



Introducing Repository Stability

Giuseppe Destefanis
giuseppe.destefanis@brunel.ac.uk
Brunel University of London
London, UK

Silvia Bartolucci
s.bartolucci@ucl.ac.uk
University College London
London, UK

Daniel Graziotin
graziotin@uni-hohenheim.de
University of Hohenheim
Stuttgart, Germany

Rumyana Neykova
rumyana.neykova@brunel.ac.uk
Brunel University of London
London, UK

Marco Ortu
marco.ortu@unica.it
University of Cagliari
Cagliari, Italy

ABSTRACT

Drawing from engineering systems and control theory, we introduce a framework to understand repository stability, which is a repository activity capacity to return to equilibrium following disturbances - such as a sudden influx of bug reports, key contributor departures, or a spike in feature requests. The framework quantifies stability through four indicators: commit patterns, issue resolution, pull request processing, and community engagement, measuring development consistency, problem-solving efficiency, integration effectiveness, and sustainable participation, respectively. These indicators are synthesized into a Composite Stability Index (CSI) that provides a normalized measure of repository health proxied by its stability. Finally, the framework introduces several important theoretical properties that validate its usefulness as a measure of repository health and stability. At a conceptual phase and open to debate, our work establishes mathematical criteria for evaluating repository stability and proposes new ways to understand sustainable development practices. The framework bridges control theory concepts with modern collaborative software development, providing a foundation for future empirical validation.

CCS CONCEPTS

• **Software and its engineering** → **Software development process management**; *Software development methods*; *Software version control*; • **Computer systems organization** → Maintainability and maintenance.

KEYWORDS

Stability, repository, development

ACM Reference Format:

Giuseppe Destefanis, Silvia Bartolucci, Daniel Graziotin, Rumyana Neykova, and Marco Ortu. 2025. Introducing Repository Stability. In *33rd ACM International Conference on the Foundations of Software Engineering (FSE Companion '25)*, June 23–28, 2025, Trondheim, Norway. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3696630.3728498>



This work is licensed under a Creative Commons Attribution 4.0 International License. *FSE Companion '25, June 23–28, 2025, Trondheim, Norway*
© 2025 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-1276-0/2025/06
<https://doi.org/10.1145/3696630.3728498>

1 INTRODUCTION

Understanding when a software project is in a healthy state remains a critical yet unsolved challenge in software development. While repositories provide extensive data about project activities, from code changes to community interactions, current approaches struggle to convert this wealth of information into actionable insights about project health [1, 2]. This gap affects both practitioners managing projects and researchers studying software development.

In this paper, we propose a new perspective: viewing project health through the lens of **stability** [3, 4]. We envision repositories as dynamic systems whose health manifests in their capacity to handle disturbances while maintaining consistent development practices.

Drawing from control theory [5], we introduce a framework that reimagines repository analysis. Just as engineering systems naturally seek equilibrium after physical perturbations, we propose that healthy software projects demonstrate stability when facing disruptions - from sudden increases in bug reports to key contributor departures to spikes in feature requests. This perspective opens new possibilities for understanding and maintaining project health and sustainability [6]. Our framework improves repository analysis through four measurable indicators: commit frequency, issue resolution rate, pull request merge rate, and community activity, synthesized into a **Composite Stability Index (CSI)**. This mathematical foundation provides a new perspective in how we evaluate and predict project health, moving beyond traditional statistical approaches to capture the dynamic nature of software development.

The potential impact of our stability-based vision is threefold: (1) it enables systematic, quantitative evaluation of repository health, (2) it paves the way for data-driven project management through clear stability metrics, and (3) it opens new research directions in empirical software engineering. By bridging control theory and repository analysis, we offer a rigorous yet practical approach to a fundamental challenge.

We address the research question: *How can we systematically define and measure repository stability to reflect and predict project sustainability?*

Our goal is to spark a paradigm shift in how we analyze and maintain software projects, establishing stability as a fundamental lens for understanding project health.

2 RELATED WORK

The concept of stability has evolved across multiple domains, from control theory to software engineering [5], yet a unified approach to repository stability remains elusive. While Lyapunov's work [7] and Leigh's contributions [8] established rigorous mathematical foundations for analyzing dynamic systems, their frameworks have not been adapted to the specific challenges of software repositories, which serve as the technical infrastructure where code and related artifacts of software projects are stored and managed.

