Endogenous RNA-Triggered CRISPR Gates

1. Summary

CRISPR systems let a protein be aimed at a chosen site in a genome. A guide RNA carries a 20-nucleotide spacer that matches the target, and the protein follows wherever the spacer points. The protein used here is dCas9, which does not cut: fused to an activator or a repressor, it turns the target gene up or down. Its limitation is that it is always on. Once the guide is in the cell the protein acts, whatever else is happening around it.

This section describes a CRISPR gate that removes that limitation - a guide that acts only when a chosen RNA is present in the same cell. The guide carries, besides its spacer, a sequence that folds back and pairs with it. In that folded state the spacer is occupied and the protein cannot use it: the gate is OFF. A trigger RNA complementary to the fold invades it, the structure opens, the spacer is released and the protein acts: the gate is ON. Because the trigger can be any transcript the cell makes, the gate converts the presence of an RNA marker into the regulation of a gene.

We built it in three generations. The first reproduced the published architecture on our own target and trigger and added two controls of our own design - a gate built to be as closed as the architecture allows, and one that cannot close at all - so that any measurement could be read between a floor and a ceiling. The three constructs did not separate from one another, which pointed at the closed state. The second generation added a second lock, a module that pairs the constant scaffold the protein binds, together with a longer trigger-binding window, in four graded strengths on each of two backbones. That lock tightened the fully closed gate but not the switch, placing the remaining leak in the hairpin - the one part both generations had inherited unchanged from the published design, where it had been chosen for a different spacer and a different trigger.

The third generation is the result reported here. Rather than adjust an inherited design by hand, it states the problem the gate poses and solves it: given a gene to regulate and a set of candidate marker transcripts, which spacer and which trigger should be paired, and what molecule should be built around them. The algorithm works on two levels. The inner level designs the molecule for one spacer–trigger pair, choosing the blocking sequence, the extensions, the loop and the scaffold lock to maximise an objective that rewards an open spacer when the trigger is present, penalises an open spacer when it is absent, and rewards strong trigger binding - each quantity read from the equilibrium ensemble of the molecule rather than from a single predicted structure. The outer level ranks the pairs themselves, multiplying how good a spacer–trigger pair is by the best switch that pair allows, so that a promising target cannot compensate for a molecule that does not switch. The output is a ranked list of complete designs (Figure 1).

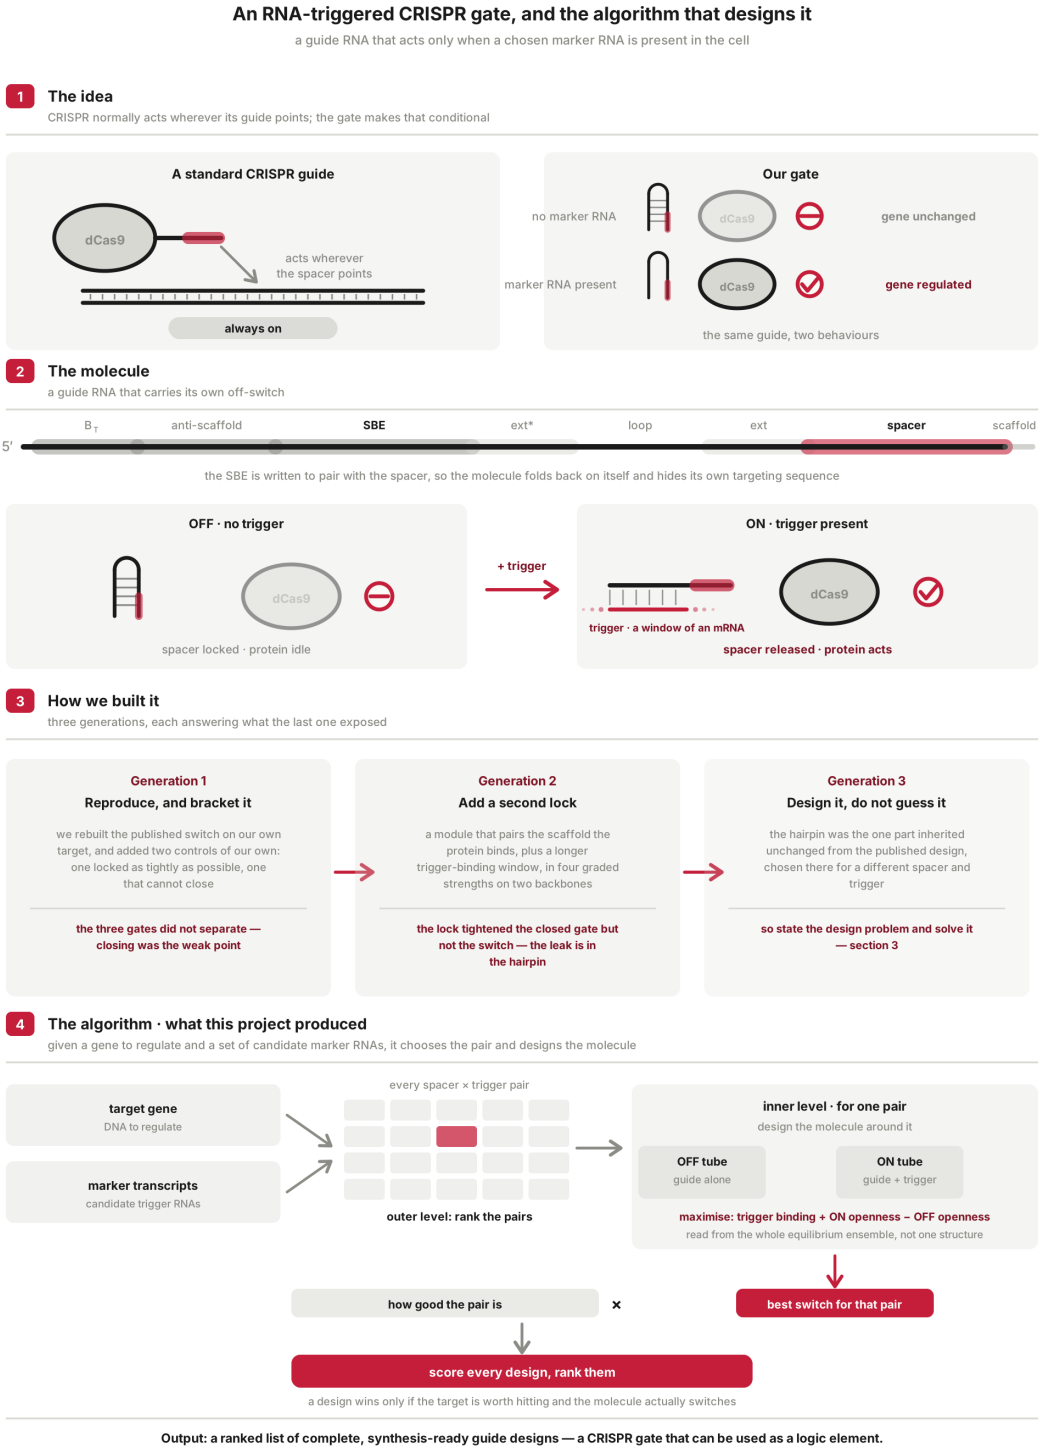

Figure 1. The project in one view: why a conditional CRISPR gate is needed, the molecule and its two states, the three generations of design, and the two-level algorithm that the work produced.

2. Introduction

2.1 Why control CRISPR at the guide

A circuit that distinguishes between cell states must read specific RNA markers and convert them into an action. Acting on the guide RNA is the least invasive point of control in a CRISPR system: no extra promoter, no transcription factor, no intermediate protein. The requirement is that the guide stay inactive while the marker is absent and become active once it appears. The marker is a natural transcript and cannot be edited, so the guide must be built around whatever sequence it happens to have.

2.2 Signal conversion: what the gate does

The gate converts an RNA-level input into a DNA-level regulatory output. The input is a trigger RNA, an endogenous transcript associated with a defined cellular state; the output is dCas9 activity at a predefined DNA target, producing transcriptional activation or repression.

The system separates into three layers (Figure 2), and this separation is what makes it reprogrammable:

1. **Input recognition** - set by the trigger-binding region of the engineered guide.
2. **Guide activation** - set by the trigger-induced structural transition.
3. **Output regulation** - set by the spacer, the DNA target site, and the dCas9 effector domain.

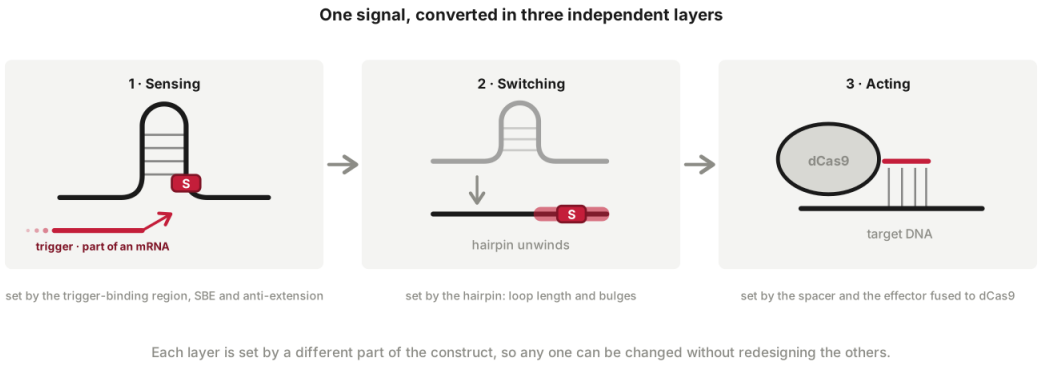

Figure 2. Sensing, switching and acting, with the sequence element that sets each layer.

2.3 How the gate is built, and how it switches

The parts of a native guide. A native single-guide RNA has two components. The *spacer* is the 20-nt sequence complementary to the CRISPR target site, and determines where dCas9 is directed. The *scaffold* is the structured remainder, which binds and stabilises the protein. A Cas nuclease can act only where the target site is immediately followed by a *protospacer-adjacent motif* (PAM), a short DNA sequence of typically two to six base pairs; this work uses SpCas9 throughout, whose PAM is 5′-NGG-3′. PAM availability is therefore what constrains where a spacer may be chosen at all. An iSBH-sgRNA extends the 5′ end of such a guide with a loop and a sequence partially complementary to the spacer, so that the guide folds back on itself and that fold holds the spacer out of use until a trigger arrives (Figure 3).

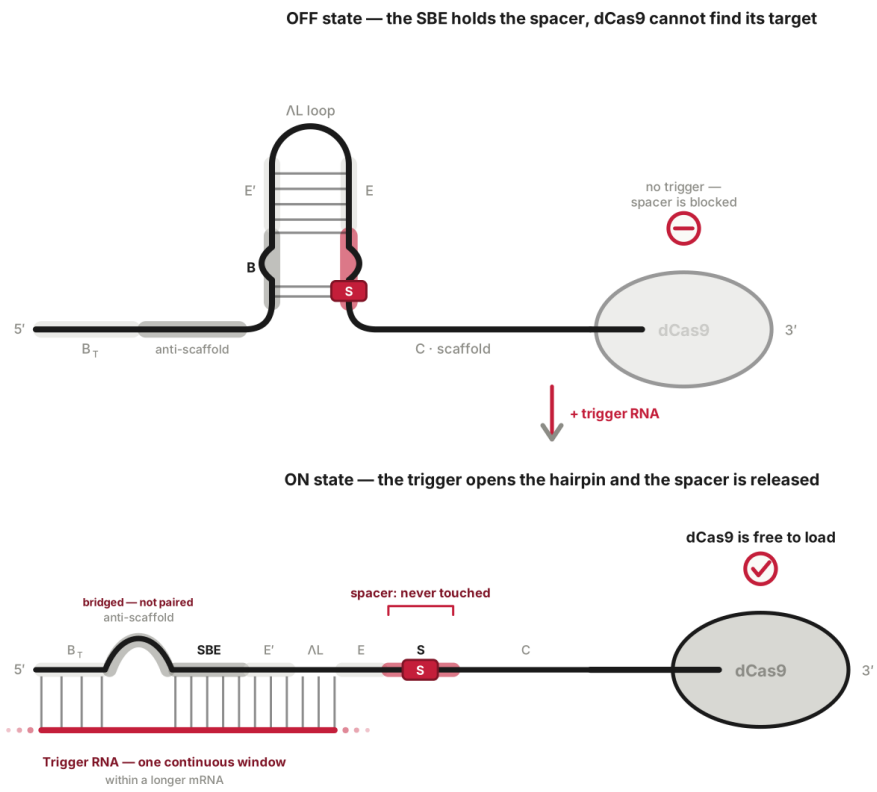

Figure 3. The iSBH-sgRNA in its OFF and ON states. A short anti-scaffold module sits between the trigger-binding region and the SBE. The trigger pairs with the trigger-binding region, the SBE, the anti-extension and the loop, bridging over the anti-scaffold; the spacer is released without ever being contacted.

The OFF state. The engineered sequence - the Spacer-Blocking Element (SBE), written spacer* in the source literature - base-pairs with the spacer, so the spacer is not accessible for DNA recognition and the guide cannot direct dCas9 productively, even where dCas9 is present in the cell. Intentional mismatches at defined positions tune the stability of the closed form. The OFF state must be stable enough to keep leak low while remaining openable by the trigger; that trade-off is the subject of stage 3.

The ON state. The trigger binds the accessible outer region of the closed structure - the complementary extension, the loop, and part of the SBE - then drives strand displacement and branch migration inward, progressively destabilising the closed conformation (Figure 4). As the outer stem opens, the constraint on the spacer is relieved and the spacer becomes accessible. Once exposed, it forms a functional complex with dCas9, which is directed to a DNA sequence complementary to the spacer and adjacent to a compatible PAM.

