

# 1 Efficient Abstraction of Clock Synchronization at 2 the Operating System Level

3 **Alessandro Sorrentino** ✉

4 DEIB, Politecnico di Milano, Italy

5 **Federico Terraneo** ✉ 

6 DEIB, Politecnico di Milano, Italy

7 **Alberto Leva** ✉ 

8 DEIB, Politecnico di Milano, Italy

## 9 — Abstract —

10 Distributed embedded systems are emerging and gaining importance in various domains, including  
11 industrial control applications where time determinism – hence network clock synchronization – is  
12 fundamental. In modern applications, moreover, this core functionality is required by many different  
13 software components, from OS kernel and radio stack up to applications. An abstraction layer  
14 devoted to handling time needs therefore introducing, and to encapsulate time corrections at the  
15 lowest possible level, the said layer should take the form of a timer device driver offering a *Virtual*  
16 *Clock* to the entire system. In this paper we show that doing so introduces a nonlinearity in the  
17 dynamics of the clock, and we design a controller based on feedback linearization to handle the issue.  
18 To put the idea to work, we extend the Miosix RTOS with a generic interface allowing to implement  
19 virtual clocks, including the newly designed controller that we call FLOPSYNC-3 after its ancestor.  
20 Also, we introduce the resulting virtual clock in the TDMH [20] real-time wireless mesh protocol.

21 **2012 ACM Subject Classification** Computer systems organization → Real-time operating systems;  
22 Networks → Time synchronization protocols; Computer systems organization → Embedded software

23 **Keywords and phrases** Clock synchronization, Real-time operating systems, Embedded software,  
24 Real-time control

25 **Digital Object Identifier** 10.4230/OASICS.NG-RES.2023.

## 26 **I** Introduction

27 The world of embedded systems is evolving from isolated to distributed systems. This move  
28 can be observed in several research and market trends such as the Industrial Internet of  
29 Things (IIoT) [16]. As a result, clock synchronization is becoming a key technology to enable  
30 both real-time industrial applications as well as low energy wireless protocols [23]. At the  
31 application level, synchronization in distributed embedded systems allows the execution of  
32 coordinated tasks among multiple devices [13], allows to perform sensing and reconstruction  
33 of spatially distributed phenomena [10, 2], while the availability of synchronization at the  
34 network level enables the use of TDMA protocols [1, 20, 7], being thus fundamental for  
35 real-time communication among devices. Since in modern embedded operating systems  
36 both applications and OS components – such as the radio stack – can benefit from clock  
37 synchronization, an abstraction layer that handles time correction directly at the OS level is  
38 therefore needed. Moreover, from a software engineering perspective, the presence in the  
39 OS codebase of both corrected and uncorrected time values is a potential source of errors.  
40 Therefore, it becomes desirable to encapsulate time correction at the lowest possible level,  
41 such as the timer device driver.

42 However, using corrected times in the entire OS codebase introduces an issue: the  
43 uncorrected time is usually required by the clock synchronization algorithm itself. Efficient  
44 clock synchronization schemes such as FLOPSYNC-2[22] require uncorrected timestamps



© Alessandro Sorrentino and Federico Terraneo and Alberto Leva;  
licensed under Creative Commons License CC-BY 4.0

4th edition of the Workshop on Next Generation Real-Time Embedded Systems.

Editors: John Q. Open and Joan R. Access; Article No. ; pp. :1–:11

OpenAccess Series in Informatics



OASICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

45 of received synchronization packets, and performing clock synchronization using corrected  
46 timestamps is challenging as it makes the model of the clock synchronization problem  
47 nonlinear.

48 This work introduces a new clock synchronization scheme, FLOPSYNC-3, that is capable  
49 of operating with timestamps corrected by the previous iteration of the algorithm itself. As  
50 a result of this improved capability, the Miosix RTOS was extended with a generic interface  
51 allowing to implement clock correction at the hardware timer level.

52 The FLOPSYNC-3 controller is here tested both in simulation and on a network of nodes  
53 running the Miosix operating system and the TDMH [20] real-time wireless mesh protocol.

54 This paper is organized as follows: Section II presents a brief overview on the state of the  
55 art in clock synchronization for distributed embedded systems. Section III discusses how the  
56 Miosix OS has been extended with a virtual clock abstraction that enables transparent clock  
57 corrections. Section IV briefly mentions the design of the Miosix subsystem for performing  
58 timestamp measurements, a key feature used to precisely timestamp clock synchronization  
59 packets. Section V presents the FLOPSYNC-3 clock synchronization scheme that can perform  
60 clock corrections using the virtual clock as actuator while operating on corrected timestamps  
61 only. Finally, Section VI presents simulation and experimental results, and Section VII  
62 outlines future research directions.

