EDITOR: Schahram Dustdar, dustdar@dsg.tuwien.ac.at

DEPARTMENT: INTERNET OF THINGS, PEOPLE, AND PROCESSES

Equilibrium-Driven Antifragility in Computing Continuum Systems

Nefeli-Marina Rouska, Víctor Casamayor-Pujol, and Ildefons Magrans de Abril, *Pompeu Fabra University, 08018, Barcelona, Spain*

Schahram Dustdar, *Pompeu Fabra University, 08018, Barcelona, Spain, and TU Wien, 1040, Vienna, Austria*

In this article, we propose a novel framework inspired by neuroscience to define and maintain equilibrium in computing continuum systems (CCSs). To do so, we conceptualize equilibrium not as a static state but as a dynamic condition of predictive regulation that supports resilience and adaptation. Our methodology quantifies distance from equilibrium by modeling system behavior using Bayesian networks and computing Kullback–Leibler divergence between expected and observed system states under perturbation. Conceptually, this approach is grounded in the Fluctuation–Dissipation Theorem, which describes how a system's internal variability predicts its response to external changes. While this is a conceptual article without empirical implementation, it provides a principled foundation based on the Free Energy Principle (FEP) and active inference for future research aimed at managing and potentially designing CCSs that are not only resilient but also antifragile, capable of learning from disruptions to improve their performance over time.

As computing systems have become increasingly distributed and dynamic, ensuring that their performance will be consistent under changing conditions has become a major challenge. This is especially true for computing continuum (CC) systems (CCSs), which are large-scale distributed architectures composed of heterogeneous resources, typically integrating Internet of Things, edge, fog, and cloud components.[1] Traditional approaches in resource management are often not enough in such environments, particularly when they have to deal with uncertainty and heterogeneity.[1] To address these limitations, we begin to explore concepts from neuroscience, where complex systems like the human brain are maintaining stability through predictive and self-organizing behavior.[2,3] In this article, we examine how the notion of *equilibrium* which is a key principle in biological and cognitive regulation, can offer a powerful framework for managing resilient and adaptive CCS.

In CCS, the computation is being distributed in various computational tiers to meet diverse application requirements, including the fulfillment of service-level objectives (SLOs). SLOs are specific, quantifiable, and predefined goals related to performance that systems are expected to meet, such as thresholds for latency, availability, CPU usage, or error rates.[1] However, consistently fulfilling the SLOs is challenging across the CC. That is because of its inherent heterogeneity—nodes differ widely in computational capabilities, memory, energy consumption, and network connection—as well as the constantly changing conditions in workload and infrastructure.[1,4]

Traditional approaches to resource management often use threshold-based mechanisms, where predefined and static rules are used to reconfigure the system after a violation of fixed boundaries.[1] They do

Digital Object Identifier 10.1109/MIC.2025.3597479
Date of current version 23 January 2026.

not typically support probabilistic modeling, continual learning, or knowledge transfer, which blocks the seamless integration of new or heterogeneous devices. These approaches lack flexibility and cannot adapt effectively in environments like the CC where conditions change continuously and the behavior of the different computing units or devices is interdependent. Moreover, they often assume a centralized control with full knowledge of the system. In reality, however, devices usually need to work on their own and make decisions without having information about the whole system. Therefore, there is a clear need for more adaptive, distributed, and predictive frameworks that can support resilient and adaptive behavior across the CC.[4]

To overcome these limitations, recent work proposes the concept of equilibrium as a new regulatory principle for distributed computing systems.[4] In this article, we incorporate the concept of equilibrium in our systems by first focusing on its origins in neuroscience and motivate the relevance of equilibrium-seeking behavior in computation. We are focusing on how adopting this perspective can support the resilience and—a step further—antifragility[5] of systems that can not only withstand disruption and perturbation, but also adapt and improve. Finally, we bridge the gap between the neuroscience perspective and the needs of distributed computing by showing how concepts from neuroscience can inform the design of CC systems that are not only resilient to change, but potentially antifragile—improving their structure and behavior in response to environmental stressors. The broader goal of this work is to explore how the concept of equilibrium can be defined and operationalized across distributed systems.

In general terms, *equilibrium* refers to a condition of balance among competing forces or influences in a system. In classical mechanics, this may indicate a state in which total forces are zero, resulting in stability. In thermodynamics, equilibrium is reached when large-scale properties like temperature and pressure stay constant over time, as long as no external factors disturb the system. These basic definitions show that equilibrium is a condition where the system experiences little or no internal change or tension.[6] From our perspective, we are considering equilibrium as a form of dynamic balance where different parts of the system work in order to maintain overall stability and performance, even when external conditions are changing.