The point that is routinely mischaracterised: **the trigger does not bind the spacer and does not compete with the SBE for it.** Spacer exposure is an indirect consequence of outer-stem opening.

This dictates where the trigger-binding interface may be placed - in the outer structural region, not where the SBE pairs with the spacer.

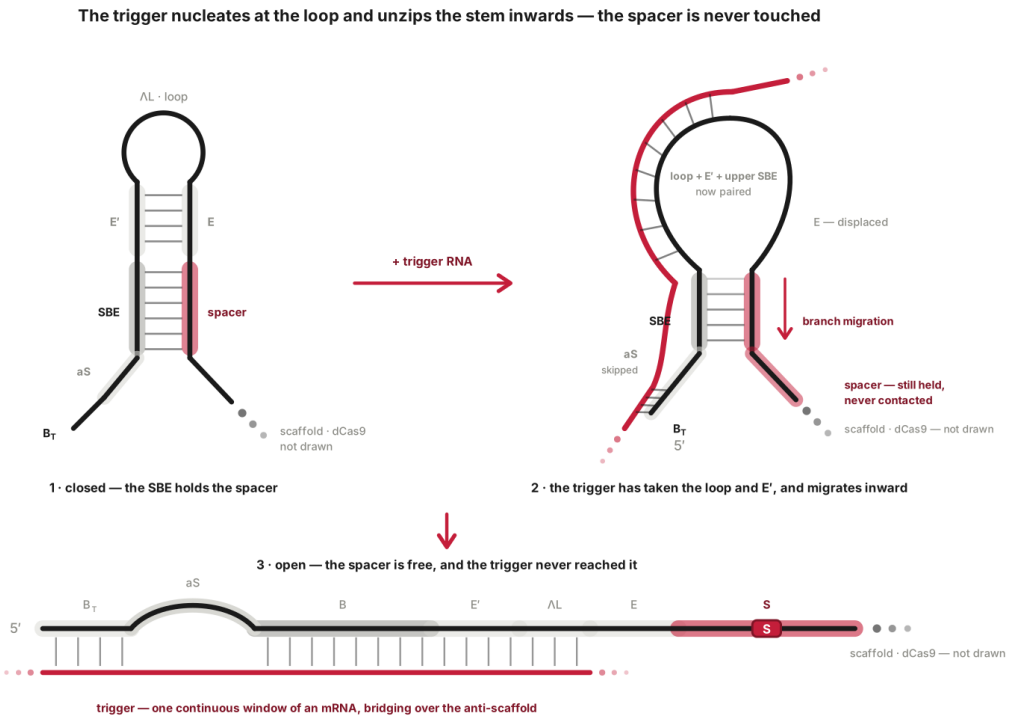

The trigger nucleates on the single-stranded loop, pairs E′ and the SBE, and unzips the stem inwards; the trigger-binding region anchors it at the 5′ end.
The spacer is released because its partner is taken, not because the trigger binds it. Faded dots mark chain that continues beyond the drawing.

Figure 4. Strand displacement in three steps. The trigger nucleates on the single-stranded loop, pairs the anti-extension and the SBE, and is anchored at the 5' end by the trigger-binding region, skipping the anti-scaffold; its footprint never reaches the spacer. Faded dots mark chain that continues beyond the drawing - the scaffold and dCas9 on the guide, the rest of the transcript on the trigger.

2.4 Logical behaviour

The logical identity of the gate is fixed before its molecular architecture is described, because the same architecture yields two different logical outputs depending on the effector. The gate's logical output depends on whether CRISPR activity or downstream gene expression is the measured quantity, since a CRISPRi configuration inverts the gene-expression output relative to CRISPR activity.

Trigger RNA (input)	iSBH-sgRNA state	dCas9 targeting	Target gene expression (output)
0 (absent)	Closed - spacer masked by the SBE	None	0 (basal)
1 (present)	Open - spacer exposed	Active	1 (activated)

Table 1. CRISPRa gate - trigger-dependent gene activation.

Trigger RNA (input)	iSBH-sgRNA state	dCas9 targeting	Target gene expression (output)
0 (absent)	Closed - spacer masked by the SBE	None	1 (expressed)
1 (present)	Open - spacer exposed	Active	0 (repressed)

Table 2. CRISPRi gate - trigger-dependent gene repression.

The gate itself behaves identically in both configurations: trigger present means spacer exposed. At the level of gene expression, CRISPRa implements YES logic and CRISPRi implements NOT logic (Figure 5).

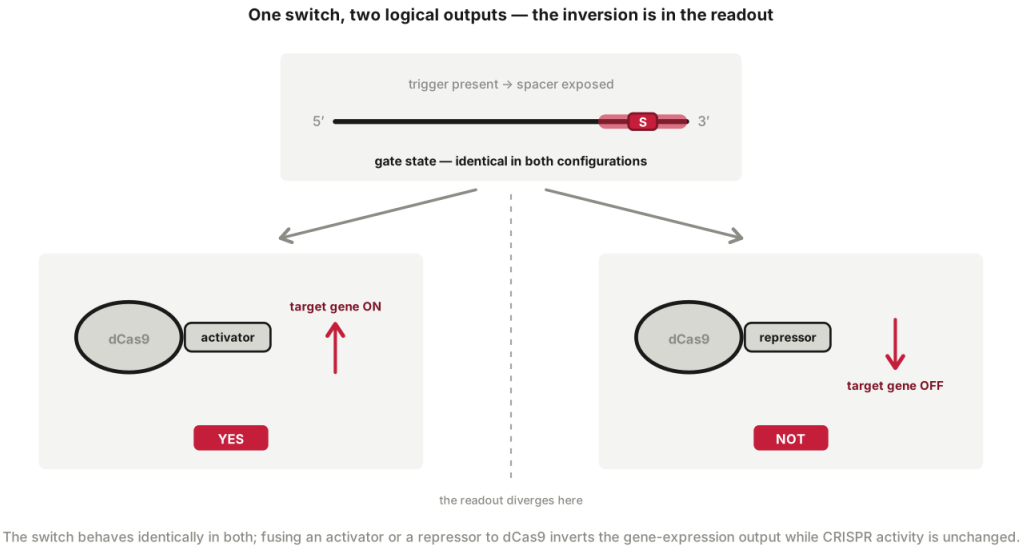

Figure 5. The same switch with an activating or a repressing effector; the logical inversion is in the readout, not in the gate.

2.5 The four variants of the gate

The gate has four configurations (Figure 6), set by two independent choices: the effector, a dCas9 fused to an activator or a repressor, which fixes the direction of the output as in section 2.4; and the DNA target, the organism's own genome or a plasmid carrying a synthetic construct.

The target choice is the one that changes the design problem, because it decides whether the spacer is an output of the pipeline or an input to it. With an **endogenous target** the gate acts on a site in the host genome, so the pipeline receives the target DNA and must find a spacer for it - PAM scan, candidate generation, on-target scoring, off-target checking against the reference genome. With a **plasmid target** the gate acts on a synthetic construct, for example a plasmid expressing GFP under defined conditions; the spacer is an engineered sequence with a known ON target, so the pipeline receives it rather than designing it, and only the off-target check remains. The trigger is an endogenous transcript in both cases: what differs is where Cas9 is sent, not what the gate listens to.

Endogenous target	Plasmid target

	Endogenous target	Plasmid target
Inputs	Target DNA sequence; RNA trigger sequence; Cas9 variant; organism reference DNA for off-target checking	Engineered spacer (known ON target); RNA trigger sequence; organism reference DNA for off-target checking
Spacer	Designed by the pipeline	Supplied as an input
Output, CRISPRa effector	Activating iSBH-sgRNA sequence	Activating iSBH-sgRNA sequence
Output, CRISPRi effector	Repressing iSBH-sgRNA sequence	Repressing iSBH-sgRNA sequence

Table 3. The four gate configurations and what the design pipeline takes in and returns for each.

This section covers the activating endogenous configuration in stages 1 to 3, and the plasmid configuration in design only, in stage 4.

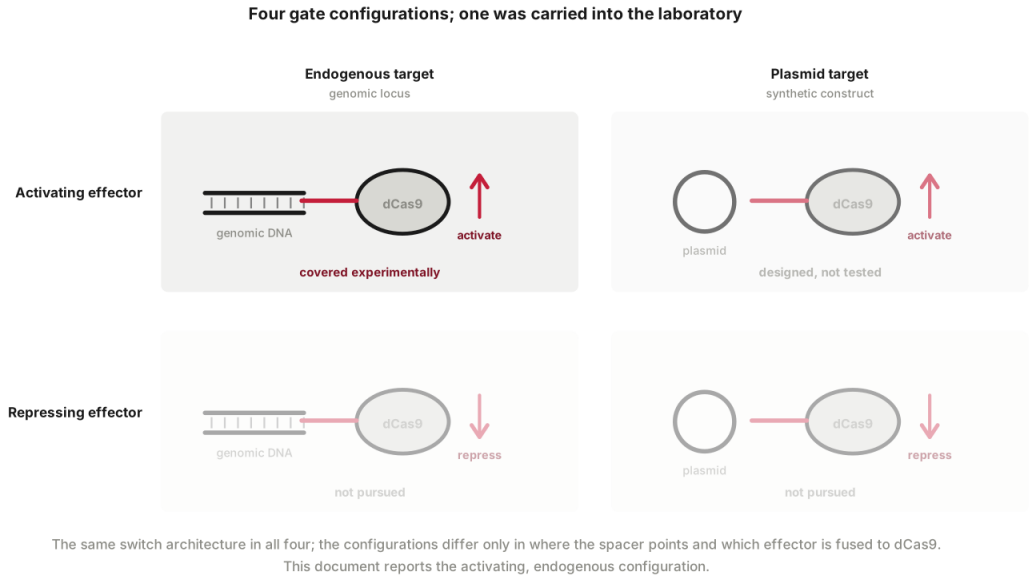

Figure 6. The four configurations of target and effector. Only the activating, endogenous configuration was carried into the laboratory.

2.6 Prior art

The architecture originates with Ferry et al. and was developed by Pelea et al., who demonstrated RNA-triggered iSBH-sgRNAs in HEK293T cells and zebrafish embryos, ran a design-build-test cycle across first-generation, second-generation and modular designs, and released a design pipeline. Two things in that work define the starting point here: residual activation in the OFF state, which the authors attributed to a fraction of molecules not adopting the intended fold; and the absence of any treatment of how Cas9 binding itself affects the closed structure.

2.7 How we got here: two engineering cycles

The framework in section 3 was not designed in the abstract. It is the output of two complete design–build–test cycles (Figure 7), each closing on a measurement that changed the next design.

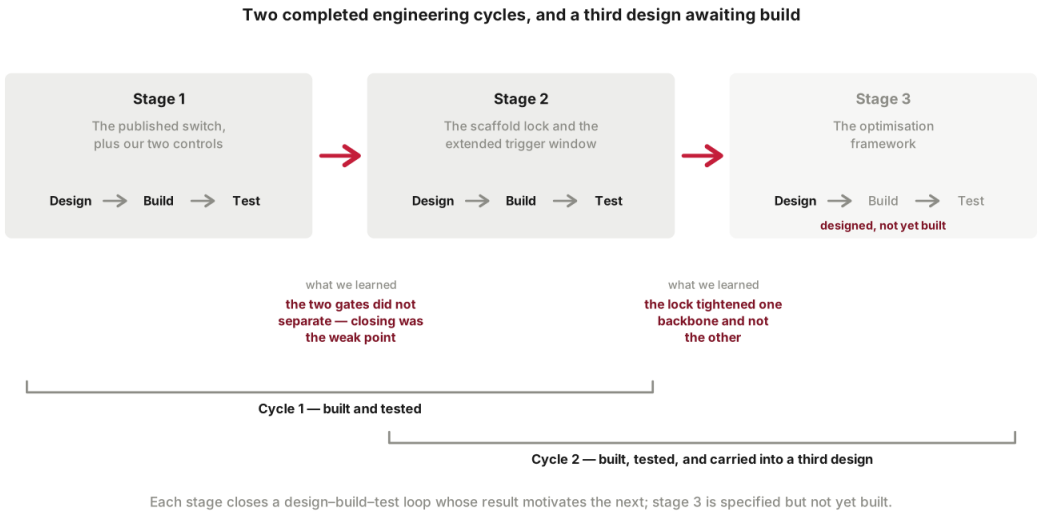

Figure 7. The three stages as design–build–test loops, each with the finding that motivated its successor. Two cycles were completed; stage 3 is designed.

The first cycle reproduced the published switch on our own spacer and trigger, and added two controls of our own: a gate built to be as closed as the architecture allows, and one that cannot close at all. Without them a single switch is uninterpretable, since there is no floor and no ceiling to read it against. The three constructs did not separate from one another, and our reading was that the gates were not closing tightly enough.

The second cycle acted on that reading. It added a second lock - an anti-scaffold that pairs the dCas9 scaffold *in cis* - together with a longer trigger-binding window, in four graded variants on each of two backbones. The lock tightened the fully-closed gate but not the switch, which still opened without its trigger.

That placed the remaining leak in the hairpin rather than the scaffold, and the hairpin is the one part neither cycle had varied: its bulge pattern was inherited from the published design, chosen there for a different spacer and a different trigger. Section 3 states that design problem explicitly and solves it. The two cycles are set out in full, with the designs they produced, in the Methods.

3. Results: a two-level optimisation framework for endogenous CRISPR gates

What this project produced is a design framework, and that framework is the result reported here. It states the endogenous gate as an optimisation problem, defines the quantities that score a candidate, and fixes the constraints that admit one. Its weights and thresholds are carried as free parameters rather than fixed by assumption, to be calibrated against measured designs as the project accumulates them (section 3.8); that calibration is part of the next steps in section 4.5. The algorithm is given first in outline, as pseudocode, so that the shape of the search is visible before the quantities inside it are defined; sections 3.2 to 3.9 then develop each piece in turn.

