

# Research Statement

Peyman Faratin  
peyman@mit.edu

## Executive Summary

Protocols are mechanisms that engineers use to commit entities in a system to certain codes of behavior. To date Internet engineers have been fortunate to autonomously make protocol design choices that often embodied most of their technical and social values — ”private provisioned version of contract law”. But public concerns are likely to become increasingly dominant, as the Internet continues to become a critical input into larger number of industries, leading to a state where the design of the Internet may become a multi-laterally negotiated outcome, involving engineers, operators, vendors, regulators, lawyers, anti-trust and advocacy groups. Different stakeholders are becoming interested in the design of the Internet because of not only the social values the Internet embodies but also the many social consequences that have resulted from earlier design choices. For example, the Internet is designed assuming a unicast point-to-point communication between end-hosts whose IP addresses are known by each parties. This implicit assumption about trust among end-hosts has resulted in host of socially undesirable outcomes including spam and other unwanted traffic. Preferentially revealing an end-host identity, a property right of that host, is also costly in today’s architecture. Property rights tensions have also resulted from mixing the resource-addresses of a web object with its Domain Name Server name, creating copyright tensions. At a higher level the architecture has also given incumbent operators, sometimes anti-competitive, discrimination powers. Both houses are currently debating whether blocking, shaping and preferential treatments of applications and IP addresses, functionalities afforded by the current architecture, are anti-competitive discrimination, and whether a “network neutrality” rule is needed. Whether desirable or not, such social behaviors can be directly attributed to the underlying architectural choices that themselves were based on strong social values.

These economic and non-economic behaviors are likely to contribute to an increasing involvement by other stakeholders in the architectural design decisions. Even a *threat* of involvement is sufficient to create incentives to design with other stakeholders’ interest in mind. In either case, architects of tomorrow’s Internet must therefore be cognizant of both the social values they adopt and how their design choices affects outcomes that are of importance to the public at large. We believe that the engineering abstractions offered by the concept of architectures will become even more central information object towards such a goal. However, there is a methodological gap that must be addressed. Multi-stakeholder consensus requires us to be able to evaluate different architectural proposals along *both* technical and social criteria whereas today’s architectural designs forgo the latter and are evaluated only along technical dimensions. There is a real need for a multi-disciplinary science of an architecture that can be clearly evaluated for decision purposes and has clear shared semantics between different stakeholders.

To date a number of uncertainties have however created barriers to such a goal. Uncertainty caused by methodological and conceptual complexities has been the primary barrier for developing a theory of architecture that includes social outcomes because there is no consensus on a) what the relevant outcomes are and b) how design choices affect them. Even technical outcomes themselves are sometimes not enumerated in the design of Internet architectures and the community coordinates itself by following the principle of “rough consensus and running code”. Even if the consensus forming process is “perfect” and result in unambiguous and universally accepted technical and social outcomes many of the technical communities believe that the

consequences of the design choices on outcomes are themselves often uncertain because there exists a complex system of temporal causal interdependencies that percolates through many layers of autonomous entities to result in (desirable/undesirable) outcomes. For these reasons the technical community has at worst ignored these methodological issues and at best reached a position not too dissimilar to evolutionary biologist counter-arguments on the famous “blind watchmaker” problem—a stance on design that requires the architecture to maximally accommodate change and not to pre-commit to any design that implements some static notion of outcome. Generative designs, or templates, are then produced based on these principles.

Social scientists on the other hand face structurally similar types of problems as Internet architects—design of (enforceable) institutions (laws, norms, rules, contracts) that, like protocols, regulate behaviors of strategic entities so as to achieve desirable social outcomes such as efficiency and stability. However, the difference is that the design process is guided by a theory with commitment to individual and social primitives that jointly place a measure on a well defined outcome space. The designer can then not only objectively evaluate institutional design choices through sensitivity analysis, but also build testable empirical models. Note, a subtle point. Uncertainty is constrained to specific primitives and distributed asymmetrically. Designers of institutions are assumed to be uninformed about the preferences of the strategic entities, not the possible set of outcomes. Therefore there is uncertainty on whether the designed institution will achieve the desired outcome, given strategic agents. There is no uncertainty over the outcomes (theory of incomplete contracts does address institutional design when the outcomes themselves are uncertain). These information economics frameworks have been used extensively to help the legal, regulatory and anti-trust authorities in reasoning about both evaluation (mostly by anti-trust and legal professions) as well as design (regulatory rules) choices.

The long-term goal of our research is to show how primitives, models and measures from social sciences, such as economics and law, can help the engineering community to move towards a multi-disciplinary science of architectural design and to better understand the social and individual tradeoffs involved in design choices. We plan to continue our efforts towards this novel and ambitious long-term goal. To this end we will focus on the concrete problem of Internet interconnection, a problem instance that demonstrates the interaction between technical design choices affecting the social outcomes of operators that at times are not desirable. The interconnection problem is an important problem to study because in the Internet, like any distributed supply-chain, competitors must cooperate to deliver a coherent end-to-end service, resulting in a system where networks carry the traffic of other networks until the final destination is reached. Stability of this coordination is therefore critically conditional on the cost-redistribution mechanisms that give the correct transport incentives to competing networks. Without these stable interconnection settlement mechanisms the Internet would not have scaled, in spite of adoption externalities and economies of scope and scale resulting from a standardized IP protocol. We will use the interconnection problem to show how the commitments made in the original Internet architecture, primarily motivated by requirements of the military, have induced a cost-distribution settlements whose stability can be sensitive depending on the strategic choices and gaming options of contracting parties. We plan to work with colleagues at Harvard, MIT and Akamai Technologies to gather data and build explanatory interconnection models that can not only contribute towards a science of architecture design and inform the design of next generation Internet architectures, but also inform economic, policy and legal debates currently being carried out in both houses over interconnection problems.