Synchronized State in Networks of Digital Phase-Locked Loops

J. M. Akré, J. Juillard
SSE, SUPELEC
Gif-sur-Yvette, France
jerome.juillard@supelec.fr
jean-michel.akre@supelec.fr

D. Galayko
LIP6, UPMC
Paris, France
dimitri.galayko@lip6.fr

E. Colinet
CEA-LETI, Minatec
Grenoble, France
eric.colinet@cea.fr

***Abstract*—Clock distribution networks of synchronized oscillators are an alternative approach to classical tree-like clock distribution methods. Each node of the network may consist of a phase-locked loop (PLL) trying to match the phase of its neighbors. Then a network of independent oscillators takes the place of the centralized clock source, providing separate clock signals to the physically distant parts of the system. In the discrete case, the digital filter is necessarily operated asynchronously: each operation is triggered by a rising edge of the locally-generated clock, the frequency and phase of which vary as the whole system tries to synchronize. The locking behavior, the synchronous state and the stability conditions of such a system are analyzed. Similarly, the synchronization of an autonomous network of two self-sampled PLLs is studied. Surprisingly, its analysis is much simpler than that of the single PLL.**

I. INTRODUCTION

Large scale synchronous systems-on-chip (SoCs) require reliable clock distribution systems to guarantee correct temporal order in the information processing. Traditionally, two topologies of clock distribution are used (H-tree or clock grid) to distribute the output of a central reference oscillator [1]. To circumvent the difficulties associated with these conventional techniques (constraints of symmetry and design), an alternative approach consists in dividing the SoC into several isochronous zones, each with its own reference oscillator, and letting each zone share some information with its neighbors so that a form of consensus (synchronization) is reached. This may be achieved by inserting in each zone a phase-locked loop (PLL) in order to cancel out the phase error between the local clock and the neighboring ones. This concept of distributed synchronous clocking was introduced in 1995 by Pratt and Nguyen [1], an implementation of which was proposed by Gutnik and Chandrakasan in 2000 [2]. However, this architecture had little success with designers of digital circuits, probably because it was based on analog techniques. The HODISS project, funded by the ANR ARFU program, aims at pursuing the seminal work of [1] and [2] into the digital domain, in order to benefit from the noise-immunity and the greater flexibility of digital components.

French National Agency of Research (ANR)

In this paper, the behavior of an all digital PLL (ADPLL) [3] which constitutes the basic building block of our PLL network is investigated in terms of convergence to a desired synchronized state. To maximize the autonomy of the system and thus avoid the use of an external clock, the loop filter of the PLL is driven by the rising edges of the output of its own digitally controlled oscillator (DCO). We show in section II that this "self-sampled" PLL can be modeled as a discrete autonomous piecewise-linear system and investigate its synchronized state by analytical and numerical means. In section III, an autonomous (without an absolute reference clock) network of two such PLLs is analyzed. The case when there is an absolute reference can be treated as an extension of the previous case. It is remarkable that, for such a network, the synchronized state and the locking behavior of the system are much simpler to establish than for the single PLL case.

II. ANALYSIS OF A SELF-SAMPLED PLL

A simple self-sampled PLL is represented in Fig. 1. It is composed of a digital phase detector (DPD), a proportional integral (PI) filter and a DCO. The PI filter is driven by the rising edges of the DCO output. The DPD is linear and outputs the (non-dimensionalized) time difference between a rising edge of the local clock and a rising edge of the reference clock, in much the same way as flip-flops comparators in [7]. Supposing that the frequency of the DCO is not far from that of the external clock, the self-sampled PLL can be described by the model of section II-A.

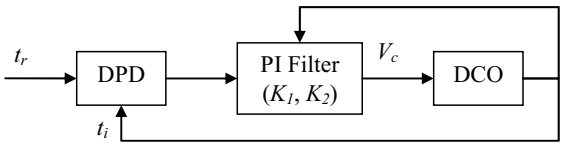

Figure 1. Block-diagram of a self-sampled PLL.