## 63 **II Related Works**

64 Clock synchronization is a classical problem in distributed systems [11, 15], but also one where  
65 research is still ongoing to produce clock synchronization schemes fine-tuned to changing appli-  
66 cation requirements and hardware capabilities. Many works related to clock synchronization  
67 in distributed embedded systems come from the Wireless Sensor Network research community,  
68 focusing on several aspects including low power synchronization [18, 22], propagation delay  
69 compensation [12, 19], efficient synchronization information dissemination [14, 8].

70 When considering accuracy, a major differentiating factor is whether a clock synchroniza-  
71 tion scheme only performs offset corrections or it also performs skew corrections. Simple  
72 schemes such as TPSN [9] and DMTS [17] only correct for offset. When implemented at the  
73 OS level, this correction can be efficiently performed by overwriting the hardware counter  
74 with the required correction at every synchronization [9]. The disadvantage is however that  
75 after each correction the hardware clock keeps counting at the incorrect frequency, and thus  
76 a time error accumulates over the synchronization period, which reaches its maximum value  
77 immediately before the next correction. Another issue is that the value returned by the clock  
78 exhibits a discontinuity at every synchronization [22], a matter that can introduce errors in  
79 interval measurements, especially for short intervals.

80 More advanced clock synchronization schemes perform skew (also known as frequency  
81 or rate) corrections. The synchronization scheme produces both an offset and a frequency  
82 correction value at every synchronization. As altering the frequency of a crystal oscillator  
83 requires additional hardware [5] which is usually unavailable in off-the-shelf boards, the  
84 frequency correction is preferably performed by applying an algorithm every time the OS or  
85 applications request the time. In this paper we refer to such algorithm as a virtual clock. For  
86 a given synchronization period, frequency correction allows for lower synchronization errors  
87 compared to offset correction. Additionally, clock synchronization schemes that perform  
88 frequency correction can make the corrected clock continuous and monotonic [22], thus  
89 avoiding clock jumps.

90 Real-time embedded systems also face increasing power and energy constraints [3],

91 especially if battery operated. Clock synchronization may thus be required to operate also  
 92 when the processor enters a deep sleep state which includes turning off the main oscillator.  
 93 In such cases, time is kept using a low power Real-Time Clock (RTC), and this introduces  
 94 the need to synchronize both the RTC and high-frequency timebase [21], a matter that we  
 95 account for by designing our virtual clock to support multiple corrections.

96 In this paper we address the clock synchronization problem from the perspective of  
 97 implementing it at the real-time OS level. Software engineering considerations suggest us to  
 98 completely encapsulate time correction, and since this makes uncorrected time unavailable to  
 99 the clock synchronization scheme, we design a new scheme that can operate with corrected  
 100 timestamps.

### 101 **III Virtual Clock**

102 A real-time OS typically requires two main time-related primitives: one to get the current  
 103 time, whose use is obvious, and one to set an interrupt in a given future time instant, to be  
 104 used to handle context switches as well as sleeping tasks wakeup. This chapter describes the  
 105 design and implementation of a virtual clock to make these primitives synchronization-aware.

#### 106 **III.1 Design**

107 An uncorrected clock  $t_{nc}$  fed by an oscillator with nominal frequency  $f_0$ , affected by (possibly  
 108 time-varying) frequency error  $\delta_s$  will progressively diverge from an ideal one as

$$109 \quad t_{nc}(t) = \int_0^t \frac{f_s(\tau)}{f_0} d\tau = t + \int_0^t \frac{\delta_s(\tau)}{f_0} d\tau \quad (1)$$

111 where  $f_s$  is the instantaneous oscillator frequency, and the integral accounts for the accu-  
 112 mulated frequency error. Accordingly, the accumulated frequency error  $\Delta(k)$  over one clock  
 113 synchronization period  $k$  of duration  $T$  is

