Study of the Effectiveness of Adaptive Rate Control for Slow Arrival Processes *

Mohamed Abdelaziz†[†] Ioannis Stavrakakis ‡

†Cascade Communications
5 Carlisle Road,
Westford, MA 01886

‡Dept. of Electrical and Computer Engineering
Northeastern University
Boston, MA 02115

Abstract

In this paper, a queueing model is developed and analyzed to evaluate the performance of a simple rate control scheme in the presence of propagation delays. Unlike previous discrete-time studies of adaptive rate control, the arrival processes are assumed to be slow and deliver cells with a certain time constant. An efficient technique is developed for the computation of the cell loss probabilities. The obtained numerical results show that the cell loss probabilities depend on the relationship between the number of the controlled sources and the available buffer capacity. The results point to the existence of a batch-absorption buffer capacity value beyond which the effectiveness of the adaptive rate control will increase as the source time constant increases.

1 Introduction

Adaptive, feedback-based mechanisms were first introduced to dynamically control the user's traffic (sending windows at the transport layer level) in classical packet switched networks. Examples include the DECBit rate adjustment scheme [1] and the slow start enhancement to TCP/IP [2]. Motivated by the potential for increased network utilization, similar concepts have been recently proposed ([3], [4],[5],[6]) to directly control the rate of traffic sources in ATM networks where rates are either determined explicitly by some network controllers [4] or are incremented and decremented in an adaptive fashion based on some fed-back network information. The ATM Forum has also recently completed the implementation agreement of the Available Bit Rate (ABR) service, which allows the delay tolerant traffic sources to dynamically share the available network resources through the use of adaptive rate control schemes [3].

Due to the large bandwidth propagation delay product characterizing high-speed networks (such as ATM), network feedback information may be outdated by the time it reaches a source. Thus, it is necessary to evaluate the performance of such adaptive control schemes in high-speed networking environments. Several studies have shown that adaptive rate control schemes are adversely affected by propagation delays and several authors have limited the applicability of such schemes to the realm of local area networks [6], [7] where propagation delays are no longer than several tens of the service time.

In addition to the large bandwidth propagation delay products characterizing high-speed networks, an often overlooked — by earlier discrete-time studies such as [7], [8] — characteristic is the potentially large disparity between the relative speeds of a network link and an input source. This disparity or speed-up factor represents the time *in network time slots* required for a source to generate enough bits to form one cell and is referred to as the source time constant [9]. Notice that both the bandwidth propagation delay product and the time between consecutive cell transmissions from the same source (in network time slots) increases as the network link speed increases. And while the former indicates the existence of larger traffic volume in the pipe (in transit), the latter indicates the existence of a lower traffic volume in the pipe coming from the same source. While both these factors clearly influence the performance of an adaptive rate control scheme, only the first has been considered in most previous studies.

In this paper, a queueing model is developed and analyzed to study the impact of the source time constant on the performance of an adaptive rate control system in the presence of propagation delays. The considered model may be adopted to model a switching node of a high speed net-

*Research supported by the Advanced Research Project Agency under Grant F49620-93-1-0564 monitored by the Air Force Office of Scientific Research (AFOSR).

†This work was done while the author was at the University of Vermont

work receiving traffic from lower speed networks or from sources with slow cell generating mechanisms compared to the network time constant (slot or cell transmission time at the network line rate). The considered rate control scheme is presented in the next section. In section 3, a queueing study is presented and in section 4 some numerical results are presented.

2 Adaptive Rate Control

The model considered in this work is shown in Fig. 1 where a number of sources share a common network access node. The access node is modeled by a finite buffer receiving deterministic service. The network node state is reviewed periodically over intervals of length T_u, referred to as update intervals. At the end of each update interval, the network state is communicated to the traffic sources. Individual sources receive the network state information after a time interval equal to their respective propagation delays and respond by modulating their transmission rates. The modulated traffic is received at the network node a round-trip propagation delay time after the network state information is sent.

In this work, a two-level (Binary) adaptation scheme is considered, where sources switch to a high arrival rate (α_1) if the network is underloaded and switch to a low arrival rate (α_0) if the network is overloaded. The two level scheme is primarily chosen for analytical tractability. The precise description of the arrival process is considered in the following subsection.

