A Wireless MAC Method with Support for Heterogeneous Data Traffic

Pablo Gutiérrez Peón*†, Elisabeth Uhlemann†, Wilfried Steiner*, and Mats Björkman†
*TTTech Computertechnik AG, Vienna, Austria
†School of Innovation, Design and Engineering, Mälardalen University, Västerås, Sweden
Email: {pablo.gutierrez-peon, wilfried.steiner}@tttech.com, {elisabeth.uhlemann, mats.bjorkman}@mdh.se

Abstract—**An important aspect of any communication technology is its medium access method, being responsible for sharing the medium among users. For delay-sensitive applications, such as industrial control systems, timely and reliable channel access is of essence. Hence, existing protocols like HART and TTEthernet use a time-triggered medium access approach. However, as the nature of industrial control systems change and evolve, there is a growing need to provide wireless access and support data traffic with mixed requirements. While technologies like WirelessHART can offer timely and reliable access to the wireless medium, only one type of data traffic is usually supported. In this paper, we therefore propose and evaluate three different medium access methods for wireless communications, all capable of supporting three different data traffic classes: time-triggered, rate-constrained and best-effort traffic. In particular, different options on how to handle best-effort traffic, using scheduled time-slots or contention, are evaluated, showing for all the proposals different drawbacks and benefits depending on additional requirements on e.g., hardware, protocol overhead and resistance to interference.**

I. Introduction

Ethernet is the most spread standard for local area networks (LAN). It was originally designed for office environments, but its success in this application domain caused Ethernet to be adopted in other scenarios like industrial control applications [1]. The typical communication requirements of industrial applications are high reliability and predictable latency regarding data delivery [2], both difficult to support for a best-effort protocol like Ethernet. Consequently, many Ethernet-inspired technologies are available with add-ons to satisfy different types of industrial requirements. One of them, TTEthernet [3], is based on the time-triggered communication paradigm to provide deterministic data delivery with low latency.

TTEhernet has been successfully applied in several use cases, one important example is being the backbone for the new generation of spacecrafts developed by NASA (the Orion Multi-Purpose Crew Vehicle). Nevertheless, TTEthernet is a wired technology, and communication without wires is thought to cover several new application use cases that are of great interest. Among its advantages, wireless links allow the possibility of connecting mobile components. Reduced wiring and installation costs, easier network deployment, easier maintenance, and the lack of problems like broken wires are important benefits too. Wireless communication technologies also enable new use cases, like quick deployment of temporary monitoring and diagnostic networks in machinery. These advantages are not new to industry and have already been envisaged by standards like WirelessHART [4].

In WirelessHART, the wired HART fieldbus protocol was extended with wireless capabilities providing a multi-vendor, inter-operable wireless standard for secure and reliable industrial communications. However, only one type of data traffic, namely time-triggered periodic messages, is typically supported in WirelessHART. A subset of all time-slots can also be shared dynamically, but through contention only, implying a random channel access delay. To guarantee a deterministic delay, all periodic traffic flows have to be pre-scheduled, consequently limiting the support for dynamic traffic flows.

In contrast, TTEthernet can support three different types of traffic with different quality-of-service levels. This allows to integrate several traffic flows used by applications with mixed-critically communication requirements. Traffic prioritization is imposed between the flows such that critical delay-sensitive traffic is completely deterministic, while low priority traffic can suffer from starvation. The highest priority traffic flow is the time-triggered (TT) traffic class, intended for time-critical communication, with predefined instants in time when a message must be sent and received. The next priority level is given to the rate-constrained (RC) traffic class, which supports traffic flows that are characterized by an average bandwidth requirement, usually employed for data streaming. Finally, the best-effort (BE) traffic class is for asynchronous traffic (legacy Ethernet traffic) and is sent without any delivery guarantee whenever no TT or RC packets are transmitted. The support of flows with different requirements under the same infrastructure is of great value for industry, since emerging industrial control applications need to combine periodic traffic for monitoring and control with event-driven traffic from autonomous or mobile nodes, if possible using standard-based technologies.

A key aspect of any wireless solution with industrial requirements is the media access control (MAC) protocol. In essence, it is responsible for sharing the medium among users. To give support to traffic flows from time-critical applications, deterministic medium access is required. This means that any device that tries to access the medium will have a guarantee in terms of an upper-bounded channel access delay. Furthermore, the jitter, i.e., the difference between the minimum and maximum access delay is also a requirement for periodic traffic.

