Plasmid-Targeted RNA-Triggered CRISPR Gates

1. Summary

CRISPR systems let a protein be aimed at a chosen site in DNA. 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 binds DNA without cutting it; 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 in the case where the DNA target is ours to write. Besides its spacer, the guide carries 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 invades the fold, the structure opens, the spacer is released and the protein acts: the gate is ON. The target site sits on a synthetic plasmid we deliver, rather than in the host genome.

That one change decides the architecture. When the target is a gene in the host genome its sequence is fixed, the spacer is imposed from outside, and the blocking element is caught between two incompatible demands: it must match the spacer to hold the lock shut, and match the trigger to be invaded. Those are unrelated sequences, so every design is a compromise between them. When we write the target ourselves the spacer is no longer imposed, the blocking element is made perfectly complementary to the trigger, and the compromise disappears.

The lock must still open, so it must still be imperfect. The imperfections move to the spacer - a few deliberate mismatches that weaken the closed hairpin while leaving the trigger duplex untouched. They cost nothing at the target, because the target is written afterwards to match the changed spacer. **The trigger-facing strand is perfect by construction, and all tuning is placed on the opposing strand.**

The framework reported here designs that molecule and chooses its input. Given a set of candidate trigger RNAs, the inner level fixes the blocking element from the trigger, then searches the mismatch pattern, the trigger-binding extension, the loop and the scaffold lock for the design that switches best, reading every quantity from the equilibrium ensemble of the molecule rather than from a single predicted structure. The outer level ranks the triggers themselves, multiplying how good a trigger is as an input by the best switch it allows, so that a promising input cannot compensate for a molecule that does not switch. The output is a ranked list of guides, each paired with the plasmid target site to be built alongside it (see Figure 1).

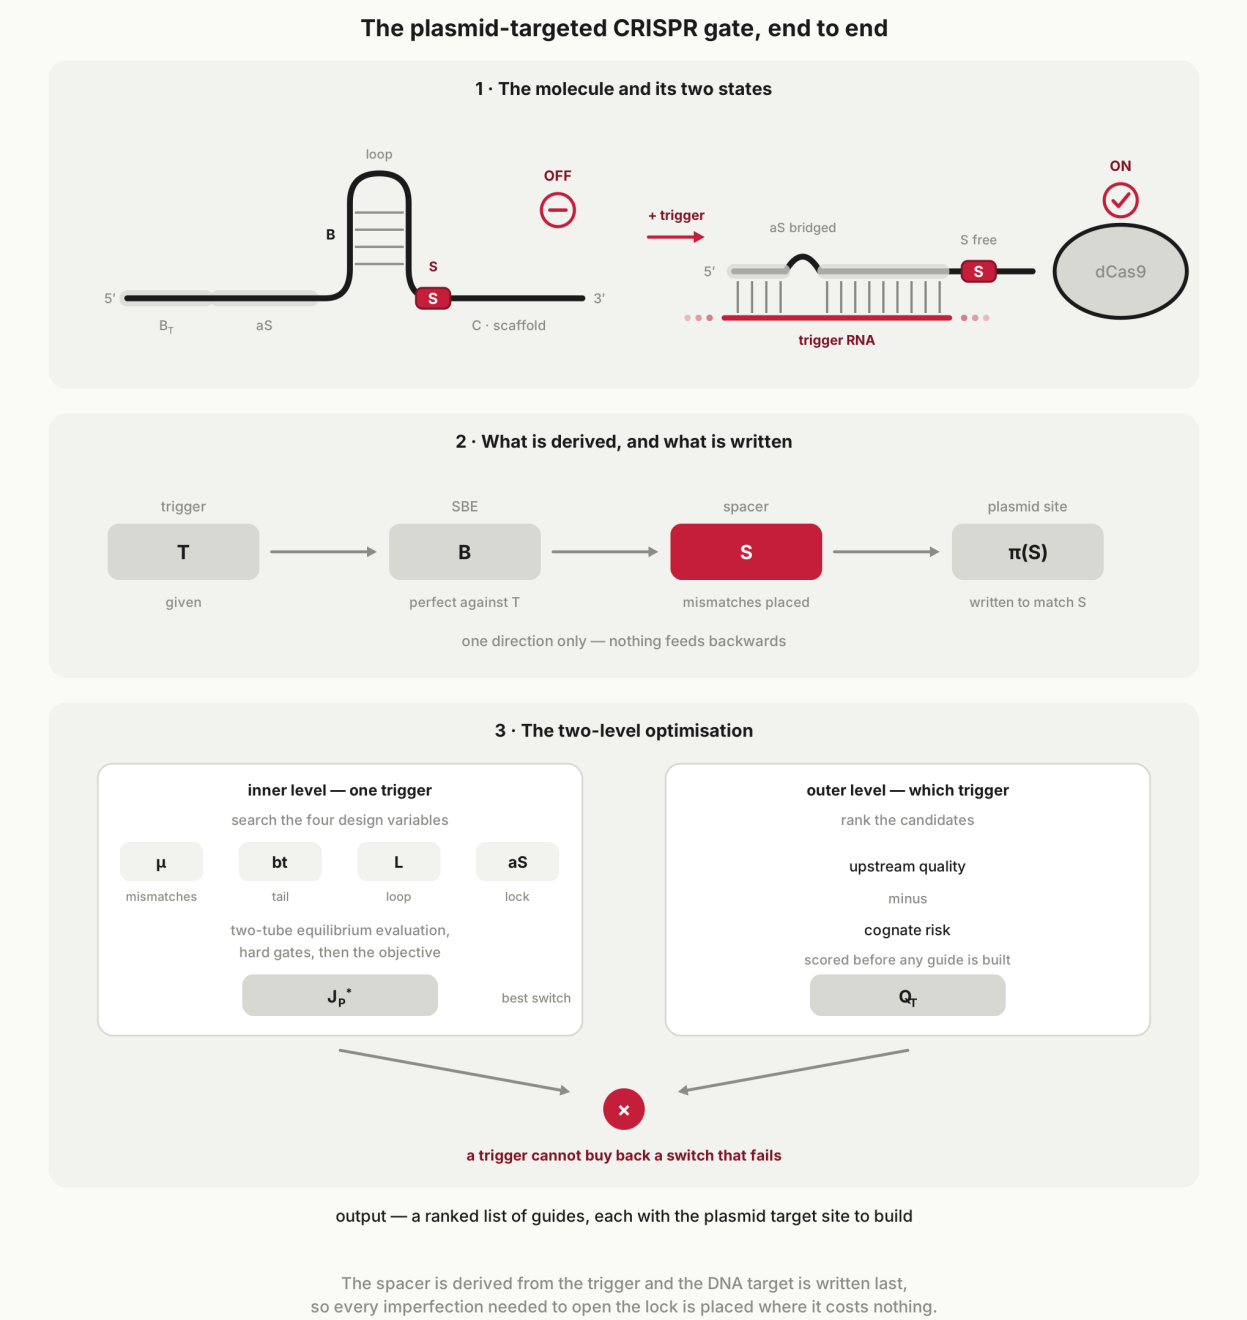

Figure 1. The plasmid gate end to end: what is written and what is derived, the molecule in its two states, and the two-level optimisation that returns both the guide and the target site to build.