2.1 The Arrival Process

Arrivals from the l^{th} controlled source can be represented by a two state process $\{r^l\}, r^l \in \{0,1\}$ and is specified by the set of parameters $\{T_c^l, f_i^l(k) : 0 \le i \le 1\}$. T_c^l is an integer constant which represents the minimum time between consecutive arrivals (potential arrival instants) from the l^{th} source and will be referred to as the time constant of the l^{th} source; $f_i^l(k)$ denotes the probability mass function (PMF) for the number of cells generated at potential arrival instants (boundaries separated by the time constant) when the arrival process of the l^{th} source is in state i, $i \in \{0,1\}$.

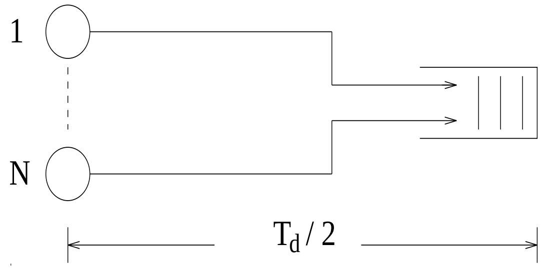

Figure 1: Queueing model for a network access node.

State transitions for the arrival process are governed by the two-level adaptive rate control mechanism and occur at time instants coinciding with the boundaries of the update intervals.

In addition to the above parameters, let c^l, $1 \le c^l \le T_c$, denote the time instant at which the l^{th} source is activated. c^l represents the phase difference between the different arrival processes and will be referred to as the initialization time of the l^{th} source. N such sources with the same time constant T_c are said to be fully synchronized if $c^1 = \cdots = c^N$.

3 Performance Analysis

In this section, the queueing analysis is carried out in discrete-time. Time is slotted and a slot duration is equal to the transmission time of one cell. Although the provided analysis can be used to study sources with different time constants, all sources will be assumed to have the same time constant T_c and the same propagation delay T_d in order to facilitate the presentation. The length of the update interval (T_u) is taken to be equal to the round-trip propagation delay of the sources (T_d). T_c, T_d, T_u are measured in network time slots. These assumptions are relaxed in [10] where the impact of the difference in propagation delays and update intervals is considered.

The round-trip propagation delay T_d is assumed to be an integer multiple of the time-constant T_c, that is, $T_d = M \times T_c$, $M \in Z^+$, where Z^+ is the set of positive integers. Similar analysis can be adopted for the case $T_c = M \times T_d, M \in Z^+$.

Time is divided into frames of length T_d slots; each frame is divided into M subframes, each of length T_c slots. Cell arrivals from the controlled sources arrive at a finite size network buffer of capacity Q and are served according to the FIFO service policy. Arrivals from individual sources occur at subframe boundaries and are assumed to follow a Bernoulli PMF with rate α_1 when in state 1 and rate α_0 when in state 0. Therefore, the aggregate arrival process at subframe boundaries is Binomial with parameters N, α_1 (N, α_0) when in state 1 (0). The aggregate arrival rate from state 1 (0) is $\lambda_1 = \frac{N\alpha_1}{T_c}$ ($\lambda_0 = \frac{N\alpha_0}{T_c}$).

Let $\{q_k\}_{k\ge 0}$ be the buffer occupancy process embedded at subframe boundaries with state space $S^Q = \{0, 1, \cdots, Q\}$. Let q_{th} be the buffer threshold value indicating the network state as overloaded ($q \ge q_{th}$) or underloaded ($q < q_{th}$). Let $\{r_k\}_{k\ge 0}$ be the arrival process embedded at the subframe boundaries with state space $S = \{0,1\}$. Note that since the duration of the frame is equal to the duration of the update interval (T_d slots), state transitions of the process $\{r_k\}_{k\ge 0}$ occur only at *frame* boundaries. Due to the propagation delay, the state of process $\{r_k\}_{k\ge 0}$ at the beginning of frame k is determined by the state of the buffer process at the beginning of frame $k-1$. Thus, the arrival process evolves as follows:

$$r_{k+1} = \begin{cases} 1 & \text{if } q_k < q_{th} \\ 0 & \text{if } q_k \geq q_{th} \end{cases}$$

