ICT-MobileSummit 2009 Conference Proceedings
Paul Cunningham and Miriam Cunningham (Eds)
IIMC International Information Management Corporation, 2009
ISBN: 978-1-905824-12-0

Tracking a large number of tokens: a customised prioritised medium-access protocol for target responses

Virgilio Rodriguez,[1] Friedrich Jondral, Holger Jaekel
Institut für Nachrichtentechnik
Universität Karlsruhe (TH)
Karlsruhe, Germany
email: vr <at> ieee.org ; {jondral, jaekel}<at> int.uni-karlsruhe.de

Abstract: An organisation is interested in tracking a large number of potentially moving "targets", each fitted with a simple tracking-assisting device, and roaming over a known, limited geographical area. To assist in tracking, a well-designed network of fixed "anchors", with plentiful energy and computing power, and a reliable (say wired) communication channel for coordination is available. Anchors take turns (with spatial reuse as possible) sending a "ranging request" (RR) and receiving responses from targets. A target response is "heard" by several anchors, which enables the system to estimate the target's current location. The key issue is to coordinate the targets responses. At a given moment some targets may be inactive, asleep, or simply static, while others may be moving at various speeds. Potentially, many active targets may congregate in certain subarea. A simple time-division solution is inappropriate because many time-slots would be needed to accommodate the maximum number of targets, and many would be wasted at a given time. A simple ALOHA-style solution would be problematic when many active targets are near each other. Below, we provide a customised priority-based medium-access solution, which works for a large (or small) number of targets under the assumed scenario. The present report covers conceptual/qualitative aspects only.

1. Introduction and scenario description

An organisation is interested in tracking a large number of potentially moving entities (things, machines, animals and/or people), or "targets" that roam over a known, limited geographical area. An appropriate approach for such scenario is to fit each target with a simple tracking-assisting device ("token" or "tag"), and to deploy a network of fixed devices ("anchors") to aid in the tracking. The anchors are assumed to have plentiful energy and computing power, and a reliable (possibly wired) communication channel for inter-anchor coordination. Reference [1] is a recent survey of the relevant literature.

The basic idea of the scheme is that anchors take turns (with spatial reuse as possible) sending a "ranging request" (RR) and receiving responses from targets. The anchor network is such that when a target responds, it is "heard" by a sufficient number of anchors for the system to estimate the target's current location. Figure 1 shows an idealised one-dimensional scenario. Since many targets may hear the RR from a given anchor, a critical problem is how to coordinate the targets responses. A medium-access control (MAC) protocol is needed.

Some aspects of this scenario complicate the choice of the MAC protocol. The movement pattern and activity level of a given target vary widely through the day. Thus, at a given moment some targets may be inactive, asleep, or simply static, while others may be experiencing various degrees of mobility. It is possible also that a relatively large number of

[1] V. Rodriguez is now with the Institute for Theoretical Information Technology, at the Rhein-Westfalia Technological University (RWTH) in Aachen, Germany.

targets are active near each other in certain area. Because the number of targets is "large", but only a fraction of them is expected to be mobile/active at a given time, and those active may be in any location within the area of interest, a simple time-division protocol is inappropriate (many time-slots would be needed to accommodate the maximum number of targets, and many would be wasted at a given time). When a large number of targets are active near a given anchor, as in figure 2, the system would like to give priority access to certain targets following an appropriate rationale. Thus, a simple ALOHA-style solution would not be appropriate.

Below, we provide a customised medium-access protocol for target responses in the situation of interest. The protocol is an adaptation of the Dutch auction, proposed for general medium access allocation in [2], in the context of infra-structureless ("ad-hoc") wireless networks. It offers a collision-free (to a high probability), prioritised medium-access protocol for target responses. The key priority idea is that an RR includes anchor identification, and the target can look up the distance between the present anchor and the anchor to which it last "spoke" in order to give itself a priority. If the present and previous anchors are far from each other, the target gives itself a high priority. At the other extreme, if the present anchor is the same as the previous one, the target gives itself minimal priority. (Notice that total distance travelled between RR is *not* the key determinant of the priority because a fast-moving target may have travelled a lot but may have returned near its previous location, or may be moving "in circle" in a relatively small area).

The remainder of the paper is organised as follows. We first describe the assumed geographical model. Then we discuss briefly the application of auctions to general medium-access problems, emphasising the advantages of the Dutch format for such purpose. Subsequently, we outline our MAC proposal, and address some of the practical implementation challenges. The paper ends with a brief discussion section.