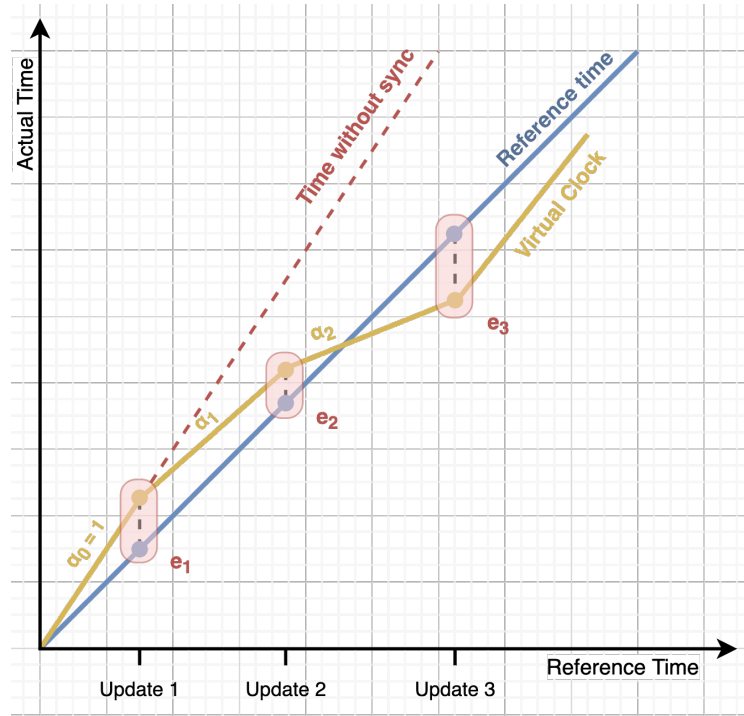
$$114 \quad \Delta(k) = \int_{(k-1)T}^{kT} \frac{\delta_s \tau}{f_0} d\tau \quad (2)$$

115 A virtual clock  $VC$  is a piece wise linear function (Figure 1) that applied to the uncorrected  
 116 clock  $t_{nc}$  produces a corrected one. Virtual clocks allow to perform not only offset corrections,  
 117 but also frequency corrections. Said otherwise, it is possible to control a virtual clock to  
 118 count time faster or slower than the underlying hardware clock to better approximate a  
 119 reference clock. A virtual clock is however just an actuator, it provides the *means* to correct  
 120 a hardware clock, but requires at every synchronization period updated parameters. A clock  
 121 synchronization scheme uses a controller and time information from an external reference to  
 122 adjust the virtual clock rate trying to align it to the reference clock. By defining the virtual  
 123 clock rate separately on each synchronization interval, it can be demonstrated by induction  
 124 that the value of the virtual clock (that is the corrected time  $t_c$ ) on a generic time  $t_{nc}$  inside  
 125 a synchronization interval  $[kT, (k+1)T]$  can be expressed as

$$126 \quad t_c = VC(t_{nc}) = VC(k) + \dot{VC}(k)(t_{nc} - t_{nc}(k)) \quad (3)$$

127 where  $\dot{VC}(k)$  is the rate of the virtual clock. More specifically, if  $t_{nc} = t_{nc}(k+1)$ , its  
 128 definition simplifies as

$$129 \quad VC(k+1) = VC(k) + \dot{VC}(k)(T + \Delta(k)) \quad (4)$$



■ **Figure 1** Virtual clock correcting clock rate to align itself to a reference clock

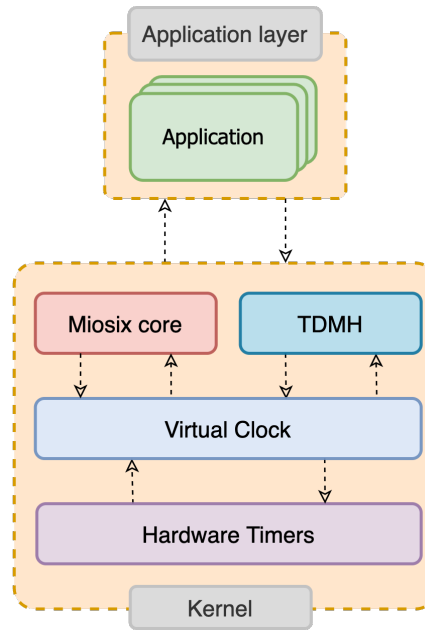
130 We can further generalize the virtual clock expressing (3) as  $f = a_k x + b_k$  by algebraic  
 131 manipulation

$$\begin{cases} a_k &= \dot{VC}(k) \\ b_k &= VC(k) - \dot{VC}(k)t_{nc}(k) \end{cases} \quad (5)$$

134 where  $a_k$  is the rate correction of the clock in the synchronization period  $k$  and  $b_k$  is the  
 135 offset. This rate is adjusted by the controller to align the current clock to the reference, and  
 136 is related to the mean skew over the synchronization period.

