

Impact of Network Conditions on Delay-Stable Communications in Closed Industrial Control Networks

David A. Miller and Ahmed E. Kamal

Dept. of Electrical and Computer Eng., Iowa State University, Ames, IA 50011, U.S.A.

E-mail: {davidam, kamal}@iastate.edu

Abstract—To properly operate closed industrial control networks, it is required that communication with nearly constant delay bounds be supported. In [9] the authors introduced FlexTDMA in order to provide this support, and with minimal delay-jitter in an asynchronous network, under unicast communication. In this paper we consider how periodic on-off traffic and frame loss are managed in the presence of network component clock drifts and bandwidth loads, and introduce the FlexTDMA+ protocol. This protocol includes three improvements over FlexTDMA: baseline preemption, partial baselining and baseline deadline density control. In this paper we consider the relative value of each of these improvements, individually and in combination, has in handling these network conditions.

Index Terms—Industrial control networks; constant delay; delay-jitter; asynchronous

I. INTRODUCTION

Closed industrial control networks require that data be delivered from the source, which is typically a central controller, to target nodes with nearly constant delay bounds, and with minimal delay-jitter. The problem with such systems is that networked components are not synchronized, they may be distributed in a wide area, they do not use the same clock, and their clocks may exhibit drifts, in different amounts, and with different polarities. The literature contains a number of solutions to achieve bounded delay periodic traffic, and constant delayed periodic traffic. Bounded delay periodic traffic is supported in [10] using a probabilistic model, and in [4] by exchanging messages for synchronous operation. References [3] [6] [7] [8] support nearly constant delayed traffic in a synchronous network, and require message exchange to maintain a synchronous state. These solutions require a synchronous state, have probabilistic delay bounds, require message exchanges for synchronous operation, or are not suitable for sub millisecond message exchanges.

Asynchronous Ethernet networks can offer bounded delays when each flow is rate constrained and a packet scheduling policy is used in the switch emission process. Stable delay bound performance can only be insured when each switch is able to discern the age of received data. Each switch of an asynchronous Ethernet network has no knowledge of the age of received Ethernet frames, and therefore will not insure stable delay bounds.

In [9] we introduced FlexTDMA to provide minimal delay-jitter with nearly maximal delays in an asynchronous network. Here asynchronous refers to the lack of clock coordination between network components. FlexTDMA works by periodically transmitting a maximally delayed frame on each flow allowing downstream switches to establish a maximal eligibility time (ET) basis, where ET is the time at which an arriving frame is in conformance with the original traffic envelope of the transmitting node. Each FlexTDMA switch traffic shapes arriving frames using this ET basis. The FlexTDMA protocol shares maximal delay bound transmission opportunities in a process called baselining using a dedicated flow called $flow_{01}$.

Here we expand the consideration of FlexTDMA to include end node behavior (periodic on-off traffic), network conditions and improvements to FlexTDMA. Periodic on-off node transmissions occur when end nodes discontinue the flow of periodic messaging traffic due to reset, maintenance, or entering a different mode. Three network conditions are considered: clock drift, frame loss due to bit errors, and bandwidth load. Clock drift of 10 to 100 ppm [5] is common. When the components experience relative drift the accumulated clock error degrades the ability to deliver data at the maximal delay bound.

The improvements included in FlexTDMA+ are baseline preemption, partial baselining and baseline deadline density control. These improvements are motivated by poor performances, in FlexTDMA, resulting from insufficient coordination of baselining opportunities. Each improvement makes FlexTDMA+ more tolerant of concurrent flow baselines. Baseline preemption allows a flow more urgently requiring baselining to preempt a flow scheduled for baselining. Partial baselining, which is used to manage concurrent flow baselines, baselines at a time prior to the deadline of the frame. Baseline deadline density control utilizes unused baseline opportunities to prematurely baseline flows which have a baseline deadline in close proximity to other flows. These improvements included in FlexTDMA+ enhance the delay-jitter and data delivery at maximal delay performance in the presence of frame loss in the network, periodic on-off traffic, and clock-drift.

The remainder of this paper is organized as follows. In Section II, we review FlexTDMA. In Section III, we discuss

the improvements offered in FlexTDMA+. In Section IV, we introduce the operational details of FlexTDMA+. Section V includes a performance study of FlexTDMA+, which is followed by a summary of findings and conclusions in Section VI.

