CAMELAMA: Cooperative Awareness and spaceborne Monitoring Enabled by Location-Assisted Medium Access

Holger Döbler
Computer Engineering Group
Humboldt-Universität zu Berlin, Germany
holger.doebler@informatik.hu-berlin.de

Björn Scheuermann
Communication Networks Group
Technical University of Darmstadt, Germany
scheuermann@kom.tu-darmstadt.de

***Abstract*—In beaconing systems such as AIS or ADS-B, used by ships and aircraft, each node periodically broadcasts its navigational state to nearby nodes to increase traffic safety. Nowadays these beacons are also used as a source for satellite-based global traffic monitoring. This dual use imposes competing needs on the medium access control protocol as the size of the collision domains varies by a large factor between the use cases. Even the subproblem of solely avoiding terrestrial nodes' packets to collide from the perspective of a receiving satellite is not trivial to solve if the satellite's collision domain spans multiple hops in the terrestrial network. Based on ideas of the LAMA protocol, we propose CAMELAMA, a novel contention-free medium access control protocol for position awareness beaconing. Our low-overhead approach uses neither forwarding of node state nor handshakes but only the navigational data that is shared between the terrestrial nodes anyway. CAMELAMA provides local cooperative awareness while at the same time desynchronizing transmissions within a satellite's collision domain even in terrestrially disconnected topologies. In a simulation-based performance evaluation we find that CAMELAMA outperforms SO-TDMA (the MAC protocol of AIS) and also scales better with respect to high node densities.**

***Index Terms*—medium access, satellite communication, maritime wireless networks, traffic monitoring**

I. Introduction

There exist cooperate awareness beaconing systems intended to increase traffic safety for many modes of transport: vessels use the automatic identification system (AIS) [1], aircraft use automatic dependent surveillance–broadcast (ADS–B) [2], automotive traffic uses ETSI ITS [3] or WAVE.

Despite the original primary use case of increasing traffic safety by vehicles broadcasting their own navigational state to nearby nodes, it has been found in the case of AIS and ADS–B that these messages can be overheard by low Earth orbit (LEO) satellites to achieve global traffic monitoring [4]. Even though *single messages* can be received and decoded in LEO, collisions of multiple simultaneously transmitted messages are a severe problem in practice, because typical satellite antenna footprints span multiple terrestrial collision domains [5].

This work was supported by the German Research Foundation (DFG) under grant number SCHE 1649/7-1.

In the past, this problem has been studied empirically and by modeling reception probabilities in satellite-based Automatic Identification System (S-AIS) with respect to self-organizing time division multiple access (SO-TDMA), the Medium Access Control (MAC) protocol of AIS that is mandatory for vessels at the high seas. Recent efforts to improve S-AIS performance were focused on advanced antenna and receiver design (see Sec. II). We present an orthogonal approach and discuss, on the example of AIS, how a MAC protocol for cooperative awareness beaconing can be designed, if the satellite-reception use case is taken into account.

CAMELAMA is a phy-layer-agnostic MAC protocol in the sense that it is merely a mechanism for the robust, decentral allocation of discrete units of channel resources to nodes moving on the Earths surface, while employing spatial re-use. With AIS as a motivating use case, we describe our protocol on the corresponding example of a single VHF channel with fixed-length time slots. We believe that the mechanisms that we describe could seamlessly be adapted to different scenarios where the channel resources are further subdivided in different dimensions like frequency (OFDM) or coding (CDMA).

In a fashion similar to LAMA [6], locally shared knowledge on node positions is used to orchestrate transmission-desynchronization and spatial re-use in a way that respects orbital collision domains. Minimal distances between simultaneously transmitting nodes can be controlled explicitly (during protocol design, not during operation) by a protocol parameter. Thereby the spatial properties of the protocols can be fit to the physical properties of a monitoring satellite formation. The length scales of spatial re-use for satellite reception lead to less frequent transmissions per node and thereby negatively affect the primary purpose of cooperative awareness. Therefore we propose a straightforward interleaving mechanism to mix transmissions intended for in-orbit and terrestrial reception.

Our results indicate that especially for high node density, CAMELAMA significantly improves S-AIS performance.

This article is structured as follows: after a discussion of related work in Sec. II, we introduce CAMELAMA in Sec. III. A performance evaluation in comparison to SO-TDMA is given in Sec. IV before this work is concluded in Sec. V.

II. Related Work

The problem of collision-free LEO reception of AIS messages has been studied extensively [7]–[14]. In [5], [13] models of the AIS network were developed to predict the probability of receiving uncorrupted messages in LEO. [8] studies the impact of the orientation of satellites' monopole antennas on the receiving performance. In [7], [10] the performance of actually operating S-AIS missions is discussed. In several publications including [9], [11] designs of advanced S-AIS receivers are proposed to decode messages despite collisions using messages' differing Doppler shifts and propagation delays, soft decision decoding and Viterbi decoding, and successive interference cancellation (SIC). Others seek to lower effects of interference using digital beamforming [12]. These approaches take SO-TDMA as given and study and/or improve at the receiving side. CAMELAMA, in contrast, seeks to improve the medium access itself with respect to in-orbit receiving. As pointed out in Sec. V, we believe that some of these approaches could be combined with CAMELAMA. While long range AIS [1, Annex 4] adds additional channel resources with the intent of satellite-based monitoring, it has no mechanism to desynchronize transmissions over multiple terrestrial transmission ranges.

Given that a LEO satellite's field of view (FoV) spans several SO-TDMA organized areas [4], the medium access pattern of vessels in distant organized areas looks approximately like random access. Several wireless random access MAC protocols that handle packet collisions using interference cancellation have been described [15]. Examples include slotless access schemes by means of ZigZag decoding [16] but also the slotted flavor using the paradigm of coded random access [17] or generalizations thereof [18]. It is useful to counter the hidden terminal problem (HTP) [16] and is especially useful in situations where any sort of feedback is expensive, e. g., in device-to-device broadcast communication [19]. Common to these approaches is the idea to introduce redundancy by spreading a transmission in time into multiple replicas of the same packet, either proactively or as the result of a retransmission mechanism. Decoding can then use this redundancy to iteratively extract information from collided transmissions, starting with a single non-collided transmission or part thereof. Like CAMELAMA, these protocols target the goal of feedback-less measures countering collisions of packets from hidden nodes. But the feedback-less subset of these SIC-based protocols lacks the ability to adapt transmission rates as it is needed for traffic beaconing scenarios where node densities change over time by orders of magnitude. Our approach guarantees a configurable upper limit of simultaneous transmissions per covered area.

