuclear localization signal. TIP binds and inactivates TetR, which represses a GAL1-UAS/tetO2 hybrid promoter (pGAL1-D12). That promoter drives the reporter, an sfGFP–mCherry tandem fluorescent timer. It also drives a second copy of TetR, closing the negative feedback loop intended to linearize an otherwise switch-like response.

The system reports two quantities: **how much** stress occurred, read from reporter abundance, together with how long ago, read from the mCherry/sfGFP ratio.

2. Design decisions that reduce risk

These were made during design, for the reasons given. They are the foundation of the safety case. We list them first because each one closes a failure mode that would otherwise appear later in this document.

Separate strains rather than a combined construct. An early design placed two sensing modules on one plasmid to reduce synthesis cost. We assessed it against two specific risks and rejected it. First, two identical copies of the TIP sequence in the same orientation on one molecule are a substrate for recombination by single-strand annealing, which would delete the intervening promoter. Second, Cu^{2+} participates in Fenton-type chemistry and generates ROS, so a copper control sharing a cell with the oxidative module could cross-activate it. Keeping the modules in separate strains cost more to synthesise and bought a signal attributable to one input. This decision pre-empts the cross-activation failure mode entirely and reduces the recombination surface. A related point belongs here and is worth stating generally rather than only about copper: oxidative

stress and ER stress are not independent conditions in a cell. Oxidative damage impairs disulphide formation and protein folding and so triggers the UPR, while a loaded or failing ER generates reactive oxygen species in turn. The two sensing modules would therefore have reported partly overlapping states even with copper out of the picture. Separating them into different strains does not abolish that biology, which belongs to the pathways and not to our design, but it does mean each reading comes from a cell given one deliberate input, which is the only arrangement in which the overlap can be characterised rather than merely suspected.

Genomic integration rather than episomal plasmids. Integration at a defined locus is intended to remove the copy-number variation that plasmid maintenance introduces, which would otherwise propagate directly into every quantitative measurement. It also means the construct is not carried on a mobilisable element. We state the aim rather than the achievement deliberately. The HO locus can carry more than one integrated copy and our copy number is not yet established, so the stability this choice is designed to deliver remains an expectation supported by the design rather than a measured property. That is why copy number appears in Section 4 as the outstanding measurement it is. The alternatives we set aside would have been worse on exactly this axis. A 2µ plasmid varies in copy number from cell to cell and across a culture as it grows. A CEN/ARS vector, while considerably more stable, still segregates imperfectly and has to be held under continuous selection to be retained at all. We are careful not to overstate this: integration is a measurement-stability and genetic-stability choice, not a containment device and we do not present it as one.

Auxotrophic selection with no antibiotic resistance marker in the yeast. Selection is by histidine prototrophy, supplied by the HIS3 carried on the cassette, which complements the his3Δ1 of the parental strain. CEN.PK2-1C is auxotrophic for histidine, leucine, uracil and tryptophan, so the transformed strain retains three of those four requirements, leu2, ura3 and trp1. The companion document sets out what those three do and do not contribute to containment and is careful not to present them as a barrier to survival outside the laboratory. No antibiotic is used anywhere in the yeast workflow, which removes both the selective pressure for resistance and the reagent hazard. The *bla* marker is present only in the plasmid vector of one custom gene (Fragment 5) and serves *E. coli* propagation only. Transformation used linear PCR fragments that do not include it, although trace carry-over of that template plasmid cannot be excluded. The companion document sets out that reasoning and its remaining verification step.

A reporter with no biological activity. The output is fluorescence. It is not an enzyme, not a toxin, not a signalling molecule and does nothing outside the cell that produces it. This is the single largest contributor to the low hazard of the system and it is a consequence of the architecture, not an accident.

Identifying the dominant parameter uncertainty before the bench. Our model's sensitivity analysis showed that output amplitude depends most strongly on TIP–TetR affinity and an independent structural estimate suggested the true value could be substantially weaker than first assumed. We found this computationally, before committing experimental time, which is why it appears in this document as a planned measurement rather than as an unexplained result. The response shape (linearity) proved structurally robust to parameter perturbation in the same analysis. It is the signal magnitude that is sensitive.