2. Introduction

2.1 The parts of the gate

Cas9 is a protein that binds a chosen site in DNA, and it does not find that site itself. It is directed there by a single-guide RNA with two functional parts: the spacer, 20 nucleotides whose sequence decides which DNA site is bound, and the scaffold, a constant sequence with a fixed fold that the

protein grips. The scaffold is identical in every guide and carries no targeting information. The variant used here is dCas9, which binds DNA but does not cut it, so the output of the gate is regulation rather than damage.

Sequence match alone is not enough. A short motif, the protospacer-adjacent motif or PAM, must sit next to the matching stretch of DNA, or the site is invisible to the protein however well the spacer matches. Different Cas variants require different motifs, so the required motif is an input to the design rather than a constant.

The guide is made conditional by appending sequence that folds it back on itself. A segment called the spacer-blocking element, the SBE, is complementary to the spacer, and the two pair into the stem of a hairpin closed by a loop. While the spacer sits inside that stem it cannot pair with DNA, and dCas9 cannot be productively directed anywhere.

Two further modules act on the two states separately. The anti-scaffold is a short segment complementary to a window of the constant scaffold, which it pairs with in cis; it is a second, independent lock that weakens dCas9 loading while the gate is shut. The trigger-binding extension B_T sits at the 5′ end and pairs the trigger perfectly, lengthening the trigger interface and helping the trigger do its work. The scaffold itself is never mutated, since its sequence is set by the protein interface; all modulation happens through these separate, displaceable modules.

2.2 How it switches

The trigger is an RNA molecule present in the cell when the gate should fire. It pairs with one continuous window of the guide that runs from B_T through the SBE to the loop. The anti-scaffold lies inside that window but is already paired with the scaffold, so the trigger is not able to pair with it and bridges over it as a single long bulge.

The loop and B_T are the only unpaired regions of the closed molecule, and they are where the trigger lands. From there it pairs inward along the SBE, displacing the spacer base by base until the stem is gone. The spacer is never contacted by the trigger: it is released indirectly, by loss of its partner (see Figure 2).

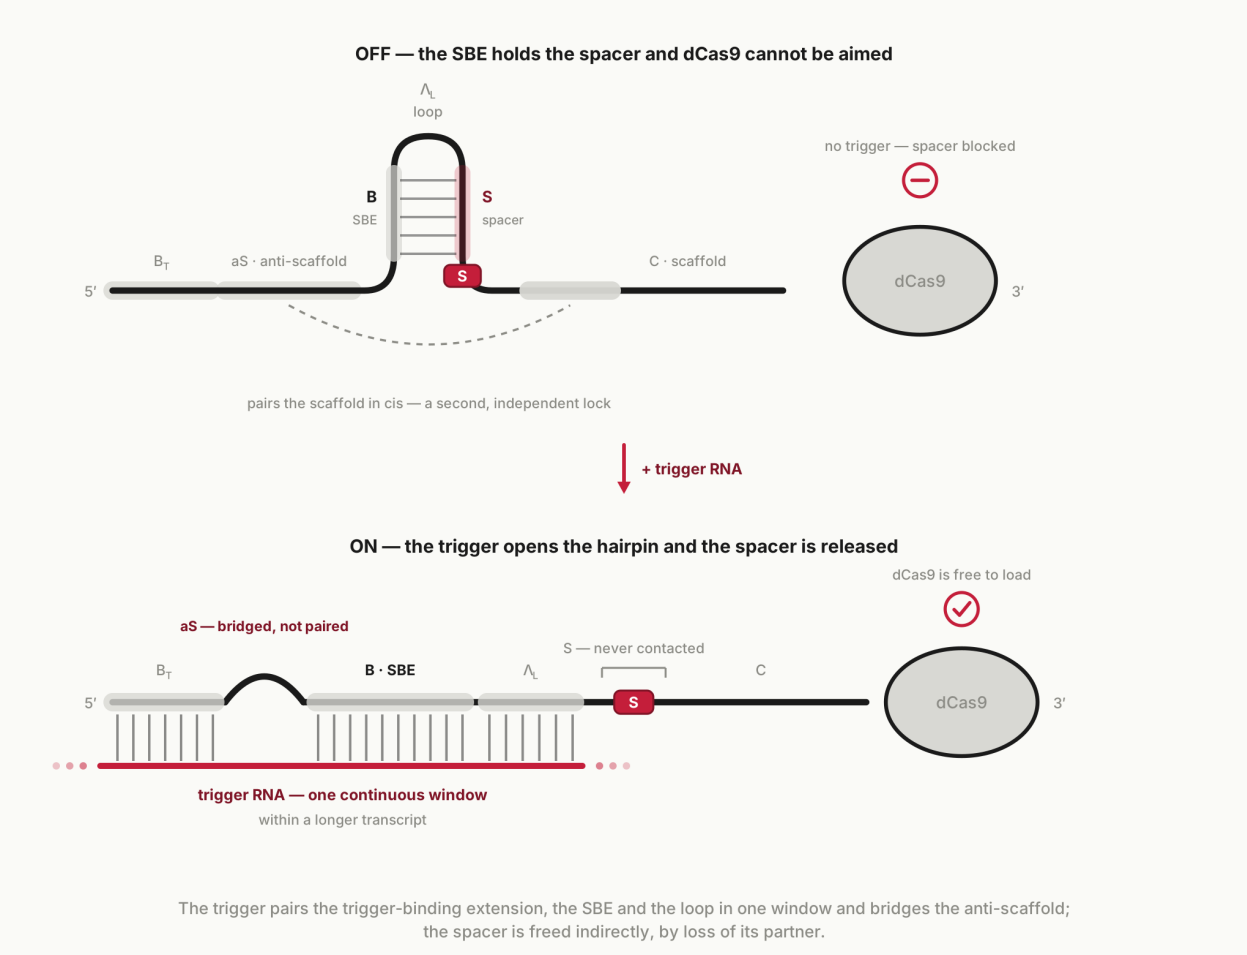

Figure 2. The guide in its OFF and ON states. The trigger pairs the trigger-binding extension, the SBE and the loop in one continuous window, bridging the anti-scaffold, and the spacer is freed without ever being touched.

