Project Summary

This proposal is submitted to the Theory of Computing program, under the direction of Dr. Richard Beigel, within the Theoretical Foundations cluster of the NSF by Andréa W. Richa, Arizona State University. The proposal describes a three year research effort to perform rigorous, theoretical research in the area of self-stabilizing overlay networks.

We are living in the information age. The ability to exchange information quickly and reliably is a vital part of our everyday life. Therefore, it is of utmost importance to provide a communication infrastructure that is highly available and robust.

Due to the rapid rise of social networks and other peer-to-peer systems, sensor networks, and mobile ad hoc networks, logical networks or overlay networks are becoming more and more wide spread. A major complication in these networks is that they are usually dynamic in nature. The topology and the link rates may frequently change, which requires highly efficient update mechanisms to preserve desirable network properties. Minimum requirements for overlay network protocols to be useful in practice are that they be *local*, *simple*, and *self-stabilizing*. Locality is important for fast response times and for minimizing the impact of topology changes on the overlay network properties, simplicity is important so that the protocols can be used in a wide range of systems and for a formal verification of their effectiveness, and self-stabilization is important for automatic recovery from any illegal state since protocols requiring human intervention will not scale to systems potentially spanning millions of sites.

Since peer-to-peer overlay networks are a relatively young area, only a few results are known so far that *rigorously* study self-stabilization issues for overlay networks, though there are many papers that evaluate the fault-tolerance and robustness of overlay network designs through experiments and simulations. Moreover, those few rigorous results that are known do not really provide self-stabilization from *arbitrary* states (from which this would still, in principle, be possible) but only consider certain degraded forms of overlay networks that appear to be most relevant.

Our goal will be to significantly extend the knowledge in this field. We will focus here on overlay networks for the Internet though we expect our results to be also useful for sensor networks and ad hoc networks as we will concentrate on *local, simple* protocols. In this context, we will address the following important questions: Which mechanisms allow overlay networks to self-stabilize from an *arbitrary* connected state in an *efficient* and *robust* way? Which mechanisms allow overlay networks to self-stabilize from an arbitrary state even under *adversarial* behavior of some of the nodes? Which mechanisms allow overlay networks to *preserve* certain properties, even under a high rate of changes (such as join and leave events and faults)? As we will demonstrate, research on these issues is still in its infancy.

Intellectual merit Overlay networks (specially in the form of social networks) are revolutionizing the way we group and exchange information, but not much is known about self-stabilization mechanisms for these networks. This project will introduce transformative approaches to view and design overlay networks under the light of self-stabilization. Since overlay networks and self-stabilization are used in many contexts, we believe that this proposal will be interesting for a number of research communities within and outside of computer science. Our original results will advance the understanding of the subject and close some of our knowledge gaps. Our new algorithms, new techniques, and new methodologies may even be of independent interest to the theoretical computer science community at large. The proposed research is a logical consequence of years of theoretical research of the PI in the area of network communication.

Broader impact. We anticipate that the proposed research will have an impact in several respects, as for instance in: (i) *bridging the gap between theory and practice*, in the sense that our self-stabilizing algorithms will be local and simple enough to have a high impact in practice, with immediate applications to NSF SING and GENI; (ii) *international collaboration*, since we will have fruitful interactions with Prof. Scheideler and the Tech. U. of Munich, Germany; (iii) *multidisciplinary activities*, since communication networks and self-stabilization span many different areas; (iv) *advancing education and enhancing diversity at ASU*.