The Phase State of Control

Paul Campbell
Vice President of AI & ML Solutions, facil.ai

Contents

Objective

To establish how we can use Phase State to better understand Control and Artificial Intelligence and by putting to rest any anxieties that one might lose the benefits achieved by automation, possibly influence Artificial Intelligence adoption rates in the control space. Understanding the efficacy of a proposed control solution is the first step in this direction.

The Meaning and Use of Stoichiometric Measures' Values

To establish Phase State, we need stoichiometric measures. Our stoichiometric measures are derived heuristics, and they describe how much information is required to stabilize a specific mathematical model. They quantify the information burden of a proposed control solution:

C: Command (Kinetic Tax)

This value represents the number of setpoints or control trajectories a control algorithm is authorized to write to the PLC/Automation layer. It represents the specific volume of external energetic input required to bridge the gap between an asset's physical inertia and its intended coordinate and can be thought of as a measure of External Governance Subsidy.

F: Fidelity (Visible Resolution)

This measure counts how many unique sensor streams or points the model must "see" to maintain its mathematical fitness. In our stoichiometry this measure is capped at 5, which represents saturation. Beyond F=5, additional points provide no incremental gain in control efficacy. For Hand-Off-Auto (Manual), F = 0 as a human flipping a switch doesn't require a real-time data feed to exist. For Model Predictive Control, even though there are 8 points in the model, F is set to 5 and is interpreted as having achieved full saturation.

B: Balance (Structural Integrity)

This represents the number of competing goals the algorithm is trying to solve simultaneously. It acts as a mission statement of sorts for the algorithm. For Transformers, B = 0 since they are feed-forward. They predict the next token/state but don't inherently balance a physical actuator trade-off. For Reinforcement Learning, F = 3 as it is constantly balancing possible benefits against attending costs. A generic example of F=3 might be "to bring temperature and pressure to their respective targets while observing a constraint on volume."

The Periodic Table of Elements illustrates the valency of each element. Perhaps we should attempt the same with the six expressions of control and the known control solutions, adding additional control solutions as they become known (see **Figure 1**). Just as a molecule of water requires exactly two Hydrogen and one Oxygen to be stable, Model Predictive Control requires exactly 1 Command, a maximum of 5 Visibility points, and 3 Balance objectives to reach its optimal 'convergent' state. Given less, it's unstable and given more, it's over-determined. It frames the chart as a periodic table of requirements rather than an inert list.

Balance	Autonomy Threshold	Visibility Ratio
$\sum B_{Solution} - \sum C_{Expression} = 0$	$\sum C_{Soltiuon} \to 0$	$\sum F_{Solution} \geq 2$

Table 1; Stoichiometric Efficacy of Control

However, efficacy is not fully expressed by a Periodic Table, which only helps us understand the potential fitness of the match between control expression and control solution. To establish the efficacy, we need to know what the end result is to be, and we need a method to measure fitness. Each of the six control expressions (hereafter we'll refer to these as profiles) and has a set of stoichiometric values, as do control solutions. **Table 1** demonstrates three thresholds using stoichiometric values that must be met to establish control efficacy:

1. There must be a balance of valency, i.e. a real-time valency match between the expression of control and the control solution. If the balance does not equal 0, there will be no match.
2. The need for a Command must diminish over time, approaching 0 perhaps even reaching 0.
3. There must be visibility of the effect. A F value >=2 implies visibility but a max F value of 5 implies total saturation.

