

SIMPLE GOSSIPING WITH BALLS AND BINS

Boris Koldehofe *

Abstract: Recent research suggests decentralised probabilistic protocols on support for multipeer communication. The protocols scale well and impose an even load on the system. They provide statistical guarantees for the reliability, i.e. an information sent from an arbitrary source will reach all its destinations. Analysing the reliability is based on modelling the propagation of events as an epidemic process often referred to as gossiping or rumour spreading.

This work provides a new method for analysing such protocols, by representing the propagation of information as a balls and bins game. The method gives a simple relation between the number of hops a gossip message is propagated and the reliability provided. This way it can facilitate the analysis of the multiple delivery problem i.e. to prevent multiple deliveries of the same message to the application layer. By introducing a new protocol it is shown how existing approaches can be adapted to the balls and bins approach. Furthermore, the proposed method is applied to analyse the performance of this protocol.

Keywords: Multipeer Communication, Gossiping, Analytical Evaluation.

1 Introduction

Multipeer communication applies to all scenarios where multiple senders and multiple receivers, associated with physically distributed processes, communicate information of common interest. Processes sharing a common interest form a group, in which they exchange messages on information. Every piece of information is embedded in an *event*, which processes deliver to the application layer of the protocol. Supporting collaboration for large groups of users can be seen as one application relying on multipeer communication. Group members require a fast and reliable propagation of events within such an environment. Moreover, the group needs mechanisms that allow an efficient handling of membership, i.e. ongoing joining and leaving of group members should not add a significant overhead to the group communication.

General aspects of interest in group communication concern the *reliability* of the communication and the *scalability* of the protocol. Reliable group communication protocols provide guarantees that group events reach all their destinations. They may also be required to give ordering guarantees, which are needed in many cases where shared data objects are modified with respect to some consistency model. Reliable group communication protocols must also be robust against faults and bursts in the communication traffic. Apart from providing guarantees that an event will be delivered, some models are also concerned with prohibiting multiple deliveries of events to the application. Protocols that lack this property are said to suffer from the *multiple delivery* problem. The *scalability* of a group communication protocol is determined by how many group members can be supported and which reliability mechanisms are provided.

A common categorisation used in [WZ99] divides protocols supporting group communication services in three different classes.

Reliable group services guarantee that if an event is delivered to one destination it will eventually be delivered to all operating destinations. Further, ordering guarantees for events are provided. However, these requirements imply a high synchronisation overhead for such protocols. The high communication cost let those protocols scale only to a low number of processes.

Unreliable group services focus on best effort strategies which scale well. Under ideal conditions some protocols even give reliability guarantees. However, under perturbed conditions the protocols can show unpredictable behaviour and thus cannot be used for time critical applications.

The focus of this paper is on group services that provide *predictable reliability guarantees* which scale well and hold even under perturbed conditions. In particular, we focus on those protocols which are based on *gossiping* or *rumour spreading* where processes exchange in every round messages on events according to a probabilistic scheme. By processes communicating with random destinations, hot spots are avoided, compared to the typically hierarchical flow of events within reliable group services. In the context of gossiping protocols, an event is also called a *rumour*. This notation is commonly used in the analysis of gossiping protocols.

*Department of Computing Science, Chalmers University of Technology, S-412 96 Göteborg, Sweden
email:khofer@cs.chalmers.se

Organisation of this paper. Section 2 introduces the *epidemic model*, which is the mathematical foundation for most existing gossiping protocols, followed by a presentation of known theoretical results for this model. Further, an overview is presented on related protocols based on gossiping in the area of multi-peer communication. Some of the protocols use an approximation of the theoretical results presented. In the next sections it is shown how to modify the epidemic model so that the balls and bins analysis can be applied. A definition of a *balls-and-bins-compliant gossiping protocol* is introduced in section 3, where the balls and bins approach is described. In section 4 a new protocol with this property is described. The framework of the protocol is based on recently proposed protocols for this problem, such as *pbcast* [BHO⁺99] and *lpbcast* [EGH⁺01]. Moreover, section 5 presents an analysis based on the balls and bins model. Hereby, the reliability of the protocol is determined as a statement with high probability that all processes will receive an event propagated within such a system after gossiping this event for a fixed number of rounds. The event can originate from any arbitrary process. Section 6 discusses the applicability of the given approach in the context of the *multiple delivery problem*.