The distributed systems community, through works like Tabuada et al. [9] and Stankovic [10], advanced our understanding of stability through bounded disturbances and performance metrics. However, these approaches, focused on technical states, have not captured the rich socio-technical dynamics that characterize modern repository evolution as part of broader project ecosystems.

Software engineering has explored stability from various perspectives. Jazayeri et al. [11] pioneered architectural stability measures for evolution, while Yau et al. [12] developed foundational metrics for code stability. Yet these valuable approaches address only isolated aspects of repository dynamics without considering repositories as complete dynamic systems.

In parallel, significant research has addressed project health, which relates to but differs from repository stability. A project involves the full software development effort—people, processes, goals, and resources—while a repository is the technical artifact storing code and development history. The CHAOSS project represents a major initiative for standardizing health metrics for open source communities¹. As Goggins et al. [13] note, open source project health fundamentally concerns "a project's ability to continue to produce quality software," but current approaches struggle to convert repository trace data into actionable insights about overall project health.

Project health has been conceptualized through various lenses: activity-based measures [14], community growth and diversity [15], and operational sustainability [16]. Crowston et al. [14] defined success through metrics related to project output, process quality, and team outcomes, while Daniel et al. [15] emphasized the importance of social diversity in distributed communities.

Jansen [17] expanded the scope beyond individual projects to ecosystem-level health through measures of productivity, robustness, and niche creation. The evaluation of project health often distinguishes between sustainability (the long-term ability to maintain development momentum) and survivability (resilience when facing disruptions). Raja and Tretter [18] introduced a viability index measuring vigor, resilience, and organization to assess a project's capacity to overcome challenges. Despite these advances, as Goggins et al. [13] argue, effective health assessment requires considering comparison, transparency, trajectory, and visualization principles when analyzing repository data.

Traditional repository mining, as demonstrated by Kagdi et al. [19] and Hammad et al. [20], has revealed valuable historical patterns [21]. However, these approaches often treat repositories as static artifacts, missing the opportunity to understand their dynamic, evolving nature as systems that exhibit stability properties analogous to those in control theory.

Salama et al. [4] surveyed stability across software artifacts; their work revealed an important gap: the absence of a unified theoretical foundation for repository stability. Our work aims at filling this gap by integrating control theory principles with repository dynamics, enabling systematic analysis of both structural and behavioral stability through well-defined metrics and thresholds.

While existing frameworks emphasize narratives and contextual understanding across multiple dimensions of project health, our stability-based approach offers a complementary mathematical perspective by applying control theory principles to quantify a repository's ability to maintain equilibrium after experiencing disturbances. This foundation enables systematic evaluation of stability as a key dimension of overall repository health, bridging theoretical rigor with practical utility.

3 THEORETICAL FOUNDATION FROM CONTROL THEORY

In control theory, stability is fundamentally concerned with a system's behavior over time and its response to perturbations. A system is considered stable if, when disturbed from equilibrium, it tends to return to its equilibrium state. More formally, for a dynamical system described by

$$\dot{x} = f(x, t), \quad (1)$$

where x represents the state vector and t represents time, stability is often characterized using Lyapunov stability theory. A system is considered stable if for any $\epsilon > 0$, there exists a $\delta > 0$ such that:

$$\|x(t_0)\| < \delta \implies \|x(t)\| < \epsilon, \quad \forall t \geq t_0. \quad (2)$$

3.1 Repository as a Dynamical System

The conceptualization of a repository as a dynamical system emerges from key observations about software development patterns. Software repositories exhibit continuous evolution through time, with state changes driven by developer interactions, mirroring classical dynamical systems in physics or engineering. These repositories feature complex feedback mechanisms where code changes trigger reviews and subsequent modifications, creating interconnected activity cycles. They encompass both deterministic elements, such as automated workflows, and stochastic components such as varying developer activity patterns, resembling mixed deterministic-stochastic systems. Repositories demonstrate equilibrium-seeking behavior, alternating between periods of intense development and stabilization, analogous to classical dynamical systems. The measurable metrics—commits, issues, pull requests, and branch activities—serve as state variables that evolve according to rules governed by technical and social factors. This conceptualization enables the application of dynamical systems analysis techniques to quantify and understand repository stability.

