Accepted for presentation in: IEEE International Conference on Robotics and Automation (ICRA), Montreal, Canada, May 2019.

ATLAS FaST: Fast and Simple Scheduled TDOA for Reliable Ultra-Wideband Localization

Janis Tiemann[1], Yehya Elmasry[2], Lucas Koring[1] and Christian Wietfeld[1]

***Abstract*— The ever increasing need for precise location estimation in robotics is challenging a significant amount of research. Hence, new applications such as wireless localization based aerial robot control or high precision personal safety tracking are developed. However, most of the current developments and research solely focus on the accuracy of the required localization systems. Multi-user scalability, energy efficiency and real-time capabilities are often neglected. This work aims to overcome the technology barrier by providing scalable, high accuracy, real-time localization through energy-efficient, scheduled time-difference of arrival channel access. We could show that simultaneous processing and provisioning of more than a thousand localization results per second with high reliability is possible using the proposed approach. To enable wide-spread adoption, we provide an open source implementation of our system for the robot operating system (ROS). Furthermore, we provide open source access to the raw data created during our evaluation.**

I. Introduction and Related Work

Recent developments in IEEE 802.15.4a ultra-wideband (UWB) transceiver hardware challenged a significant amount of research due to new and precise time of arrival (TOA) estimation capabilities [1]. Especially for robotics and automation, UWB systems yield great potential due to their low deployment cost and few constraints [2]. In contrast to optical motion capture systems, laser scanner or structured light based approaches, no special lighting conditions have to be created for proper operation. Therefore, current research focuses on using UWB systems for indoor mobile robot operation, see [3]. Nevertheless, the predominant way of achieving localization is by using the two-way ranging (TWR) approach, in which two independent nodes obtain the range to each other by measuring the round-trip time of a packet exchange. However, due to the clock-drift during the exchange the localization results of this simple approach tend to be erroneous, see [4]. Therefore, symmetric (or asymmetric) double-sided two-way ranging (SDS-TWR) adds another packet to the exchange and exploits the symmetry of the processing times to eliminate the corresponding measurement error [5]. Other approaches use multiple-acknowledgements (TWR-MA) to calculate the relative drift and compensate for it, see [6]. Although all TWR-based approaches are capable of providing range-measurements required for localization, the linear relationship in terms of channel utilization and energy consumption with the used amount of anchors leads to significant problems when considering multi-user, battery powered applications. Therefore, recent research involved utilizing the time-difference of arrival (TDOA) of the infrastructure-transmitted (T-TDOA) signals at the mobile nodes for highly scalable localization, see [7] [8]. Although this, global navigation satellite system (GNSS) similar, topology scales greatly with increasing user count, the energy consumption at the mobile nodes is increasing with the used amount of anchors. In this work, we propose building upon the TDOA topology in which the mobile nodes transmit a single message and the infrastructure-receives (R-TDOA) as presented in [9]. Wireless clock synchronization is used to achieve a common time base among the clocks of the static infrastructure-based anchor nodes. Yet, random access is incapable of providing guaranteed update rates required by many robotic applications. Furthermore, with an increase of mobile nodes, the quality of the localization will degrade due to missed synchronization frames, see [10]. [11] shows that time-division multiple access (TDMA) operation of the mobile nodes enables high user densities. Therefore, based upon previous work [12] we propose a novel lightweight scheduling protocol that seamlessly integrates with the wireless clock synchronization needed for proper operation. As depicted in Fig. 1 our approach aims to support scalable, low-power, reliable and real-time capable wireless localization. In the course of the following sections we will document the approach and experimentally evaluate it's capabilites in an industrial setting. The open-source implementation provided alongside this work will enable usage of scalable UWB localization in the robotics and automation community.

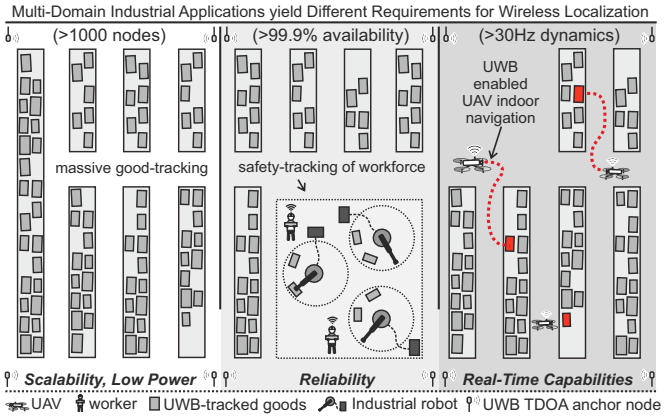

Fig. 1. Illustration of the three main features of the proposed ATLAS FaST scheme. More than 1000 simultaneously active nodes are supported while allowing high update-rates for selected nodes with real-time requirements.