A dedicated functional control. The copper-inducible strain exists to test the shared downstream module, the TIP–TetR–reporter chain, independently of whether stress sensing itself works. Without it, a null result would be uninterpretable: a failure to respond could lie anywhere in the chain. This is a deliberate piece of experimental design, not a spare strain.

3. Biological hazard: what the assessment excludes

This section records findings, not open questions. We state them positively because they are the substantive conclusion of the hazard half of this assessment.

No module encodes a hazardous function. There is no toxin, no virulence factor, no pathogenicity determinant and no antibiotic resistance gene in the integrated cassettes. Several parts are of bacterial or viral origin and we name them rather than leave the question hanging. TetR is the tetracycline repressor of the bacterial transposon Tn10, routinely propagated in laboratory E. coli. The SV40 nuclear localization signal is a seven-residue targeting peptide, not an intact virus or a transforming viral protein. GB1 is a small bacterial folding domain used throughout structural biology as a solubility tag. Two non-coding elements are also of bacterial or phage origin: the tet operator (*tetO*) sites in pGAL1-D12 come from the same Tn10 system, and the short attB linker sequences between parts in the assembled cassettes derive from bacteriophage lambda. Neither encodes a protein. Origin alone does not assign infectivity or hazard. The claim that carries weight is about function: none of these parts is a toxin, an effector or a virulence determinant. A part being bacterial or viral in origin says nothing on its own about whether it causes harm. We note this because the names involved can prompt the question.

The circuit confers no survival or growth advantage beyond its selection marker. The *HIS3* marker restores histidine prototrophy, which helps growth only where histidine is limiting. Nothing else in the architecture extends the host's environmental range, improves nutrient acquisition or increases stress resistance by design. The system detects stress. It does not defend against it. Apart from histidine prototrophy, the only phenotype the construct adds is fluorescence.

The chassis is well characterised and low hazard. S. cerevisiae is Risk Group 1. The companion document gives the full chassis assessment, including the honest limits of the auxotrophy argument and the point that autoclaving, not genotype, is what actually prevents release. The cycle is what gives that claim content: 121 °C under 1–1.5 atm for 15 minutes, with a chemical process-indicator label on every load that is read and confirmed before the load leaves the autoclave.

Dual-use potential is low, for structural reasons. ICHNOS is a recorder, not an effector: it makes an internal state externally visible and secretes no product that acts outside the cell. The modularity that is our central engineering claim deserves scrutiny, since a plug-and-play architecture is by definition repurposable, so we state the boundary precisely. Substituting the *sensing* promoter changes what the system detects and every plausible substitution yields another sensor. Producing harm would require replacing the *reporting* module with an effector payload, which is not a modification of ICHNOS but a replacement of the part that makes it ICHNOS. The recording architecture is not the difficult component of such a project, so publishing it offers no meaningful uplift. We identified nothing warranting restriction on safety grounds.

The limits of these findings. They are assessments of the intended design, supported by direct inspection of the delivered construct sequences as described in the companion document. They are not an exhaustive screen against toxin and pathogenicity databases. And none of them removes ordinary exposure pathways: spills, contaminated surfaces, aerosols, breakages and viable cells in waste are properties of working with a live microorganism and are unaffected by whether the circuit functions. Those are controlled by the handling, transport, inactivation and incident procedures in the companion document and Section 9 adds nothing to them.

4. Construct identity and host physiology

What is settled. Selection by histidine prototrophy worked as designed: transformants grew on histidine-free medium, which requires a functional *HIS3* gene. A no-DNA control reaction went through the same transformation and was plated on selective and on non-selective medium. What is not settled is the integration itself. Colony PCR across the left junction (12 October) gave no product, and the integrated constructs have not been sequenced, as the companion document records.

Cassette quality before transformation. The cassette fragments were amplified by PCR, purified on a silica column and quantified on a NanoDrop before transformation, with the 260/280 and 260/230 absorbance

ratios recorded alongside the concentration. It is worth being exact about what that establishes, because a spectrophotometric reading invites being over-read. The NanoDrop gives concentration and two purity ratios. It does not give fragment size and it does not give identity, so it cannot distinguish the intended amplicon from a product of the wrong length or from primer-dimer carryover. Nor do acceptable ratios certify a clean preparation even after a column clean-up, since several of the substances that would matter absorb weakly or not at all in that region. What the ratios do support is the narrower and still useful claim that the sample is reasonably free of the contaminants the clean-up is designed to remove, chaotropic salts above all, at concentrations that would interfere with transformation. The fragments were not run on a gel after amplification, so their length was not checked independently, and their identity has not been established by sequencing. Both remain open.