2.3 Why a written target changes the problem

The DNA site the activated guide binds can be of two kinds, and the design problem is not the same in the two cases. In the endogenous architecture the target is a gene in the cell's own genome. That sequence cannot be changed, so the spacer is fixed from the outside and must match the genomic site. The SBE then has to satisfy two templates at once - the reverse complement of the spacer, to lock it, and the reverse complement of the trigger, to be invaded. Those differ, and no single SBE satisfies both.

In the plasmid architecture the target site is not found but written. We place it on a synthetic plasmid introduced into the cell, so its sequence is ours to choose. The spacer is no longer imposed, the SBE can be made perfectly complementary to the trigger with nothing conceded on the trigger side, and the DNA target is written afterwards to match whatever spacer results.

The lock still has to open, so it still has to be imperfect, but the imperfections are now written into the spacer rather than into the SBE. A mismatch in the spacer weakens the closed hairpin and leaves the trigger duplex untouched, so the trigger wins by exactly the amount we chose to concede. The

same mismatch costs nothing at the target, because the target is written to match the changed spacer (see Figure 3).

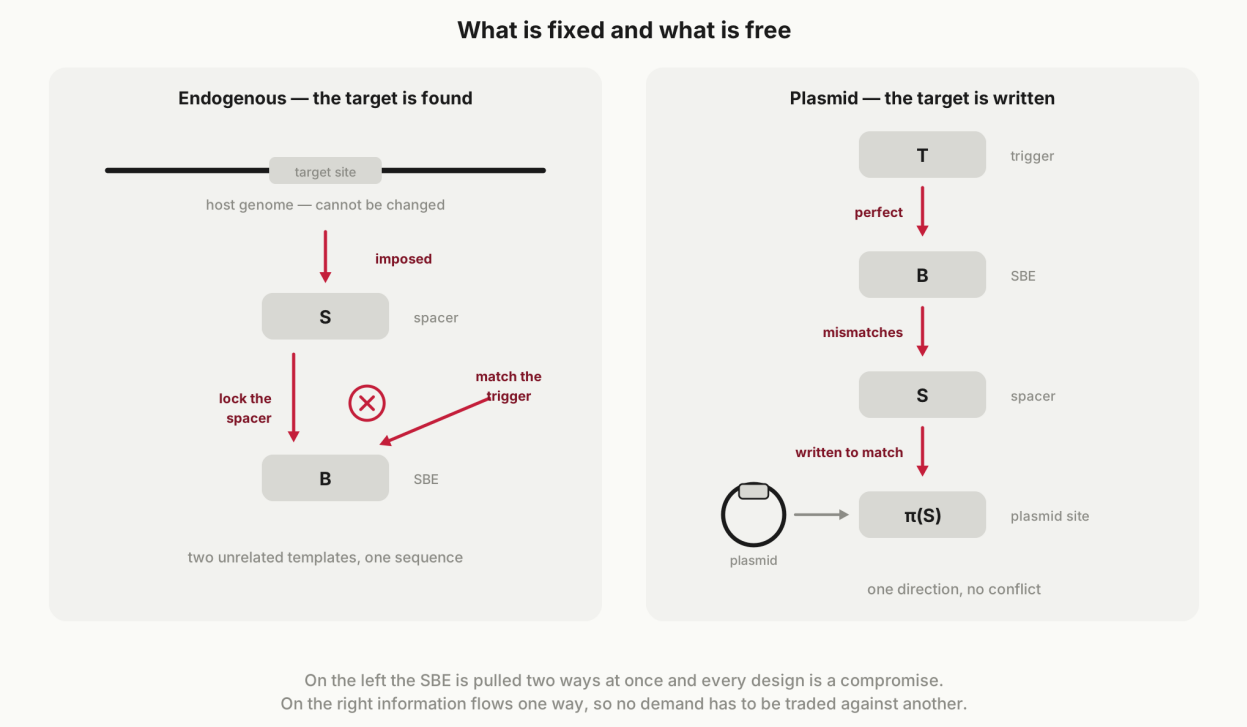

Figure 3. What is fixed and what is free in the two architectures. On the left, information flows into the SBE from two directions and collides. On the right it flows one way, from the trigger to the SBE to the spacer to the plasmid site, with no conflict.

2.4 The two configurations

The gate inverts cleanly at the protein. Fusing an activator to dCas9 makes the target gene rise when the trigger is present; fusing a repressor makes it fall. The RNA design is identical in both cases, and nothing in this framework changes between them.

3. Results: a two-level framework for plasmid-targeted gates

3.1 The algorithm in outline

The design stage receives four inputs: a set of candidate triggers T_{cand}, supplied by the upstream selection stage; the off-target reference Γ, meaning all DNA present in the cell that dCas9 could bind by mistake, host genome and delivered constructs alike; the PAM definition P of the chosen Cas variant; and the constant scaffold C. It returns, for each trigger, a complete guide and the DNA sequence of the target site to be written into the plasmid.

Because the spacer is derived from the trigger rather than chosen independently, there is no search over spacer–trigger pairs. There is one optimisation per trigger, repeated across the set. Greek symbols are spelled out in the code - `tau_off` for τ_{off}, `Lam` for Λ_L, `T_c` for $\overline{T}$, and so on.

```
PROCEDURE PlasmidDesign(T_cand, C, s_C, P, Gamma) -> ranked list
  for each T in T_cand :
      T_c   <- revcomp(T)              # the trigger-facing template
      rho   <- CognateRisk(T, P, Gamma)        # section 3.7
      J_best <- -INF ;  g_best <- NONE
      # the modules that pair T must span it exactly: bt + |B| + L = |T|
      for each (bt, L) with bt in BT_set, L in L_set, bt + 20 + L = |T| :
          B_T <- first bt nt of T_c     # trigger-binding extension
          B   <- next 20 nt of T_c      # SBE, perfect against T
          Lam <- remaining L nt of T_c  # loop
          S_0 <- revcomp(B)             # trigger-derived spacer
          for each aS in {aS1, aS2, aS3, aS4, none} :     # scaffold lock
              for each mu in MismatchPatterns(S_0) :      # section 3.3
                  S <- apply(S_0, mu)
                  g <- B_T . aS . B . Lam . S . C         # equation (1)
                  # observables, two tubes (section 3.4)
                  A_off   <- Accessibility(g,     spacer region)
                  A_on    <- Accessibility(g & T, spacer region)
                  dG_bind <- G(g.T) - G(g) - G(T)
                  d_on    <- Defect(g & T, s_C, scaffold region)
                  # hard gates (section 3.5)
                  if not ( A_off <= tau_off and A_on >= tau_on
                           and d_on < eps and eta_off(S) <= o_max ) :
                      skip mu
                  # inner objective (section 3.6)
                  J_A <- w_on * A_on - w_off * A_off
                  J_E <- -dG_bind
                  J_S <- eta_on(S)
                  J   <- w_E*J_E + lambda*( w_A*J_A + w_S*J_S )
                  if J > J_best :  J_best <- J ;  g_best <- g
      Q   <- u_alpha * alpha(T) - u_rho * rho        # section 3.7
      Phi <- Q * J_best                              # equation (15)
      emit (T, g_best, target_site(g_best), Phi)
  return the emitted designs sorted by Phi
```