II. BACKGROUND

Algorithm 1: FlexTDMA

```

switch switch frame event on flow k do
  case Frame Arrival
    frame  $\leftarrow$  retrieve frame from input port
    AT[k]  $\leftarrow$  now
    ET[k]  $\leftarrow$  max(AT[k], ET[k] +
       $X_{\min} \cdot (1 - clockDrift)$ )
    frame deadline  $\leftarrow$  ET[k] + flow(k) delay bound
    enqueue(Eligibility Queue, frame)
  case Frame Eligibility
    frame  $\leftarrow$  dequeue(Eligibility Queue)
    store frame in FIFO or Flow01
  case Frame Transmission Completion
    if flow01 queue head scheduled time > now then
      frame  $\leftarrow$  dequeue(Flow01 Queue)
      Baselined[flow k]  $\leftarrow$  true
      transmit frame
    else if not empty(FIFO Queue) then
      frame  $\leftarrow$  dequeue(FIFO Queue)
      transmit frame

```

In this section we describe the key properties of FlexTDMA. The FlexTDMA frame processing stages, shown in Algorithm 1, are arrival, eligibility, and transmission. When a frame arrives the flow ET is determined. The frame deadline is determined from ET and the flow delay bound computed at the output port. A frame, held until eligibility, is scheduled for transmission on flow₀₁ or in a FIFO queue. In [9] we defined a *Baselined flow* as: *A flow on which a frame has been recently transmitted at the delay bound, and each frame has been received before its eligibility time.* A newly baselined frame causes a series of eligibility times in subsequent switches. Once a maximally delayed frame k is received from switch j-1, the eligibility time ET_j^k will be computed as

$$\max(AT_j^k, ET_j^{k-1} + X_{\min} \cdot (1 - clockDrift_{ppm})) \quad (1)$$

where X_{\min} is the minimum frame inter-arrival time on the flow and AT_j^k is the arrival time of frame k at switch j. This generates a series of eligibility times separated by the frame period reduced by the clock drift so that the eligibility times will not out pace arriving frames.

A virtual flow called flow₀₁ provides high priority transmission opportunities for baselining [9]. Flows are scheduled for baselining by reserving a transmission opportunity in flow₀₁ constrained by the allocated transmission period. Flow₀₁ scheduled transmissions are queued in ascending order.

III. FLEXTDMA+

In this section we discuss FlexTDMA+ and show how it improves over FlexTDMA.

A. Baseline Deadline Density Control

The FlexTDMA+ protocol maintains a schedule of baseline deadlines for each active flow. Each active flow must be re-baselined by the established baseline deadline in order to bound the effect of clock drift to planned limits. When a flow is re-baselined, the baseline deadline is updated in the baseline deadline schedule by the minimum baseline interval for the flow. Baseline deadline density control attempts to maintain a uniform distribution for the scheduled baseline deadlines of active flows. The result is that flows tend to reach their baseline deadlines at a steady rate rather than bursting collections of flows requiring baselining. Each frame arriving to a baselined flow is considered for baseline deadline density control. When the flow is not yet baselined, density control is not an issue, instead the goal is to reach a baselined state as soon as possible. The baseline density of the flow is computed relative to all flows with a baseline deadline (all active flows which have reached a baselined state). When the density of the flow is greater than the average, the flow is eligible for baselining at the flow deadline. When flow₀₁ has a baselining opportunity at the flow deadline the flow is scheduled in flow₀₁ queue at the deadline, otherwise it is not scheduled for baselining.




Fig. 1. Baseline deadline density flow i computation

Figure 1 shows the FlexTDMA+ switch baseline deadline schedule for all active flows. The baseline deadline for flow i is shown. The distance between the baseline deadline preceding and following flow i is a, and the distance from the first baseline deadline to the last is b. Therefore the density assigned to flow i is $(a/2)/(b/(n-1))$ when there are n flows.

