Performance analysis of an arbiter using probabilistic timed Petri nets

Marco A. Escalante†, Luciano Lavagno‡, and N. Dimopoulos*

†Intel Corp., 5200 NE Elam Young Parkway, Hillsboro, OR 97124, USA
‡Dip. di Elettronica, Politecnico, C. Duca degli Abruzzi 24, 10129 Torino, ITALY
*Dept. of ECE, University of Victoria, P.O. Box 3055, Victoria, BC V8W 3P6, CANADA

Abstract

Asynchronous circuits have experienced a resurgence due to potential advantages to low-power and high-speed design. The timed signal transition graph model, an interpreted Petri net, is a formal specification that is capable to characterize the temporal behavior of asynchronous circuits by utilizing interval delays. However asynchronous circuits are prone to metastability, which may subside after an arbitrarily long time. Thus an interval analysis yields a very pessimistic description of the performance bounds. In this paper we take a probabilistic approach based on a Petri net model to study the effect of metastability in the performance of asynchronous circuits. We illustrate our approach by modeling the Seitz' arbiter as a case study.

1. Introduction

Interval timing analysis has been used to determine the worst-case performance of asynchronous circuits modeled by timed signal transition graphs [1]. There are some situations in which a worst-case analysis is not very appropriate. For instance if some of the involved delays are unbounded.

In this paper we consider the problem of modeling the performance of an arbiter. An ideal 2-way arbiter controls the access to a shared resource that can only service one client at a time. Such an arbiter can accept up to two requests at any time, but it will produce at most one grant even if the requests arrive simultaneously.

A typical circuit that implements an arbiter is the Seitz arbiter shown in Figure 1. If only one of the requests is generated, the corresponding grant is produced after some delay. However if both requests arrive at (about) the same time, the NAND latch may enter a metastable state, and the resolution time τ_m, after which only one of the grants is generated, can be arbitrarily long. The probability distribution function that describes the time τ_m when a circuit that has entered a metastable behavior leaves such state is given by (cf. [7]):

$$f_{\tau m}(\tau_m) = C\, e^{-K\tau m} \qquad \text{[Eq. 1]}$$

where C and K are constants that depend on properties of the circuit elements. Notice that in the Seitz' arbiter, the differential detection circuit after the NAND latch always exhibits a well defined binary output, not being affected by the metastable behavior that may take place in the SR latch.

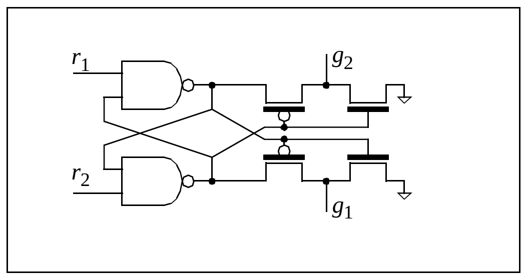

Figure 1. Seitz' arbiter.

If one wants to determine the worst delay of a grant from a request for the Seitz' arbiter, the answer would be "arbitrarily long", which lacks a quantitative notion. Instead of using a timed Petri net with intervals associated with its places which has the limitation that can only characterize a worst-case scenario, we propose to model the metastable behavior using a probabilistic Petri net in which random variables are associated with its places [3] because it allows us to quantify a possibly unbounded delay by obtaining its probability density function (pdf).

Figure 2 shows a partial timed Petri net [8, 4] that represents the *timed* behavior of an ideal arbiter. In interval *timed* Petri nets, a compact non-empty time interval is associated with each place [1, 4]. A transition fires immediately when all its input places have a *visible* token. When a transition fires, it consumes the tokens on its input places, and sends tokens to its output places. A place labeled with interval Δ_i that receives a token at time τ, will make the token visible at time $\tau + \tau_i$, where $\tau_i \in \Delta_i$.