The system state can be defined at frame boundaries in terms of the quantities (r_k, q_k). Given that $r_k = i$ at the beginning of frame k, the finite queue will evolve over frame k under the i.i.d. batch arrival process with PMF $F_i(\cdot)$. It is easy to establish that the stochastic process $\{r_k, q_k\}_{k\geq 0}$, embedded at frame boundaries, is a finite state Markov chain with state space $S_m = \{(r_1, q_1) : 0 \leq r_1 \leq 1, 0 \leq q_1 \leq Q\}$

The probability that the Markov chain moves from state (r_1, q_1) at the beginning of the k^{th} frame to state (r_2, q_2) at the beginning of the $(k+1)^{st}$ frame will be denoted by $p(r_1, q_1, r_2, q_2)$ and can be evaluated recursively as follows:

$$p(r_1, q_1, r_2, q_2) = \begin{cases} 0 & \text{if } (q_1 < q_{th}), (r_2 = 0) \\ 0 & \text{if } (q_1 \geq q_{th}), (r_2 = 1) \\ p^M_{r_1}(q_1, q_2) & \text{otherwise} \end{cases} \qquad (1)$$

where $p^M_{r_1}(q_1, q_2)$ denotes the probability that the process $\{q_k\}_{k\geq 0}$ moves from state q_1 to state q_2 in M subframes under the assumption that process $\{r_k\}_{k\geq 0}$ remains unchanged over this transition. The i.i.d. arrivals — occurring at the subframe boundaries — will be determined by the PMF $F_{r_1}(\cdot)$. $p^M_{r_1}(q_1, q_2)$, $0 \leq q_1, q_2 \leq Q$ — called the M-subframe-step transition probability under state r_1 — are easily calculated in terms of the 1-subframe-step transition probabilities $p^1_{r_1}(q_1, q_2)$, $0 \leq q_1, q_2 \leq Q$ which are derived as follows. Consider the following definitions:

$\tilde{p}^1_{r_1}(q_1, l)$: denotes the probability that the buffer occupancy process moves from state q_1 to state l over the first slot of a subframe under state r_1, $r_1 \in \{0, 1\}$, $0 \leq q_1, l \leq Q$.

$\tilde{p}^2(l, q_2)$: denotes the probability that the buffer occupancy process moves from state l at the end of the first slot of the current subframe to state q_2 at the end of the current subframe $(0 \leq l, q_2 \leq Q)$. Notice that no arrivals can occur during this transition due to the source time constant and that the buffer occupancy can only decrease.

The 1-subframe-step transition probabilities are given by:

$$p^1_{r_1}(q_1, q_2) = \sum_{l=0}^{Q} \tilde{p}^1_{r_1}(q_1, l)\tilde{p}^2(l, q_2)$$

where $\tilde{p}^2(l, q_2)$, $0 \leq l, q_2 \leq Q$, are given by:

$$\tilde{p}^2(l, q_2) = \begin{cases} 0 & \text{if } q_2 \geq 1, (l - q_2 \neq T_c - 1) \\ 1 & \text{if } q_2 \geq 1, (l - q_2 = T_c - 1) \\ 1 & \text{if } q_2 = 0, (l \leq T_c - 1) \end{cases}$$

and $\tilde{p}^1_{r_1}(q_1, l)$, $r_1 \in \{0, 1\}$, $0 \leq q_1, l \leq Q, \max(0, q_1 - 1) \leq l \leq Q - 1$ are given by:

$$\tilde{p}^1_{r_1}(q_1, l) = F_{r_1}(l - q_1 + 1) + F_{r_1}(0) \cdot \mathbf{1}_{\{q_1 + l = 0\}}$$

$$\tilde{p}^1_{r_1}(q_1, Q) = \sum_{l=Q-q_1+1}^{N_{r_1}} F_{r_1}(l) \quad 0 \leq q_1 \leq Q$$

where $\mathbf{1}_{\{q_1+l=0\}}$ is an indicator function that takes the value 1 if $q_1 + l = 0$ and takes the value 0 otherwise, and N_{r_1} denotes the maximum number of arrivals, as determined by the state of the aggregate cell arrival process, r_1.

The stationary probabilities of the Markov Chain $\{r_k, q_k\}$ $(\pi(r_1, q_1), r_1 \in \{0, 1\}, 0 \leq q_1 \leq Q)$ can be easily computed by solving the system of linear equations $\{P\Pi = \Pi, \Pi e' = 1\}$, where P is the transition probability matrix with elements given by (1), Π is the stationary probability vector and e is the unity vector.

