Review of P. Alvaro and S. Tymon, 2018, "Abstracting the geniuses away from failure testing",

Communications of the ACM **61** (1): 54–61

T.H. Tse
Department of Computer Science
The University of Hong Kong
Pokfulam, Hong Kong

Large-scale distributed systems are difficult to test using traditional failure-testing or fault-injection techniques. Even recent approaches such as chaos engineering rely on experienced experts who can observe the system, propose hypotheses of its behavior, and formulate experiments to validate the results of variations. The process assumes the availability of human expertise, a formal specification, and the source code.

This article presents a lineage-driven fault injection (LDFI) approach that automates the process, starting with successful outcomes and reasoning backward through call-graph traces and data provenance. It was successfully applied at Netflix. I strongly recommend this excellent introductory article if you are new to chaos engineering. It gives enlightening ideas to novices. The writing is smooth and interesting.

If you are a practicing software tester, however, you may want more than just bedtime reading. For example, in order to apply LDFI, we still need an executable specification and a correctness specification, including invariant definitions. The invariants are based on homeostatic states, which are often mistaken as steady states in chaos engineering literature. The former refers to a relatively stable state of equilibrium such as our body temperature of 37 degrees Celsius under normal circumstances, whereas the latter refers to an unvarying condition such as our bodies at room temperature after death. Furthermore, we need to work around nonreplayability and nondeterminism in real-life distributed systems. I suggest that readers refer to Rosenthal et al.[1] for precise details of chaos engineering and Alvaro et al.[2] for technical assumptions and consequences of LDFI.

References

[1] Rosenthal, C., Hochstein, L., Blohowiak, A., Jones, N., and Basiri, A., *Chaos Engineering: Building Confidence in System Behavior through Experiments*, O'Reilly, Sebastopol, CA, 2017.

[2] Alvaro, P., Andrus, K., Sanden, C., Rosenthal, C., Basiri, A., and Hochstein, L., Automating failure testing research at Internet scale, in: *Proceedings of the 7th ACM Symposium on Cloud Computing* (*SoCC '16*), ACM, New York, NY, 2016, pp. 17–28.