Remarks on the notation. Throughout the paper \log and \ln denote the logarithms in basis of 2 and basis of e , respectively. For any other basis, say a , the logarithmic function is expressed as \log_a . In the analysis the expression “with high probability”, i.e. with probability $1 - O(n^{-k})$ for a constant $k > 0$, is abbreviated by writing w.h.p.

2 Related work and models

The epidemic model. The idea of applying the mathematical theory of general epidemics [Bai75] to rumour spreading in distributed systems was first adapted in [DGH⁺87]. In this approach, the protocol models the spreading of a rumour as an epidemic process, with the purpose to infect as many other processes as possible. Processes can be

1. *susceptible* to the rumour, i.e. they have not received the rumour yet,
2. *infected* by the rumour, i.e. the rumour is received and processes are spreading this rumour to other processes,
3. *resistant*, i.e. the process has received the rumour, but does not participate in the propagation of the rumour.

Let s, i, r denote the fraction of susceptible, infected and resistant processes at time t such that $s + i + r = 1$. Assume that within the time fraction Δt the infectious rate denoted by α is proportional to the number of infected processes and to the number of susceptible processes, respectively. Then, the fraction of new infection is given by $\alpha i s \Delta t$. On the other hand, the constant removal rate denoted by β needs to be considered, such that the fraction of new resistant processes is given by $\beta i \Delta t$. The resulting differential equations are given by

$$\begin{aligned}\frac{\delta s}{\delta t} &= -\alpha s i \\ \frac{\delta i}{\delta t} &= \alpha s i - \beta i \\ \frac{\delta r}{\delta t} &= \beta i.\end{aligned}$$

By $i(t) = i(t(s))$ we can express i as a function of s . Hence,

$$\begin{aligned}\frac{\delta i}{\delta s} &= \frac{\alpha s i - \beta i}{-\alpha s i} \\ &= -1 + \frac{\beta}{\alpha} \frac{1}{s}.\end{aligned}$$

Let s_0 denote the initial fraction of susceptible processes while i_0 denotes the initial fraction of infected processes. The anti-derivative for $-1 + \frac{\beta}{\alpha} \frac{1}{s}$ is given by $-s + \frac{\beta}{\alpha} \ln s + C$. Applying $i_0 = -s_0 + \frac{\beta}{\alpha} \ln s_0 + C$, we obtain

$$i(s) = i_0 + s_0 - s + \frac{\beta}{\alpha} \ln \frac{s}{s_0}.$$

The maximum number of infected processes at the same time is given by $i(\frac{\beta}{\alpha})$. The term $\frac{\beta}{\alpha}$ also represents the threshold for the ratio of susceptible processes needed to start an epidemic. The epidemic process terminates if $i(s) = 0$, which is the case if

$$s = s_0 e^{\frac{\alpha}{\beta}(s - i_0 - s_0)}.$$

Initially, only one process knows about the rumour and $s_0 + i_0 = 1$. Thus for large groups,

$$s \approx e^{-\frac{\alpha}{\beta}(1-s)}.$$

Theoretical bounds. Theoretical work presents results on how to bound the number of rounds and the number of transmissions needed to infect every process with high probability.

B. Pittel [Pit87] analyses a simplified model, in which processes never lose interest, i.e. the removal rate is zero. Let n denote the group size and assume that an infected process informs only one other process in a round. In probability the number of rounds needed to infect all group members is determined by

$$\log n + \log_{10} n + O(1) \quad \text{as } n \rightarrow \infty$$

Karp et al. [KSSV00] introduce a combined push and pull scheme in which processes cease to push a rumour when the expansion of the rumour has exceeded a certain threshold. Since the basic scheme is based on an exact estimation on the number of informed processes, an algorithm for terminating the spreading of the rumour is introduced. The complete scheme is address-oblivious, i.e. it does not depend on the state of neighbour processes. The number of transmissions of the scheme is $O(n \log \log n)$. Further, it is shown that any such scheme needs at least $\Omega(n \log \log n)$ transmissions independent of the number of rounds.