3.1 Cell Loss Probability

In this subsection, the cell loss probability of an arbitrary arriving cell is derived.

In the following, let r_1 be the state of the arrival process at the beginning of the current frame. Let q_1 be the buffer state at the beginning of the k^{th} subframe in the current frame. Let $w_f(r_1, i, k), 1 \leq k \leq M$ be a random variable describing the number of cells lost over the period starting from the beginning of the k^{th} subframe till the end of the current frame (over a period of $M - k$ subframes). Clearly, $w_f(r_1, q_1, 1)$ describes the number of losses over a period of a frame.

It can be easily shown that the loss probability (L) can be derived in terms of the expected value of $w_f(.,.,1)$, $\overline{w_f}(.,.,1)$, as follows:

$$L = \frac{\sum_{(r_1,q_1)\in S_m} \pi(r_1, q_1)\overline{w_f}(r_1, q_1, 1)}{\sum_{(r_1,q_1)\in S_m} \pi(r_1, q_1)\lambda_{r_1} T_d}$$

The random variables $w_f(r_1, q_1, k)$, $1 \leq k \leq M$, $(r_1, q_1) \in S_m$ and their expected values can be derived recursively by noting that the number of cells lost starting from the k^{th} subframe in the current frame till the end of the frame is equal to the number of cells lost over the k^{th} subframe plus the number of cells lost from the beginning of the $(k+1)^{st}$ subframe till the end of the frame. Let $w_s(r_1, q_2)$ be a random variable describing the number of cells lost during the k^{th} subframe at the beginning of which the buffer occupancy is q_2, $(r_1, q_2) \in S_m$. Then:

For $q_1, q_2 \in \{0, \cdots, Q\}, r_1 \in \{0, 1\}, 1 \leq k \leq M - 1$:

$$w_f(r_1, q_1, k) = w_s(r_1, q_1) + w_f(r_1, q_2, k + 1) \quad \text{with probability } p^1_{r_1}(q_1, q_2),$$

and

$$w_f(r_1, q_1, M) = w_s(r_1, q_1).$$

The probabilities $p^1_{r_1}(q_1, q_2)$ were computed in section 3. The random variables $w_s(.,.,.)$ needed for the computation of $w_f(.,.,.)$ can be evaluated as follows. Let a_1 denote a random variable describing the number of arrivals over the first slot of the subframe.
For $r_1 \in \{0, 1\}, q_1 \in \{0, \cdots, Q\}$,:

$$w_s(r_1, q_1) = \begin{cases} (l - Q + q_1 - 1) & \text{if } a_1 = l, \\ & (Q - q_1 + 1) \leq l \leq N_{r_1} \\ 0 & \text{otherwise} \end{cases}$$

By considering the expectation of the above equations, computing $\overline{w}_f(r_1, q_1, M)$, and back substitution, $\overline{w}_f(r_1, q_1, 1)$ can be obtained.

4 Numerical Results

The numerical results presented in this section are derived assuming 2-state Markov modulated Binomial arrivals at subframe boundaries with parameters $(N, \lambda_0/N)$ under state 0 and $(N, \lambda_1/N)$ under state 1.

The induced cell loss probability, as a function of the round-trip propagation delay T_d, is shown in Fig. 2 for various values of the common source time constant T_c. The results have been obtained for $\lambda_1 = 0.9$, $\lambda_0 = 0$, $Q = 70$, $q_{th} = 40$ and $N = 40$. As expected (and established in past research), the effectiveness of the adaptive feedback-based rate control decreases as T_d increases, as indicated by an increased value of the induced cell loss probability. The results also indicate that for a fixed value of T_d, the effectiveness of this control scheme increases as the source time constant increases. This implies that the performance of an adaptive feedback-based rate control scheme may be substantially underestimated if the time constant of the arrival process is assumed to be one ($T_c = 1$).