B. Partial Baselining

Partial baselining of a flow occurs when a baseline is scheduled for a flow prior to the deadline time of the flow. This occurs when several flows are contending for baselining having nearly the same deadline and results in approximated baselining for those flows not baselined at their flow deadline. This allows flows that are not yet baselined to deliver data with a delay that approximates the maximal delay bound, improving the overall performance of the flow set supported by the protocol. Each arriving frame on a flow is considered for partial baselining when the flow₀₁ queue has no baselining opportunity. This will establish an eligibility time basis based on this partial baseline in the downstream switch, and will make the end-to-end delay bound approximate the maximal delay bound reduced by the error in partial baselining. When a frame is received on a flow, and flow₀₁ queue has no baselining opportunity the flow may be partially baselined

at the next earlier transmission opportunity with one important constraint: the partial baseline must extend the baseline deadline - that is, cause the baseline deadline of the flow to increase. When the error (the amount of time the partial baseline precedes the flow deadline) is large enough the new baseline deadline resulting from the partial baseline will be earlier than the current baseline deadline (or current time when the flow is non-baselined). The purpose of baselining a baselined flow is to update the frame eligibility time in the downstream switch given the relative drift that may occur between baselines. The new baseline deadline will be $now + baseline\ interval - error / maximum\ drift\ rate$, where $maximum\ drift\ rate$ is the maximal allowed clock drift for each node and switch. When $baseline\ interval$ is less than $error / maximum\ drift\ rate$, the new baseline deadline would precede the current. In this case the partial baseline will not be performed.

A disadvantage of partial baselining is that it requires another baselining in downstream switches once the baselining event occurs in the current switch since the baselining frame will be transmitted with a larger eligibility time basis than the partial baselining frame. This requires all downstream switches to baseline again.

C. Baselining Preemption

When a frame is received and $flow_{01}$ has no baselining opportunity at the arriving flow deadline, that flow may preempt the current frame scheduled in the $flow_{01}$ queue when three conditions are met. First, the current scheduled frame in $flow_{01}$ must be from a baselined flow. No preemption will be allowed when the current scheduled frame is non-baselined as it is important to allow that flow to attain baselined status. Second, either the preempting flow is not baselined or the baseline deadline of the preempting flow is less than the baseline deadline of the preemption candidate. Thus a non-baselined flow is allowed to preempt a baselined flow since it is urgent to baseline the non-baselined flow, and a baselined flow may preempt another baselined flow when the baseline deadline is less, and so on. Third, as shown in Figure 2, in order for a frame scheduled in $flow_{01}$ to be a candidate for preemption it must also be the case that the preempted frame would meet its deadline when placed in the FIFO queue. Each frame under FlexTDMA+ must be transmitted at or before its deadline. To determine this, the laxity of the frame to its deadline is compared with the existing workload in the FIFO queue. Figure 2 shows a frame on flow 1 preempting a frame from flow 2 scheduled for baselining using $flow_{01}$. The deadline time of flow 1 corresponds the deadline time of flow 2. When flow 2 is preempted from $flow_{01}$ to the FIFO queue, the transmission opportunity created in $flow_{01}$ will be used by flow 1 to schedule a baseline. The workload of the FIFO queue is less than the scheduled transmission time of flow 2 in the $flow_{01}$ queue. Therefore, flow 2 can be preempted from $flow_{01}$ and meet its transmission deadline when placed at the tail of the FIFO queue.

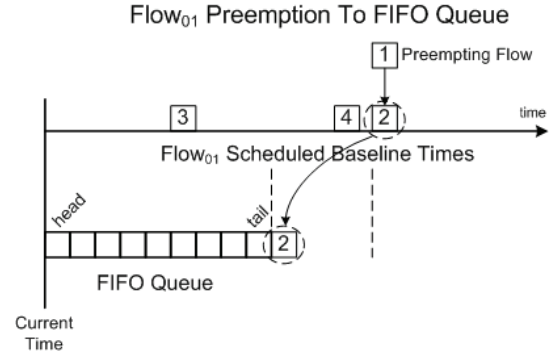


Fig. 2. Scheduling Baselines Preempted From $Flow_{01}$ to FIFO Queue.

