2. Clock distribution networks

2.1. Synchronous digital systems

In a synchronous digital system, data signals are usually stored in a clocked register. This register is waiting for the arrival of the clock signal. When it happens, data signal leaves the bistable register and propagates through the combinatorial network and, for a properly working system, enters in the next register and is fully latched into that register before the next clock signal appears. In conclusion, the components that make up a general synchronous system are the following three subsystems:

1) Memory storage elements.
2) Logic elements.
3) Clocking circuitry and distribution network.

Clock signals are usually considered as simple control signals. However, they have some characteristics and attributes very special. Clock signals must be provided with a very clean and sharp waveform because they are the time reference of any data signal in a synchronous digital system. Furthermore, technology scaling particularly affects to clock signals. For example, the interconnect line resistivity is hugely increased when the line dimensions are decreased (width and thickness) and, consequently, the delay time grows considerably. This increased line resistance is one of the main reasons of the growing interest in the study of clock distribution network. Finally, differences in the delay of clock signals when they arrive at their destinations, the clock skew, can seriously affect system performance, as well as create catastrophic race conditions in which a data signal may be incorrectly latched in a register. The proper design of the clock distribution network ensures that critical timing requirements are satisfied and no race conditions exist.

Two registers are sequentially adjacent if there is at least a sequence of logic blocks that connects the output of the first register R_i (initial) with the input of the second register R_f (final). Otherwise, any event at the output of R_i does not affect the input of R_f.

We can define a local data path in the case of two sequentially adjacent registers (figure 2.1).

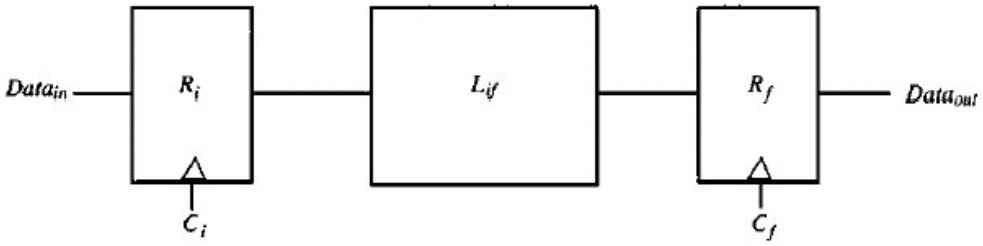

Figure 2.1: Local data path components [FRI-00].

The minimum clock period T_{CP} (min) between any two registers in a sequential data path is given by:

$$T_{CP}(\min) = \frac{1}{f_{CLK}(\max)} = T_{PD}(\min) + T_{skew} \qquad (2.1)$$

where

$$T_{PD} = T_{C-Q} + T_{logic} + T_{int} + T_{setup} \qquad (2.2)$$

The total path delay T_{PD} of a data path is the sum of:

- $T_{C\text{-}Q}$: The maximum time required for the data signal to leave the initial register once clock signal C_i arrives at the initial register.

- T_{logic} and T_{int}: The time necessary for the data signal to propagate through the logic block L_{if} and interconnection lines.

- T_{setup}: The time required for the data signal to successfully propagate and latch within the final register of the data path.

The sum of delays components in (2.2) must satisfy the timing constraints of (2.1) in order to support the minimum clock period T_{CP} (min), which is the inverse of the maximum clock frequency f_{clk} (max). It is important to note that T_{skew} can be positive or

negative depending which clock signal, C_i or C_f, arrives earlier. The waveforms depicted in figure 2.2 show the timing requirement of (2.1) being just satisfied (the data signal arrives at R_f just before the clock signal arrives at R_f).