The impact of the buffer capacity Q and the source time constant T_c on the induced cell loss probability is shown in Fig. 3 for a fixed value of $T_d = 30$; the other parameters were set to $\lambda_1 = 0.94$, $\lambda_0 = 0.9$, $q_{th} = 20$ and $N = 40$. Notice that for $T_c > 1$, the effectiveness of the adaptive rate control scheme can either increase or decrease compared to that for $T_c = 1$. When the buffer capacity is less than N (approximately) the effectiveness of the adaptive rate control scheme decreases as the source time constant increases. The opposite behavior is observed for buffer capacity values exceeding (approximately) N. In view of the fact that the number of sources coincides with the maximum value of the cell batch size, it may be concluded that as long as sufficient buffer space is available to absorb a single batch (called *batch-absorption buffer capacity value*), the effectiveness of the adaptive rate control will increase as the source time constant increases.

The same conclusions can be reached from Fig. 3 where the impact of the number of controlled sources N and the source time constant T_c on the induced cell loss probability is shown for $T_d = 30$, $\lambda_1 = 0.94$, $\lambda_0 = 0.9$, $q_{th} = 30$ and $Q = 50$. One conclusion from these results is that for a given cell loss probability objective and (sufficiently large) fixed buffer capacity, a larger number (N) of slow sources (compared to the number of sources with $T_c = 1$) can be supported, provided that N remains below the buffer capacity.

Notice that for the results presented thus far, it is assumed that the controlled sources are synchronized ($c^1 = \cdots = c^N$) i.e: arrivals from the aggregate process occur in Binomially distributed batches with time constant equal to that of the individual sources. This synchronization represents a worst case in terms of cell loss performance as can be seen from Fig. 5 where $N = 60$, $T^j_c = 30, 1 \leq j \leq 60$, $\lambda_1 = 0.94$, $\lambda_0 = 0.9$, $q_{th} = 30$, $T_d = 30$ and the following distributions of the different sources initialization instants are considered:

- All sources are assumed to be synchronized, thus leading to an aggregate arrival process with $T_c = 30$ and $N = 60$.
- Initialization times of the sources are assumed to be $c^j = 1$, $1 \leq j \leq 30$ and $c^j = 15$, $31 \leq j \leq 60$, thus leading to an aggregate arrival process with $T_c = 15$ and $N = 30$.
- Initialization times of the sources are assumed to be $c^j = 1$, $1 \leq j \leq 20$, $c^j = 11$, $21 \leq j \leq 40$ and $c^j = 21$, $1 \leq j \leq 60$, thus leading to an aggregate arrival process with $T_c = 10$ and $N = 20$.
- Initialization times of the sources are assumed to be $c^j = j$, $1 \leq j \leq 30$ and $c^j = j - 30$, $31 \leq j \leq 60$ thus leading to an aggregate arrival process with $T_c = 1$ and $N = 2$.

Notice that the last case ($T_c = 1$ and $N = 2$) should not be confused with the classical assumption of unity time constant ($T_c = 1$ and $N = 60$) shown in the figure. The overlap in effectiveness of the rate control for $T_c = 30$ is as explained before. As expected spreading the initialization times of the different sources leads to lower loss probabilities and consequently increased applicability of the adaptive rate control.

References

[1] K. K. Ramakrishnan and R. Jain. A binary feed back scheme for congestion avoidance in computer networks. *ACM Transactions on Computer Systems*, pages 158–181, May 1990.

[2] V. Jacobson. Congestion avoidance and control. *ACM Computer Communication Review*, pages 314–329, Aug 1988.

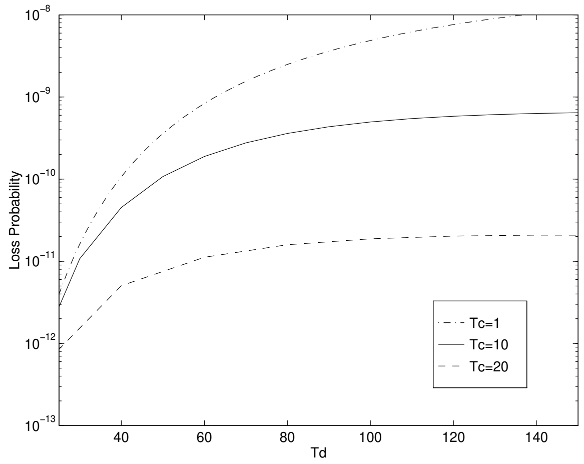

Figure 2: Loss Probability versus the propagation delay for various source time constants

[3] F. Bonomi and K. W. Fendick. The rate-based flow control framework for the available bit rate ATM service. *IEEE Network Magazine*, March 1995.