### 137 III.2 Implementation

138 The virtual clock was implemented in C++ as part of the Miosix RTOS, as shown in Figure 2.  
 139 To support clock synchronization as well as deep sleep operation which entails transitions  
 140 from a board RTC to an high resolution clock (a technique called VHT [21]), a virtual  
 141 clock may need to perform multiple clock corrections  $f_i$  combined. The software design of  
 142 the virtual clock thus supports multiple corrections as a *Variable Length Correction Stack*  
 143 (VLCS). This design allows for an arbitrary number of *Correction Tiles*, each with their own  
 144 correction parameters  $a_{k,i}$  and  $b_{k,i}$ . For performance reasons, the number of correction tiles  
 145 is configured at compile time as a template parameter. Having  $n$  distinct corrections chained



■ **Figure 2** Virtual clock interface

146 together as  $f_0 \circ \dots \circ f_n$ , the combined correction parameters can be calculated as

$$147 \quad a_{k,vc} = \prod_{i=0}^{n-1} a_{k,i} \quad (6)$$

$$148 \quad b_{k,vc} = \sum_{i=0}^{n-2} \left\{ b_{k,i} \cdot \prod_{j=i+1}^{n-1} a_{k,j} \right\} + b_{k,n-1} \quad (7)$$

150 where index 0 is the correction closer to the hardware timer, and  $n$  the furthest.

151 As reading the current time is a more frequent operation than changing the correction  
 152 coefficients, the combined parameters are precomputed when a new clock correction is  
 153 produced (Figure 3). Conversely, to set a time interrupt the corrected time coming from the  
 154 OS will need to be back-converted as the hardware timer still works using uncorrected time.

$$155 \quad t_c = a_{k,vc} \cdot t_{nc} + b_{k,vc} \quad (8)$$

$$156 \quad t_{nc} = (t_c - b_{k,vc}) / a_{k,vc} \quad (9)$$

158

### 159 III.3 Optimization

160 As the typical skew of quartz clocks is in the order of tens of parts per million (*ppm*), the  
 161  $a_{k,vc}$  coefficient should be very close to 1. Since many microcontrollers lack a Floating Point  
 162 Unit (FPU), we need an efficient way (exploiting the range of  $a_{k,vc}$  as just identified) to  
 163 perform the multiplication  $a_{k,vc} \cdot x$ , as the time retrieval is one of the most critical path  
 164 of the operating system. For this purpose, a template class *Fixed* was designed. This is  
 165 capable of representing a fixed point number with an arbitrary number of bits for the decimal  
 166 part. Given a few compile-time optimized functions able to handle multiplication between a  
 167 64-bit integer and a fixed point 32.32, a specialization of the said class – called *fp32\_32* and

168 representing a fixed point number as a 64-bit integer with 32-bit for both decimal and integer  
 169 part – was used. With `fp32_32`, the multiplication  $a_{k,vc} \cdot x$  can be performed in just 60 clock  
 170 cycles bringing the total time to get the current time to 170 clock cycles, a 37% improvement  
 171 compared to the previous implementation. Regarding the uncorrection, we can note from (9)  
 172 that a division by  $a_{k,vc}$  is needed. There is no nice properties to perform fast division using  
 173 fixed point, so it was implemented as a multiplication for the inverse. The inverse value is  
 174 precomputed using 64-bit floating point numbers at every update of the  $a_{k,vc}$  parameter and  
 175 converted to `fp32_32`. The pre-computation is optimized using a modified version of the  
 176 fast inverse square root algorithm [6], adapted to perform  $1/x$  instead of  $1/\sqrt{x}$  as follows.  
 177 The optimization relies on the fact that an *IEEE754* double precision number is very similar  
 178 to an Logarithmic Number System (LNS) number, as they never differ for more than a  
 179 small factor. An interesting property of LNS numbers is that it makes implementations for  
 180 multiplications and divisions very efficient. In particular, the inverse of an LNS number  $v$   
 181 is  $-v$ . The bit representation of a floating point number  $u$  can approximated as the LNS  
 182 number  $x = 2^{u/2^{52}-1023}$ , and using this representation the inverse  $q$  can be computed as  
 183 follows

