

Amplification at Equilibrium: Structural and Thermodynamic Limitations, and Implementation

Hamidreza Akef 

Department of Electrical and Computer Engineering, The University of Texas at Austin, TX, USA

Chia-Yu Sung 

Department of Molecular Biosciences, The University of Texas at Austin, TX, USA

Aneesh Vanguri 

Department of Molecular Biosciences, The University of Texas at Austin, TX, USA

David Soloveichik 

Department of Electrical and Computer Engineering, The University of Texas at Austin, TX, USA

Abstract

Amplifying weak molecular signals is essential in both natural and engineered biochemical systems. While most amplification schemes operate out of equilibrium, relying on kinetic barriers and fuel-driven cascades, it is also possible to amplify at thermodynamic equilibrium by shifting the energy landscape upon addition of an analyte. Equilibrium amplification is appealing because, in principle, the system can remain indefinitely in the untriggered state. In this work, we establish fundamental structural and thermodynamic limits on equilibrium-based amplification. We first prove that dimerization networks – systems restricted to complexes of at most two monomers – are inherently incapable of equilibrium amplification. This no-go theorem explains the absence of amplification in prior undercomplementary “strand commutation” designs. We then show that allowing trimeric complexes breaks this barrier. We propose an isometric trimer-based equilibrium amplifier whose output preserves the size of the input, enabling modular composition, and validate it experimentally, achieving an amplification factor close to the expected $2\times$. Finally, we derive universal thermodynamic bounds applicable regardless of complex size: the maximum amplification factor scales linearly with the free energy of interaction between the analyte and the amplifier components. For nucleic acid systems, this implies that the analyte length must grow linearly with the desired amplification factor, and that composing modular amplifiers yields diminishing returns for a fixed analyte. Together, these results delineate the structural and energetic boundaries of equilibrium amplification and rigorously justify the necessity of out-of-equilibrium approaches for achieving high gain.

2012 ACM Subject Classification Applied computing → Systems biology; Computer systems organization → Molecular computing; Mathematics of computing → Mathematical analysis

Keywords and phrases Equilibrium amplification, Dimerization networks, Strand commutation, Thermodynamics of amplification

Digital Object Identifier 10.4230/LIPIcs.DNA.32.8

Related Version *Full Version:* <https://arxiv.org/abs/2604.04285> [1]

Funding Schmidt Sciences Polymath Award to D.S., DOE grant DE-SC0024248, NSF SemiSynBio III: GOALI grant 2227578.

Acknowledgements We thank Antoni Szegłowski for performing experiments and Marko Vasić for sequence design on an earlier variation of the trimeric equilibrium amplifier.



© Hamidreza Akef, Chia-Yu Sung, Aneesh Vanguri, and David Soloveichik;
licensed under Creative Commons License CC-BY 4.0

32nd International Conference on DNA Computing and Molecular Programming (DNA 32).

Editors: Dominic Scalise and Robert Schweller; Article No. 8; pp. 8:1–8:24

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 Introduction

In both natural and engineered systems, amplifying weak molecular signals is crucial for reliable function. For example, signal transduction cascades employ multiple mechanisms – including second messengers, protein phosphorylation, and enzyme cascades – to convert low-level extracellular inputs into robust intracellular responses [25]. Similarly, regulatory RNAs act as ultrasensitive switches to tightly govern gene expression [11]. This broad capacity for signal amplification is essential for maintaining cellular specificity and responsiveness across diverse biological contexts, from hormone signaling to receptor-mediated pathways that drive cellular differentiation and metabolism [2, 13].

Parallel to these natural architectures, amplification in molecular diagnostics and synthetic biology enables the detection of scarce biomarkers and the construction of complex biochemical circuits. Enzyme-free DNA circuits have demonstrated highly sophisticated functions: Seelig *et al.* implemented modular logic gates with signal restoration and amplification [23], and Qian and Winfree scaled strand-displacement cascades into multi-layer logic devices [22]. Likewise, dynamic amplification schemes such as the hybridization chain reaction and catalytic hairpin assembly rely on cascading strand displacements to robustly boost signal levels [12].

Chemical amplifiers typically rely on bridging an energy barrier in the presence of the analyte and thus necessarily operate out of equilibrium. For example, most DNA-based amplifiers rely on fuel-driven kinetic cascades. In the hybridization chain reaction, two DNA hairpins remain kinetically trapped until an initiator strand opens one, setting off a chain of strand-displacement reactions [5]. Likewise, catalytic hairpin assembly uses a series of metastable hairpin fuels: analyte binding to one hairpin exposes a toehold that catalytically opens the next, releasing the analyte to repeat the cycle and yielding hundreds-fold signal amplification [10].

Nonetheless, it is possible to perform signal amplification at thermodynamic equilibrium. In such systems, prior to the addition of the analyte, the thermodynamic equilibrium has a relatively small amount of output (detection species). The addition of the analyte shifts the energy landscape, and subsequent relaxation to the new equilibrium increases the output concentration by more than the amount of analyte added. An important potential advantage of equilibrium amplifiers is that they do not rely on a kinetic barrier, and thus can indefinitely remain in the untriggered state. Equilibrium amplifiers also necessarily avoid complex enzymes typical of natural signaling cascades since the only role an enzyme can play is to modify a kinetic barrier. Equilibrium amplifiers also support dynamic signal processing: changing the input (e.g., swapping input strands) causes the system to re-equilibrate accordingly [24], whereas kinetically controlled systems are generally single-use, requiring their metastable complexes to be externally regenerated for each new input. A family of equilibrium amplifiers was recently proposed in the Thermodynamic Binding Network (TBN) model [6, 20].

Equilibrium-based chemical computation recently received a significant boost with the work of Nikitin on “strand commutation” networks, who showed that diverse computations can be performed by weak hybridization of DNA strands [15]. Specifically, the inputs are DNA strands and the computation arises via the competition of multiple possible binding partners between pairs of strands. The thermodynamic equilibrium of the input together with the rest of the system then contains the desired amount of output strands. Undercomplementarity appears to allow significantly more compact constructions. For example, Nikitin contrasts prior work on the square root circuit using 130 strands [22] with his undercomplementary implementation using only nine strands (plus four inputs). Nonetheless, the computations Nikitin demonstrated suffer from significant signal attenuation, especially in the more complex constructions, which exhibit output concentrations many orders of magnitude lower than

the input concentration (e.g. analog computation). In none of the proposed systems in [15] is the amount of output greater than the amount of input, suggesting that some aspect of Nikitin’s paradigm prevents amplification.

In the current paper, we study amplification at thermodynamic equilibrium, using both theoretical arguments and experimental realization. Our theoretical results are general and apply to DNA-based systems as well as other substrates. In Section 2, we consider the natural classification of equilibrium systems according to maximum complex size, and show that dimerization networks (maximum complex size two) cannot amplify. This result explains the lack of amplification in Nikitin’s work despite the otherwise extensive functional diversity of dimerization networks.

In contrast to dimerization networks, allowing trimers permits equilibrium amplification. In Section 3, we develop a novel undercomplementary trimeric $2\times$ DNA amplifier, and demonstrate that it is able to achieve close to ideal amplification in wet-lab experiments. In contrast to the previously proposed equilibrium amplifier in the TBN model [20], our undercomplementary implementation is substantially more compact in terms of the number of strands overall, the size of the largest complexes, and the number of logical domains. Similar to the TBN amplifier, our amplifier is “composable”: the output strand has the same form as the input strand, and could act as an input to a downstream amplifier of the same design. Composability suggests a method to achieve overall increased amplification.

In the final part of this paper (Section 4), rather than focusing on structural constraints that limit equilibrium amplification, we derive bounds on the amplification factor based on the free energy of interaction between the analyte and the rest of the system. In the context of nucleic acid systems, our bound implies that the length of the detected analyte must scale linearly with the desired equilibrium amplification factor. These results formalize the intuition that a larger difference between the triggered and untriggered configurations – necessary for a greater degree of amplification – necessarily incurs an energetic cost, and we show that this cost scales linearly with the amplification factor. Since the analyte is the only strand added to the system, and all other strands remain unchanged, with only the complexes they form being altered, the requisite energy must originate from the direct interaction of the analyte with the other amplifier strands. For modular equilibrium amplifiers such as our trimeric system and the TBN amplifier [20], the impossibility result implies that the composition of more amplifier modules yields diminishing returns for a fixed analyte.

Thermodynamic limits on biochemical response have been studied most extensively for nonequilibrium sensitivity and ultrasensitivity. For example, Qian showed that the free energy from ATP/GTP hydrolysis controls the achievable ultrasensitivity in phosphorylation-dephosphorylation switches [21], and subsequent work derived universal thermodynamic response bounds for nonequilibrium steady states and tight energy bounds for covalent-modification switches [16, 18]. A related line of work bounds Hill-type sensitivity by structural quantities such as the number of binding sites or the support of the perturbation [17]. These results are complementary to the question addressed here, but do not directly address equilibrium amplification as we define it.

Our results establish fundamental principles of equilibrium-based amplification based on structural network topology (i.e., dimerization versus trimerization), as well as thermodynamic constraints. Our trimeric equilibrium amplifier could potentially be the missing piece for Nikitin-style undercomplementary schemes, allowing computation with less signal loss than is possible in dimerization-only equilibrium networks. More generally, however, our thermodynamic impossibility results rigorously justify out-of-equilibrium amplifiers as unavoidable for greater amplification in the next-generation ultrasensitive molecular sensors and signal processing architectures.


2 Structural Limitations: Dimerization Networks Cannot Amplify

In this section, we analyze the amplification capacity of a chemical reaction network limited to dimerization reactions (i.e., complexes of at most two monomers). By amplification, we mean that an increase in the input monomer leads to a larger-than-proportional increase in the concentration of some monomer species or dimer species in the network. We rigorously demonstrate that dimerization networks are structurally incapable of signal amplification, regardless of the thermodynamic parameters. Specifically, we prove that the sensitivity of any monomer or dimer species in the network to an input signal is bounded by unity.