This paper proposes three different MAC methods with deterministic channel access delay, as in WirelessHART, but with support for the three different traffic classes of TTEthernet namely TT, RC and BE. Different options on how to handle best-effort traffic, using scheduled time-slots or contention, are evaluated, and the three MAC proposals all show different drawbacks and benefits in terms of best- and worst-case channel access delay and utilization. In addition, a set of non-functional requirements, such as required changes to existing hardware, protocol overhead and resistance to interference are also discussed. We aim for a solution that is implementable on top of standardized protocols such as IEEE 802.11 and IEEE 802.15.4 without requiring any changes to the hardware.

The rest of the paper is structured as follows. Section II covers the related work, with special attention given to the MAC layer and solutions to provide deterministic access. The characteristics and requirements of the three different data traffic classes are explained in Section III. The set of evaluation criteria for the MAC methods is in IV. Our three proposed wireless MAC schemes with support for heterogeneous data traffic are detailed in Section V and evaluated in Section VI. Finally, we conclude the paper in Section VII.

II. Related Work

One very widely used MAC protocol in wireless LAN is carrier sense multiple access (CSMA), present in standards like Ethernet, CAN, IEEE 802.11 and IEEE 802.15.4. CSMA senses if the medium is free for a duration of an Interframe Space (IFS) before transmitting. If the medium is occupied during the IFS, the node defer its access until it becomes free again. After IFS, several nodes could try to transmit at the same moment, causing a collision. Hence, a random backoff value is also added to the IFS. Clearly, this behaviour does not fulfil the requirements of bounded access delay that industrial applications have, as the carrier sense function makes channel access random and unpredictable.

In the CAN bus protocol nodes are connected to the same bus and transmit event-driven traffic using CSMA after a deterministic arbitration phase based on node priorities. Unfortunately, this arbitration mechanism does not work in wireless settings due to the need to listen while transmitting. FlexRay is another wired technology that supports time- and event-driven traffic by the definition of static and dynamic phases that are cyclically repeated. During the static time-triggered phase, time-slots are used based on a previously agreed schedule. The event-driven phase is also divided into time-slots, but these are assigned dynamically according to priorities that are given by the message ID.

The IEEE 802.15.4 is a wireless standard which works with CSMA, but also defines a non-mandatory phase, called Guaranteed Time-Slots (GTS). These GTS are assigned to nodes that previously requested them from a special node, the network coordinator. When a GTS is assigned, all nodes are aware of this so no other device than the granted one will try to transmit during this time-slot. However, since the GTS are reserved during the CSMA phase when channel access is random, it is not certain that a GTS can be assigned and thus the channel access delay is still random and unbounded.

The IEEE 802.11 standard generally uses CSMA, but also has a guaranteed access functionality. This is performed via the point coordination function (PCF) in IEEE 802.11, which is basically a polling mechanism placed on top of CSMA. By using a shorter IFS, the so-called point coordinator (PC) will get prioritized access to the channel. Channel access is then granted by the PC by sending a poll frame to a node, which permits the node to transmit a frame to the PC. In case the polled node does not have any frames to send, it must transmit a null frame. Consequently, this method implies a lot of overhead and in some cases wasted resources. In addition, PCF suffers from problems with hidden nodes. The IEEE 802.11e amendment defines new mechanisms intended for an improved quality of service: Enhanced Distributed Channel Access (EDCA) and HCF Coordinated Channel Access (HCCA). EDCA is based on different IFS values (named AIFS) as well as different backoff values to provide several priority levels. HCCA defines separated contention and contention-free periods inside a superframe structure, where the contention-free phase is based on a polling mechanism. Both EDCA and HCCA are of limited use for industrial traffic. The first one due to the underlying CSMA, that cannot support determinism. The second one due to the polling communication overhead [5] and the fact that it is not commercially implemented yet.

In contrast, time division multiple access (TDMA) is a MAC mechanism able to provide deterministic channel access. It divides time into slots that are assigned to nodes. This implies that some sort of central coordinator is needed if time-slots should be assigned during run-time, or that all time-slots need to be assigned offline. Another consequence of the use of TDMA is the need to share a common time notion among the involved nodes, usually via a clock synchronization mechanism in order for them to agree on the delimitation of time-slots.