3.1 The algorithm in outline

The algorithm compiles an endogenous DNA target and a set of pre-selected trigger RNAs into a ranked list of engineered guides, and it is a nested search. Scanning the target DNA gives many candidate spacers and scanning the transcripts many candidate trigger windows, so the **outer** problem is which spacer–trigger pair to use. For a fixed pair, the **inner** problem is what the best guide for it is: it enumerates the architecture and the blocker sequence, folds each candidate, and keeps the one whose OFF→ON switch scores highest. A pair cannot be judged without knowing how good a guide can be built for it, so the inner problem is solved for every candidate pair and its optimum becomes one factor in the outer score. That dependency is what makes the formulation bilevel rather than a single ranked list (Figure 8).

Three quantities are read as given rather than computed here: the on-target activity $\eta_{\text{on}}(S)$ and off-target risk $\eta_{\text{off}}(S)$ of a spacer, which come from the genomic scan, and the trigger score $Q_T(T)$, which comes from the trigger-selection stage of the CERNAL pipeline (section 3.7).

Outer loop - choose the spacer–trigger pair.

```
PROCEDURE Main(D, window, triggers, C, s_C) → ranked designs
    spacers ← ScanSpacers(D, window)      # SpCas9, NGG, both strands
    results ← empty list

    for each spacer S in spacers :
        if η_off(S) > o_max :  skip S              # outer off-target gate
        Q_S ← v_on·η_on(S) + v_off·(1 − η_off(S))  # spacer quality in [0,1]

        for each trigger T in triggers :
            (J*, g*) ← InnerDesign(S, T, C, s_C)   # best guide for the pair
            if J* = −∞ :  skip T                   # no feasible guide

            Φ ← (W_S·Q_S + W_T·Q_T(T)) · max(0,J*) # outer pair score
            append (S, T, g*, Φ) to results

    return top-k of results ranked by Φ
```

Inner loop - design the best guide for one pair. The inner search set is the architecture - the loop length L and the trigger-binding tail length bt - times the anti-scaffold choice, times the blocker completions. Its feasible subset $F(S,T)$ is the designs that pass the hard gates, and the guide returned is the maximiser of J over that subset.

```
PROCEDURE InnerDesign(S, T, C, s_C) → (J*, g*)   # max of J over F(S,T)
    T_c ← revcomp(T)            # trigger complement = sensing template
    rcS ← revcomp(S)            # OFF-lock template for the blocker B
    J* ← −∞ ;  g* ← NONE

    # the loop L and the tail bt together span the trigger window
    for each (L, bt) with L in L_set, bt in BT_set, bt+|B|+|E′|+L = |T| :

        # cut T_c into consecutive sensing modules, 5′→3′
        B_T ← first bt nt of T_c          # trigger-binding tail
        rcT ← next |B| nt of T_c          # trigger template for B
        E′  ← next |E′| nt of T_c         # anti-extension
        Λ_L ← remaining L nt of T_c       # loop
        E   ← near-complement of E′       # extension, 2 bulges; E ≠ E′

        # the blocker, base by base: pair the spacer, or pair the trigger.
        # identical where S and T agree; where they differ, BOTH are tried.
        choices ← [ {rcS[i]} if rcS[i] = rcT[i] else {rcS[i], rcT[i]}
                    for i in 1..|B| ]

        for each aS in {aS1, aS2, aS3, aS4, none} :   # anti-scaffold
            for each B in product(choices) :  # 2^(#disagreements)

                g ← B_T·aS·B·E′·Λ_L·E·S·C     # the guide, 5′→3′
                if |g| > L_max :  skip B

                # observables - two tubes (section 3.4)
                A_off   ← Accessibility(g,     spacer region)   # OFF tube
                A_on    ← Accessibility(g & T, spacer region)   # ON tube
                ΔG_bind ← G(g·T) − G(g) − G(T)                 # binding
                d_on    ← Defect(g & T, s_C, scaffold region)   # ON scaffold

                # hard gates (section 3.5) - is this design in F(S,T)?
                if not ( A_off ≤ τ_off and A_on ≥ τ_on and d_on < ε ) :
                    skip B

                # inner objective (section 3.6)
                J_A ← w_on·A_on − w_off·A_off   # structural switch
                J_E ← −ΔG_bind                  # thermodynamic drive
                J   ← w_E·J_E + w_A·λ·J_A       # λ rescales J_A onto J_E
                if J > J* :  J* ← J ;  g* ← g   # argmax over F(S,T)

    return (J*, g*)
```

Two steps carry most of the design decisions. The **architecture loop** is constrained rather than free: the trigger window has a length, and the modules that pair with it must add up to exactly that length, which is why bt, L and the fixed $|B| = 20$ and $|E'| = 10$ are tied together by one equation instead of varied independently. Each admissible (L, bt) then cuts the trigger complement into the sensing modules, so B_T, E' and Λ_L are read off the trigger rather than searched.

The **blocker loop** is where the competition of section 3.2 becomes concrete. At each of the 20 blocker positions the base may be chosen to pair the spacer, which locks the OFF state, or to pair the trigger, which helps the ON state open. Where the spacer and the trigger agree the two demands name the same base and there is nothing to decide; where they disagree both are tried, so the search is over 2^n blockers for n disagreements. The trade-off is therefore not resolved by a rule but enumerated and settled by the objective.

Everything these two procedures call is defined in the sections that follow: the construct and its modules in section 3.2, the design variables in 3.3, `Accessibility`, `ΔG_bind` and `Defect` in 3.4, the three hard gates in 3.5, J in 3.6, and Φ, Q_S and the outer gates in 3.7. The architecture sets `L_set` and `BT_set`, the length ceiling $L_{\max}$ and the output size k are listed with the other parameters in Table 5.

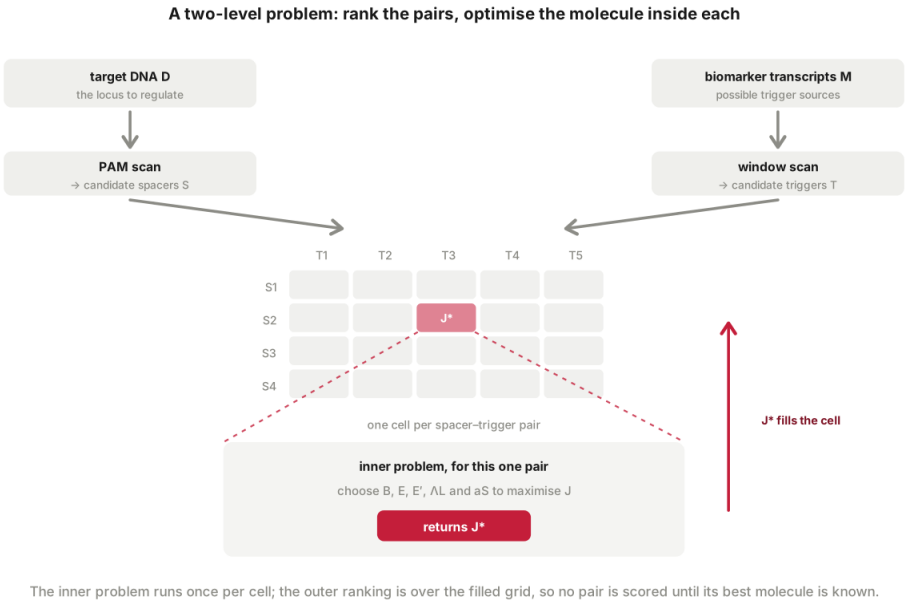

Figure 8. The bilevel problem: the inner optimisation runs once per spacer–trigger pair and returns the value the outer ranking uses.

Stages 1 and 2 fixed one spacer and one trigger in advance, produced candidates by enumeration, admitted them by structural filters and ranked them heuristically. Stage 3 removes those assumptions and states the design problem explicitly: given a gene to regulate and a set of biomarker transcripts, find the guide that switches best. Four things change. The spacer is no longer given - the algorithm receives an endogenous target DNA and generates candidate spacers itself. The trigger is no longer given - it receives biomarker transcripts and scans candidate windows along them. Design is no longer enumerate-and-filter - the quantity to be maximised is written down, and the design is the maximiser. And the molecule works differently: where the published architecture confines the trigger to the outer stem, never touching the duplex that blocks the spacer, here the trigger invades inward across the SBE itself, so that binding competes directly with the closed hairpin and the open state releases the spacer more completely.

3.2 The construct, and where the competition sits

$$g = B_T \cdot aS \cdot B \cdot E' \cdot \Lambda_L \cdot E \cdot S \cdot C \tag{1}$$

B is the SBE, E and E' the extension and anti-extension, Λ_L the loop of length L, B_T the trigger-binding extension of the SBE, aS the anti-scaffold, S the spacer and C the scaffold; Appendix A defines every symbol. In the OFF state the molecule is held shut by a stem with two distinct parts:

$$\text{OFF stem} = \underbrace{S \cdot B}_{\text{overlap (competition)}} + \underbrace{E \cdot E'}_{\text{free zone}} \tag{2}$$

The trigger pairs the $B_T \cdot B \cdot E'$ strand together with the unpaired loop, where it nucleates, and displaces the $E \cdot S$ strand, and the two parts behave differently under that attack. In the **overlap**, the same SBE bases that lock the spacer are the ones the trigger must invade - this is where the two demands compete directly. The **free zone** and B_T are never occupied by the spacer, so they can be

matched to the trigger at no cost to the lock. Hence the central difficulty: one SBE is asked to satisfy two templates at once, the reverse complement of the spacer in order to lock it and the reverse complement of the trigger in order to be invaded. Those differ, so no SBE satisfies both and every design is a position on that trade-off - which the objective resolves, with the drive to open supplied mainly by the free zone, where it is free (Figure 9).

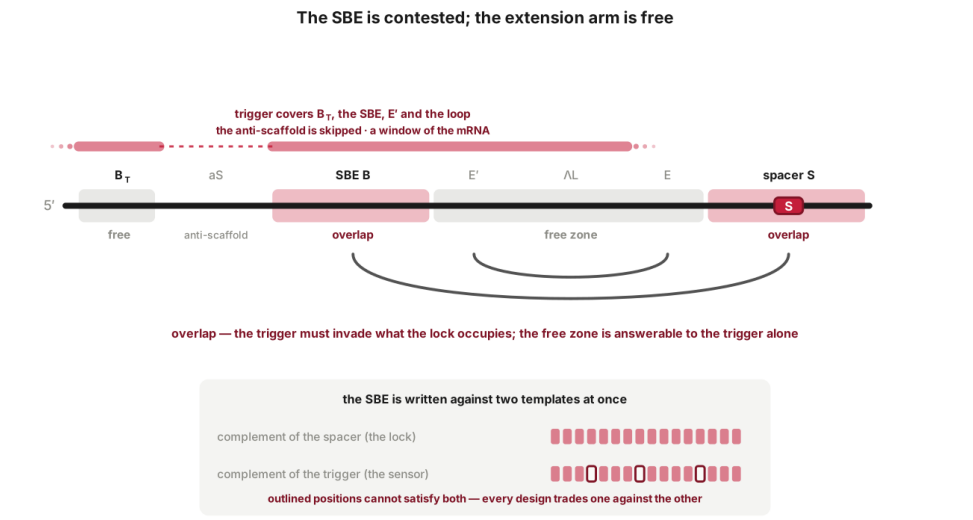

Figure 9. The SBE must serve the lock and the sensor at once; the remaining modules answer to the trigger alone.

3.3 Design variables

Variable	What it is	How it is set
B	SBE - the spacer-blocking element	Sequence; the only module searched freely, position by position
E'	Anti-extension	Read off the trigger window
E	Extension	Near-complement of E', carrying two bulges
Λ_L	Loop, of length L	Sequence read off the trigger window; the length L is chosen from `L_set`
aS	Anti-scaffold module	Discrete: one of four variants, or none
B_T	Trigger-binding extension of the SBE	Read off the trigger window; the length is chosen from `BT_set`

Table 4. The inner design variables and how each is set. The spacer, trigger window and scaffold are fixed per iteration and are not optimised. Of the engineered modules only the blocker is searched over sequences; the sensing modules are read off the trigger complement once the loop and tail lengths are chosen, and the extension follows from the anti-extension.

The anti-scaffold is the four-variant panel of section 5.2.2 unchanged, but its role differs: in stage 2 the variants were four separate design categories compared side by side, whereas here the choice among them - including none - sits inside the search. It changes neither the inner nor the outer

objective directly; its effect is structural, entering through the equilibrium ensembles and hence through the observables and the gates. Because it is *meant* to disrupt the scaffold in the OFF state, the OFF-state scaffold gate is deliberately not imposed (section 3.5).

3.4 Two tubes, four observables

Everything the objective needs is a difference between the guide alone and the guide in the presence of its trigger (Figure 10). Each is a test-tube ensemble at equilibrium, with the trigger in excess in the ON tube:

$$\Omega_{\text{off}} = \langle\{g\}\rangle, \qquad \Omega_{\text{on}} = \langle\{g, T\}\rangle \tag{3}$$

and from the equilibrium pair-probability matrix, $q_i(\Omega)$ is the probability that nucleotide i is unpaired:

$$q_i(\Omega) = Pr[\,\text{nt } i \text{ unpaired} \mid \Omega\,] \tag{4}$$

Spacer accessibility is the central observable, and it is not whether the most stable structure has an open spacer but what fraction of the whole ensemble is usable - the probability that a molecule drawn at random from the tube could load dCas9. For an ensemble X the partition function sums the Boltzmann weight of every structure the molecule can adopt,

$$Z_X = \sum_{i \in \Omega_X} e^{-\Delta G_i / RT} \tag{5}$$

where Ω_X is the set of all possible secondary structures in X and ΔG_i the free energy of structure i. The denominator RT is the molar thermal energy - the gas constant times the absolute folding temperature, 303.15 K throughout this work - and is written as a single quantity because the symbol T on its own always denotes the trigger window. The same sum restricted to the spacer-accessible subset $\Omega_X^{\text{acc}} \subseteq \Omega_X$ - the structures in which the spacer is accessible - is

$$Z_X^{\text{acc}} = \sum_{i \in \Omega_X^{\text{acc}}} e^{-\Delta G_i / RT} \tag{6}$$

and accessibility is their ratio:

$$A_X = \frac{Z_X^{\text{acc}}}{Z_X} \qquad \Rightarrow \qquad A_{\text{off}} = \frac{Z_{\text{off}}^{\text{acc}}}{Z_{\text{off}}}, \quad A_{\text{on}} = \frac{Z_{\text{on}}^{\text{acc}}}{Z_{\text{on}}} \tag{7}$$

A structure counts as accessible when the spacer meets a stated structural condition; the strict form used here requires all 20 spacer nucleotides to be unpaired, deliberately conservative and relaxable during calibration if partial exposure proves sufficient for dCas9 function. This is the definition of spacer accessibility for the gate. The 80 % threshold in section 5.1.5 is a different thing and should not be read as a competing definition: it is an acceptance criterion for screening one control construct, not a statement about how the gate behaves. A good design has OFF-state accessibility near zero and ON-state accessibility near one.