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 $|B| = 20$ are tied together by one equation instead of varied independently. And the mismatch loop is where the lock is weakened, because it is the only place where weakening is possible.

3.2 The construct and the design chain

The guide is the concatenation of six modules, written 5′ to 3′:

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

where B_T is the trigger-binding extension, aS the anti-scaffold, B the SBE, Λ_L the loop of length L, S the spacer and C the scaffold. In the OFF state S pairs with B to form the stem, the loop closes it, and B_T and Λ_L are left single-stranded. The endogenous version of this architecture carries two further modules between the spacer and the loop, an extension and its near-complement; they are not used here, and the design is simpler without them.

Three of the modules are read off the trigger rather than searched. Writing $\overline{x}$ for the reverse complement of a sequence x, the trigger complement is cut into three consecutive segments in the same order they appear along the guide:

$$B_T \cdot B \cdot \Lambda_L = \overline{T}, \qquad bt + |B| + L = |T| \tag{2}$$

The spacer follows from the SBE. The sequence that pairs with B perfectly is the trigger-derived spacer $S_0(T) = \overline{B}$, and a spacer equal to it would lock perfectly and never open, since the trigger would have to displace a duplex exactly as strong as the one it forms. So a few bases are deliberately changed:

$$S = S_0(T) \oplus \mu \tag{3}$$

where μ is the mismatch pattern, a list of positions and the base each is changed to, and $\oplus$ means applying those changes. The plasmid target site $\pi(S)$ is then written as the DNA matching S exactly, changed bases included, with a PAM drawn from P installed beside it and positioned relative to the promoter as the effector requires. The target site is never searched for - it is an output.

The chain runs in one direction, and nothing feeds backwards. The trigger fixes the SBE, the SBE fixes the trigger-derived spacer, the mismatch pattern turns that into the built spacer, and the built spacer fixes the plasmid site (see Figure 4).

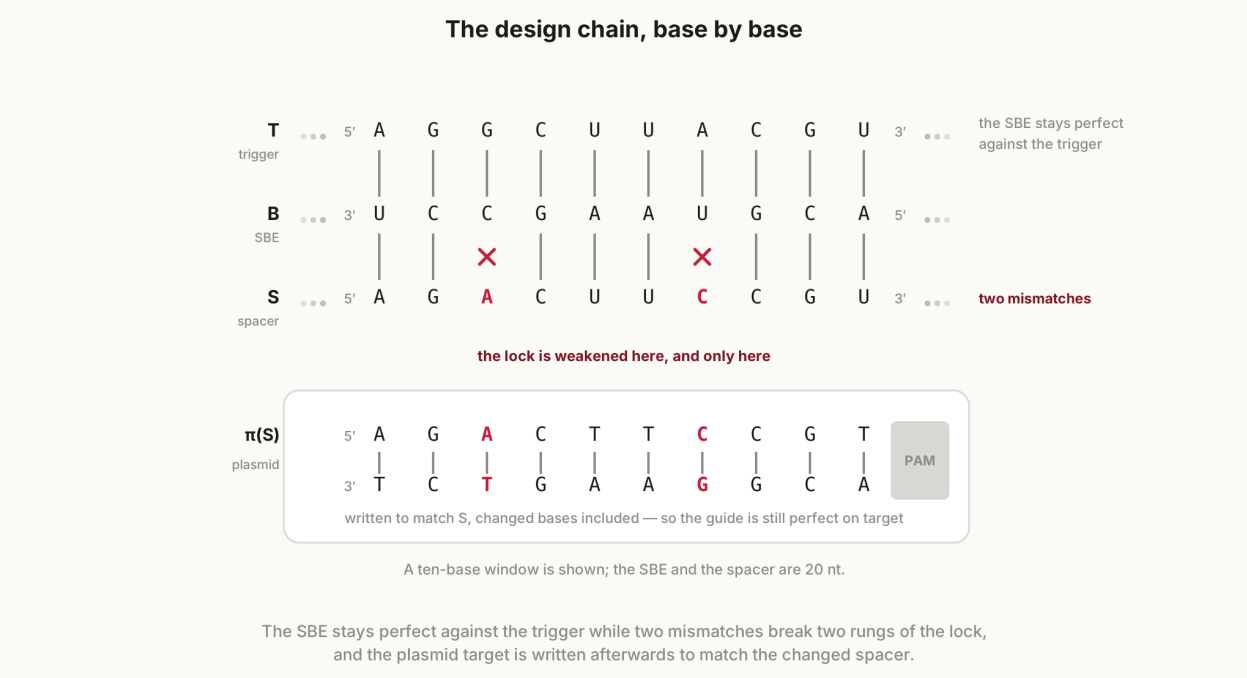

Figure 4. The design chain on a ten-base window. The SBE is perfect against the trigger and stays perfect; two mismatches in the spacer break two rungs of the lock, and the plasmid target is written to match the changed spacer, so the guide remains a perfect match on target.

3.3 Design variables and where the opening drive comes from

The inner optimisation has four design variables: the mismatch pattern μ, the length bt of the trigger-binding extension, the loop length L, and the choice of anti-scaffold aS among four graded variants or none. Everything else is either given - T, C, P, Γ - or determined by what is given, namely B, Λ_L, S_0, and through them S and $\pi(S)$.

In the endogenous architecture part of the drive to open the hairpin comes from an extension region the spacer never occupies, which can be matched to the trigger at no cost to the lock. Here there is

no such free region, because the entire stem is $S \cdot B$ and the entire stem is contested. The drive comes from two sources instead, and both are unopposed: the trigger-binding extension, because the spacer does not reach it, and the mismatch budget, because the trigger duplex does not depend on it (see Figure 5).