A. Governing equations

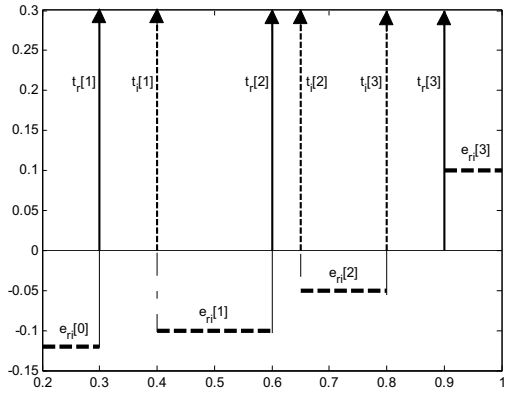

Figure 2. Error signal e_{ri} and rising edges of the reference clock t_r and local clock t_i.

Let $t_i[n]$ (respectively $t_r[n]$) be the n^{th} rising edge of the DCO output (respectively of the reference clock). These quantities are governed by the linearized equations:

$$t_r[n+1] = t_r[n] + T_r\,,\ t_i[n+1] = t_i[n] + T_i + V_c[n] \tag{1}$$

where T_r is the period of the reference clock, V_c the control voltage of the DCO and T_i is the central period of the DCO. Since the PI filter is driven by the rising edges of the local clock, it cannot read $e_{ri}[n] = t_r[n] - t_i[n]$, the n^{th} value of the time difference output by the DPD, if the reference clock lags behind the local clock (because $t_r[n]$ is not known yet), (Fig.2). Instead, another value is read, some obvious choices being $e_{ri}[n-1]$, 0 or a prediction of $e_{ri}[n]$ based on previous observations. All these choices lead to qualitatively similar problems; in the present paper, for the sake of brevity, only the first choice is investigated. Hence V_c is governed by:

$$V_c[n] = V_c[n-1] + K_1 \cdot \varepsilon[n] + K_2 \cdot e_{ri}[n-1], \tag{2}$$

where
$$\varepsilon[n] = \begin{cases} e_{ri}[n] & \text{if } e_{ri}[n] \le 0 \\ e_{ri}[n-1] & \text{otherwise} \end{cases}. \tag{3}$$

Note that the value of $\varepsilon[n]$ could also be propagated in the filter so that $V_c[n]$ would be expressed as a function of $\varepsilon[n]$ and $\varepsilon[n-1]$ rather than $\varepsilon[n]$ and $e_{ri}[n-1]$. Once again, these two choices lead to qualitatively similar problems. From (1-3), the system can be shown to switch between S_{lag} :

$$e_{ri}[n+1] - (2-K_1)e_{ri}[n] + (1+K_2)e_{ri}[n-1] = 0\,, \tag{4}$$

if $e_{ri}[n] \le 0$ and S_{lead} :

$$e_{ri}[n+1] - 2e_{ri}[n] + (1+K_1+K_2)e_{ri}[n-1] = 0\,, \tag{5}$$

if $e_{ri}[n] > 0$. Thus, the self-sampled PLL behaves as an autonomous 2nd-order piecewise-linear switched system [4]. Depending on the values of K_1 and K_2, this system may be globally stable and converge to a synchronized state. This is studied in sub-section B.

B. Convergence and locking behavior analysis

The convergence and the global locking behaviour of the self-sampled PLL depend on the roots of the characteristic polynomials of (4) and (5). First, it should be noticed that S_{lead} is always unstable. If its roots are real, at least one of them is positive and larger than 1 in magnitude. Thus, depending on the initial conditions, the solution of (5) may be monotonically increasing, leading to a globally unstable self-sampled PLL. A necessary condition for global convergence is then:

$$K_1 + K_2 > 0\,. \tag{6}$$

If (6) is satisfied, S_{lead} has two complex conjugate roots. Because of the oscillatory nature of the solution of (5), there always exists n so that $e_{ri}[n] \le 0$ and the self-sampled PLL switches to S_{lag}. Two cases must now be considered. If S_{lag} has real roots, convergence is ensured if they are positive and smaller than 1 (because the solution of (4) then decreases exponentially without ever crossing back to S_{lead}). This translates into the following conditions:

$$\begin{cases} K_1^{\,2} - 4(K_1+K_2) > 0 \\ 1 + K_2 > 0 \\ 0 < K_1 < 2 \end{cases}, \tag{7}$$