Trigger binding is the free-energy change on forming the complex, more negative meaning stronger, where $G(\cdot)$ is the free energy of the ensemble of the species in brackets:

$$\Delta G_{\text{bind}} = G(g \cdot T) - G(g) - G(T) \tag{8}$$

Scaffold fidelity is the ensemble defect of the scaffold against its reference structure, counting how far the scaffold region deviates on average from the fold dCas9 requires; lower is better. Writing s_C

for that reference structure - the scaffold fold dCas9 binds, in dot-bracket notation - and $n(\cdot \mid s_C, \Omega)$ for the ensemble defect, the expected number of nucleotides of a region that are incorrectly paired relative to s_C across the ensemble Ω:

$$d(\Omega) = n(C \mid s_C, \Omega) \tag{9}$$

The anti-scaffold deliberately raises this in the OFF ensemble - that is its purpose. What must be preserved is its value in the ON ensemble, so that dCas9 loads once the switch has opened.

The realised switch. What drives opening is the difference between the energy gained by binding the trigger and the cost of prying the stem apart:

$$\Delta\Delta G \approx \Delta G_{\text{bind}} - \Delta G_{\text{open}}(S \cdot B,\ E \cdot E') \tag{10}$$

The ensemble provides no clean "open guide" reference from which the opening cost could be computed, so the realised switch - ON-state minus OFF-state accessibility - serves as its operational proxy, and it is the right one: it is exactly where binding drive and stem resistance net out. A hairpin too stable to open shows up as low ON-state accessibility even when the binding energy looks favourable, which is the failure a binding energy alone would miss.

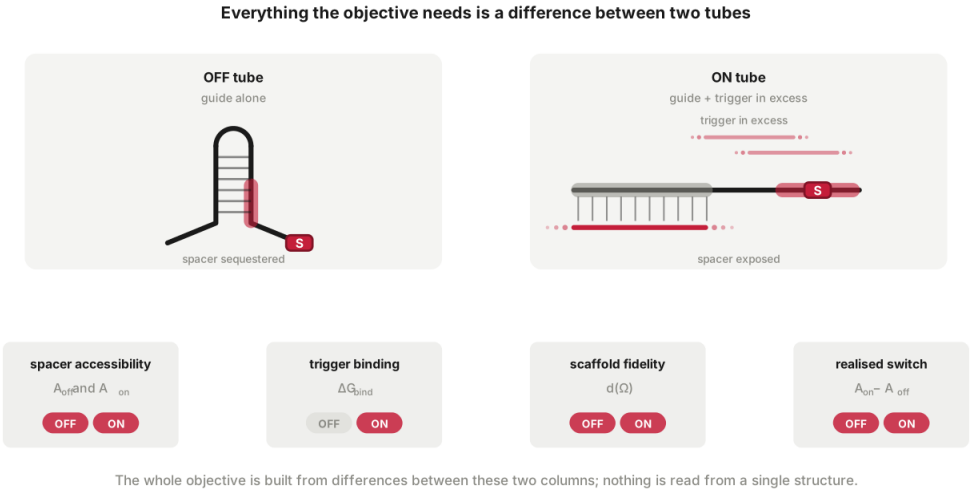

Figure 10. The OFF and ON ensembles, and which observable is read from each.

3.5 Hard gates

As in stages 1 and 2, some conditions remove a design outright rather than scoring it down. Three of them define the feasible set:

$$F = \{\,(B, E, E', L, aS) : A_{\text{off}} \le \tau_{\text{off}},\quad A_{\text{on}} \ge \tau_{\text{on}},\quad d(\Omega_{\text{on}}) < \varepsilon\,\} \tag{11}$$

The OFF state must be locked, the ON state must open, and the scaffold must be intact in the ON state so that dCas9 can load; a design failing any one of these never reaches the objective. There is deliberately no OFF-state scaffold gate: there the spacer is already occluded, so scaffold competence is immaterial, and the anti-scaffold may be disrupting it on purpose - imposing fidelity would reject exactly the designs that module is meant to produce.

3.6 The inner objective

For designs in F - and minus infinity for everything outside it - the objective combines two complementary descriptions of switching, with non-negative free coefficients (Figure 11). The **structural component** rewards ON-state accessibility and penalises residual OFF-state accessibility; the **thermodynamic component** rewards favourable trigger binding:

$$J_A = w_{\text{on}} A_{\text{on}} - w_{\text{off}} A_{\text{off}} \qquad J_E = -\Delta G_{\text{bind}} \tag{12}$$

The two are expected to correlate, since both reflect the trigger's ability to shift the guide from closed to open, but are kept separate because they capture different aspects of that transition: one asks whether the molecule ends up open, the other whether the trigger is held strongly enough to get it there. They occupy different numerical ranges - a probability difference against an energy - so the structural component is first rescaled by a factor determined from a calibration set of designs, written cal, over which the averages $\langle \cdot \rangle_{\text{cal}}$ are taken:

$$\lambda = \frac{\langle |J_E| \rangle_{\text{cal}}}{\langle |J_A| \rangle_{\text{cal}}} \tag{13}$$

This scaling is purely a matter of units and says nothing about relative biological importance, which the weights control separately:

$$J = \underbrace{w_E J_E}_{\text{trigger binding}} + \underbrace{w_A \lambda J_A}_{\text{ON + OFF score}}, \qquad w_E + w_A = 1 \tag{14}$$

Keeping the scaling and the weights separate is what makes the weights interpretable: they are calibrated against experimentally characterised designs rather than fixed by assumption. The selected design is the constrained maximiser:

$$(B^*, E^*, E'^*, L^*, aS^*) = \underset{(B,E,E',L,aS) \in F}{argmax} \; J \tag{15}$$

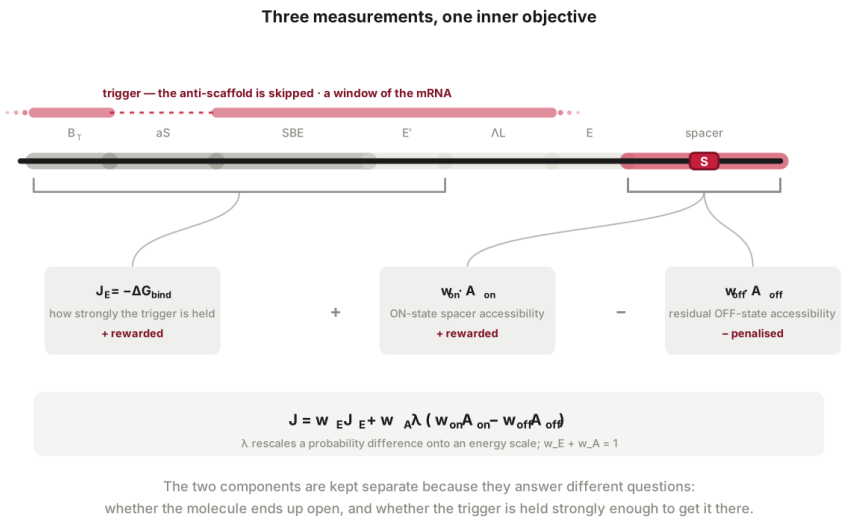

Figure 11. Where the three terms of the inner objective are measured, and how they combine.

3.7 The outer objective

Candidate generation. The candidate spacers $\mathcal{S}_{\text{cand}}$ come from a PAM scan of the target DNA, filtered to the allowed targeting window; the candidate trigger windows $\mathcal{T}_{\text{cand}}$ from scanning the supplied biomarker transcripts. One constraint ties the two scans together: trigger-window length is set by the loop length rather than optimised independently, so for each admissible loop length only windows of the corresponding length are considered. Each combination is then evaluated by solving the inner problem for its optimum.

Spacer quality. A spacer must direct dCas9 efficiently to its intended site and not be confused for anywhere else in the genome. Let the normalised on-target activity and off-target risk scores both lie in [0,1], larger being better and worse respectively:

$$Q_S(S) = v_{\text{on}}\,\underbrace{\eta_{\text{on}}(S)}_{\text{on-target}} + v_{\text{off}}\big(1 - \underbrace{\eta_{\text{off}}(S)}_{\text{specificity}}\big), \qquad v_{\text{on}}, v_{\text{off}} \geq 0, \quad v_{\text{on}} + v_{\text{off}} = 1 \tag{16}$$

Subtracting the risk from one converts it into a specificity reward, so both terms point the same way; as a convex combination of quantities in [0,1] the spacer score lies genuinely in [0,1], larger values marking spacers that combine high predicted activity with high genomic specificity. Both metrics are method-agnostic; CRISPOR is one implementation.

Trigger score. The trigger score, in [0,1], quantifies how suitable a candidate trigger window is as an endogenous RNA input, with larger values representing more favourable candidates. It is calculated from the trigger-selection metric used by the pipeline, at the trigger-selection stage.

Outer gates. A pair is admissible only if a working guide exists for it - a non-empty inner feasible set - and the spacer is genomically unique enough:

$$\mathcal{G} = \{\,(S,T) \in \mathcal{S}_{\text{cand}} \times \mathcal{T}_{\text{cand}} : F(S,T) \neq \varnothing, \quad \eta_{\text{off}}(S) \leq o_{\max}\,\} \tag{17}$$

The objective. Targeting quality and trigger quality may trade against each other, but neither should buy back a device that does not switch. The inner optimum therefore enters as a multiplicative gain rather than another addend, taken as its non-negative part:

$$\Phi(S,T) = \underbrace{\big(W_S Q_S(S) + W_T Q_T(T)\big)}_{Q_{\text{pair}} = \text{targeting and openness quality}} \cdot\; J_+^*(S,T), \qquad W_S, W_T \geq 0, \quad W_S + W_T = 1 \tag{18}$$

The inner factor is the performance of the optimised switch and the pair-quality factor, in [0,1], determines how much of it survives once spacer and trigger quality are accounted for: a good pair leaves the inner score largely intact, a poor one scales it down, and a pair whose switch is not net-positive receives zero and ranks last. Pairs failing the inner constraints are removed before outer scoring, so quality cannot compensate for a non-functional switch at all. The program returns a ranked top-k list of pairs with their optimal guides, for wet-lab validation:

$$(S^*, T^*) = \underset{(S,T)\in\mathcal{G}}{argmax}\ \Phi(S,T) \tag{19}$$

3.8 Free parameters and calibration

Parameter	Where it appears	What it controls	Current value

Parameter	Where it appears	What it controls	Current value
w_{on}, w_{off}	Structural component	Opening versus staying shut	0.5, 0.5
w_E, w_A	Inner objective	Thermodynamic against structural evidence	0.5, 0.5
λ	Inner objective	Unit scaling between the two components	computed per batch
τ_{off}, τ_{on}	Inner gates	How locked the OFF and how open the ON state must be	0.05, 0.50
ε	Inner gate	Permitted scaffold deviation in the ON state	0.10
v_{on}, v_{off}	Spacer score	Activity against specificity	0.4, 0.6
W_S, W_T	Outer objective	Spacer quality against trigger quality	0.6, 0.4
$o_{\max}$	Outer gate	Maximum tolerated off-target risk	0.2
`L_set`	Architecture	Admissible loop lengths	14–32 nt, even
`BT_set`	Architecture	Admissible trigger-tail lengths	10, 15, 20 nt
$L_{\max}$	Architecture	Ceiling on the assembled guide length $\|g\|$	given constraint
k	Output	How many ranked designs are returned	given

Table 5. The free parameters, with the values the implementation currently carries.

Those values are defaults and chosen settings, not calibrated ones. The gate thresholds τ_{off}, τ_{on} and ε, together with the architecture sets, are the codebase defaults. The weight sets and $o_{\max}$ have no default in code and must be supplied; the values above are the project's chosen settings. The four inner weights sit at 0.5 as neutral placeholders - they encode no judgement about the relative importance of the two components and are exactly what calibration against measured designs is meant to replace.

3.9 How to read the output

The outer score carries an energy scale, inherited from the thermodynamic term in the inner objective, but structural accessibility, spacer quality and trigger quality all enter through calibrated dimensionless factors. It should therefore be read as an energy-scaled composite ranking score rather than a physical free-energy difference. Keeping it on that scale is deliberate: it lets the endogenous gate be ranked consistently against the other CERNAL gate-design models, which are represented on the same calibrated scale.

4. Discussion

4.1 Two engineering cycles

The work reported here is two complete design–build–test–learn cycles rather than one. Each closed on a measurement, and each measurement changed the next design (Figure 7).