What remains to confirm and why it matters quantitatively. Histidine selection selects for a functional *HIS3*, not for a complete circuit. Because the construct was assembled by *in vivo* homologous recombination of overlapping fragments, a clone can carry the marker alongside a truncated or partially assembled payload. Our colony PCR across the left junction gave no product, so even an integration event at the *HO* locus has not been demonstrated, and sequencing would also be needed for internal architecture, copy number and absence of retained backbone.

The consequence is quantitative rather than hazardous. Clone-to-clone differences would be inherited by every dose-response curve and the fitted Hill coefficient, our central quantitative claim, would then describe the clone set rather than the circuit. Copy number matters for the same reason and in a sharper way: because the reporter and the TetR feedback copy share a promoter, additional copies do not simply scale output, they change loop gain and therefore the linearity itself. Copy number is a circuit parameter, not a housekeeping detail.

Planned verification. Sequencing across the integration locus in the working clone, with copy-number determination. We record this as the single highest-value outstanding measurement in this document, because several other interpretations depend on it. The HO status of the parental stock should be confirmed in the same pass. Integration at this locus disrupts HO, so in an HO-proficient background it removes the capacity for mating-type switching, while any cell in the stock that retains a functional HO could still switch, mate and diploidise during propagation, changing construct dosage. Confirming the status tells us which of those two situations we are in.

Host physiology. Expressing TetR, the TIP fusion and a two-fluorophore reporter consumes transcriptional, translational and folding capacity. There is a circularity worth stating plainly: a stress reporter imposes a load and load on proteostasis is the kind of stressor the ER module is built to detect, so in principle the circuit could report its own presence. The payload is cytosolic and nuclear rather than secretory, which makes a direct ER folding load unlikely, but “unlikely” is a hypothesis and the control that tests it is cheap: uninduced engineered cells against the parental and a marker-matched strain, with growth rate and viability recorded in the same experiment for the reasons in Section 7. Elevated baseline in uninduced cells is the signature. If growth or stress tolerance differed measurably from the parental strain, that would be a change in the organism’s characteristics and would send us back to the organism-level assessment. We do not expect it. The presence of a stress sensor gives no reason to assume it.

5. What the sensing modules report

All three inputs share one property that shapes interpretation, so we treat it once rather than three times.

5.1 Adaptation is a property of the pathways, not a defect of the circuit

Yap1, the UPR and copper homeostasis are all adaptive systems. The Yap1 response raises the cell's thiol-based antioxidant capacity, principally the glutathione and thioredoxin systems, which are the larger contributors to clearing peroxide in this organism. Yap1 does also drive catalase induction, but specifically in response to H_2O_2 rather than to oxidants in general. The principal regulator of the catalase genes is the Msn2/Msn4 general stress response rather than Yap1. The distinction matters for interpretation, because it means our oxidative module reports the Yap1 arm of the defence, the one proximal to peroxide and not the antioxidant response as a whole. The UPR restores folding capacity, with BiP tuning Ire1 sensitivity so that the response is adaptive at low and moderate DTT doses and sustained only at high ones. Copper homeostasis downregulates uptake and sequesters copper in metallothionein. In each case, when the cell succeeds, the stimulus disappears and the promoter switches off, while external exposure may continue.

The interpretive consequence. What the circuit records is the duration of the cell's unbuffered exposure, not the duration of exposure as such. This is a property of using native stress pathways as inputs. Any sensor built this way inherits it, which is precisely why we think the honest formulation of our output is that ICHNOS records episodes of unbuffered stress. Stated that way the claim is defensible and still useful. Stated as "total exposure" it would not be.

This matters most in the repeated moderate-exposure regime, which is the application we are interested in, so it is the first thing a characterisation should test: a sustained input against repeated pulses of comparable total dose. If they are indistinguishable at the readout, the temporal claim does not hold for that regime and we would say so.