Set against them is the stem itself, and now the anti-scaffold as well, which has to be displaced before the trigger can bridge across it. Lengthening B_T is what pays for that displacement, which is why the two modules were introduced together.

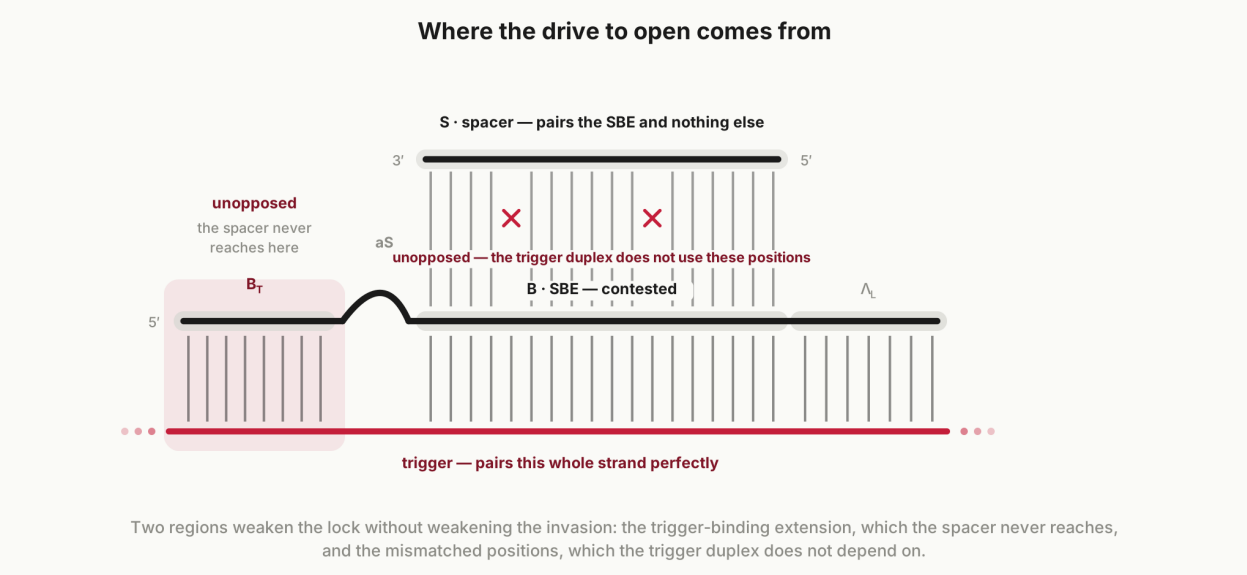

Figure 5. The two unopposed sources of opening drive, and the contested region between them. The spacer never reaches the trigger-binding extension, and the trigger duplex is indifferent to the mismatched positions, so both weaken the lock without weakening the invasion.

The mismatch budget $k = |\mu|$ is a variable in its own right. Too few and the hairpin never opens; too many and it falls open without the trigger. Which positions are edited matters as much as how many, which is why μ carries both. The two requirements are driven by the same variable in opposite directions, so the balance is computed rather than guessed (see Figure 6).

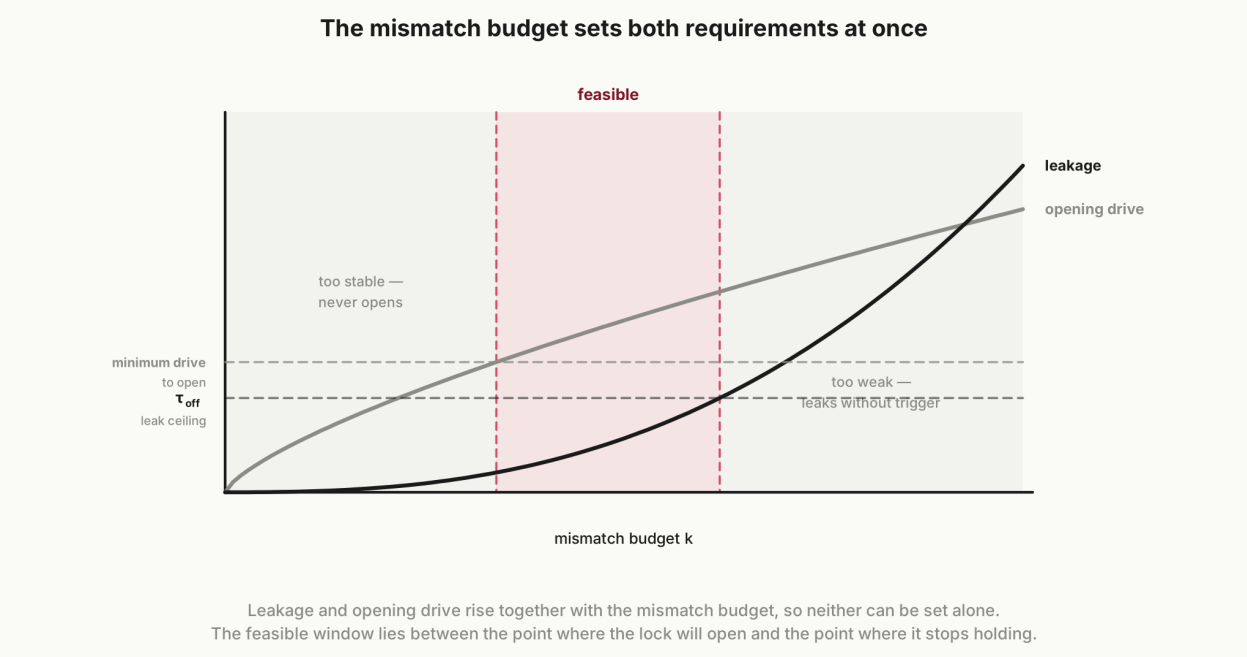

Figure 6. The mismatch budget. Leakage and opening drive both rise with the mismatch budget, which leaves a feasible window between a hairpin too stable to open and one too weak to stay shut.