Sedlak et al.[4] defined equilibrium in CCSs as a state where all SLOs are satisfied across the system. In their approach, individual edge devices make autonomous decisions to minimize their own prediction error, while a central cluster agent facilitates coordination and assists in integrating new devices. This combination enables the system to maintain overall balance. Building on this, Sedlak et al. explored how the broader concept of equilibrium could be applied to CCS, which integrate diverse and distributed resources across the edge, fog, and cloud.[4]

Despite recent advances in applying equilibrium-inspired frameworks to CCSs, several challenges are unaddressed. One key limitation is the lack of a formal and operational definition of equilibrium in the environment of CCSs. Prior work does not identify specific and measurable indicators that would allow systems to assess their state relative to an equilibrium. Additionally, there is no established method for quantifying the "distance" from equilibrium or determining when a system has entered a nonequilibrium state. These limitations make it difficult to manage CCSs effectively in dynamic and resource-constrained environments. Furthermore, existing models rarely address practical constraints such as partial observability, external compliance requirements, or the simultaneous optimization of multiple, sometimes competing SLOs.

In this article, we aim to build on this foundation by 1) clarifying what equilibrium means in computational terms, 2) proposing Kullback–Leibler (KL) divergence (DKL) between expected and observed system states as a measurable indicator of deviation from equilibrium, and 3) proposing how systems can approach and maintain equilibrium by drawing insights from neuroscience concepts. Our approach reframes equilibrium not as a passive condition, but as an active, emergent property of self-regulating intelligent CC systems.

EQUILIBRIUM IN NEUROSCIENCE: FROM HOMEOSTASIS TO THE FREE ENERGY PRINCIPLE

In order to better understand the concept of equilibrium in CCSs, it is important to explore its origins in neuroscience.

In biological and cognitive systems, equilibrium is a dynamic process that adjusts accordingly in order to maintain internal stability in a constantly changing interacting environment. The brain offers a good biological example of a system that is regulating itself continuously, taking into account internal and external change. In this section, we will examine how equilibrium is maintained in the brain through foundational concepts in neuroscience like homeostasis, allostasis, and the Free Energy Principle (FEP). These concepts provide insights into how intelligent systems can predict, adapt, and therefore remain stable in dynamic environments.

The term *Homeostasis* was introduced by Walter Cannon and refers to the ability of biological systems

to maintain stable internal conditions (like body temperature, pH, and glucose levels) through reactive mechanisms independently of the constantly changing external conditions. If something changes inside or outside the body, the body responds in ways that will eventually bring it back to its normal state.[7] *Allostasis* on the other hand, introduced by Sterling and Eyer and later elaborated by Sterling,[3] expands the concept of homeostasis by adding the idea of predictive regulation (PR). PR describes the continuous process of adjusting beliefs and actions to minimize prediction error.[8] So instead of simply reacting to changes, organisms predict upcoming demands and adjust internal states in advance. In simpler terms, homeostasis keeps things stable by responding to change, while allostasis prepares the body in advance by predicting what will happen next.

This move from reactive to predictive regulation provides more flexible and efficient adaptation especially under uncertainty or in dynamically changing environments. A similar distinction can be seen in CCSs: Homeostasis-like mechanisms resemble traditional threshold-based resource management, where systems react only after specific boundaries or thresholds are violated. In contrast, allostasis aligns more closely with approaches like active inference, where CCS components anticipate future states and adjust their behavior accordingly, enabling for a more robust and adaptive performance across the continuum.

Building on these ideas, the FEP offers a formal theory of biological self-regulation. Proposed by Karl Friston, FEP proposes that biological systems maintain their structure and function by minimizing a quantity known as *variational free energy*, which acts like an upper bound to the amount of surprise or prediction error that can arise from incoming sensory information.[2] In this framework, the goal is to reduce the difference between the expected and the actual outcomes, something that supports self-preservation and long-term adaptive behavior.

Homeostasis, allostasis, and FEP illustrate together how equilibrium in the brain and body is not about *stasis*, but about maintaining functional integrity through continuous interaction with a dynamic environment. These principles offer powerful insights for designing distributed systems that are not only stable, but also adaptive, predictive, and robust.

Neuroscience of Predictive Regulation: Examples from the Brain

In the human brain, PR plays an important role in managing both body functions and mental processes. A clear example is how the brain handles signals from inside the body, like hunger, thirst, or heart rate. Brain areas such as the anterior insula and anterior cingulate cortex combine current sensory signals with past experiences to predict what's happening inside the body and adjust things like breathing or heart rate accordingly.[9]

One well-known case is thermoregulation. The brain continuously integrates sensory signals, such as changes in temperature, and expects future demands of the body before they occur. For instance, on a hot day, even if the body is currently well-hydrated, the brain can predict that water will soon be lost through sweat. It integrates temperature signals and hormonal levels to shift the body into conservation mode by for example reducing urine output and promoting cooling responses like vasodilation or sweating before any measurable imbalance occurs. This anticipatory adjustment reflects the core principle of allostasis: maintaining stability through prediction rather than reaction.[3]

The FEP also explains how the brain works in perception and movement. In vision, for example, the brain builds models that constantly try to guess what you'll see next. When something looks different from what was expected, this prediction error is sent upwards to update the brain's beliefs. Meanwhile, better predictions are sent downwards to reduce the error.[2] This helps the brain react in smart, flexible ways and deal with uncertainty.