[4] D.D. Clark A. Charny and R. Jain. Congestion control with explicit rate indication. In *ICC'95*, pages 1954–1963, Seattle, WA, 1995.

[5] N. Yin and M. Hluchyi. On closed-loop rate control for ATM cell relay networks. In *Infocom'94*, pages 99–108, Toronto, Canada, 1994.

[6] P. Newman. Backward explicit congestion notification for ATM local area networks. In *GLOBECOM'93*, pages 719–723, Switzerland, 1993.

[7] Y. T. Wang and B. Sengupta. Performance analysis of a feedback congestion control policy under non-negligible propagation delay. In *Sigcomm'91*, pages 149–157, Zurich, Switzerland, Sep 1991.

[8] M. Murata K. Kawahara, Y. Oie and H. Miyahara. Performance analysis of reactive congestion control for ATM networks. *IEEE Trans. Commun.*, pages 651–661, May 1995.

[9] R. Landry and I. Stavrakakis. Non-deterministic periodic packet streams and their impact on a finite-capacity multiplexer. In *Infocom'94*, Toronto, Canada, 1994.

[10] M. Abdelaziz and I. Stavrakakis. Study of an adaptive rate control scheme under unequal propagation delays. In *ICC'95*, pages 1964–1968, Seattle, WA, 1995.

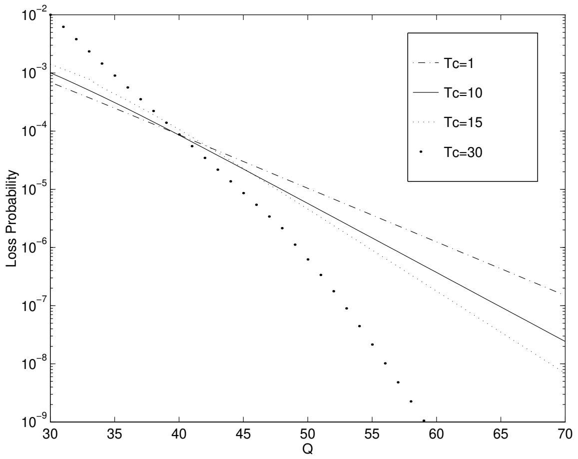

Figure 3: Loss probability versus buffer capacity for various source time constants ($N = 40$)

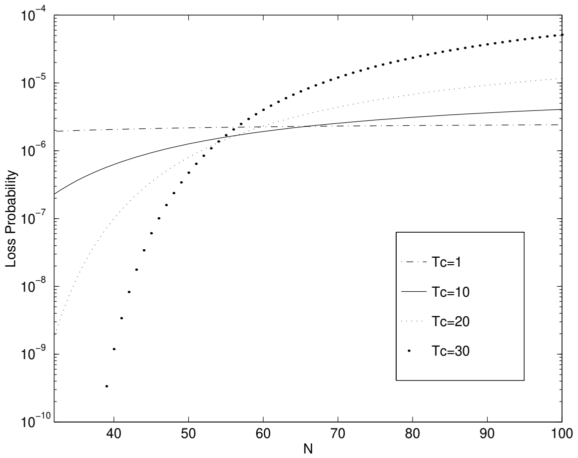

Figure 4: Loss probability versus the number of controlled sources for various source time constants ($Q = 50$)

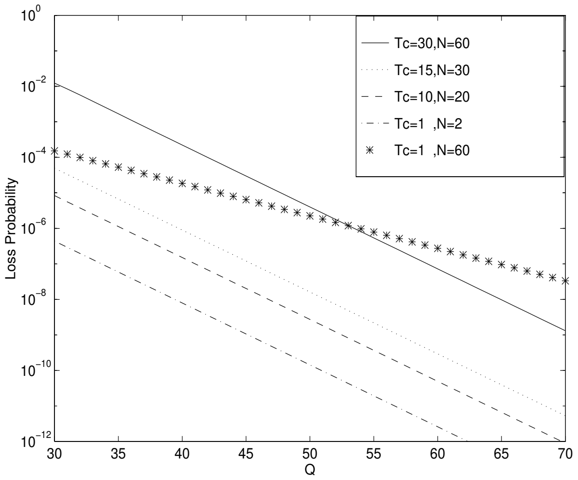

Figure 5: Impact of Source Synchronization