While previous work [14] used numerical simulations to argue that dimerization networks cannot amplify, rigorous systems-level understanding of these architectures remains elusive [19]. This gap persists even as researchers increasingly push the boundaries of these architectures, such as utilizing competitive dimerization networks as analog physical computers trained via directed evolution [26].

2.1 Model Definition and Equilibrium

Consider a system comprising n distinct monomeric species $\{X_1, \dots, X_n\}$. These monomers can interact to form dimeric species D_{ij} according to the reversible reactions:



with the convention $D_{ji} = D_{ij}$. The system is assumed to be in thermodynamic equilibrium. The equilibrium concentrations, denoted by lowercase variables x_i and d_{ij} , satisfy the mass-action law:

$$d_{ij} = K_{ij} x_i x_j, \tag{1}$$

where $K_{ij} \geq 0$ is the equilibrium association constant. The case $K_{ij} = 0$ represents an absent dimer, while $K_{ij} > 0$ corresponds to an allowed dimerization reaction with finite formation free energy.

We define the total concentration of monomer X_i , denoted by C_i , as the sum of the free monomer concentration and the concentration of X_i sequestered in dimers. According to the conservation of mass:

$$C_i = x_i + \sum_{j \neq i} d_{ij} + 2d_{ii}. \tag{2}$$

2.2 Sensitivity Analysis

We examine the system's equilibrium response to adding total concentration σ of the input monomer X_1 . Specifically, we aim to derive an equation for the relative sensitivity of each monomer species, defined as the derivative of its equilibrium concentration with respect to the added input amount, normalized by its concentration. Formally, we say the *sensitivity* of species X_i is $x'_i = \frac{dx_i}{d\sigma}$, and its *relative sensitivity* is $\frac{x'_i}{x_i}$.

For added input amount σ , the total concentration of monomer X_i is $C_i = C_i^{\text{init}} + \sigma \delta_{1i}$, where δ_{1i} is the Kronecker delta. Substituting Equation (1) into Equation (2), we obtain the closed algebraic mass-balance equations:

$$x_i + \sum_{j \neq i} K_{ij} x_i x_j + 2K_{ii} x_i^2 = C_i^{\text{init}} + \sigma \delta_{1i}.$$

To determine the sensitivity of the system, we differentiate with respect to the signal σ to obtain:

$$x'_i + \sum_{j \neq i} K_{ij}(x'_i x_j + x_i x'_j) + 4K_{ii} x_i x'_i = \delta_{1i}.$$

Using the equilibrium relationship $K_{ij}x_i x_j = d_{ij}$, we can rewrite $K_{ij}x_j = d_{ij}/x_i$ and $K_{ij}x_i = d_{ij}/x_j$ to relate the relative sensitivities to concentrations of dimeric species:

$$\frac{x'_i}{x_i} \left(x_i + \sum_{j \neq i} d_{ij} + 4d_{ii} \right) + \sum_{j \neq i} d_{ij} \frac{x'_j}{x_j} = \delta_{1i}. \quad (3)$$

2.3 The Upper Bound on Amplification

The following theorem shows that no monomer species in the dimerization network can exhibit a relative sensitivity magnitude greater than that of the input monomer, and that the input monomer species itself cannot be amplified beyond the added signal. The theorem is proven in Appendix A.

► **Theorem 1** (The Input Domination Theorem). *Let σ represent the total added concentration of input monomer X_1 . In a dimerization network at equilibrium, the sensitivity of the free concentration x_1 to changes in the total input is bounded by 1, i.e., $0 < \frac{dx_1}{d\sigma} \leq 1$. Furthermore, for any monomer species X_j , the relative sensitivity satisfies $|\frac{x'_j}{x_j}| \leq |\frac{x'_1}{x_1}|$.*

2.4 System-Wide Attenuation of Equilibrium Responses

Having established that $|\frac{x'_1}{x_1}|$ is the global maximum for relative sensitivity, we now analyze the network-wide response. At the fixed equilibrium under consideration, keep the input species indexed as 1 and re-index the remaining species, breaking ties arbitrarily, so that the relative sensitivities are ordered by magnitude:

$$\left| \frac{x'_1}{x_1} \right| \geq \left| \frac{x'_2}{x_2} \right| \geq \dots \geq \left| \frac{x'_n}{x_n} \right|. \quad (4)$$

Now, let $d'_{ij} = \frac{dd_{ij}}{d\sigma}$ denote the sensitivity of the dimer D_{ij} .

► **Lemma 2** (Sign Coherence). *Given the ordering in Equation (4), for any dimer species D_{ij} with $i < j$, the sign of the sensitivity d'_{ij} is determined by the monomer species with the lower index (higher sensitivity). Specifically, if $d'_{ij} \neq 0$, then $\text{sgn}(d'_{ij}) = \text{sgn}(x'_i)$.*

Proof. From Equation (1), $d'_{ij} = d_{ij}(\frac{x'_i}{x_i} + \frac{x'_j}{x_j})$. Since $i < j$, our ordering implies $|\frac{x'_i}{x_i}| \geq |\frac{x'_j}{x_j}|$. Thus, the first term weakly dominates the sum, and the sign of d'_{ij} follows the sign of x'_i . In the symmetric case where $\frac{x'_i}{x_i} = -\frac{x'_j}{x_j}$, the terms cancel, yielding $d'_{ij} = 0$. ◀

We can now rewrite the derivative of the mass conservation equation (Equation (2)) for species i by splitting the sum into terms with indices $j > i$ and $j < i$:

$$x'_i + \sum_{j > i} d'_{ij} + 2d'_{ii} = \delta_{1i} - \sum_{j < i} d'_{ji}. \quad (5)$$

By Lemma 2, for all $j > i$, d'_{ij} has the same sign as x'_i or is zero. Similarly, d'_{ii} clearly has the same sign as x'_i . Therefore, all non-zero terms on the left-hand side of Equation (5) share the same sign. This allows us to take the absolute value of the sum as the sum of absolute values:

$$|x'_i| + \sum_{j>i} |d'_{ij}| + 2|d'_{ii}| = \left| \delta_{1i} - \sum_{j<i} d'_{ji} \right| \leq \delta_{1i} + \sum_{j<i} |d'_{ji}|. \quad (6)$$

This inequality bounds the changes in monomer species i and its dimers with higher-indexed monomers ($j > i$) by the changes in its dimers with lower-indexed monomers ($j < i$). We now generalize this to show system-wide attenuation.

► **Theorem 3** (No-Go Theorem for Dimerization Amplification). *For any monomer species X_i in the ordered dimerization network, the sensitivities of X_i and its dimers D_{ij} with $j \geq i$ satisfy:*

$$|x'_i| + \sum_{j>i} |d'_{ij}| + 2|d'_{ii}| \leq 1 - \sum_{j<i} \left(|x'_j| + 2|d'_{jj}| + \sum_{k>j} |d'_{jk}| \right). \quad (7)$$

Here, the left-hand side represents the total concentration sensitivity directly involving monomer i , including its free form and its dimers with higher-indexed species ($j \geq i$). The right-hand side shows that this quantity is bounded by the remaining portion of the unit input sensitivity after the lower-indexed terms are accounted for; the proof of Theorem 3 is given in Appendix A. The main takeaway of the theorem is that the sensitivities of both monomer species and dimers are individually bounded: $|x'_i| \leq 1$ and $|d'_{ij}| \leq 1$ for all i, j .

► **Corollary 4** (Total Variation Bound). *The sum of the absolute sensitivities of the monomer species is bounded by unity; homodimer formation strictly tightens this bound.*

$$\sum_{i=1}^n |x'_i| \leq 1 - 2 \sum_{i=1}^n |d'_{ii}|. \quad (8)$$

Proof. This follows directly by substituting $i = n$ into Equation (7). ◀


This corollary can be thought to strengthen Theorem 1 in two ways. First, Theorem 1 bounds the relative sensitivities of non-input monomers, but a small relative sensitivity could still correspond to a large absolute concentration change for an abundant species. The total variation bound rules this out. Second, even if no single monomer changes by more than the input, one might imagine a large distributed response spread across many monomer species. The corollary rules this out as well.

3 Breaking the Limit: Trimerization-Based Amplification

In Section 2, we established a fundamental structural constraint: reaction networks limited to dimerization cannot amplify an input signal, regardless of the underlying thermodynamic parameters. In this section, we show that by relaxing the constraint to allow for the formation of higher-order complexes (specifically, only trimers are needed), signal amplification becomes realizable. We propose concrete trimerization-based architectures, and present wet-lab experimental data showing amplification with a trimeric amplifier design.


3.1 Entropy-Driven Trimer-Based Amplifier

A simple construction for a trimer-based amplifier is the entropy-driven design illustrated in Figure 1. This system leverages the thermodynamic drive toward greater global entropy (in the sense of the number of separate complexes) following the introduction of the input X_0 .



■ **Figure 1** Reaction schematic of the entropy-driven amplifier.

Intuitively, the system implements the following abstract reaction:



Such idealized reactions, when composed, implement amplification exponential in the number of modules: one copy of upstream input X_i can release two copies of downstream input X_{i+1} , doubling at every module.¹ Thus idealized, the overall stoichiometric transformation from initial reactants to final products is given by

$$X_0 + \sum_{i=0}^{n-1} 2^i F_i \rightarrow \sum_{i=0}^{n-1} 2^i W_i + 2^n X_n$$

and the theoretical (maximum) amplification factor of this network is $\beta = 2^n$.

There are two fundamental problems with this idealized analysis that limit amplification. The first is the obvious structural limitation: If the system is implemented with DNA, then, as described, each signal X_i is half the length of the previous one.² The second, deeper limitation is based on general thermodynamic considerations, follows from Section 4, and will be explored later.

In the following section, we overcome the structural limitation by introducing an *isometric* amplifier – a design that preserves the size of the output relative to the input.


3.2 Isometric Trimer-Based Amplifier

To overcome the obvious scaling limitation of the above entropy-driven cascades, we introduce a minimal – still with maximum complex size three – network designed to amplify an input signal while strictly preserving the size of the output. This *isometric* amplification relies on a carefully engineered partial ordering of binding affinities among the system’s underlying complexes.