Motor control gives another good example. When we move, the brain doesn't just wait to see what happens. It predicts what the movement should feel like and compares that with what actually happens. If there is a difference, the brain uses it to improve future movements. This is key function for learning how to move better and staying coordinated.

A clear example of the aforementioned in robotics comes from Lanillos et al., who used the FEP in a humanoid robot called iCub. The robot learns to predict what it should see—like where its hand is—and compares this to what it actually sees. If there is a difference, it updates its beliefs and moves (like turning its head) to reduce the error. This creates a smart loop of seeing, predicting, and adjusting.[10]

These examples illustrate that the brain operates not just through reactive mechanisms, but also through anticipatory regulation, continuously updating internal models and initiating actions aimed at minimizing expected surprise or prediction error. This continuous process of belief updating and action selection lies at the core of the FEP: that biological systems act to minimize uncertainty by reducing variational free energy and maintaining functional integrity.[2,9]

These neuroscience principles have inspired frameworks for artificial systems. For instance, Sedlak et al. proposed an innovative approach to achieving equilibrium in CC systems through the application of the Active Inference Framework (AIF). AIF builds upon the FEP and is a concept originally from neuroscience that describes how the brain continuously predicts and evaluates sensory information to minimize long-term uncertainty or surprise. This framework has been adapted to computing systems to enable self-regulating behavior.[4] Sedlak et al. introduced an AIF-based framework to support collaboration between edge devices. In this model, individual edge devices develop a causal understanding of how to enforce their SLOs and can share learned knowledge to speed up the integration of new or heterogeneous devices and help them adapt to system requirements. Their implementation demonstrated that edge devices were able to satisfy four different SLOs after only ten training sessions.[4]

These ideas are now being used to design smarter artificial and distributed systems. In such systems, the ability to predict problems and adapt early is becoming essential for staying stable and working well in changing or unpredictable conditions.

Hierarchical and Adaptive Predictive Control in the Brain

The brain is not only a predictive system; it is also hierarchically structured and dynamically adaptive. Predictive coding theories propose that the cortex is organized in layers, where higher-order areas generate top-down predictions and lower-order sensory areas compute bottom-up prediction errors.[11] This hierarchical exchange allows the brain to integrate abstract goals with sensory information, continuously refining internal models to minimize uncertainty at multiple levels of abstraction.

This hierarchical architecture enables the brain to efficiently coordinate and maintain its functionality.[12] Low-level circuits in the brainstem and hypothalamus regulate immediate physiological states, such as heart rate and respiration, by generating predictions about internal bodily signals and initiating rapid autonomic responses.[9] In contrast, higher cortical areas such as the prefrontal cortex and anterior cingulate cortex are involved in longer-term planning, cognitive control, and behavioral adaptation in response to changing goals or contexts.[2] Across this hierarchy, prediction errors are propagated upwards to update internal models, while predictions move downwards to suppress expected inputs—enabling both short-term regulation and long-term learning.[11] This dynamic exchange allows the system to respond adaptively to perturbations while preserving overall functional stability.

In the following sections, we explore how this biological perspective on predictive equilibrium and dynamical regulation can inform the design of distributed computing systems. We will examine how these concepts contribute to resilience and antifragility in computing environments, and in the "Methodology: Quantifying Equilibrium In CCSs" section we will introduce a methodology for quantifying equilibrium and analyzing equilibrium-seeking behavior in artificial systems.

FROM RESILIENCE TO ANTIFRAGILITY: RETHINKING ADAPTATION IN DISTRIBUTED SYSTEMS

Defining Resilience

In both biological and artificial systems, *resilience* is understood as the capacity of the system to absorb disturbances or recover from them, and as a result to have the ability to restore the normal functionality without having to undergo a fundamental change in structure or performance.

In neuroscience, resilience is connected with the brain's ability to maintain stable functioning despite any internal or external perturbations. Mechanisms like homeostasis and allostasis support this stability by regulating physiological and cognitive functions even during stress or injury. The brain's ability to preserve core functions such as autonomic control, perception, or functional decision-making despite trauma, fatigue, or disease is an example of biological resilience. Importantly, this resilience does not imply a static way of functioning but it includes dynamic reorganization across neural circuits to preserve the important functional goals.[9]

A concrete example of neural resilience can be found in individuals recovering from localized brain injury, such as stroke. Despite the loss of function in specific regions, nearby or opposite-side brain areas can reorganize to counterbalance for the damage that has occurred, enabling partial or even full recovery of motor or language functions that had been affected. This plastic reorganization reflects not only the brain's ability to maintain functional goals but also to dynamically reassign resources under constraints.[13]

Similarly, resilience in CCSs, which are composed of heterogeneous and dynamically changing edge, fog, and cloud nodes, faces the challenge of maintaining performance when disruptions or local failures occur. These systems operate under unpredictable conditions with

limited central control, where devices have variations in capability, observability, and network reliability.[14] Current approaches, such as elastic computing,[15] address this challenge by dynamically scaling the resources in response to workload changes. However, elastic computing mainly reacts to changes, rather than predicting them or adjusting the system on its own. In this context, resilience requires mechanisms that not only maintain SLO compliance but also adapt to variability in the environment and internal system structure more proactively.

