The Peer Sampling Service in Perspective

Rachid Guerraoui
EPFL, Switzerland

Márk Jelasity
University of Szeged, Hungary

Anne-Marie Kermarrec
Inria, France

Maarten van Steen
VU University, The Netherlands

In the early 2000s, the foundations of peer-to-peer (P2P) technology were being layed. Various P2P distributed service abstractions promoted compositionality and helped focus research efforts along well-defined problems. The most well-known such abstraction is the distributed hashtable. Peer sampling joined the set of these abstract services in 2004. The peer sampling service (PSS) targeted the selection of peers to gossip with in each computation round. Basically, a vanilla PSS returns uniform random samples from the set of peers that cooperate in a gossip protocol.

In 2004 we proposed a specification, along with an implementation of a PSS [1]. An extended version is available [2] as well as a short video introducing the PSS [3]. The key features of PSS were scalability and fault tolerance: no full membership was required, nor a stable and reliable membership, not even in in an "eventual" sense. The PSS and its implementation were gossip-based. Nodes had local partial membership views that they kept exchanging and mixing with each other according to various strategies. The implementation leveraged large dynamic random graphs defined by the nodes and the partial views they had. The area of complex network analysis was also taking off about the same time, which offered us useful tools to analyze our networks.

The PSS enabled other researchers to focus on the functionality of a large family of gossip protocols and to implement these in very large and unreliable networks. The PSS not only supports gossip multicast, but also any other protocol based on periodic information exchange among random peers. The PSS promoted an entire family of distributed protocols for many applications including information aggregation, distributed data mining, proximity network formation, and distributed load balancing. These protocols could also be combined to form more complex applications such as video streaming, or recommender systems in file sharing networks, to eventually end up in commercial and open source systems today. Similar ideas were applied in cloud computing, in datacenters for network management tasks, in asynchronous fault tolerant algorithms for high performance computing, or in pervasive computing systems and sensor networks.

Obviously, researchers applying the PSS in their system designs quickly identified some shortcomings challenging the assumptions that the PSS relies on. One example is security, where one would like to get guarantees on the the randomness of the samples in the presence of peers that disrupt the algorithm on purpose. Besides, due to NAT devices nodes cannot freely communicate with each other. Also, correlated failures on the Internet could translate in a bias towards certain nodes. The list goes on. It has been exciting to watch how these problems have been uncovered and then, to a certain extent, solved over the past ten years.

MIDDLEWARE '14, December 08 - 12 2014, Bordeaux, France

http://dx.doi.org/10.1145/2663165.2675740 [dx.doi.org] ...$15.00.

Short Bios.

Rachid Guerraoui is professor of Computer Science at the Swiss Federal Institute of Technology in Lausanne where he leads the Distributed Programming Laboratory. Rachid is fellow of the ACM and is PI of an advanced ERC grant and a Google focused award. He has also been affiliated in the past with the Research Center of Ecole des Mines de Paris, CEA, HP Labs and MIT.

Mark Jelasity is a senior research scientist in the MTA-SZTE Research Group on AI, University of Szeged, Hungary. He obtained his DSc degree in 2014 from the Hungarian Academy of Sciences. In the past he worked at the VU University in Amsterdam, the University of Bologna, and he visited Cornell University in 2013 as a Fulbright scholar.

Anne-Marie Kermarrec is a research director at Inria where she leads a research group on large-scale dynamic distributed systems. Before that she has been with VU University Amsterdam and Microsoft Research Cambridge. She was the PI of an ERC Starting Grant (2008-2013) and an ERC Proof of Concept Grant (2013).

Maarten van Steen is professor and head of the Department of Computer Science at VU University Amsterdam He concentrates on large-scale (wired and wireless) distributed systems. He has been appointed Scientific Director of the Center for Telematics and Information Technology of University Twente from 2015 onwards.

1. REFERENCES

[1] M. Jelasity, R. Guerraoui, A.-M. Kermarrec, and M. van Steen. The peer sampling service: Experimental evaluation of unstructured gossip-based implementations. *Middleware 2004*.

[2] M. Jelasity, S. Voulgaris, R. Guerraoui, A.-M. Kermarrec, and M. van Steen. Gossip-based peer sampling, *ACM TOCS 25(3),2007*.

[3] The Peer Sampling Service: A Primer. http://wandida.com/en/archives/1699