[1]Janis Tiemann, Lucas Koring and Christian Wietfeld are with the Communication Networks Institute (CNI), TU Dortmund University, Dortmund, Germany {janis.tiemann, lucas.koring, christian.wietfeld}@tu-dortmund.de

[2]Yehya Elmasry is with the Department of Electrical Engineering and Computer Science, Northwestern University, Evanston, IL, USA yehyaelmasry2020@u.northwestern.edu

II. Proposed ATLAS FaST Scheme

This work proposes the Fast and simple Scheduled Time-difference of arrival (FaST) scheme as an extension for the previously introduced ATLAS localization system [12]. The localization scheme is based on the R-TDOA topology. By synchronizing and scheduling the mobile nodes, real-time capabilities through guaranteed positioning rates are achieved, which were not available in R-TDOA topologies so far.

A. Physical Layer Channel Access

The basic structure of the scheme is depicted in Fig. 2. The mobile unit is required to synchronize to the static infrastructure in order to align to a common time-base for the whole system.

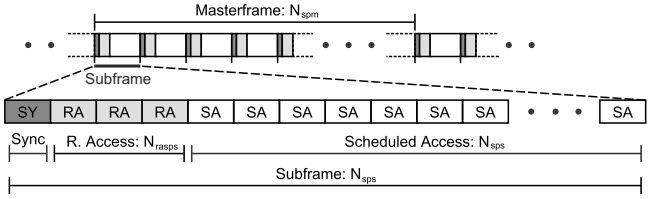

Fig. 2. Schematic illustration of the masterframe structure. Each subframe may hold a sync frame of different sync anchors to enable cascaded clock sync. Each subframe consists of a random and a scheduled access period.

The overall medium access scheme is structured in masterframes which correspond to the wireless clock synchronization period of a single anchor τ_m, see (1). One masterframe holds N_{spm} subframes. Similar to IEEE 802.15.4a beacon superframes, a subframe consists of N_{sps} slots with a duration of τ_s from which N_{rasps} are random access slots that tags can use to register themselves to the system and $N_{sasps} = N_{sps} - N_{rasps}$ are scheduled access slots. For the considered system, each tick has a duration τ_t of 2.0032 ns. Using the exemplary parameter set listed in Tab. I the synchronization period is precisely 125 ms.

$$\tau_m = N_{spm} N_{sps} \tau_s \qquad \tau_s = N_{tps} \tau_t \tag{1}$$

A typical registration procedure for an unassociated tag would follow a simple pattern as depicted in Fig. 3. The tag starts listening until it receives a synchronization frame. Once a synchronization frame is received it accesses one of the provided random access slots within the same subframe using a uniformly distributed random integer $\lfloor U(0,1) N_{rasps} \rceil$. Hence, the waiting period τ_{ra} is calculated as denoted in (2). A typical number for N_{rasps} would be 3

TABLE I
Overall system configuration parameters

type	default	name	description
uint32_t	0x1dc13	N_{tps}	Number of ticks per slot
uint8_t	0x10	N_{sps}	Number of slots per subframe
uint8_t	0x10	N_{spm}	Number of subframes per masterframe

$$\tau_{ra} = (\lfloor U(0,1) N_{rasps} \rceil) \tau_s \tag{2}$$

To lower the overall system complexity, the synchronization request is the same as any other positioning frame transmitted over the UWB channel. It consists of the tags extended unique identifier (EUI) and a sequence number that is increased with every new message as listed in Tab. II. In some cases, this message may be extended using inertial measurement unit (IMU) data or battery status depending on the application.

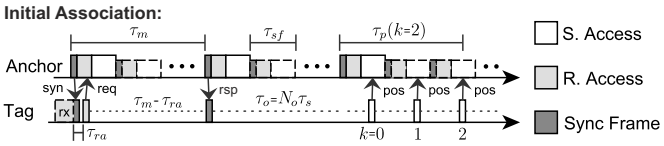

Fig. 3. Schematic illustration of the association procedure. Note that the reception of the two sync frames is also used to calculate and compensate the mobile node's relative clock drift.

After accessing the channel, the tag immediately goes back to sleep and waits for the next sync frame. In the mean time the anchor nodes receive and process the requests. Successful random access requests are propagated to the scheduling engine that has a database of pre-known default period configurations and priorities from which it can prepare a response through the requesting tag's EUI. Through dynamic reconfiguration, parameters can be adjusted in run-time. This procedure simplifies the configuration overhead as the system is infrastructure-based anyway. It also allows for seamless, graceful degradation of the update rate if the overall system capacity limit is reached.