Figure 2 illustrates the proposed reaction pathway. Unlike the entropy-driven amplification scheme above, where the output species is smaller than the input, our isometric design generates an output molecule O of comparable size to the input I . This size preservation is a critical feature for building deeper, multi-stage cascading amplifiers so that the output of

¹ If higher-order complexes (larger than trimers) are allowed, such complexes can form by exchanging sequestered signals X_j among different fuel molecules F_i . As an example, $F_i + F_{i+1} \rightleftharpoons Q + W_{i+1}$, where Q is formed of one copy of X_i^* , one copy of X_{i+1} , and two copies of X_{i+2} . However, in the absence of input, the system still remains trapped in an equilibrium distribution without the final product X_n .

² Because each step in the cascade effectively “divides” the original input to generate multiple outputs, the amplification factor is limited by the total length of the input in base pairs $|X_0|$ and the minimum domain size κ – the minimum strand length necessary for reliable molecular recognition and binding: $\beta \leq \frac{|X_0|}{\kappa}$.



■ **Figure 2** Schematic of the proposed isometric trimerization-based amplifier. **(Left)** The thermodynamically favored configuration in the absence of the input signal I. Strand W exhibits a higher binding affinity for X than for P. Consequently, the output strands O remain sequestered in a stable trimeric complex with the “locking” molecule P, forming O:O:P. **(Right)** The thermodynamically favored configuration following the introduction of the input. Input I cooperatively binds with “helper” species A and X to form the I:A:X complex, displacing W from the W:X complex. The newly liberated strand W then drives the dissociation of the O:O:P complex by binding to P (forming W:P), thereby releasing the free output O. The forward directionality of this cascade is governed by the thermodynamic hierarchy of the complex affinities: $K_{I:A:X} > K_{W:X} > K_{W:P} > K_{O:O:P}$. Crucially, in the absence of I, species A possesses insufficient affinity to displace X or P. *Note: An increased density of white gaps within the gray binding regions visually denotes progressively weaker binding affinities.*

one stage can serve as an input for the subsequent stage. Thermodynamically, the system is designed such that enough energy is released upon the formation of the input complex I:A:X to effectively “pay” the energetic penalty required to dissociate the output complex O:O:P.

In order to avoid leak (release of O without I), we must limit the affinities in the affinity hierarchy noted in the figure caption. Specifically, without input I, molecule A can displace half of the W:X complex. Then this free half of W may displace an output molecule O from the O:O:P complex, i.e., $A + W:X + O:O:P \rightarrow A:W:X:O:P + O$.³ This reaction pathway is inhibited because it reduces the number of separate complexes by one, but we need to ensure that the increased bonding does not compensate for the loss of entropy.

By strategically engineering sequence mismatches, we can tune the free energies of the complexes, establishing the required partial order of binding affinities and their relationship to separate complex formation entropy. Furthermore, because of partial complementarity, the sequence of O can be sufficiently different from I that several identical isometric amplification modules could be composed, where the output of one is input to another.

Experimental Results

To validate the theoretical framework of the isometric trimer-based amplifier, we implemented an experimental model using synthetic DNA. First, we verified the theoretical performance of our designed sequences using NUPACK [8]. As observed in Figure 3a, while the simulated system exhibited some leakage (release of output in the absence of input) at a higher temperature (28°C), this leak was effectively eliminated by decreasing the temperature. Overall, the NUPACK-simulated response closely follows the ideal 2× slope until reaching saturation.

³ While our theoretical model restricts complex size to three, mitigating the formation of higher-multimericity off-model polymers remains a practical engineering goal.




Figure 3 Experimental results of the isometric amplifier. The input concentration is varied from 0 to 200 nM, the total concentration of the output strand is 300 nM, and all other circuit strands are set at 150 nM. (a) Experimental data plotted against NUPACK estimated behavior. (b) Experimental data plotted against the ideal 2× response. Linear regression of the response regime yields an experimental amplification factor of 1.7×.

The experimental data shown in Figure 3(a)-(b) was collected at 4°C, but in contrast to NUPACK predictions, it still exhibits non-negligible leakage in the absence of the input (0 nM). This minor discrepancy may stem from several factors: imprecise stoichiometry of the added components, DNA synthesis impurities (as only the fluorophore-tagged output strand was HPLC-purified, while the rest were purified via standard PAGE), or fluorophore-induced energetic perturbations not captured by NUPACK.

As shown in Figure 3b, to quantify the experimental amplification factor, we performed linear regression. While the theoretical maximum amplification factor of this specific system is 2×, our experimental data demonstrates 1.7× signal amplification. The regression analysis excluded the saturated data points (input concentrations of 150 nM and 200 nM) to capture the true linear amplification regime. (If the leakage point is also excluded from the fit to strictly evaluate the active response slope, the calculated amplification factor rises to 1.77×.) Complete details of the experimental protocols, DNA sequences, and statistical methods are provided in the full version of this paper [1].

Since there is no physical limit imposed by shrinking molecule sizes as in the naive construction of Section 3.1, one might intuitively expect that arbitrary signal amplification is physically realizable. However, escaping structural limitations does not grant unlimited amplification. In the following section, we establish that regardless of a system’s specific structural architecture or local mechanisms, there exists a fundamental, inescapable thermodynamic limit on signal amplification for *all* thermodynamic equilibrium-based networks.

4 Thermodynamic Bounds on Amplification

In the previous sections, we demonstrated that fundamental structural differences (dimerization versus trimerization) can disallow or permit amplification. We have also seen that the trimeric amplifier could in principle be composed to seemingly increase its amplification factor. In this section, we turn to general and unavoidable limits on amplification at thermodynamic equilibrium determined by thermodynamic considerations, regardless of

structural assumptions. Unlike prior thermodynamic bounds on nonequilibrium biochemical response [21, 16, 18], the systems considered here have no persistent chemical-potential driving: both the pre-input and post-input configurations are equilibria. The relevant resource is therefore the binding free energy gained when the added input analyte interacts with amplifier components.

There is a strong intuition that chemical amplification should require significant energy expenditure since the system must traverse substantial “configurational distance” from a state with little output to one with a lot of output. Since we are requiring the system before and after the addition of input to be at thermodynamic equilibrium, the only source of this energy must be from the input itself. This section formalizes this reasoning, showing that the free energy with which the input binds amplifier components must scale linearly with the amplification factor. In the context of DNA nanotechnology, this implies that for an analyte (DNA molecule) of a fixed length, there is only so much amplification that can be realized. Moreover, we show that the achievable amplification is further discounted by the *output contrast*, defined as the ratio of the triggered output (signal) to the untriggered output (leak). For a fixed input free energy, demanding a sharper contrast leaves less room for amplification.

4.1 Model

We represent the monomer composition of any complex as a vector tuple $(u, v) \in \mathbb{N}^m \times \mathbb{N}$, where u is the vector of stoichiometric counts of non-input monomers, and v is the count of the input monomer I . Given a complex (u, v) , we define its *residual core* to be just its non-input monomers. The L_1 -norm $|u| = \sum_{i=1}^m u_i$ of (u, v) is the size of its residual core. Similarly, we define the *input core* of (u, v) as its constituent input monomers, with its size naturally given by v .

The concentration of a specific complex is denoted $x_{u,v}$. This vector representation naturally partitions the network into three distinct pools: pure residual complexes ($v = 0, u \neq \mathbf{0}$), pure input complexes ($u = \mathbf{0}, v \geq 1$), and mixed complexes ($u \neq \mathbf{0}, v \geq 1$).⁴ All concentrations are dimensionless mole fractions – the ratio of complexes to solvent molecules (i.e., water) – and are thus smaller than one. (The derivation of the energy function (Equation (9)) only holds in the regime of substantially more dilute complexes than solvent [4].)

To model the equilibrium behavior of these systems, we adopt the free-energy formulation utilized by Dirks *et al.* [4]. Let $g(x)$ denote the dimensionless free energy of the overall system configuration (formula given below), let $G_{u,v}$ represent the dimensionless free energy of a specific complex, and let $\mathcal{P}$ denote the set of all valid complexes in the network. We define the sign convention for $G_{u,v}$ such that more negative values correspond to more thermodynamically favorable complexes.

Let σ represent the total concentration of the input monomer I . The equilibrium concentrations are obtained by minimizing the system’s total free energy (corresponding to the pseudo-Helmholtz free energy utilized extensively throughout the chemical reaction network theory literature [9, 7]):⁵

$$g(x) = \sum_{(u,v) \in \mathcal{P}} x_{u,v} (\log x_{u,v} - 1 + G_{u,v}) \quad (9)$$

⁴ There is no complex ($u = \mathbf{0}, v = 0$).

⁵ In this work, $\log$ denotes the natural logarithm.

This minimization is subject to the mass conservation constraint $Ax = x^0(\sigma)$, where x is the vector of complex concentrations, $x^0(\sigma)$ is the vector representing the total concentrations of all monomers (including the amount σ of added input I), and A is the mass conservation matrix mapping the concentrations of complexes to their constituent monomers.

To formally manipulate the system configurations, we define two projection operators that isolate the cores from the overall complex concentration vector x .

The **Input-Core operator**, denoted $IC(x)$, extracts the input cores by stripping away all residual components. In other words, it removes all residual monomers, and for every $v \geq 1$, creates a complex containing v input monomers whose concentration is the sum of the concentrations of all complexes containing exactly v input monomers in x . Formally, we define the resulting vector $w = IC(x)$:

$$w = IC(x) \implies \begin{cases} w_{(0,v)} = \sum_u x_{(u,v)} \\ w_{(u,v)} = 0 \end{cases} \quad \text{for } u \neq 0 \quad (10)$$

Conversely, the **Residual-Core operator**, denoted $RC(x)$, extracts the residual cores by stripping away all input monomers. Analogously to Equation (10), we define the resulting vector $w = RC(x)$:

$$w = RC(x) \implies \begin{cases} w_{(u,0)} = \sum_v x_{(u,v)} \\ w_{(u,v)} = 0 \end{cases} \quad \text{for } v \geq 1 \quad (11)$$

4.2 Energy Change upon Input Addition

We now investigate the total change in system free energy upon the introduction of the input. To quantify this transition, we compare the thermodynamic equilibria of the network before and after the input is introduced – configurations we formalize as the *Resting* and *Active* states, respectively. To establish a rigorous mathematical bound on this free-energy difference, we will also define a specific configuration termed the *Drained* state.