Rumour spreading in group communication. Rumour spreading (gossiping) has received a lot of attention in group communication, especially after the introduction of *pbcast* [BHO⁺99]. The arguments in favour of rumour spreading are that those protocols introduce an even load on the system and are robust against bursts in the network traffic. Further, they usually provide mechanisms that support fault tolerance and scale better than deterministic approaches. The reliability is based on an approximation of the epidemic model, in which each event associated with a rumour will be spread to all other processes with high probability. However, approaches like *pbcast* rely on global knowledge of all group members and often on a centralised organisation of membership. This way the organisation of a dynamic membership, i.e. processes can join a group or leave a group, requires a lot of synchronisation overhead and a lot of memory resources affecting the scalability. In contrast, recent work [GKM01b, EGH⁺01] suggests protocols that support decentralised handling of membership where processes only have a partial view on the system.

In *lightweight probabilistic broadcast (lpbcast)* [EGH⁺01] the size of the view, i.e. the set of group members known to a process, as well as the size of buffers storing recent events, are statically fixed. When propagating an event, a process communicates only with a fixed number of processes randomly chosen from the local view. The reliability is derived by an approximation of the theoretical result presented in [Pit87]. With respect to the approach introduced in this paper, it provides the interesting property that the view converges to a situation where it is as if members of the local view have been selected uniformly randomly among all group members.

Similar to *lpbcast SCAMP* [GKM01b] presents a scheme supporting decentralised group services. However, instead of using a fixed view size, the membership scheme takes care that the view automatically adapts depending on the group size. The local view size converges to approximately $\log n + C + 1$ where n denotes the group size and C is a constant determining the fault tolerant behaviour of the protocol. The reliability of the event transmission is related to the connectivity in the random graph model and is analysed in [GKM01a].

3 The balls and bins model

The occupancy problem considers a random allocation of m indistinguishable objects denoted by balls on n destinations denoted as bins. Many problems in computer science can be modelled using the occupancy problem. For example, in distributed computing the random allocation problem has received a lot of attention. One of the many issues in the analysis is to bound the maximum number of balls that can fall into a single bin.

In the setting of spreading a rumour among n processes, we associate each process with one bin. Let $H_n = \sum_{k=1}^n \frac{1}{k}$ denote the n th harmonic number. If the rumour is carried within balls which are randomly placed into bins, it is known that nH_n balls are expected to be placed until every process has received at least one ball. To obtain a statement with high probability, i.e. with probability at least $1 - n^{-k}$ for some positive constant k , we look at the random variable X denoting the number of empty bins and the respective expectation value

$$\mathbf{E}[X] = n \left(1 - \frac{1}{n}\right)^m.$$

Since for $n \geq 1$ the inequality $\left(1 - \frac{1}{n}\right)^n \leq e^{-1}$ holds, the expectation value can be bounded by

$$\mathbf{E}[X] \leq ne^{-m/n}.$$

Indeed, we can also write $\mathbf{E}[X] \sim ne^{-m/n}$. By choosing m and k such that $m = cn \log n$ and $k \geq c + 1$ for a constant $c > 1$, one can derive $\mathbf{E}[X] \leq n^{-k}$. From Markov's inequality it is known that $\Pr[X \geq 1] \leq \mathbf{E}[X]$, which gives

$$\Pr[X = 0] \geq 1 - n^{-k}.$$




Figure 1: A scheme for spreading a rumour that terminates after r rounds and guarantees that at least $cn \log n$ balls are placed into random bins.

This reasoning, summarised in theorem 1, is the basis for the approach of analysing the rumour spreading as a balls and bins game.

Theorem 1 *Let $c > 1$ be a constant and let m denote the number of balls that are placed into n bins chosen uniformly and independently at random. If $m \geq cn \log n$ then every bin contains at least one ball with high probability.*

Now, the balls and bins model is introduced by defining what a *balls-and-bins-compliant gossiping protocol* is. Hereby, it is required to send a rumour to $O(n \log n)$ destinations which are uniformly and independently chosen at random. Since balls representing a rumour can only be placed by processes that have received the rumour before and initially only one process knows the rumour, it is allowed that the rumour is propagated in multiple rounds. This way, processes which received balls in previous rounds can create themselves new balls. As long as balls are placed in each round independently of the round the ball was created, the protocol can terminate after overall $O(n \log n)$ balls have been placed. The resulting scheme of a balls-and-bins-compliant protocol, as defined in definition 1, is illustrated in Figure 1.