The synchronization node that received the positioning request builds its response into the next sync message as listed in Tab. III. The sync frame holds a dynamic number of response structures. Through the reception of the sync message the tag can see if its association request was successful. In case it was not successful, it needs to start another request. In case of success the tag can read its repetition period and its offset from the structure of the sync message. The period is encoded in through N_p which defines how frequently the tag is allowed to send, see (5). The highest localization frequency that is possible using this scheme is once every subframe at $N_p = 4$. In the default configuration listed in Tab. I this leads to selectable position update rates f_p as defined in (3), starting at 128 Hz.

TABLE II
Structure of the UWB localization frame payload.

type	name	description
uint64_t	eui	EUI of the transmitting node
uint8_t	seq	Sequence number locally increased by the tag
uint8_t	bat	Battery status [1-100] in [%], 0 if no battery is used
float	quat[4]	Orientation in quat. referenced to the magnetic north
float	accel[3]	Acceleration in [m/s^2]

$$f_p = 1/(2^{N_p} N_{sps} \tau_s) \tag{3}$$

In order to tell the tag at which point to send, the offset N_o is transmitted by the sync anchor. This offset value holds two components, the offset within a subframe N_{of} and the offset within the masterframe N_{os} which can be extracted as of (4).

$$N_{of} = N_o \mod N_{sps} \qquad N_{os} = \lfloor N_o / N_{sps} \rfloor \tag{4}$$

Knowing its offset and access rate, the tag is enabled to calculate its positioning schedule offsets from the reception of the last sync message $\tau_p(k)$ for each position k transmitted after synchronization, see (5).

$$\tau_p(k) = ((2^{N_p} N_{sps})k + N_o)\tau_s \tag{5}$$

B. Implications of the Local Clock Drift

The overall system design is optimized for low complexity at the mobile unit and on the UWB PHY. Only two different fundamental UWB messages have to be distinguished in the proposed scheme allowing for simple and battery efficient mobile unit implementation. However, simple and cost-efficient mobile unit implementations are limited in the quality of the associated time-source, see [13]. Typically crystal oscillators with a long term stability of 1 ppm to 10 ppm are used. On the long run, the corresponding clock drift leads to a loss of synchronization within the scheduling scheme. The proposed scheme allows for an elegant way of countering this effect. The difference between the initial reception τ_{sy-1} of a sync message and the response sync message τ_{sy} is measured and compared against the expected synchronization period τ_m as denoted in (6).

$$\tau_{kc} = (1 + \dot{\varepsilon}_s)\tau_k \qquad \dot{\varepsilon}_s = \frac{\tau_{sy} - \tau_{sy-1}}{\tau_m} \tag{6}$$

Through the use of this inherently available estimation of the clock drift $\dot{\varepsilon}_s$ between mobile node and static sync anchor, the major source for de-synchronization can be eliminated up to a certain degree, see [4]. The compensated transmission interval τ_{kc} is calculated as denoted in (6). The difference between slot duration τ_s and packet duration τ_{ppdu} determines the allowed margin within one slot. Hence, the duration τ_{ns} which the mobile node is allowed to send without re-synchronization is calculated as denoted in (7), assuming the drift component $\dot{\varepsilon}_s$ can be compensated and the remaining error can be modeled as the change of the drift over time $\ddot{\varepsilon}_s$.

$$\tau_{ns} = \sqrt{\frac{\tau_s - \tau_{ppdu}}{\ddot{\varepsilon}_s}} \qquad N_{ns} = \lfloor 2^{N_p} N_{sps} \tau_s / \tau_{ns} \rfloor \tag{7}$$

The maximum number of messages that the mobile node is allowed to send after synchronization N_{ns} using compensation is calculated as defined in (7). In a practical system $\ddot{\varepsilon}_s$ will strongly depend on environmental factors such as temperature change. Therefore, the number of allowed position transmission $N_{tx} \leq N_{ns}$ is configured through the sync anchor frame to account for different environmental circumstances, system load and estimated node dynamics, see Tab. III.

C. Reliable Re-association

One masterframe period before completion of the allowed position transmissions, the anchor will automatically generate a response in the upcoming sync frame. If received, the mobile node is able to update it's estimated clock drift and continue sending without having to send and receive another message pair. If missed, a new association has to be triggered by the mobile node. Depending on the requirements of the application, the blind sync response may be sent multiple times before the end of the allowed transmissions to avoid a positioning outage during re-association.

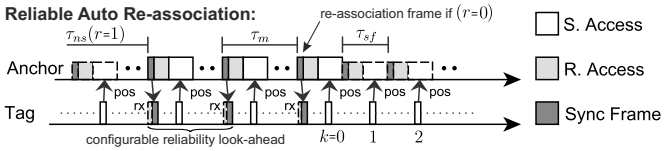

Fig. 4. Schematic illustration of the re-association procedure. Note that a variable number of look-ahead sync frames is used to improve the reliability of the re-association and avoid new associations using random access.