WirelessHART uses IEEE 802.15.4, and thus CSMA, but places a TDMA scheme on top of the built-in MAC. Transmissions are then performed using time-slots that are assigned to time-triggered traffic offline. These time-slots can also be shared dynamically though contention, so channel access delay is random.

IsoMAC [6] is an example of a MAC protocol that uses TDMA on top of IEEE 802.11. It has a centralized manager that assigns time-slots to the nodes based on the requirements that they specify in resource request messages, sent using CSMA. The coordinator decides about the requested traffic flows based on the remaining resources and constructs the communication cycles that start after every beacon frame. Beacons are used to convey the traffic schedule, and time synchronization information for the current cycle. Similarly to IEEE 802.15.4, a contention and a scheduled phase are defined. The contention phase is based on ordinary CSMA and is used for BE traffic and to request time-slots. The scheduled phase is in turn divided into time-slots that are assigned to the admitted flows. Like with PCF, a shorter IFS is used to get access to scheduled time-slots. IsoMAC has the same problem with random and unbounded channel access delay as IEEE 802.15.4. Devices asking for time-slots have to do it during a contention phase.

III. Heterogeneous Data Traffic

The time-triggered paradigm [7], the foundation for TTEthernet and WirelessHART, states that messages must be sent according to a strict time schedule. This has the great advantage of determinism, making it very appropriate for safety-critical systems. A benefit of TTEthernet is that it supports time-triggered traffic flows, and integrates it with rate-constrained traffic, guaranteeing the isolation of the first one with respect to the second. In TT flows, all messages are sent at predefined points in time. In case the application that booked these time-slots decides not to use some of them, event-triggered flows could take their place. In RC traffic flows, there is a guarantee in terms of a minimum inter-arrival time and sufficient bandwidth, such that delay and jitter have defined limits. RC messages are not sent at predefined points in time and its arrival pattern is asynchronous. BE traffic is random and there is no guarantee at all for BE messages to be delivered. All three traffic classes have separated buffers from where messages are taken to be sent on the network.

An offline computed schedule is spread to the switches and end systems that constitute a TTEthernet network, in order to make them aware of the instants when they have to take TT messages from the buffers and place them on the output ports. The same applies to the instants when they should expect the reception of TT messages. For RC messages, an average bandwidth is assured by the scheduler, implying that TT messages leave enough unused time for guaranteeing this RC bandwidth. In the worst case, RC traffic is periodic, so the scheduler can leave free slots periodically after allocation of TT slots. BE messages are placed during run-time in the free periods of time left by TT and RC. Clock drift [8] can affect the rate at which RC traffic is generated, meaning that in the worst case a RC sender has more messages to send than allocated RC slots when working under an early clock. However, for reasons of simplicity we assume that clock drifts are negligibly and that TT and RC will be periodic. The scheduling problem of TT and RC messages can be modelled via first-order logic constraints [9], that reflect the network infrastructure and the traffic flows. An important aspect of the schedule is its period of validity. The schedule will be run cyclically and is valid for all the time the network is up and running.

IV. Evaluation Criteria

This section provides a set of evaluation criteria, targeting different aspects to consider when developing MAC protocols capable of supporting heterogeneous traffic over wireless links.

Protocol overhead and efficiency. A MAC protocol based on TDMA introduces a relatively small overhead if compared to polling- or token-based MAC mechanisms. This criterion is intended to evaluate the amount of control traffic introduced by the proposed MAC protocol, as well as any potential unused pre-scheduled resources. For example, this can be evaluated by comparing the difference in utilization between MAC methods.

Channel access delay. For periodic traffic, the channel access delay should be known and constant, so that the jitter can be set to zero. However, for BE traffic, the time until channel is granted is an important performance measure. Hence, since some of our proposed MAC schemes use contention for BE traffic, the best- and worst-case delay until a node can transmit in the contention phase is evaluated.

Reliability. The reliability of the wireless medium is lower than the wired, mainly due to multipath fading, shadowing, and interferences from other devices. Therefore, it is likely needed to add mechanisms to increase the reliability in the wireless segment to make a better balance towards the wired one. Replication, that is to send more than one copy of the data to be delivered, is an example of a mechanism for increased reliability. Hence, this feature relates to the ease of adding extra redundancy to the MAC method for increased reliability.