The Active state is the thermodynamic equilibrium of the entire system with concentration σ of the input I . Formally:

► **Definition 5.** For any input concentration $\sigma \geq 0$, the **Active state** $\mathbf{x}$ is the concentration vector that minimizes the total system free energy $g(x)$ subject to mass conservation:

$$\begin{aligned} \mathbf{x} = \operatorname{argmin}_{x \geq 0} \quad & g(x) \\ \text{s.t.} \quad & Ax = x^0(\sigma) \end{aligned} \quad (12)$$

Next we define the Drained state, which represents “ripping” the input out of the Active state without allowing the system to relax to a new equilibrium. Formally, the Drained state is constructed by separating the input cores and residual cores of the Active state:

► **Definition 6.** *Drained state* $\mathbf{y}$ is defined by:

$$\mathbf{y} = IC(\mathbf{x}) + RC(\mathbf{x}) \quad (13)$$

Note that the dissociation into input cores and residual cores preserves mass conservation $A\mathbf{y} = x^0(\sigma)$.

Finally, we define the Resting state, which represents the thermodynamic equilibrium without input to which the input is added without being able to interact with the rest of the system. More precisely, the Resting state is the thermodynamic equilibrium of the system without the input, plus $IC(\mathbf{x})$. Formally:

► **Definition 7.** *Resting state $\mathbf{z}$ is defined by:*

$$\begin{aligned} \mathbf{z} &= IC(\mathbf{x}) + \operatorname{argmin}_{x \geq 0} g(x) \\ \text{s.t. } Ax &= x^0(0) \end{aligned} \quad (14)$$

► **Lemma 8 (Free Energy Hierarchy).** *The free energies of the Active, Resting, and Drained states satisfy the following inequality:*

$$g(\mathbf{x}) \leq g(\mathbf{z}) \leq g(\mathbf{y})$$

The proof of Lemma 8 is given in Appendix B.

We now define the following quantities with respect to the Active state $\mathbf{x}$. Let $\mathcal{U} = \{u \neq \mathbf{0} \mid \exists v \geq 1: \mathbf{x}_{u,v} > 0\}$ denote the set of residual stoichiometries u that can form complexes with the input in the Active state. Let $X_{\text{inp}} = \sum_{v \geq 1} \mathbf{x}_{\mathbf{0},v}$ be the total concentration of the pure input complexes, $X_{\text{res}} = \sum_{u \in \mathcal{U}} \mathbf{x}_{u,0}$ be the total concentration of the pure residual complexes, and let $X_{\text{mix}} = \sum_{u \in \mathcal{U}} \sum_{v \geq 1} \mathbf{x}_{u,v}$ be the total concentration of the mixed complexes.

Next, we define the binding free energy, $\Delta G_{u,v}^{\text{bind}} := G_{u,v} - G_{u,0} - G_{\mathbf{0},v}$, which quantifies the thermodynamic cost of forming the complex (u, v) from its isolated residual and input cores. Finally, we let $\Delta G_{\text{stg}}^{\text{inp}} = \min_{u,v} \Delta G_{u,v}^{\text{bind}}$ denote the most favorable binding energy achievable within the network upon interaction with the input. Note that we use the convention that $\Delta G_{\text{stg}}^{\text{inp}} < 0$, and the more negative it is, the more favorable the interaction.

Having formally defined the Active, Drained, and Resting states, our primary objective is to quantify the total change in the pseudo-Helmholtz free energy of the system upon the addition of an input species with concentration σ . The Drained state serves as a crucial analytical intermediate. Starting from the Active state, where the input has already propagated through the network to trigger the targeted outputs, we conceptually decouple all bound input cores from the residual network. If we then allow the residual cores of this Drained state to thermodynamically relax while remaining strictly isolated from the input pool, the network equilibrates into the Resting state.

Altogether, by comparing the Active, Drained, and Resting states, and incorporating the strongest potential energy released by input binding ($\Delta G_{\text{stg}}^{\text{inp}}$), we can rigorously bound the total pseudo-Helmholtz free energy change of the system as shown below. We show that the bound depends on the following dimensionless parameter, which we term amplification capacity. Intuitively, the upper bound on amplification capacity is the maximum energy with which the input binds ($-\Delta G_{\text{stg}}^{\text{inp}}$). From this upper bound, we subtract the log of “concentration of amplifier components” (as captured by $X_{\text{res}} + X_{\text{mix}}$) – with smaller concentrations leading to smaller amplification capacity.⁶

► **Definition 9.** *We define the amplification potential as $\psi(X, \Delta G) := -\Delta G + \log(X)$. Hence, the input amplification potential is denoted by $\psi_{\text{inp}} := \psi(X_{\text{res}} + X_{\text{mix}}, \Delta G_{\text{stg}}^{\text{inp}})$. We then define the amplification capacity ϕ as follows:*

$$\phi := \begin{cases} e^{\psi_{\text{inp}}} & \text{if } \psi_{\text{inp}} \leq 0, \\ \psi_{\text{inp}} + 1 & \text{if } \psi_{\text{inp}} > 0. \end{cases} \quad (15)$$

⁶ Recall that all concentrations are strictly less than one, as they are expressed as dimensionless mole fractions (the ratio of complexes to solvent molecules). Consequently, $X_{\text{res}} + X_{\text{mix}} < 1$.

Below we will derive upper bounds on the amplification factor as a function of the amplification capacity, justifying this terminology.

► **Theorem 10** (Input-Driven Energy Change). *Let σ be the total added concentration of the input species. Let $g(\mathbf{y})$ denote the free energy of the Drained state, let $g(\mathbf{x})$ denote the free energy of the Active state. Then, the amplification capacity ϕ satisfies:*

$$g(\mathbf{y}) - g(\mathbf{x}) \leq \sigma \phi \quad (16)$$

The proof of Theorem 10 is given in Appendix B.

► **Corollary 11.** *A consequence of the free energy hierarchy (Lemma 8) is that the intermediate free energy transitions are likewise bounded. Specifically, because $g(\mathbf{x}) \leq g(\mathbf{z}) \leq g(\mathbf{y})$, it holds that both $g(\mathbf{z}) - g(\mathbf{x}) \leq \sigma \phi$ and $g(\mathbf{y}) - g(\mathbf{z}) \leq \sigma \phi$.*

To connect these thermodynamic constraints to system performance, we now translate the bounded free energy transitions into concentrations, yielding an explicit upper bound on the amplification factor.

4.3 Amplification Factor Upper Bound

Fix an input amount σ . Intuitively, we want the output amount to be much larger than the leak, which is the output amount in the absence of input. For actual amplification, we also want the output amount to be much larger than the input amount σ . The bound on energy release as a function of amplification capacity reported in Theorem 10 can be recast into an explicit relationship between these quantities, properly formalized.

Let O denote the designated free-output species, and define its concentration at the Active and Resting states as:

$$a := \mathbf{x}_O, \quad r := \mathbf{z}_O$$

In other words, a is the signal amount of output when the amplifier is activated, and r is the leak amount. We define the amplification factor β and the output contrast γ as:

$$\beta := \frac{a - r}{\sigma}, \quad \gamma := \frac{a}{r}.$$

Thus the amplification factor β represents the amount of output increase between leak and signal, relative to the input amount. Output contrast γ represents the ratio of the free output species between leak and signal states.

► **Theorem 12** (Thermodynamic Bound on Amplification). *Under the natural assumption for amplification that the free-output concentration in the Active state exceeds its Resting leak ($\gamma > 1$), the amplification factor β is upper-bounded by amplification capacity ϕ discounted by the logarithmic output contrast:*

$$\beta \leq \frac{\phi}{\log \gamma}. \quad (17)$$

In the regime relevant for amplification ($\psi_{\text{inp}} > 0$), the amplification capacity $\phi = \psi_{\text{inp}} + 1$ grows linearly with the strength of the input interaction $-\Delta G_{\text{stg}}^{\text{inp}}$. A key takeaway of this theorem is therefore that the input must bind the amplifier components with a free energy that scales (at least) linearly with the desired amplification factor (at fixed contrast).

8:14 Amplification at Equilibrium: Limitations and Implementation

► **Remark 13.** In the positive input potential regime ($\psi_{\text{inp}} > 0$), expanding ϕ in Theorem 12 reveals:

$$\beta \leq \frac{-\Delta G_{\text{stg}}^{\text{inp}} + \log(X_{\text{res}} + X_{\text{mix}}) + 1}{\log \gamma} < \frac{-\Delta G_{\text{stg}}^{\text{inp}} + 1}{\log \gamma}. \quad (18)$$

Proof of Theorem 12. The Resting and Drained states share the same isolated input pool, so their free-energy difference is entirely a residual-subsystem difference. The residual part of $\mathbf{z}$ is the equilibrium minimizer subject to the isolated residual mass constraint. Hence the pseudo-Helmholtz Bregman identity gives [3]:

$$g(\mathbf{y}) - g(\mathbf{z}) = \sum_i \left(\mathbf{y}_i \log \frac{\mathbf{y}_i}{\mathbf{z}_i} - \mathbf{y}_i + \mathbf{z}_i \right),$$

where the sum is over the pure residual complexes. Each summand is nonnegative. Keeping only the output O :

$$g(\mathbf{y}) - g(\mathbf{z}) \geq \mathbf{y}_O \log \frac{\mathbf{y}_O}{\mathbf{z}_O} - \mathbf{y}_O + \mathbf{z}_O.$$