The first inequality is the condition that must be satisfied for S_{lag} to have real roots. The condition on the modulus of these roots translates into the other two inequalities. Equations (6) and (7) define the grey striped area in Fig. 3. There is no simple way to determine the convergence of the PLL when the roots of S_{lag} are real but one of them is negative and one should then turn to the techniques described in [4-5] in that case. Now, if S_{lag} has complex conjugate roots, i.e. if

$$K_1^{\,2} - 4(K_1+K_2) < 0 \tag{8}$$

is satisfied, the self-sampled PLL keeps switching from S_{lag} to S_{lead} and back again. The conditions under which such a system is globally stable are notoriously difficult to determine. For a given couple $\{K_1, K_2\}$, one may demonstrate stability by constructing a piecewise-continuous Lyapunov function [4-5]. The problem of synthesis, in which we try to determine all $\{K_1, K_2\}$ couples which ensure convergence and stability, is much more complex. Intuitively, a necessary condition for

global convergence is that the positive damping of S_{lag} should exceed the negative damping of S_{lead}. Two approaches may be used to translate this intuition into mathematical terms.

First of all, one may recast (4) and (5) into one equation:

$$\begin{aligned} & e_{ri}[n+1]-(2-\frac{K_1}{2})e_{ri}[n]+(1+\frac{K_1}{2}+K_2)e_{ri}[n-1] \\ & =\frac{K_1}{2}\text{sign}(e_{ri}[n])(e_{ri}[n]-e_{ri}[n-1]) \end{aligned} \quad (9)$$

and one may consider the self-sampled PLL as a perturbation of a linear averaged system S_{avg}:

$$e_{ri}[n+1]-(2-\frac{K_1}{2})e_{ri}[n]+(1+\frac{K_1}{2}+K_2)e_{ri}[n-1]=0\,. \quad (10)$$

It would then be reasonable to assume that a sufficient condition for the global convergence of the self-sampled PLL is that S_{avg} should be convergent. This results in the following condition on K_1 and K_2:

$$K_1+2K_2<0\,. \quad (11)$$

Another approach that may be used is to assume that K_1 and K_2 are small. The roots of S_{lag} and S_{lead} are then close to the unit circle and it is rational to transform the two systems to continuous time (using impulse invariance, for example). Supposing the initial state of the continuous system is:

$$e(0)=0\,,\ \dot{e}(0)=\dot{E}_0>0\,, \quad (12)$$

it is then very simple to determine the times t_n when the zero-crossings of $e(t)$ occur. The successive values of $\dot{e}(t_{2n})$ are a geometric sequence: a sufficient condition for the convergence of the transformed system is that $\dot{e}(t_2)<\dot{E}_0$. The calculations are straightforward and, in the limit of small K_1 and K_2, the condition for the convergence of the continuous system is found to boil down to (11).

Simulations of the self-sampled PLL are performed (Fig.3). For every couple $\{K_1,K_2\}$, we note if the PLL synchronizes or not. The simulations show that, for small values of K_1 and K_2, (11) is a condition for convergence, as predicted. For moderate values of K_1 and K_2, the frontier defined by (11) is more or less respected. Furthermore, we remark that there are values of K_1 and K_2 for which S_{lag} has at least one negative real pole and for which synchronization occurs. The rigorous determination of the convergence domain of the self-sampled PLL is still an open problem. In the next section, we show that the convergence domain of an autonomous network of two self-sampled PLLs is, much simpler to determine.