Due to the additional cost introduced through the extended payload of the sync message, the reliability look-ahead r is configurable individually depending on the criticality of the given tag. A schematic illustration of the reliable auto re-association using $r = 1$ is depicted in Fig. 4. In the association response, the highest order two bits are used to set r to either 0,1,2 or 3. For $r > 0$ the number of actual additional receptions are $N_r = 2^r$.

D. Preparation of Distributed Wireless Clock Sync

In order to distribute a common clock among different areas using wireless clock synchronization, a multi-hop clock synchronization algorithm is inherently enabled through ATLAS FaST. Since a masterframe lasts over the period sufficient for wireless clock sync [10], every subframe can be used by a different anchor and hence create a hop in the distributed clock sync. In a practical system the global clock would be reconstructed as depicted in Fig. 6. Each anchor, active or passive, is assigned a sync master. Through

TABLE III
STRUCTURE OF SYNC ANCHOR UWB FRAME PAYLOAD.

type	name	description
uint64_t	eui	EUI of the transmitting node
uint8_t	seq	Sequence number locally increased by the anchor
uint8_t	per	Period in 2^{per} slots at which the sync anchor transmits
uint8_t	nra	Number of random slots N_{rasps} for coming subframe
uint32_t	reui_i	Reduced EUI of i^{th} node from last random access period
uint8_t	$\text{r\|per}_\text{i}$	Reliability look-ahead r and Period N_p and for i^{th} node
uint16_t	off_i	Offset N_o for i^{th} node from last random access period
uint16_t	num_i	Number of allowed pos transmissions N_{tx} for i^{th} node

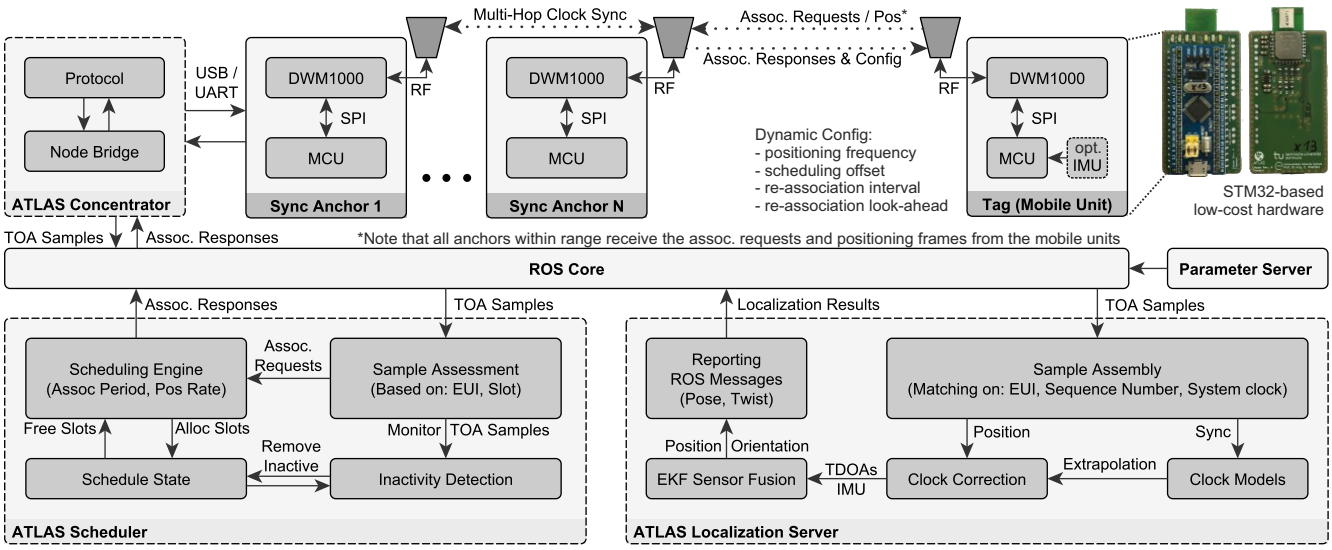

Fig. 5. Schematic illustration of the system architecture for the ATLAS localization system using FaST scheduling. Note that the system is built upon the Robot Operating System (ROS) and is highly modular having no inderdependencies between localization and scheduling.

that assignment a clock sync graph can be constructed. For processing, each master anchor will generate a slave list. Based on the graph, a clock correction path can be built for each anchor. Through this path the global clock can be re-obtained by iterating through the clock offsets and drifts of all anchors in the list. In this step, the individual absolute offsets and drift-induced extrapolated offsets would be accumulated.

