This paper proposes a new way to determine the business relationships between ASes. The proposed idea collects data from various vantage points and the collected data from each vantage point is used the extrapolate the relationships. Each vantage point's AS paths provide some insight into the relationships between the ASes (via ranking) in the path - using the rankings mined from several vantage points, the algorithm infers the business relationships between the ASes. In additions to inferring business relationships, the paper creates a hierarchy of those ASes based on their relationships with each other. Their ranking system as 5 categories; dense core, transit core, core ,small isp, stub as.

I found the paper to be generally well written, with one flaw. I consider the evaluation technique used to validate their algorithm as being the one flaw. As evaluation, the authors chose to examine the cases where the algorithm failed to accurately classify the AS relationships. While this approach made the evaluation step less time consuming, i do not believe that the method used was enough to truly evaluate the algorithm. In addition to what was done, i would have verified the relations defined by the algorithm against relations publicized by various ASes. This way we could verify both instances when the algorithm appeared to be right and when it appeared to be wrong. Appearing to be right does not always mean it's right; it is very easy for type I paths to become type II paths and still be validate and acceptable paths but there who path would infer different AS relationships one of which would be wrong.

The clean slate argument should not be used to classify the core issues of this paper. This paper is about ways to observe and infer the policies used to govern the underlying Inter-domain mechanism: the clean slate argument argues for the definitions of new mechanism. However, if one argues that the algorithm used to infer and observe policies is itself a mechanism then the core of the paper could be critiqued from a clean slate approach. The AS policy inferring mechanism that existed before this idea used the node degree and a global topology to characterize the relationship and hierarchy of an AS. This paper defers from the norm, by collecting data from various vantage points, using the data to create several topology maps rooted at each vantage point, and using data from the AS paths from the various topology to define relationships.

As stated in the paper, ISPs and stub AS can use the topology and relationship map to determine the economically efficient peering or customer provider relationships but this does not affect end users as their options are usually limited. I am at a loss as to how the relationships between AS can be practically exploited by end users or by developers of network mechanisms. Theoretically, a new routing mechanism could be developed that allow end users to perform AS level source routing with the use of paths that do not violate business policies of the transit ASes. AS level source routing would be an alternative to using the default interdomain routing based on BGP; source routing will use the hierarchy, topology and relationship maps provided by this paper's algorithm, A user will know which ASes are between it's provider and it's destination AS and which paths would satisfy existing relationships. The AS level source routing can allow end hosts to avoid ASes that drop packets, have security issues, or that are hostile to certain types of traffic i.e. illegal p2p apps.