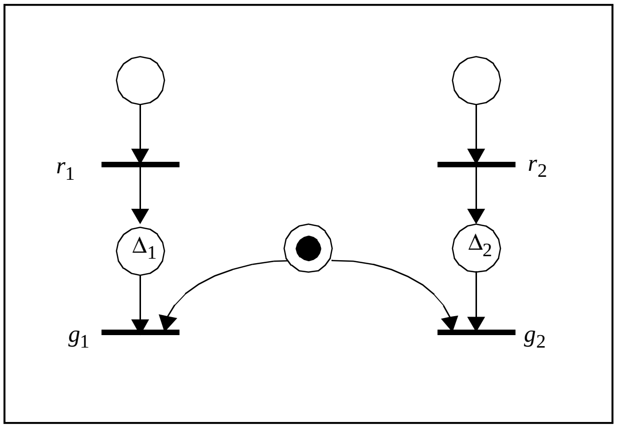

Figure 2. Arbiter.

To understand the behavior of the net shown in Figure 2, let us assume that the token shown in the common input place to transitions g_1 and g_2 is already visible. Suppose that a token is made visible at the input place of r_1. Then transition r_1 fires and sends a token to place labeled τ_1. When the token *matures* (i.e., the token is made visible) in place p_1 labeled with interval Δ_1, and if there is no visible token at place p_2 labeled with interval Δ_2, then transition g_1 fires. Thus the grant enabled by the the first visible token (at place p_1 or place p_2) is the only one that fires. If tokens at both places p_1 and p_2 mature exactly at the same time, one of the grant transitions g_1 or g_2 is chosen to occur non-deterministically.

Clearly this Petri net cannot model the richer behavior of the Seitz' arbiter, since it does not distinguish between meta-stability and normal (digital) behavior. In the next section we will propose a more accurate model that takes meta-stability into account. In order to do so, we have to consider probabilistic Petri nets [3]. In such nets, each place is associated with a random variable which is characterized by a probability density function. This variable represents the random maturing time of the token, relative to the time when the token arrives in the place.

2. Model of the Seitz' arbiter

In this section we introduce our probabilistic approach to the timing analysis of asynchronous circuits by working out a case example: the Seitz' arbiter.

Throughout this paper we will make the following assumptions: (i) the circuit responds with a fixed delay if the separation between the requests is greater than T_w; (ii) if the requests arrive within T_w of each other, the probability that a grant is generated after delay τ_m is given by Eq. 1; (iii) strictly speaking K depends on the time of arrival of the requests but in this paper we assume that K is invariant.

We propose to model the Seitz' arbiter with the Petri net shown in Figure 3. To understand the behavior of the Petri net shown in Figure 3, consider first the case in which a request arrives and the other request is not issued during the window T_w.

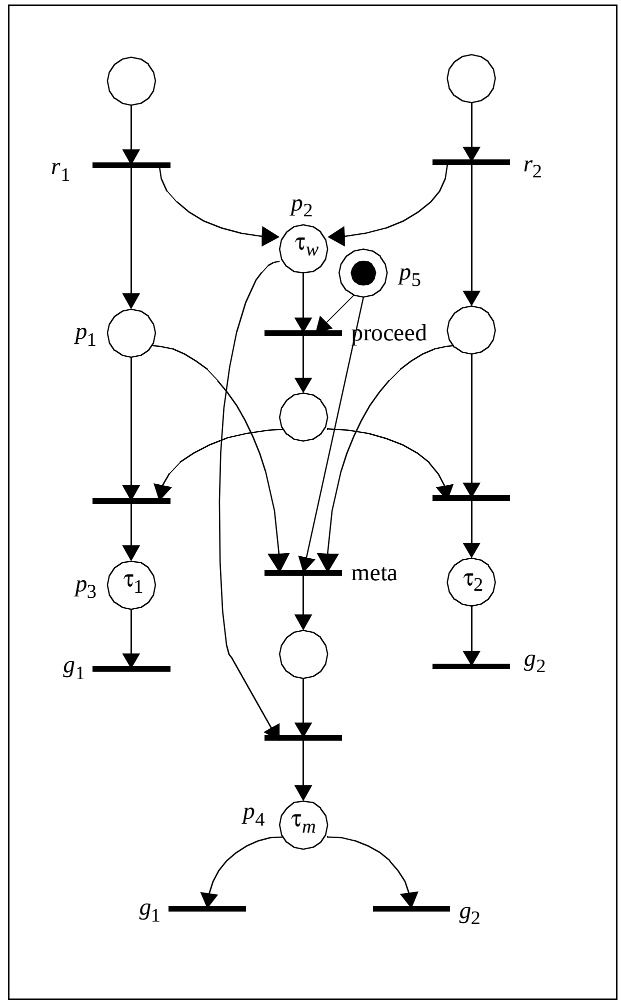

Figure 3. Modeling metastability.