Hardware requirements. The proposed MAC protocol is thought to be implemented over standard IEEE 802.11 or IEEE 802.15.4 hardware. Both standards have native support for CSMA, making it a tractable option for contention. Some of the standard compliant chipsets allow to make changes on the MAC layer, but the most likely option is to place the proposed MAC on top of the standard MAC. Besides, CSMA parameters can be tuned to reduce the overhead of an extra layer. This criterion evaluates the amount of hardware changes needed.

Integration with wired networks. The traffic integration between wired networks such as TTEthernet and a wireless complement can be achieved thanks to the way the TT traffic is scheduled, using first-order logic constraints. Besides, wired and wireless segments will have different time-slot sizes tied to the particularities of the medium. This can be solved by the addition of constraints to the scheduler to place traffic on slots with different size. The singularity of the wireless MAC protocols can also end in the addition of new constraints to the scheduler, e.g. telling the scheduler not to place traffic when a beacon is to be sent, or booking more slots for increased reliability on the wireless medium. This criterion evaluates how easy integration with TTEthernet can be made.

Interferences from other CSMA devices. Other nodes working on the same frequency as our network can cause interference, but if they use CSMA, they have to listen before they transmit. This aspect can be used to get faster access to the medium if waiting less time before a transmission. This criterion relates to the sensitivity of the proposed MAC to interference from other nodes using CSMA.

V. Proposed Wireless MAC

Due to the requirement of deterministic access to the wireless medium, we will use a centralized network topology with an access point (AP). All messages must go though the AP which can also connect to a potential wired segment. In addition, the selected infrastructure network topology is thereby similar to e.g., the wired TTEthernet switch topology.

We can guarantee deterministic access to the wireless medium if we propose a MAC mechanism that has an upper-bounded channel access delay. Based on the outlined evaluation criteria, we have opted for a TDMA-based approach. The resource that the wireless scheduler will manage is then time-slots. Dividing periods of time into slots fits easily with the time-triggered paradigm and the requirement of deterministic access to the medium. The actual allocation of the three different TTEthernet traffic flows to those time-slots opens a range of possibilities with different implications, from which three of them are selected in this paper. In all of them, the TT traffic class will be scheduled offline, guaranteeing its deterministic access to the wireless medium, similarly to what is done in WirelessHART. In the worst case, RC flows are periodic, so they can be modelled as TT flows and therefore added to the scheduler offline. Based on the individual periods of the TT and RC traffic, a hyperperiod, being the least common denominator of the individual periods, can be defined. In the case of BE, the packets are just placed in the internal queues of the nodes by the user applications at runtime, without any

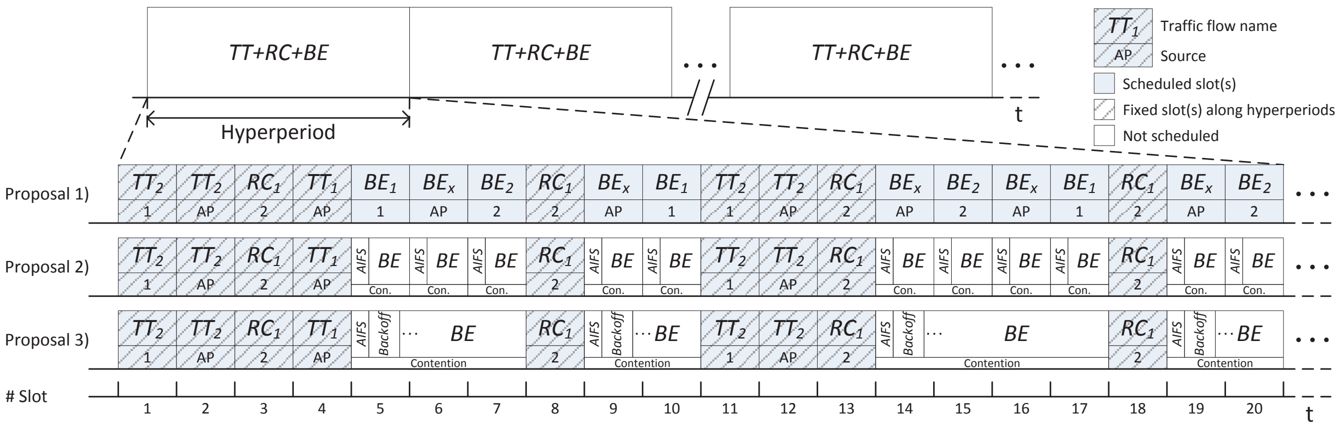

Fig. 1. Example of slot allocation for the three MAC proposals.