IV. FLEXTDMA+ SWITCH OPERATION

In this section we review the operational details of the FlexTDMA+ protocol. Figure 3 shows the stages of frame

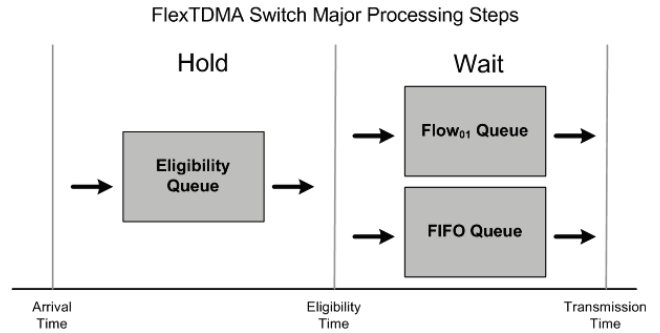


Fig. 3. Major processing phases and events of the FlexTDMA switch.

processing in a FlexTDMA+ switch: 1) arrival, 2) eligibility, and 3) selection for transmission. Arriving frames are held a *hold* period until the ET of the frame [1] [2]. Once eligible, the frame is scheduled for transmission at the output port in either $flow_{01}$ or FIFO queue [9]. The period of time the frame exists in a queue pending transmission is called the *wait* time [1] [2].

Stage 1) Frame Arrival At frame arrival the frame is stored pending eligibility.

Stage 2) Frame Eligibility Once a frame reaches its eligibility time it is processed for each forwarding output port. There are three steps to frame processing at eligibility time.

Eligibility Step 1 - Determine Frame Deadline: The frame deadline is set to ET plus the port delay bound.

Eligibility Step 2 - Determine $Flow_{01}$ Availability and Preemption Potential (if needed): The next step is to determine the $flow_{01}$ availability and preemption potential when another flow is scheduled at the deadline time of the current flow. Figure 4 shows the minimal interval period p between scheduled baseline times in the $flow_{01}$ queue. This limits the bandwidth utilization to $(1\ frame) / (period\ p)$. Baseline times are located in $flow_{01}$ in accordance with this constraint, but as close to the deadline as possible, ideally at the deadline, which exactly establishes the ET basis of the flow at the maximal delay bound of the flow. $Flow_{01}$ queue will be traversed starting




Fig. 4. Flow₀₁ queue is serviced using a minimal frame period to limit bandwidth utilization.

at the deadline of the flow and continuing until the duration limit. The duration limit is set to $maximum\ drift\ rate * ((now + baseline\ interval) - current\ baseline\ deadline)$ insuring the extent of early baselining is limited to the maximal drift that may occur until the baseline deadline. When a transmission opportunity is found in flow₀₁ the flow is marked as a candidate for baseline scheduling.

Eligibility Step 3 - FlexTDMA+ Queuing Decision The actions taken by the FlexTDMA+ protocol at frame eligibility are characterized in Tables I and II. These tables are used to determine which queue an eligible frame will be stored pending transmission (either the flow₀₁ or FIFO queue). Table I shows the queuing decision logic for the ‘Non-Baselined’ state, and Table II shows the queuing decision logic for the ‘Baselined’ state of the FlexTDMA+ Protocol. When baseline deadline is exceeded *density above average* is assumed, but no preemption is allowed.

TABLE I
FLEXTDMA+ QUEUING DECISION LOGIC FOR NON-BASELINED STATE

Baseline Deadline Density Above Average	Flow ₀₁ Availability at Deadline	Flow ₀₁ Availability Within Duration Limit	Preemption Eligible	Queue
NA	Yes	NA	NA	Flow ₀₁ at deadline
NA	No	No	No	FIFO
NA	No	NA	Yes	Preempt Flow ₀₁ at deadline
NA	No	Yes	No	Flow ₀₁ at latest available location

Stage 3) Frame Transmission A ready frame is selected from flow₀₁ if any, otherwise from FIFO queue.

V. FLEXTDMA+ EVALUATION

In this section we describe the evaluation of FlexTDMA+. There are two FlexTDMA+ evaluation phases: Periodic On-Off traffic and Frame Loss. There are three protocol parameters of testing: Baseline Density Control, Partial Baselining and Preemption. For each test run these are either enabled or disabled. The FlexTDMA+ protocol was evaluated by configuring the system and collecting data for each key performance criterion. Table III shows the parameters configured for each