Let $R(t)$ represent the state of a repository at time t , defined as a vector:

$$R(t) = [c(t), i(t), p(t), a(t)]^T, \quad (3)$$

where:

- $c(t)$: Commit frequency function;

¹<https://chaoss.community/>

- $i(t)$: Issue resolution rate function;
- $p(t)$: Pull request merge rate function;
- $a(t)$: Activity engagement function.

These four components have been specifically chosen based on available repository metrics that represent fundamental dimensions of repository activity and health. The commit frequency function $c(t)$ captures the development momentum and intensity of code changes, providing insights into the project's active development patterns through commit timestamps and frequencies. The issue resolution rate function $i(t)$ reflects the project's ability to handle and resolve problems, calculated through the analysis of issue creation and closure timestamps. The pull request merge rate function $p(t)$ measures the effectiveness of the code review and integration processes, derived from pull request lifecycle data. Finally, the activity engagement function $a(t)$ provides insight into the overall repository engagement through comment activity and interaction patterns, replacing the branch lifetime metric with a more readily available measure of repository vitality. Each component is defined as follows:

Commit Frequency Function.

$$c(t) = \frac{N_c(t, t + \Delta t)}{\Delta t}, \quad (4)$$

where $N_c(t, t + \Delta t)$ represents the number of commits in the time interval $[t, t + \Delta t]$, directly measurable from our commit history data.

Issue Resolution Rate.

$$i(t) = \frac{N_i^{\text{closed}}(t, t + \Delta t)}{N_i^{\text{total}}(t)} \cdot \frac{1}{1 + \bar{T}_{\text{resolution}}(t)}, \quad (5)$$

where N_i^{closed} represents closed issues, N_i^{total} represents total issues, and $\bar{T}_{\text{resolution}}(t)$ is the average resolution time for issues closed in the interval $[t, t + \Delta t]$.

Pull Request Merge Rate.

$$p(t) = \frac{N_p^{\text{merged}}(t, t + \Delta t)}{N_p^{\text{total}}(t)} \cdot \frac{1}{1 + \bar{T}_{\text{review}}(t)}, \quad (6)$$

where N_p^{merged} represents merged pull requests, N_p^{total} represents total pull requests, and $\bar{T}_{\text{review}}(t)$ is the average review time for pull requests in the interval $[t, t + \Delta t]$.

Activity Engagement Function.

$$a(t) = \frac{N_{\text{comments}}(t, t + \Delta t)}{N_{\text{issues}}(t) + N_{\text{prs}}(t)} \cdot \frac{N_{\text{active_users}}(t, t + \Delta t)}{N_{\text{total_users}}(t)} \quad (7)$$

where $N_{\text{comments}}(t, t + \Delta t)$ represents the number of comments in the interval $[t, t + \Delta t]$, $N_{\text{issues}}(t) = N_i^{\text{total}}(t) - N_i^{\text{closed}}(t)$ and $N_{\text{prs}}(t) = N_p^{\text{total}}(t) - N_p^{\text{merged}}(t)$ represent the number of open issues and pull requests at time t respectively, and $N_{\text{active_users}}(t, t + \Delta t)$ represents users who have engaged with the repository through comments, commits, or pull requests in the interval $[t, t + \Delta t]$.

4 REPOSITORY-STABILITY DEFINITION

We define the stability of a repository through a framework that considers both the individual metrics and their interrelationships. A repository's stability is characterized by consistent development patterns, efficient issue resolution, effective pull request management, and sustained community engagement.

Definition 4.1 (Stability). A repository $R(t)$ is considered stable if it satisfies all of the following criteria over an observation period $[t_0, t_0 + T]$.

1. Commit Pattern Stability:

$$\left| \frac{dc(t)}{dt} \right| \leq \alpha_c, \quad \forall t \in [t_0, t_0 + T], \quad (8)$$

where α_c represents the maximum allowable rate of change in commit frequency (the highest acceptable value that the rate of change in commit frequency can have for a repository to still be considered stable). This criterion ensures that development activity maintains a consistent rhythm without extreme fluctuations that could indicate project's instability. The threshold is defined as

$$\alpha_c = \frac{\sigma_{\text{daily commits}}}{\mu_{\text{daily commits}}} \leq 0.5. \quad (9)$$

The coefficient of variation threshold α_c provides a measure of commit frequency stability. We propose $\alpha_c = 0.5$ as an initial value because this would indicate that the standard deviation remains less than half the mean, thus allowing for natural variations, while potentially identifying problematic patterns such as long periods of inactivity followed by a burst of commits. This relative measure accommodates projects of different sizes and activity levels, although empirical validation is needed to confirm its effectiveness.

