13-th IEEE International Conference on Peer-to-Peer Computing

Towards GROUP protocol formalization

Matteo Mordacchini
IIT-CNR
m.mordacchini@iit.cnr.it

Patrizio Dazzi
ISTI-CNR
p.dazzi@isti.cnr.it

Ranieri Baraglia
ISTI-CNR
r.baraglia@isti.cnr.it

Laura Ricci
University of Pisa
ricci@di.unipi.it

I. INTRODUCTION

Over recent years, we experienced a huge diffusion of internet connected computing devices. As a consequence, this leaded to research for efficient and scalable approaches for managing the burden caused by the highly increased volume of data to be exchanged and processed. Efficient communication protocols are fundamental building blocks for realizing such approaches [1], [2]. Thus, several peer-to-peer protocols have been proposed. Gossip protocols [3]–[7] are a family of peer-to-peer protocols that proved to be well-suited for supporting a scalable and decentralized strategy for peer and data aggregation and diffusion. However, one of the typical limitation of Gossip protocols consists in the selfish behavior adopted by peers in defining their neighborhood and, as a consequence, the topology of the overlay they build. GROUP [8] is a Gossip protocol we conceived to overcome this limitation. It builds explicit defined communities of peers that are identified by their leaders, each one elected in a distributed fashion. This protocol experimentally proved to be efficient and effective with respect to its aim. Anyhow, no analytical study has been realized so far. This work presents a currently ongoing work we are conducting for exploring the properties of GROUP in a more formal way. We conduct this preliminary investigation using a formalization based on Markov chains.

II. GROUP PROTOCOL

The GROUP protocol aims at the election of a set of communities representatives that peers select through a distributed voting phase, driven by the consensus a peer profile can obtain from the other peers. As it happens in democracy, each peer in the system expresses a number of votes that assigns to the peers it considers good representative for itself. With GROUP, each peer elected, together with the peers that contributed to its election, constitutes a community. Once created, such community is identified with the elected peer. The rationale behind this choice is that such peer has been considered a good representative by many peers, hence it would be a good representative for the whole community of peers that contributed to its election. The detection of the representative of each community is structured in the following three stages: i) Potential Candidates Selection, ii) Representative Selection, iii) Community Leader Election. Each phase correspond to a different step in the voting protocol. For each peer, the first two of these phases imply an exchange of information with targeted peers, properly selected among the node's neighbors. The last phase regards the identification of the peer to be elected as community leader. In previous works, the GROUP protocol has been extensively tested using different network topologies and representative selection policies. It demonstrated to be efficient and effective with respect to the aim it has been conceived for. Besides the ability of matching its goals, some additional interesting properties of the protocol can be observed. For instance, its ability of keeping stable the number of peers in each community, independently from the number of peers in the network. Figure 1 depicts the average community size varying the network size from 800 to 2800 peers.

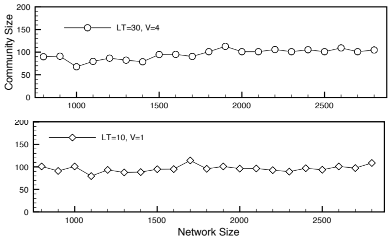

Fig. 1. Average size of communities

III. FORMALIZATION OF GROUP PROTOCOL

As we mentioned, the GROUP protocol is organized in three distinct steps. The output of steps consists, respectively, in the identification of candidates, representatives and, eventually, a leader.

TABLE I. SYMBOLS AND THEIR DEFINITIONS

Symbol	Definition
N	Numb. of nodes in the network
T_{cand}	Number of votes required to became a leader candidate
T_{repr}	Number of votes required to became a leader representative
k	Number of *candidate-votes* that each peer can express
R	Number of *representative-votes* that each peer can express
v	The *view* size of peers
$\alpha,\ \beta$	Aging parameters

As for typical gossip-based solutions (e.g. [9], [10]), the various steps of the GROUP protocol can be formally described by means of correlated Discrete Time Markov chains (DTMC). In the following, we describe the first two steps of GROUP using two correlated Markov chains. The states of the last Markov chain are then used to compute the probability of each node to be elected as a leader of a community.