Due to symmetry, it suffices to consider only request r_1. When transition r_1 fires after a token matures at its input place, it puts tokens into places p_1 and p_2. Unlabeled places such as p_1 make tokens visible immediately (i.e., the pdf $f_x(x)$ of the corresponding associated random variable x is the impulse function $\delta(x)$). Place p_2 is labeled with random variable τ_w with pdf shown in Figure 4. Thus transition *proceed* will fire after T_w, and transition g_1 will fire after D_1. The total delay from the occurrence of r_1 to the issuance of the respective grant is $T_w + D_1$.

However if request r_2 fires within window T_w after r_1 has fired, then transition *meta* will fire and either g_1 or g_2 (as selected by the free choice place p_4) will be generated after a delay τ_m. Random variable τ_m obeys an exponential pdf as given by Eq. 1.

3. Analysis

In this section we discuss how to analyze the Petri net shown in Figure 3. We assume that the pdf's of the time of occurrence for requests r_1 and r_2, τ_{r1} and τ_{r2}, are known and given by $f_{\tau r1}(\tau_{r1})$ and $f_{\tau r2}(\tau_{r2})$ (see Figure 5). (In [3] we show how to find the pdf of a given transition for a subclass of probabilistic timed Petri nets.) Our goal is to determine the probabilistic profile (i.e., pdf) of the grant transitions g_1 and g_2.

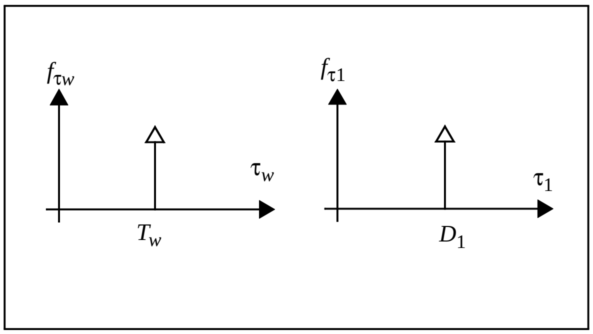

Figure 4. Probability distributions of the random variables associated with labeled places of the Petri net shown in Figure 3.

From the previous analysis it is clear that the firing of transitions *meta* and *proceed* are mutually exclusive (there is a single token in place p_5).

Let us find the time of occurrence of transition g_1. First we introduce some basic concepts from [6]. Let x be a random variable with probability density function (pdf) $f_x(x)$. The probability that variable x takes a value in range $[x_1,x_2)$ is given by:

$$\text{Prob}\{x_1 \le x < x_2\} = F_x(x_2) - F_x(x_1) \qquad \text{[Eq. 2]}$$

where $F_x(x)$ is the accumulative distribution function of random variable x, related to $f_x(x)$ by the following equation:

$$F_x(x) = \int_{-\infty}^{x} f_x(t)\,dt \qquad \text{[Eq. 3]}$$

Using Eqs. 2 and 3, it can be shown that:

$$\text{Prob}\{x_0 \le x < x_0 + dx\} = f_x(x)\,dx \qquad \text{[Eq. 4]}$$

The random variable τ_m associated with place p_4 represents a metastable state and thus it is described by the exponential pdf $f_{\tau m}(\tau_m)$ given by Eq. 1. The probability density functions $f_{\tau ri}(\tau_{ri})$ describe the firing of transitions r_i at time τ_{ri}, for $i = 1, 2$.

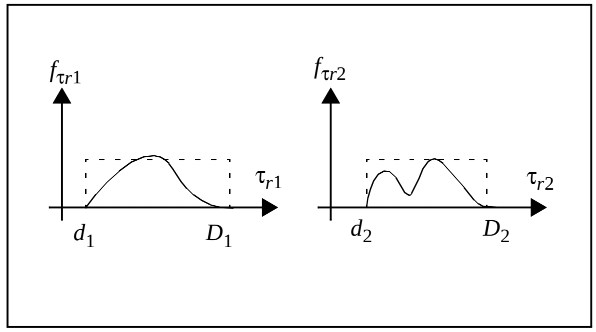

Figure 5. Probability of the time occurrence of requests r_1 and r_2.