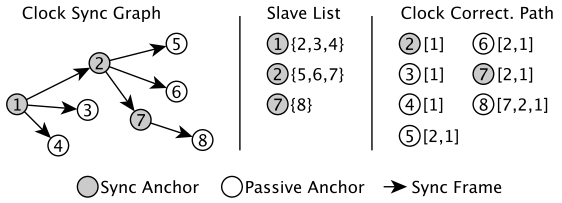

Fig. 6. Schematic illustration of the distributed wireless clock synchronization. A sync graph is used to obtain the slave list of each sync anchor and hence the clock correction path of all participating anchors.

E. Implementation and Architecture of ATLAS FaST

Based on the previously developed ATLAS wireless localization system [12], this work presents the second release by building upon the widely used robot operating system (ROS) [14]. In contrast to previous work which focussed mainly on plain localization aspects, this work provides a significant increase in scalability and real-time capabilities, presenting a method that allows for scalable localization without degradation in the performance of the localization results. The proposed system architecture is depicted in Fig. 5. The ATLAS Concentrators are capable of connecting to multiple anchors and sync anchors over a proprietary binary protocol over USB. To allow for widespread adoption, the hardware design files, for an easy to solder UWB node used in this work are provided as a supplement to this work, see [15]. Furthermore, the ROS-based source-code for the system implementation is provided, see [16].

III. ATLAS FaST Evaluation

In order to assess the performance of the proposed approach a combination of analytical, simulative and experimental evaluation is conducted. The raw data and system configuration of the experimental analysis are provided alongside this work [17]. Using the open-source implementation, the raw data allows for an interactive demo, which will be further highlighted in section III-E.

A. Scalability Evaluation

Due to the slotted approach and the competition-free channel access in the scheduled slots, the successful positioning throughput scales linearily with the positioning frame load as depicted in Fig. 7. Non-reception due to noise induced frame error rates were neglected in this analysis as this will highly depend on the link budget the wireless localization system planner will provide for the given setup. However, the variation of the number of random access slots will define the upper bound of the systems capacity. To compare the multi-user scalability of the proposed approach to the previous

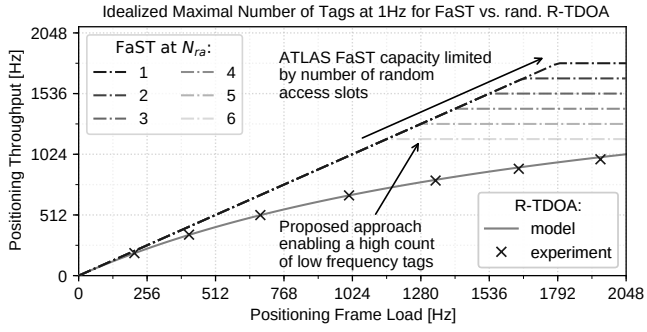

Fig. 7. Exemplary analytical analysis of the successful positioning throughput under a varied positioning frame load. Note that the throughput of the random R-TDOA approach is based on experiments conducted in previous work [10].

R-TDOA system with random channel access, an analytical model for R-TDOA obtained in a scaled experiment with competing tags is depicted in Fig. 7 alongside the throughput of the FaST approach. It is clearly visible that through the scheduled approach, a higher throughput is possible compared to the same underlying positioning load using random access based R-TDOA. Also, the analytical model for the R-TDOA throughput matches the experiment closely. The current implementation only supports static adjustment of N_{rasps}, for optimal performance a dynamic adjustment would be neccesary. However, even when using a static set of six random access slots, FaST easily supports more than 1000 mobile units at 1 Hz.

B. Experimental Evaluation of the Achieved Availability

Due to the high availability requirements of mobile robots using wireless localization as control feedback, keeping the tags associated is a core-design-goal of ATLAS FaST. Therefore, the effect of the reliable re-association setting on the overall mean de-association rate in challenging channel conditions is of high interest. In order evaluate this effect a simple experiment is conducted using a sync-anchor and a set of four tags. Tag A was kept close to the sync-anchor, while B, C and D where placed in different locations, artificially inducing higher frame error rates as depicted in Fig. 8. Here, the mean frame-drop probability is shown over the number of consecutively dropped frames.

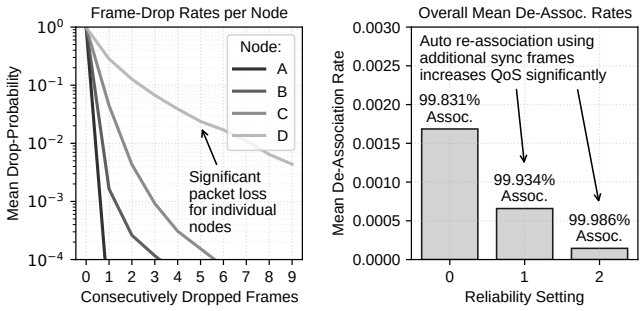

Fig. 8. Experimental analysis of the effect of the reliability setting on the overall mean de-association rates using a set of four nodes with different channel quality characteristics. Note that the extra re-association frames increase the quality of service (QoS) significantly.