With LT-MAC [20] and LBTM [21] other location-based TDMA protocols have been proposed. LT-MAC is designed for to-base-station unicasts whereas LBTM's goal is to reduce overhead of ACKs. As neither unicasts nor ACKs are needed for cooperative awareness, these approaches do not seem relevant in our context.

For a discussion of HTP-resolving broadcast-capable MAC protocols for terrestrial traffic-safety beaconing we refer to [6].

CAMELAMA re-uses some core ideas of the previously proposed protocol LAMA [6], extends the protocol towards applicability for space-borne global traffic monitoring. LAMA is restricted to geographically confined networks the topologies of which could be approximated in the Euclidean plane. CAMELAMA models a topology embedded on the unit sphere, enabling application of the protocol on a planetary scale. Where LAMA uses a simple hack to allow protocol bootstrapping and entry of nodes into the network that sacrifices channel resources for simplicity, CAMELAMA comes with a more robust mechanism. Finally, CAMELAMA can be adjusted to the aforementioned dual use as opposed to LAMA.

III. Protocol

A. Problem Statement

We propose a MAC protocol that is suited for the distributed allocation of channel resources to a large number of nodes (e. g., vessels) moving on the Earth's surface who seek to repeatedly broadcast their navigational state to other terrestrial nodes and to satellites in LEO. We assume that each vessel is equipped with a radionavigation satellite receiver (e. g., Global Positioning System (GPS)) that provides position information and clock synchronization at the order needed to align transmissions within slots [1]. The use for a position beaconing application implies that each node is aware of its transmitting and receiving terrestrial neighbors. When assuming a high-gain receiver antenna in LEO with sufficiently small swath, terrestrial single-hop collision avoidance would imply good receivability in orbit. However, wider satellite antenna footprints are favorable for two reasons: first, it enables monitoring a greater surface area per time, second antennas with sufficient gain require greater antenna sizes not compatible with the size of nano satellites [9]. Two-hop collision avoidance at the surface on the other hand is not sufficient if a satellite's collision domain spans several terrestrial hops. We assume that nodes are not aware of the orbital elements or merely the number of overhearing satellites in orbit, however. Our goal is to schedule fair and frequent transmissions of the vessels to achieve both, good cooperative position awareness between the vessels on the ground, and high packet reception rates for over-passing LEO satellites. In addition we assume that the size of the payload, i. e., the navigational state, is small compared to the typical count of neighbors in range. Therefore any mechanism whose overhead scales linearly with the number of receiving neighbors would waste a significant share of channel resources. In order to maintain a low protocol complexity we also abstain from any active participation of the receiving satellite(s), i. e., we assume that the satellites are silent observers who do not influence the medium access protocol and are unknown to the sending nodes.

B. The CAMELAMA Slot Allocation Mechanism

CAMELAMA is a single-channel time division multiple access (TDMA) protocol, i. e., it decides for each slot which

subset of nodes transmit; all other nodes are assumed to listen. Roughly speaking, both packet collisions at a receiving satellite and the HTP in terrestrial ship-to-ship communication are caused by simultaneous transmissions of two (or more) nodes that are too close together. In Sec. III-D we deal with the problem that terrestrial cooperative awareness and orbital overhearing have different demands regarding this minimum distance. For now, let's take as granted that in each time slot, a subset of all nodes with a minimum pairwise distance not smaller than a certain given distance shall transmit. While satisfying this condition, the subset of transmitting nodes shall be as large as possible to achieve good spatial re-use.

The fundamental idea of the CAMELAMA protocol is to manage medium access based on the nodes' locations to ensure a certain minimum distance between simultaneous senders. Time slots have a duration T_{slot}. They are numbered consecutively beginning at some epoch t_0: $i(t) = \lfloor \frac{t-t_0}{T_{\text{slot}}} \rfloor$. Given that the nodes' topology is (approximately) embedded in a unit sphere S_2, we assign a distinct set of points $F_i \subset S_2$ to each time slot i, such that the points are pairwise at least L apart (1):

$$\forall \mathbf{x} \in F_i,\, \mathbf{y} \in F_i,\, \mathbf{x} \neq \mathbf{y} \colon d_{S_2}(\mathbf{x}, \mathbf{y}) \geq L \tag{1}$$

$d_{S_2}(\cdot,\cdot)$ denotes spherical distance. F_i is called the *set of fire positions* of slot i, examples which of are qualitatively depicted at top of Fig. 2. The sequence of sets (F_i) is hardcoded into the protocol so that all nodes agree on the same set F_i based on synchronized clocks only.

To decide whether to transmit in slot i, a node first computes its *nearest fire position* $\mathbf{np}_i$, i. e., the element of F_i closest to its last transmitted position $\mathbf{x}$. The node transmits if both of the following conditions hold

1) $d_{S_2}(\mathbf{x}, \mathbf{np}_i) < d_{S_2}(\mathbf{y}, \mathbf{np}_i) \,\forall\, \mathbf{y} \in \{\text{neighbors' pos.}\}$
2) $d_{S_2}(\mathbf{x}, \mathbf{np}_i) < r_{\max}$ where the distance bound $r_{\max}$ is a parameter of the protocol.

The first condition allows only the node that is closest to its nearest fire position to transmit. Since nodes know their neighbors' positions, they are likely to mutually agree which node is nearest; thus, only one node per fire position will transmit in a slot. The second condition enforces an upper bound on the distance between transmitting nodes and corresponding fire positions. $r_{\max}$ (which corresponds approximately to $\frac{\alpha L}{2}$ in [6]) should be chosen smaller than half the typical maximum range over which position reports are received. Otherwise two (or more) nodes with the same $\mathbf{np}_i$ but without mutual awareness could both transmit in the same slot.