We know the output O is a pure residual species and that draining the input from the Active state may increase its concentration $\mathbf{y}_O \geq \mathbf{x}_O = a$. This holds because a fraction of the total output mass might be bound with input species in the Active state. Letting $r = \mathbf{z}_O$, define

$$F(t) := t \log \frac{t}{r} - t + r.$$

Since $F'(t) = \log(t/r)$, the function F is increasing for $t \geq r$. Applying $\mathbf{y}_O \geq a$ along with $a > r$ yields:

$$g(\mathbf{y}) - g(\mathbf{z}) \geq a \log \frac{a}{r} - a + r. \quad (19)$$

A sharper estimate is obtained by also retaining the output contribution between the Resting and Active states. Since the Active state is the equilibrium minimizer for the full system under the same total mass constraint, the pseudo-Helmholtz Bregman identity similarly gives:

$$g(\mathbf{z}) - g(\mathbf{x}) \geq r \log \frac{r}{a} - r + a. \quad (20)$$

Adding Equations (19) and (20) gives:

$$\begin{aligned} g(\mathbf{y}) - g(\mathbf{x}) &\geq \left(a \log \frac{a}{r} - a + r \right) + \left(r \log \frac{r}{a} - r + a \right) \\ &= (a - r) \log \frac{a}{r}. \end{aligned}$$

Substituting $\gamma = a/r$ and $a - r = \beta\sigma$ gives:

$$g(\mathbf{y}) - g(\mathbf{x}) \geq \beta\sigma \log \gamma.$$

► **Remark 14.** One of the reasons Theorem 12 is not strictly tight is the truncation of terms in Equations (19) and (20). Nonetheless, looking at the entire sum may be appropriate to generate design strategies and effective amplifiers. One such heuristic is apparent: limit the variety of distinct chemical species that do not contain the output, which decreases the number of non-zero terms in the sum.

Combining this lower bound with Theorem 10,

$$g(\mathbf{y}) - g(\mathbf{x}) \leq \sigma\phi,$$

and cancelling σ completes the proof. $\blacktriangleleft$

To illustrate the order of magnitude of these thermodynamic limits, we evaluate the upper bound using the parameters of the isometric trimeric amplifier from Section 3.2. Our input strand has a maximum binding affinity of roughly -49 kcal/mol at the operating temperature of 4°C (277.15 K). Assuming an overestimate of $X_{\text{mix}} + X_{\text{res}} < 1 \mu\text{M}$, we obtain the following dimensionless parameter values⁷:

$$\Delta G_{\text{stg}}^{\text{inp}} \approx -89, \quad X_{\text{res}} + X_{\text{mix}} \lesssim 1.8 \times 10^{-8}.$$

Substituting these values into the definition for the amplification capacity yields:


$$\phi = \psi_{\text{inp}} + 1 = -\Delta G_{\text{stg}}^{\text{inp}} + \log(X_{\text{res}} + X_{\text{mix}}) + 1 \approx 72.$$

Under these conditions, requiring a modest 10-fold target output contrast ($\gamma = 10$) restricts the maximum achievable amplification factor to $\beta \lesssim 31$. This concrete example demonstrates how a seemingly vast input free-energy budget (hybridization energy of -49 kcal/mol is typically considered all but irreversible) is naturally attenuated by the system's concentration scales and signal contrast requirements, leaving a ceiling on the final equilibrium gain.⁸

4.4 Output Contrast Upper Bound

In the previous section, Theorem 12 bounds the amplification factor β as a function of amplification capacity ϕ and the contrast γ . In this section, our results address the following intuition in turn: to have high contrast, the output must be capable of binding strongly with something in order to have low leak. Thus, we turn from the free energy of binding the input to that of binding the output.

In order to link contrast to the underlying thermodynamics of output sequestration in the Resting state, suppose a set of pure residual complexes $\mathcal{Q}$ sequesters the output species in the Resting state. Each complex $Q \in \mathcal{Q}$ forms via the reversible binding of one free output molecule O to a distinct residual partner P_Q ⁹:



We define the sequestration free-energy gain for each complex as $\Delta G_Q^{\text{out}} := G_Q - G_O - G_{P_Q}$. Let $\Delta G_{\text{stg}}^{\text{out}} := \min_{Q \in \mathcal{Q}} \Delta G_Q^{\text{out}}$ denote the most favorable binding energy achievable within the network upon the output sequestration. Note that similar to $\Delta G_{\text{stg}}^{\text{inp}}$, we use the convention that $\Delta G_{\text{stg}}^{\text{out}} < 0$, and the more negative it is, the more favorable the interaction.

By the law of mass action, the Resting equilibrium configuration satisfies:

$$\mathbf{z}_Q = \mathbf{z}_O \mathbf{z}_{P_Q} e^{-\Delta G_Q^{\text{out}}} \quad (21)$$

⁷ Using the molarity of water (55.49 M at 4°C) as the solvent to non-dimensionalize the system, and adopting the ideal gas constant $R = 0.00198717$ kcal/mol.K.

⁸ As discussed in Section 5, the tightness of our results remains open. While we have not demonstrated multi-stage composition and the bound need not be tight, it leaves open the possibility that 4 layers of the isometric trimeric amplifier could be composed, since each contributes a factor of 2 and $2^4 = 16 \leq 31 < 32 = 2^5$.

⁹ If Q includes more than one copy of O , then P_Q includes O as well, which does not cause a problem for the subsequent reasoning steps.

Let $Z_{\text{res}}^{\text{out}} := \sum_{Q \in \mathcal{Q}} \mathbf{z}_{P_Q}$. This quantity serves as the output-side counterpart to X_{res} : whereas X_{res} represents the total concentration of complexes capable of binding the input species in the Active state, $Z_{\text{res}}^{\text{out}}$ represents the total concentration of complexes P_Q capable of sequestering the output species in the Resting state. Furthermore, let ν_{max} be the largest output stoichiometry of any complex Q .

Finally, the output potential ψ_{out} is defined as follows:

$$\psi_{\text{out}} := \psi(\nu_{\text{max}} Z_{\text{res}}^{\text{out}}, \Delta G_{\text{stg}}^{\text{out}}) = -\Delta G_{\text{stg}}^{\text{out}} + \log(\nu_{\text{max}} Z_{\text{res}}^{\text{out}}).$$

► **Lemma 15** (Contrast Upper Bound). *The output contrast γ of an amplifier is bounded by: $\gamma \leq 1 + e^{\psi_{\text{out}}}$.*

Proof. The net increase in the free output concentration, $a - r$, cannot exceed the total amount of sequestered output available in the Resting state. Recalling that the Resting free output concentration is $\mathbf{z}_O = r$, we apply Equation (21) to bound this difference:

$$a - r \leq \nu_{\text{max}} \sum_{Q \in \mathcal{Q}} \mathbf{z}_Q = \nu_{\text{max}} r \sum_{Q \in \mathcal{Q}} \mathbf{z}_{P_Q} e^{-\Delta G_Q^{\text{out}}} \leq r \left(\nu_{\text{max}} Z_{\text{res}}^{\text{out}} e^{-\Delta G_{\text{stg}}^{\text{out}}} \right).$$

Recognizing that the term in the parentheses is exactly equivalent to $e^{\psi_{\text{out}}}$, we obtain $a - r \leq r e^{\psi_{\text{out}}}$. Dividing the entire inequality by r and rearranging for $\gamma := a/r$ completes the proof. ◀

Lemma 15 together with Theorem 12 reveals the following big picture. Because the output contrast γ acts as the denominator in the overall amplification bound (Theorem 12), enforcing a high signal contrast inherently restricts the maximum achievable amplification factor β . Conversely, designing a system with weaker output sequestration (a less negative $\Delta G_{\text{stg}}^{\text{out}}$) mathematically expands the potential for amplification, but leaves the system susceptible to leak and poor contrast. Therefore, to increase the amplification upper bound *without* sacrificing signal clarity, one must instead drive the system with a more favorable input free energy $\Delta G_{\text{stg}}^{\text{inp}}$, thereby raising the overall amplification capacity ϕ .

► **Remark 16.** Lemma 15 implies that to achieve a target output contrast $\gamma > 1$, the output sequestration free-energy must be favorable enough (recall $\Delta G_{\text{stg}}^{\text{out}} < 0$ by convention):

$$\Delta G_{\text{stg}}^{\text{out}} \leq \log \left(\frac{\nu_{\text{max}} Z_{\text{res}}^{\text{out}}}{\gamma - 1} \right). \quad (22)$$

In practice, bounding the leakage is critical to avoid prematurely triggering downstream cascades or violating diagnostic sensitivity thresholds. To guarantee a target contrast γ for these applications, Equation (22) provides a direct design heuristic: the intrinsic sequestration affinity $\Delta G_{\text{stg}}^{\text{out}}$ must be driven sufficiently negative, taking into account the total concentration of available residual partners $Z_{\text{res}}^{\text{out}}$ and their maximum output sequestration size ν_{max} .

To complete our analysis of the isometric amplifier example, we examine the thermodynamic constraints on the output side of the network. In our experiment, the hybridization energy of the output O is -24 kcal/mol under our conditions. While this is substantially less favorable than the hybridization of the input, it is in principle compatible with very large contrast depending on $Z_{\text{res}}^{\text{out}}$. Although $Z_{\text{res}}^{\text{out}}$ is not directly under our control, we could imagine that adding a reasonable excess of strand P results in nanomolar quantities of free P in equilibrium. Then the corresponding dimensionless parameters are evaluated as:

$$\Delta G_{\text{stg}}^{\text{out}} \approx -44, \quad \nu_{\text{max}} = 2, \quad \text{and} \quad Z_{\text{res}}^{\text{out}} \lesssim 1.8 \times 10^{-11}.$$

Substituting these values into the expression for the output amplification potential yields:

$$\psi_{\text{out}} = -\Delta G_{\text{stg}}^{\text{out}} + \log(\nu_{\text{max}} Z_{\text{res}}^{\text{out}}) \approx 19.$$

By Lemma 15, this potential yields a bound on the output contrast of $\gamma \lesssim e^{19}$. While operating near this extreme contrast limit would be highly effective for minimizing leak, it heavily penalizes the system's amplification β by Theorem 12. Alongside an amplification capacity of $\phi \approx 72$ from the previous section, we see that the maximum achievable amplification factor becomes constrained to:

$$\beta \leq \frac{\phi}{\log \gamma} \approx \frac{72}{19} \approx 3.8.$$

5 Conclusion

In this work, we first provided a rigorous proof demonstrating that equilibrium dimerization networks are fundamentally incapable of signal amplification, a problem previously addressed only through numerical approximations due to its inherent algebraic complexity. Our results help explain the lack of signal amplification in undercomplementary designs proposed by Nikitin [15] as they all rely on a maximum complex size of two. We subsequently evaluated the amplification potential of equilibrium trimerization networks. Our investigation of an entropy-driven trimerization network confirmed its capacity to amplify, but revealed a size-shrinking feature that inherently limits its amplification power. To resolve this, we designed an isometric trimer-based equilibrium amplifier that successfully achieved signal amplification while strictly preserving the length of the output strands, potentially allowing for module composition. We validated this architecture through the wet-lab implementation of a $2\times$ amplifier, achieving an experimental amplification factor that closely matches this theoretical target.