a) Candidate selection: This is the first step of the protocol. Its goal is to give an initial identification of the peers that can potentially become leaders. During this phase a peer can receive votes from the other nodes that have it in their gossip views. We consider that, at a time t, in order to pass this phase and being considered as a *candidate*, a peer should receive a number of votes equal or greater than a T_{cand} threshold. Consider that D_c is the in-degree of a peer, i.e. the number of nodes that have that peer in their local

978-1-4799-0521-8/13/$31.00 ©2013 IEEE

views, that v is the cardinality of each peer's view, and that each of the D_c nodes can express k candidate votes. Thus, a peer pass this phase **iff** $D_c \frac{k}{v} \geq T_{cand}$. As a consequence, the minimum in-degree needed is then $D_c^{min} = T_{cand} \frac{v}{k}$. Therefore, the probability of passing the candidate selection is given by $P_c = P[indeg \geq D_c^{min}]$. We can define the state transition matrix in the following way:

$$p_{i,j}^c = \begin{cases} P_c & \text{if } i=0, j=1 \\ 1-P_c & \text{if } i=j=0 \\ \alpha & \text{if } i=1, j=0 \\ 1-\alpha & \text{if } i=j=1 \end{cases} \tag{1}$$

where $\alpha \in [0,1]$ is an aging factor for the votes collected till $t-1$. Consider that the initial condition at $t=0$ for the state probability vector π^c is $\pi^c = \{1, 0\}$.

b) Representative Selection: In this second phase, candidates are voted by their neighbors in order to further select the nodes that can be reputed as good representatives of a community by the other peers of the system. Note that π_1^c is the probability of being a candidate, computed using the DTMC of the previous step, D_r is the in-degree of the node and that each of its neighbors has R representative votes that are given to the $\pi_1^c\ v$ candidates present in its view. Similarly to the previous step, to pass this phase, a node should satisfy the condition of $D_r(1-\pi_1^c)\frac{R}{\pi_1^c v} \geq T_{repr}$. Then, the minimum required in-degree is $D_r^{min} = T_{repr}\frac{\pi_1^c v}{(1-\pi_1^c)R}$ and the probability of become a representative is $P_r = P[indeg \geq D_r^{min}]$. The state transition matrix can be defined as:

$$p_{i,j}^r = \begin{cases} \pi_1^c\ P_r & \text{if } i=0, j=1 \\ (1-\pi_1^c) + \pi_1^c(1-P_r) & \text{if } i=j=0 \\ \beta & \text{if } i=1, j=0 \\ 1-\beta & \text{if } i=j=1 \end{cases} \tag{2}$$

where $\beta \in [0,1]$ is an aging parameter similar to one defined at the previous phase. The initial condition for the state probability vector π^r is $\pi^r = \{1, 0\}$.

c) Community Leader Election: The probability of a single node to become leader depends on its probability of being a representative computed at the phase described above. Moreover, we consider that a node is a leader if *at least one of its neighbors* chooses it as its community leader among the $v\ \pi_1^r$ representatives in the neighbor's view. Note that, in order to become leader, we consider that the in-degree of a node should be at least D_r^{min}. The probability E to become a leader is then:

$$\begin{aligned} E &= \pi_1^r \left[\sum_{i=D_r^{min}}^{D_{max}} P[X=i] \left(1 - \prod_{j=1}^{i} \left(1 - \frac{1}{v\ \pi_1^r} \right) \right) \right] = \\ &= \pi_1^r \left[\sum_{i=D_r^{min}}^{D_{max}} P[X=i] \left(1 - \left(1 - \frac{1}{v\ \pi_1^r} \right)^{i} \right) \right] \end{aligned} \tag{3}$$

A. Matching notable properties

As an initial result, we can note that the probabilities at each phase depend on the in-degree probability distribution function. In our system, this function depends on the GROUP underlying gossip protocols, i.e. Cyclon and Vicinity. Note that, for Cyclon, the in-degree distribution depends mainly on the view size [11], rather than factors like the network size. Hence, since the expected number of leaders (and communities) in the network is $N\ E$, the mean number of nodes per community is $\frac{N}{NE} = \frac{1}{E}$. This number does not vary when the network grows, and, thus, this result is in accordance with the experimental findings reported in Fig. 1.