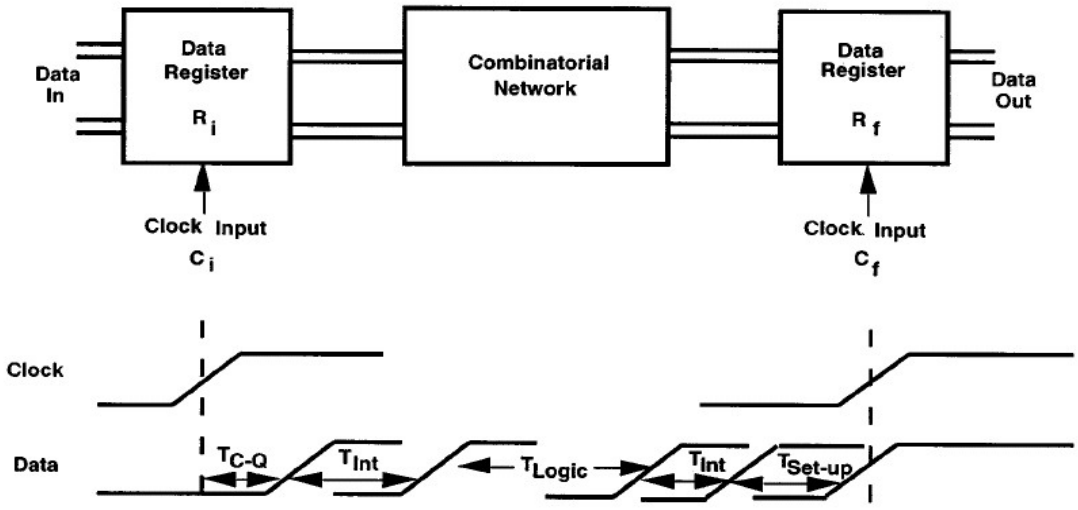

Figure 2.2: Timing diagram [FRI-00].

2.2. Clock distribution networks design

In a synchronous digital system, the clock signal is used to define a time reference for the movement of data signals within it. The clock distribution network distributes the clock signal from the source point to all the elements that need it. Since this function is vital to the operation of a synchronous system, much attention has been given to the characteristics of these clock signals and the electrical network used in their distribution. Clock signals are often regarded as simple control signals, however, these signals have some very special characteristics and attributes that must be fulfilled.

Many approaches for designing clock distribution networks exist, from customized design techniques to automatic algorithms. The requirement of distributing a strictly controlled clock signal to each synchronous register, which are located on a large hierarchically structured integrated circuit, within some specific temporal bounds is difficult and problematic. Furthermore, there are important tradeoffs among system

speed, physical die area and power dissipation when the clock distribution network is planned [FRI-00].

The most common strategy when a clock distribution network is designed is to divide it into different hierarchical levels. The difficulty of applying symmetric clock distribution strategies is that they do not easily support the ability of dividing large VLSI systems into hierarchically structured functional blocks (in figure 2.3, the floorplan of a structured custom VLSI circuit is depicted). Preferably, each large functional block would contain its own locally optimized clock distribution network (lower hierarchical levels) to satisfy the local timing of that particular functional block.

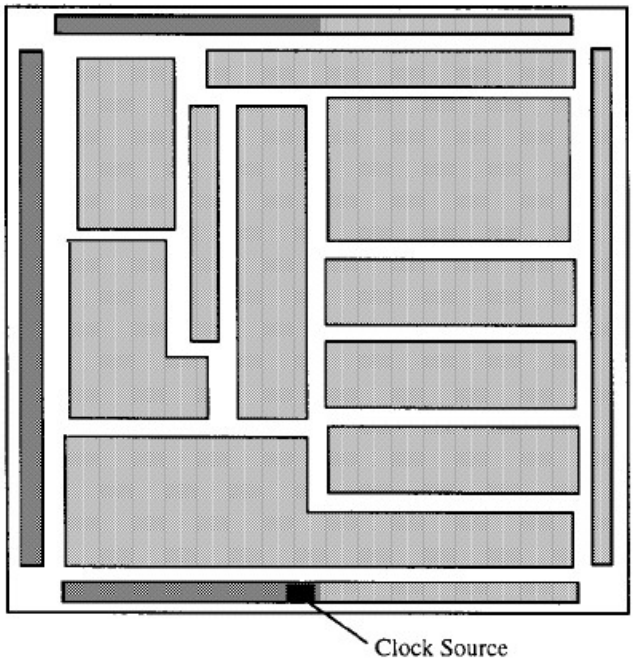

Figure 2.3: Floorplan of a structured custom VLSI circuit [FRI-00].