We finally broadened our scope to investigate fundamental limits on equilibrium amplification in a more general case based on thermodynamic constraints, regardless of the multimericity of the network. We demonstrated that the free energy released by the input binding determines the theoretical ceiling of the amplification power; the network design then dictates how much of this capacity is attenuated. Specifically, we established that enforcing a high signal contrast fundamentally limits the maximum achievable amplification factor. Furthermore, we proved that the theoretical ceiling of this output contrast is itself dictated by the free energy of output sequestration. Ultimately, these thermodynamic bounds reveal a universal engineering trade-off: securing a robust, high-contrast signal requires deep output sequestration to suppress leak, which in turn demands a proportionally larger input free-energy budget to successfully drive molecular amplification.

There are two natural directions for future work. First, we aim to experimentally compose multiple amplifier modules to achieve greater than a factor-of-two amplification, using our component-level heuristics (e.g., Remarks 14 and 17) to guide the physical design. Second, we must investigate the tightness of our theorems by testing how closely these physical implementations can approach our established thermodynamic limits for amplification (Theorem 12) and contrast (Lemma 15). Ultimately, this leads to a fundamental open question: subject to these thermodynamic limits, does an “optimal” equilibrium amplifier or amplifier family exist?

References

- 1 Hamidreza Akef, Chia-Yu Sung, Aneesh Vanguri, and David Soloveichik. Amplification at equilibrium: Structural and thermodynamic limitations, and implementation, 2026. doi:10.48550/arXiv.2604.04285.
- 2 Hui Chen, Shaohong Zhou, Kleins Ngocho, Jing Zheng, Xiaoxiao He, Jin Huang, Kemin Wang, Hui Shi, and Jianbo Liu. Oriented triplex DNA as a synthetic receptor for transmembrane signal transduction. *Nature Communications*, 15(1):9789, 2024. doi:10.1038/s41467-024-53960-5.
- 3 Imre Csiszár. I-divergence geometry of probability distributions and minimization problems. *The Annals of Probability*, 3(1):146–158, 1975. doi:10.1214/aop/1176996454.
- 4 Robert M Dirks, Justin S Bois, Joseph M Schaeffer, Erik Winfree, and Niles A Pierce. Thermodynamic analysis of interacting nucleic acid strands. *SIAM review*, 49(1):65–88, 2007. doi:10.1137/060651100.
- 5 Robert M. Dirks and Niles A. Pierce. Triggered amplification by hybridization chain reaction. *Proceedings of the National Academy of Sciences*, 101(43):15275–15278, 2004. doi:10.1073/pnas.0407024101.
- 6 David Doty, Trent A. Rogers, David Soloveichik, Chris Thachuk, and Damien Woods. Thermodynamic binding networks. In Robert Brijder and Lulu Qian, editors, *DNA Computing and Molecular Programming*, pages 249–266, Cham, 2017. Springer International Publishing. doi:10.1007/978-3-319-66799-7_16.
- 7 Martin Feinberg. *Foundations of chemical reaction network theory*. Springer, 2019.
- 8 Mark E. Fornace, Jining Huang, Cody T. Newman, Avinash Nanjundiah, Nicholas J. Porubsky, Marshall B. Pierce, and Niles A. Pierce. NUPACK: Analysis and design of nucleic acid structures, devices, and systems. ChemRxiv preprint, 2025. Version 2. doi:10.26434/chemrxiv-2022-xv981-v2.
- 9 Fritz Horn and Roy Jackson. General mass action kinetics. *Archive for rational mechanics and analysis*, 47:81–116, 1972.
- 10 Yu (Sherry) Jiang, Bingling Li, John N. Milligan, Sanchita Bhadra, and Andrew D. Ellington. Real-time detection of isothermal amplification reactions with thermostable catalytic hairpin assembly. *Journal of the American Chemical Society*, 135(20):7430–7433, 2013. PMID: 23647466. doi:10.1021/ja4023978.
- 11 Erel Levine, Zhongge Zhang, Thomas Kuhlman, and Terence Hwa. Quantitative characteristics of gene regulation by small RNA. *PLOS Biology*, 5(9):e229, 2007. doi:10.1371/journal.pbio.0050229.
- 12 Bingling Li, Andrew D. Ellington, and Xi Chen. Rational, modular adaptation of enzyme-free DNA circuits to multiple detection methods. *Nucleic Acids Research*, 39(16):e110–e110, June 2011. doi:10.1093/nar/gkr504.
- 13 Yanhao Liu, Wan Zhao, Yingming Zhao, Xiangxiang Zhang, Jingjing Zhao, Shubin Li, Yongshuo Ren, and Xiaojun Han. An artificial cell capable of signal transduction mediated by adrb2 for the regulation of glycogenolysis. *Nature Communications*, 17(1):1795, 2026. doi:10.1038/s41467-026-68503-3.
- 14 Sergei Maslov and I. Ispolatov. Propagation of large concentration changes in reversible protein-binding networks. *Proceedings of the National Academy of Sciences*, 104(34):13655–13660, 2007. doi:10.1073/pnas.0702905104.
- 15 Maksim P. Nikitin. Non-complementary strand commutation as a fundamental alternative for information processing by dna and gene regulation. *Nature Chemistry*, 15(1):70–82, 2023. doi:10.1038/s41557-022-01111-y.
- 16 Jeremy A. Owen, Todd R. Gingrich, and Jordan M. Horowitz. Universal thermodynamic bounds on nonequilibrium response with biochemical applications. *Physical Review X*, 10(1):011066, 2020. doi:10.1103/PhysRevX.10.011066.
- 17 Jeremy A. Owen and Jordan M. Horowitz. Size limits the sensitivity of kinetic schemes. *Nature Communications*, 14(1):1280, 2023. doi:10.1038/s41467-023-36705-8.

- 18 Jeremy A. Owen, Pranay Talla, John W. Biddle, and Jeremy Gunawardena. Thermodynamic bounds on ultrasensitivity in covalent switching. *Biophysical Journal*, 122(10):1833–1845, 2023. doi:10.1016/j.bpj.2023.04.015.
- 19 Jacob Parres-Gold, Matthew Levine, Benjamin Emert, Andrew Stuart, and Michael B. Elowitz. Contextual computation by competitive protein dimerization networks. *Cell*, 188(7):1984–2002.e17, 2025. doi:10.1016/j.cell.2025.01.036.
- 20 Joshua Petrack, David Soloveichik, and David Doty. Thermodynamically Driven Signal Amplification. In Ho-Lin Chen and Constantine G. Evans, editors, *29th International Conference on DNA Computing and Molecular Programming (DNA 29)*, volume 276 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 8:1–8:22, Dagstuhl, Germany, 2023. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.DNA.29.8.
- 21 Hong Qian. Thermodynamic and kinetic analysis of sensitivity amplification in biological signal transduction. *Biophysical Chemistry*, 105(2–3):585–593, 2003. doi:10.1016/S0301-4622(03)00068-1.
- 22 Lulu Qian and Erik Winfree. Scaling up digital circuit computation with DNA strand displacement cascades. *Science*, 332(6034):1196–1201, 2011. doi:10.1126/science.1200520.
- 23 Georg Seelig, David Soloveichik, David Yu Zhang, and Erik Winfree. Enzyme-free nucleic acid logic circuits. *Science*, 314(5805):1585–1588, 2006. doi:10.1126/science.1132493.
- 24 Tristan Stérin, Abeer Eshra, Janet Adio, Constantine Glen Evans, and Damien Woods. A thermodynamically favoured molecular computer: Robust, fast, renewable, scalable. *bioRxiv*, 2025. doi:10.1101/2025.07.16.664196.
- 25 Jimeng Su, Ying Song, Zhipeng Zhu, Xinyue Huang, Jibiao Fan, Jie Qiao, and Fengbiao Mao. Cell–cell communication: new insights and clinical implications. *Signal Transduction and Targeted Therapy*, 9(1):196, 2024. doi:10.1038/s41392-024-01888-z.
- 26 Alexei V. Tkachenko, Bortolo Matteo Mognetti, and Sergei Maslov. Evolutionary chemical learning in dimerization networks, 2025. doi:10.48550/arXiv.2506.14006.

A Proofs for Section 2 (Structural Limitations: Dimerization Networks Cannot Amplify)

Proof of Theorem 1. Let k be the index of the monomer species with the maximum absolute relative sensitivity. That is $\left| \frac{x'_k}{x_k} \right| = \max_j \left| \frac{x'_j}{x_j} \right|$.