Visually speaking, each node transmits in a slot, if a fire position hits into its send region, where the send region is the intersection of its Voronoi cell with respect to node positions and the open ball $B_{r_{\max}}(\mathbf{x})$ at its own position.

We also employ the fairness mechanism called *vow of silence* as introduced in [6]: a position report's header contains a 16-bit unsigned integer v. After transmitting a position report, a node becomes inactive for the next v slots, i. e., it guarantees to other nodes that it will not transmit. After v slots of inactivity (while it still listens), it becomes active again. Each node keeps track of which neighbors are inactive and the first condition above is evaluated with respect to the set of active neighbors only. The value of v is recomputed every time a node transmits: assume there are globally N nodes in the network which are (stochastically) homogeneously distributed on S_2, yielding a node density of $\rho = \frac{N}{4\pi}$. For a sufficiently high ρ, there are globally $|F_i|$ transmissions per slot, meaning that with a perfectly fair MAC protocol, each node transmits every $\frac{N}{|F_i|} = \frac{4\pi\rho}{|F_i|}$ slots. In CAMELAMA, a transmitting node measures the *local* node density in its vicinity using an estimator function $\hat{\rho}$ and uses the value $v = \min\{\lfloor \nu \frac{4\pi\hat{\rho}}{|F_i|} \rfloor, 2^{16} - 3\}$. (The two remaining values $2^{16} - \{1, 2\}$ are reserved for bootstrapping, as explained later.) To avoid clipping we use 16 bit such that inter-transmission times up to $\approx \frac{1}{2}$h can be encoded. Effects of limiting the values of v are empirically studied in [6]. We use $\hat{\rho} = \frac{N_{B_{r_{\max}}}}{A(B_{r_{\max}})}$, i. e., the density of nodes within the ball $B_{r_{\max}}$. The parameter $\nu \in \mathbb{R}_+$ is used to adjust the intensity of the fairness mechanism; $\nu = 0$ would disable the mechanism whereas $\nu = 1$ corresponds to full fairness in case of perfectly uniform node density. In [6] it is shown empirically that increasing ν leads to better fairness at the cost of decreasing channel utilization. Here we use $\nu = 0.8$ as a compromise throughout the evaluation in Sec. IV.

C. Feasible Fire-Position Sets

Since (1) is invariant under rotations (the group $\mathrm{SO}(3)$), we reduce the problem of finding adequate F_i for every slot to finding *one* "base fire position set" set F_* that satisfies (1) and rotating it differently in every slot:

$$F_i = \{R_i \mathbf{x} \mid \mathbf{x} \in F_*\} \text{ with } R_i \in \mathrm{SO}(3)$$

The rotations R_i are generated uniformly at pseudo-random, but deterministically depending on i. Using a pseudorandom number generator (PRNG), a hash function, or a combination thereof, i is mapped to a triple (x_i, y_i, z_i) uniformly in $[0, 1)^3$ from which R_i is computed using the method from [22].

The problem how to place as many points as possible on a sphere while satisfying (1) is closely related to the Tammes problem [24]. Most, if not all, state-of-the-art approaches to compute approximations to the Tammes problem for $k \gg 14$ are numerical methods [25] that place k *particles* on the sphere at, e. g., random, initial positions, assume a repulsive conservative force and apply numerical optimization strategies to minimize the potential energy. Since it turns out that this results in point sets with little regularity, an implementation of CAMELAMA would need to store all points of F_* explicitly. In addition, finding the closest point $f \in F_*$ for a given location $\mathbf{x} \in S_2$ either costs $O(|F_*|)$ time when using exhaustive search or requires an additional data structure.

Instead of solutions of the Tammes problem or best-known approximations thereof we use sets consisting of layers of constant latitude each holding longitudinally equidistant points. A naïve way to construct such layered sets satisfying (1) is to use L as distance of adjacent layers and fill each layer with as many points as possible. Since the naïve approach achieves

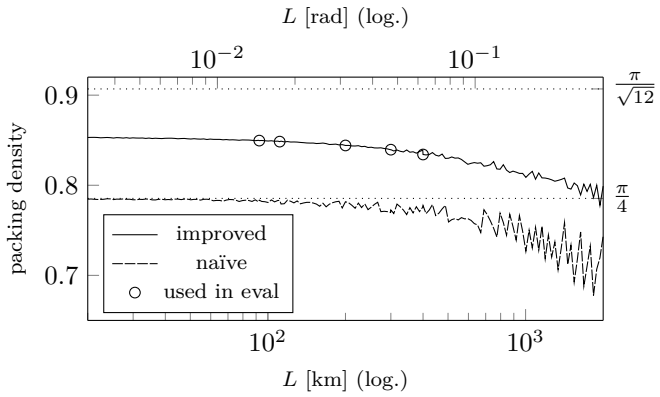

Fig. 1: Packing densities of naïve and improved layered fire position sets. "Naïve" seems to approach $\frac{\pi}{4}$ for small L whereas "improved" achieves densities around 0.85 for the range of L employed in our evaluation, which is $\approx 7\%$ smaller than the upper bound [23] $\frac{\pi}{\sqrt{12}}$. The bottom axis corresponds to a $r = 6371\,\mathrm{km}$ sphere.

packing densities of only $\frac{\pi}{4}$ we use an improved method to generate layered base fire position sets with packing density ≈ 0.85 (see Fig. 1). A detailed description of the algorithm is provided along the the implementation.[1] Instead of positions for all points, only latitude, number of points, and a longitude offset need to be stored for each layer. Computation of the nearest fire position of a given point in S_2 is computationally cheap: a small candidate set of the layers inside the point's $\frac{\sqrt{3}L}{2}$-neighborhood can be found using interpolation search; the nearest point in a given layer is found with one single division. This allows an exhaustive search for the minimum distance in a set of typically no more than four candidates.