In this scenario, the frame drop-rate of tag D is very high at more than 20 % drop-rate per frame. Although the scenario is highly challenging, the reliabile re-association mechanism introduced by FaST is capable of minimizing the frame-drop-induced de-association. As depicted, the mean de-association rates are decreasing significantly with additional re-association frames introduced by a higher reliability parameter. Therefore, it can be stated that the reliability of greater 99.9 % is fully achieved by the proposed approach.

C. Stability under Massive Initial Association

Due to the hybrid nature of the proposed protocol, a random access period is required for initial association of new nodes. The proposed approach makes use of UWB-specific properties that were analyzed in [10]. The sparse

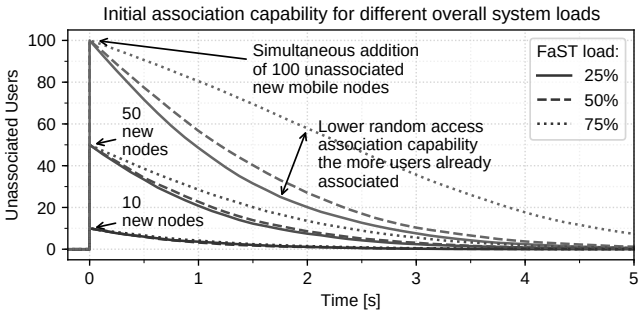

Fig. 9. Simulative analysis of massive initial association, the results show the system's response to simultaneous addition of 100, 50 and 10 unassociated new mobile nodes. Note that a worst-case scenario is evaluated with only one master anchor within the visibility of the new nodes.

nature of UWB frames reduces the probability of destructive collisions on the physical layer. In order to evaluate the effect of massive initial association simulative analysis of the system's response to simultaneous addition of several new unassociated nodes is conducted. The simulation analyzes the behavior for three different load situations of the overall system. A load of 25 % corresponds with N_{rasps} of 12. A total of eight random access slots are available at a load of 50 % and four at 75 % load respectively. A total of 100 simulation runs over 10 s were conducted using a uniformly distributed access to the random access slots for each unassociated node at 1 Hz. Only collision-induced failures were considered using the model developed in [10]. In order to determine the worst-case behavior, a single-cell scenario with visibility of all nodes to one sync anchor with a sync rate of 8 Hz was considered. The results of simultaneous addition of 10, 50 and 100 nodes are depicted in Fig. 9. It is clearly visible, that due to the special characteristics of UWB, the proposed approach is able to handle massive initial association in the random access phase. However, dynamic variation of N_{rasps} is required to obtain a stable system behavior. In realistic deployments, though, the unassociated mobile nodes will mostly be, within visibility of multiple sync anchors and will therefore, be distributed among the available subframes.

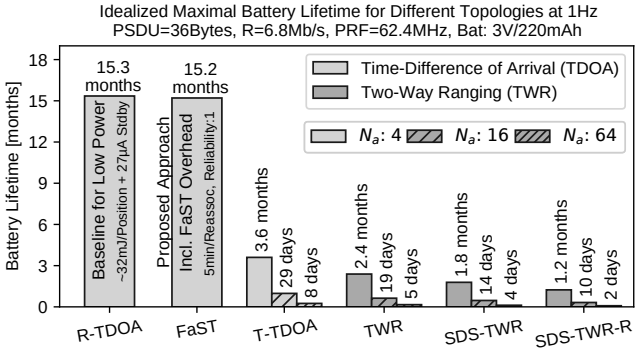

Fig. 10. Exemplary analysis of the battery lifetime for the different topologies. Note that only the transceiver-induced power consumption during tranmission and reception is considered. Peripheral power-losses are neglected due to the strong dependency on firmware-level optimization.