2. Geographical model

In a common tracking scenario, targets are assumed to move in a two or three dimensional region. However, since this paper is primarily about the medium-access solution, we assume a "corridor" model, which is idealised as an interval in the horizontal axis. Thus, at a given time, a target position is determined by a single coordinate, interpreted as an abscissa (horizontal axis position). The corridor model enables us to highlight more clearly the medium access issues. The extension to a more general 2-dimensional situation is straightforward (from the MAC point of view).

Figure 1 shows the idealised one-dimensional scenario. The "tall", dark-coloured triangles denote anchors, identified as A_i. Even-numbered anchors (green) can all be simultaneously active for a period T, after which they switch role with odd-numbered anchors (red), and so on. The ovals indicate transmission range. Anchors power levels are set such that an anchor's range is half of the range of a target. The distance between adjacent anchors is such that, if a target that is between A_J and A_{J+1} answers a ranging request (RR), its responses reaches both A_J and A_{J+1}. Each of these anchors can then estimate the distance between itself and the target (for example, the anchor may know the transmit power utilised by the responding target, and can measure the power received, from which it can estimate the distance between itself and the target). With these two measurements "the system" can determine where the target is.

For example, t_3 can be heard by both A_1 and A_2, while both A_3 and A_4 are in listening

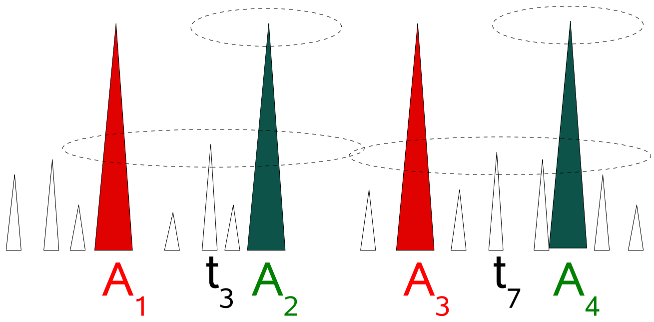

Figure 1: A simplified tracking scenario

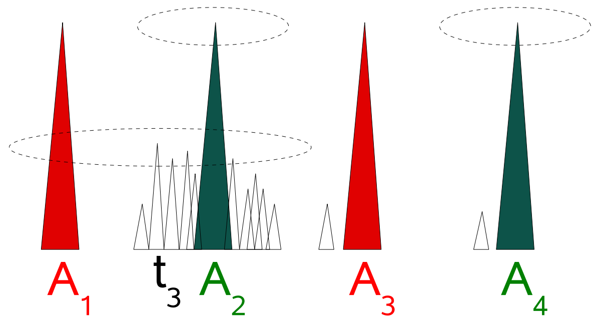

Figure 2: An unfavourable situation in which many targets "congregate" in a small region

range of t_7. Notice that, if figure 1 represents the entire system, a target to the left of A_1 (respectively, to the right of A_4) can only be heard by A_1 (respectively, A_4). But the system can still locate such target, because, for example, if a target is heard by A_4 and *not* heard by A_3, it must be to the right of A_4 (if it was between A_3 and A_4, it would be heard by both A_3 and A_4).

3. The Dutch auction

For organising the target's responses, we propose a procedure inspired on the Dutch auction. This section briefly discusses the use of auctions for *general* medium-access control (MAC), and highlights the advantages of the Dutch auction as a *general* medium-access protocol (that is, *not* specific to the scenario we study).

Since time immemorial, auctions have been employed as a practical mechanism for the transfer of ownership of articles of value, for such reasons as: (i) speed of allocation, (ii) discovery of the true "value" of the offered object, and (iii) transaction "transparency" (fraud prevention)[3]. For medium-access control (MAC), auctions provide a form of "prioritised access" in that the channel is allocated to the terminal that most values access. A terminal's valuation of access could either (a) represent the "true" monetary "willingness to pay" of a (selfish) human user, or (b) be a "priority" index computed/adjusted by software inside the terminal using local information. A terminal's priority may be "adaptive", depending on such factors as its "importance", packet type, location, channel state, distance travelled, battery status, etc.

A practical auction-based MAC protocol must be relatively simple, and produce a winner rapidly, because access must be granted quickly, and repetitively. Thus, previous auction-based MAC proposals resemble "sealed bid" auctions: each bid is independently submitted

Figure 3: A real-life Dutch flower auction

in a "sealed envelope", the auctioneer opens all envelopes simultaneously, the highest bidder wins, and pays as pre-specified by the rules. A sealed-bid auction requires an auctioneer (controller), a role that could be played by an anchor, in our scheme. Unfortunately, they also need an alternate MAC protocol to receive the bids. This protocol may be problematic with a large, possibly variable number of bidders (as in the situation we study). If it is contention-free, such as a time-division solution, it may be wasteful of resources; and if it is contention-based, as an aloha variant, the highest-value terminals may be unable to make a bid, and, consequently, a suboptimal allocation may result.