We analyze Equation (3) for the index k with maximum relative sensitivity:

$$\frac{x'_k}{x_k} \left(x_k + \sum_{j \neq k} d_{kj} + 4d_{kk} \right) + \sum_{j \neq k} d_{kj} \frac{x'_j}{x_j} = \delta_{1k}. \quad (23)$$

We first establish the sign of x'_k . Assume for the sake of contradiction that $x'_k < 0$. By the definition of k , we know $\left| \frac{x'_j}{x_j} \right| \leq \left| \frac{x'_k}{x_k} \right|$. It follows that:

$$\sum_{j \neq k} d_{kj} \frac{x'_j}{x_j} \leq \sum_{j \neq k} d_{kj} \left| \frac{x'_k}{x_k} \right|.$$

Substituting this bound into Equation (23) yields:

$$\frac{x'_k}{x_k} \left(x_k + \sum_{j \neq k} d_{kj} + 4d_{kk} \right) + \sum_{j \neq k} d_{kj} \left| \frac{x'_k}{x_k} \right| \geq \delta_{1k}.$$

Since we assumed $x'_k < 0$, we have $\frac{x'_k}{x_k} = - \left| \frac{x'_k}{x_k} \right|$. The terms involving $\sum_{j \neq k} d_{kj}$ cancel out, simplifying the inequality to:

$$\frac{x'_k}{x_k} (x_k + 4d_{kk}) \geq \delta_{1k}.$$

8:20 Amplification at Equilibrium: Limitations and Implementation

The left-hand side is strictly negative (as $x_k > 0$ and $d_{kk} \geq 0$), while the right-hand side is non-negative ($\delta_{1k} \in \{0, 1\}$). This is a contradiction. Therefore, x'_k must be non-negative.

Having established $x'_k \geq 0$, we now derive the upper bound. Rearranging Equation (23):

$$\frac{x'_k}{x_k} \left(x_k + \sum_{j \neq k} d_{kj} + 4d_{kk} \right) = \delta_{1k} - \sum_{j \neq k} d_{kj} \frac{x'_j}{x_j}.$$

Taking the absolute value and applying the triangle inequality:

$$\left| \frac{x'_k}{x_k} \right| \left(x_k + \sum_{j \neq k} d_{kj} + 4d_{kk} \right) \leq \delta_{1k} + \left| \sum_{j \neq k} d_{kj} \frac{x'_j}{x_j} \right| \leq \delta_{1k} + \sum_{j \neq k} d_{kj} \left| \frac{x'_j}{x_j} \right|.$$

Subtracting $\sum_{j \neq k} d_{kj} \left| \frac{x'_j}{x_j} \right|$ from both sides yields:

$$\left| \frac{x'_k}{x_k} \right| (x_k + 4d_{kk}) \leq \delta_{1k}.$$

Since $x_k > 0$ and $d_{kk} \geq 0$, the term $(x_k + 4d_{kk})$ is strictly positive. If $k \neq 1$, then $\delta_{1k} = 0$, which implies $|x'_k/x_k| \leq 0 \implies x'_k = 0$. However, since k is the index of maximum sensitivity, this would imply all sensitivities are zero, which contradicts the existence of the perturbation σ . Therefore, the maximum relative sensitivity must occur for the input species, i.e., $k = 1$.

Substituting $k = 1$ back into the inequality gives $x'_1 (x_1 + 4d_{11}) \leq x_1$. Solving for the sensitivity x'_1 :

$$x'_1 \leq \frac{1}{1 + 4\frac{d_{11}}{x_1}} \leq 1.$$

Thus, the concentration of the free monomer species X_1 cannot increase more than the concentration of the added signal σ . $\blacktriangleleft$

Proof of Theorem 3. To prove the theorem, we must bound the term $\sum_{j < i} |d'_{ji}|$. Applying Equation (6) to each lower-indexed species $j < i$ and isolating the term d'_{ji} gives:

$$|d'_{ji}| \leq \delta_{1j} + \sum_{k < j} |d'_{kj}| - |x'_j| - \sum_{\substack{k > j \\ k \neq i}} |d'_{jk}| - 2|d'_{jj}|.$$

Summing this over all $j < i$:

$$\sum_{j < i} |d'_{ji}| \leq \sum_{j < i} \left(\delta_{1j} + \sum_{k < j} |d'_{kj}| - |x'_j| - \sum_{\substack{k > j \\ k \neq i}} |d'_{jk}| - 2|d'_{jj}| \right). \quad (24)$$

We now simplify the double summations over the interaction terms d . The expression contains two sums: one over $k < j$ (complexes made of lower-indexed species) and one over $k > j$ (complexes made of higher-indexed species).

$$\sum_{j < i} \left(\sum_{k < j} |d'_{kj}| - \sum_{\substack{k > j \\ k \neq i}} |d'_{jk}| \right) = - \sum_{j < i} \sum_{k > i} |d'_{jk}|.$$

This equality holds because for any pair $\{u, v\}$ both less than i , the term $|d'_{uv}|$ appears once positively (when $j = v, k = u$) and once negatively (when $j = u, k = v$). These internal terms cancel out, leaving only the negative terms where $k > i$.

Substituting this simplification back into Equation (24):

$$\sum_{j < i} |d'_{ji}| \leq \sum_{j < i} \left(\delta_{1j} - |x'_j| - 2|d'_{jj}| - \sum_{k > i} |d'_{jk}| \right). \quad (25)$$

Finally, we merge Equation (25) into Equation (6):

$$|x'_i| + \sum_{j > i} |d'_{ij}| + 2|d'_{ii}| \leq \delta_{1i} + \sum_{j < i} \delta_{1j} - \sum_{j < i} \left(|x'_j| + 2|d'_{jj}| + \sum_{k > i} |d'_{jk}| \right).$$

Since $\sum_{j \leq i} \delta_{1j} = 1$ (as the sum includes the only perturbed species X_1), we arrive at the inequality stated in Equation (7). ◀

B Proofs for Section 4 (Thermodynamic Bounds on Amplification)

Proof of Lemma 8. First, we evaluate the relationship between the Active and Resting states. By definition, $\mathbf{x}$ is the absolute global minimizer of the free energy $g(x)$ over the entire feasible concentration space subject to the total mass constraint $Ax = x^0(\sigma)$. The Resting state, defined as $\mathbf{z} = \text{IC}(\mathbf{x}) + \text{argmin}_{Ax=x^0(0)} g(x)$, is also a physically valid configuration in this same feasible space; its decoupled input mass and minimized residual mass exactly sum to the total network mass $x^0(\sigma)$. Because $\mathbf{x}$ is the unconstrained minimizer over the entire valid space, and $\mathbf{z}$ represents just one specific state within that space (specifically, one restricted to have no mixed complexes), it necessarily follows that $g(\mathbf{x}) \leq g(\mathbf{z})$.

Second, we compare the Resting and Drained states. Based on Definition 6 and Definition 7, both states share the exact same isolated input pool, $\text{IC}(\mathbf{x})$. Therefore, any difference in their total free energies depends entirely on their respective residual pools.

The residual pool of the Drained state is given by the projection $\text{RC}(\mathbf{x})$. Because this projection conserves all residual monomers from the Active state, it perfectly satisfies the isolated residual mass constraint $A \cdot \text{RC}(\mathbf{x}) = x^0(0)$. Thus, $\text{RC}(\mathbf{x})$ is a valid concentration vector in the isolated residual space. However, the residual component of the Resting state is explicitly defined as the global free energy minimizer over this exact same isolated space (i.e., $\text{argmin}_{Ax=x^0(0)} g(x)$). Because the Resting state pairs the shared input pool with the thermodynamically optimal residual configuration, whereas the Drained state pairs it with the unoptimized active projection, we must conclude that $g(\mathbf{z}) \leq g(\mathbf{y})$. ◀

Proof of Theorem 10.

Step 1: Expressing $g(\mathbf{y}) - g(\mathbf{x})$. Expanding the free energy function for the Drained and Active states yields:

$$\begin{aligned} g(\mathbf{y}) - g(\mathbf{x}) = & \sum_{v \geq 1} \left[\mathbf{y}_{\mathbf{0},v} (\log \mathbf{y}_{\mathbf{0},v} - 1 + G_{\mathbf{0},v}) - \mathbf{x}_{\mathbf{0},v} (\log \mathbf{x}_{\mathbf{0},v} - 1 + G_{\mathbf{0},v}) \right] \\ & + \sum_{u \neq \mathbf{0}} \left[\mathbf{y}_{u,0} (\log \mathbf{y}_{u,0} - 1 + G_{u,0}) - \sum_v \mathbf{x}_{u,v} (\log \mathbf{x}_{u,v} - 1 + G_{u,v}) \right] \end{aligned}$$

By definition, the Drained state uses RC and IC operators mapping all mixed complexes onto their respective pure residual and pure input pools. Thus, for any pure residual core u (where $u \neq \mathbf{0}$), its concentration in the Drained state is exactly the sum over all its possible input components v , which means $\mathbf{y}_{u,0} = \sum_v \mathbf{x}_{u,v}$.

Since all extracted input cores get mapped onto the pure input pool complexes of the same size, the concentration of pure input complexes of size v (where $v \geq 1$) becomes $\mathbf{y}_{\mathbf{0},v} = \sum_u \mathbf{x}_{u,v}$.

Substituting these distributions into the free energy function $g(x)$ and grouping the terms by their input cores v and residual cores u , the energy difference resolves to:

$$\begin{aligned} g(\mathbf{y}) - g(\mathbf{x}) = & \sum_{v \geq 1} \left[\sum_u \mathbf{x}_{u,v} (\log \sum_u \mathbf{x}_{u,v} - 1 + G_{\mathbf{0},v}) - \mathbf{x}_{\mathbf{0},v} (\log \mathbf{x}_{\mathbf{0},v} - 1 + G_{\mathbf{0},v}) \right] \\ & + \sum_{u \neq \mathbf{0}} \left[\sum_v \mathbf{x}_{u,v} (\log \sum_v \mathbf{x}_{u,v} - 1 + G_{u,0}) - \sum_v \mathbf{x}_{u,v} (\log \mathbf{x}_{u,v} - 1 + G_{u,v}) \right] \end{aligned}$$

By expanding the equations and canceling the constant -1 terms (noting that these terms in the second line perfectly cancel each other out, while the terms of the first line leave the total scalar concentration of all mixed complexes), the energy difference resolves to:

$$\begin{aligned} g(\mathbf{y}) - g(\mathbf{x}) = & \underbrace{\sum_{v \geq 1} \left[\left(\sum_u \mathbf{x}_{u,v} \right) \log \left(\sum_u \mathbf{x}_{u,v} \right) - \mathbf{x}_{\mathbf{0},v} \log \mathbf{x}_{\mathbf{0},v} \right]}_{\mathcal{A}_{\text{inp}}} \\ & + \underbrace{\sum_{u \neq \mathbf{0}} \left[\left(\sum_v \mathbf{x}_{u,v} \right) \log \left(\sum_v \mathbf{x}_{u,v} \right) - \sum_v \mathbf{x}_{u,v} \log \mathbf{x}_{u,v} \right]}_{\mathcal{A}_{\text{res}}} \\ & - \sum_{v \geq 1} \sum_{u \neq \mathbf{0}} [\mathbf{x}_{u,v} (1 + \Delta G_{u,v}^{\text{bind}})] \end{aligned} \quad (26)$$

where the quantities $\mathcal{A}_{\text{inp}}$ and $\mathcal{A}_{\text{res}}$ capture the Shannon-like entropy of mixing for the pure input complexes and the pure residual complexes, respectively.