Definition 1 *A protocol for spreading a rumour is said to be balls-and-bins-compliant if for some fixed integer $r > 0$ and a constant $c > 1$*

- *the protocol provides a scheme that creates $m_i > 0$ balls associated with the rumour in consecutive rounds $1 \leq i \leq r$,*
- *the protocol guarantees that $\sum_{i=1}^r m_i \geq cn \log n$,*
- *the destination of each ball is chosen uniformly at random and independently of the destination of the other balls,*
- *and the protocol terminates after r rounds, i.e. $m_i = 0$ for all $i > r$.*

Corollary 2 Any balls-and-bins-compliant protocol terminating after r rounds informs all processes about a rumour with high probability.

With respect to a rumour and a set of balls associated with the rumour, in this scheme a process is

- *susceptible* when no ball has fallen into the bin associated with the process for the first time,
- *infectious* whenever it creates balls that are placed randomly into some bins,
- and *resistant* when the process knows about the rumour, but does not create balls.

Compared to the state description of the epidemic model shown, the description based on balls and bins is more flexible regarding the state changes of a process. The balls and bins model allows processes to change several times between being infectious and being resistant without affecting the underlying analysis.

In the next section a scheme implementing a balls-and-bins-compliant protocol is introduced. Further, the number of rounds the protocol needs to terminate is analysed for a fixed fanout value.

4 A balls-and-bins-compliant protocol

The balls-and-bins compliant protocol obtained in this section considers a framework which also applies to epidemic protocols. Note that such a framework is not only intended to propagate new events fast, but must also be able to organise the membership of a group, i.e. it must deal with members joining and leaving the group.

Hereby, a group $G = \{p_1, \dots, p_n\}$ is defined as a set of processes. Processes exchange in every round *gossip* messages where each gossip message includes information about members that joined the group, members that left the group, and events to be delivered to the application layer of the gossip protocol. In order to keep track of possible communication partners, each process p_i maintains a view $V_i \subset G$. Let K denote the fanout, i.e. the number of communication partners chosen in a round. Each process p_i communicates with K communication partners chosen uniformly at random from its view V_i . The framework of the protocol is based on every process doing the following computation in every synchronous round (see also the basic protocol in Figure 2):

1. evaluate received gossip messages by processing joining and leaving members and events
2. create a new gossip message and send it to K randomly chosen neighbours known from the view

For the later analysis it may be helpful to keep in mind that the destination of a gossip message will determine the destination of a ball associated with an event.

The first part of the framework is concerned with the maintenance of the local view of processes. The evaluation of members which either join the group or leave the group determines the new view of the process. In the following we assume that the view of a process allows the selection of K randomly chosen elements from the whole group. If the process maintains only a partial view, we consider only schemes as [GKM01b, EGH⁺01] where the views converge to a uniform random distribution.

The second part of the framework deals with the propagation of events such that every process can receive the event. Therefore, we provide a scheme in order to place sufficiently many balls associated with an event into distinct bins. By using a set denoted by *newevents*, a process keeps track of the events which have been delivered in the current round. Another set denoted by *history* contains the events which have been delivered to the application layer. An event e with $e \in \text{newevents}$ and $e \notin \text{history}$ will be added to *history* and delivered to the application layer.

In order to provide termination for the propagation of an event, each event is associated with a *tag value*, which indicates the number of rounds the event has been propagated. A new gossip message contains, in addition to the information on members joining and leaving the group, all events $e \in \text{newevents}$ whose tag value is smaller than a parameter r , indicating the maximum number of rounds an event needs to be propagated.

The complete scheme including the basic protocol and the evaluation of events is illustrated in Figure 2. Only the evaluation of members joining and leaving the group is omitted since the protocol is intended to serve many different ways to organise the membership.

Now, one can look at the spreading of each event as a separate balls and bins game, where the placing of the balls is determined by the random destination of the gossip message. Whenever a process receives an event which has stayed less than r rounds in the system, the process will create K balls, which are sent with K different gossip messages to their destinations. Hence, by choosing $r = cn \log n$ for $c > 1$ we can derive theorem 3.