Cycle	Designed	Built and tested	What it showed	What it changed

Cycle	Designed	Built and tested	What it showed	What it changed
1	Stage 1 - the published switch reproduced, plus our fully-closed and fully-open controls	9 guides	The three gates did not separate from one another	The scaffold lock and the longer trigger window of stage 2
2	Stage 2 - anti-scaffold and B_T on both gated backbones	8 guides	The lock tightened the fully-closed gate but not the switch	The SBE optimisation of stage 3

Table 6. The two cycles. Stage 3 is designed but not yet built, so the second cycle is still open at its test step.

4.2 What the first cycle established

Stage 1 asked a narrow question: reproduced faithfully and run on our spacer and trigger, does the published architecture give a gate that separates from its own floor and ceiling? The two controls exist so that this question can be answered at all - without them, an absent output cannot be distinguished from a non-functional construct, and no measurement can be normalised into a fold-activation (section 5.1.1).

Testing, found the three designs too similar to one another. Our reading is that the guides the model produced were not locked tightly enough. One suspicion we carried forward was that dCas9 binding itself destabilises the closed conformation, so that a guide can be correctly folded and still leak.

That reading pointed at a gap in the architecture rather than in the sequences. Stage 1 varied only how firmly the spacer stem closes, and in all three designs the scaffold was left entirely free. Stage 2 followed directly from that.

4.3 What the second cycle established

Stage 2 added the anti-scaffold and B_T and applied them to both gated backbones, holding the hairpin constant so that the two new modules were the only variable (section 5.2.7).

The outcome was asymmetric, and the asymmetry is the informative part. On the fully-closed gate the anti-scaffold did what it was designed to do: the OFF state tightened and the gate improved. On the source switch it did not - in the absence of trigger the gate was still too open. Two conclusions follow.

- **The scaffold lock works as a mechanism.** This is not a null result. On a backbone whose spacer stem is already perfectly paired, adding a second, independent lock on the scaffold still tightened the OFF state measurably. That supports the premise of stage 2, and indirectly the suspicion that dCas9 loading contributes to leak.
- **It is not sufficient on its own, and the residual leak tracks the stem.** The two backbones differ in exactly one respect: the fully-closed gate has a perfectly complementary SBE and a fully paired extension stem, while the source switch carries four inherited bulges across those two stems. The same scaffold module tightened the first and not the second. Our reading is that what still leaks in the switch comes from its stem, not from its scaffold.

That is a specific and testable diagnosis, and it points at the one part of the molecule stages 1 and 2 never varied.

4.4 Why we expect the optimisation to help

The case for it rests on three properties of the earlier stages that we now think were wrong for our system.

- **The bulge pattern was inherited, not designed.** The four bulges of the source switch sit at positions fixed in advance by the published target fold (section 5.1.3), chosen for the designs in that paper. Nothing ties those positions to our spacer or our trigger. Stage 3 makes the SBE, the extension arms and the loop length search variables rather than inherited constants.
- **The ON state was never scored.** The stage-1 score (section 5.1.4) rewards OFF-state folding probability cubed, spacer identity and trigger accessibility. None of its four terms asks whether the trigger can actually open the molecule, so a design can rank first while being effectively unopenable - which is consistent with what the lab saw. Stage 3 scores both states and gates on both: the OFF state must be locked and the ON state must open, or the design is discarded before scoring (section 3.5).
- **The trade-off was never resolved explicitly.** Lock tightness and trigger responsiveness pull in opposite directions through the same bases (section 3.2). Stages 1 and 2 handled this by inheriting one fixed compromise; stage 3 writes the trade-off down and searches for its optimum, with the drive to open supplied mainly by the free zone, where it costs the lock nothing.

The hypothesis is therefore specific: the switch leaks because it carries a compromise tuned for a different spacer–trigger pair, and an SBE optimised for general input should sit closer to the point where the lock holds and the trigger can still open it. Whether that is right is an experimental question, and completing the second cycle is what answers it.

4.5 Next steps

In order of value:

- Implement the stage-3 optimisation and generate candidates for the *hph* spacer and the mCherry trigger, so that the new designs can be compared against the stage-1 and stage-2 guides on the same pair.
- Calibrate the weights and thresholds against the published Pelea designs, then against our own measured fold changes once they exist.
- Test the stage-3 designs and measure the ON/OFF ratio against the fully-closed and fully-open controls, which remain the floor and the ceiling.
- Extend the plasmid variant experiments, where the spacer is supplied rather than designed.
- If the optimised designs still leak, widen the inner search: relax the strict accessibility criterion to partial spacer exposure.

5. Methods

The two cycles that produced the framework of section 3, in the order they were run.

5.1 Stage 1 - reproduction and control designs

5.1.1 What this stage was for

CERNAL needs a CRISPR gate it can treat as a logic element: a guide that stays inert until a chosen RNA is present. The published iSBH-sgRNA [1] is exactly that - a guide whose spacer is held shut by a hairpin and released when a trigger RNA invades it - so the first question was not whether to invent an architecture, but whether we could build this one in our hands and measure how well it worked.

Reproducing the published design answers the first half. It does not answer the second, because a single switch measured on its own is uninterpretable: a construct showing little activity may be properly closed or simply broken, one showing a great deal may be properly open or merely leaky, and neither reading can be turned into a fold-change without something to normalise against.

Stage 1 therefore produced three constructs rather than one: the published switch, and two controls of our own design - a gate built to be as closed as the architecture allows, and a gate built so that it cannot close at all. The controls are this stage's contribution. They fix the floor and the ceiling against which every later measurement is read.

5.1.2 How the guide is built

All three constructs share one architecture and differ in their first module and extension*. Reading 5′ to 3′, the designed region is

```
5′ - SBE (20) - extension* (10) - loop (L) - extension (10) - spacer (20) - scaffold - 3′
```

The spacer is the 20-nt targeting sequence. The SBE - the spacer-blocking element - is written to pair with it, folding the molecule into a hairpin whose two stems are the SBE·spacer duplex and the extension·extension* duplex, with the loop unpaired between them (Figure 12). In that OFF state the spacer is occupied and dCas9 cannot use the guide. The loop and extension* present a single-stranded window to the trigger; when the trigger binds there and migrates inward, both stems open and the spacer is released. The constant dCas9 scaffold is appended to every design and never altered.

Everything in this stage was folded with the RNA energy model at 30 °C, the temperature used throughout the project. One sequence rule applies everywhere: no designed region may contain TTTT, since the constructs are transcribed from a human U6 promoter and four consecutive T residues terminate polymerase III.

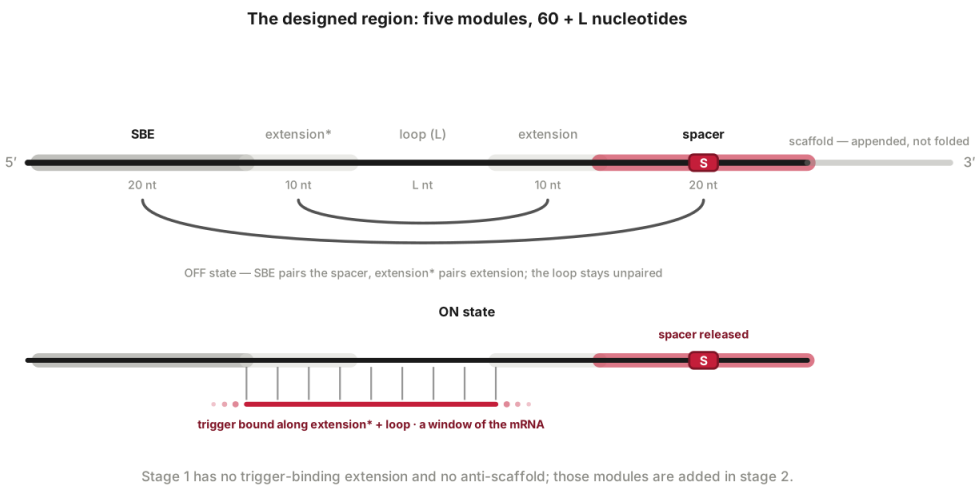

Figure 12. The five modules of the stage-1 designed region, with the OFF-state pairings and the ON-state trigger duplex.

5.1.3 The three designs

Design	SBE	Intended behaviour	Role in the set
Source switch	Near-complement of the spacer, with the bulges of the published design	Closed until the trigger arrives, then openable	The gate under test
Fully-closed gate	Exact reverse complement of the spacer, no bulges at all	Locked as tightly as the architecture allows	Negative control - the OFF reference
Fully-open control	Scrambled 20-mer with no complementarity to the spacer	Spacer accessible with or without trigger	Positive control - the ON reference

Table 7. The three stage-1 designs. Only the 20-nt SBE differs between them.

Because only the SBE differs, the three sit on a single axis: how strongly the hairpin can close (Figure 13).

The source switch is the published architecture, reproduced through the published toolkit; its design rules and generation procedure are those of Pelea et al. [1]. The one point worth restating is why its SBE is not a perfect complement of the spacer. A perfectly paired hairpin is too stable for the trigger to open, so the published design places four deliberate mismatches - two in the SBE and two more in the extension*, at fixed positions - to bring the OFF state into the range where strand displacement is favourable while the spacer stays blocked.

The fully-closed gate is ours. We removed the bulges entirely and made the SBE the exact reverse complement of the spacer, so that both stems are perfectly paired. This is the tightest OFF state the architecture permits, and it answers a specific question: if a gate with no destabilising mismatches still shows activity, the leak is not coming from the stems.

The fully-open control is also ours, and is the inverse - an SBE with no complementarity to the spacer at all, so the spacer stays accessible whether or not the trigger is present. It is derived from the fully-

closed construct rather than designed independently: the 34-nt inner body and the spacer are taken verbatim and only the SBE is replaced with optimized non-paring SBE, so that length, composition, promoter and vector context are held constant while a single variable changes.

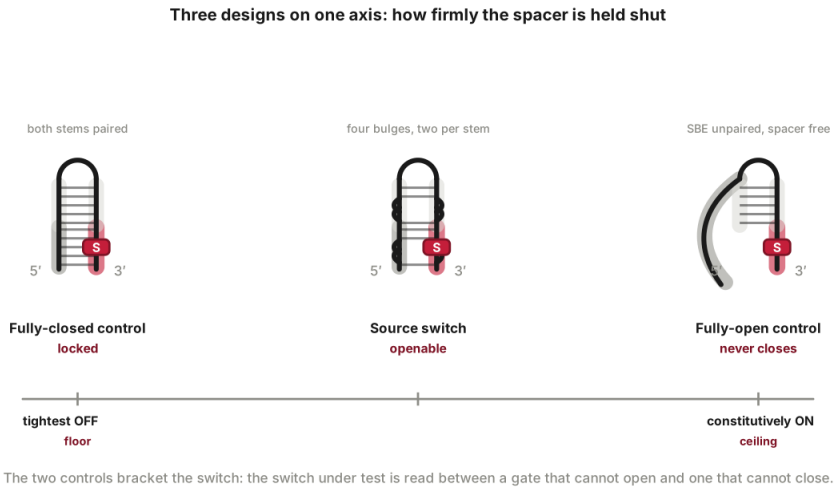

Figure 13. The fully-closed control, the source switch and the fully-open control, ordered by how firmly the spacer is held shut.

5.1.4 Generating and ranking candidates

The source switch and the fully-closed gate were produced by the same pipeline, which differs between them only in the target fold it is asked to match - bulged for the switch, fully paired for the closed gate.

Loop length is the principal tunable parameter, since it sets the size of the single-stranded window offered to the trigger: too short and the trigger cannot nucleate, too long and the OFF state leaks. Rather than fix it by assumption, ten lengths were swept, 14 to 32 nt even, as parallel design classes. For each length a window of the matching size is slid one nucleotide at a time across the whole mCherry transcript and the reverse complement of each window taken as a candidate sensing module, so that every register is tested and the choice deferred to thermodynamics. Each candidate is then assembled with the fixed spacer and SBE.

Four conditions act on the assembled candidate (Figure 14). The MFE structure must equal the intended fold exactly, a strict equality rather than a similarity threshold and the main reason long loops yield few designs; the sequence must contain no TTTT; the intended fold must hold at least 40 % of the equilibrium ensemble; and the trigger interaction must be independently re-derivable and locatable in the trigger, with duplicates collapsed.

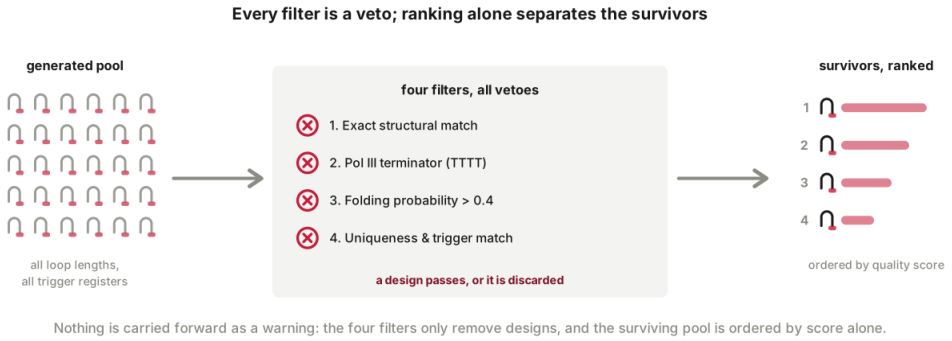

Figure 14. The four filters act only as vetoes; the surviving pool is ordered by score alone.

Survivors are then ranked, because structural validity is necessary but not sufficient: a design can fold correctly and still sense poorly, if its trigger site is buried inside the transcript's own structure or if assembly has drifted from the intended spacer. Four terms are scored.