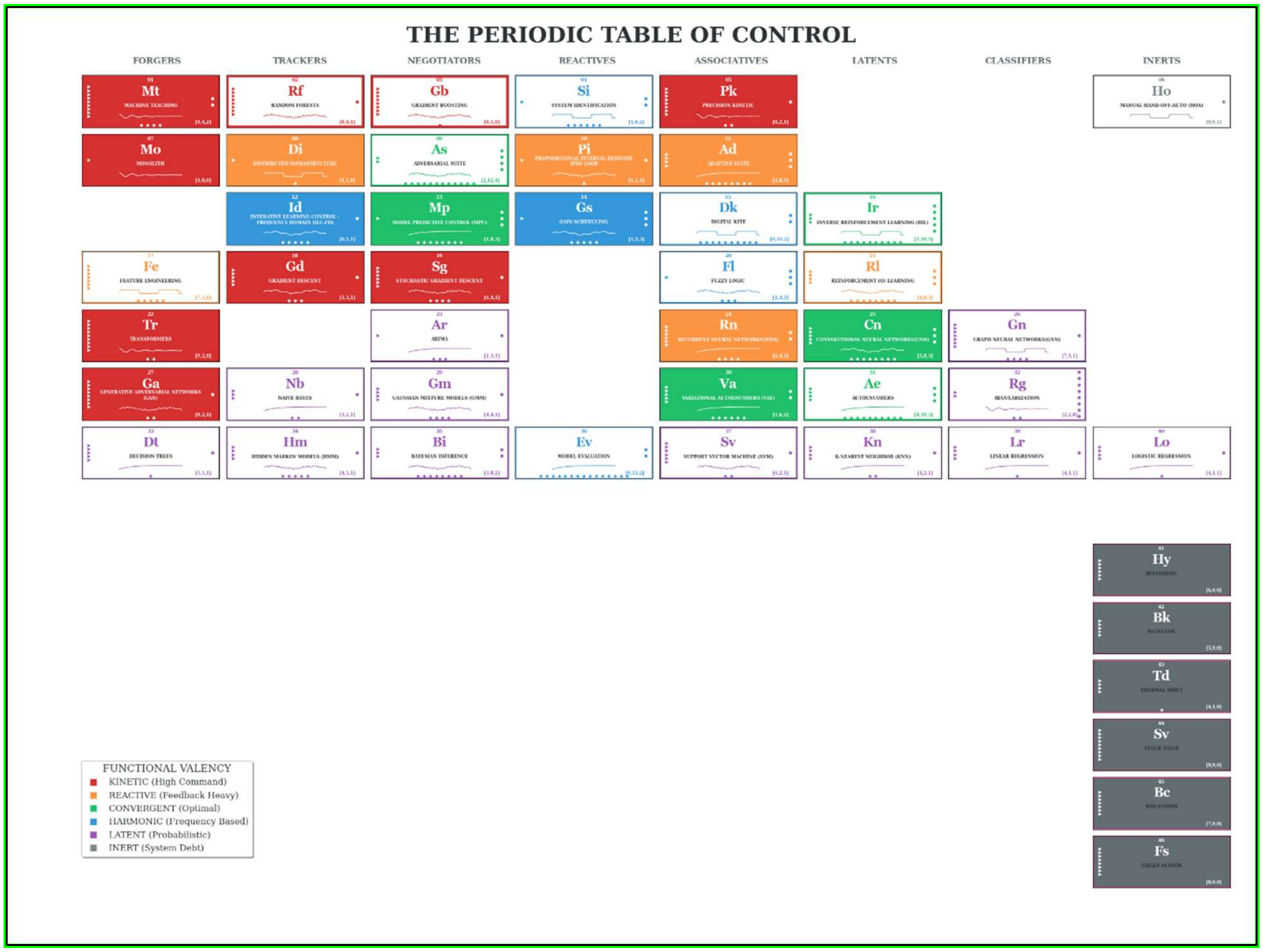

Figure 1: Periodic Table of Control

Thermodynamics of Control & the Latent Debt of Phase Transitions

The time required to reconcile latent phase debt is a controllable variable. In Solid and Liquid states (Automation), the system incurs a debt of mechanical and informational stress due to the imbalance of High Command (C) and Low Fidelity (F). This is illustrated in **Figure 2**. As a measure of control, rather than trying to quantify the latent phase debt, we can direct our energies to minimizing the reconciliation period where kinetic instability prevails and the system is paying for its lack of inherent governance. The time required to reach the debt-free gas phase state can be our measure: by saturating the stoichiometry at the point of phase change, the solution eliminates the latent debt, resulting in immediate mission reconciliation with no settling interval.

