

Data Storage in Unreliable Multi-agent Networks

Filip Miletic
Delft University of Technology
Faculty of Electrical Engineering,
Mathematics and Informatics
Mekelweg 4
2628 CD Delft, Netherlands
f.miletic@ewi.tudelft.nl

Patrick Dewilde
Delft University of Technology
Faculty of Electrical Engineering,
Mathematics and Informatics
Mekelweg 4
2628 CD Delft, Netherlands
p.m.dewilde@ewi.tudelft.nl

ABSTRACT

The distributed data storage on unreliable devices, connected by a short-range radio network is analyzed. Failing devices incur loss of data. To prevent the loss, the data is split and distributed across the network. The graph-based connectivity model assuming independent erasures is given, and the capacity of such graph is computed. It is shown by an information-theoretic argument that multi-agent systems with unreliable connectivity can increase overall data storage reliability through cooperation.

Categories and Subject Descriptors

C.2.1 [Computer Systems Organization]: Network Architecture and Design—*wireless communication*; C.2.4 [Computer Systems Organization]: Distributed Systems—*network operating systems*; I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence—*multiagent systems*

General Terms

Reliability, Performance, Design

Keywords

Distributed Coding, Erasure Graph, Partition Encoding

1. INTRODUCTION

Agent programs in multi-agent systems (MAS) often use an *agent platform* as middleware [3]. The reliability of data storage in an agent platform has received little attention *in practice*, though it is recognized as a MAS issue. A simple networked failure model for an agent platform is given and its properties are analyzed. It is assumed that no reliable persistent storage exists, so reliability must be derived from inter-agent cooperation instead. The model describes agents living in chaotic environments. Examples are the

The research result described here is conducted within the scope of Project COMBINED, run by DECIS Lab in Delft, Netherlands.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

AAMAS'05, July 25-29, 2005, Utrecht, Netherlands.
Copyright 2005 ACM 1-59593-094-9/05/0007 ...\$5.00.

agents in mobile wireless networks, where the dynamics can prevent the communication and thus agent platforms to work together. An example application is a network supporting emergency services in search-and-rescue operations (e.g. police, fire department, ambulance). The agents live on nodes connected by a short range radio-based network. Nodes in the network cooperate in relaying each others' data packets. Let a network graph be formed when nodes drawn from a set V are placed uniformly in a compact two-dimensional domain. Two nodes v_i and v_j from V can communicate directly when closer than a fixed range r . Further nodes can be reached by relaying. Each node stores precious information (*tag*), N bits long. Agents can be destroyed at random times, thus their tags lost. To prevent the loss, agents cooperate with neighbors, contracting a *keep* for some tags. We investigate the strategies that guarantee longest tag survival.

1.1 Motivation and Related Work

Using error correction coding in distributed systems for fault tolerance has not been widely applied in multi-agent systems. Current multi-agent platforms like JADE and COUGAAR use replication. Replication needs a lot of storage and bandwidth support, and information theoretic results tell us that we can do better. Some information theoretic ideas are used by the CFS cooperative storage [1] and the OceanStore [2]. Peer-to-peer layers such as Kademia [4], Tapestry [7] and Chord [6] form the building blocks of these applications. Contrasting to given examples, we are interested in localized networks that suffer failures. For each node, knowing the neighborhood structure entails best data distribution. As a difference from the classical peer-to-peer, the environment puts an upper bound to the number of neighbors.

The exposition uses the framework given in [5]. The *erasure graph* model is used for the connectivity, with independent, identically distributed erasures. Erasures give rise to disjoint *chains* of connected keepers that can re-distribute tags or contract new keepers. A keeper i can belong to different chains, giving rise to its *availability* w_i , and number of bits s_i it stores. The environment is able to switch off any network node, with a given probability ϵ . We consider a partition of a string of binary data, of length N is made. Each component of the partition is stored at a different node. We give the capacity result for this storage schema in Section 2. The notation is explained in Section 3.

2. RESULT STATEMENT

THEOREM 1. *Let G be a graph, and \mathcal{D}_i the corresponding i.i.d degradation model. Let \mathcal{P}_e be a set of erasure patterns, with a fixed number of erasures e . The capacity of the channel \mathcal{C} , defined on an erasure-graph $\mathcal{G} = (G, \mathcal{D}_i)$, under a given partition*

encoding Π and assuming uniform connection probability, is obtained by solving a linear program: $C = \max_{\Pi} \sum_{1 \leq i \leq |V|} w_i |i|$, $\sum_{1 \leq i \leq |V|} |i| = n$ with w_i for $1 \leq i \leq |V|$ and $\iota(x)$ is an indicator function:

$$w_i = \sum_{e=0}^{|V|-1} \frac{\varepsilon^e (1-\varepsilon)^{|V|-e}}{|V|} \sum_{E_p \in \mathcal{P}_e, G \in \mathcal{Q}_{E_p}} |V(G)|_{\iota}(v_i \in V(G)).$$

THEOREM 2. *Let the channel be defined as in Theorem 1. Let $w = \max_{1 \leq i \leq |V|} w_i$. Then the solution to the linear program of Theorem 1 is: $C = nw$.*

3. NOTATION AND DEFINITIONS

Let V be a set of vertexes, and E be a set of pairs (v_1, v_2) , where $v_1, v_2 \in V$. Let the graph G be defined as $G = (V, E)$. Interpret this in the usual sense: V is a set of vertexes, and E is a set of edges, constructed by connecting pairs of vertexes from V . Erasure is the construction of a new graph G' from G , by picking a subset V' of V , with respect to some probability distribution.