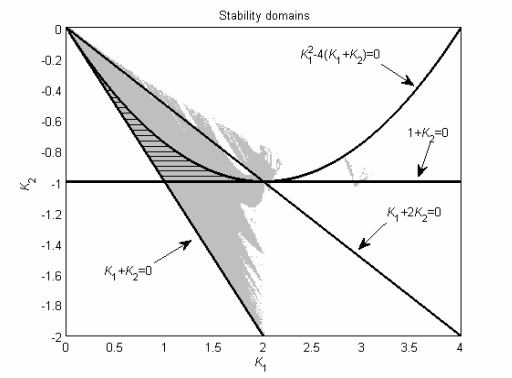

Figure 3. Convergence domain of a self-sampled PLL. The shaded area corresponds to values of filter coefficients for which simulations show that the PLL synchronizes.

III. SYNCHRONISATION OF A NETWORK OF TWO PLLS

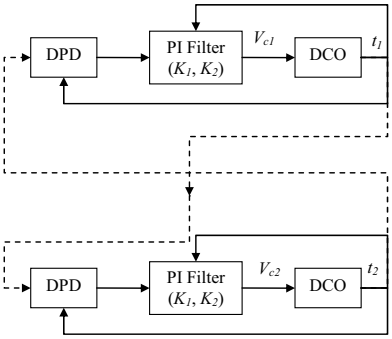

Figure 4. Autonomous network of two self-sampled PLLs

Let us now consider a very simple network of two PLLs, as depicted in Fig. 4. The filters are assumed to be identical, with coefficients K_1 and K_2. The same notations as in the previous section are used: t_1 and t_2 denote the moments when rising edges occur, V_{c1} and V_{c2} are the control voltages of the DCOs and T_1 and T_2 are their central periods:

$$t_k[n+1]=t_k[n]+T_k+V_{ck}[n],\ k\in\{1,2\} \quad (13)$$

It is clear that the phase differences satisfy: $e_{21}[n]=-e_{12}[n]$

From (13) and (2), it is then simple to establish that:

$$e_{12}[n+1]-2e_{12}[n]+e_{12}[n-1]=K_1(\varepsilon_1[n]-\varepsilon_2[n])-2K_2e_{12}[n], \quad (14)$$

where $$\varepsilon_i[n]=\begin{cases} e_{ji}[n] \text{ if } e_{ji}[n]\le 0 \\ e_{ji}[n-1] \text{ otherwise} \end{cases} \quad i,j\in\{1,2\},\ i\ne j\,. \quad (15)$$

From (15), it follows that: $\varepsilon_1[n]-\varepsilon_2[n]=-e_{12}[n]-e_{12}[n-1]$

and (14) reduces to:

$$e_{12}[n+1]-(2-K_1)e_{12}[n]+(1+K_1+2K_2)e_{12}[n-1]=0\,. \quad (16)$$

Thus, the network of 2 PLLs behaves as an autonomous linear system S_{net}. Its convergence is guaranteed provided its roots are inside the unit circle. S_{net} has real roots if

$$K_1^{\,2}-8(K_1+K_2)>0 \quad (17)$$

and they are stable provided:

$$\begin{cases} K_1+K_2>0 \\ 0<K_1<4 \\ -2<K_2 \end{cases}. \quad (18)$$

If the roots are a complex conjugate pair, they are stable if:

$$K_1+2K_2<0\,. \quad (19)$$

The corresponding areas are represented in Fig. 5. Thus, provided K_1 and K_2 are chosen inside the grey area in Fig. 5, the PLLs synchronize fully, in phase and in frequency.

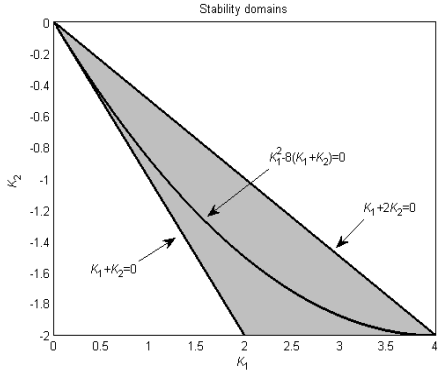

Figure 5. Convergence domain of an autonomous network of 2 self-sampled PLLs. The shaded area corresponds to values of filter coefficients for which the PLLs synchronize.