$$184 \quad 2^q/2^{52-1023} = 2^{-(u/2^{52}-1023)} \quad (10)$$

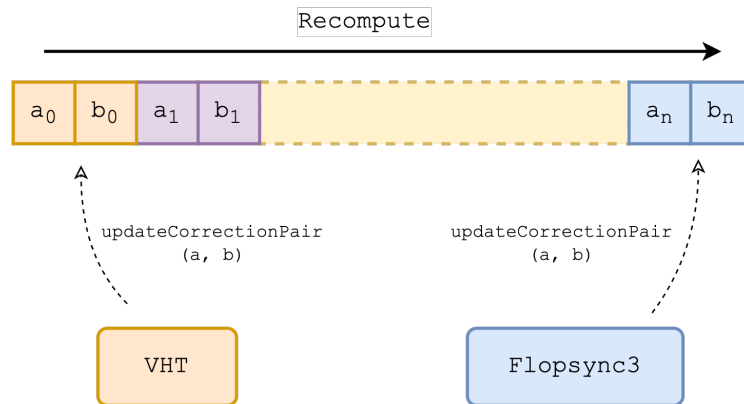
186 which solving the implicit equation results in

$$187 \quad q = 0x7FE0000000000000 - u \quad (11)$$

189 Performing a sweep with sufficient precision, it was possible to elaborate a quadratic regression  
 190 model to approximate faster and with more precision the hexadecimal value. Having a closer  
 191 approximation, less Newton steps are necessary making the inversion faster. This whole  
 192 inversion process is called *optimizedFastInverse*.

193 **IV Hardware Events**

194 Although what presented above is sufficient to support the time-related requirements of  
 195 an OS, performing clock synchronization requires accurate timestamping of received radio  
 196 packets. Moreover, advanced radio transmission techniques such as constructive interference  
 197 require accurate packet transmission times [22]. To abstract hardware-accelerated event  
 198 timestamping and generation, Miosix was extended with an *Eventstamping* interface. This



■ **Figure 3** Virtual clock recomputing aggregated parameters  $a_{vc}$  and  $b_{vc}$

199 abstraction introduces the concept of *event channels* that abstract the event sources or sinks  
 200 of a given platform. Every event channel can be configured as *input*, for external event  
 201 timestamping, or as *output* to trigger events. When configured in input mode, a thread  
 202 can block and wait for an event to happen on the chosen channel, with an optional timeout  
 203 (Figure 4). When configured as output, a thread can generate a hardware event in the future,  
 204 blocking until that time point. This design simplifies the realization of TDMA networking  
 205 protocols. Since events are measured/generated in hardware, the achievable time granularity  
 206 is that of the hardware timer (in our implementation 21ns), and is unaffected by software  
 207 interrupt latencies.

## 208 V FLOPSYNC-3

209 The redesign of the Miosix OS timing subsystem in order to only operate in terms of corrected  
 210 time in the entire OS –except for the timer driver– required the design of a new clock  
 211 synchronization scheme. Previously, the clock synchronization packets were timestamped  
 212 using the uncorrected clock, as this was needed by the FLOPSYNC-2 algorithm [22]. The  
 213 previous approach required to deal with both corrected and uncorrected times and was causing  
 214 code maintainability issues from a software engineering standpoint. However, performing clock  
 215 synchronization using timestamps corrected by the previous round of clock synchronization  
 216 makes the problem nonlinear. The FLOPSYNC-3 controller was designed to address the  
 217 aforementioned nonlinearity, and implemented at the OS level.

### 218 V.1 Design

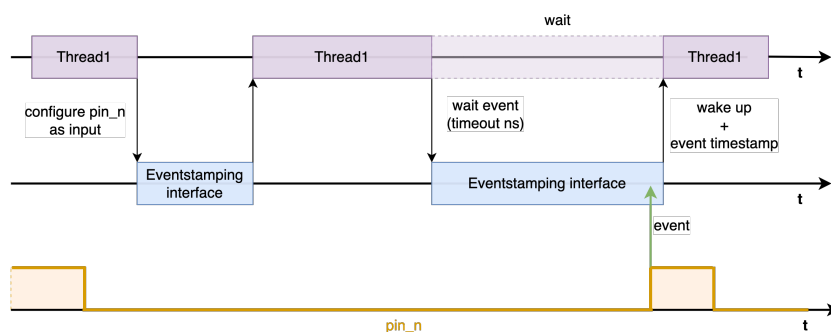
219 Given (4), we can define the clock synchronization error at the end of each synchronization  
 220 period as

$$221 \quad e(k) = VC(k) - kT \quad (12)$$

222 To observe the evolution of the error across synchroniation periods, we can compute the next  
 223 error as a function of the previous, resulting in

$$224 \quad e(k+1) = e(k) + T(1 - \dot{VC}(k)) - \Delta(k)\dot{VC}(k) \quad (13)$$