As an alternative to the sealed-bid formats, the Dutch auction utilises a public "clock" which displays a progressively falling price (see figure 3). Each participant watches the clock while waiting for the price to reach a desired level. At some point, a participant indicates its willingness to pay the current price (the first participant to do so is the one that most values the object) [3].

For MAC purposes, the Dutch auction retains the relative simplicity and allocation speed of other simple auction schemes, and add several fundamental advantages: (i) A built-in bid-processing protocol that automatically and simply prioritise the highest bid(s); (ii) the possibility of a distributive (auctioneer-free) implementation (start times, initial price, and rate of decrease can all be pre-specified, so that a terminal can determine from its own clock the current status of the auction); (iii) Confirmation of transmitter-receiver pairs at auction time, with smooth continuation if the pair is infeasible; (iv) exceptional signalling economy (only one bid signal (the winner's) is strictly necessary in a single channel scenario). The Dutch auction and its application to medium access allocation is discussed further in [2], where it is proposed for medium-access allocation in an infra-structureless ("ad hoc") synchronised wireless network. More recently, [4] extends [2] to consider network-layer issues.

4. MAC for targets' responses

In this section we discuss how to organise targets' responses following a procedure inspired on the Dutch auction. The key issue is how to determine a target's "bid". An outline of the algorithm follows.

If there are N anchors, let $1, 2, \ldots, N$ be the basic "priority" levels.

1. Each anchor includes its ID and location as part of the ranging request (RR).
2. Each target keeps in memory the ID and location of the anchor to which it last "spoke"
3. After hearing an RR, each target calculates a priority level and then a bid, as follows:
   (a) the target's priority is $p = 1 + d$, where d is the distance between the anchor sending the RR, and the last anchor to which it spoke, with the unit of length equal to the distance between adjacent anchors (recall that we are considering a "corridor" (linear) geography; the formula would be slightly more complex for a bi-dimensional region). For example, the priority of a target that last "spoke" to A_5 and now hears an RR from A_2 is simply 1+(5-2)=4.
   (b) The target's bid equals $p + r$ where r is a random number between -0,5 and 0,5 (with as many significant digits as allowed by the "tick" of the auction). Thus, if the target in the previous example randomly draws -0,2450, its bid is 4-0,2450= 3,7550.
4. Shortly after the RR is sent, the Dutch auction clock starts "ticking". The initial "price" equals the highest possible bid, that is, $(N-1)+1+0,5 = N+0,5$. The price drops at every tick by a pre-specified amount. If targets are synchronised, each target can silently determine the current price from its internal clock, while waiting for the "right price".
5. The first terminal whose bid is reached by the current price "takes the channel", to respond to the RR (Example, suppose there are 10 anchors, the price drops 0,1 per "tick", and that the highest bids for the targets in range of a given anchor are 9,3 , 7,0 , 5,5 and 4,0. Each target starts the synchronised silent countdown at 10,5, a "tick" later they reach 10,4 then 10,3 etc. When the price reaches 9,3 the corresponding target "takes the channel". Notice that a target need NOT know the bid of any other.)
6. The anchor acknowledges the response. If there is time left in this anchor period, the clock ticking and countdown resume (from 9,3 in this example: 9,3 9,2 9,1 etc). The process moves back to step 5 above.
7. Steps 5-6 are repeated until the current anchor period expires (in the preceding example, a second target speaks when the price reaches 7,0).
8. When this anchor period runs out, control is transferred to the other set of anchors (see figure 1).
9. Optionally, one can give a responding target an opportunity for a second response within the same anchor period, by telling it to generate a new bid at the lowest priority level (for example, 1 plus a random number here), and to wait until the price reaches this new (low) bid (if ever). It is, however, unlikely that a target's location change significantly within the same anchor period.

5. Implementation challenges

As any engineering solution, the MAC protocol must be judiciously designed. In particular, the parameters of the protocol (the "tick" of the clock, "price" reduction per "tick", and the length of the anchor period) should be chosen judiciously. Processing and signal travel time, and the clock "drift" are among the factors to be considered in choosing the protocol parameters.

The proposed scheme is (mostly) "collision free", in the sense that in order for a collision to occur two or more targets must (i) have the same priority AND (ii) draw the same random number, AND (iii) respond to the same anchor. It is, hence, reasonable to assume that the probability of such event is negligible. It would be a simple matter to introduce in the anchor's behaviour a procedure to handle such unlikely event.