In conclusion, in a hierarchical clock distribution network, there will be a global clock distribution network in the highest level that distributes the clock signal to each lower level where different clock distribution strategies can be applied in order to optimize the performance of a particular functional sub-block. Each level of the structure can be implemented in a different metal level, preferably on the higher levels where the line sizes are bigger and their resistance lower. This hierarchical structure is depicted in figure 2.4:

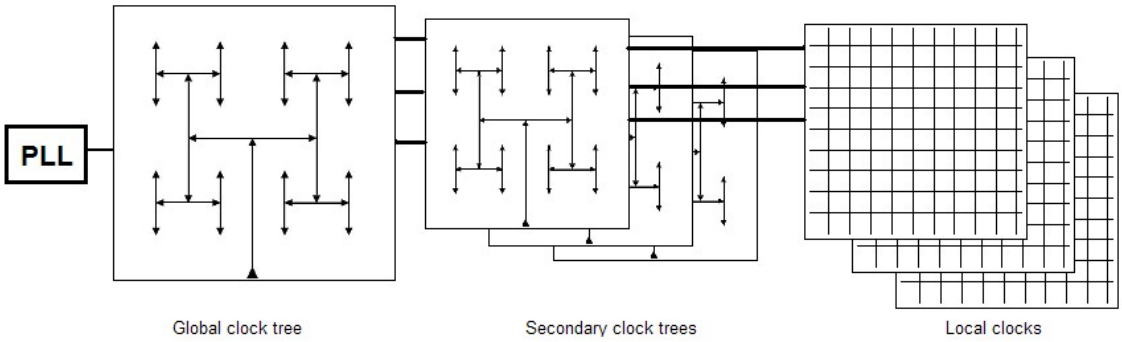

Figure 2.4: Hierarchical clock distribution system.

With regard to the clock generation, there are many issues to consider in the clock generator design. Both power dissipation and jitter impose important constraints when the generation system is chosen, but operating frequency plays the major role, especially in today's systems. A phase locked loop clock generator is nowadays the best solution. Figure 2.5 illustrates the basic block diagram of a PLL-based clock generator.

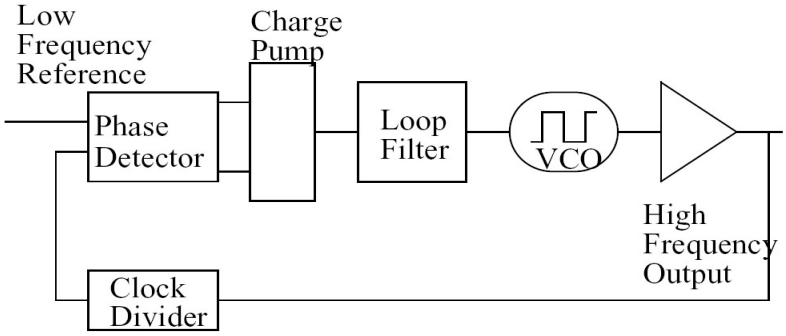

Figure 2.5: Generic PLL block diagram [STE-97].

In a synchronous system, the use of a PLL clock generator can eliminate the clock skew caused by active device variations. Since the output of the PLL is in-phase with the input, if clock distribution buffers of the distribution network are included within the loop, their delay is essentially removed from the system. This does not eliminate the chip-to-chip or on-chip clock skew due to passive device variations and interconnect delay [STE-97].

Various clock distribution strategies have been developed. The most common and general approach to equipotential clock distribution is to use buffered trees (section

2.2.1), especially for custom VLSI circuits and in the final levels of a hierarchical clock distribution network. Furthermore, there are different topologies for these buffered trees that can be applied in some cases to satisfy specific conditions (section 2.2.2). A mesh version exits. Also, there is a grid-like structure. Other very important strategy is the symmetric H-tree, which is generally used in the global part of clock distribution networks, distributing the clock signal to each different functional block. Compensation techniques to control clock signal properties within the clock distribution network exist and are presented in section 2.2.3.

2.2.1. Buffered clock trees

The most common strategy to distribute the clock signal in VLSI systems is to insert buffers, either at the clock source and/or along clock paths, forming a tree structure as shown in figure 2.6. The single clock source is called root. The trunk is the initial tree portion. Branches are the individual paths driving each register. These registers are called leaves.

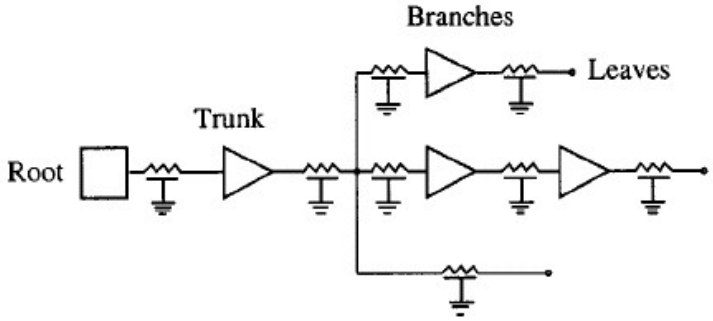

Figure 2.6: Buffered clock tree [FRI-00].