225 where the  $\Delta(k)\dot{VC}(k)$  term makes the model nonlinear. To perform the control synthesis we  
 226 used *Feedback Linearization* [4] to linearize this process using the output  $u(k)$  of a linear



227 ■ **Figure 4** Eventstamping, wait event

## XX:8 Efficient Abstraction of Clock Synchronization at the Operating System Level

227 controller and express the new error as

$$228 \quad e(k+1) = \beta e(k) + (1 - \beta) u(k) \quad (14)$$

229 and as a consequence have the new output  $u(k)$  of the controller provide  $\dot{VC}(k)$  from

$$230 \quad \dot{VC}(k) = \frac{e(k)(1 - \beta) + u(k)(\beta - 1) + T}{T + \Delta(k)}. \quad (15)$$


231 The mean skew value at the synchronization period  $k$  is of course not available and needs  
232 to be approximated with the previous one ( $k - 1$ ), i.e.,

$$233 \quad \hat{\Delta}(k) = \Delta(k - 1) = \frac{VC(k) - VC(k - 1)}{\alpha(k - 1)} - T \quad (16)$$

234 We can then obtain the transfer function  $\mathcal{H}(z)$  of the imposed dynamic (14) as

$$235 \quad \mathcal{H}(z) = \frac{E(z)}{U(z)} \Rightarrow \frac{1 - \beta}{z - \beta} \quad (17)$$

The controller  $C(z)$  used to work in conjunction with the feedback linearization is a pro-



■ **Figure 5** FLOPSYNC-3 Control scheme

236  
237 portional one having a gain of 0.15. The full FLOPSYNC-3 control scheme is shown in  
238 Figure 5.

## 239 **VI** Simulation and Experimental Results


240 The operation of the FLOPSYNC-3 controller and virtual clock were first assessed through  
241 simulations performed using the *Modelica* language.

242 Figure 6 shows one such simulation, where the clock synchronization period  $T$  was set to  
243 10 seconds and  $\beta$  was chosen to be 0.025. The left part of the figure shows the simulated  
244 clock skew profile, that starts from 10 ppm and increases to 50 ppm from  $T = 150$  seconds,  
245 approximating in the simulation the effect of an ambient temperature change. The right part  
246 of the figure shows the clock synchronization error. The blue line is the instantaneous error  
247 of the virtual clock, thus the time error exposed to the operating system and application. As  
248 a node in the network can measure its error only at discrete intervals, corresponding to when  
249 synchronization packets are received, the red line shows the measured error that feeds the  
250 FLOPSYNC-3 controller.

251 As can be seen, the initial 10 ppm skew causes a  $100 \mu\text{s}$  error that is quickly corrected by  
252 the FLOPSYNC-3 controller. The frequency change caused by the simulated temperature  
253 change, although higher in amplitude than the initial skew causes a lower peak error, less  
254 than  $75 \mu\text{s}$ , due to its slower nature.

255 The FLOPSYNC-3 controller as specified by equation (15) and (16) has been implemented  
256 in Miosix acting on the the variable length correction stack of the virtual clock. Since deep






■ **Figure 6** Simulated clock skew profile (left) and clock synchronization error (right). The blue line shows the instantaneous synchronization error, while the red line shows the measured error.

257 sleep support was not implemented, the correction stack was configured to perform the  
 258 FLOPSYNC-3 correction only. Synchronization parameters  $T$  and  $\beta$  were configured as in  
 259 the simulations. The clock synchronization error measurement was taken from the TDMH  
 260 networking stack using the eventstamping interface to provide the timestamps of received  
 261 synchronization packets. FLOPSYNC-3 was implemented using the `fp32_32` type to perform  
 262 intermediate calculations efficiently. Because of the limited range of this type, pre-scaling  
 263 was necessary to avoid overflows.