5.2 Each promoter reports its pathway, not a chemical

The synthetic oxidative promoter reports Yap1 activity. Yap1 responds to thiol-reactive electrophiles, diamide, some heavy metals and lipid peroxidation products as well as to peroxide, so the supported statement is "Yap1 was active" and attributing a signal to a specific agent needs independent evidence. The KAR2 promoter is in effect a readout of ER folding health, since KAR2 encodes the yeast BiP chaperone and its induction tracks the folding load the compartment is carrying. That is what makes it informative and it is also why its specificity has to be stated carefully. The promoter carries heat-shock elements alongside the UPR element and is therefore inducible by thermal stress independently of the UPR, which makes incubator excursions, warm rooms and microscope-stage warming genuine inputs to that module. The practical answer is mundane and effective: log temperature and include a deliberate thermal control in the characterisation set.

DTT deserves one note because it cuts both ways. It was chosen for compatibility with the most detailed available kinetic literature, which is a good reason. But it reduces disulphides throughout the cell, so it perturbs the same cytosolic thiol balance the Yap1 module senses. As a reducing agent it could also suppress an oxidative signal. The two inputs are not orthogonal perturbations of the cell even though they are read by separate strains. Experiments combining or sequencing them should account for that. Two practical consequences follow from the same chemistry. The first is that DTT and hydrogen peroxide react with each other directly, so the two are never combined in a culture, in a stock or in a waste container. The companion document carries the pairing as a named incompatibility and here it is also the reason a combined-input experiment cannot be done by simply adding both. The second is that neither inducer is stable on the bench. Hydrogen peroxide decomposes and DTT oxidises over time, so a solution loses potency between preparation and use. The delivered dose is therefore uncertain in a direction that always understates the stress applied, which makes it a systematic error running through every dose-response curve rather than a random one. Fresh dilutions prepared from aliquoted stock on the day of use, with the preparation time recorded, are what keep the nominal dose close to the delivered one.

Baseline is growth-state dependent in both modules: respiring cells generate more mitochondrial ROS than fermenting ones and stationary phase mounts a general stress response overlapping Yap1 targets. This interacts with Section 6.3, since the reporter promoter requires galactose and galactose growth is more respiratory. Fixing growth phase and carbon source across compared conditions resolves it.

5.3 The copper control: what it establishes

As set out in Section 2, the separate copper strain is a deliberate control for the downstream module. Its boundary is worth stating so that it is not over-read: a successful copper response shows the downstream module works at the expression level copper produces, which may be well above what the stress promoters produce. Titrating copper to a comparable expression level, rather than using a single saturating dose, converts a qualitative control into a quantitative one. Copper tolerance also depends on endogenous CUP1 copy number, which varies between strains, so induction and toxicity thresholds should be established in CEN.PK2-1C rather than taken from a protocol written for another background. One point deserves more emphasis than a control section usually gives it. Copper is not an inert inducer. Cu^{2+} participates in Fenton-type chemistry and generates reactive oxygen species, which is the same property that Section 2 gives as the reason for keeping the copper strain away from the oxidative one. The choice of dose is therefore part of the control's validity and not a matter of convenience. We work between 10 and 100 µM, below the roughly 200 µM at which copper becomes toxic to this chassis, precisely because the cell has substantial capacity to handle copper at those levels: uptake is downregulated and metallothionein sequesters what gets in. Within that range the copper strain exercises the TIP–TetR–reporter module at an expression level the cell reaches without being oxidatively stressed into it. Pushed higher, the same strain would stop being a clean downstream control and start being a second oxidative input, which is the failure mode the separate-strain decision was taken to avoid.

6. Signal transmission: TIP and TetR

6.1 The fusion is the unit that matters

TIP is expressed as GB1–linker–peptide–NLS. Every property that matters (abundance, solubility, localization, binding) belongs to that fusion rather than to the peptide. Two size relationships follow from the design and are worth having thought about before anyone asks.

GB1 is roughly 6 kDa. The peptide it carries is under 2 kDa. The solubility tag is about three times the mass of its cargo, which is a reasonable trade, since the peptide sequence is hydrophobic and aromatic and would be expected to aggregate unassisted, but it means tag interference belongs on the list of candidate explanations if measured affinity comes in below prediction. Separately, the assembled fusion at roughly 9–10 kDa is well below the ~40 kDa threshold for passive diffusion through the nuclear pore, so it equilibrates across the envelope in both directions. The NLS biases that equilibrium rather than trapping. Nuclear concentration may therefore be closer to cytoplasmic than the design assumes. Since the effective signal depends on nuclear TIP concentration multiplied by affinity, this compounds with the affinity uncertainty rather than adding to it.