2. Issue Management Stability:

$$i(t) \geq \beta_i \quad \text{and} \quad \bar{T}_{\text{resolution}}(t) \leq \tau_i, \quad \forall t \in [t_0, t_0 + T], \quad (10)$$

where

- β_i is the minimum acceptable issue resolution rate (proposed value: 0.3);
- τ_i is the maximum acceptable average resolution time (proposed value: 14 days);
- $\bar{T}_{\text{resolution}}(t)$ is the moving average of issue resolution time.

The thresholds for issue management stability reflect considerations in modern software development practices. We propose $\beta_i = 0.3$ for the minimum resolution rate, acknowledging that not all issues require immediate resolution – some may be feature requests, duplicates, or issues that soon become obsolete. For the maximum resolution time, here we consider 14 days, aligning with sprint cycles of two weeks in agile development. These are initial proposed values, but empirical validation is required to determine their effectiveness across different project sizes and development intensities.

3. Pull Request Processing Stability:

$$p(t) \geq \beta_p \quad \text{and} \quad \bar{T}_{\text{review}}(t) \leq \tau_p, \quad \forall t \in [t_0, t_0 + T], \quad (11)$$

where

- β_p is the minimum acceptable pull request merge rate (threshold: 0.4);

- τ_p is the maximum acceptable average review time (threshold: 5 days);
- $\bar{T}_{review}(t)$ is the moving average of pull request review time.

The proposed $\beta_p = 0.4$ pull request merge rate acknowledges that, while some requests warrant rejection due to experimental nature or needed revisions, maintaining an overly low rate may unnecessarily impede contributions. The five-day maximum review period strikes a balance between enabling thorough code evaluation and maintaining development momentum. These initial thresholds aim to optimize between quality control and velocity, though their effectiveness will need to be validated through implementation data.

4. Community Engagement Stability:

$$a(t) \geq \gamma_a \text{ and } \frac{N_{active_users}(t)}{N_{total_users}(t)} \geq \delta_a, \forall t \in [t_0, t_0 + T], \quad (12)$$

where

- γ_a is the minimum acceptable activity ratio (threshold: proposed value 0.25);
- δ_a is the minimum acceptable active user ratio (threshold: proposed value 0.15).

The activity function $a(t)$ here is the same as defined in Section 3.1 (see Eq. 7), measuring both the intensity of interactions through comment ratios and the breadth of community participation through user activity ratios. The thresholds ensure that the repository maintains both sufficient discussion density, relative to open items, and broad community involvement.

4.1 Stability Thresholds

We represent the proposed threshold values for each metric as a threshold matrix:

$$\text{T. Matrix } \Theta = \begin{bmatrix} \alpha_c & \beta_i & \tau_i & \beta_p & \tau_p & \gamma_a & \delta_a \\ 0.5 & 0.3 & 14 & 0.4 & 5 & 0.25 & 0.15 \end{bmatrix}. \quad (13)$$

These initial threshold values represent educated estimations based on the authors' experience with repository analysis and software development processes. The commit pattern threshold ($\alpha_c = 0.5$) allows for natural variations while still identifying erratic behavior. Issue and pull request thresholds balance timely response with thorough processing. Community engagement thresholds aim to ensure broad participation. We acknowledge these are initial proposals that must be empirically validated in future studies.

4.2 Composite Stability Index

To provide a single quantitative measure of stability, we define the Composite Stability Index (CSI) as a weighted sum of normalized stability metrics:

$$CSI(t) = w_c \phi_c(c(t)) + w_i \phi_i(i(t)) + w_p \phi_p(p(t)) + w_a \phi_a(a(t)). \quad (14)$$