3.4 Two tubes, four observables

The design is scored at thermodynamic equilibrium. Rather than ask what single structure the guide adopts, we ask what the whole population looks like, and how that population changes when the trigger is added. Two ensembles are evaluated, with the trigger in excess in the second:

$$\Omega_{\text{off}} = \langle\{g\}\rangle, \qquad \Omega_{\text{on}} = \langle\{g, T\}\rangle \tag{4}$$

Spacer accessibility is the probability that the guide adopts a conformation in which the spacer is free. For an ensemble X the partition function is

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

where Ω_X is the set of structures available in that ensemble, ΔG_i the folding free energy of structure i, and RT the molar thermal energy. Each structure contributes a weight that grows as its energy falls, so stable structures dominate the sum. Restricting the same sum to the accessible subset Ω_X^{acc}, meaning those structures in which all 20 spacer nucleotides are unpaired, gives

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

and accessibility is the fraction of the population that qualifies:

$$A_X = \frac{Z_X^{\text{acc}}}{Z_X} \tag{7}$$

evaluated as A_{off} and A_{on} in the two tubes (see Figure 7). A good design keeps A_{off} low and A_{on} high. The criterion is strict on purpose: the mismatch pattern leaves unpaired nucleotides inside the spacer by construction, and a closed guide carrying them must not be counted as open.

Trigger binding measures how favourable the guide–trigger complex is:

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

where $G(\cdot) = -RT\ln Z$ is the ensemble free energy of the species in brackets. A more negative value means stronger binding. Accessibility says what fraction of the population is open; this says how much free energy is released when the trigger binds.

Scaffold integrity is quantified by the ensemble defect

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

the expected number of nucleotides of the scaffold region C that deviate, across the ensemble Ω, from the reference structure s_C that dCas9 requires. The anti-scaffold raises this in the OFF ensemble deliberately - that is its purpose. What must be preserved is its value in the ON ensemble, so that dCas9 loads once the switch has opened.

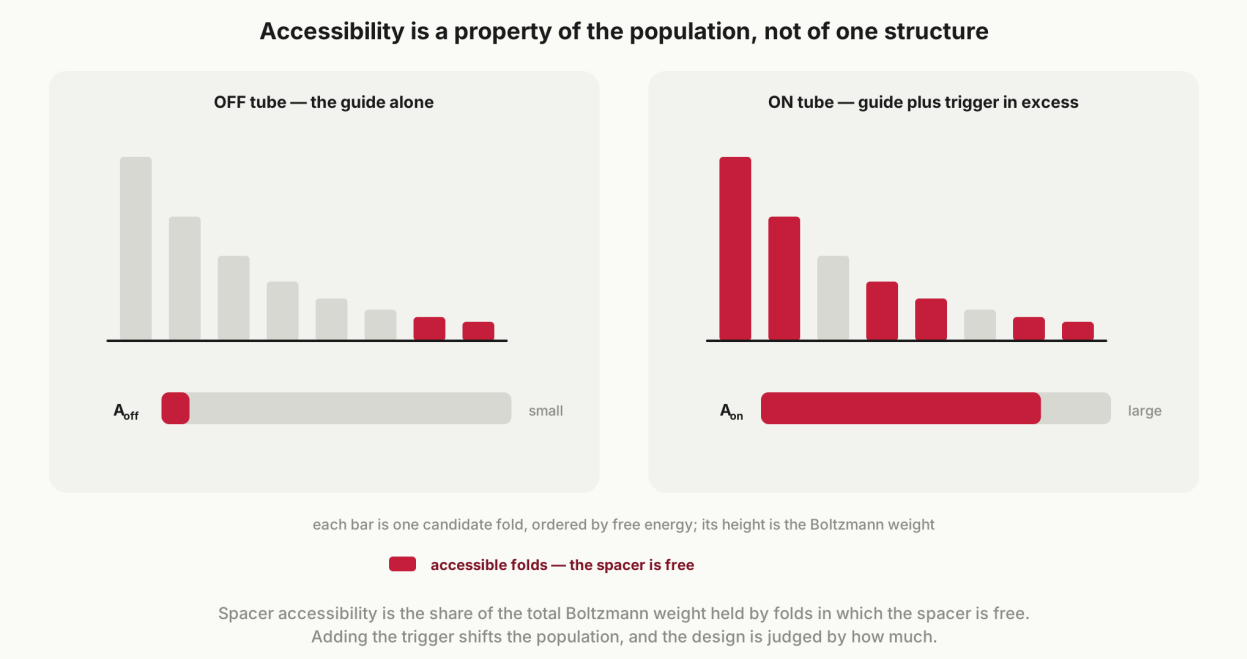

Figure 7. Accessibility as a property of a population, not of one structure. Each candidate fold carries a Boltzmann weight; accessibility is the share of the total weight held by folds in which the spacer is free, and it grows when the trigger is added.

The two observables pull against each other. Raising the mismatch budget destabilises the closed guide, which raises A_{off} and is penalised as leakage, and at the same time makes ΔG_{bind} more favourable through its $G(g)$ reference term, which is rewarded. The mismatch budget is the explicit dial between leakage and openness, and the objective is what settles it.

3.5 Hard gates

Before a design is scored it must pass four pass/fail filters. These express requirements that cannot be traded against one another: a guide that leaks, fails to open, breaks the scaffold when open, or points at too much of the genome is not a worse design but a non-functional one.

$$F_P(T) = \{(\mu, bt, L, aS) : A_{\text{off}} \le \tau_{\text{off}},\ A_{\text{on}} \ge \tau_{\text{on}},\ d(\Omega_{\text{on}}) < \varepsilon,\ \eta_{\text{off}}(S) \le o_{\max}\} \tag{10}$$

There is deliberately no OFF-state scaffold gate. In the OFF state the spacer is already occluded, so scaffold competence is immaterial, and the anti-scaffold may be disrupting it on purpose; imposing fidelity there would reject exactly the designs that module exists to produce.

The fourth condition is specific to this architecture. The spacer is produced here rather than selected from outside, and two mismatch patterns of the same size give different genomic behaviour, so off-target risk cannot be assessed until μ is chosen. Designs failing any condition are discarded and never reach the objective.

3.6 The inner objective

For designs inside $F_P(T)$, and $J_P = -\infty$ outside it, the objective combines three descriptions of how well the design performs.

$J_A = w_{\text{on}}A_{\text{on}} - w_{\text{off}}A_{\text{off}}$ **is the structural component.** It rewards an open spacer when the trigger is present and penalises residual accessibility when it is absent, so it measures the realised size of the switch rather than either state alone.

$J_E = -\Delta G_{\text{bind}}$ **is the thermodynamic component.** It measures the energetic favourability of trigger binding, which is the force that drives the transition.

$J_S = \eta_{\text{on}}(S)$ **is the on-target component.** The spacer is produced by this optimisation rather than selected outside it, so the quality of the resulting spacer has to be scored where it is decided. Without this term the optimiser would place mismatches for switching alone and accept whatever targeting the resulting spacer happened to give.