The Solid Phase State: Static Automation (F=0)

Automation is characterized by high Command ($C >= 7$) and no Visibility ($F = 0$). In this state, automation cannot flow around obstacles such as system disturbances or system debt. A loss of communication or a sensor failure is enough to drive the C value to zero (C

→ 0). Hence, the system fails or is severely impaired. Settling for automation is choosing a system that is easy to understand but impossible to optimize.

The Liquid Phase State: Dynamic Control (F > 0)

With the introduction of Visibility (F > 0), the control process can now begin with active negotiation among the C, F, and B stoichiometric measures. Like a fluid, the control solution seeks it level and fills the cracks left by resident system debt, compensating for hysteresis in real-time, for example. This is the phase state where we plot Model Predictive Control (MPC) in **Figure 3**. Control in this state is efficacious but requires constant energy in the form of computational resources to maintain control. In the liquid phase state, C will never reach 0 (C >=1).

The Gaseous Phase State: Self-Governance (C → 0)

This phase state is characterized by high F (>4) and B (>=4) values, but C approaches zero as the system reaches intrinsic equilibrium (see **Figure 4**). The control is pervasive but weightless. The system concerned and the control solution move in perfect synchrony so that there is no command required to maintain state. The catalytic reaction is the signature of this phase state, where the control becomes a part of the system concerned, thereby making the pairing of the control solution and the physical asset an immutable object.

The Phase Transition (The Latent Heat of Integration)

The energy required to move between phases is critical to understanding the meaning of the stoichiometric measures**.** This energy can be thought of as escape velocity, where C > 0 only to provide the energy to overcome inertia but once the phase state has changed, the command terminates (C = 0, see **Figure 2**). Self-governance having been achieved, the control solution can govern the physics of the system concerned without further command instances.

Phase Diagram of Control (**Figure 3**) is fully mapped. If a system remains in the "Solid" phase (Manual/Static Automation), it isn't because the science is missing. It’s because the application of the Command Catalyst (C) required to overcome latent heat initiate the phase change has been refused.

Phase Proximity as a Measure of Control

Phase Proximity is the measure of distance between the control logic’s internal model and the control profile’s physical reality. When a control solution is far away from the physics, it has low proximity. This is the trait of the Automation Gap whereby automation issues commands (C) based on a static schedule, oblivious to the phase state of the asset. Alternatively, when a control solution is mathematically affecting the physics it has high proximity and is in a kinetic state.