Limitations of Traditional Resilience

Resilience provides an important framework to understand system stability under stress; however, it has a limitation in addressing the demands of a highly dynamic and heterogeneous environment such as a CCS. More classical formulations of resilience emphasize restoring of a prior state of equilibrium after a disruption occurs. This idea assumes that disturbances are temporary and that the system's optimal configuration is fixed.

However, in environments where uncertainty is high and where conditions are constantly evolving, such assumptions may not really hold. One major limitation is the reliance on static recovery protocols or predefined strategies of responding or reacting. These are often designed based on failure patterns extracted from past experiences or expected perturbations, and therefore they lack the flexibility to handle novel or combinational disruptions.[16]

In addition, traditional resilience does not imply any learning or any change in the system. A resilient system is not necessarily one that adapts or improves in the process. This reactive approach can create stiff systems that work well for past conditions but struggle to handle new or unexpected situations. As Hollnagel et al. point out, truly resilient performance requires not only responding and monitoring, but also learning and anticipating what will may come.[16]

Moreover, classical resilience approaches often depend on centralized control or full system observability, both of which does not apply to CCSs, where devices must operate with partial knowledge and local autonomy.[14] In such settings, ensuring system-wide goals like SLO compliance requires mechanisms that go beyond simple recovery. In such settings, ensuring system-wide goals like SLO compliance requires mechanisms that go beyond simple recovery. Instead, systems must be able to adapt their internal structure, redistribute its tasks, and anticipate changes before failures occur.

These limitations suggest that while resilience is necessary, it is not sufficient for long-term adaptability in distributed, decentralized environments. More flexible and predictive approaches are required that support continual learning, efficient coordination between system components, local autonomy, and dynamic reconfiguration in response to environmental changes. Without these characteristics, systems could potentially be increasingly vulnerable to unexpected future conditions.

Some studies have suggested that resilience can become more effective when combined with adaptive mechanisms that allow improvement over time.[5] We could refer to this that as *dynamic resilience*, where the resilience of a dynamic system can change at different points in time in parallel with changing conditions. However, this does not necessarily imply that the system learns or improves, it may simply reflect a change in the system's configuration in response to changing demands. Dynamic resilience indicates some adaptability, but it remains distinct from the higher-level property of antifragility.

Beyond Dynamic Resilience: Antifragility

Antifragility, was introduced by Taleb[17] and goes beyond any type of resilience by describing systems that withstand stress and improve because of it. In his framework, systems can be robust (unaffected), resilient (able to recover), or antifragile (able to grow stronger from disruption). Antifragile systems do not return to the same state—they transform into more adaptive configurations over time.

Antifragility, in this sense, marks a shift in how we think about system design: It is not just about surviving uncertainty, it is about growing through it. In the context of CCS, antifragility means designing systems that withstand disruption while improving. These systems learn from stress, fine-tune their internal structure by reducing prediction errors, and make better decisions over time. To do that, they need feedback mechanisms that support more than just recovery, they need to support ongoing learning and the ability to reorganize themselves when needed. Achieving this requires moving from reactive systems to proactive ones, that are capable of exploring new states, adapting to change, and learning to anticipate future demands.

Equilibrium as a Foundation for Antifragility

In this work, we define equilibrium in CCSs as the condition in which the system's predictive model accurately represents its actual state under perturbation,

meaning that deviations between expected and observed behaviors remain within acceptable limits to fulfill performance goals.

This definition frames equilibrium not only as stability or balance but as a measurable correspondence between what the system predicts and how it actually behaves. This correspondence allows the system to operate effectively under uncertainty, but it is not sufficient on its own to ensure continuous equilibrium in nonstationary environments. In such contexts, systems must also be able to adapt and update their internal models over time.

Equilibrium provides the necessary foundation for antifragility. While equilibrium ensures that the system maintains performance despite disruptions, antifragility requires the system to use these deviations to update its internal model and improve future performance. Thus, equilibrium enables the stability from which antifragility—improvement through disruption—emerges.

METHODOLOGY: QUANTIFYING EQUILIBRIUM IN CCS

In this section, we present our methodology for quantifying equilibrium in CCSs following Figure 1. Our approach combines ideas from neuroscience and statistical physics with probabilistic modeling to assess how well a system's predicted behavior matches its actual behavior under perturbation.