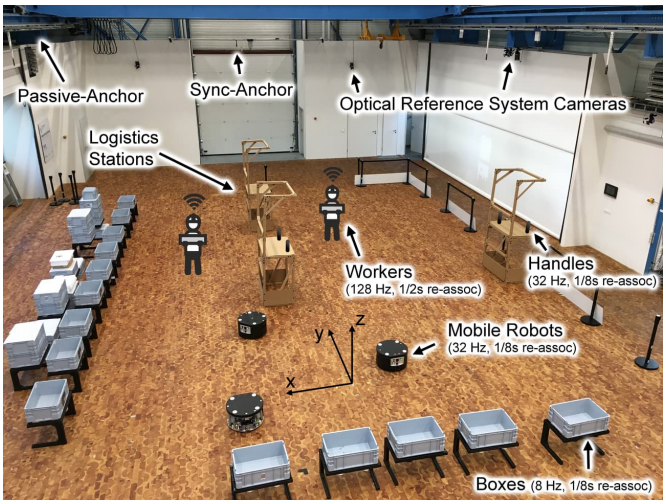

Fig. 11. Picture of the exemplary industrial scenario setup. A set of low update-rate boxes at 8 Hz, high update rate persons at 128 Hz and mobile robots at 32 Hz is tracked simultaneously. Note that lower update rates would be chosen to enable a significantly larger amount of tracked assets.

D. Trade-Off in Energy Consumption

In order to evaluate the energy consumption compared to other channel access schemes, a simple analytical study is conducted. By calculating the power per transmitted and received UWB frame for the transceiver system and the fundamental requirements for the localization topologies an idealized maximal battery lifetime can be calculated. An exemplary positioning rate of 1 Hz, a re-association period of 300 s and a reliability look-ahead $r = 1$ was chosen. The power consumption is calculated based on [18] resulting in P_{TX}=32 mJ and P_{RX}=59 mJ. The resulting battery lifetime is depicted in Fig. 10. Here, the battery-lifetime of the T-TDOA, and TWR-based approaches is lower by orders of magnitude due to the dependency on the number of anchor nodes N_a which is inherent to those topologies. For the R-TDOA based approaches, it is clearly visible that FaST results are close to those of random-access, which is the baseline for low power consumption at the mobile node due to no additional overhead. Therefore, it can be stated that through the planned scheduling with infrequent synchronization, FaST is capable of tracking many low-power devices without interfering with the real-time requirements for critical applications.

E. Exemplary Industrial Application

To evaluate the systems performance, an experiment in an industrial environment is conducted. As depicted in Fig. 11 the experimental setup comprises a set of three mobile robots at 32 Hz, four mobile handles at 32 Hz, 14 boxes at 8 Hz and two workers at 128 Hz. The area is covered by an optical reference system using 38 cameras and the proposed wireless localization system with one synchronization anchor and eight passive anchor nodes, see Tab. IV. ATLAS FaST is deployed with pre-configured settings (update rate and re-association rate) for all mobile nodes.

As a supplement to this work, an example configuration and replay-able dataset (TOA data and ground-truth) for this experiment [17] is provided as a ROS bag in addition to the source code [16] allowing for a demonstration where localization-filter parameters can be changed interactively. Furthermore, a video illustrating the core-functionality, dynamic reconfiguration of mobile nodes and showing the experiment of the provided dataset is provided. It should be noted though, that the experiment is provided as a proof of concept as the focus is on the multi-user scalability and not on the accuracy. For accuracy evaluation, which has been conducted sufficiently in [19] in a single-Tag scenario, static offsets for the TDOAs need to be considered and a translatory offset for the ground-truth due to the different origins needs to be taken into account.

TABLE IV
POSITIONS OF THE ANCHOR NODES USED IN THE EXPERIMENTS.

anchor	1	2	3	4	5	6	7	8	9
x $[m]$	15.4	10.45	4.95	-0.10	-6.10	-6.10	-0.10	4.95	10.45
y $[m]$	0.40	-5.80	-5.80	-5.80	-5.80	6.10	6.10	6.10	6.10
z $[m]$	5.00	5.00	5.00	5.00	5.00	5.00	5.00	5.00	5.00
mode	sync.	pass.	pass.	pass.	pass.	pass.	pass.	pass.	pass.

IV. Conclusion and Future Work

This paper proposes a novel approach to achieve highly scalable and energy efficient wireless localization while providing guaranteed real-time capabilities. Next to the approach, we provide a ROS-based open-source implementation of the scheduling and the second release of the underlying ATLAS localization system. We were able to analytically illustrate the systems potential and experimentally show the system performance. Furthermore, we provide the raw experimental data, a video and an interactive demonstration (by re-playing the raw data with the provided software) of multi-asset tracking in an industrial scenario showing the system's capabilities.

Acknowledgment

The work on this paper has been partially funded by Deutsche Forschungsgemeinschaft (DFG) within the Collaborative Research Center SFB 876 "Providing Information by Resource-Constrained Analysis", project A4 and was supported by the federal state of Northrhine-Westphalia and the "European Regional Development Fund" (EFRE) 2014-2020 in the course of the "CPS.HUB/NRW" project under grant number EFRE-0400008. The stay of Mr. Elmasry at TU Dortmund was supported by the German Academic Exchange Service (DAAD) as part of the Research Internships in Science and Engineering (RISE) program. We would like to thank Prof. Dr. Dr. h.c. Michael ten Hompel and Dipl.-Inform. Moritz Roidl for their kind support with the experimental setup by providing access to the Logistics Innovation Lab "Innovationslabor für Hybride Dienstleistungen in der Logistik (LIDO) in Dortmund" funded by the German Federal Ministry of Education and Research (BMBF) under grant number 02P16Z201.