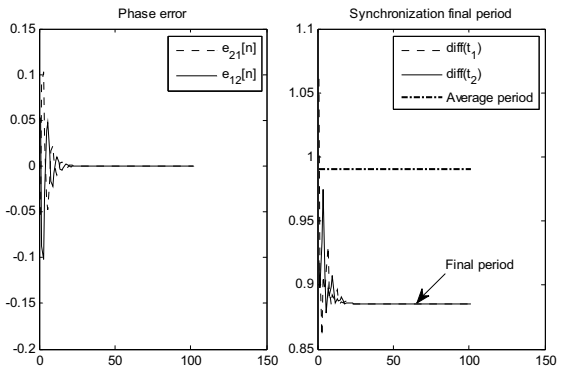

Figure 6. Synchronisation of an autonomous network of two PLLs. Phase error $e_{ji}[n]$ (left) and period $t_i[n]$-$t_i[n$-1] (right).

Note that because of the presence of an integrator in the PLL, the final synchronization period of the autonomous PLL network is not equal to the average period (T_1+T_2)/2 of the two PLLs but depends on the initial conditions of the system (Fig.6).

IV. CONCLUSION

The basic building block of an active clock distribution network based on ADPLLs was described in this paper. In spite of its apparent simplicity, the behavior of such a system is quite complex to analyze and some questions remain open as the rigorous determination of the entire synchronization domain of the system. It was shown that a self-sampled PLL behaves as a piecewise linear system and sufficient conditions for the convergence of this system were derived and verified with simulations, with good agreement. A more thorough investigation of the convergence issues, using piecewise quadratic Lyapunov functions, is in progress. Finally, it should be noted that the self-sampled PLL may also be designed so that the instantaneous phase error $\varepsilon[n]$ is set to 0 instead of $e_{ri}[n-1]$ when $e_{ri}[n]$ is positive. More generally, since it is always possible to know if a clock is ahead or behind another, one might consider replacing the phase error missing by the prediction of this error, based on measurements of the previous edges. One can thus expect a convergence of the self-sampled PLL in the stability domain of a "standard" digital PLL and to their similar behaviour around an operating point. The simple case of an autonomous network of two self-sampled PLLs was also treated. It was shown that the dynamics of such a network are much simpler to analyze than those of a single PLL. The analysis of the network can be reduced to that of a linear system. Necessary and sufficient conditions of convergence were established for this network.

The extension of this work to networks of arbitrary dimensions is the subject of ongoing work.

ACKNOWLEDGMENT

This work is supported by the French National Agency of Research (ANR) through the HODISS project.

REFERENCES

[1] G. A. Pratt, and J. Nguyen, “Distributed Synchronous Clocking”, IEEE Transactions on Parallel and Distributed System, vol. 6, 1995, pp. 314-28.

[2] V. Gutnik, and A. P. Chandrakasan, “Active GHz Clock Network Using Distributed PLLs”, IEEE Journal of Solid-State Circuits, vol. 35, 2000, pp. 1553-60.

[3] S. Al-Araji, Z. Hussain, M. Al-Qutayri, “Digital Phase Lock Loops - Architectures and Applications”, Springer, 2006.

[4] G. Feng, “Stability analysis of piecewise discrete-time linear systems”, IEEE Transactions on Automatic Control, vol. 47, 2002, pp. 1108-12.

[5] M. Johansson and A. Rantzer, “Computation of piecewise quadratic Lyapunov functions for hybrid systems,” IEEE Transactions on Automatic Control, vol. 43, 1998, pp. 555–559.

[6] R. Flynn, and O.Feely, “Limit Cycles in Digital Bang- Bang PLLs”, 18th European Conference on Circuit Theory and Design (ECCTD), Aug. 2007, pp 731-734.

[7] M. Curtin and P. O'Brien, “Phase Locked Loops for High-Frequency Receivers and Transmitters-3”, Analog Dialogue (Analog Devices), July/August,1999.