Throughout this work, we have described equilibrium in CCS from multiple complementary perspectives. Here, we unify these views into a precise definition. We define equilibrium in CCS as a dynamic condition characterized by the following:

1) *Dynamic balance*: This is where different parts of the system interact to maintain overall stability and performance, even under changing external conditions.
2) *Continuous reorganization*: This is where the system dynamically adjusts its internal

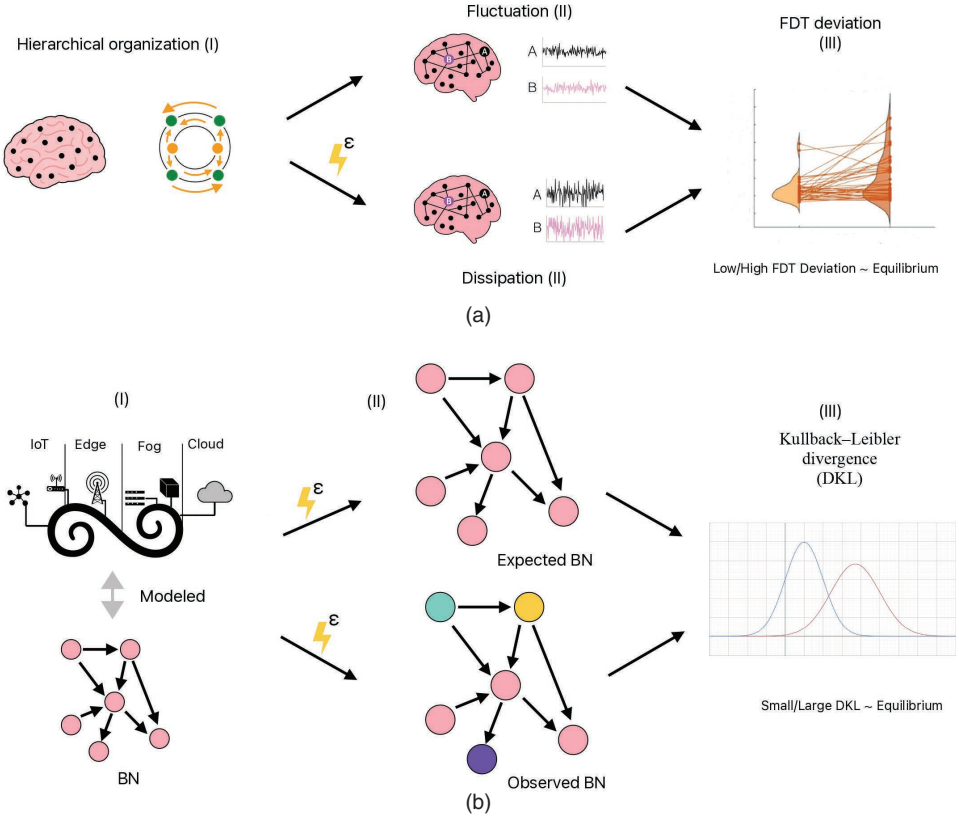

FIGURE 1. Framework comparison between CCSs and brain dynamics equilibrium analysis. (a) Conceptual pipeline for brain dynamics, showing hierarchical organization, fluctuation–dissipation processes under perturbations (ϵ), and computation of FDT deviation as an equilibrium measure. (b) Conceptual pipeline for CCSs, showing modeled BNs, expected BN under perturbations (ϵ), observed BN under perturbation (ϵ), and computation of DKL to assess equilibrium.

configurations to meet global performance targets, such as SLOs, despite internal or external disturbances.

3) *Predictive consistency*: This is where the system's internal models accurately represent its actual behavior under perturbation, enabling successful fulfillment of performance goals.

Taken together, these definitions frame equilibrium not as a static state but as an active property of CCSs, that involves continual prediction, adaptation, and regulation to maintain functional stability. This integrated perspective underpins our methodological operationalization of equilibrium as a measurable condition of predictive consistency, quantified via KL divergence between expected and observed system states. This formal definition serves as the conceptual foundation for our framework, supporting the emergence of antifragility in CCS.

Conceptually, we draw inspiration from the Fluctuation Dissipation Theorem (FDT), which describes how a system's internal variability predicts its response to external changes and how the deviations from this prediction reflect the distance from equilibrium of the system—in this case the human brain [Figure 1(a)]. We operationalize this using KL divergence, which measures the difference between the expected and observed behavior of the system [Figure 1(b)].

Mathematically, our approach aligns with the Expected Free Energy (EFE) framework from AIF, which uses KL divergence to measure prediction error and help the agent choose actions that reduce this error. Together, these foundations provide a principled way to measure when a distributed system remains in equilibrium and when it deviates, enabling improved resilience and targeted adaptation.

FDT as a Conceptual Basis

Equilibrium is a functional state where the CCSs maintain their internal coherence and are able to successfully fulfill their SLOs despite any perturbation—any change in external or internal conditions. This definition is in line with the FDT, a fundamental theorem from statistical physics which describes how a system's internal fluctuations are related to its response to external perturbations.[18] In the context of neural systems, this relationship is shaped by their hierarchical organization [Figure 1(a)/I] where different levels contribute to spontaneous neural activity—fluctuation—and determine the system's response to perturbation—dissipation—[Figure 1(a)/II]. The deviation[12] is computed as

$$D_{A,B} = \frac{\beta\langle AB\rangle_0 - \chi_{A,B}}{\chi_{A,B}} \quad (1)$$

where

- $\langle AB\rangle_0$ is the equilibrium covariance (spontaneous fluctuation) between variables A and B.
- $\chi_{A,B}$ is the susceptibility, quantifying the system's response to a small perturbation in B as measured in A.
- β is the inverse temperature parameter that scales the fluctuation-response relationship.