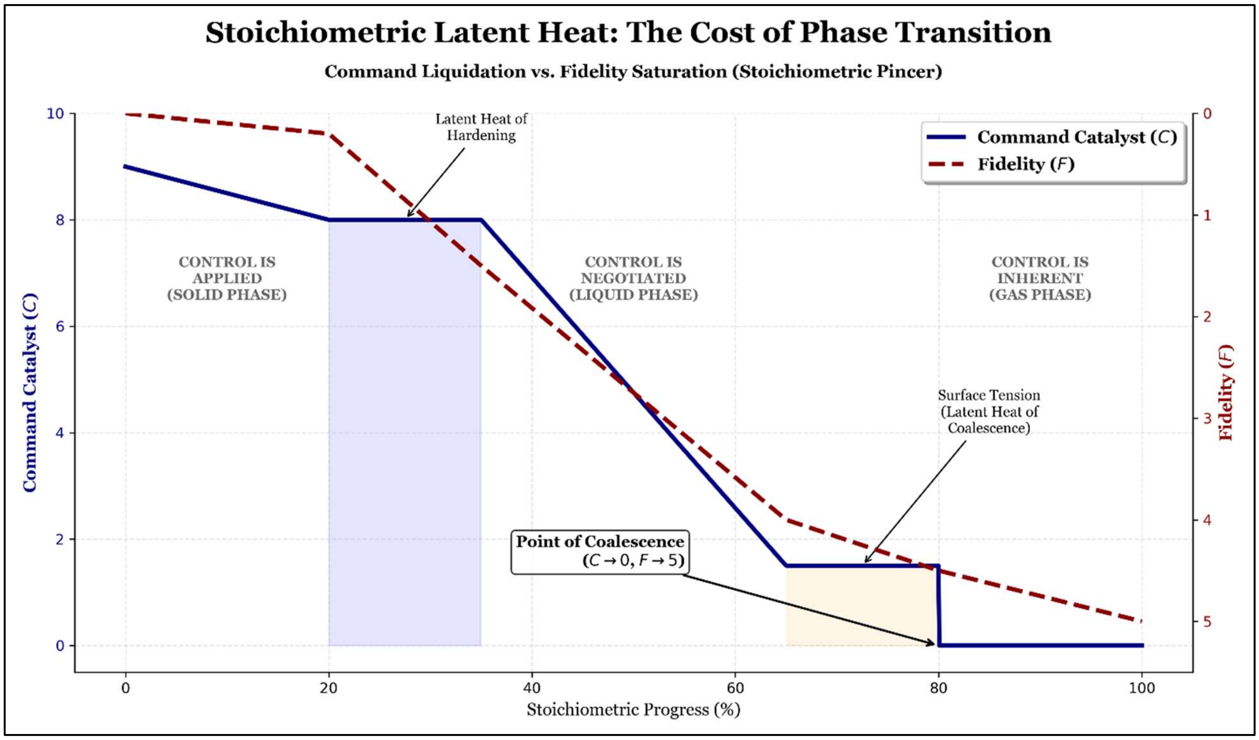

Figure 2: Latent Heat Required for Phase Transitions

Every F (Visibility) value is a real-time reflection of a physical state, and every C (Command) is a micro-adjustment that maintains equilibrium.

A Precision Kinetic profile is characterized by high command authority used with extreme temporal and spatial accuracy. When a Monolith profile is optimized, its control suite moves from a broad, sluggish tracking profile to closely approximate a Precision Kinetic profile, where it no longer hunts for the setpoint but rather it anticipates the load. Phase proximity can therefore serve as a measure of control.

One chooses to settle for automation. Automation is a choice to maintain low phase proximity, likely because it doesn't require deep integration. Ultimately, control is the pursuit of phase proximity. One can argue that once phase proximity reaches 100% – i.e. that the control solution and the reality are synced – that perfect control has been achieved. There is no more efficiency to be gained because the logic is operating at the speed of the physics.

However, if the resource requirements to reach 100% phase proximity are too taxing, we need to calibrate our understanding of control when viewed through our stoichiometric lens. As a practical matter, if a control profile and a control solution pairing reach a phase proximity of ~92%, the Fidelity and Balance have saturated the system's requirements. The unattained ~8% proximity will not change the phase state of the system concerned. This is known as phase-lock in physics, but it is critical to understand that phase-lock does not place functional perfection of control out of reach.

While the monolith profile targets the Precision Kinetic triple point in **Figure 3** to measure phase proximity, other profiles may possess triple points deeper within the coordinate space where control is inherent. However, the stoichiometric rule of phase-lock remains constant: A solution must achieve ~92% proximity to its respective triple point to trigger the transition to C=0. The distance to target is fixed. Only the complexity of the journey changes.

Phase Proximity is the ultimate metric of control efficacy. The transition from Automation to Self-Governance is essentially a journey of driving the proximity gap toward zero, delivering the exact amount of authority required at the exact moment the physics demand it.