Term	Definition	What it penalises
Folding probability	Equilibrium probability of the designed OFF-state fold	Switches that do not adopt the intended locked structure
Spacer similarity	Identity between the realised spacer and the original spacer	Drift away from the genomic target, i.e. loss of on-target activity
Trigger similarity	Identity between the sensing region and the trigger window	Imperfect complementarity to the actual trigger
Openness	Unpaired fraction of the sensed window in the trigger's own MFE structure	Sensing regions buried inside trigger secondary structure

Table 8. The four terms of the stage-1 score.

$$\text{quality score} = p^3 \times \text{spacer similarity} \times \text{trigger similarity} \times \text{openness} \tag{20}$$

Three properties of this score are deliberate. The terms multiply rather than average, so a near-zero value on any axis vetoes the design instead of being compensated elsewhere. The folding probability p enters as a cube, making OFF-state fidelity dominant, on the reasoning that a leaky gate cannot be rescued by good accessibility. And the terms keep their native scale, so the score is unnormalised and ranks only within a comparable group - designs are therefore sorted within each loop length, so that the best candidate at a given length appears alongside its true competitors.

5.1.5 Screening the fully-open control

The open control could not come from that pipeline, because its design problem is the inverse: not to find a sequence that folds into a specified structure, but one that reliably fails to. It was produced by a two-stage screen (Figure 15).

Candidate SBEs are uniform random 20-mers, filtered on sequence before any folding call, since folding is the rate-limiting cost. Three filters apply: no TTTT; GC content no higher than the original SBE's 60.0 %, a ceiling rather than a window, since the aim is only to reduce the region's propensity to pair; and no local complementarity with the spacer, rejecting any candidate containing a short

substring that is the exact reverse complement of a substring of the spacer. The third is the one that enforces “open by design”, because a scrambled sequence can share no overall identity with the original SBE and still contain a short complementary tract that nucleates a stem and sequesters part of the spacer. Since the substring length traded strictness against pool diversity, it was swept rather than fixed, and the pools combined.

Survivors are assembled into the full construct and screened thermodynamically: the spacer must be at least 80 % unpaired in the MFE structure, and only candidates passing that gate go on to the costlier partition-function step. Openness alone is insufficient, since a candidate can be open in one structure among many comparably likely ones, so the final ranking combines the count of unpaired spacer nucleotides with the equilibrium probability of the MFE fold on a common scale.

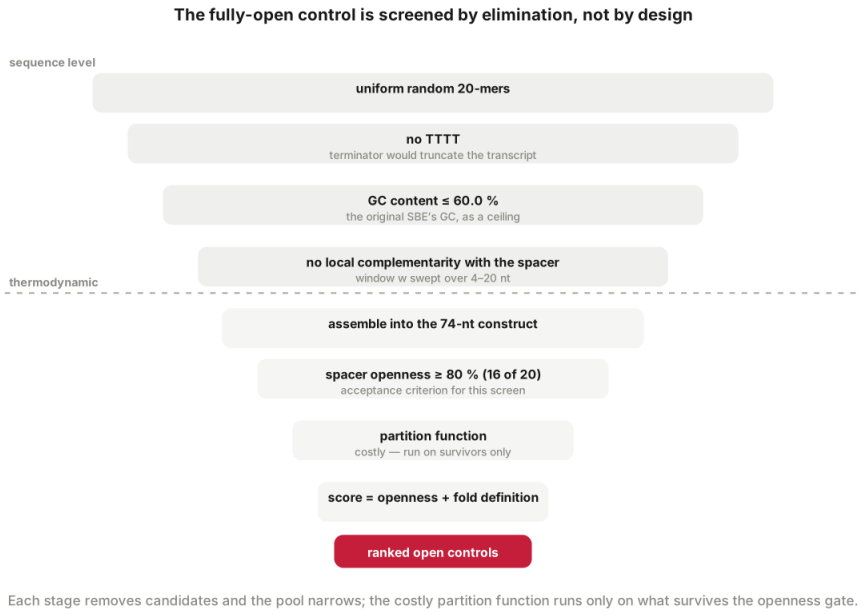

Figure 15. The two-stage screen for the fully-open control. The partition function is computed only for candidates passing the openness gate.

5.1.6 Parameters

Parameter	Source switch	Fully-closed gate	Fully-open control
Folding model	RNA, 30 °C	same	same
Spacer	`GACCTGATGCAGCTCTCGGA`	same	same
SBE	`TCCGAGAGCTAAATCTTGTC`	`TCCGAGAGCTGCATCAGGTC`	scrambled, screened
Spacer stem	20 bp, two bulges	20 bp, fully paired	does not form
Extension stem	10 bp, two bulges	10 bp, fully paired	inherited from reference
Loop lengths	14–32 nt, even	same	fixed at 14
SBE retained in	first 15 nt	all 20 nt	replaced entirely

Parameter	Source switch	Fully-closed gate	Fully-open control
module			
Spacer retained in module	last 17 nt	all 20 nt	all 20 nt
Folding-probability cut-off	0.4	0.4	not applied
Spacer-openness threshold	not applied	not applied	80 % (16/20)
Score	p^3 × spacer sim × trigger sim × openness	same	openness + fold definition

Table 9. Parameters for each stage-1 design.

5.1.7 Candidates sent for synthesis

Nine designs were taken forward: four from the source switch, two fully-closed and three fully-open. The source switch senses mCherry and is the gate under test; the other two are our own controls, fixing the floor and the ceiling against which it is read (section 5.1.1). Full sequences are in Appendix B.

Design	Group	Rank	Loop	SBE (nt 1–20)	Score	*p*(MFE)	ΔG MFE (kcal·mol^{-1})
sgRNA mCherry 1	Source switch	1	14	TCCGAGAGCTAAATCCCGTC	82,066	-	−32.49
sgRNA mCherry 2	Source switch	2	14	TCCGAGAGCTAAATCCCGTC	81,542	-	−31.71
sgRNA mCherry 3	Source switch	3	14	TCCGAGAGCTAAATCCCGTC	80,981	-	−31.08
sgRNA mCherry 4	Source switch	1	26	TCCGAGAGCTAAATCCTTGA	17,794	-	−31.22
sgRNA closed 1	Fully-closed	1	14	TCCGAGAGCTGCATCAGGTC	90,261	-	−66.89
sgRNA closed 2	Fully-closed	2	14	TCCGAGAGCTGCATCAGGTC	61,463	-	−67.47
sgRNA open 1	Fully-open	1	14	TCGGCATAAAACGCCCAGGG	18.00	0.735	−26.91
sgRNA open 2	Fully-open	2	14	CAGGTTAAAGAACCCAGGCG	17.52	0.598	−27.46
sgRNA open 3	Fully-open	3	14	CCTATATAATTAGGAGGGGA	17.50	0.480	−27.58

Table 10. The nine stage-1 candidates. ΔG MFE is the minimum free energy of the designed region at 30 °C, as reported by the visualiser. Ranks are within design group and loop length. Scores come from two different formulas and are not comparable between the gated and open groups - see below.

The two score columns measure different things. The source switch and fully-closed designs are ranked by the gated score of section 5.1.4, whose magnitude is unnormalized and typically 10^4–10^5. The fully-open controls are ranked by the two-term 0–20 score of section 5.1.5. A value of 18.00 and a value of 90,261 say nothing about each other.

The assembly rules are visible in the sequences. The fully-closed designs carry the full 20-nt reverse complement TCCGAGAGCTGCATCAGGTC unchanged, as section 5.1.3 specifies for that design; the source switch retains only the outer 15 nt, TCCGAGAGCTAAATC, with the remainder trigger-derived.

The fully-open ranking shows the trade-off the score encodes. Open 3 is the most accessible of the three at 17 of 20 spacer nucleotides unpaired, against 16 for the other two, yet ranks last because its fold is the least dominant (p = 0.480 against 0.735). Open 1 was selected as the design that is both open and thermodynamically well-defined.

The loop-14 source-switch candidates are bulge variants of one sensing module. mCherry 1–3 share their first 40 nt and differ only at one or two positions in the extension arm, where section 5.1.3 induces bulges - they are alternative bulge patterns on one design rather than three independent designs. Only one 26-nt loop design passed the filters, consistent with the yield skew described in section 5.1.4.

The MFE energies track the axis the three designs sit on (Figure 16)**.** At a 14-nt loop the fully-closed designs are by far the most stable, at −66.89 and −67.47 kcal·mol^{-1}; the source switch sits near −31 to −32, and the fully-open controls are the least stable at −26.91 to −27.58. mCherry 4 carries a 26-nt loop and is 12 nt longer than the rest, so its energy is not on the same footing as the others.

Stage 1 — the nine designs carried forward

Source switch · mCherry-sensing

sgRNA mCherry 1 — MFE −32.49 kcal·mol⁻¹



TCCGAGAGCTAAATCCCGTCCTCGAAGTTCATCACGC
GCTCCCAGAAGATCGTCGACCTGATGCAGCTCTCGGA

sgRNA mCherry 2 — MFE −31.71 kcal·mol⁻¹



TCCGAGAGCTAAATCCCGTCCTCGAAGTTCATCACGC
GCTCCCAGAAAATCGTCGACCTGATGCAGCTCTCGGA

sgRNA mCherry 3 — MFE −31.08 kcal·mol⁻¹

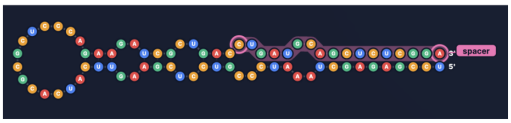

TCCGAGAGCTAAATCCCGTCCTCGAAGTTCATCACGC
GCTCCCAGAAGATCGCTGACCTGATGCAGCTCTCGGA

sgRNA mCherry 4 — MFE −31.22 kcal·mol⁻¹



TCCGAGAGCTAAATCCTTGATGATGGCCATGTTATCC
TCCTCGCCCTTGCTCACCAATGAACATATTCACTGAT
GCAGCTCTCGGA

Fully-closed control

sgRNA closed 1 — MFE −66.89 kcal·mol⁻¹

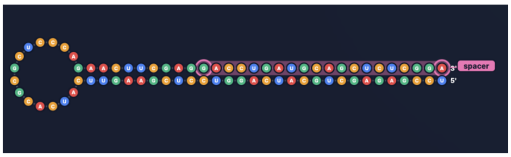

TCCGAGAGCTGCATCAGGTCCTCGAAGTTCATCACGC
GCTCCCAGAACTTCGAGGACCTGATGCAGCTCTCGGA

sgRNA closed 2 — MFE −67.47 kcal·mol⁻¹

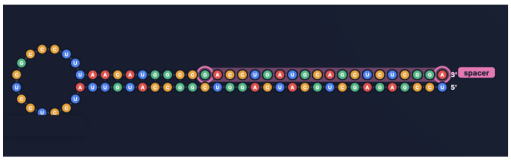

TCCGAGAGCTGCATCAGGTCGGCCATGTTATCCTCCT
CGCCCTTTAACATGGCCGACCTGATGCAGCTCTCGGA

Fully-open control

sgRNA open 1 — MFE −26.91 kcal·mol⁻¹



TCGGCATAAAACGCCCAGGGCTCGAAGTTCATCACGC
GCTCCCAGAACTTCGAGGACCTGATGCAGCTCTCGGA

sgRNA open 2 — MFE −27.46 kcal·mol⁻¹



CAGGTTAAAGAACCCAGGCGCTCGAAGTTCATCACGC
GCTCCCAGAACTTCGAGGACCTGATGCAGCTCTCGGA

sgRNA open 3 — MFE −27.58 kcal·mol⁻¹



CCTATATAATTAGGAGGGGACTCGAAGTTCATCACGC
GCTCCCAGAACTTCGAGGACCTGATGCAGCTCTCGGA

Folds and energies computed with cernal-rnaviz at 30 °C; the spacer is highlighted in each panel.

Figure 16. The nine stage-1 designs as folded by the visualiser, grouped by design. Each panel carries the construct name and its 74-nt designed region.

5.2 Stage 2 - the scaffold lock and the extended trigger window

Stage 1 did not separate the three gates from one another. Our reading was that the model may be producing guides that are not closed tightly enough in the presence of Cas9 - that binding of the protein itself supposedly destabilizes the structure and opens the gate without a trigger. Stage 1 varied only how firmly the spacer stem closes; in all three designs the scaffold was left free. Stage 2 therefore departs from the published architecture for the first time, by adding two modules at the 5′ end.

5.2.1 What changes in the molecule

```
5′ - B_T (10) - anti-scaffold (9–31) - SBE (20) - extension* (10) - loop (L) - extension (10) - spacer (20) - scaffold - 3′
```

Everything from the SBE rightward is unchanged from previous stage. The two modules address the two states separately (Figure 17): the anti-scaffold makes the OFF state tighter, and B_T makes the ON state easier to reach. There is still one trigger binding one continuous window; what changes is that the window is longer.

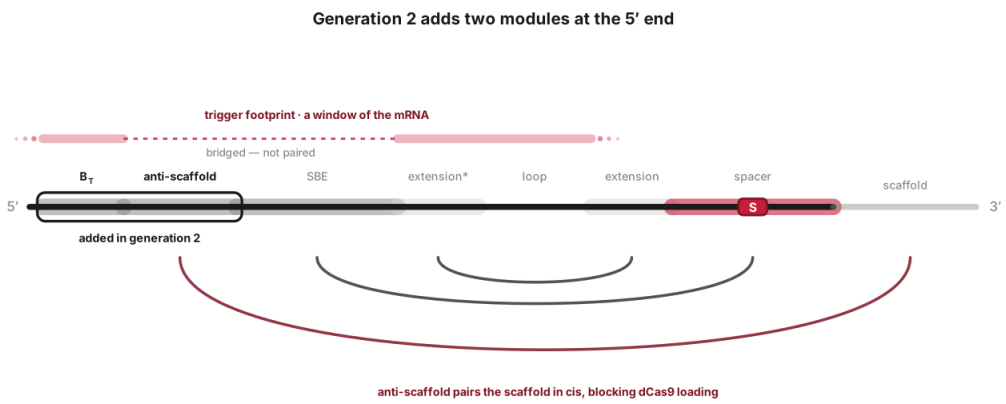

Figure 17. The generation-2 guide. The anti-scaffold pairs the scaffold in cis, and the trigger window spans the new modules but skips them.