The weights reflect the relative importance of each stability component, and their sum must equal 1, with commit pattern stability given slightly higher priority due to its direct reflection of development activity. The weight vector is defined as

$$W = [w_c, w_i, w_p, w_a] = [0.3, 0.25, 0.25, 0.2]. \quad (15)$$

Each component is normalized through a function ϕ_k that maps the raw metrics to a $[0, 1]$ scale, ensuring comparable contributions to the final index:

$$\phi_k(x) = \begin{cases} 1 - \frac{|x - \mu_k|}{\sigma_k} & \text{if } |x - \mu_k| \leq \sigma_k, \\ 0 & \text{otherwise.} \end{cases} \quad (16)$$

The target values μ_k and acceptable deviations σ_k are defined for each component based on their respective thresholds:

- For commit pattern stability (ϕ_c):
 - $\mu_c = 0.25$ (target coefficient of variation);
 - $\sigma_c = 0.25$ (allowing variation up to the 0.5 threshold).
- For issue management stability (ϕ_i):
 - $\mu_i = 0.4$ (target resolution rate);
 - $\sigma_i = 0.1$ (allowing variation down to the 0.3 threshold).
- For pull request stability (ϕ_p):
 - $\mu_p = 0.5$ (target merge rate);
 - $\sigma_p = 0.1$ (allowing variation down to the 0.4 threshold).
- For community engagement stability (ϕ_a):
 - $\mu_a = 0.35$ (target activity ratio);
 - $\sigma_a = 0.1$ (allowing variation down to the 0.25 threshold).

Target values (μ_k) and deviations (σ_k) are set to encourage improvement beyond minimum thresholds, while ensuring that meeting these minimums contributes positively to the CSI. We propose a CSI threshold of 0.7 for overall stability, requiring components to perform consistently above their minimum thresholds, while accommodating temporary fluctuations. This weighted and normalized scoring system enables both monitoring overall repository health and diagnosing specific areas requiring attention when the CSI falls below the threshold.

5 MATHEMATICAL PROPERTIES

The proposed Repository-Stability framework exhibits important theoretical properties that validate its usefulness as a measure of repository health and stability:

1. **Boundedness** The Composite Stability Index (CSI) is bounded by design:

$$0 \leq CSI(t) \leq 1, \forall t \in [t_0, t_0 + T]. \quad (17)$$

This boundedness property ensures that stability measures are normalized and comparable across different repositories and time periods.

2. **Piecewise Continuity** The stability measure should be piecewise continuous with respect to all its component metrics. Within each piece defined by $|x - \mu_k| \leq \sigma_k$, for any metrics m_1, m_2 :

$$|m_1 - m_2| < \delta \implies |CSI(m_1) - CSI(m_2)| < \epsilon. \quad (18)$$

This property ensures that small changes in repository metrics result in proportional changes in the stability measure within each regime.

3. **Long-Term Stability** If a repository's metrics stabilize over time (i.e., the system approaches an equilibrium), the Composite Stability Index (CSI) should converge to a constant value:

$$\lim_{t \rightarrow \infty} CSI(t) = \text{constant}. \quad (19)$$

This property reflects the long-term stability of a repository. In engineering, stable systems often approach equilibrium states as

disturbances diminish over time. For a repository, this could correspond to regular commit patterns, balanced issue resolution and pull request processing, and consistent community engagement.

If each metric $\phi_k(x_k(t))$ stabilizes over time, such that:

$$\lim_{t \rightarrow \infty} \phi_k(x_k(t)) = \phi_k^\infty, \quad (20)$$

then the Composite Stability Index (CSI) converges to:

$$\lim_{t \rightarrow \infty} CSI(t) = \sum_k w_k \phi_k^\infty, \quad (21)$$

where w_k are the weights assigned to each metric. This result shows that the CSI becomes a weighted sum of the asymptotic values of the individual metrics, providing a clear mathematical representation of the long-term behavior of the repository.

Conditions for Convergence: Metrics such as commit frequency, issue resolution rate, and pull request merge rate must exhibit bounded fluctuations or a decaying trend toward equilibrium. For instance, commit frequency $c(t)$ stabilizes if:

$$\frac{dc(t)}{dt} \rightarrow 0 \quad \text{as } t \rightarrow \infty. \quad (22)$$

While theoretical at this stage, this property could be empirically tested by analyzing historical data from repositories over long periods. Stable repositories should display convergence trends in their CSI values, whereas repositories with irregular or unsustainable practices may show divergence or high variability.