It is evident that timing plays a critical role in the Dutch auction, and any scheme based on it, such as our proposal. However, this is not particular to our proposal. Infrastructure-based wireless networks typically utilise synchronous, contention-free multiple-access schemes, such as time-division and code-division multiple access in the data channel (and random access for the uplink control channel). The system under study has a network of well-connected anchors as infra-structure, and it is perfectly reasonable to assume that the anchors are synchronised, as we implicitly do when we specify that the anchors "take turn" sending ranging-requests (with spatial reuse). Target synchronicity is somewhat more controversial, because the tokens may be very inexpensive and simple devices.

Asynchronous targets can in principle be accommodated, by having the anchor broadcast the new "price" at every "tick" of the auction "clock". However, the anchor may need to also supply some data to enable the target to account for the propagation delay. Also, the energy spent by the targets decoding the "price" may be non-negligible. As an alternative, the targets may be equipped with very accurate clocks, which may be synchronised before deployment. All that is needed is that the synchronicity of these clocks outlasts the duration of the battery charge. While a target's battery is recharged or exchanged, the clock may be re-synchronised against a "master clock". Finally, the system could achieve target synchronisation by means similar to those employed by radio-controlled clocks, which rely in timing radio-signals provided by public service agencies [5].

Ultimately, synchronisation is not a MAC issue. Indeed, ultra-wide-band communication (the technology we target) may utilise extremely narrow pulses (of the order of nano-seconds in width), which necessitate extremely accurate transmitter-receiver synchronisation [6]. If the system can satisfy such requirements, it should find a way to also satisfy the synchronisation requirements of our MAC scheme.

6. Discussion

We have proposed a medium-access control scheme for a tracking scenario in which the number of targets is very large, and the mobility pattern and activity level of each target varies widely. With a small number of target, a time-division based approach could work quite well. But under the current scenario, time division would require a very large number of "slots" most of which would go wasted under typical operating conditions. The other extreme would be an ALOHA-based protocol, which would be problematic if many targets congregate near the same anchor (as in figure 2). In such case, a purely random back-off interval would be inefficient because a target that has barely moved may gain access, while another target that is currently very far from the location where it last "talked" to an anchor may be unable to “talk” in the present anchor period.

Our proposal works well for the scenario envisioned, regardless of the number of active targets, and their location. For instance, if a very large number of active targets “congregate” near an anchor (as in figure 2), the target that are farthest from the location they had the last time they “talked” to an anchor would get high priority for access. If the situation continues into the next anchor period, the targets that previously spoke would give themselves a low priority, which would facilitate channel access to the other targets, and so on. On the other hand, when few targets are active in range of a given anchor, all will get (virtually) collision-free access to the channel (a given target could even “speak” more than once, in a given anchor period).

Presently, our analysis is entirely *qualitative/conceptual*. Quantitative performance experiments are clearly desirable. However, to fully evaluate our proposal we need a simulator that fully captures the scenario we have addressed (a very large number of targets with various mobility patterns). Otherwise, the benefits of our proposal may remain hidden. The simulator should also account for some of the implementation challenges at the physical level. Such simulator is not, however, available as of this writing. On the other hand, the *qualitative* benefits of our proposal for the scenario we address are, in our view, evident. Thus, we hope that the dissemination of our proposal, even without numerical data, can serve a useful purpose.

Acknowledgement

We thank the European Commission for financial support through PULSERS-II project. This material should *not* be construed as official position of any project or agency.

References

[1] E. D. Manley, H. Al Nahas, and J. S. Deogun, “Localization and tracking in sensor systems,” in *Sensor Networks, Ubiquitous, and Trustworthy Computing, IEEE Inter. Conf. on*, vol. 2, pp. 237–242, Jun 2006.

[2] V. Rodriguez and F. Jondral, “Simple adaptively-prioritised spatially-reusable medium access control through the dutch auction: Qualitative analysis, issues, challenges,” in *Communications and Vehicular Tech. in the Benelux, 14th IEEE Symp. on*, vol. Supp., pp. 1–5, Nov 2007.

[3] E. Wolfstetter, “Auctions: An introduction,” *Journal of Economic Surveys*, vol. 10, pp. 367–420, Dec 1996.

[4] C. Lima and G. de Abreu, “Game-theoretical relay selection strategy for geographic routing in multi-hop wsns,” in *Positioning, Navigation and Communication, 5th Workshop on*, pp. 277–283, Mar 2008.

[5] R. E. Beehler and D. W. Allan, “Recent trends in NBS time and frequency distribution services,” *Proceedings of the IEEE*, vol. 74, pp. 155–157, Jan. 1986.

[6] L. Yang and G. Giannakis, “Ultra-wideband communications: an idea whose time has come,” *IEEE Signal Processing Magazine*, vol. 21, pp. 26–54, Nov. 2004.