Consequence and response. Weak transmission compresses the dynamic range and raises the limit of detection. Within our own research that means a harder measurement. The mitigation is the one already planned: direct measurement of the interaction where instrumentation allows and otherwise a functional proxy, the copper strain titrated across TIP expression levels, which bounds the input–output relationship without requiring purified protein. Localization is answered by imaging the fusion. If affinity is confirmed weak, targeted mutagenesis to strengthen binding is the designed next step, which is why finding this computationally rather than experimentally was worth doing.

One measurement caveat to carry into interpretation: tracking expression through the GB1 portion of the transcript shows the message is made. Full-length protein, correct localization, solubility and binding competence are four further claims that the transcript does not support.

6.2 TetR and the loop

TetR functions as a homodimer, so whether inactivation requires one bound TIP per dimer or two sets the cooperativity of that step and works with or against the intended linearization. It is an assumption of the model and should be labelled as one.

The loop's speed is set by how fast TetR disappears (degradation plus dilution) and TetR is not rapidly turned over. A slow loop means reporter production continues for some time after the input has gone, widening the production window. That is the specific reason this section connects to the timer: Section 7 depends on the production window being short and known. The comparison that answers it is a recovery time course after a defined pulse, with and without the feedback arm. A steady-state comparison alone demonstrates linearization but reveals nothing about the time constant.

6.3 The carbon source is an input to the circuit

pGAL1-D12 is glucose-repressed and galactose-induced. This is a hard operating constraint rather than a subtle effect: in glucose-containing medium the reporter arm is off regardless of stress, so routine rich medium is not a valid readout condition. Three consequences follow. Residual glucose or diauxic transition during a long time course moves output independently of input. Galactose-grown cells differ metabolically, with a different ROS baseline and growth rate. And because one promoter drives both reporter and feedback TetR, a carbon-source shift moves both arms and changes loop gain. A fixed, documented carbon source recorded as an experimental parameter resolves all three. It is worth recording one further operating fact plainly, because it is easy to state loosely: the reporter arm requires two conditions at the same time, the absence of glucose and the presence of galactose. Neither on its own is sufficient to read the circuit, so the carbon source has to be specified as a pair of conditions rather than as a single ingredient. The requirement is therefore conjunctive rather than additive. It is worth naming as such, because the two conditions are controlled by different mechanisms: glucose repression acts through the Snf1/Mig1 route and on GAL4 levels, while galactose induction acts through Gal3 relieving the Gal80 inhibition of Gal4. In a non-repressing carbon source without galactose the reporter stays off, so the two are independently settable rather than one variable seen twice. Taken together with the TetR node, output requires the absence of glucose, the presence of galactose and an inactivated TetR. We describe this as a property of the GAL1 promoter we inherited rather than as a logic gate we designed, since we have not characterised the behaviour across those states. Whether the D12 variant retains the direct glucose-repression elements or depends instead on repression at the level of the regulon should be read off the annotated sequence before any mechanistic claim is made about where the repression acts.

For completeness in any future application rather than for present risk: tetracyclines act on TetR and would derepress the reporter independently of stress. Irrelevant in defined culture. Relevant immediately if the system were ever applied to a sample of unknown composition.

7. The timer and measurement integrity

7.1 What the ratio reports

sfGFP matures quickly and mCherry slowly, so a freshly synthesised reporter population is predominantly green and shifts red as the same molecules age. The mCherry/sfGFP ratio reports the **average age of the existing reporter population**. Between a stress event and that population lie transcription, translation,

fluorophore maturation and a production window set by the feedback loop, so reporter age maps onto time-since-stress only when that window is short and known.