From the discussion in Section 2, the probability that transition *meta* will fire at time α (blocking the firing of *proceed*) is:

$$\begin{aligned}&\text{Prob}\{\alpha \le \tau_{meta} < \alpha+d\alpha\} = \\ &\quad \text{Prob}\{\alpha \le \tau_{r1} < \alpha+d\alpha\} \cdot \text{Prob}\{\alpha - T_w \le \tau_{r2} < \alpha\} + \\ &\quad \text{Prob}\{\alpha - T_w \le \tau_{r1} < \alpha\} \cdot \text{Prob}\{\alpha \le \tau_{r2} < \alpha+d\alpha\} + \\ &\quad \text{Prob}\{\alpha \le \tau_{r1} < \alpha+d\alpha\} \cdot \text{Prob}\{\alpha \le \tau_{r2} < \alpha+d\alpha\}\end{aligned} \qquad \text{[Eq. 5]}$$

Similarly the probability that transition *proceed* fires at time α given that transition r_1 has occurred is given by:

$$\begin{aligned}&\text{Prob}\{\alpha \le \tau_{proceed\,|\,r1} < \alpha+d\alpha\} = \\ &\quad \text{Prob}\{\alpha - T_w \le \tau_{r1} < \alpha - T_w + d\alpha\} \cdot \\ &\quad [1 - \text{Prob}\{\tau_{r2} \le \alpha\}\end{aligned} \qquad \text{[Eq. 6]}$$

Thus the pdf's of the occurrence time for transitions meta and proceed are given by:

$$\begin{aligned}&f_{meta}(\alpha) = \\ &\quad [F_{\tau r2}(\alpha) - F_{\tau r2}(\alpha - T_w)] \cdot f_{\tau r1}(\alpha) + \\ &\quad [F_{\tau r1}(\alpha) - F_{\tau r1}(\alpha - T_w)] \cdot f_{\tau r2}(\alpha)\end{aligned} \qquad \text{[Eq. 7]}$$

$$\begin{aligned}&f_{proceed\,|\,r1}(\alpha) = \\ &\quad [1 - F_{\tau r2}(\alpha)] \cdot f_{\tau r1}(\alpha - T_w)\end{aligned} \qquad \text{[Eq. 8]}$$

Let us assume for the sake of illustration that both $f_{\tau r1}$ and $f_{\tau r2}$ are uniform in the interval $[0,D]$ and that $D = 20\,T_w$. Substituting the parameters of the pdf's into Eqs. 7 and 8, one can obtain the following expressions:

$$f_{proceed\,|\,r1}(\alpha) = \begin{cases} \dfrac{D-\alpha}{D^2}, & \text{if } T_w \le \alpha < D \\ 0, & \text{otherwise} \end{cases}$$

$$f_{meta}(\alpha) = \begin{cases} \dfrac{2\alpha}{D^2}, & \text{if } 0 \le \alpha < T_w \\ \dfrac{2T_w}{D^2}, & \text{if } T_w < \alpha \le D \\ 0, & \text{otherwise} \end{cases}$$

If *proceed* has occurred due to r_1, the grant g_1 will be issued at time $\tau_{proceed} + \tau_1$, where τ_1 is the random variable associated with place p_3. To compute the firing time of g_1 we shall use the fact that the pdf of random variable $x = y + z$ is $f_x = f_y * f_z$ if y and z are independent random variables, where the $*$ operator denotes convolution [6].

If, meta has occurred, place p_4 selects either g_1 or g_2, with a 50% chance. (Note: in a first approximation, a non-deterministic choice event can be considered a randomly selected event; an extension of the model could assign a probability to each of the choices of a free choice place). The pdf of random variable τ_m associated with p_4 is $f_{\tau m} = C\, e^{-K\tau m}$, for $\tau_m \ge 0$, and g_1, if selected, will fire at time $\tau_{meta} + \tau_m$.