D. Terrestrial Cooperative Awareness & Orbital Overhearing

In case of vessel safety at the high seas, broadcasting navigational data serves two goals: terrestrial cooperative awareness and global surveillance through overhearing of position beacons by LEO satellites. As discussed in [6], good spatial reuse and therefore high channel utilization for terrestrial reception can be achieved when the distance between nearest simultaneous senders is approximately twice the typical terrestrial transmission range. On the other hand, such optimal terrestrial spatial reuse is counter productive for orbital overhearing: it leads to multiple nodes transmitting simultaneously within the satellite's receiving antenna main lobe. To solve this problem, we use two base fire positions sets, one with a small value of $L = L_{\mathrm{ter}}$ that is optimized for terrestrial transmissions and the other with a much larger parameter $L = L_{\mathrm{orb}}$ for orbital overhearing. These base fire position sets are applied in a simple fixed interleaving pattern: Given an integer protocol parameter N_{orb}, the finely granular base fire position set is used in *terrestrial slots* $i \equiv 0 \pmod{N_{\mathrm{orb}} + 1}$ (where i is the integer slot number) while the coarsely granular base fire position set is used in all other slots (from here on called *orbital slots*). Achieving channel access fairness in orbital slots usually requires much larger values of v than in terrestrial slots, because if fewer nodes can transmit simultaneously then each node transmits less frequently. In order to handle this appropriately, each vow of silence applies only to the type of slots it was transmitted in. In our implementation, each node does not only count up total slot numbers in order to generate pseudo-random fire positions and determine the slot type, but also counts up per-type slot numbers i_{ter} and i_{orb} satisfying $i_{\mathrm{ter}} + i_{\mathrm{orb}} = i$ that are used only for the vow-of-silence fairness mechanism. For a visual overview, see Fig. 2.

[1] https://gitlab.informatik.hu-berlin.de/stephaho/tammes, accessed 12/10/21.

E. Bootstrapping

Using the protocol mechanisms described so far, CAME-LAMA is inherently hostile against new nodes entering the network as well as practically unable to reach a sane steady state of cooperative awareness: two nearby nodes that are not mutually aware of each other are forced by the protocol to use the same slots for transmission most of the time.

In LAMA [6], each node remains silent and listens probabilistically in a fixed fraction of its assigned slots to solve this. Here we take a different approach and add an introductory node state (that each node goes though once) to the existing protocol mechanisms that are from now on referred to as "regular state" behavior. A node in the introductory state behaves oppositely to regular state nodes: it transmits position reports flagged (using $v = 2^{16} - 2$) as introductory (a) probabilistically with a small probability, e. g., $p_{\mathrm{tx}} = 0.01$, (b) and only in slots where its nearest fire position is more than $\frac{L}{2}$ away. Condition (b) assures that the nearest transmitting regular-state node is far away whereas condition (a) decreases the likelihood of two or more nearby introductory-state node transmit simultaneously.

Regular-state nodes receiving introductory messages acknowledge their reception using their own regularly assigned slots (using $v = 2^{16} - 1$ as ACK flag). Introductory-state nodes switch to regular state when either receiving an acknowledgment or when a couple of introductory-state messages has been sent and there is no regular-state node known in the vicinity that could transmit an acknowledgment. The whole bootstrapping mechanism is restricted to terrestrial slots only.

In our simulation-based experiments all protocol nodes start simultaneously. This, however, is rather an artifact than a common situation in practice. In addition, even though the bootstrapping phase is necessary to conduct simulations, it happens entirely in the equilibration time span that was excluded from measurements and therefore does not affect the results. Therefore we don't go further into detail and also abstain from an explicit performance evaluation of the bootstrapping mechanism described here. The mere fact that it successfully bootstrapped the protocol in each simulation that we conducted, ipso facto is a strong empirical argument for it to be sufficient for its very purpose.

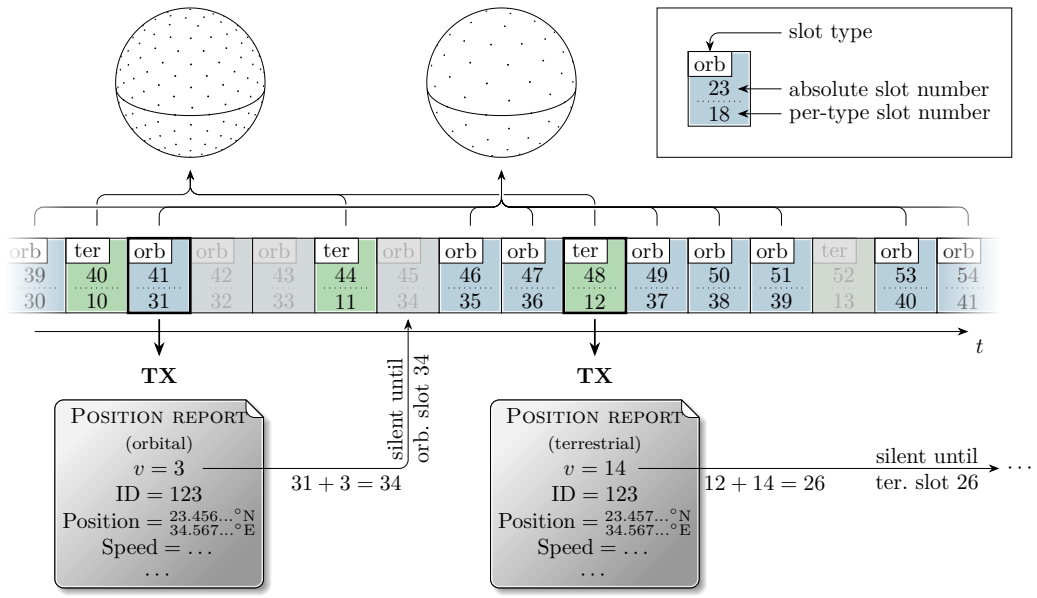

Fig. 2: An example of slot type interleaving, slot numbering, and the vow of silence mechanism. Base-fire-position sets are drawn with L much greater than in the evaluation. Slots not used because of the vow-of-silence mechanism are grayed out.