Reported maturation times and why ours are expected to differ. The published values for this timer pair, measured at 30 °C, are a maturation half-time of about 5.6 minutes for sfGFP and two sequential steps of roughly 17 and 30 minutes for mCherry (Khmelinskii et al., Nature Biotechnology, 2012). Those figures are what make the ratio interpretable in principle and the model takes them as its starting point. They are not constants of nature. Fluorophore maturation is temperature dependent and oxygen dependent. Both are experimental conditions rather than fixed properties, so a different incubation temperature moves the whole calibration. More to the point for us, the cells being measured are under stress by design. Stress alters translation, proteostasis and growth rate, each of which bears on how fast a newly made fluorophore matures and how long it persists. The times governing our own results are therefore expected to differ from the published ones. The position we take is that the literature fixes the form of the relationship while our own calibration has to fix the numbers.

7.2 Growth rate couples the two channels

This is the most important item in the document and the one we would most want the team able to discuss unprompted.

The ratio reflects how fast the existing population is replaced. In a dividing cell, protein is removed by two processes: degradation and dilution by division. For a stable fluorescent protein in a growing yeast culture, dilution typically dominates. This is not a flaw we introduced. The published tandem-timer formalism accounts for turnover including dilution, which is why the method is normally applied under controlled growth.

For ICHNOS the implication is specific. **Stress slows growth.** A stressed culture dilutes its reporter more slowly and therefore shows an older-looking population at the same true elapsed time. So the project's two readouts, intensity for how much and ratio for how long ago, are not independent unless growth rate is held constant or measured and corrected for.

We accept three consequences, which are design requirements rather than obstacles:

- Growth rate becomes a **required measurement**, recorded alongside every fluorescence measurement.
- Valid temporal comparisons are restricted to **matched growth conditions**.
- Any claim of the form "this stress occurred *n* hours ago" must be demonstrated at known growth rates rather than inherited from the published formalism, which was calibrated for a different question.

Characterising the timer at deliberately varied growth rates is the experiment that converts this from a confound into a calibration.

7.3 Interpretation and artefacts

Non-uniqueness. One ratio value can arise from several input histories: a recent weak episode and an older strong one can converge, as can a single pulse and a series of smaller ones. This is a property of summarising a multi-dimensional history in one number. Until identifiability is demonstrated for the histories of interest, the defensible endpoint is a **bounded classification, "recent"** versus "not recent", with stated bounds rather than an elapsed time. Histories the readout cannot separate should be reported as ambiguous. Delivering known pulsed, sustained and repeated inputs tests this directly.

Physical confounders. Fluorophore maturation requires molecular oxygen, so aeration and culture density are timer parameters rather than incidental conditions. Because the two fluorophores mature at different rates, oxygen limitation does not cancel in the ratio. The ratio is also only meaningful while the fusion

remains intact and is removed as one unit, so linker integrity should be checked rather than assumed. The FRET correction used in the model is taken from the literature rather than measured on this construct, while linker length, fluorophore order and instrument settings all affect it. It is a modelling assumption and should be labelled as one. FRET deserves more than a correction factor here, because in this particular pair it is entangled with the quantity being measured. sfGFP is the donor and mCherry the acceptor. Only a matured mCherry chromophore can accept the transferred energy. Transfer efficiency therefore rises as the mCherry matures, which is exactly the process the timer uses to tell age. Energy transfer suppresses donor emission and adds to acceptor emission, so FRET moves the measured ratio in the same direction as ageing does and a single fixed correction cannot separate the two. Treating the correction as a constant builds part of the signal into the baseline. The consequence we accept is that the FRET contribution has to be bounded experimentally on this construct, with a donor-only strain and an acceptor-only strain under the same acquisition settings, before any elapsed-time statement is allowed to rest on the ratio.

Optical confounders and the control set. Spectral bleed-through, autofluorescence (predominantly green, so it biases the ratio in one direction), photobleaching and unstable ratios where both signals approach background. Widefield imaging gives field-level rather than single-cell measurements, so a field average cannot distinguish a bimodal population from a uniform intermediate one, a limitation to state wherever heterogeneity is discussed. The control set that addresses these is standard and we intend to run it in full: parental autofluorescence, each fluorophore alone, the uninduced tandem reporter, fixed acquisition settings, documented channel correction, predefined exclusion rules for saturated and near-background measurements and retention of raw images with metadata. Biological replicates are independent cultures, not repeated images of one culture.