Theorem 3 Let K denote the fanout and let r denote the maximum number of rounds an event is propagated with respect to the provided gossiping scheme. Then, for any $K \geq 1$ there exists a fixed r such that the provided scheme becomes balls-and-bins-compliant.

Basic Protocol:

```

for all rounds in which  $p_i \in G$  do
  evaluate received gossip messages by processing joining and leaving members and events
  create a new gossip message  $M$ 
  for  $l = 1$  to  $K$  do
    choose  $j \in V_i$  uniformly at random
    send  $M$  to  $p_j$ 
  end for
end for

```

Evaluation of events:

```

for all events  $e$  received in a round do
  if  $e \notin \text{history}$  then
    deliver( $e$ )
  end if
  tag = gettag( $e$ )
  if tag <  $r$  then
     $M = M \cup (e, \text{tag} + 1)$  /* add the event to the gossip message  $M$  */
  end if
end for

```

Figure 2: The balls-and-bins compliant protocol consists out of the basic protocol and the evaluation of events. It can be adapted to various schemes determining the membership of a group.

Note that the introduced scheme makes explicit use of the possibility to allow processes to change states between being infectious and being resistant several times. In the following section a performance analysis is presented, which bounds the value of r depending on the fanout K .

5 Performance analysis

The focus of this analysis is to bound the number of rounds indicated by parameter r such that the introduced balls-and-bins-compliant protocol of the previous section can be applied for practical purposes. The reliability of the protocol is expressed as a guarantee that an event created by a process will reach all other processes with high probability. As a main result it is shown that for a small fanout $K = \lceil 2e \ln n / \ln \ln n \rceil$, it is sufficient to choose $r = O(\log n)$ to guarantee the delivery of an event to all destinations w.h.p. Further, an analysis for constant fanout proves that $r = O(\sqrt{n})$ is sufficient.

The analysis is based on the balls and bins model introduced in section 3 and uses in particular corollary 2. We must ensure that within r rounds sufficiently many balls associated with an event are created. This is achieved by examining the number of balls which are placed within a round i denoted by m_i and determining r such that

$$\sum_{i=1}^r m_i \geq cn \log n \quad \text{w.h.p.}$$

Hereby, we aim to relate the number of balls which are placed into random bins within two consecutive rounds, i.e. we are looking for a statement such that for a positive constant $d > 1$

$$m_{i+1} \geq d \cdot m_i \quad \text{w.h.p.} \quad (1)$$

From such a statement it is shown how one can derive the number of rounds until a minimum number of balls is created in every round, where the number of rounds needed is required to be logarithmic in the minimum number of balls. Therefore, similar to the random allocation problem we look into the maximum number of balls that may fall into any single bin.

Theorem 4 *Let m denote the number of balls, which are placed into n bins chosen uniformly and independently at random. If $m \leq n$ holds, then every bin contains no more than $e \frac{\ln n}{\ln \ln n}$ balls with probability $\frac{1}{n}$.*

The proof is not presented since a detailed proof of this theorem is given in [MR95] (see pages 44 ff.). A generalisation of theorem 4 can be found in [RS98].

Let now m_i denote the number of balls placed into random bins in the i th round. Applying theorem 4 leads to the following corollary, which gives an estimate of how much time is needed until we can expect at least n balls to be placed into bins within a round.

Corollary 5 *If $K \geq 2e \frac{\ln n}{\ln \ln n}$ and $m_i \leq n$, then $m_i \geq 2m_{i-1}$ with probability $1 - \frac{1}{n}$.*

Finally, from theorem 6 we can conclude that the protocol needs $O(\log n)$ rounds to terminate.

Theorem 6 *Let K denote the fanout of the presented balls-and-bins compliant protocol. If $K \geq 2e \frac{\ln n}{\ln \ln n}$, then after $O(\log n)$ rounds every process of the group is informed about an event w.h.p.*

Proof. We show that after $O(\log n)$ rounds an overall of $cn \log n$ balls have been placed into bins.

Let ξ_i denote the event that the number of balls did not double in round i for which $m \leq n$. Then, by corollary 5 it follows that

$$\Pr[\xi_i] \leq \frac{1}{n}.$$

Therefore, the occurrence of any event can be expressed as the union of $\log n$ events ξ_i :

$$\Pr[\cup_{i=1}^{\log n} \xi_i] \leq \sum_{i=1}^{\log n} \xi_i \leq n^{-1/2}$$

and thus after $\log n$ rounds the number of balls placed into random bins in a round is at least n with probability $1 - n^{-1/2}$.