264 Clock synchronization experiments were performed with a network of nodes running the  
 265 TDMH networking stack on top of Miosix. Figure 7 shows the clock synchronization error of  
 266 one such node. The top part of the figure shows the measured clock synchronization error in  
 267 the first three minutes after synchronization. The initial clock synchronization error of  $40 \mu\text{s}$   
 268 occurs when the node is booted and joins the network. This value is the accumulation of the  
 269 oscillator frequency error over the entire first synchronization period, as the FLOPSYNC-3  
 270 algorithm, being a feedback one, requires a first error measure to compute a correction. The  
 271 bottom part of the figure shows the error after the initial synchronization, over a period of  
 272 approximately 24 hours, to better appreciate the error dynamics after the initial skew is  
 273 corrected. The observed stochastic nature of the clock synchronization error, not present  
 274 in the simulations, is caused by the measurement noise of packet timestamps. The error  
 275 standard deviation, excluding the first transient, is 137 ns.

## 276 VII Conclusions

277 This work addressed abstracting clock synchronization at the operating system level. To  
 278 achieve this goal a virtual clock was introduced as an efficient abstraction allowing a hardware  
 279 timer driver to provide a time reference whose rate can be changed compared to the one  
 280 of the underlying oscillator. Support for multiple corrections sources was accounted for,  
 281 allowing the implementation of deep sleep solution such as VHT [21]. Encapsulating time  
 282 correction allows reducing bugs and problems during development since all components are  
 283 just using the same time source (corrected), but makes the uncorrected synchronization  
 284 packet timestamps unavailable to the clock synchronization algorithm. The FLOPSYNC-3  
 285 controller was thus designed specifically to overcome this issue.



■ **Figure 7** Clock synchronization error during experimental evaluation. Top plot includes the initial clock skew compensation, bottom plot shows the synchronization error after the initial transient.

286 The Miosix real-time OS was extended with a flexible, efficient and modular timing  
 287 subsystem based on the virtual clock design, capable of internalizing the clock correction and  
 288 only exposing corrected time to all kernel and application tasks. This new timing subsystem  
 289 was designed from the start to be general allowing to easily port the Miosix to different  
 290 microcontrollers.

291 Future research directions will focus on further improving clock synchronization resilience  
 292 to temperature variations, while future improvements of the Miosix timing subsystem will  
 293 address completing the support for maintaining clock synchronization during deep sleep  
 294 periods using the variable length correction stack.

## 295 — References —

- 296 1 Diogo Almeida, Miguel Gaitán, Pedro d’Orey, Pedro Santos, Luis Ramos Pinto, and Luís  
 297 Almeida. Demonstrating RA-TDMAs+ for robust communication in WiFi mesh networks. In  
 298 *RTSS@work workshop co-located with the 42nd IEEE Real-Time Systems Symposium*, 2021.
- 299 2 Riad Azzam and Nabil Aouf. Visual information to enhance time difference of arrival based  
 300 acoustic localization. In *2014 IEEE International Conference on Aerospace Electronics and*  
 301 *Remote Sensing Technology*, pages 77–82, 2014. doi:10.1109/ICARES.2014.7024400.
- 302 3 Ashikahmed Bhuiyan, Federico Reghenzani, William Fornaciari, and Zhishan Guo. Optimizing  
 303 energy in non-preemptive mixed-criticality scheduling by exploiting probabilistic information.  
 304 *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 39(11):3906–  
 305 3917, 2020. doi:10.1109/TCAD.2020.3012231.
- 306 4 Roger W Brockett. Feedback invariants for nonlinear systems. *IFAC Proceedings Volumes*,  
 307 11(1):1115–1120, 1978.