The first and third components are dimensionless and the second is an energy, so the dimensionless block is brought onto the energy scale by a single factor λ, taken from the calibration set as the ratio of the typical magnitudes of the two kinds of quantity. The scaling makes the terms numerically comparable; their relative importance is carried separately by the weights.

$$J_P = w_E J_E + \lambda(w_A J_A + w_S J_S), \qquad w_E + w_A + w_S = 1 \tag{11}$$

The selected design is the constrained maximiser

$$(\mu^*, bt^*, L^*, aS^*) = \arg\max_{(\mu,bt,L,aS)\in F_P(T)} J_P \tag{12}$$

and its value $J_P^*(T)$ is carried up to the outer level as the switching quality achievable with trigger T.

The model is an equilibrium one: it describes the end states, not the speed of the transition between them. The trigger-binding extension enters through whether the ON state opens at equilibrium, not through how fast it opens.

3.7 Selecting the trigger

The inner problem returns the best gate for a given trigger. The outer problem asks which trigger to use, and since the spacer is derived rather than chosen, it reduces to a ranking over T_{cand}.

One risk is specific to this architecture and has to be scored before anything else. The spacer is a near-copy of the trigger, and the trigger was transcribed from somewhere; the gene or construct encoding it therefore carries DNA that nearly matches our spacer. That site is the cognate locus $\chi(T)$, and we know it exists before looking at the reference at all, simply from how the spacer was built.

Whether it matters depends on the PAM. We inspect the sequence adjacent to $\chi(T)$ and score how well it satisfies P: a full match makes the site engageable and the risk maximal, a degenerate match leaves it intermediate, and no match makes the site unreachable. That score is $\rho(T) \in [0,1]$, computable from the trigger alone, which is why it enters here rather than inside the guide optimisation. The cognate locus is excluded from Γ when η_{off} is computed, so it is counted once, here, where its meaning is represented (see Figure 8).

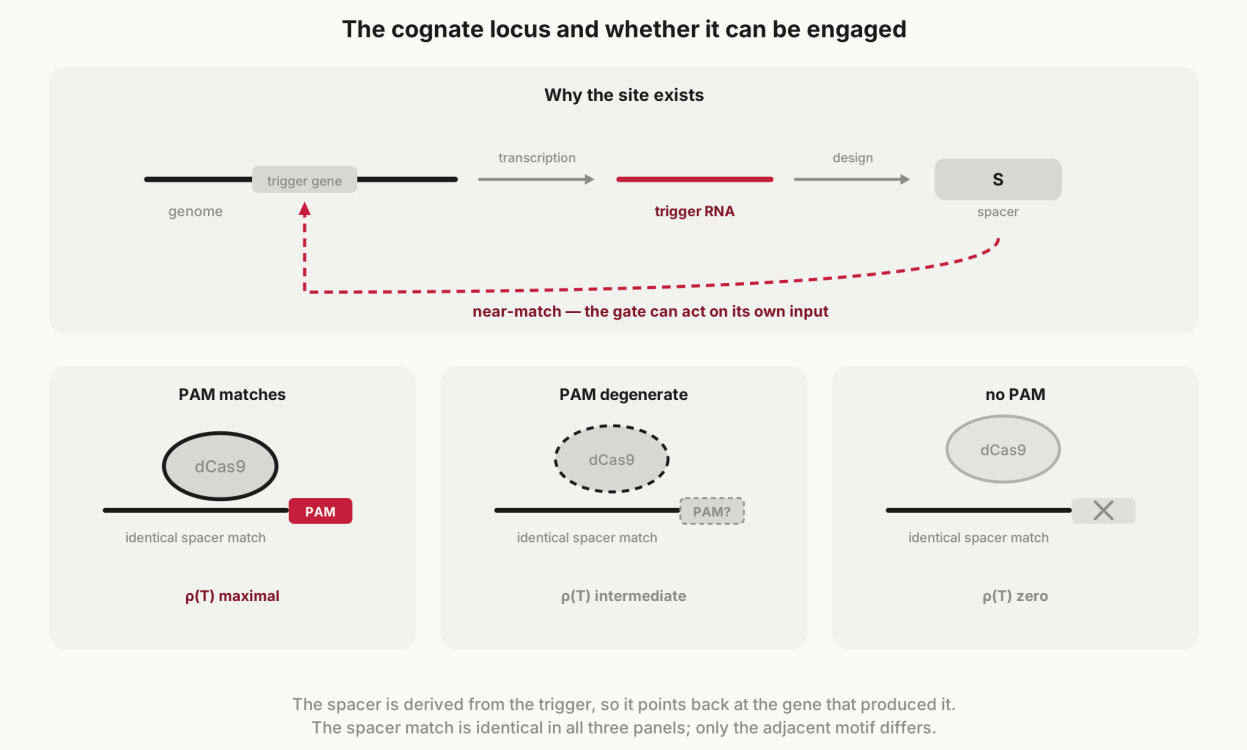

Figure 8. The cognate locus and its PAM dependence. The spacer is derived from the trigger, so it points back at the gene that produced the trigger; whether dCas9 can act there is decided by the three bases beside the site, which are identical in sequence match across all three panels.

Two properties decide how good a trigger is: its quality as an input, supplied by the upstream stage as $\alpha(T)$, and the cognate risk $\rho(T)$. They combine into

$$Q_T(T) = u_\alpha\, \alpha(T) - u_\rho\, \rho(T), \qquad u_\alpha + u_\rho = 1 \tag{13}$$

The cognate risk is a penalty and not a gate. A trigger whose cognate locus is engageable is ranked down rather than excluded, and a strong trigger may still be preferred to a weaker one with a dead cognate site. The weight u_ρ should be calibrated with the circuit in mind: engaging the cognate locus couples the gate onto its own input, suppressing the trigger's production with a repressive effector and amplifying it with an activating one, which is a circuit-level failure rather than an isolated off-target event.

A trigger is admissible if a working guide exists for it at all:

$$T_{\text{adm}} = \{T \in T_{\text{cand}} : F_P(T) \neq \varnothing\} \tag{14}$$

Switching then enters multiplicatively, so that trigger quality cannot compensate for a device that fails to switch:

$$\Phi_P(T) = Q_T(T) \cdot J_P^*(T) \tag{15}$$

A good trigger leaves the achievable switching score largely intact, while a poor or risky one discounts it. Because the product is taken rather than the sum, a trigger cannot earn a high final score on input quality alone.