TABLE II
FLEXTDMA+ QUEUING DECISION LOGIC FOR BASELINED STATE

Baseline Deadline Density Above Average	Flow ₀₁ Availability at Deadline	Flow ₀₁ Availability Within Duration Limit	Preemption Eligible	Queue
No	NA	NA	NA	FIFO
Yes	Yes	NA	NA	Flow ₀₁ at deadline
Yes	No	No	No	FIFO
Yes	No	No	Yes	Preempt Flow ₀₁ at deadline
Yes	No	Yes	No	Flow ₀₁ at latest available location
Yes	No	Yes	Yes	Flow ₀₁ at latest available location

TABLE III
FLEXTDMA+ PARAMETERS TESTED

FlexTDMA+ Phase	FlexTDMA+ Phase: Periodic On-Off	FlexTDMA+ Phase: Frame Loss
Periodic On-Off	4 Probabilities	-
Frame Loss	-	4 Probabilities
Preemption	2 (On/Off)	2 (On/Off)
Partial Baselining	2 (On/Off)	2 (On/Off)
Baseline Density Control	2 (On/Off)	2 (On/Off)
Clock Drift Modes	4 (none, increasing, decreasing, mixed)	4 (none, increasing, decreasing, mixed)
Bandwidth Load	3 (20, 50, 90%)	3 (20, 50, 90%)
Total combinations	384	384

test run. For each test run the event occurrence probability is configured. For phase *Periodic On-Off* this is the probability, at each frame transmission, the node will pause periodic transmissions. For phase *Frame Loss* this is the probability, at each frame transmission from the node or switch, the frame is lost to simulate bit errors. Four probability values were chosen for each phase so that, given the expected inter-frame arrival of events and resulting switch egress port baseline demands, the baselining capacity of flow₀₁ was 20%, 50%, 75% and 90% loaded. Four clock drift modes were used in the simulation evaluation: none, increasing, decreasing and mixed. Under clock drift mode *none* no clock drift was applied. Under clock drift mode *increasing* (*decreasing*) nodes 0 to 9 are configured with increasing (decreasing) drift values from 0 to +50 ppm (0 to -50 ppm), and switches 0 to 5 are configured with increasing (decreasing) drift values from 0 to +50 ppm (0 to -50 ppm). Under clock drift mode *mixed* nodes 0 to 9 are configured with alternating drifts of +50 ppm and -50 ppm, and switches 0 to 5 are configured with alternating drifts of +50 ppm and -50 ppm. The key performance criteria of the FlexTDMA+ protocol are: time-to-baseline, delay-jitter and laxity. The time-to-baseline criterion is the time needed for an active flow to achieve a baselined state. The delay-jitter criterion is the

extent of compression between any two arriving frames [1]. When a flow is baselined the FlexTDMA+ switch is able to nearly reconstruct the traffic envelope of the arriving flow, which minimizes the delay-jitter values achieved for frames delivered to destination nodes. The laxity criterion is the extent to which data is delivered to a destination node prior to the delay bound of the flow. The topology in Figure 5 shows the




Fig. 5. FlexTDMA+ Testing Topology

simulated physical topology used to demonstrate the use of FlexTDMA+. The topology is configured so that nodes 0 to 9 transmit to node 10. In this way the loading on the switches increase from switch 0 to 4. Switch 0 carries traffic from nodes 0 and 1, while switch 4 carries traffic from nodes 0 to 9. The total bandwidth is divided among the flows allocated to each node so that the bandwidth allocation varies, causing the periodic transmissions to not maintain a fixed relative offset, but instead to phase so that collisions occur. The delay bounds for the flows of this testing topology range from 227 μs to 1133 μs . The magnitude of these delay bounds is important to keep in mind relative to the performance criteria evaluation.

A. Results

TABLE IV
CLOCK DRIFT EFFECT ON FLEXTDMA+

FlexTDMA+ Phase	FlexTDMA+ Phase: Periodic On-Off	FlexTDMA+ Phase: Frame Loss
Time To Baseline (range of observed values)	0.977 to 7.05 ms	0.935 to 7.12 ms
Time To Baseline (Variation between drift types)	1 to 11%	3 to 11%
Time To Baseline (Average per bandwidth load)	20% - 1 ms 50% - 1.5 ms 90% - 6 ms	20% - 1 ms 50% - 2.2 ms 90% - 6 ms
Delay Jitter	0.035 to 0.05 μs	0 to 0.371 μs
Laxity	0.169 to 7.21 μs	2.7 to 133 μs

