Review of “Active network vision and reality” by David Wetherall

Sreenivasa Pavan Kuppili

Active networks form a novel architecture, in which customized programs can be executed at the various network elements. It gives power to the source node to decide what routines (for example, forwarding routines) are to be executed at various active nodes along the path to the destination. The advantages of an active network are it facilitates easy development and experimentation of new Internet services. The main potential disadvantage is the execution of mobile and untrusted code which leads to serious performance and security issues. ANTS tries to alleviate overhead by using caching and on-demand loading of code. A hash of the code implementation is used as a name for the code, and it also serves as a security check for the integrity of the code. For performance issues, the code is limited to 16KB and is not allowed to run for too long. Besides, for security reasons, the forwarding routine cannnot be changed on the path, and the routine to be executed needs to be certified by a third party authority.

One disadvantage is the limitation on the kinds of services which can be developed. There are various restrcitions on the size and runtime of the code. It might also not be feasible for every router to be an active node, and so the services need to be incrementally deployable. Besides, it is difficult for different routines to share state. The hierarchical fingerprint mechansim suggested for sharing state seems a bit ad hoc and is not easily extensible to more than two routines sharing state. It is also not clear how well caching would work when active nodes are deployed in an actual Interent, especially in the routers handling a large number of flows. The experimenatal results in the paper hardly say anything about the performance under heavy load. There are security issues as well. One routine can potentially change the soft state associated with the same routine of a different flow. The reliance on a third party for certifying code is not a very clean design. Further, as pointed by the authors, end-to-end encryption is a problem for active networks. Further, the incentives for the routers across various ISP’s to become active nodes is not clear.

Active networks may be considered a reasonably good clean slate idea in theory. It is like a natural generalization of the functionality of various network elements. For instance, instead of a single forwarding routine, active networks give the source node the power to choose the particular forwarding routine which is best suited for its application. However, active networks violate the end-to-end principle. It is difficult to monitor what the code does, and even if the code is not malicious, it can just be a drain on the various network resources. Security will always be a major concern with active networks, as mobile code is being executed. There will always be some loophole to be exploited, which would raise the possibility of a coordinated attack to bring down a significant part of the network infrastructure simultaneously.