5.2.2 The anti-scaffold

The published architecture blocks one thing: the spacer. The anti-scaffold adds a second, independent lock on the scaffold. Placed between B_T and the SBE, it is an exact reverse complement of a window of the constant scaffold, so that it base-pairs with it *in cis*. The intent is to weaken dCas9 loading, so that a guide whose spacer is already held shut also cannot assemble a functional ribonucleoprotein - keeping the gate closed even with Cas9 present. The scaffold itself is never mutated, because its sequence is constrained by the protein interface; all modulation happens through this separate, displaceable module.

Four designs were made, each the exact reverse complement of a window of the scaffold given in section 5.1.2:

Design	Sequence (RNA, 5′→3′)	Length	Scaffold window	Invasion
anti_scaffold_1	CUCUAAAAC	9 nt	1–9	Medium, first stem
anti_scaffold_2	AUUUCUAGCUCUAAAAC	17 nt	1–17	Deepest, first stem
anti_scaffold_3	AACGGACUAGCCUU	14 nt	31–44	Bridge into the second, small stem
anti_scaffold_4	CACUUUUUCAAGUUGAUAACGGACUAGCCUU	31 nt	31–61	Deeper bridge into the second stem

Table 11. The anti-scaffold panel. Coordinates are 1-based on the scaffold of section 5.1.2, whose sequence is given in Appendix B.

Designs 1 and 2 form a graded depth series within the first stem; designs 3 and 4 instead reach across into the second, smaller stem (Figure 18), testing whether invading a different structural element works better at comparable length. Length is confounded with position across the two sub-series but not within them, so the two variables can be partly separated.

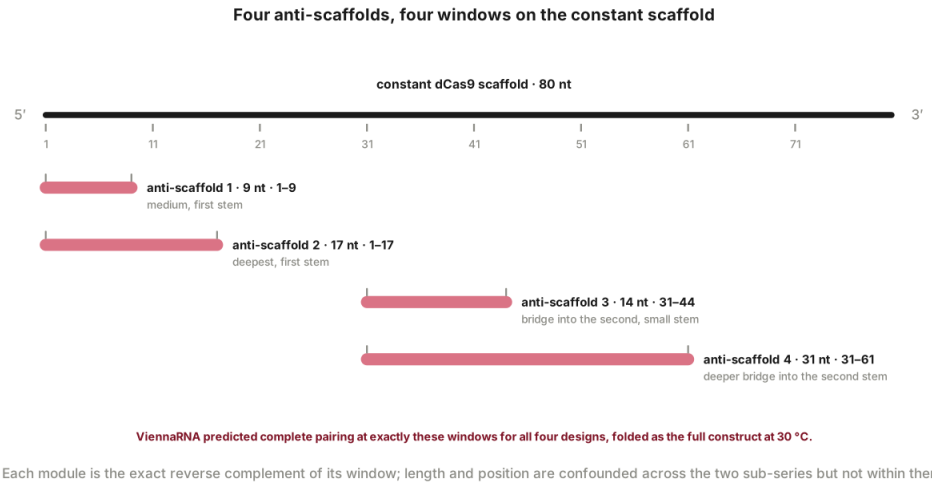

Figure 18. The four anti-scaffold designs and the windows they are complementary to on the 80-nt scaffold.

5.2.3 B_T and the extended trigger window

B_T is a 10-nt module at the extreme 5′ end of the guide. It is perfectly reverse-complementary to a stretch of the same trigger the hairpin already senses, so its sequence is taken from the trigger window the algorithm is examining. The length was fixed at 10 nt as a compromise - long enough to anchor specifically, short enough to be unlikely to form structure on its own.

Its purpose is to improve the ON state by lengthening the trigger interface. In the published architecture the trigger pairs with the loop and extension* only. Here the window is extended so that

the trigger also pairs with B_T, and that pairing is perfect rather than bulged. The extra anchoring helps the trigger open the lock, offsetting the energy the anti-scaffold costs to displace.

The window remains **one continuous stretch of one trigger**. Between B_T and the loop lies the anti-scaffold segment, which is already paired with the scaffold and so is not available to the trigger: along the duplex the trigger simply skips it, bridging over it as a single long bulge (Figure 19). In stage 2 the SBE is likewise outside the trigger interface, so the bridge spans both.

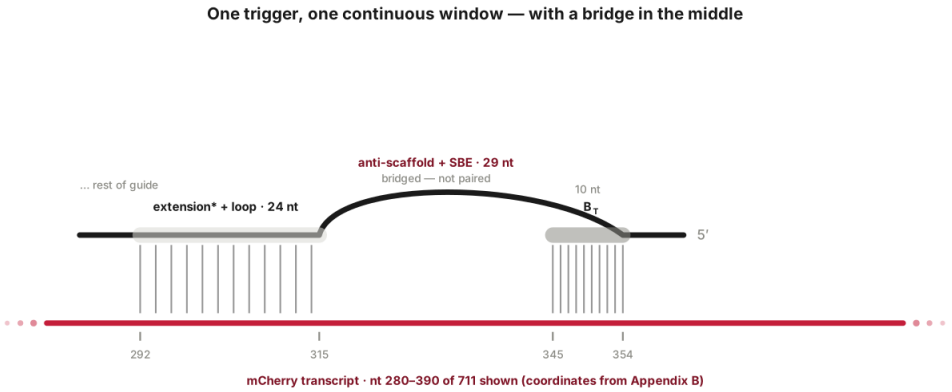

Figure 19. The extended trigger window and its bridge.

5.2.4 Changes to the pipeline

Generation, the loop sweep, the sliding-window scan and the four filters of section 5.1.4 are unchanged, and operate on the designed region only - the anti-scaffold and B_T are appended to designs that have already passed. The structural filters therefore admit the same internal folds as in stage 1; the four categories differ in the assembled construct and in where B_T is taken from, not in which hairpins are admissible.

Two things change in scoring. Trigger similarity is computed over the paired stretches only, with the bridge excluded by construction since it is not meant to pair. And a non-scoring consistency check re-derives B_T independently from each design's own sequence and rejects any design where the two derivations disagree - which happens when the loop + extension* match is not unique in the trigger and the two scripts anchor to different occurrences. The score formula itself is unchanged.

5.2.5 Verifying the anti-scaffold

That check was run separately, on the full assembled construct folded as one molecule rather than as an isolated duplex - the stronger test, since it asks not only whether the anti-scaffold can reach its window but whether it does so in preference to pairing elsewhere in the guide. The calculation used `cernal-rnaviz`, our own ViennaRNA-based folding and visualisation tool, at 30 °C. All four designs were predicted to pair with the scaffold completely and at exactly their intended windows; the structures are reported in section 5.2.7.

The check is structural, not functional: it shows the anti-scaffold occupies its window without disrupting the rest of the fold, not that doing so weakens dCas9 loading, since no protein appears in either folding model.

5.2.6 Backbones

The generation-2 modules were applied to the source switch and the fully-closed gate. The fully-open control was left unchanged, because stage 1 gave no reason to alter it - the problem it exposed was in closing, not opening.

5.2.7 Designs produced

Eight candidates: two backbones - the source switch and the fully-closed gate - each carrying the four anti-scaffolds of section 5.2.2, all at a 14-nt loop (Figures 20 and 21).

Version	Anti-scaffold (DNA)	Length	B_T	Full guide
1	CTCTAAAAC	9 nt	CAGGGAGGAG	173 nt
2	ATTTCTAGCTCTAAAAC	17 nt	CCGTCCTGCA	181 nt
3	AACGGACTAGCCTT	14 nt	TCCTGCAGGG	178 nt
4	CACTTTTTCAAGTTGATAACGGACTAGCCTT	31 nt	TGTAGATGAA	195 nt

Table 12. The four versions, applied identically to both backbones. Full guide length is B_T (10) + anti-scaffold + designed region (74) + scaffold (80). Sequences are written in the DNA alphabet; the panel is specified as RNA in section 5.2.2.

The hairpin is unchanged from stage 1. Each backbone carries its stage-1 rank-1 design verbatim - closed 1 and mCherry 1 - and all four versions of a backbone share it. This is what section 5.2.4 describes: the filters act on the designed region, and the two new modules are appended to designs that have already passed.

The score is identical across the four versions of a backbone - 82,066 for the source switch, 128,112 for the fully-closed gate. Neither the anti-scaffold nor B_T enters the score, and every B_T is a perfect match to its trigger window, so the four versions are indistinguishable to the ranking. They are separated by a structural question the score does not address, which is why the panel exists and why section 5.2.5 verifies it by folding instead.

The B_T sites step along the trigger by exactly the anti-scaffold length. The hairpin senses the same 24-nt window in all eight designs, TGGGAGCGCGTGATGAACTTCGAG, but each version takes B_T from a different position: version 1 to version 3 is a 5-nt step for an anti-scaffold 5 nt longer, version 3 to version 2 a 3-nt step for 3 nt.

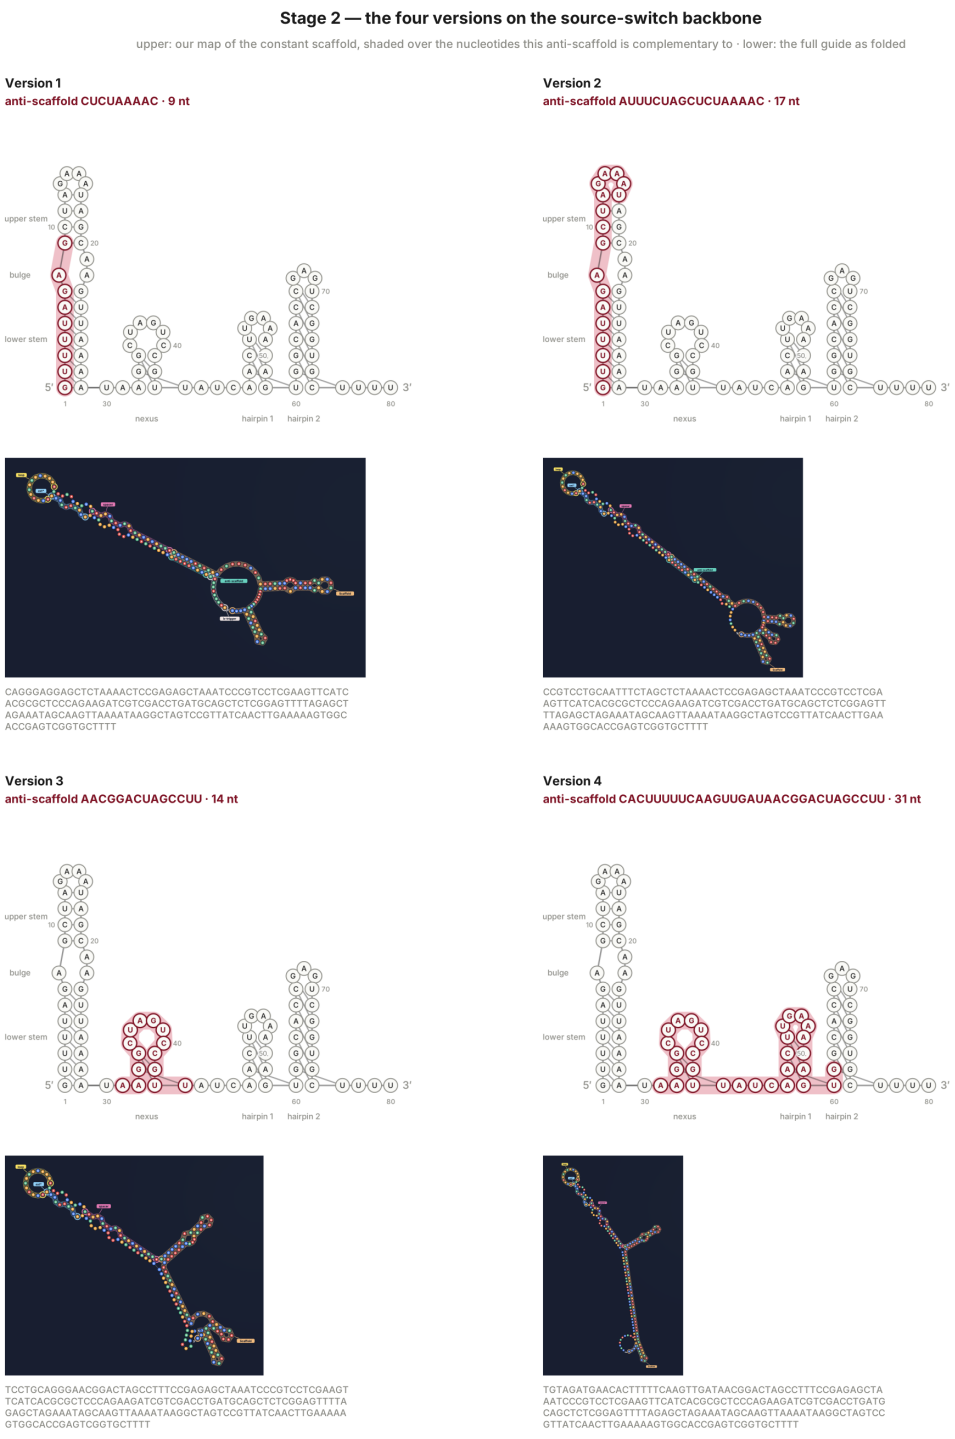

Figure 20. The four generation-2 versions on the source-switch backbone. For each version, our map of the constant scaffold is shaded over the nucleotides the anti-scaffold of that version is complementary to, and the full guide is shown as folded.