Practical Implications: Repositories with erratic behavior (e.g., bursts of activity followed by dormancy) may not exhibit convergence, signaling instability, while stable repositories demonstrate convergence, reflecting healthy long-term development practices. This property provides a diagnostic tool for identifying repositories requiring intervention to achieve sustained stability. When stability metrics consistently fall below their thresholds, automated monitoring systems can identify specific areas requiring attention. For commit pattern instability, the system would highlight development cycle issues; for degrading pull request processing, it would indicate review process inefficiencies. Developers can then implement targeted interventions, which may include restructuring sprint planning to stabilize commit patterns, or redistributing review responsibilities to improve pull request processing. These metric-specific corrections allow teams to address stability issues before they compromise overall project health, transforming theoretical stability measures into practical maintenance tools within existing development workflows.

6 CONCLUSION AND OPEN CHALLENGES

We introduced a theoretical framework for analyzing repository stability through the lens of control theory. The framework's four core components — commit patterns, issue resolution, pull request processing, and community engagement — provide a complete view of repository stability, while remaining computationally tractable. Through the newly introduced Composite Stability Index, we offer a tool to measure and monitor repository health through its stability characteristics. Our framework makes several fundamental assumptions: that repositories exhibit equilibrium-seeking behavior similar to physical systems, that stability can be meaningfully quantified through our chosen metrics, and that these measures

capture the essential dynamics of software development. These assumptions, while grounded in software engineering practices, open critical questions for the research community:

1) *To what extent can universal stability thresholds be established, or do they inherently depend on repository context?* This question builds on prior work by Jansen [17], who observed that metrics may have different meanings for different projects. While our framework proposes initial thresholds (e.g., $\alpha_c = 0.5$ for commit pattern stability), empirical research across diverse repositories is needed to determine whether universal thresholds exist or if contextual adaptation is necessary. A promising hypothesis is that repositories within similar domains (e.g., system libraries vs. web applications) share similar optimal stability thresholds.

2) *Under what conditions does the control theory analogy hold for software repositories, and which differences could be brought by open-source and proprietary software development?* This question extends Salama et al.'s [4] work on stability concepts in software engineering, exploring whether open-source repositories follow different equilibrium patterns than proprietary ones. Factors such as governance structures [18] and corporate engagement [15] likely influence stability dynamics. We hypothesize that repositories with corporate backing exhibit different recovery patterns following disturbances than community-driven projects.

3) *What mathematical models could better represent the complex interactions between stability components while maintaining practical applicability?* Building on Filieri et al.'s [5] application of control theory to software engineering, future work could explore non-linear models that capture emergent interactions between stability components. For instance, the relationship between issue resolution and community engagement might be better modeled through coupled differential equations that capture feedback mechanisms.

4) *How do different socio-technical factors — from team structure to development methodologies to community norms — affect our notion of stability?* This question connects with Goggins et al.'s [13] emphasis on social context in health assessment. We hypothesize that repositories with explicit governance mechanisms (e.g., defined contribution processes) exhibit higher stability in the face of similar disturbances compared to those with ad-hoc structures. Future research could categorize repositories based on these socio-technical factors and analyze their CSI patterns.

5) *How can this theoretical framework inform practical tools for project maintenance while accounting for the social nature of software development?* This bridges theory with practice. Stability metrics could be integrated into project dashboards that provide early warnings of declining stability, with visualizations that will help interpret complex stability data in context. These questions point to rich opportunities in both theoretical and empirical directions — from validating the framework through large-scale repository analysis, to empirically defining threshold values, to investigating non-linear stability metrics — potentially transforming how we understand and evaluate software projects' health.

Future work should directly compare our stability framework with traditional repository health metrics to validate whether our approach captures dynamic aspects of development that static metrics miss. User studies with development teams could reveal how our framework translates from a theoretical model to practical monitoring tools.