This measures how far the brain's response $\chi_{A,B}$ deviates from what would be expected based on spontaneous fluctuations $\langle AB\rangle_0$ normalized by the measured susceptibility [Figure 1(b)/III]. Deviations of $D_{A,B}$ from zero indicate nonequilibrium behavior and as a more hierarchical structure in the system.[12]

In our case, the KL divergence between expected and observed Bayesian Networks (Figure 1(b)/II) plays an analogous role, offering a domain-specific way to quantify deviation from predictive equilibrium in distributed systems. This means that equilibrium is defined as a condition where the system's predicted behavior (directly related to its internal fluctuations) matches its observed response under perturbation (analogous to dissipation). Any mismatch, in the KL divergence that we will introduce between expected and observed Bayesian networks, indicates a breakdown in predictive regulation, conceptually similar to an FDT violation. Therefore, our methodology can be seen as a computational analogue of FDT adapted to complex, distributed systems.

System Representation With Bayesian Networks

We propose a methodology that models our systems using Bayesian networks (BNs), which offer interpretable, probabilistic representations of system dependencies. This choice supports explainability and allows us to meaningfully trace how system components relate and respond under perturbation [Figure 1(b)/I], while also enabling us to measure deviations from equilibrium. A BN is a probabilistic graphical model that uses a directed acyclic graph to represent a set of variables and their conditional dependencies. In the context of CCSs, each node corresponds to a metric of the system, and each edge represents the probabilistic dependencies from one component to the other. These networks describe how likely it is to have specific outcomes based on the relationships between the components.

To compare the expected and actual behavior of the system, we first need to define a perturbation, a change introduced to simulate a dynamic environment

(e.g., failing a connection, altering a probability, or modifying a timing response). This perturbation is then applied to the model of the system and we get an expected and its actual behavior.

This results in two BNs [Figure 1(b)/II]:

- BN_{exp}: The expected network, built from the model's internal knowledge—its structure and probabilities—showing how it believes its components should interact.
- BN_{obs}: The observed network, built by applying the same perturbation to the actual system, showing how the system behaves in practice.

Comparing these two allows us to quantify the predictive error introduced by the perturbation and assess how far the system deviates from equilibrium.

Quantifying Deviation From Equilibrium

On intelligent systems, adaptive behavior is thought to be maintained by minimizing *surprise*, a quantity defined as the negative log-probability of an observation under the agent's internal model.[19] This is expressed as $-\ln p(o)$, where o is an observation and $p(o)$ its predicted likelihood. However, this quantity is typically difficult to compute directly. Therefore, in the AIF an alternative measure called EFE is used. EFE provides an approximation of surprise that is conditioned on a sequence of actions, and enables agents to select actions that not only achieve their goals but also reduce uncertainty about future states. EFE acts like an upper bound of surprise.

Although our approach does not use the EFE formula directly, it does have a conceptual and mathematical similarity. In EFE, as shown in the (2), the agent evaluates its future actions by minimizing expected free energy, which combines the pragmatic value and the expected information gain. This connection becomes clearer when we consider the formal structure of the EFE. Formally, it is written as

$$EFE_\tau(\pi) \approx -E_{Q(o_\tau,x_\tau|\pi)}[\ln \tilde{p}(o_\tau)] - E_{Q(o_\tau|\pi)} D_{KL}(Q(x_\tau \mid o_\tau) \mid Q(x_\tau \mid \pi)). \quad (2)$$

Similarly, in our method, we compare the distributions produced by BN_{exp} and BN_{obs} using a probabilistic divergence measure. Specifically, we compute the DKL [Figure 1(b)/III].

$$D_{KL}(BN_{obs}(x) \mid BN_{exp}(x)) = \sum_{x \in BN} P(x_{obs}) \log\left(\frac{P(x_{obs})}{P(x_{exp})}\right). \quad (3)$$

This divergence quantifies how much the observed behavior deviates from the expected. A small divergence indicates that the system remains close to equilibrium, meaning that its predictions are valid, while a big divergence signals a breakdown in predictive regulation. Just like in EFE minimization, this KL-based comparison reflects the mismatch between the predicted and the observed distributions. Minimizing it supports adaptive, goal-consistent behavior in the presence of uncertainty. In active inference, surprise minimization is central to adaptation; analogously, a lower KL divergence in our framework suggests predictive alignment and functional equilibrium. A higher divergence, on the other hand, indicates structural mismatch, failed prediction, or poor adaptability.