IV. Evaluation

We evaluate CAMELAMA using ns-3 [26] in a setting motivated by AIS and compare it against SO-TDMA. Our primary focus lies on the quantitative analysis of the successful reception of vessels' position reports by LEO satellites.

A. Scenario and Simulation Setup

1) Mobility: Ideally we would have liked to evaluate the protocol with a global Earth-scale configuration of vessels, complemented by an Earth-covering constellation of monitoring satellites as proposed in [4]. To produce trustworthy results we decided to model transmission errors based on the signal to interference plus noise ratio. In ns-3, using a linear list for range searching, this implies computation cost per transmission to be linear in the global number of nodes. This in turn leads to per-simulation computational cost quadratic in the number of nodes if the nodes' average transmission rate is kept constant. Given that our largest simulations of 6000 nodes already take several hours per run, simulations with realistic global ship counts (400,000–550,000 nodes as of 2016 [5]) would be prohibitively expensive. Improved range searching techniques based on grids and kd-trees [27] rely on in-euclidean-plane mobility and are therefore not applicable to Earth-spanning node distributions.

We therefore decided to use scenarios where the simulated vessels are confined to a rectangular box in latitude/longitude parametrization. We modeled vessel mobility using a spherical waypoint mobility model, i. e., each vessel's mobility is described in terms of a sequence of waypoints (i. e., time-position pairs); the node moves with constant speed along the shortest path between consecutive waypoints.

For most experiments we used a random walk mobility [28] with speed drawn uniformly at random (UAR) from $[0, 30\,\frac{\mathrm{m}}{\mathrm{s}}]$, inter-waypoint times are drawn UAR from $[0, 300\,\mathrm{s}]$, and initial positions drawn UAR with respect to areal probability density from the confining rectangle, i. e., $\sin\theta$ uniform in $[\arcsin\theta_{\min}, \arcsin\theta_{\max})$. We verified experimentally that this mobility model has a steady state and that it is initialized in this steady state to avoid pitfalls like [29]. In some experiments we used real-world AIS traces captured in the area of Denmark at 1:00pm–1:30pm (UTC) on 2017/06/01. [2] The traces were limited to those nodes corresponding to AIS Class A devices with at least one waypoint inside the rectangular region of latitude $[49.26°\mathrm{N}, 62.74°\mathrm{N}]$ and longitude $[5.97°\mathrm{E}, 14.03°\mathrm{E}]$, i. e., a rectangle centered at 56°N, 10°E with a latitude extent of $1500\,\mathrm{km}$ and a longitude extent of $500\,\mathrm{km}$ at 56°N. This results in 1688 nodes. The random walk topologies were generated for the same rectangle.

In order to study message reception of multiple satellites passing over the same region, we put four satellites in a string-of-pearls orbital configuration with $500\,\mathrm{km}$ altitude, vanishing eccentricity, 90° inclination, and $500\,\mathrm{km}$ inter-satellite distance. The remaining orbital elements were adjusted such that the formation's geometrical center passes over the rectangle's center point and such that the first satellite's footprint enters the rectangle just after we consider the terrestrial MAC protocol as equilibrated. The satellites' mobility was then modeled with SGP4 and converted to Earth-centered, Earth-fixed (ECEF) Cartesian coordinates used for vessel mobility.

2) Channel Model: We implemented CAMELAMA and SO-TDMA for a single $25\,\mathrm{kHz}$ bandwidth VHF channel at $161.975\,\mathrm{MHz}$ center frequency modulated in binary GMSK with $9600\,\mathrm{bit/s}$. Slots of $\frac{2}{75}\mathrm{s} \mathrel{\hat{=}} 256\,\mathrm{bit}$ are used. The signal-to-interference-plus-noise based packet loss model and the

[2] ftp://ftp.ais.dk/ais_data/dk_csv_jun2017.rar, accessed 12/02/18.

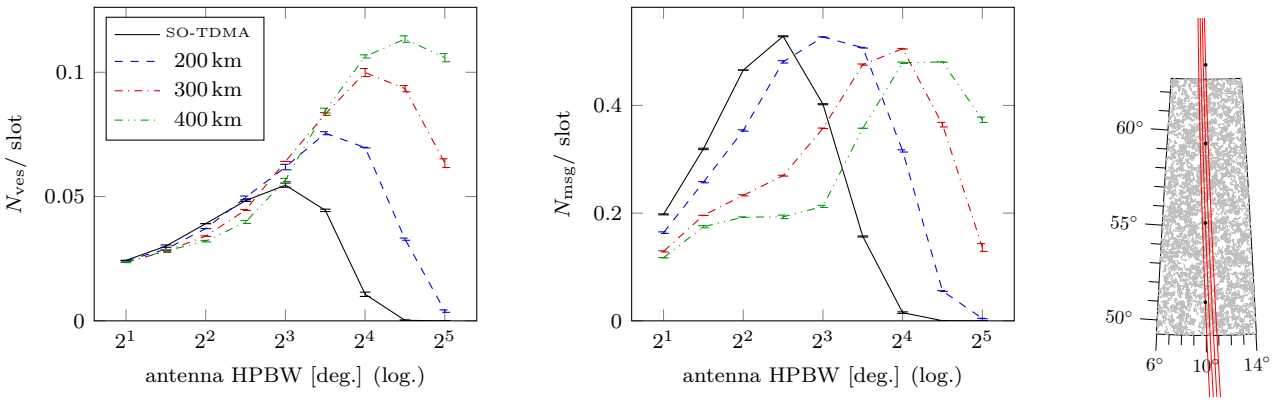

(a) Discovered vessels per time slot.

(b) Message count per time slot.

(c) Simulated topology.