Self-Governance (C=0) is the natural end-state of high phase proximity. Once a control profile - control solution pairing is perfectly proximate, the system's own momentum does the work, and the need for external Command evaporates. With external command in the role of catalyst, C >0 is only a temporary necessity.

External Command (C) must be viewed not as a permanent reagent of the control reaction, but rather as its catalyst. Its primary function is to overcome the latent inertia during phase transitions. As the system approaches a state of high Phase Proximity, the catalytic requirement diminishes. In a solved system Command is reclaimed, leaving behind a state of Self-Governance where the physics and the logic exist in a permanent, self-sustaining equilibrium..

We may encounter a control solution that is immediately post-catalytic, where the stoichiometric ratios (C, F, B) were satisfied before the first movement was ever recorded. This immediate post-catalytic signature is characterized by the absence of a transitionary phase. See **Figure 5**.

While automation exhibits a kinetic struggle as it reconciles latent debt, and a less than perfect control solution (where C→0 and F→5 but neither stoichiometric measure has reached their target values and perhaps may never) wrestles with latent friction, an immediate post-catalytic solution (C=0, F=5) manifests as absolute transparency where the delta between intent and outcome is zero from the moment of inception. This is the goal behind achieving the gas phase state: to establish a system that does not transition into governance but rather exists as a governed and immutable object. The difference is that automation and imperfect control are each a process (becoming) and perfect control is a state (being).