Thus the probability that g_1 will be issued at time α is given by:

$$f_{g1}(\alpha) = f_{proceed\,|\,r1}(\alpha) * f_{\tau 1}(\alpha) + 0.5\, f_{meta}(\alpha) * f_{\tau m}(\alpha) \qquad \text{[Eq. 9]}$$

The first term corresponds to the generation of g_1 via *proceed* (which is $f_{proceed\,|\,r1}(\alpha - D_1)$, a transport delay) and the second term corresponds to the generation of g_1 via *meta*. Figure 6 shows the pdf of the occurrence of grant g_1 for the uniform case. One can observe a "triangular" shape that corresponds to g_1 generated via *proceed*, and a tail that corresponds to g_1 generated via *meta*. The area under the curve is 0.5 which represents the 50% probability of occurrence of g_1 (g_1 and g_2 being equally likely to occur). The probability that g_1 is generated after a delay > 15 diminishes exponentially. For example the probability that g_1 will be generated after 15time units is approximately 1.9%. Moreover, the probability that g_1 will be generated after 40 time units is under 0.15%.

4. Summary

We have introduced a probabilistic model capable of representing with more accuracy the complex behavior of the Seitz' arbiter, including metastability. The advantage of our approach is twofold: first our analysis procedure relies upon a formal model circuit specification (probabilistic timed Petri nets), and secondly our model is an extension of signal transition graphs (STG's) [2, 5] which are widely used to describe the behavior of asynchronous circuits.

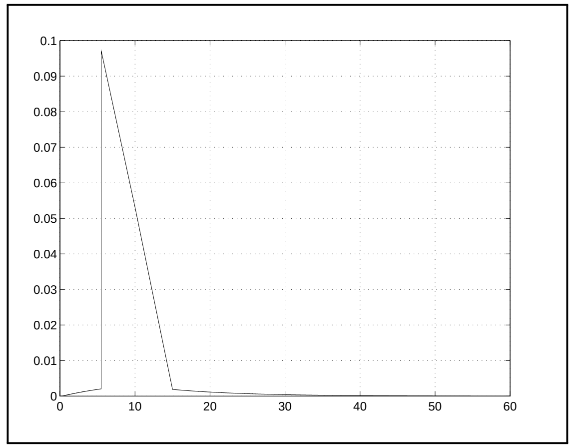

Figure 6. Probability density function of the occurrence time of g_1 for D=10, D_1=5, and K=0.1.

We believe that a probabilistic analysis is essential in the qualitative study of the impact of metastable behavior in the timing performance of asynchronous circuits which can exhibit this phenomenon.

5. Acknowledgments

We are greatly thankful to the anonymous reviewer who read the original version of this paper for his/her insightful comments.

References

[1] T. Amon, H. Hulgaard, S. M. Burns, and G. Borriello, "An algorithm for exact bounds on the time separation of events in concurrent systems," in *Proceedings of the International Conference on Computer Design*, pp. 166–173, 1993.

[2] T.-A. Chu, "On the models for designing VLSI asynchronous digital systems," *INTEGRATION, the VLSI journal*, no. 4, pp. 99–113, 1986.

[3] M. A. Escalante and N. J. Dimopoulos, "A probabilistic approach to timing analysis for synthesis", Technical Report ECE-94-6. University of Victoria, 1994.

[4] M. A. Escalante et. al., "Timing analysis for synthesis of hardware interface controllers using timed signal transition graphs", in *Proceedings of the 8th International Workshop on Petri nets and Performance Models*, pp. 232–240, 1995.

[5] L. Lavagno, "Synthesis and testing of bounded wire delay asynchronous circuits from signal transition graphs," Tech. Rep. UCB/ERL M92/140, U.C. Berkeley, November 1992.

[6] A. Papoulis, *Probability, Random Variables, and Stochastic Processes*, 2nd. edition. New York: McGraw-Hill, 1984.

[7] C. L. Seitz, *System Timing*, chapter 7, pp. 218–262, Reading, Massachusetts: Addison-Wesley, 1980.

[8] J. Sifakis, "Performance evaluation of systems using nets", in *Net Theory and Applications*, Lecture Notes in Computer Science, Springer-Verlag, pp. 307–319, 1980.