However, the core contribution of this conceptual article lies in the synergy between FDT and EFE. In this framework, FDT serves as a quality indicator of the internal model, while EFE computes the agent's drive by balancing pragmatic value and information gain. Therefore, we propose using the FDT signal (or derivative) as a modulatory signal of the components of the EFE formulation. We believe this extension of the classical EFE formulation, with a modulatory component based on the FDT, will enhance the agent's recovery in face of internal and external disruptions. This is because we will be able to define an FDT-based meta-controller able to generate *tradeoff paths* (a sequence of tradeoff values of the EFE) that for instance will first focus on the information gain component and gradually shift toward the nominal balance of the EFE formulation. We believe this enhancement will enable EFE-based agents to behave more optimally in continuously changing environments such as CCSs. We plan to elaborate on this in future work, focusing on architectural design, implementation, and empirical evaluation of the proposed components.

CONCLUSION

This article brings several important benefits for managing and potentially designing CCS. Our main goal was to define what equilibrium means in CCSs by bringing together ideas from neuroscience and statistical physics. This interdisciplinary perspective offers a new way to think about regulation and stability in these systems, going beyond traditional engineering approaches. We also introduce a practical way to measure system stability, using DKL between expected and observed BNs, providing a clear and interpretable metric that can guide more adaptive management strategies. This discussion leads us to the key conceptual contribution of this article, the enhancement of the EFE drive with a modulatory signal generated by an FDT-based component that would be in charge of dynamically optimizing the tradeoff between the

pragmatic and information gain terms. Because it is broadly formulated, this framework can be applied to different CCS architectures, allowing us to assess equilibrium in varied contexts. Finally, it lays down an important conceptual foundation for antifragility, opening up possibilities for building systems that not only stay stable but also learn and improve when faced with disruptions, enhancing their long-term performance and adaptability.

We used AIF and the FEP as our main inspiration. These ideas describe how the brain reduces uncertainty by predicting what will happen and updating its behavior when there is a mismatch. We applied this to CCSs by measuring the system's expected behavior and comparing it to what actually happens after a perturbation. This mismatch, or prediction error, was measured using DKL between two BNs, one expected and one observed. A small KL divergence suggests that the system is close to equilibrium, a large one indicates a breakdown in predictive alignment. Whether a divergence is considered large depends on the system's normal variability.

Our definition of equilibrium is based on PR where systems can act in advance to reduce surprise rather than only reacting after failure. Instead of viewing equilibrium as a static point, we consider it a dynamic condition in which the system adapts continuously to maintain functional stability. To support this idea, we used the FDT. In our approach, fluctuations represent the expected system behavior (based on internal variability), and dissipation represents how the system responds to changes. The difference between them reflects the system's deviation from equilibrium.

Altogether, this framework connects concepts from AIF, FEP, PR, EFE, FDT, DKL, and probabilistic modeling through BNs to support systems that are not only *resilient*—capable of recovering from disruption, but also potentially *antifragile*—capable of improving due to disruption. This is particularly valuable in CCSs, where unpredictability and local autonomy are common.

However, this work has its limitations. Most notably, the framework has not yet been implemented or tested on real CCS deployments since this is a conceptual article, so its practical feasibility, computational cost, and scalability still need to be assessed. The FDT-based modulatory signal adds an extra layer of complexity due to the added flexibility which encourages an even stricter validation phase. Additionally, other system properties, like energy consumption or tradeoffs between multiple objectives, can also affect equilibrium in practice and need to be included. Finally, because this definition of equilibrium depends on each system and its predictive models, its generalizability requires careful consideration of the specific system in each case.

Future work should aim to address these limitations. Testing and implementing the framework in simulated or real CCS environments will be key to understanding its effectiveness and computational performance. Expanding the methodology to include other indicators, such as increasing action space and increasing the monitorable diversity and device types, could provide a more complete picture of equilibrium. Since this adds more variables to monitor and react to, it's not just about checking if the method works (feasibility), but whether it still works well when applied to more complex situations (suitability). This step is important for understanding how far the method can go.

Developing mechanisms that not only measure deviations from equilibrium but also use them to update internal models and improve system performance over time will be crucial for achieving antifragility in practice. Future work could also examine the existence of different equilibrium states that the system can shift between, depending on what is most important in each moment. Finally, another promising direction for future work could also be to explore how this framework could also be used to infer or measure hierarchical structure within CCSs. Since the FDT has been used to reveal hierarchy in brain networks, applying similar principles to CCSs may allow us to assess not just their stability, but also their internal organization and their coordination patterns.

Overall, this conceptual framework is a first step toward defining, measuring, and potentially designing equilibrium-driven antifragile computing systems. Putting it into practice and validating it with real data will be essential to test its full potential.

REFERENCES

1. S. Dustdar, V. C. Pujol, and P. K. Donta, "On distributed computing continuum systems," *IEEE Trans. Knowl. Data Eng.*, vol. 35, no. 4, pp. 4092–4105, Apr. 2023, doi: 10.1109/TKDE.2022.3142856.
2. K. Friston, "The free-energy principle: A unified brain theory?" *Nature Rev. Neurosci.*, vol. 11, no. 2, pp. 127–138, 2010, doi: 10.1038/nrn2787.
3. P. Sterling, "Allostasis: A model of predictive regulation," *Physiol. Behav.*, vol. 106, no. 1, pp. 5–15, 2012, doi: 10.1016/j.physbeh.2011.06.004.
4. B. Sedlak, V. C. Pujol, P. Kumar Donta, and S. Dustdar, "Equilibrium in the computing continuum through active inference," *Future Gener. Comput. Syst.*, vol. 160, pp. 92–108, Nov. 2024, doi: 10.1016/j.future.2024.05.056.