Using the same argument for the next $c \log n$ rounds, the number of created balls in a round will not decrease below n with probability $1 - cn^{-1/2}$.

Hence, with probability $1 - o(1)$, $cn \log n$ balls have been created and placed into random bins after $(c + 1) \log n$ rounds. \square

Note that $\ln n / \ln \ln n$ grows very slowly, so that the fanout value can be regarded as sufficiently small for all practical purposes. Further, the maximum number of messages a process receives in a round is expected to be smaller than $O(\frac{\log n}{\ln \ln n})$ since the number of balls which are created within one round is clearly bounded by Kn , and placing $O(n \log n)$ balls randomly into bins gives a maximum of $O(\log n)$ balls in each single bin.

If we assume that the fanout K is constant, one needs $\Omega(\log n)$ rounds as we cannot expect to place more than Kn balls in a round. Further, as stated in theorem 8, it is shown in the following that $O(\sqrt{n})$ rounds are sufficient. To determine the result, we examine once again for which values it is possible to achieve a statement as given by equation 1.

Lemma 7 *If $m \leq \sqrt{n} \log_{20}(n)$ holds, then every bin contains no more than 6 balls with probability $\frac{1}{n}$.*

Proof. For $m = \sqrt{n} \log_{20}(n)$ the probability that bin i receives exactly j balls is given by

$$\begin{aligned} & \binom{m}{j} \left(\frac{1}{n}\right)^j \left(1 - \frac{1}{n}\right)^{m-j} \\ & \leq \left(\frac{em}{jn}\right)^j \\ & = \left(\frac{e \log_{20}(n)}{j\sqrt{n}}\right)^j. \end{aligned}$$

Let $\xi_i(k, m)$ denote the event that bin i has k or more balls in a bin after m balls have been tossed. Then for $k \geq 6$,

$$\begin{aligned} \Pr[\xi_i(k, m)] & \leq \sum_{j=k}^n \left(\frac{e \log_{20}(n)}{j\sqrt{n}}\right)^j \\ & \leq \left(\frac{e \log_{20}(n)}{k\sqrt{n}}\right)^k \frac{1}{1 - \left(\frac{e \log_{20}(n)}{k\sqrt{n}}\right)} \\ & \leq 2 \left(\frac{e \log_{20}(n)}{k\sqrt{n}}\right)^k. \end{aligned}$$

In particular, for the event $\xi_i(6, m)$ we obtain

$$\Pr[\xi_i(6, m)] \leq 2 \left(\frac{e \log_{20}(n)}{6\sqrt{n}} \right)^6.$$

By applying $\log_{20}(n) \leq n^{1/3}$, it follows that

$$\Pr[\xi_i(6, m)] \leq n^{-2}.$$

Taking the union of the events $\xi_i(6, m)$, for $i = 1, \dots, n$, yields the desired result. \square

Using the same proof technique as for theorem 6 one can derive theorem 8.

Theorem 8 *Let K denote the fanout of the presented balls-and-bins compliant protocol. If $K \geq 12$ holds, the protocol needs $O(\sqrt{n})$ rounds in order to inform every process of the group about an event w.h.p.*

6 Discussion

Although the epidemic model has been shown to be useful to estimate the performance of various protocols, it is a hard problem to obtain a precise bound if one assumes limitations on the resources available to processes. As motivated by *lpcast* [EGH⁺01] and *SCAMP* [GKM01b] the next generation of protocols will exactly deal with these requirements.

To cope with limited resources, the epidemic model needs to consider the possibility that a resistant process may become susceptible again. Although in general this scenario is undesirable, it occurs when protocols limit the size of the history buffer. The history buffer itself determines whether a process can stay resistant against an event. If a process deletes an event too early from its history buffer it is susceptible again and potentially delivers the event multiple times to the application layer. In this case the protocol is said to suffer from the *multiple delivery problem*. Even worse an epidemic process which almost stopped could be started once over again, thus preventing the epidemic process to terminate. In fact, simulation results confirm that even a small infection rate of new events can lead to a significant overhead.