REFERENCES

- [1] Slinger Jansen. Measuring the health of open source software ecosystems: Beyond the scope of project health. *Information and Software Technology*, 56(11):1508–1519, 2014.
- [2] Tianpei Xia, Wei Fu, Rui Shu, Rishabh Agrawal, and Tim Menzies. Predicting health indicators for open source projects (using hyperparameter optimization). *Empirical Software Engineering*, 27(6):122, 2022.
- [3] Mohammed E Fayad and Adam Altman. Thinking objectively: an introduction to software stability. *Communications of the ACM*, 44(9):95, 2001.
- [4] Maria Salama, Rami Bahsoon, and Patricia Lago. Stability in software engineering: Survey of the state-of-the-art and research directions. *IEEE Transactions on Software Engineering*, 47(7):1468–1510, 2019.
- [5] Antonio Filieri, Martina Maggio, Konstantinos Angelopoulos, Nicolas d’Ippolito, Ilias Gerostathopoulos, Andreas Berndt Hempel, Henry Hoffmann, Pooyan Jamshidi, Evangelia Kalyvianaki, Cristian Klein, et al. Software engineering meets control theory. In *2015 IEEE/ACM 10th International Symposium on Software Engineering for Adaptive and Self-Managing Systems*, pages 71–82. IEEE, 2015.
- [6] Colin C Venters, Rafael Capilla, Stefanie Betz, Birgit Penzenstadler, Tom Crick, Steve Crouch, Elisa Yumi Nakagawa, Christoph Becker, and Carlos Carrillo. Software sustainability: Research and practice from a software architecture viewpoint. *Journal of Systems and Software*, 138:174–188, 2018.
- [7] Aleksandr Mikhailovich Lyapunov. The general problem of the stability of motion. *International journal of control*, 55(3):531–534, 1992.
- [8] James R Leigh. *Control theory*, volume 64. Iet, 2004.
- [9] Paulo Tabuada, Ayca Balkan, Sina Y Caliskan, Yasser Shoukry, and Rupak Majumdar. Input-output robustness for discrete systems. In *Proceedings of the tenth ACM international conference on Embedded software*, pages 217–226, 2012.
- [10] John A Stankovic. Stability and distributed scheduling algorithms. In *Proceedings of the 1985 ACM thirteenth annual conference on Computer Science*, pages 47–57, 1985.
- [11] Mehdi Jazayeri. On architectural stability and evolution. In *Reliable Software Technologies—Ada-Europe 2002: 7th Ada-Europe International Conference on Reliable Software Technologies Vienna, Austria, June 17–21, 2002 Proceedings 7*, pages 13–23. Springer, 2002.
- [12] Stephen S Yau and James S Collofello. Some stability measures for software maintenance. *IEEE Transactions on Software Engineering*, (6):545–552, 1980.
- [13] Sean Goggins, Kevin Lombard, and Matt Germonprez. Open source community health: Analytical metrics and their corresponding narratives. In *2021 IEEE/ACM 4th International Workshop on Software Health in Projects, Ecosystems and Communities (SoHeal)*, pages 25–33. IEEE, 2021.
- [14] Kevin Crowston, James Howison, and Hala Annabi. Information systems success in free and open source software development: Theory and measures. *Software Process: Improvement and Practice*, 11(2):123–148, 2006.
- [15] Sherae Daniel, Ritu Agarwal, and Katherine J Stewart. The effects of diversity in global, distributed collectives: A study of open source project success. *Information systems research*, 24(2):312–333, 2013.
- [16] Charles M Schweik and Robert C English. *Internet success: a study of open-source software commons*. MIT Press, 2012.
- [17] Slinger Jansen. Measuring the health of open source software ecosystems: Beyond the scope of project health. *Information and Software Technology*, 56(11):1508–1519, 2014.
- [18] Uzma Raja and Marietta J Tretter. Defining and evaluating a measure of open source project survivability. *IEEE Transactions on Software Engineering*, 38(1):163–174, 2012.
- [19] Huzefa Kagdi, Michael L Collard, and Jonathan I Maletic. A survey and taxonomy of approaches for mining software repositories in the context of software evolution. *Journal of software maintenance and evolution: Research and practice*, 19(2):77–131, 2007.
- [20] Maen Hammad, Michael L Collard, and Jonathan I Maletic. Automatically identifying changes that impact code-to-design traceability during evolution. *Software Quality Journal*, 19:35–64, 2011.
- [21] Nicholas M Synovic, Matt Hyatt, Rohan Sethi, Sohini Thota, Shilpika, Allan J Miller, Wenxin Jiang, Emmanuel S Amobi, Austin Pinderski, Konstantin Läufer, et al. Snapshot metrics are not enough: Analyzing software repositories with longitudinal metrics. In *Proceedings of the 37th IEEE/ACM International Conference on Automated Software Engineering*, pages 1–4, 2022.