5. C. Axenie et al., "Antifragility in complex dynamical systems," *Sci. Rep.*, vol. 13, no. 1, 2023, Art. no. 12652, doi: 10.1038/s44260-024-00014-y.
6. H. B. Callen, *Thermodynamics and an Introduction to Thermostatistics*. New York, NY, USA: Wiley, 1985. [Online]. Available: https://www.academia.edu/35013052/Herbert_B_Callen_Thermodynamics_and_an_Introduction_to_Thermostatistics
7. W. B. Cannon, *The Wisdom of the Body*. New York, NY, USA: W. W. Norton, 1932. [Online]. Available: https://digital.library.cornell.edu/catalog/chla3117174
8. P. Sterling, "Predictive regulation and human design," *eLife*, vol. 7, Jun. 2018, Art. no. e36133, doi: 10.7554/eLife.36133.
9. L. F. Barrett and W. K. Simmons, "Interoceptive predictions in the brain," *Nature Rev. Neurosci.*, vol. 16, no. 7, pp. 419–429, 2016, doi: 10.1038/nrn3950.
10. P. Lanillos et al., "Active inference in robotics and artificial agents: Survey and challenges," 2021, *arXiv:2112.01871*.
11. A. M. Bastos, W. M. Usrey, R. A. Adams, G. R. Mangun, P. Fries, and K. J. Friston, "Canonical microcircuits for predictive coding," *Neuron*, vol. 76, no. 4, pp. 695–711, 2012, doi: 10.1016/j.neuron.2012.10.038.
12. G. Deco, C. W. Lynn, Y. Sanz Perl, and M. L. Kringelbach, "Violations of the fluctuation-dissipation theorem reveal distinct nonequilibrium dynamics of brain states," *Phys. Rev. E*, vol. 108, no. 6, 2023, Art. no. 064410, doi: 10.1103/PhysRevE.108.064410.
13. S. C. Cramer, "Repairing the human brain after stroke: I. Mechanisms of spontaneous recovery," *Ann. Neurol.*, vol. 63, no. 3, pp. 272–287, 2008, doi: 10.1002/ana.21393.
14. V. C. Pujol, P. K. Donta, A. Morichetta, I. Murturi, and S. Dustdar, "Edge intelligence—Research opportunities for distributed computing continuum systems," *IEEE Internet Comput.*, vol. 27, no. 4, pp. 53–74, Jul./Aug. 2023, doi: 10.1109/MIC.2023.3284693.
15. J. R. Wensing, and G. Stitt, "Elastic computing: A portable optimization framework for hybrid computers," *Parallel Comput.*, vol. 38, no. 8, pp. 438–464, 2012, doi: 10.1016/j.parco.2012.05.001.
16. E. Hollnagel, D. D. Woods, and N. Leveson, *Resilience Engineering: Concepts and Precepts*. Boca Raton, FL, USA: CRC Press, 2017. [Online]. Available: https://www.researchgate.net/publication/50232053
17. N. N. Taleb, *Antifragile: Things That Gain from Disorder*. New York, NY, USA: Random House, 2012.
18. R. Kubo, "The fluctuation-dissipation theorem," *Rep. Prog. Phys.*, vol. 29, no. 1, 1966, Art. no. 255, doi: 10.1088/0034-4885/29/1/306.
19. N. Sajid, P. J. Ball, T. Parr, and K. J. Friston, "Active inference: Demystified and compared," *Neural Comput.*, vol. 33, no. 3, pp. 674–712, 2021, doi: 10.1162/neco_a_01357. [Online]. Available: https://www.researchgate.net/publication/348263054

NEFELI-MARINA ROUSKA is a master's student in the Brain and Cognition program and a researcher in the Engineering Department at Pompeu Fabra University, 08018, Barcelona, Spain. She is the corresponding author of this article. Contact her at nefelimarina.rouska@upf.edu.

VÍCTOR CASAMAYOR-PUJOL is an assistant professor in the Engineering Department at Pompeu Fabra University, 08018, Barcelona, Spain. Contact him at victor.casamayor@upf.edu.

ILDEFONS MAGRANS DE ABRIL is a researcher in the Engineering Department at Pompeu Fabra University, 08018, Barcelona, Spain. Contact him at ildefons.magrans@upf.edu.

SCHAHRAM DUSTDAR is a full professor of computer science, heading the Research Division of Distributed Systems at TU Wien, 1040, Vienna, Austria, and Pompeu Fabra University ICREA, 08018, Barcelona, Spain. Contact him at dustdar@dsg.tuwien.ac.at.