REFERENCES

[1] S. Gezici, Zhi Tian, G. B. Giannakis, H. Kobayashi, A. F. Molisch, H. V. Poor, and Z. Sahinoglu. Localization via ultra-wideband radios: a look at positioning aspects for future sensor networks. *IEEE Signal Processing Magazine*, 22(4):70–84, July 2005.

[2] A. Alarifi, A. Al-Salman, M. Alsaleh, A. Alnafessah, S. Al-Hadhrami, M. A. Al-Ammar, and H. S. Al-Khalifa. Ultra wideband indoor positioning technologies: Analysis and recent advances. *Sensors (Basel, Switzerland)*, 16(5):707, May 2016.

[3] K. Guo, Z. Qiu, C. Miao, A. H. Zaini, C. Chen, W. Meng, and L. Xie. Ultra-wideband-based localization for quadcopter navigation. *Unmanned Systems*, 04(01):23–34, 2016.

[4] A. Dongare, P. Lazik, N. Rajagopal, and A. Rowe. Pulsar: A wireless propagation-aware clock synchronization platform. In *2017 IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS)*, pages 283–292, April 2017.

[5] D. Neirynck, E. Luk, and M. McLaughlin. An alternative double-sided two-way ranging method. In *2016 13th Workshop on Positioning, Navigation and Communications (WPNC)*, pages 1–4, Oct 2016.

[6] M. W. Mueller, M. Hamer, and R. D'Andrea. Fusing ultra-wideband range measurements with accelerometers and rate gyroscopes for quadrocopter state estimation. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1730–1736, May 2015.

[7] A. Ledergerber, M. Hamer, and R. D'Andrea. A robot self-localization system using one-way ultra-wideband communication. In *Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ International Conference on*, pages 3131–3137, Sep 2015.

[8] M. Hamer and R. D'Andrea. Self-calibrating ultra-wideband network supporting multi-robot localization. *IEEE Access*, 6:22292–22304, 2018.

[9] C. McElroy, D. Neirynck, and M. McLaughlin. Comparison of wireless clock synchronization algorithms for indoor location systems. In *2014 IEEE International Conference on Communications Workshops (ICC)*, pages 157–162, June 2014.

[10] J. Tiemann, F. Eckermann, and C. Wietfeld. Multi-user interference and wireless clock syncronization in TDOA-based UWB localization. In *2016 International Conference on Indoor Positioning and Indoor Navigation (IPIN)*, Alcalá de Henares, Madrid, Spain, Oct 2016.

[11] M. Ridolfi, S. Van de Velde, H. Steendam, and E. De Poorter. Analysis of the scalability of UWB indoor localization solutions for high user densities. *Sensors (Basel, Switzerland)*, 18(06):1875, 2018.

[12] J. Tiemann, F. Eckermann, and C. Wietfeld. ATLAS - an open-source TDOA-based ultra-wideband localization system. In *2016 International Conference on Indoor Positioning and Indoor Navigation (IPIN)*, Alcalá de Henares, Madrid, Spain, Oct 2016.

[13] D. W. Allan. Time and frequency (time-domain) characterization, estimation, and prediction of precision clocks and oscillators. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 34(6):647–654, Nov 1987.

[14] M. Quigley, K. Conley, B. P. Gerkey, J. Faust, T. Foote, J. Leibs, R. Wheeler, and A. Y. Ng. ROS: an open-source robot operating system. In *ICRA Workshop on Open Source Software*, 2009.

[15] J. Tiemann. ATLAS second revision hardware design files, http://dx.doi.org/10.5281/zenodo.1419607. Sep 2018.

[16] J. Tiemann and Y. Elmasry. ATLAS localization framework ROS package, http://dx.doi.org/10.5281/zenodo.2579935. Feb 2019.

[17] J. Tiemann, Y. Elmasry, and L. Koring. Example settings and raw data, http://dx.doi.org/10.5281/zenodo.2579940. Feb 2019.

[18] DecaWave Ltd. *DWM1000 IEEE 802.15.4-2011 UWB Transceiver Module Datasheet 1.7*, 2016.

[19] J. Tiemann, L. Koring, P. Gorczak, and C. Wietfeld. Improving the robustness of control-grade ultra-wideband localization. In *3rd IFAC Conference on Embedded Systems, Computer Intelligence and Telematics (CESCIT 2018)*, Faro, Portugal, Jun 2018.