1) *Clock Drift Effect on FlexTDMA+*: Table IV shows a summary of the effect of clock drift on FlexTDMA+ for each key performance criterion and for each phase of evaluation. Time-to-baseline had the highest impact from increasing drift. This follows as increasing drift accelerates the number of baselines needed. All time-to-baseline differences resulting from different drift modes are small compared to differences generated from other run parameters. The delay-jitter performance was sub μs . Frame delay bound laxity ranged from 3% to 59% of minimal flow delay bounds, when considering all potential improvements. FlexTDMA+ managed clock drift

efficiently for all modes of operation for each key performance criteria, and clock-drift had a minimal but consistent effect on performance.

2) *Bandwidth Load on FlexTDMA+*: The total bandwidth loading of the network is set for each test. Table V shows the average performance relating to the three bandwidth loads of 20%, 50%, and 90%. The delay-jitter performance was sub μs . The bandwidth load has a consistent effect on the resulting key performance criteria, as time-to-baseline and laxity were reduced as bandwidth loads were reduced. This is the result of reduced demand for baselining opportunities in flow₀₁. As demands lower, the protocol is more responsive to re-baselining needs. Further evaluation will be limited to a single bandwidth load of 90%.

TABLE V
IMPACT OF BANDWIDTH LOAD ON FLEXTDMA+

FlexTDMA+ Phase	FlexTDMA+ Phase: Periodic On-Off	FlexTDMA+ Phase: Frame Loss
Time To Baseline (average)	1.06, 1.81 and 5.70 ms	1.05, 2.05 and 6.02 ms
Delay Jitter	0 to 0.03 μs	0 to 0.37 μs
Laxity (average)	1.6, 2.0 and 7.2 μs	12.6, 22.1, and 28.6 μs

3) *Effect of Probability of Periodic On-Off and Frame Loss on FlexTDMA+*: Table VI shows the increase achieved by modifying the probability of periodic on-off and frame loss. The delay-jitter performance was sub μs . The periodic on-off probability and frame loss probability have a consistent effect of increasing both time-to-baseline and frame delay bound laxity. This follows as a higher periodic on-off probability and frame loss probability means that the density of demanded baselines is increased. Further evaluation uses maximum probability values.

TABLE VI
IMPACT OF THE PROBABILITY OF FLOW INTERRUPTION ON FLEXTDMA+

FlexTDMA+ Phase	FlexTDMA+ Phase: Periodic On-Off	FlexTDMA+ Phase: Frame Loss
Time To Baseline	9.6%	7.6%
Delay Jitter	0 to 0.244 μs	0 to 0.371 μs
Laxity	24.1%	28.8%

4) *Improvements to FlexTDMA+*: There were three improvements to FlexTDMA: 1) baseline deadline density control, 2) partial baselining and 3) baseline preemption. We consider the effect these improvements, used individually and in combination, have on the key performance criteria. Table VII shows the results of testing all combinations of the three protocol improvements to FlexTDMA. The table includes those combinations of the three protocol improvements that resulted in improved performance compared to no protocol modification, and, for comparison to FlexTDMA, includes the results when no improvements were made. The time-to-baseline performance criterion had no clear performance improvement for any of the protocol improvement

TABLE VII
PERFORMANCE OF FLEXTDMA IMPROVEMENTS

Performance Criteria	Improvements	Performance
Time-to-Baseline	All Combinations	5.6 - 7.1 ms
	none	5.8 - 6.6 ms
Laxity	Partial Baselining and Baseline Preemption	2 - 10 μs
	Partial Baselining	5 - 10 μs
	none	10 - 14 μs
Delay Jitter	Baseline Density	0.06 - 0.23 μs
	Baseline Density, Partial Baselining and Baseline Preemption	0.08 - 0.23 μs
	Partial Baselining and Baseline Preemption	0.1 - 0.24 μs
	Baseline Density and Partial Baselining	0.12 - 0.22 μs
	none	0.17 - 0.27 μs