DEFINITION 1 (ERASURE PATTERN). *Let $G = (E, V)$ be a graph. An erasure pattern is a mapping $E_p : V \mapsto \{0, 1\}$, which associates each vertex of G with 1 if this node is to be erased from the graph, and 0 if this node is not to be erased.*

DEFINITION 2 (DEGRADATION MODEL). *Let a set of erasure patterns 2^V and an induced probability mass function P be given, such that for $e_p \in 2^V$, $P(e_p) = \Pr(e_p)$, the probability that a given pattern e_p occurs. A degradation model is given by $\mathcal{D} = 2^V, P$.*

DEFINITION 3 (ERASURE GRAPH). *Let there be given a graph $G = (V, E)$, and the degradation model \mathcal{D} . The tuple $\mathcal{G} = (G, \mathcal{D})$ is called the erasure-graph.*

As a special case, consider the degradation model where all single vertex erasures are independent identically distributed (i.i.d) with probability ε that each component of V disappears in V' . $G' = (V', E')$ is now a partial subgraph of G induced by a subset of vertexes $V' \subseteq V$. Assume that G can be transformed into any of the subgraphs G' that can result from removing a number of vertexes from V . The transition probability depends upon ε , the outage probability. It defines the probability that a node will go offline. Since outages are i.i.d., the probability of $k \leq |V|$ nodes going offline is:

$$\Pr(k \text{ nodes fail}) = \binom{|V|}{k} \varepsilon^k (1-\varepsilon)^{|V|-k} \quad (1)$$

DEFINITION 4 (I.I.D. DEGRADATION MODEL). *Let there be given an erasure graph \mathcal{G} and an outage probability ε . An i.i.d. degradation model is a degradation model $\mathcal{D}_i = 2^V, P_i$, where P_i is a probability mass function, given by equation (1).*

DEFINITION 5 (ERASURE-GRAPH DEGRADATION). *Let \mathcal{G} be an erasure-graph. A degraded graph, with respect to a degradation model $\mathcal{D} = 2^V, P$ is an erasure graph $\mathcal{G}' = (G', \mathcal{D}')$, where $G'(E', V')$ is a partial subgraph of G induced by a set of vertexes $V' = \{v : E_p(v) = 0\}$, and $\mathcal{D}' = 2^{V'}, P'$, where P' is a probability mass function obtained by marginalizing P over elements of $V \setminus V'$. A state transition of the graph \mathcal{G} to \mathcal{G}' by an erasure pattern E_p is denoted as: $\mathcal{G} \xrightarrow{E_p} \mathcal{G}'$.*

DEFINITION 6 (CONNECTED COMPONENTS). *Given a graph G , the set of connected components of G is a set of graphs $\mathcal{Q}_G = \{G_1, \dots\}$, where each G_i is a maximal connected partial sub-graph of G .*

Let a binary string X be written into a communication channel, and let a string Y be retrieved from it. Let there be employed the partition coding defined by Π . The mutual information on X and Y is given as: $I(X, Y) = H(X) - H(X|Y)$, where $H(X)$ and $H(X|Y)$ are, in order, the entropy of a random variable X and the conditional entropy of a random variable X given Y . For the erasure graph, we define Π to be the partition, i.e. the set with elements $|i|$ for $1 \leq i \leq |V|$.

DEFINITION 7 (CAPACITY OF AN ERASURE GRAPH). *The capacity of a communication channel is the maximum of mutual information over tunable channel parameters Π , given the components: $C = \max_{\Pi} I(X, Y)$.*

In this case $I(X, Y) = H(Y)$. For the ‘‘channel’’ formed by the erasure graph, the channel transition matrix depends on G , ε , and Π . Only Π can be chosen freely, thus by a suitable choice of Π , with other parameters fixed, it is possible to maximize C .

4. CONCLUSION

This paper presented an information-theoretic approach to improving multi-agent platform reliability. It was shown that there is substantial space for improvement by considering qualities emerging from cooperation. The information-theoretic arguments employed here come from the vast information and coding theory literature but have so far received comparatively little attention in multi-agent platforms despite potential gains.

5. REFERENCES

- [1] Frank Dabek, M. Frans Kaashoek, David Karger, Robert Morris, and Ion Stoica. Wide-area cooperative storage with CFS. In *Proceedings of the 18th ACM Symposium on Operating Systems Principles (SOSP '01)*, Chateau Lake Louise, Banff, Canada, October 2001.
- [2] John Kubiawicz, David Bindel, Yan Chen, Patrick Eaton, Dennis Geels, Ramakrishna Gummadi, Sean Rhea, Hakim Weatherspoon, Westly Weimer, Christopher Wells, and Ben Zhao. Oceanstore: An architecture for global-scale persistent storage. In *Proceedings of ACM ASPLOS*. ACM, November 2000.
- [3] Luck et al. *Agent Technology: Enabling Next Generation Computing*. The Agentlink Community, 2002.
- [4] P. Maymounkov and D. Mazieres. Kademlia: A peer-to-peer information system based on the xor metric, 2002.
- [5] Filip Miletic and Patrick Dewilde. Distributed coding in multi-agent systems. In *IEEE Conference on Systems, Man and Cybernetics*. IEEE, October 2004.
- [6] Robert Morris, David Karger, Frans Kaashoek, and Hari Balakrishnan. ‘‘chord: A scalable peer-to-peer lookup service for internet applications’’. In *ACM SIGCOMM 2001*, San Diego, USA, September 2001.
- [7] Ben Y. Zhao, Ling Huang, Jeremy Stribling, Sean C. Rhea, Anthony D. Joseph, and John D. Kubiawicz. Tapestry: A resilient global-scale overlay for service deployment. *IEEE Journal on Selected Areas in Communications*, 22(1):41–53, January 2004.