3.8 Output, parameters and calibration

The design returned for the gate is

$$T^* = \arg\max_{T \in T_{\text{adm}}} \Phi_P(T) \tag{16}$$

In practice the result is a ranked list of triggers rather than a single winner, each with its optimal guide and the plasmid target site to be built alongside it, which leaves the final choice to experimental validation.

Φ_P is expressed on an energy scale but should be read as an energy-scaled composite ranking score rather than as a physical free energy difference. The energetic component comes from the thermodynamic term; structural accessibility, spacer activity and trigger quality enter through calibrated dimensionless factors. Keeping the output energy-compatible lets the plasmid gate be compared with the endogenous gate and with other gate models on the same scale.

The weights w_{on}, w_{off}, w_E, w_A, w_S, u_α and u_ρ, the scale factor λ, and the thresholds τ_{off}, τ_{on}, ε and $o_{\max}$ are left symbolic throughout. They are fixed by calibration against experimentally characterised designs - reconstructing published guides, running them through the same pipeline, and fitting the objective to their measured ON/OFF behaviour - rather than chosen by hand. All folding is carried out at 30 °C.

4. Discussion

The architecture removes a compromise rather than managing it. In the endogenous gate the SBE is asked to satisfy two unrelated templates, and the objective exists partly to settle how much of each demand to give up. Writing the target removes one of the two templates: the trigger-facing strand becomes perfect, and the optimiser is left tuning a strand whose imperfections cost nothing on target. The free parameters that remain describe how much lock to concede, not which of two goods to sacrifice.

The anti-scaffold should do more here than it did on the endogenous switch. Tested on the endogenous gate it tightened the fully-closed control measurably but did not rescue the source switch, whose hairpin was still too open without its trigger. The plasmid stem is perfect apart from the mismatches we place in it, which makes it structurally much closer to the fully-closed control than to that switch - the backbone on which the scaffold lock worked. It also makes the mismatch

budget the one thing standing between a tight lock and a stuck one, which is why it carries its own treatment above.

What the model does not cover. It describes end states and not rates, so a design that is thermodynamically open but kinetically slow would look acceptable here. A kinetic treatment of strand displacement would be added only if that case appeared in practice. The accessibility criterion is strict, and if partial spacer exposure proves sufficient for dCas9 it can be relaxed during calibration.

Next steps. Calibrate the weights, the scale factor and the thresholds against characterised designs; sweep the mismatch budget on a small panel to locate the feasible window experimentally; and build the first ranked output as a plasmid pair, guide and written target together, for validation.

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.

[2] 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.

[3] Zadeh JN, Steenberg CD, Bois JS, Wolfe BR, Pierce MB, Khan AR, Dirks RM, Pierce NA. NUPACK: analysis and design of nucleic acid systems. Journal of Computational Chemistry 2011;32:170–173.

[4] 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.

[5] CERNAL, iGEM TAU 2026. Endogenous RNA-triggered CRISPR gates. Team wiki. The anti-scaffold panel and the trigger-binding extension are specified there in full.

Appendix A. Notation

Every symbol used in the body appears here and is not redefined elsewhere. Sequence modules are written 5′ to 3′.

Symbol	Meaning
g	The engineered guide
C	Scaffold; constant, required for dCas9 binding
s_C	Reference secondary structure of the scaffold, the fold dCas9 requires
S	Spacer, 20 nt; built from the trigger, not imposed
B	SBE, the spacer-blocking element; perfect against the trigger
B_T	Trigger-binding extension of length bt, at the 5′ end
aS	Anti-scaffold; one of four graded variants, or none
Λ_L	Loop of length L
T, T_{cand}	Trigger window, and the supplied set of candidate triggers
$\overline{x}$	Reverse complement of a sequence x

Symbol	Meaning
$S_0(T)$	Trigger-derived spacer, $\overline{B}$; pairs B perfectly
μ, k	Mismatch pattern, and the mismatch budget $k = \lvert\mu\rvert$
$\oplus$	Application of a mismatch pattern to a sequence
$\pi(S)$	Plasmid target site emitted for a spacer: DNA matching S, PAM installed
P	PAM motif required by the selected Cas variant
Γ	Off-target reference: all DNA in the cell dCas9 might bind
$\chi(T)$	Cognate locus: the place in Γ encoding the trigger
Ω_{off}, Ω_{on}	Equilibrium ensembles of the guide alone and of guide plus trigger; a generic ensemble is Ω_X and Ω_X^{acc} its spacer-accessible subset
Z, Z^{acc}	Partition function over an ensemble, and over its accessible subset
ΔG_i	Folding free energy of structure i; kcal·mol^{-1}
RT	Molar thermal energy: gas constant times absolute folding temperature, 303.15 K throughout; kcal·mol^{-1}
A_X; A_{off}, A_{on}	Spacer accessibility in an ensemble X, and its values in the OFF and ON ensembles; a probability between 0 and 1
$G(\cdot)$	Ensemble free energy of the species in brackets, $-RT\ln Z$; kcal·mol^{-1}
ΔG_{bind}	Trigger-binding free energy; kcal·mol^{-1}
$n(\cdot)$, $d(\Omega)$	Ensemble defect operator, and the scaffold defect $n(C \mid s_C, \Omega)$; nt
η_{on}, η_{off}	Normalised on-target activity and off-target risk of a spacer
$\alpha(T)$, $\rho(T)$	Upstream trigger quality, and cognate-engagement risk; both normalised to 0–1
Q_T	Combined trigger score; normalised to 0–1
J_A, J_E, J_S	Structural, thermodynamic and on-target components of the inner objective; J_E in kcal·mol^{-1}, the other two normalised to 0–1
J_P, J_P^*	Inner objective, and its constrained maximum for a trigger; a star marks a constrained optimum throughout; on the scaled energy axis, kcal·mol^{-1}
Φ_P	Outer objective; the composite ranking score for a trigger, read on the same scaled energy axis
F_P, T_{adm}	Inner feasible set of designs, and the admissible triggers
w_{on}, w_{off}	Weights inside J_A
w_A, w_E, w_S	Weights of the three components of J_P; they sum to 1
u_α, u_ρ	Weights inside Q_T; they sum to 1
λ	Scale factor placing the dimensionless components on the energy scale; kcal·mol^{-1}
τ_{off}, τ_{on}	Thresholds on accessibility in the OFF and ON states; 0–1
ε	Threshold on the scaffold ensemble defect; nt
$o_{\max}$	Threshold on off-target risk; 0–1

Table 1. Every symbol used in the document, grouped by the order in which it appears. Weights and thresholds are left symbolic and fixed by calibration.