combinations. All combinations offered performances in the same range. The time-to-baseline performance was typically 5 to 26 times the flow delay bound depending on the delay bound of the flow. Partial baselining only approximates a baselining frame transmission and requires additional utilization of baselining transmission opportunities once a full baseline state is achieved. A performance improvement from baseline preemption was expected, but the improvement was mitigated by low concurrent baselining demand as most flows do not usually require re-baselining at the same time. Baseline deadline density control had little effect on time-to-baseline as concurrent re-baselining is low. The laxity performance criterion was improved in two combinations. The first combination was partial baselining and baseline preemption and the second was partial baselining alone. The laxity performance was typically 0.04% to 2% of the flow delay bound, when using optimal improvements. Partial baselining allows each flow to approximate the baselining timing prior to being baselined allowing more frames to be delivered nearly at the deadline. Baseline preemption allows flows requiring a baseline status or flows having experienced more clock drift to be baselined first keeping the computed maximal delay bounds more accurate. The delay-jitter performance criterion was improved in four combinations. The first combination was baseline deadline density control, the second was baseline deadline density control, partial baselining and baseline preemption, the third was partial baselining and baseline preemption, and the fourth was baseline deadline density control and partial baselining.

We conclude that the two improvements, partial baselining and baseline preemption, should be defined in FlexTDMA+. This improvement combination showed the same performance for time-to-baseline, the best performance for laxity and sub μs delay-jitter performance.

VI. SUMMARY

In this paper we proposed and evaluated 3 improvements to the FlexTDMA protocol: partial baselining, baseline preemption and baseline deadline density control. In our evaluation we first characterized the ability of FlexTDMA+ to manage clock

drift. We determined the effect flow transmission interruption has on the FlexTDMA+ performance. We performed full evaluation of the proposed improvements to FlexTDMA using 90% bandwidth loading. We demonstrated that two improvements, partial baselining and baseline preemption, together offered the most improvement in performance compared to FlexTDMA. Partial baselining improves performances by approximating a baselined state until the flow can actually be baselined. Baseline preemption insures baseline opportunities are granted to those flows in most need.

REFERENCES

- [1] H. Zhang and D. Ferrari, "Rate-Controlled Static-Priority Queueing," *In Proc. IEEE Infocom*, pp. 227-236, 1993.
- [2] H. Zhang and D. Ferrari, "Rate-Controlled Service Disciplines," *Journal of High Speed Networking*, 1994.
- [3] J. Ferreira, P. Pedreiras, L. Almeida and J. A. Fonseca, "The FTT-CAN protocol for flexibility in safety-critical systems," *Micro, IEEE*, vol. 22, no. 4, pp. 46-55, Jul/Aug 2002.
- [4] D. M. Cuong, M. K. Kim and H. C. Lee, "Supporting Hard Real-time Communication of Periodic Messages over Switched Ethernet," *The 1st International Forum on Strategic Technology*, vol., no., pp. 419-422, Oct. 2006.
- [5] H. Zhou, C. Nicholls, T. Kunz and H. Schwartz, "Frequency Accuracy and Stability Dependencies of Crystal Oscillators," *Carleton University, Systems and Computer Engineering, Technical Report SCE-08-12*, November 2008.
- [6] R. Santos, R. Marau, A. Oliveira, P. Pedreiras and L. Almeida, "Designing a costumized Ethernet switch for safe hard real-time communication," *IEEE International Workshop on Factory Communication Systems, WFCS*, vol., no., pp. 169-177, May 2008.
- [7] M. Jakovljevic, "Synchronous/asynchronous Ethernet networking for mixed criticality systems," *Digital Avionics Systems Conference, DASC '09. IEEE/AIAA 28th*, vol., no., pp. 1.E.3-1-1.E.3-10, Oct. 2009.
- [8] R. Santos, R. Marau, A. Vieira, P. Pedreiras, A. Oliveira and L. Almeida, "A synthesizable ethernet switch with enhanced real-time features," *35th Annual Conference of IEEE Industrial Electronics, IECON 09*, vol., no., pp. 2817-2824, Nov. 2009.
- [9] D. A. Miller and A. E. Kamal, "FlexTDMA for Delay-Stable Communications in Asynchronous Industrial Control Networks," *Proceedings of the IEEE Local Computer Networks (LCN) Conference*, 2010.
- [10] G. Y. Keung, B. Li and Q. Zhang, "Message Delivery Capacity in Delay-Constrained Mobile Sensor Networks: Bounds and Realization," *IEEE Transactions on Wireless Communications*, vol. 10, no. 5, pp. 1552-1559, May 2011.