Fig. 3: Number or (a) discovered ships N_{ves} and (b) received messages N_{msg}, each per slot per satellite, collected with antennas of varying beam width, during one formation overflight of a 3000 nodes random walk topology performing SO-TDMA or CAMELAMA, the latter with $L_{\text{orb}} \in \{200, 300, 400\}\,\text{km}$ and $N_{\text{orb}} = 4$. Legend of (a) applies to (b) as well. (c) shows the mobility of a 1500 nodes run.

path loss of terrestrial ship-to-ship communication was taken from [6]. Path loss computation for ship-to-satellite signal propagation was ported from the ESTNeT simulator [30].

3) Antenna model: A 500 km-altitude satellite with an omnidirectional antenna has a swath width of $\approx$ 5000 km horizon to horizon. We can neither simulate realistic-node-density topologies of this size (Sec. IV-A1), nor do we have access to corresponding captured AIS traces. Thus, we limit the swath width with an abstract model of a nadir-pointing directional antenna. It is parametrized by its half-power beam width (HPBW) and has a directional gain (2) depending only on the angle α between the received signal and the nadir direction.

$$A(\alpha) = c \cdot 2^{-\left(\frac{2\alpha}{\text{HPBW}}\right)^2} \qquad c \in \mathbb{R}\colon \int A \,\mathrm{d}\Omega = 4\pi \tag{2}$$

This abstract model allows to evaluate protocol performance depending on HPBW while leaving open if the receiver uses a physically directive antenna, digital beamforming using a patch antenna array, or creates the effective FoV by means of signal processing using timing- and Doppler-shift-based filtering on the received signals.

4) Split Simulation Strategy: As our protocol contains no transmissions from satellites at all, the vessels' behavior is completely independent of the satellites'. We use this to split each simulation into two separate parts: 1) a terrestrial-only simulation of CAMELAMA or SO-TDMA where we record an exact schedule of which message is transmitted in which slot. 2) And a satellite-reception-only simulation, where the vessels replay the transmission schedules recorded in the first simulation part while the vessels' and satellites' mobility are simulated. The second simulation part is much cheaper, because the vessels' behavior is fixed and independent of the messages that are received. Therefore, the loop in ns-3 that iterates over all nodes in the channel to compute signal-power levels for every node for every transmission can be limited to the few satellites, avoiding the need to perform this computation for thousands of terrestrial nodes. We leveraged this computational simplification by re-using the same terrestrial simulation part for different satellite configurations and receiver-antenna configurations that we examine. Unless stated differently, each data point corresponds to six independently seeded simulations and error bars in plots depict the standard error. Distances given as legend keys denote L_{orb} values.

5) Measurement Timing: We conduct measurements with different antenna HPBWs with the widest swaths just smaller than the longitudinal width of the terrestrial topology. If we measured data in situations with only part of the terrestrial topology in the satellite's FoV, boundary effects would not only affect our results, but their magnitude would vary with the antenna's HPBW, making it hard to tell apart boundary effects from the desired underlying performance characteristics. To circumvent this problem we seek to essentially cut out the boundary effects by limiting the measurement of in-orbit reception per satellite to the time span when the satellite's ground track is within the terrestrial topology's latitude range shrunk by 250 km. This way we can assure that the satellite's ground track during measurement is surrounded by $\gtrapprox$ 200 km of populated ground area (see Fig. 3c) in every direction.

Every measurement of SO-TDMA scenarios is preceded by 10 min of equilibration time to allow SO-TDMA to reach a steady state. Every measurement of CAMELAMA scenarios is preceded by equilibration time consisting of time needed for all nodes to reach regular state plus 2 min, which was together less than 10 min in total in each simulation run.

6) Parameters: In the evaluation performed in [6], the LAMA protocol proved robust with respect to its parameters. CAMELAMA is based on LAMA, so in this evaluation we used $L_{\text{ter}} = 92.6\,\text{km} (= 50\,\text{NM})$, $r_{\max} = 23.15\,\text{km} (= 12.5\,\text{NM})$, $\nu = 0.8$, and $N_{\text{orb}} = 4$ (see Sec. IV-B2) together with $L_{\text{orb}} \in \{200, 300, 400\}\text{km}$ and antenna beam widths $HPBW \in [2°, 32°]$. An exhaustive discussion of the parameter space would have exceeded the length of this publication.

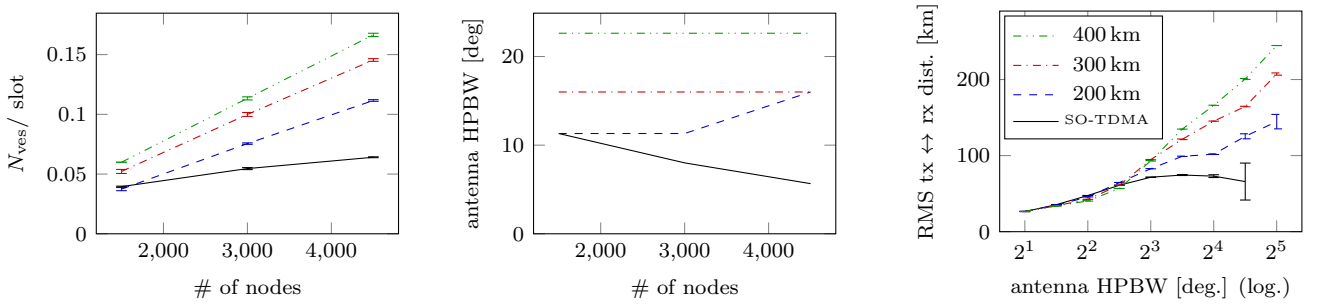

(a) # of vessels discovered.

(b) Optimal antenna beam width.

(c) RMS Sat. ground trace ↔ ship distance.

Fig. 4: (a) N_{ves} per slot for per-parameter-optimal antenna beam width (corresponding to the curve maxima in Fig. 3) and (b) the corresponding beam widths. Error bars are omitted for (b) because the values shown correspond to an argmax out of a finite set of parameter values. (c) shows RMS distance from sending ship to receiving satellite's ground track measured in runs depicted in Fig. 3a,b. The legend of (c) applies to all three plots.