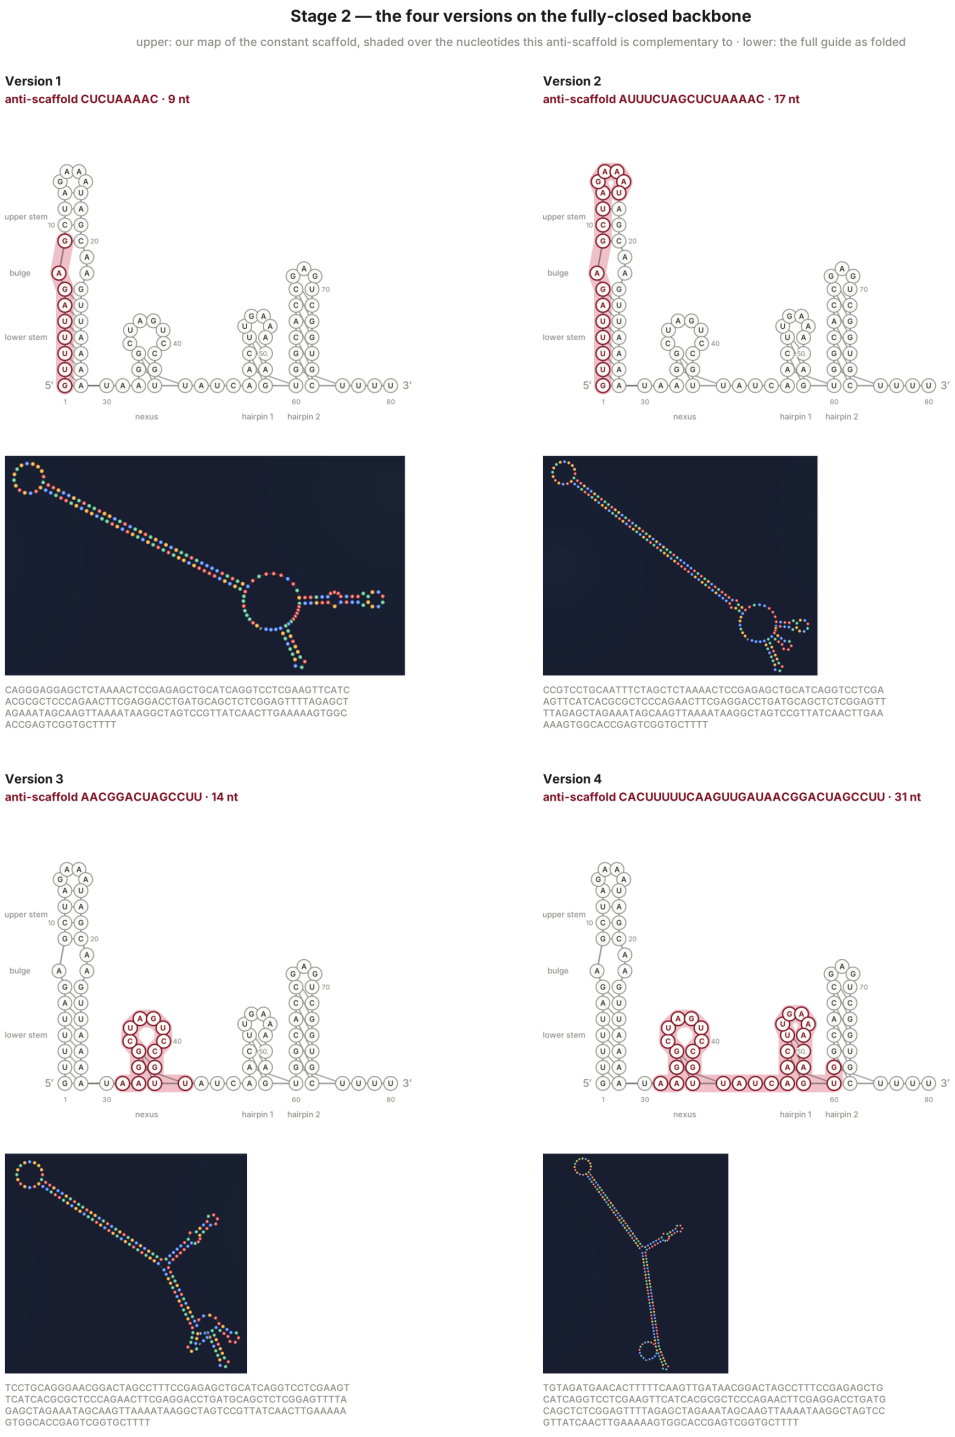

Figure 21. The same four versions on the fully-closed backbone, drawn to the same conventions as Figure 20.

References

[1] Pelea O, Mayes S, Ferry QRV, Fulga TA, Sauka-Spengler T. Specific modulation of CRISPR transcriptional activators through RNA-sensing guide RNAs in mammalian cells and zebrafish embryos. *eLife* 2023;12:RP87722. doi:10.7554/eLife.87722

[2] Lorenz R, Bernhart SH, Höner zu Siederdissen C, Tafer H, Flamm C, Stadler PF, Hofacker IL. ViennaRNA Package 2.0. *Algorithms for Molecular Biology* 2011;6:26. doi:10.1186/1748-7188-6-26

[3] Ferry QRV, Lyutova R, Fulga TA. Rational design of inducible CRISPR guide RNAs for de novo assembly of transcriptional programs. *Nature Communications* 2017;8:14633. doi:10.1038/ncomms14633

A fourth entry is needed for `cernal-rnaviz` *once it has a wiki or repository URL. The ViennaRNA package itself is cited above as [2].*

Appendix A. Notation

Every symbol used in the body appears here and is not redefined elsewhere. Sequence modules are written 5′ to 3′. Free parameters, thresholds and the architecture sets are listed separately in Table 5.

Symbol	Meaning	Units
S	Spacer, fixed by the target locus; length 20 nt	-
T	Trigger window, fixed per iteration; a subsequence of a biomarker transcript	-
C	Scaffold, fixed; required for dCas9 binding	-
s_C	Reference secondary structure of the scaffold, in dot-bracket - the fold dCas9 requires	-
B	SBE - spacer-blocking element; engineered	-
E, E'	Extension and anti-extension; engineered	-
Λ_L	Loop of length L; engineered	-
B_T	Trigger-binding extension of the SBE; sequence follows from the trigger window	-
aS	Anti-scaffold module; one of four graded variants, or none	-
g	The engineered guide	-
D	Endogenous target DNA containing the gene to be regulated	-
M	Set of biomarker transcripts supplied as possible trigger sources	-
$\mathcal{S}_{\text{cand}}$	Candidate spacers from the PAM scan of the target DNA	-
$\mathcal{T}_{\text{cand}}$	Candidate trigger windows from the transcript scan	-
Ω_{off}, Ω_{on}	Equilibrium ensembles of the guide alone, and of guide plus trigger. A generic ensemble is written Ω_X, and Ω_X^{acc} its spacer-accessible subset	-
Z, Z^{acc}	Partition function over an ensemble, and over its spacer-accessible subset	-

Symbol	Meaning	Units
RT	Molar thermal energy: the gas constant times the absolute folding temperature, 303.15 K throughout	kcal·mol^{-1}
A_{off}, A_{on}	Spacer accessibility in the OFF and ON ensembles	Probability, 0–1
q_i	Probability that nucleotide i is unpaired in a given ensemble	Probability, 0–1
$G(\cdot)$	Free energy of the ensemble of the species in brackets	kcal·mol^{-1}
ΔG_{bind}	Trigger-binding free energy, $G(g \cdot T) - G(g) - G(T)$	kcal·mol^{-1}
ΔG_{open}, $\Delta\Delta G$	Cost of opening the OFF stem, and the net drive to open	kcal·mol^{-1}
$n(\cdot)$	Ensemble defect operator $n(\cdot \mid s_C, \Omega)$: the expected number of nucleotides of a region incorrectly paired relative to s_C across the ensemble Ω	nt
$d(\Omega)$	Scaffold ensemble defect, $n(C \mid s_C, \Omega)$	nt
F	Inner feasible set - designs passing all three inner gates	-
J_A, J_E	Structural and thermodynamic components of the inner objective	0–1; kcal·mol^{-1}
J	Inner objective; J^* its constrained maximum for a given pair, and J_+^* its non-negative part	kcal·mol^{-1} (scaled)
$\langle \cdot \rangle_{\text{cal}}$	Average over the calibration set of characterised designs	-
η_{on}, η_{off}	Normalised on-target activity and off-target risk of a spacer	0–1
Q_S, Q_T	Spacer-quality and trigger-quality scores	0–1
$\mathcal{G}$	Outer feasible set - admissible spacer–trigger pairs	-
Φ	Outer objective; the composite ranking score for a pair	Energy-scaled, see section 3.7
p	Equilibrium probability of the designed OFF-state fold (stage-1 score)	Probability, 0–1

Appendix B. Sequences

Full designed-region sequences, 5′ → 3′, in the DNA alphabet. Each loop-14 design is 74 nt and runs SBE → extension* → loop → extension → spacer; the constant scaffold of section 5.1.2 is appended unchanged to all of them and is listed first below.

Constant dCas9 scaffold (80 nt)

```
scaffold   GTTTTAGAGCTAGAAATAGCAAGTTAAAATAAGGCTAGTCCGTTATCAACTTGAAAAAGTGGCACCGAGTCGGTGCTTTT
```

Appended unchanged to the 3′ end of every design listed below and never mutated. Its terminal TTTT is the Pol III terminator. The anti-scaffold windows in Table 11 are 1-based on this sequence.

Trigger - mCherry coding sequence (711 nt)

```
  1  ATGGTGAGCAAGGGCGAGGAGGATAACATGGCCATCATCAAGGAGTTCATGCGCTTCAAG
 61  GTGCACATGGAGGGCTCCGTGAACGGCCACGAGTTCGAGATCGAGGGCGAGGGCGAGGGC
121  CGCCCCTACGAGGGCACCCAGACCGCCAAGCTGAAGGTGACCAAGGGTGGCCCCCTGCCC
181  TTCGCCTGGGACATCCTGTCCCCTCAGTTCATGTACGGCTCCAAGGCCTACGTGAAGCAC
241  CCCGCCGACATCCCCGACTACTTGAAGCTGTCCTTCCCCGAGGGCTTCAAGTGGGAGCGC
301  GTGATGAACTTCGAGGACGGCGGCGTGGTGACCGTGACCCAGGACTCCTCCCTGCAGGAC
361  GGCGAGTTCATCTACAAGGTGAAGCTGCGCGGCACCAACTTCCCCTCCGACGGCCCCGTA
421  ATGCAGAAGAAGACCATGGGCTGGGAGGCCTCCTCCGAGCGGATGTACCCCGAGGACGGC
481  GCCCTGAAGGGCGAGATCAAGCAGAGGCTGAAGCTGAAGGACGGCGGCCACTACGACGCT
541  GAGGTCAAGACCACCTACAAGGCCAAGAAGCCCGTGCAGCTGCCCGGCGCCTACAACGTC
601  AACATCAAGTTGGACATCACCTCCCACAACGAGGACTACACCATCGTGGAACAGTACGAA
661  CGCGCCGAGGGCCGCCACTCCACCGGCGGCATGGACGAGCTGTACAAGTAG
```

The single trigger used throughout stages 1 and 2, supplied as one 711-nt window. Numbering is 1-based from the A of the initiator ATG. The sequence carries no TTTT, so no part of it conflicts with the Pol III terminator filter of section 5.1.2.

Source switch (mCherry-sensing)

```
mCherry 1  TCCGAGAGCTAAATCCCGTCCTCGAAGTTCATCACGCGCTCCCAGAAGATCGTCGACCTGATGCAGCTCTCGGA
mCherry 2  TCCGAGAGCTAAATCCCGTCCTCGAAGTTCATCACGCGCTCCCAGAAAATCGTCGACCTGATGCAGCTCTCGGA
mCherry 3  TCCGAGAGCTAAATCCCGTCCTCGAAGTTCATCACGCGCTCCCAGAAGATCGCTGACCTGATGCAGCTCTCGGA
mCherry 4  TCCGAGAGCTAAATCCTTGATGATGGCCATGTTATCCTCCTCGCCCTTGCTCACCAATGAACATATTCACTGATGCAGCTCTCGGA
```

Fully-closed control

```
closed 1   TCCGAGAGCTGCATCAGGTCCTCGAAGTTCATCACGCGCTCCCAGAACTTCGAGGACCTGATGCAGCTCTCGGA
closed 2   TCCGAGAGCTGCATCAGGTCGGCCATGTTATCCTCCTCGCCCTTTAACATGGCCGACCTGATGCAGCTCTCGGA
```

Fully-open control

```
open 1     TCGGCATAAAACGCCCAGGGCTCGAAGTTCATCACGCGCTCCCAGAACTTCGAGGACCTGATGCAGCTCTCGGA
open 2     CAGGTTAAAGAACCCAGGCGCTCGAAGTTCATCACGCGCTCCCAGAACTTCGAGGACCTGATGCAGCTCTCGGA
open 3     CCTATATAATTAGGAGGGGACTCGAAGTTCATCACGCGCTCCCAGAACTTCGAGGACCTGATGCAGCTCTCGGA
```

mCherry 4 has a 26-nt loop; at 86 nt it is longer than the rest.

Stage 2 - the added modules

Generation-2 guides are assembled as B_T + anti-scaffold + designed region + scaffold. The designed region is `mCherry 1` for the source-switch backbone and `closed 1` for the fully-closed backbone, both listed above, and the scaffold of section 5.1.2 is appended unchanged. Only the two added modules differ between versions, and both backbones use the same four pairs:

```
version 1   B_T  CAGGGAGGAG   anti-scaffold  CTCTAAAAC
version 2   B_T  CCGTCCTGCA   anti-scaffold  ATTTCTAGCTCTAAAAC
version 3   B_T  TCCTGCAGGG   anti-scaffold  AACGGACTAGCCTT
version 4   B_T  TGTAGATGAA   anti-scaffold  CACTTTTTCAAGTTGATAACGGACTAGCCTT
```

Full guide lengths are 173, 181, 178 and 195 nt for versions 1 to 4.