- 308 5 Maxim Buevich, Niranjini Rajagopal, and Anthony Rowe. Hardware assisted clock synchrono-  
309 nization for real-time sensor networks. In *2013 IEEE 34th Real-Time Systems Symposium*,  
310 pages 268–277, 2013. doi:10.1109/RTSS.2013.34.
- 311 6 John Carmack. Fast inverse square root. URL: [https://blog.timhutt.co.uk/  
312 fast-inverse-square-root/](https://blog.timhutt.co.uk/fast-inverse-square-root/).
- 313 7 Julius Degeysys, Ian Rose, Ankit Patel, and Radhika Nagpal. DESYNC: Self-Organizing  
314 Desynchronization and TDMA on Wireless Sensor Networks. In *Proceedings of the 6th  
315 International Conference on Information Processing in Sensor Networks*, page 11–20, 2007.  
316 doi:10.1145/1236360.1236363.
- 317 8 F. Ferrari, M. Zimmerling, L. Thiele, and O. Saukh. Efficient network flooding and time  
318 synchronization with glossy. In *Information Processing in Sensor Networks (IPSN)*, 2011.
- 319 9 S. Ganeriwal, R. Kumar, and M. Srivastava. Timing-sync protocol for sensor networks. In  
320 *International Conference on Embedded Networked Sensor Systems*, 2003.
- 321 10 Grzegorz Krukar, Marco Wenzel, Piotr Karbownik, Norbert Franke, and Thomas von der Grün.  
322 Proof-of-concept real time localization system based on the UWB and the WSN technologies.  
323 In *2014 International Conference on Indoor Positioning and Indoor Navigation (IPIN)*, pages  
324 756–757, 2014. doi:10.1109/IPIN.2014.7275559.
- 325 11 L. Lamport. Time, clocks, and the ordering of events in a distributed system. *Communications  
326 of the ACM*, 1978.
- 327 12 Roman Lim, Balz Maag, and Lothar Thiele. Time-of-flight aware time synchronization for  
328 wireless embedded systems. In *Proceedings of the 2016 International Conference on Embedded  
329 Wireless Systems and Networks, EWSN '16*, page 149–158, USA, 2016. Junction Publishing.
- 330 13 L. Maillet and C. Fraboul. Scheduling complex real-time tasks in an embedded distributed  
331 system. In *Proceedings Seventh Euromicro Workshop on Real-Time Systems*, pages 62–65,  
332 1995. doi:10.1109/EMWRTS.1995.514293.
- 333 14 M. Maroti, B. Kusy, G. Simon, and A. Ledeczi. The flooding time synchronization protocol.  
334 In *Conference On Embedded Networked Sensor Systems*, 2004.
- 335 15 D.L. Mills. Internet time synchronization: the network time protocol. *IEEE Trans. on  
336 Communications*, 39(10):1482–1493, 1991.
- 337 16 Dinh C. Nguyen, Ming Ding, Pubudu N. Pathirana, Aruna Seneviratne, Jun Li, Dusit Niyato,  
338 Octavia Dobre, and H. Vincent Poor. 6G Internet of Things: A Comprehensive Survey. *IEEE  
339 Internet of Things Journal*, 9(1):359–383, 2022. doi:10.1109/JIOT.2021.3103320.
- 340 17 Su Ping. Delay measurement time synchronization for wireless sensor networks. In *Intel  
341 Research*, 2003.
- 342 18 A. Rowe, V. Gupta, and R. Rajkumar. Low-power clock synchronization using electromagnetic  
343 energy radiating from ac power lines. In *Sensys*, pages 211–224, 2009.
- 344 19 Federico Terraneo, Alberto Leva, Silvano Seva, Martina Maggio, and Alessandro Vittorio  
345 Papadopoulos. Reverse Flooding: Exploiting Radio Interference for Efficient Propagation  
346 Delay Compensation in WSN Clock Synchronization. In *2015 IEEE Real-Time Systems  
347 Symposium*, pages 175–184, 2015. doi:10.1109/RTSS.2015.24.
- 348 20 Federico Terraneo, Paolo Polidori, Alberto Leva, and William Fornaciari. TDMH-MAC:  
349 Real-Time and Multi-hop in the Same Wireless MAC. In *2018 IEEE Real-Time Systems  
350 Symposium (RTSS)*, pages 277–287, 2018. doi:10.1109/RTSS.2018.00044.
- 351 21 Federico Terraneo, Fabiano Riccardi, and Alberto Leva. Jitter-Compensated VHT and Its  
352 Application to WSN Clock Synchronization. In *2017 IEEE Real-Time Systems Symposium  
353 (RTSS)*, pages 277–286, 2017. doi:10.1109/RTSS.2017.00033.
- 354 22 Federico Terraneo, Luigi Rinaldi, Martina Maggio, Alessandro Vittorio Papadopoulos, and  
355 Alberto Leva. FLOPSYNC-2: Efficient Monotonic Clock Synchronisation. In *2014 IEEE  
356 Real-Time Systems Symposium*, pages 11–20, 2014. doi:10.1109/RTSS.2014.14.
- 357 23 Hüseyin Yiğitler, Behnam Badihi, and Riku Jäntti. Overview of time synchronization for iot  
358 deployments: Clock discipline algorithms and protocols. *Sensors (Switzerland)*, 20(20):1 – 58,  
359 2020. doi:10.3390/s20205928.