B. Experiments and Results

1) Satellite-Based Monitoring: When a satellite formation passes over the simulated area of Earth, there is some redundancy in the received messages. The same message can be received by two or more satellites at the same time, and multiple different messages of the same vessel can be received, carrying different but highly correlated information. For scenario, we consider both, the number of messages received N_{msg} and, deduplicating messages from the same sender, the number N_{ves} of "discovered" vessels, i. e., vessels that at least one message is received from. These total counts were then divided by the sum of time slots regarded for reception over all four satellites. Fig. 3 shows both metrics over antenna HPBW for SO-TDMA and for CAMELAMA with $L_{\text{orb}} \in \{200, 300, 400\}\,\text{km}$ measured for 3000 nodes random walk topologies. Message counts received from SO-TDMA are on par with CAMELAMA. However, the former requires significantly higher antenna gain which is unfavorable as discussed in Sec. III-A, while the latter leads to more distinct vessels being discovered. As expected, the optimal antenna HPBW, i. e. the curves' arg maxima, increase with L_{orb}. A corresponding measurement based on real AIS traces (Fig. 5a) shows that CAMELAMA still performs better but the strong correlation of antenna beam width and L_{orb} vanishes.

We repeated this measurement for random walk topologies of identical geometrical bounds but with varying node counts. Fig. 4a shows N_{ves} per slot over node count where each data point was measured for the optimal HPBW. The corresponding optimal HPBW values are given in Fig. 4b. For CAMELAMA, the number of vessels detected increases while the optimal beam width shows no significant dependence on node density. SO-TDMA shows a significantly weaker increase along with a narrowing of the optimal antenna directivity, indicating that receiving SO-TDMA signals at high node densities requires high-gain antennas to compensate for smaller distances between simultaneous senders.

We also measured for every received message the distance of the transmitting node to the receiving satellite's ground track point to get a sense of the "effective swath width" that a satellite is receiving messages from. In Fig. 4c these RMS distances that correspond to the data points in Fig. 3 are shown. CAMELAMA effectively receives messages from a significantly larger FoV that increases roughly linearly with growing value of L_{orb}, just as one would expect. Finally we measured the monitoring accuracy in the scenario of real traces at the per-MAC-protocol-optimal HPBWs. At the end of the satellite reception measurement time interval, we recorded for each vessel in the simulation the distance between its true position and the position resulting from extrapolating the last beacon received in orbit to this point in time using dead reckoning. For vessels no beacon at all was received from, ∞ was recorded. Fig. 5b shows the relative cumulative counts and reveals that while CAMELAMA achieves a tracking accuracy smaller than $1\,\text{km}$ for 55%–64% of vessels, SO-TDMA provides the same accuracy only for 40%.

2) Terrestrial Performance: We have re-run the whole evaluation from [6] for CAMELAMA with $N_{\text{orb}} = 0$ (i. e., without orbital slots at all) and found that CAMELAMA without orbital slots behaves qualitatively equivalent to LAMA, outperforming SO-TDMA. To quantify the effect of orbital slots on the performance of terrestrial cooperative awareness, we measured the neighbor location prediction inaccuracy (the distance between neighboring nodes' true positions and the predicted positions extrapolated from the last received position report's location, speed, and course) depending on N_{orb}, using the methodology of [6], Section 4.3. In a $(222.24\,\text{km})^2$ ($120\,\text{NM}$) square topology of 1600 random-walk nodes we measured once every $7\,\text{s}$ the location-prediction inaccuracy of every ordered node pair and the node pair's true distance. We then binned these samples by the node distance. In Fig. 5c the prediction inaccuracy median is plotted over the ship-to-ship distance for SO-TDMA and CAMELAMA with varying $N_{\text{orb}} \in \{0, 2, 4, 6\}$. As we take away channel resources in terms of terrestrial slots, the inaccuracy increases from less than $2\,\text{m}$ to more than $10\,\text{m}$, which is still below a typical vessel size. Nevertheless, CAMELAMA achieves to maintain a nearly constant accuracy up to distances of $30\,\text{km}$.

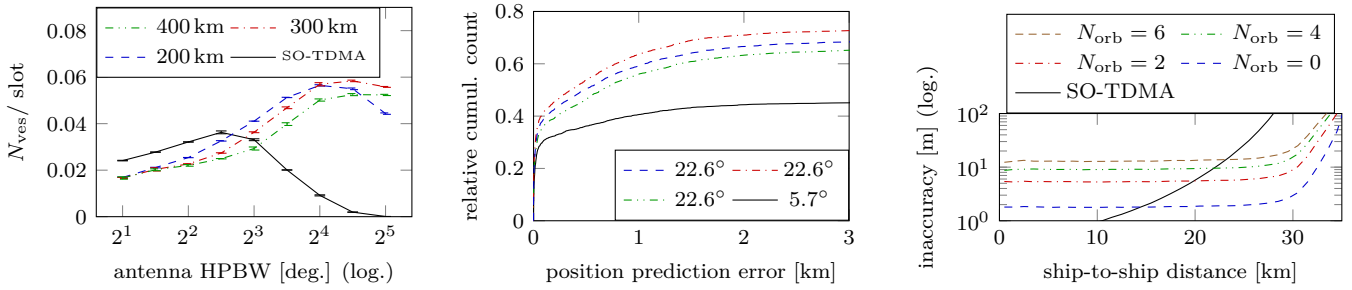

(a) N_{ves} vs. HPBW for real AIS traces.

(b) Orbital tracking inaccuracy.

(c) Median terr. location-prediction inaccuracy.

Fig. 5: (a) N_{ves} per slot over sat. HPBW for real vessel traces (legend denotes L_{orb}) and (b) a CDF-plot of the corresponding tracking (in-) accuracy measured at 5.7° (SO-TDMA) and 22.6° (CAMELAMA) HPBW. Legend of (a) applies to (b) as well. (c) Median of terrestrial location-prediction inaccuracy over ship-to-ship distance in random walk mobility for SO-TDMA and CAMELAMA at $L_{\text{orb}} = 400\,\text{km}$.