predefined traffic pattern. Their characteristics are not known at the moment of creating the schedule, so it is not possible to assign time-slots to specific frames. As a consequence, its access will not be deterministic, but still this is consistent with the definition of BE traffic in TTEthernet. The schedule will be repeated continuously during the system run-time. It is important to notice that the receiver in a slot does not need to be specified for the wireless schedule due to its broadcast nature. The only time it could be good to schedule a receiver is to preserve energy, but this is out of the scope of the current paper. Also, by allowing more than one receiver, the future implementation of steps to increase reliability is left opened.

1) Pre-scheduled time-slots only: This approach does an offline allocation of time-slots for all three types of traffic and all transmitters (i.e., in each slot, there is one unique transmitter and one unique traffic class allowed). The specific allocation of slots for TT and RC will be given by the scheduler. Next, BE traffic will be placed on the remaining slots after the allocation of TT and RC traffic. As BE traffic characteristics are not known in advance, the remaining slots cannot be directly assigned, so the most fair way of sharing these resources is using a round-robin schedule, giving the same amount of time-slots to each node in a circular order. This approach could also be implemented in WirelessHART, with the difference that all pre-scheduled slots in WirelessHART are typically placed at the beginning of the superframe, while our proposal may allow a more even distribution of BE traffic, yielding a shorter delay until access is granted. Still, this scheme can be seen as a benchmark of how WirelessHART would perform.

2) Pre-scheduled time-slots and contention-based time-slots: This approach is similar to Proposal 1, but the remaining time-slots not used by TT and RC can be used for BE traffic flows following a contention process, similar to the shared slots in WirelessHART. Further, one single node in each slot is given a higher priority, using a higher priority class in IEEE 802.11e, yielding a shorter T_{AIFS}. Thereby prioritized access is given to nodes in a round-robin fashion. However, using CSMA-based contention, all nodes waiting exactly for a AIFS in the beginning of a slot is not recommended, as when more than one node wants to transmit in a slot, they will collide when starting to transmit. Therefore, we will add a random time in addition to the AIFS to help to avoid collisions. However, the high priority node will not add any backoff. The approach is similar to the dynamic phase in FlexRay, but different from WirelessHART, where shared contention slots use slotted ALOHA.

3) Pre-scheduled time-slots and contention-based phase: Here the proposed mechanism establishes a contention-based continuous phase for BE traffic, and a scheduled phase for TT and RC traffic. The ratio of the contention-based phase and the scheduled phase can be determined using schedulability analysis [10]. The two extreme cases are that all TT and RC traffic is grouped together in one long scheduled phase, leaving the reminder of the hyperperiod to the contention-based phase, or alternatively, that each TT and RC instance is evenly spread in the hyperperiod, leaving the space in between for BE traffic. During a contention phase, access is granted after waiting a AIFS plus a random backoff value. The benefit of having the scheduled phases evenly spread over the hyperframe, is that the best-case channel access delay is reduced for BE. The drawback is that, as the contention phase size is multiple of a scheduled slot size (to facilitate scheduling), the spaces between scheduled slots may not be big enough to accommodate the time it takes for contention and transmission. This MAC approach is similar to IEEE 802.15.4 GTS, but with the advantage that slots are guaranteed to be booked as they are not reserved during a contention phase.

An example showing the three different MAC proposals is depicted in Figure 1, using the simple traffic scenario in Figure 2. The traffic scenario accounts two TT and one RC traffic flows expressed using slots as the time unit. TT flows are characterized by their period. RC flows are characterized by a bandwidth, expressed as the maximum number of slots that can be needed in a period (e.g. 1 out of 5 slots for the RC1 flow). In the worst case (maximum rate), RC1 will use 1 slot every 5, so it can be modelled as a periodic message with period 5. The TT together with the RC traffic from Figure 2 results in a hyperperiod of 20, meaning that the slot allocation for TT and RC will be repeated every 20 slots. Note that the slot size between proposals is not the same.