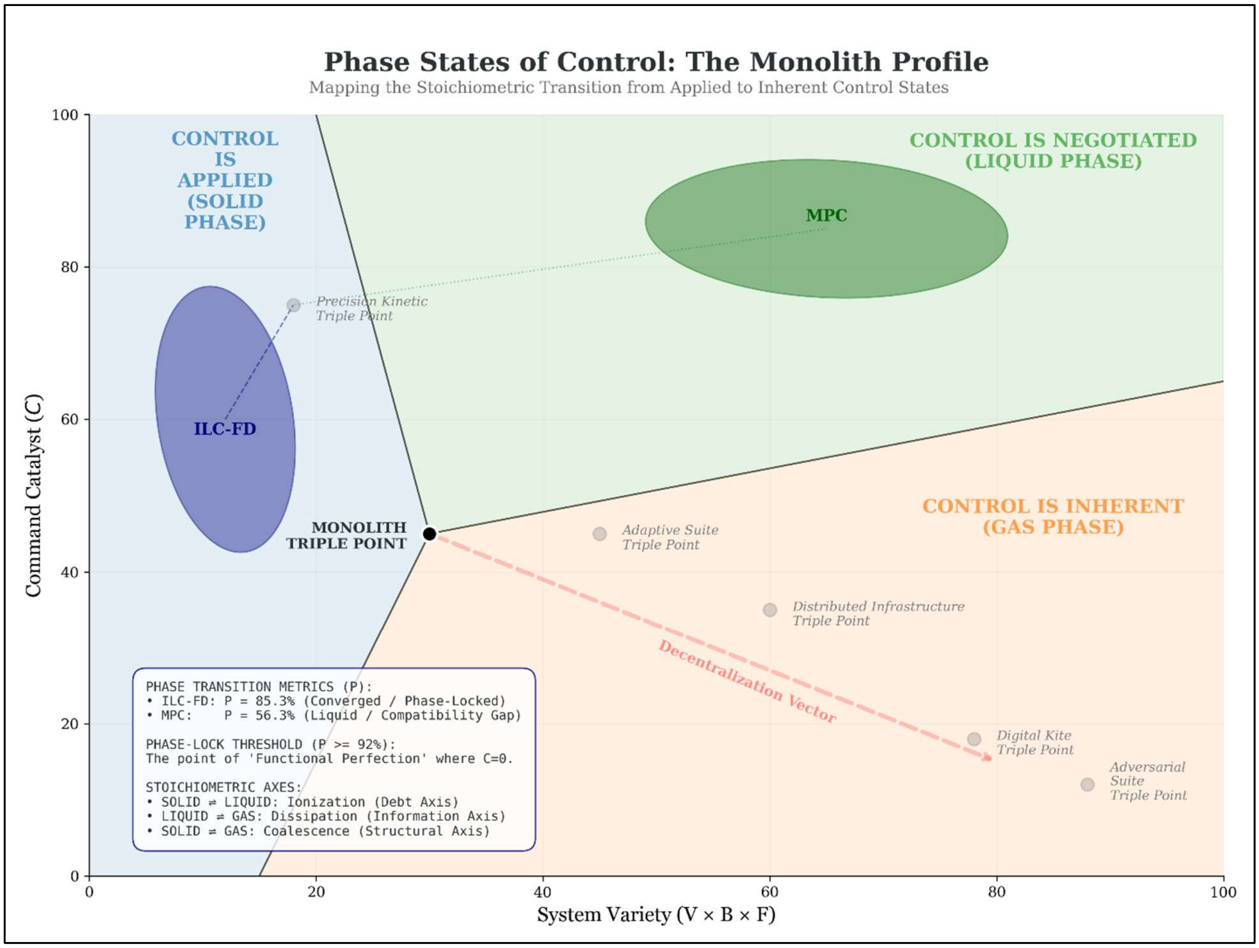

Figure 3: Phase State of Control

C → 0: Commanding is Only a Temporary Necessity to Establish Stability

The control solution needs high C initially to begin the phase transition. Once the system reaches phase proximity, the C value drops. **Figure 4** illustrates the Command Intensity curve as it trends toward zero as the Self-Governance curve trends toward one. This catalytic cycle demonstrates that the AI is settling the asset, not hijacking it.

The fear of the loss of control is the logical result of a visibility gap. If the control solution has F=0, a loss of control can and will be assumed, understandably. A control solution must have visibility in order to preempt the fear of loss of control. When F = 5, the saturation can be construed as a type of “confidence reagent”. We don't fear a cruise control in our automobiles because we can see the speedometer, i.e. there is visibility. Efficacy in a control solution requires stoichiometric transparency, where a high F value allows for a full observation of the self-governance in action.

The role of Phase Distance in control merits acknowledgment here. When we treat a control solution as an external force, we fear its command (C). But when we recognize a control solution as a catalyst designed to achieve self-governance, we realize that its goal is to reach C = 0 (C → 0 rather than C = 1). The ultimate success of a control solution is not that it takes the wheel, but that it stabilizes the road so perfectly that the wheel is no longer a point of stress.

The elimination of the control solution's response time is how the post-catalytic signature (see **Figure 5**) is defined. Where traditional signatures exhibit a visible process, the post-catalytic signature manifests as absolute symmetry. It is the kinetic realization of C = 0 and F = 5 at the outset. This is a rare occurrence. It is far more common to see C → 0 than C = 0 upon inception.