Step 2: Establishing an Upper Bound for the Free Energy Difference. To establish the upper bound on the free energy difference $g(\mathbf{y}) - g(\mathbf{x})$, we must evaluate the entropic terms $\mathcal{A}_{\text{inp}}$ and $\mathcal{A}_{\text{res}}$. We rely on a fundamental algebraic inequality for logarithms: for any $A, B > 0$, the relation $A \log(1 + B/A) \leq B$ holds true. Consequently, we can bound the entropy of a sum as follows:

$$(A + B) \log(A + B) - A \log A = A \log \left(1 + \frac{B}{A} \right) + B \log(A + B) \leq B + B \log(A + B)$$

We apply this inequality directly to the sums defining $\mathcal{A}_{\text{inp}}$ and $\mathcal{A}_{\text{res}}$. For $\mathcal{A}_{\text{inp}}$, let $A = \mathbf{x}_{\mathbf{0},v}$ and $B = \sum_{u \neq \mathbf{0}} \mathbf{x}_{u,v}$, such that $A + B = \mathbf{y}_{\mathbf{0},v}$. This yields:

$$\mathcal{A}_{\text{inp}} \leq \sum_{v \geq 1} \left[\left(\sum_{u \neq \mathbf{0}} \mathbf{x}_{u,v} \right) (\log \mathbf{y}_{\mathbf{0},v} + 1) \right] = \sum_{u \neq \mathbf{0}} \sum_{v \geq 1} [\mathbf{x}_{u,v} (\log \mathbf{y}_{\mathbf{0},v} + 1)]$$

Similarly, for $\mathcal{A}_{\text{res}}$, we set $A = \mathbf{x}_{u,0}$ and $B = \sum_{v \geq 1} \mathbf{x}_{u,v}$, yielding $A + B = \mathbf{y}_{u,0}$. Expanding the subtracted sum over v , we find:

$$\begin{aligned} \mathcal{A}_{\text{res}} & \leq \sum_{u \neq \mathbf{0}} \left[\left(\sum_{v \geq 1} \mathbf{x}_{u,v} \right) (\log \mathbf{y}_{u,0} + 1) - \sum_{v \geq 1} \mathbf{x}_{u,v} \log \mathbf{x}_{u,v} \right] \\ & = \sum_{u \neq \mathbf{0}} \sum_{v \geq 1} [\mathbf{x}_{u,v} (\log \mathbf{y}_{u,0} + 1 - \log \mathbf{x}_{u,v})] \end{aligned}$$

Substituting these bounded entropic quantities back into Equation (26), the total subtracted scalar concentration of the mixed complexes (the isolated sum in the third line) perfectly absorbs one of the +1 constants generated by our inequalities. The expression simplifies into a unified single summation strictly over the mixed complexes:

$$\begin{aligned} g(\mathbf{y}) - g(\mathbf{x}) &\leq \sum_{u \neq \mathbf{0}} \sum_{v \geq 1} \mathbf{x}_{u,v} \left[\log \mathbf{y}_{\mathbf{0},v} + 1 + \log \mathbf{y}_{u,0} + 1 - \log \mathbf{x}_{u,v} - 1 - \Delta G_{u,v}^{\text{bind}} \right] \\ &= \sum_{u \neq \mathbf{0}} \sum_{v \geq 1} \mathbf{x}_{u,v} \left[\log \left(\frac{\mathbf{y}_{\mathbf{0},v} \mathbf{y}_{u,0}}{\mathbf{x}_{u,v}} \right) + 1 - \Delta G_{u,v}^{\text{bind}} \right] \end{aligned}$$

Because each term in the preceding expression is multiplied by $\mathbf{x}_{u,v}$, any complex (u, v) with a zero Active-state concentration vanishes and can be excluded. Consequently, instead of summing over all $u \neq \mathbf{0}$, we can restrict the summation to $u \in \mathcal{U}$. Recall that this set is defined as: $\mathcal{U} = \{u \neq \mathbf{0} \mid \exists v \geq 1: \mathbf{x}_{u,v} > 0\}$.

$$g(\mathbf{y}) - g(\mathbf{x}) \leq \sum_{u \in \mathcal{U}} \sum_{v \geq 1} \mathbf{x}_{u,v} \left[\log \left(\frac{\mathbf{y}_{\mathbf{0},v} \mathbf{y}_{u,0}}{\mathbf{x}_{u,v}} \right) + 1 - \Delta G_{u,v}^{\text{bind}} \right] \quad (27)$$

► **Remark 17.** Before applying simplifications and decoupling steps to derive a closed-form tractable bound, evaluating the component-level dependencies of the preceding inequality may be appropriate to generate design strategies and effective amplifiers. One such heuristic is apparent: maximize the concentration of unbound “input partners” ($\mathbf{x}_{u,0}$) for any residual core u where the corresponding mixed complex ($\mathbf{x}_{u,v}$) is significant, which increases the drained concentration $\mathbf{y}_{u,0}$ and mitigates the negative logarithmic penalty associated with complex formation.

The formulation in Equation (27) casts the free energy transition strictly as an expectation over the mixed complexes. To globally decouple the specific complex concentrations, we apply the Log-Sum Inequality: $\sum a_i \log(b_i/a_i) \leq (\sum a_i) \log(\sum b_i / \sum a_i)$.

For the numerator sum $\sum b_i$, we evaluate $\sum_{u \in \mathcal{U}} \sum_{v \geq 1} \mathbf{y}_{\mathbf{0},v} \mathbf{y}_{u,0}$. Because all concentration terms are non-negative, we can create a separable upper bound, $\sum_{v \geq 1} \mathbf{y}_{\mathbf{0},v} \cdot \sum_{u \in \mathcal{U}} \mathbf{y}_{u,0}$. Crucially, because the Drained state is defined as sums over the Active state complexes, summing these over their respective domains exactly partitions them into the pure and mixed Active state pools:

$$\begin{aligned} \sum_{v \geq 1} \mathbf{y}_{\mathbf{0},v} &= \underbrace{\sum_{v \geq 1} \mathbf{x}_{\mathbf{0},v}}_{X_{\text{inp}}} + \underbrace{\sum_{u \in \mathcal{U}} \sum_{v \geq 1} \mathbf{x}_{u,v}}_{X_{\text{mix}}} & \sum_{u \in \mathcal{U}} \mathbf{y}_{u,0} &= \underbrace{\sum_{u \in \mathcal{U}} \mathbf{x}_{u,0}}_{X_{\text{res}}} + \underbrace{\sum_{u \in \mathcal{U}} \sum_{v \geq 1} \mathbf{x}_{u,v}}_{X_{\text{mix}}} \end{aligned}$$

Substituting this separated numerator back into the Log-Sum framework, the upper bound resolves to a thermodynamic evaluation dependent solely on the Active state partitioning:

$$g(\mathbf{y}) - g(\mathbf{x}) \leq X_{\text{mix}} \left[\log \left(\frac{(X_{\text{inp}} + X_{\text{mix}}) \cdot (X_{\text{res}} + X_{\text{mix}})}{X_{\text{mix}}} \right) + 1 - \Delta G_{\text{stg}}^{\text{inp}} \right]$$

To make the expression easier to work with, we observe that $X_{\text{inp}} + X_{\text{mix}} \leq \sigma$, so:

$$g(\mathbf{y}) - g(\mathbf{x}) \leq \underbrace{X_{\text{mix}} \left[\log \left(\frac{\sigma (X_{\text{res}} + X_{\text{mix}})}{X_{\text{mix}}} \right) + 1 - \Delta G_{\text{stg}}^{\text{inp}} \right]}_{h(X_{\text{mix}})}$$

8:24 Amplification at Equilibrium: Limitations and Implementation

To find the global maximum of this bound, we treat $(X_{\text{res}} + X_{\text{mix}})$ as a fixed parameter. Under this condition, $h(X_{\text{mix}})$ is a strictly concave function with respect to X_{mix} . By setting the derivative to zero, the unconstrained maximum of h occurs at $X_{\text{mix}} = \sigma (X_{\text{res}} + X_{\text{mix}}) e^{-\Delta G_{\text{stg}}^{\text{inp}}} = \sigma e^{\psi_{\text{inp}}}$.

Because X_{mix} is physically constrained by the total input concentration ($X_{\text{mix}} \leq \sigma$), we must evaluate the boundary conditions. If $\psi_{\text{inp}} \leq 0$, the optimal value lies within the valid physical domain, yielding a maximum of $\sigma e^{\psi_{\text{inp}}}$. If $\psi_{\text{inp}} > 0$, the unconstrained optimal value exceeds σ ; because h is concave, the constrained maximum must occur at the boundary $X_{\text{mix}} = \sigma$, which evaluates to $\sigma(\psi_{\text{inp}} + 1)$. Consequently, Equation (16) is proven:

$$g(\mathbf{y}) - g(\mathbf{x}) \leq \begin{cases} \sigma e^{\psi_{\text{inp}}} & \text{if } \psi_{\text{inp}} \leq 0 \\ \sigma (\psi_{\text{inp}} + 1) & \text{if } \psi_{\text{inp}} > 0 \end{cases} \quad \blacktriangleleft$$