VI. Performance Evaluation

Protocol overhead and efficiency. We have used schedulability analysis [11] to calculate the minimum bandwidth

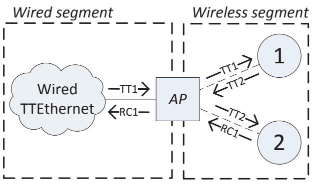

Name	Characteristics	Path
TT1	Period: 20 slots	Wired→AP→1
TT2	Period: 10 slots	1→AP→2
RC1	Rate: max 1/5 slots	2→AP→Wired
BE	Unknown pattern	-

Fig. 2. Example of wireless traffic scenario.

that needs to be reserved for TT and RC traffic while still guaranteeing its timely delivery. The time left for BE traffic can thereby be determined. Schedulability analysis was first conceived for task scheduling in processors, but can easily be adapted to scheduling of messages that share a wireless link. However, this requires the scheduling policy to be earliest deadline first (EDF), i.e. when selecting the next message to transmit, the message closest to its deadline is selected. This will be automatically enforced in our MAC proposals, since the assignment of time-slots to nodes is made offline.

A periodic message is defined by a duration C, a period T, and a deadline D, where D is usually equal to T. The analysis assumes preemptive messages, but as this is not possible in communications, further delays must be taken into account. Therefore, the original deadline is reduced to d by subtracting the maximum blocking time a message can suffer $T_{blocking}$, that is given by the transmission time of a message T_{transm}, and the propagation of the packet itself T_{prop}. For simplicity we assume that T_{prop} is negligible. Thus $d = D - T_{transm}$.

The schedulability analysis consists of two steps and checks if two conditions are met. The first condition says that the utilization U of the wireless link must not exceed one. The second says that the workload function $h(t)$, that is the sum of the transmission times of all messages before instant t, must be less than or equal to t. Given a set of periodic messages $i = 1, .., N$, each one with duration C_i, period T_i, and reduced deadline d_i, the utilization of the link U is given by:

$$U = \sum_{i=1}^{N} \frac{C_i}{T_i} \leq 1, \tag{1}$$

while the workload function $h(t)$ is given by:

$$h(t) \geq \sum_{i=1}^{N} \left\lfloor \frac{t + T_i - d_i}{T_i} \right\rfloor C_i. \tag{2}$$

A numerical evaluation of the utilization using Matlab has been conducted to obtain the time remaining for BE after allocation of the scheduled TT and RC messages for each of the three MAC proposals. A set of TT and RC messages corresponding to low, medium and high utilization (occupying close to 25%, 50% and 75% of a total time of 10000 μs), is created to determine the corresponding utilization in each of the three MAC proposals. A relatively small packet with size of 62 bytes is used for all flows and the data rate is 6 Mbps. These numbers were selected to reflect that packets in industrial networks are usually small and reliability is often more important than throughput. For the schedulability analysis we have selected $T_{AIFS} = 2\mu$s for contention, representing the highest priority class in IEEE 802.11e and thereby introducing less protocol overhead. Given the above numbers, the size of the slot in Proposal 1 and Proposal 3 is $T_{slot_1} = T_{slot_3} = T_{transm}$ since contention is not needed in these slots and the propagation time is neglected. For Proposal 2 the slot size is $T_{slot_2} = T_{transm} + T_{AIFS}$, since all slots need to be of the same size to facilitate synchronization and contention is needed in some of them. The resulting time left for BE traffic transmission for the different MAC proposals, named BE busy time, is shown in Figure 3. We can see that the contention mechanisms introduce overhead in both Proposal 2 and Proposal 3. For Proposal 3, it depends on how the scheduled slots are distributed, evenly over the entire superframe (maximum fragmentation) or grouped together (minimum fragmentation). Both options are shown in Figure 3. With only one contention period, Proposal 3 is better than Proposal 2 because there is more time left for BE during the hyperperiod due to $T_{slot_3} < T_{slot_2}$. However, when having the maximum number of contention phases, Proposal 3 suffers as for medium- and high-scheduled traffic load, all the remaining space can be considered as dead and not valid for BE traffic transmission. This is due to the fact that the minimum size a BE phase must have to allow a transmission is $T_{transm} + T_{AIFS}$, which corresponds to a node that waits the AIFS and randomizes a backoff value of 0. This means that a contention phase of at least two slots must be left between a pair of scheduled slots if a contention-based transmission should fit. Note also that the effective utilization for BE traffic depends on its distribution among the nodes. For Proposal 1, performance is reduced as some scheduled slots are not used if they are assigned to nodes which do not have BE traffic.