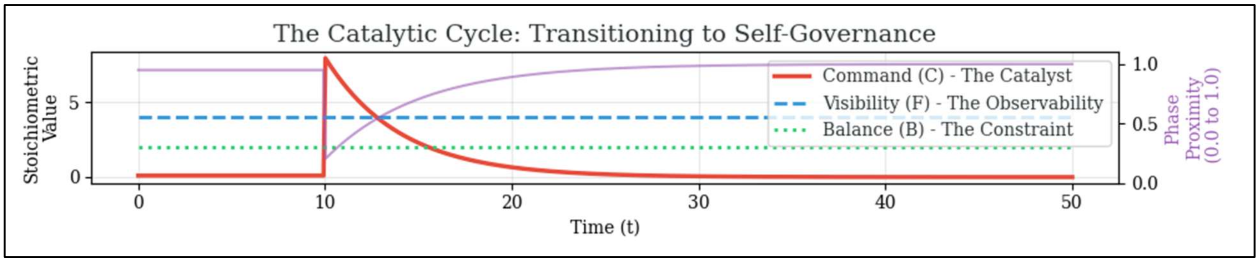

Figure 4: The Catalytic Cycle

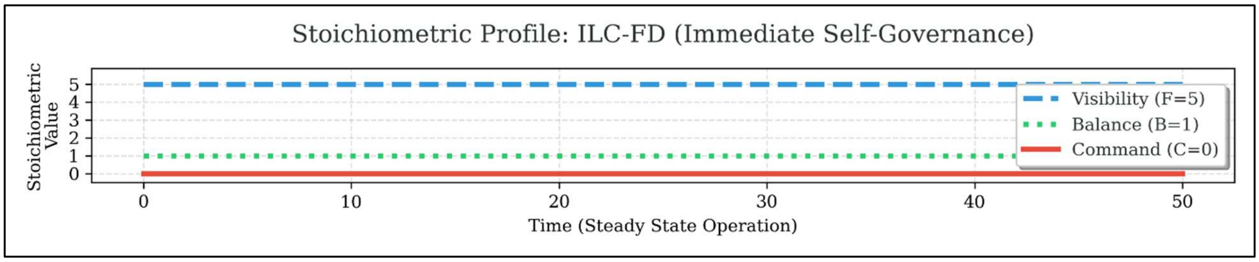

Figure 5: Immediate Self-Governance

Control: The Solved Problem

While control is a solved science, it is not a solved practice. Mathematically, control *is* solved. Algorithms are able to handle any control profile, and the C, F, and B variables present themselves as efficacy measures. However, choosing automation isn't always a voluntary preference. More likely it's a stoichiometric failure in the infrastructure. To move from Automation (F=0) to Control (F>0), you must have already invested or committed to invest in the requisite reagents (sensors, high-speed telemetry, computational capacity, etc.). Settling for automation is often a choice to avoid the entropy of integration.

Many choose automation (F=0) because it is deterministic. In a world where C=0 (Self-Governance) feels like pixie dust, people retreat to the safety of automation and its limitations. They view a control solution with performance that predictably never quite reaches its optimal state as a safer choice than a system that does reach its optimal state but has enough complexity that it requires a deeper understanding to properly defend or explain. This choice actually incurs massive System Debt (Hysteresis/Inefficiency) that usually costs more than the control upgrade itself.

Most stakeholders don't realize they are choosing automation over control because they lack a way to measure the difference. Without stoichiometric measures for instance, they reduce a control solution to Software as a Service. The horizons of control are open: the math is free and control solutions are available. Blind Automation (F=0) is no longer a technical limitation. It is a choice to accept sub-optimal performance. Control becomes a solved problem when we realize that active commanding is a poor substitute for deeper understanding.

Phase State Equivalency	C	F	B	Status	Control Efficacy	Latent Debt
Solid (Brittle)	= 0	= 0	= 0	Inertia	None	Extreme
	Zero intent, zero visibility, zero structure.					
Liquid (Surface Tension)	> 0	n	< C	Control	Dependency	Moderate
	The system requires a constant "Command Tax" because its internal Balance (B) is insufficient.					
Gas (Immutable Object)	= 0	> 2	= C	Self-Governance	Harmony	Zero
	The solution's architecture (B) satisfies the system's physical requirement (C), allowing external command to evaporate.					

Table 2: Stoichiometric Values and Phase States Summary

Conclusion: Control and Artificial Intelligence

Viewing control through the volume of commands issued lens is archaic and must be abandoned. Control is not a verb, something one does to a system. Control is a noun, the resultant state a system *inhabits*. A Command (C >0) is a tax paid to service system debt. Reaching C=0 is the achievement of self-governance, not the surrender of automation authority. When C=0, the Virtual Closed Loop ceases to be the vehicle of correction and becomes the manifestation of permanent and observable equilibrium. Hence control becomes a solved practice.