The stressed parental control. One control belongs in that list explicitly rather than by implication: the parental strain, carrying no cassette, put through the same stress treatments as the engineered strains. Measuring unmodified cells at rest establishes the autofluorescence baseline, but stress itself changes cellular autofluorescence, largely through flavin and NAD(P)H pools that shift with redox state and metabolism. Those changes fall mainly in the green channel where they bias the ratio in one direction. Stressing the parental strain across the same dose range is what separates a change in the reporter from a change in the cell's own background and it costs almost nothing to run alongside the real experiment.

7.4 Viability and the readable range

Reporter synthesis requires a living, translating cell, so at high doses the dose-response curve turns over: signal rises with stress, then falls as the cells that would report it lose the capacity to.

A low reading is therefore compatible with low stress and with overwhelming stress. In our own experiments this costs an experiment if unnoticed and the remedy is straightforward: measure viability at every dose rather than only where something interesting happens, define the assay's operating range as the range over which viability holds and write the invalidity rule before seeing the data. Results outside that range are reported as invalid rather than as low.

What the instrument is for, so that the operating range is not read as a shortcoming. ICHNOS is not intended to be a marker of cell death. An assay that responded only when cells were dying would be easier to build and considerably less useful, since markers of apoptosis and necrosis already exist and report an endpoint that has already passed. What we want is the range in which stress still matters to the cell and to us: where the response is engaged, the cell is adapting and a record of that adaptation is still being written. The turnover of the curve at high dose is therefore the edge of the useful range rather than a defect in the sensor. Doses beyond it fall outside the question the instrument was built to answer. Defining the range is how we state which measurements count, not an admission that the rest nearly did.

We flag this with particular care for one reason. If an output of this kind were ever used to support a claim that something is safe, which is not the present context and which Section 12 keeps out, this is the mechanism

by which it could be wrong in the reassuring direction. Writing the operating range into the method now is what keeps that boundary enforceable later.

8. Genetic stability

The decision to abandon the combined design (Section 2) removed the clearest recombination substrate and we regard it as the right call. The final design is not free of repeats: the reporter and the feedback TetR copy share the same pGAL1-D12 promoter sequence, about 445 bp that is identical in the two copies and lies in inverted orientation in the registered cassette. The three terminators (*ADH1*, *CYC1* and *TEF1*) are different sequences. The actual extent of the repeats should be assessed from the sequence of the integrated construct rather than from design intent, which the sequencing in Section 4 would deliver in the same pass.

Selection maintains *HIS3*, not the circuit. A cell that has lost or rearranged the reporter while retaining the marker grows well on selective medium and, carrying less expression burden, may outgrow intact cells over successive passages. The drift direction is toward **loss of signal**. The control is a defined maximum passage number before results are compared against an authenticated frozen stock, with genotype and functional response re-verified after a set propagation history. A stock that fails requalification is discarded rather than recalibrated.

9. Exposure and containment

The mechanism adds nothing to the containment picture established in the companion document and Section 3 explains why: there is no bioactive output and no mobilisable element. What the mechanism also cannot do is remove the ordinary exposure pathways of working with a live microorganism. A broken circuit in a spilled culture is exactly as much of a containment question as a working one. Handling, closed transport, autoclaving on the cycle given in Section 3 and incident response are covered in the companion document and are the measures that actually apply.

10. Biosecurity and data integrity

The dual-use assessment is in Section 3 and is low for structural reasons. What remains here is proportionate custody and one observation that follows from the rest of this document: because the realistic harm pathway for an instrument runs through a wrong number rather than a dangerous organism, calibration data, raw images and analysis versions deserve the same care as the strains.

- Strains and stocks in access-controlled storage with unique IDs recording construct version, integration locus, date and custodian. Transfers recorded with the recipient's intended use.
- Raw images, calibration data and analysis scripts retained with version history. Changes affecting reported results documented.
- Engineered yeast is not used in public demonstrations and is not released outside the laboratory.
- Any change of host, promoter, cargo, scale, sample matrix, intended user or decision context requires a new assessment.

On confidentiality. Some mechanistic details are currently held back pending consideration of intellectual property. We distinguish this from biosecurity explicitly, because the two are easily conflated: withholding ordinary technical information is an IP decision, not a safety measure and we do not present it as one. Part sequences will be published consistent with iGEM's openness principles once that question is settled.

11. Risk register

Priority is deliberately differentiated: a register in which everything is high priority is not a prioritisation. It combines how much rests on the answer with how cheaply it can be obtained. Entries marked (d) are addressed by a design decision already taken (Section 2) and carry the residual only.