ACKNOWLEDGMENT

This work is funded by EU projects RECOGNITION (no. 257756), EINS (no. 288021), Contrail (no. 257438), InGeoCLOUDS (no. 297300), MIDAS (no. 318786), E-CLOUD (no. 325091), the Italian PRIN 2011 project "Algoritmica delle Reti Sociali Tecno-Mediate" (2013-2014) and the Regional (Tuscany) project SECURE! (FESR PorCreo 2007-2011).

REFERENCES

[1] M. Mordacchini, P. Dazzi, G. Tolomei, R. Baraglia, F. Silvestri, and S. Orlando, "Challenges in designing an interest-based distributed aggregation of users in p2p systems," in *Proceedings of the International Conference on Ultra Modern Telecommunications & Workshops*. Hersonissos, Crete, Greece: IEEE, October 2009, pp. 1–8.

[2] P. Dazzi, P. Felber, L. B. Anh, L. Leonini, M. Mordacchini, R. Perego, M. Rajman, and E. Riviere, "Peer-to-peer clustering of web-browsing users," in *Proceedings of the 7th Workshop on Large-Scale Distributed Systems for Information Retrieval, co-located with ACM SIGIR 2009*, Boston, Massachusetts, USA, July 2009, pp. 71–78.

[3] A. J. Ganesh, A.-M. Kermarrec, and L. Massoulié, "Peer-to-peer membership management for gossip-based protocols," *IEEE Transactions on Computers*, vol. 52, no. 2, pp. 139–149, 2003.

[4] M. Jelasity and A. Montresor, "Epidemic-style proactive aggregation in large overlay networks," in *24th International Conference on Distributed Computing Systems*. IEEE, 2004.

[5] R. Baraglia, P. Dazzi, M. Mordacchini, and L. Ricci, "A peer-to-peer recommender system for self-emerging user communities based on gossip overlays," *Journal of Computer and System Sciences*, vol. 79, no. 2, pp. 291–308, March 2013.

[6] M. Mordacchini, R. Baraglia, P. Dazzi, and L. Ricci, "A p2p recommender system based on gossip overlays (prego)," in *Proceedings of the 10th IEEE International Conference on Computer and Information Technology*. Bradford, UK: IEEE, June 2010, pp. 83–90.

[7] R. Baraglia, P. Dazzi, B. Guidi, and L. Ricci, "Godel: Delaunay overlays in p2p networks via gossip," in *IEEE International Conference on Peer-to-Peer Computing*. Spain: IEEE, September 2012, pp. 1–12.

[8] R. Baraglia, P. Dazzi, M. Mordacchini, L. Ricci, and L. Alessi, "Group: A gossip based building community protocol," in *Smart Spaces and Next Generation Wired/Wireless Networking*, ser. Lecture Notes in Computer Science, S. Balandin, Y. Koucheryavy, and H. Hu, Eds. Springer Berlin / Heidelberg, 2011, vol. 6869, pp. 496–507.

[9] R. Bakhshi, D. Gavidia, W. Fokkink, and M. van Steen, "An analytical model of information dissemination for a gossip-based protocol," *Computer Networks*, vol. 53, no. 13, pp. 2288–2303, 2009.

[10] F. Bonnet, F. Tronel, and S. Voulgaris, "Brief announcement: Performance analysis of cyclon, an inexpensive membership management for unstructured p2p overlays," in *Distributed Computing*. Springer, 2006.

[11] E. Rivière and S. Voulgaris, "Gossip-based networking for internet-scale distributed systems," in *E-Technologies: Transformation in a Connected World*. Springer, 2011, pp. 253–284.

[12] E. Carlini, M. Coppola, P. Dazzi, D. Laforenza, S. Martinelli, and L. Ricci, "Service and resource discovery supports over p2p overlays," in *Proceedings of the International Conference on Ultra Modern Telecommunications & Workshops*. Hersonissos, Crete, Greece: IEEE, October 2009, pp. 1–8.

[13] P. Dazzi, M. Mordacchini, and F. Baglini, "Experiences with complex user profiles for approximate p2p community matching," in *Proceedings of the 11th IEEE International Conference on Computer and Information Technology*. Pafos, Cyprus: IEEE, 2011, pp. 53–58.