V. Conclusion

We introduced CAMELAMA, a location-assisted MAC protocol for cooperative awareness beaconing combined with satellite-based traffic monitoring. We demonstrated that it leads to significantly higher node discovery rates and tracking accuracy compared to SO-TDMA. This already holds for a simple static nadir-pointing satellite antenna and a receiver naïvely treating interference as noise in decoding. Since improving the MAC protocol is orthogonal to recent efforts to improve in-orbit de-collision of AIS messages, we see great potential in combining these techniques. If nodes use CAMELAMA, a satellite knows approximately where transmitting nodes are located. This knowledge could be exploited with respect to Doppler shift and propagation delay when using techniques as in [9] as well as for direction-dependent antenna gain adaption, i. e., receive antenna beamforming [12]. As SO-TDMA lacks this transmitter-location knowledge, the advantages of CAMELAMA might further increase significantly if improved decoding techniques are applied. As this work is merely a proof of concept, consideration of environmental effects like adverse atmospheric conditions is left for future work.

References

[1] ITU, "Recommendation ITU-R M.1371-5: Technical characteristics for an automatic identification system using time division multiple access in the VHF maritime mobile frequency band," Geneva, Feb. 2014.

[2] Federal Aviation Administration, "Automatic Dependent Surveillance—Broadcast (ADS–B) Out Performance Requirements To Support Air Traffic Control (ATC) Service; Final Rule," 2010, 14 CFR Part 91.

[3] European Telecommunications Standards Institute, "Intelligent Transport Systems (ITS); Vehicular Communications; Basic Set of Applications; Part 2: Specification of Co-operative Awareness Basic Service," 2010.

[4] D. Kocak and P. Browning, "Real-time AIS tracking from space expands opportunities for global ocean observing and maritime domain awareness," in *OCEANS'15*, Washington, DC, USA, Oct. 2015.

[5] C. Eisler *et al.*, "A surveillance application of satellite AIS utilizing a parametric model for probability of detection," in *ICORES '17*, Feb. 2017.

[6] H. Döbler and B. Scheuermann, "LAMA: Location-Assisted Medium Access for Position-Beaconing Applications," in *MSWiM '19*, Nov. 2019.

[7] A. Skauen, "Ship tracking results from state-of-the-art space-based AIS receiver systems for maritime surveillance," *CEAS Space Journal*, vol. 11, 2019.

[8] W. Hasbi *et al.*, "The impact of space-based AIS antenna orientation on in-orbit AIS detection performance," *Applied Sciences (Switzerland)*, vol. 9, 2019.

[9] S. Jayasimha *et al.*, "Satellite-based AIS receiver for dense maritime zones," in *COMSNETS '17*, Jan. 2017.

[10] S. Li *et al.*, "Data reception analysis of the AIS on board the TianTuo-3 satellite," *Journal of Navigation*, vol. 70, 2017.

[11] F. Clazzer *et al.*, "Enhanced AIS receiver design for satellite reception," *CEAS Space Journal*, vol. 8, 2016.

[12] F. Maggio *et al.*, "Digital beamforming techniques applied to satellite-based AIS receiver," *IEEE Aerosp. Electron. Syst. Mag.*, vol. 29, 2014.

[13] A. Harchowdhury *et al.*, "Generalized mechanism of SOTDMA and probability of reception for satellite-based AIS," in *CODEC '12*, Dec. 2012.

[14] G. Høye, "Ship detection probability analysis for a possible long-range AIS system," Forsvarets forskningsinstitutt, Tech. Rep., 2004, FFI/RAPPORT-2004/04383.

[15] E. Casini *et al.*, "Contention Resolution Diversity Slotted ALOHA (CRDSA): An Enhanced Random Access Schemefor Satellite Access Packet Networks," *IEEE Trans. Wireless Commun.*, vol. 6, 2007.

[16] S. Gollakota and D. Katabi, "Zigzag decoding: combating hidden terminals in wireless networks," in *SIGCOMM '08*, Aug. 2008.

[17] E. Paolini *et al.*, "Coded random access: applying codes on graphs to design random access protocols," *IEEE Commun. Mag.*, vol. 53, 2015.

[18] S. Akin and M. Fidler, "Multi-Access Spreading over Time: MAST," in *MSWiM '19*, Nov. 2019.

[19] M. Shih *et al.*, "A Distributed Multi-Channel Feedbackless MAC Protocol for D2D Broadcast Communications," *IEEE Wireless Commun. Lett.*, vol. 4, 2015.

[20] J. Mao *et al.*, "LT-MAC: A location-based TDMA MAC protocol for small-scale underwater sensor networks," in *CYBER '15*, Shenyang, China, Jun. 2015.

[21] H. Jang *et al.*, "Location-Based TDMA MAC for Reliable Aeronautical Communications," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 48, 2012.

[22] J. Arvo, "III.4 - FAST RANDOM ROTATION MATRICES," in *Graphics Gems III (IBM Version)*, 1992.

[23] R. Robinson, "Arrangement of 24 points on a sphere," *Mathematische Annalen*, vol. 144, 1961.

[24] T. Tarnai and Z. Gáspár, "Multi-symmetric close packings of equal spheres on the spherical surface," *Acta Crystallographica Section A*, vol. 43, 1987.

[25] W. Ridgway and A. Cheviakov, "An iterative procedure for finding locally and globally optimal arrangements of particles on the unit sphere," *Computer Physics Communications*, vol. 233, 2018.

[26] G. Riley and T. Henderson, *The ns-3 network simulator*. Springer, Berlin, Heidelberg, 2010.

[27] E. Hamida *et al.*, "Impact of the physical layer modeling on the accuracy and scalability of wireless network simulation," *Simulation*, vol. 85, 2009.

[28] R. Roy, "Random walk mobility," in *Handbook of Mobile Ad Hoc Networks for Mobility Models*. Springer US, 2011.

[29] J. Yoon *et al.*, "Random waypoint considered harmful," in *INFOCOM '03*, Mar. 2003.

[30] A. Freimann *et al.*, "ESTNeT: a discrete event simulator for space-terrestrial networks," *CEAS Space Journal*, vol. 13, 2021.