This scenario requires to consider more differential equations and needs a far more complex analysis. Therefore, the motivation for this work was to provide a simpler model as an alternative. The analysis as presented in section 5 is based on elementary mathematics and can this way facilitate the buffer analysis by examining the collisions of balls containing different events.

7 Conclusion

This work has shown an alternative approach of analysing gossiping based on the occupancy problem. With respect to the balls and bins model, a new protocol is presented based on existing practical approaches [EGH⁺01, BHO⁺99]. The maintenance of the group membership can be adapted to various schemes which allow processes to select uniformly random group members from its view. The reliability of the protocol is proven as a statement with high probability that an event, created by an arbitrary process, will reach all group members. The presented analysis shows that an event needs to remain for $O(\log n)$ rounds in the system if the fanout is $\lceil 2e \ln n / \ln \ln n \rceil$, and $O(\sqrt{n})$ rounds if a constant fanout is used.

The future work focuses on analysing the tightness of the presented bounds. Moreover, an analysis on the average number of gossip messages imposed by the spreading of an event is needed. Note that only a lower bound can be derived from the fact that at least an overall of $cn \log n$ messages are needed in order to ensure that every process receives the event with high probability. Finally, work in progress looks into ways to manage buffer space so as to improve the multiple delivery problem discussed in section 6.

Acknowledgements

Thanks a lot to Patrick Thomas Eugster, Anders Gidenstam, Marina Papatriantafilou and Philippas Tsigas for discussions and comments that contributed to this work.

References

- [Bai75] Norman T. J. Bailey. *The Mathematical Theory of Infectious Diseases and its Applications*. Griffin, 2nd edition, 1975.
- [BHO⁺99] Kenneth P. Birman, Mark Hayden, Ozgur Ozkasap, Zhen Xiao, Mihai Budiu, and Yaron Minsky. Bimodal multicast. *ACM Transactions on Computer Systems*, 17(2):41–88, May 1999.
- [DGH⁺87] Alan Demers, Dan Greene, Carl Hauser, Wes Irish, John Larson, Scott Shenker, Howard Sturgis, Dan Swinehart, and Doug Terry. Epidemic algorithms for replicated database maintenance. In *Proceedings of the 6th Annual ACM Symposium on Principles of Distributed Computing*, pages 1–12, Vancouver, BC, Canada, 1987. ACM Press.
- [EGH⁺01] Patrick Th. Eugster, Rachid Guerraoui, Sidath B. Handurukande, Anne-Marie Kermarrec, and Petr Kouznetsov. Lightweight probabilistic broadcast. In *Proceedings of the International Conference on Dependable Systems and Networks (DSN 2001)*, pages 443–452, Gothenburg, Sweden, July 2001.
- [GKM01a] Ayalvadi J. Ganesh, Anne-Marie Kermarrec, and Laurent Massoulié. Probabilistic reliable dissemination in large-scale systems. Technical report, Microsoft Research, 2001.
- [GKM01b] Ayalvadi J. Ganesh, Anne-Marie Kermarrec, and Laurent Massoulié. Scamp: Peer-to-peer lightweight membership service for large-scale group communication. In *Proceedings of the Third International COST264 Workshop, LNCS 2233*, pages 44–55, Berlin Heidelberg, 2001. Springer-Verlag.
- [KSSV00] Richard Karp, Christian Schindelhauer, Scott Shenker, and Berthold Vöcking. Randomized rumor spreading. In *Proceedings of the 41st Annual Symposium on Foundations of Computer Science*, pages 565–574, Redondo Beach, California, November 2000. IEEE Computer Society Press.
- [MR95] Rajeev Motwani and Prabhakar Raghavan. *Randomized Algorithms*. Cambridge University Press, Cambridge, England, June 1995.
- [Pit87] Boris Pittel. On spreading a rumor. *SIAM Journal on Applied Mathematics*, 47(1):213–223, February 1987.
- [RS98] Martin Raab and Angelika Steger. “balls into bins” — A simple and tight analysis. *Lecture Notes in Computer Science*, 1518:159–170, 1998.
- [WZ99] Ralph Wittmann and Martina Zitterbart. *Multicast Communication*. Morgan Kaufmann Publishers, 1999.