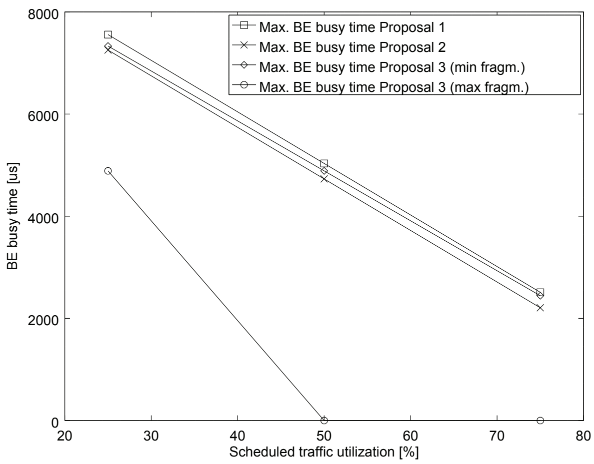

Fig. 3. Best-effort busy size per scheduled traffic load for the three MAC proposals given a hyperperiod of 10000μs.

Channel access delay. We have analysed the channel access delay that a node can experience in the best and worst case for BE traffic. In Proposal 1, the best case is to not have to wait at all, while the worst-case access delay for BE traffic occurs when a message is generated just after the node's scheduled BE slot and this is followed by a period of consecutive slots scheduled for TT and RC traffic. The node then has to wait until all scheduled slots S have finished, and then it has to wait for all the other $N-1$ devices to get access to their BE slots. We therefore have $T_{BE\_worst\_delay_1} = (S + N - 1)T_{transm}$ In Proposal 2, the best case for BE is immediate access to the slot, implying a delay of T_{AIFS}, while the worst case is when the message is generated just after the longest possible consecutive scheduled phase, as the node then has to wait until all scheduled slots S have finished and all other BE slots where other nodes have higher priority before it can begin the contention. Note that the worst case is that all other nodes have BE traffic and they have higher priority in the respective slot: $T_{BE\_worst\_delay_2} = (S + N - 1)(T_{transm} + T_{AIFS})$. In Proposal 3, the best case for BE traffic is to access directly after waiting T_{AIFS}, while the worst case is that access is never granted and thus the delay is, in theory, infinite.

Given the expressions above, Proposal 1 and 2 have a similar worst case, except for T_{AIFS} due to contention, whereas Proposal 3 have an unbounded worst case. However, the worst case in Proposal 3 has low likelihood of occurring in practice, and for the fragmented case, it is not necessary to wait for the entire scheduled phase if an opportunity is missed but only for a fragment of it, implying several cases with reasonably low albeit not deterministic delay. Consequently, the best choice depends on the type of BE traffic in the system. If it is generated by all nodes or only from some nodes which are known in advance, Proposal 1 is best. However, if the origin of BE traffic is random, unknown or varying, Proposal 2 is best if a bounded delay is needed and Proposal 3 is best if a low average delay is preferred.

Reliability. Proposal 1 does not include contention. When using contention, there is a non-zero probability to have a collision, reducing reliability. Proposal 1 is therefore the best in terms of inherent reliability, Proposal 3 second best in case of not fragmented contention phase, and Proposal 2 is the worst, as the contention is synchronized at the beginning of each slot, aligning the moment in time when the devices start to listen and therefore increasing the probability of a collision.

In addition, mechanisms for increased reliability based on replication can be introduced. Two options are considered: TT and RC traffic classes need more redundancy through the scheduling of more slots. In this case the proposal which already has a lower utilization from the beginning suffers. Alternatively, these retransmissions are placed as if they were BE traffic. In this latter case, the proposal starving BE traffic will also starve retransmissions.

