Concurrent Online Tracking of Mobile Users

Seminar of Distributed Computing
ETH Zurich

Claudio Munari
(munaricl@student.ethz.ch)

December 2003

1 Summary

B. Awerbuch and D. Peleg present in their paper [AP91] a mechanism to solve the key problem of basically any communication network: the location of its (mobile) entities in order to route traffic. They developed a strategy to keep track of mobile entities travelling within a network with static nodes, such as users in a cellular phone network. The introduced mechanism is based on a decomposition of the network into a hierarchy of regional directories, built of overlapping sets of nodes: the Read and Write sets. The construction of these sets is based on the graph-theoretic notion of a regional matching. By following this distributed approach, it can be guaranteed that local operations always cost less than global ones and worst-case communication costs can be kept low.

An m-regional matching is the collection RW of all Read and Write sets in the network with the property that the Read(u) and Write(v) sets of any two nodes u and v overlap as long as the distance[1] between u and v is smaller or equal to m. Read(v) and Write(v) set of a node v form together a Regional Directory RD that supports two operations, Find(ξ,v) and Move(ξ,v,u), the search of user ξ invoked at node v and a move operation of user ξ from v to u, respectively. As mentioned above, every Regional Directory consists of a hierarchy of levels, ranging from the lowest level RD_1 up to RD_δ, with δ as the logarithm of the diameter[2] of the network, $\delta = \lceil \log D(G) \rceil$. The purpose of the i'th level of the Regional Direcotry, RD_i, is to track any user residing within distance 2^i from a searching node.

Furthermore, a cleverly worked out mechanism of forwading pointers is added to the concept of regional directories. Thanks to this trick, the costly updates caused by Move operations can be limited. Instead of updating all levels of a Regional Directory, only the lowest $\log d$ levels are updated. As a consequence, higher levels of the Directory still point to an old address of the moving user. However, this fact does not cause a problem, because the user leaves behind at its previous location a forwarding pointer to its new location.

As a main result of the proposed tracking mechanism, it can be shown that the communication overhead of the Find and Move operations is within a polylogarithmic factor of the lower bound, polylogarithmic in the size and diameter of the network, respectively.

2 Analysis

The design of a distributed directory server can be seen as a three-way tradeoff between overhead, latency, and simplicity. The proposed hierarchy of regional directories together with

[1] dist(u,v): length (weight) of shortest path between u and v
[2] D(G): maximal distance between any two nodes in the network G

their update and forwarding mechanisms is by far not a simple method as can be guessed from the brief summary above. Moreover, a real implementation has to consider many additional mechanisms in order to tackle the problem of concurrent accesses (which can cause race problems), update algorithms (using locking mechanisms) that prevent long chains of forwarding pointers, and, last but not least, schemes which also tolerate failures in the network.

Some of these problems have been considered in the paper [AP91]. Unfortunately, most of the solutions presented are very informal and without any proofs. The reader is put off to the "full paper" that was published in 1995 by the same authors [AP95]. However, even in the "full paper", some of the basic questions are answered only superficially. The central and non trivial problem of partitioning the network in Read and Write sets with low radius and little overlap are not explained in detail and deferred to additional papers. But only thanks to achievements as in the match-making strategy of Mullender and Vitnyi [1988] or the related graph-theoretic problem of designing sparse graph covers, the presented results have become possible. These strategies play a fundamental role in the construction of Regional Directories, because the communication overhead of Move and Find operations grows as the product of the degree and the radius of the related 2^i-regional matching does.

3 Experimental Results

It would be interesting to compare the proposed tracking mechanism with existing mechanisms as in GSM networks, for instance. As the hierarchical decomposition of the network in the presented paper [AP91] is finer than in traditional solutions, it is possible that the presented tracking mechanism is more efficient. Although the presented scheme has shown low worst-case communication complexity, its implementation behavior is not known in detail. A first implementation of the scheme has been presented in [WWS00] and shows interesting results. In particular, the Read set size in randomly-generated networks is much smaller than the theoretical bound ($4 * \log_2(n)$) given by Awerbuch and Peleg. Furthermore, a surprising result was that the Read set size is a maximum of 2 for uniform balanced tree networks, independently of the node degree and of the total number of nodes. The authors of [WWS00] believe that the Awerbuch-Peleg scheme could serve as a viable alternative to conventional approaches used in cellular wireless networks and also in the Internet. In particular, the consideration of locality of Finds and Moves is seen as a main strength compared to conventional, more centralized solutions used today.

4 Related Work

Of course, the presented scheme does not apply to evolving network types like ad-hoc networks, since it is designed for networks with a static and reliable infrastructure. However, the idea of overlapping sets is not far away from the concept of quorum systems. In particular, both concepts have similar problems to solve, such as the construction of the node sets with little overlap and low radius, for example. [HL99] presents a tracking mechanism using quorum systems, thereby extending the considerations above with focus on ad-hoc networks.

References

[AP91] B. Awerbuch and D. Peleg, Concurrent online tracking of mobile users, SIGCOMM 1991

[AP95] B. Awerbuch and D. Peleg, Online Tracking of Mobile Users, J. of the ACM, 1995

[HL99] Z.J. Haas and B.L. Liang, Ad Hoc Mobility Management with Uniform Quorum Systems, IEEE/ACM TON, April 1999

[WWS00] I.P. Weerakoon, A.L. Wijesinha and D.P. Sidhu, On properties of read and write sets in the Awerbuch-Peleg scheme for tracking mobile users, Wireless Networks, May 2000