Gossip-based self-managing services for large scale dynamic networks*

Ozalp Babaoglu, Márk Jelasity, and Alberto Montresor

Università di Bologna
`babaoglu,jelasity,montreso@cs.unibo.it`

Modern IP networks are dynamic, large-scale and heterogeneous. This implies that they are more unpredictable and difficult to maintain and build upon. Implementation and management of decentralized applications that exploit these networks can be enabled only through a set of special middleware services that shield the application from the scale, dynamism and heterogeneity of the environment.

Among others, these services have to provide communication services (routing, multicasting, etc.) and global information like network size, load distribution, etc. The goal is not to provide abstractions that hide the distributedness of the system, but rather, to hide the unpleasant features of the system, such as dynamism, scale and heterogeneity. Most importantly, these services have to be self-managing: they have to be able to maintain certain properties in the face of extreme dynamism of the network. In this manner, such services can serve as the lowest layer that makes possible building more complex applications, or simply as a plugin to enhance existing systems, for example, GRID environments. Apart from self-management, we require that the services be simple and lightweight, to allow easy implementation and incur low cost.

Our approach to achieving these goals is based on the gossip communication model. Gossip protocols are simple, robust and scalable, besides, they can be applied to implement not only information dissemination, but several other functions, as we will show. So far, we have designed gossip-based protocols for maintaining random overlays, which define group membership. Based on this random overlay, we have designed gossip-based protocols to calculate aggregate values such as maxima, average, sum, variance, etc. We have also developed protocols to build several structured overlays in this framework, including superpeer, torus, ring, binary tree, etc. These protocols build on the random overlay and also on aggregate values. The gossip-based model is well suited to dynamic and large networks. Our protocols are extremely simple to implement while being robust and adaptive without adding any extra components or control loops.

Our approach also support composition at a local level. At each node in the network, the same services are available: for example, data aggregation uses the random overlay (peer sampling service) and superpeer topology construction applies aggregate values, such as maximal and average capacity. In fact, proto-

---

cols that implement the different services are heavily interconnected and form a modular system within this lighweight self-managing service layer.

While this presentation focuses on the self-managing systems services, it is clear that other application-level services can also be built at higher layers. These services can be proactive, like load balancing, that can make use of the target (average) load and overlays for optimization of load transfer, or reactive, like broadcasting or search, that can be performed on top of an appropriate overlay network (eg spanning tree or superpeer network), maintained by the lighweight self-managing systems services.