ID	What could go wrong	Consequence	How we address it	Priority
M1	Growth rate couples the intensity and timer channels	The two readouts are not independent. Temporal claims confounded by stress intensity	Growth rate recorded with every measurement. Timer characterised at matched and deliberately varied growth rates	1
M2	Integration at the *HO* locus, construct architecture, copy number or retained backbone unconfirmed	Clone variability read as biology. Loop gain and therefore linearity unknown	Sequencing across the integration locus in the working clone, with copy number and *HO* status	1
M3	Cells too damaged to report	Low reading compatible with overwhelming stress as well as none	Viability at every dose. Operating range defined. Invalidity rule written before data	1
M4	Carbon source gates the reporter promoter	Reporter silent in glucose. Medium variation appears as response variation	Fixed documented carbon source, recorded as an experimental parameter	1
M5	Pathway adaptation truncates the recorded episode	Sustained exposure recorded as brief, in the cumulative-exposure regime we target	Sustained versus pulsed inputs at matched total dose. Output framed as unbuffered-stress episodes	2
M6	TIP fusion underperforms: affinity, nuclear retention or aggregation	Compressed range and raised detection limit	Localization imaging. Affinity measurement or copper-titration proxy. Mutagenesis as the planned response	2
M7	Timer readout non-unique or distorted by optics and production window	Distinct histories reported as one elapsed time	Pulsed, sustained and repeated inputs. Recovery time course with and without feedback. Full optical control set	2
M8	Promoter specificity uncharacterised	Stress source misattributed. Baseline drift read as signal	Input panel per module. Thermal control for pKAR2. Fixed growth phase and carbon source	2

ID	What could go wrong	Consequence	How we address it	Priority
M9	Circuit burden induces the response it measures	Baseline signal attributed to external stress	Uninduced engineered cells against parental and marker-matched controls, with growth and viability	3
M10	Genetic drift over propagation (d)	Signal loss across passages, in the false-negative direction	Passage limit. Requalification against frozen stock. Combined design already rejected for this reason	3
M11	Accidental exposure or release of viable cells	Contamination. Unassessed environmental presence	Companion document: handling, transport, validated inactivation, incident response	3
M12	Stock or data integrity failure	A result traced to the wrong strain or an analysis that cannot be reproduced	Inventory reconciliation. Raw data retention. Version history	3
M13	Extrapolation beyond the yeast proof of concept	A claim the data do not support	Section 12 limits maintained in every output	2
M14	Inducer potency drifts between preparation and use	Delivered dose below nominal, biasing every dose-response curve in one direction	Fresh dilutions from aliquoted stock on the day of use, preparation time recorded. H_2O_2 and DTT never combined	2

Within contained research, M1 through M10, M12 and M14 principally affect data quality. M11 is the only direct exposure pathway and is managed under the companion document. M1, M3, M6 and M13 are the entries that would acquire safety significance in any future context where a reading supported a claim about whether something is safe.

12. Limits of use

These apply now and accompany any presentation of results.

- ICHNOS reports the response of engineered yeast under defined laboratory conditions. It does not report what a human tissue would experience.
- *Yap1* is the yeast oxidative stress pathway. The corresponding human pathway, Keap1–Nrf2–ARE, differs in activation mechanism and in chassis. Translation is an open research question, not a configuration change.
- A low reading is not evidence of a safe sample, for the reasons in Sections 6.1, 7.4 and 8.
- Temporal output is reported as a bounded classification with stated conditions, not as an elapsed time, until identifiability and growth-rate independence are demonstrated.

- No result from this system informs a product, clinical, agricultural or environmental decision. Each such application would require its own assessment, validation against an established method and institutional acceptance.
- Any change of host, promoter, reporter, marker, insertion site, scale, sample matrix or intended user invalidates this assessment.

Triggers for immediate review. A change to any construct element or to the host. Unexpected growth, survival, baseline activation or loss of response. Genotype drift or calibration failure. A spill, exposure, failed inactivation, missing stock or unexplained data change. Any proposal to use an output outside contained research.

Review record. Technical review by the team's supervisor, institutional safety review and residual-risk acceptance are pending and are recorded here as such, so that this document is not cited as though they had occurred.