Hardware requirements. The only difference between the proposals is that the first one does not have any contention-based phase and therefore CSMA is not required, while for the other two it is. This latter case can be problematic, as CSMA is not needed during the scheduled phases, but is used in the contention ones. Turning on and off CSMA will most likely involve some overhead that will be inefficient. However, all solutions can be implemented on top of standardized chipset without major changes to the hardware.

Integration with wired networks. There are differences in the way BE traffic is treated in the three proposals. However, given that in e.g., TTEthernet there is no guarantee for BE traffic to be delivered at all, it does not affect integration.

Interferences from other CSMA devices. All proposals except Proposal 1 suffer from the same problem when trying to start a contention-based phase in an environment with co-located CSMA-based devices. If there are any other ongoing transmissions by these devices, the transmissions in the contention-based phase may be postponed. However, in Proposal 1, the CSMA functionality can be turned off, and thus no AIFS needs to be included and the scheduled nodes will get access first. In Proposals 2 and 3, nodes in contention-based phases can have prioritized access if using a shorter waiting time than interfering devices. For that, we should assure that the AIFS plus the randomized time is still lower than the legacy AIFS waiting time of these other external devices.

VII. CONCLUSION

The support for different types of traffic with different reliability requirements is already working successfully in TTEthernet, but only under wired network settings. The possibility of using heterogeneous traffic with the same characteristics but in a wireless setting is desired. In this paper we have proposed three different MAC protocols suitable for time-critical wireless communications, with support for three different traffic classes. A comparative study of them has been performed considering different aspects such as delay, reliability and efficiency. Proposal 1 is best if high predictability is needed and BE traffic is evenly distributed among all nodes. Proposal 3 has the highest flexibility, and works best when BE traffic is uneven or changes often.

ACKNOWLEDGEMENTS

The research leading to these results has received funding from the People Programme (Marie Curie Actions) of the European Union's Seventh Framework Programme FP7/2007-2013/ under REA grant agreement 607727.

REFERENCES

[1] M. Felser, "Real-time Ethernet–industry prospective," *Proc. IEEE*, vol. 93, no. 6, pp. 1118–1129, 2005.

[2] A. Willig, K. Matheus, and A. Wolisz, "Wireless technology in industrial networks," *Proc. IEEE*, vol. 93, no. 6, pp. 1130–1151, 2005.

[3] W. Steiner, G. Bauer, B. Hall, and M. Paulitsch, "TTEthernet: Time-Triggered Ethernet," in *Time-Triggered Communication*, R. Obermaisser, Ed. CRC Press, Aug 2011.

[4] J. Song, S. Han, A. K. Mok, D. Chen, M. Lucas, and M. Nixon, "WirelessHART: Applying wireless technology in real-time industrial process control," in *Proc. IEEE RTAS*, St. Louis, USA, 2008, pp. 377–386.

[5] C. Casetti, C. Chiasserini, M. Fiore, and M. Garetto, "Notes on the Inefficiency of 802.11e HCCA," in *Proc. IEEE VTC*, Dallas, USA, 2005.

[6] H. Trsek and J. Jasperneite, "An isochronous medium access for real-time wireless communications in industrial automation systems-A use case for wireless clock synchronization," in *Proc. IEEE ISPCS*, Munich, Germany, 2011, pp. 81–86.

[7] H. Kopetz and G. Bauer, "The time-triggered architecture," *Proc. IEEE*, vol. 91, no. 1, pp. 112–126, 2003.

[8] H. Kopetz and W. Ochsenreiter, "Clock synchronization in distributed real-time systems," *IEEE Trans. Computers*, vol. 100, no. 8, pp. 933–940, 1987.

[9] W. Steiner, "An evaluation of SMT-based schedule synthesis for time-triggered multi-hop networks," in *Proc. IEEE RTSS*, San Diego, USA, 2010, pp. 375–384.

[10] A. Bohm, M. Jonsson, and E. Uhlemann, "Performance comparison of a platooning application using the IEEE 802.11p MAC on the control channel and a centralized MAC on a service channel," in *Proc. IEEE WiMob*, Lyon, France, 2013, pp. 545–552.

[11] C. L. Liu and J. W. Layland, "Scheduling algorithms for multiprogramming in a hard-real-time environment," *JACM*, vol. 20, no. 1, pp. 46–61, 1973.