



Brief Announcement: Induced Churn to Face Adversarial Behavior in Peer-to-Peer Systems

Emmanuelle Anceaume, F. Brasiliere, Romaric Ludinard, Bruno Sericola,
Frédéric Tronel

► To cite this version:

Emmanuelle Anceaume, F. Brasiliere, Romaric Ludinard, Bruno Sericola, Frédéric Tronel. Brief Announcement: Induced Churn to Face Adversarial Behavior in Peer-to-Peer Systems. The 11th International Symposium on Stabilization, Safety, and Security of Distributed Systems (SSS 2009), Nov 2009, Lyon, France. pp. 773-774. hal-00420559

HAL Id: hal-00420559

<https://hal.science/hal-00420559>

Submitted on 29 Sep 2009

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Brief Announcement: Induced Churn to Face Adversarial Behavior in Peer-to-Peer Systems

Emmanuelle Anceaume¹, Francisco Brasileiro⁴, Romaric Ludinard^{2,*},
Bruno Sericola², and Frederic Tronel³

¹ CNRS / IRISA, Rennes, France

² INRIA Rennes Bretagne-Atlantique, Rennes, France

³ Supelec, France

⁴ Universidade Federal de Campina Grande, LSD Laboratory, Brazil

Awerbuch and Scheideler [2] have shown that peer-to-peer overlays networks can only survive Byzantine attacks if malicious nodes are not able to predict what will be the topology of the network for a given sequence of join and leave operations. A prerequisite for this condition to hold is to guarantee that nodes identifiers randomness is continuously preserved. However targeted join/leave attacks may quickly endanger the relevance of such an assumption. Inducing churn has been shown to be the other fundamental ingredient to preserve randomness. Several strategies based on these principles have been proposed. Most of them are based on locally induced churn. However either they have been proven incorrect or they involve a too high level of complexity to be practically acceptable [2]. The other ones, based on globally induced churn, enforce limited lifetime for each node in the system. However, these solutions keep the system in an unnecessary hyper-activity, and thus need to impose strict restrictions on nodes joining rate which clearly limit their applicability to open systems.

In this paper we propose to leverage the power of clustering to design a provably correct and practically usable solution that preserves randomness under an ϵ -bounded adversary. Our solution relies on the clusterized version of peer-to-peer overlays combined with a mechanism that allows the enforcement of limited nodes lifetime. Clusterized versions of structured-based overlays [1,3] are such that clusters of nodes substitute nodes at the vertices of the graph. Nodes are grouped together according to some distance function. Our solution is based on the partitioning of the cluster population into two sets, called resp. core and spare sets. Core members are responsible for implementing the overlay. By using classical quorum-based active replication mechanisms among core members, the impact of a minority of malicious nodes is easily masked. The spare set is used to reduce the management overhead caused by the natural churn that is present in typical overlay networks. Nodes in the spare set that leave the system cause next to no overhead, and so do most of the nodes joining the system as they are inserted in the spare set of their corresponding clusters. On the other hand, whenever a core member leaves the system, new core and spare sets are possibly generated.

We propose two strategies to handle these leaves. These strategies mainly differ in the amount of randomisation they impose to introduce the unpredictability

* Supported by the Direction Générale des Entreprises - P2Pim@ges project.

required to deal with attacks. In the first one, a core member that leaves is replaced by a randomly chosen spare member, while in the second one, the departure of a core member leads to the renewal of the whole core set by randomly selecting new core members within both core and spare sets. Then we model each strategy as a game.

The long term behavior of both games is evaluated by using a homogeneous Markov chain $X = \{X_n, n \geq 0\}$ that represents the evolution of the number of malicious nodes in both core and spare sets of a cluster. Both games are played against an adversary whose strength represents the amount of induced churn at a cluster level. In its stronger version, the adversary is free to keep the nodes it manipulates forever in the cluster, while in its weakest form, manipulated nodes are forced to move (they can rejoin later if they wish). The adversary wins the game when process X reaches a set of polluted states from which it can never exit. A state is polluted if the fraction of malicious nodes in the core set exceeds ϵ (i.e., a collusion is mounted). A state that is not polluted is *safe*.

The first step of our work shows that the amount of randomization implemented at cluster level does not prevent a strong adversary from winning in both games in a bounded number of steps, however randomization together with cluster size influence the speed at which this pollution is reached. In particular, the second game alternates, for a random number of steps, between safe and polluted states.

The second step evaluates the benefit of constraining the adversary by limiting the sojourn time of its manipulated nodes in both sets, so that randomness among manipulated and honest nodes is continuously preserved. We first show that none of the games exhibit an absorbing class of states (i.e., both games never ends). Next we prove that process X reaches a stationary distribution which is surprisingly the same for both games. Specifically

Theorem 1. *For both games 1 and 2, the stationary distribution is the same. For all $x = 0, \dots, c$ and $y = 0, \dots, s$, where c (resp. s) represents the upper bound of the core (resp. spare) sets, we have*

$$\lim_{n \rightarrow \infty} \mathbb{P}\{X_n = (x, y)\} = \alpha(x, y),$$

where

$$\alpha(x, y) = \binom{c}{x} \mu^x (1 - \mu)^{c-x} \binom{s}{y} \mu^y (1 - \mu)^{s-y}. \quad (1)$$

References

1. Anceaume, E., Brasileiro, F., Ludinard, R., Ravoaja, A.: Peercube: an hypercube-based P2P overlay robust against collusion and churn. In: Procs. of the IEEE Int'l Conference on Self-Adaptive and Self-Organizing Systems (2008)
2. Awerbuch, B., Scheideler, C.: Towards scalable and robust overlay networks. In: Proceedings of the Int'l Workshop on Peer-to-Peer Systems (2007)
3. Fiat, A., Saia, J., Young, M.: Making chord robust to byzantine attacks. In: Brodal, G.S., Leonardi, S. (eds.) ESA 2005. LNCS, vol. 3669, pp. 803–814. Springer, Heidelberg (2005)