# An empirical study of the global behavior of structured overlay networks as complex systems

**Keywords : Complex systems; Overlay networks; Self healing; Resilience; Churn.**

**Ruma Rani Paul, Peter Van Roy**

***Abstract – A paradigm shift is underway in the field of distributed computing systems. From data-center based computing, we are moving toward edge computing, i.e., pushing the computation toward the logical extreme of a network. The inherent unreliability and heterogeneity of edge networks introduce a crucial challenge for distributed systems as the size of the system grows. To survive in inhospitable environments such as edge networks, the system needs to have strong self healing properties. We make empirical studies of both the behavior and design of these systems in inhospitable environments. We define the concept of a reversible system, where the system's functionality depends on the current environment and not on its history. By showing how to make a distributed system reversible, we prepare the way to build provably correct distributed systems for arbitrarily hostile environments.***

Complex systems, such as large-scale distributed systems, consist of many interacting parts. Their behavior cannot be predicted in a straightforward way from the behavior of each part. They have many operating modes depending on their environment and what they are supposed to do. This has resulted in many problems when such systems are stressed beyond where their behavior is a straightforward extrapolation of the behavior of their parts. Applications running on such infrastructures break down when there are too many node or communication failures. Ideally, the application should survive with partial functionality during arbitrary system failures and recover its full functionality when the underlying system is restored. This is not just a fringe case: mobile and ad-hoc networks, for example, have this kind of failure. Even supposedly stable parts of the Internet have peaks of unstable behavior.

In order to design provably correct complex systems, it is required to understand the behavior of such systems in inhospitable environments. We define an *inhospitable environment* as one in which certain stress parameters, e.g., churn (i.e., the rate of node failure and replacement), network partition, and communication delay, can potentially reach high values and temporarily increase without bound. The goal is to build systems that are both predictable (hence, useful in practice) and reversible (hence, they survive) in these environments. We say a system is *reversible* if it regains its original functionality as the stress recedes. Formally, reversibility means that the system's functionality depends only on the current stress and not on its history.

An in-depth study of a complex system in its entire operating space is important for three reasons. First, for practical system design it is important to explore highly stressful environments, since even systems running in so-called "stable" environments also have peaks of high stress. Second, it can open new venues for application design, such as mobile, ad-hoc networks, Internet of Things (IoT), and general edge computing. All these have high node dynamicity and intermittent connectivity, and frequently change network topology. For them, current fault-tolerance techniques are insufficient. Third, it is important for scientific reasons to understand what happens in highly inhospitable regimes.

For our practical study, we have chosen a class of *structured overlay networks* (SONs) as our representative complex systems. SONs are the third generation of peer-to-peer systems (i.e., where each node of the system is both client and server), where a structure is induced through pointers maintained by each peer of the system. SONs are a well-known approach to building decentralized distributed systems. We extend the SON architecture to make it reversible. This can be achieved by carefully designing the maintenance strategies of the SON. A *maintenance strategy* actively maintains the structural integrity of a SON under certain environmental stresses. We organize the maintenance strategies proposed in existing literature along a spectrum between *efficiency* on one side and *resilience* on the other. For example, a reactive local strategy like correction-on-change is much more efficient than gossip, whereas gossip is much more resilient. Our results show that both efficiency and resilience are essential for a system to be reversible in inhospitable environments [1].

In order to describe the qualitative behavior of a SON, we introduce the concept of *phase*. Based on the concept of phase in physical systems [2], we have drawn an analogy for a SON. All nodes in the same phase exhibit the same qualitative properties, which are different for nodes in different phases. We demonstrate the existence of *phase transitions* (i.e., a significant fraction of nodes changes phase) as the stress in operating environment varies. We empirically identify the *critical points* (i.e., when there exists more than one phase simultaneously in significant fractions of the system) observed in our experiments. We propose an API to communicate the current phase of a node to the application layer. Our results show the usefulness of reversibility for designing a SON and of phase for building applications on a SON.

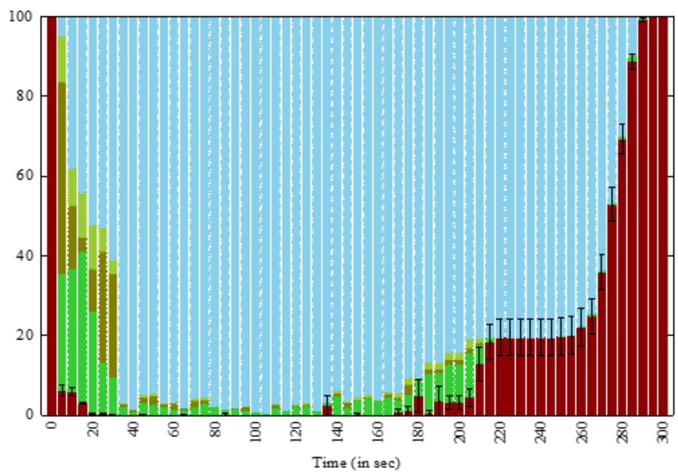


**Figure 1: Phase transitions in a reversible system under increasing and decreasing churn: red, green (different shades correspond to three liquid sub-phases) and blue areas correspond to percentage of nodes in solid (stable), liquid (unstable, loosely connected) and gaseous (isolated) phases respectively.**

## References

[1] Ruma Rani Paul, "An Empirical Study of the Global Behavior of Structured Overlay Networks as Complex Systems", Licentiate Thesis, KTH, Sweden, Oct. 2015.

[2] Wikipedia. Phase (matter). en.wikipedia.org/wiki/Phase_(matter)