If the buffer interconnect resistance at the clock source is small compared to the output buffer resistance, a single buffer is often used to drive the whole clock distribution network. The primary requirement is that the buffer must provide enough current to drive the entire network capacitance, both interconnection and fanout, while maintaining high-quality waveform (short transition times). Also, it is necessary to miniminize the interconnection resistance effects by ensuring that the output buffer resistance is much higher than the interconnection section resistance that it drives. This

strategy may be appropriated if the clock completely distributed on metal, making load balance of the network less critical [FRI-00].

An alternative to use a single buffer at the clock source is to distribute buffers throughout the clock distribution network. This approach requires additional area, but improves greatly the precision and control of the clock signal waveform and it is necessary if the interconnection line resistance is not insignificant. Furthermore, distributed buffers serve the double function of amplifying the clock signal degraded by the distributed interconnection impedances and isolating local clock distributing networks (lower levels in the hierarchal structure) from upstream load impedances. The number of buffer stages between the clock source and each clocked registers depend upon the total capacitance load and the maximum permissible clock skew. It is worth noting that buffers are the main source of clock skew within a well-balanced clock distribution network since active devices characteristics vary much more greatly than passive device characteristics.

2.2.2. Topologies

- Binary tree

Some early clock routing algorithms [RAM-89], [JAC-90], [KAH-91] defined the delay as a measure of the total wire length along a path. These algorithms attempt to equalize the lengths of each net from the root of the clock tree to each of the leaf nodes. The strategy used is to construct binary tree-like structures with the clock pins at the leaf nodes. Minimal skew clock distribution networks are created using a recursive bottom-up approach. The point where two zero-skew clock sub-nets connect is chosen such that the effective delay from that point to each clocked register is identical. This process continues up the clock distribution tree. The connection point of each new branch is chosen to satisfy the zero-skew design goal. The layout process terminates when the root (or source) of the clock tree is reached. The schematic diagram of this geometric matching process is illustrated in figure 2.7.

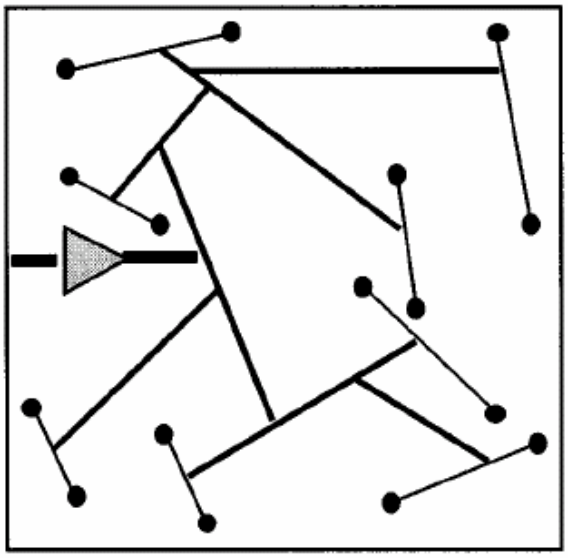

Figure 2.7: Geometric matching for zero clock skew [FRI-00].

Elmore delay instead of the path length can be other design rule in order to determine the binary tree. Here, the delay along a path is the summation of the products of the branch resistance and the downstream capacitance of every branch on the path from the root to the clock pin of the register:

$$T_{Di} = \sum_{k} R_{ki} C_k \qquad (2.3)$$

where C_k is the capacitance at node k and R_{ik} is the resistance of the portion of the unique path between the input and the output node i, that is common with the unique path between the input and node k [FRI-00].

Automated algorithms present many problems. For example, no intermediate buffers can be placed and a single clock buffer must drive the whole network. Also, the possibility of impossible wiring schemes can arise (two wires crossing in a point).

- Symmetric H-trees and X-trees

One of the most important approaches for distributing clock signals utilizes a hierarchy of planar symmetric X-tree or H-tree (figure 2.8) to ensure zero clock skew. It

is achieved by maintaining identical the distributed buffers and interconnections from the clock source to each clocked register.

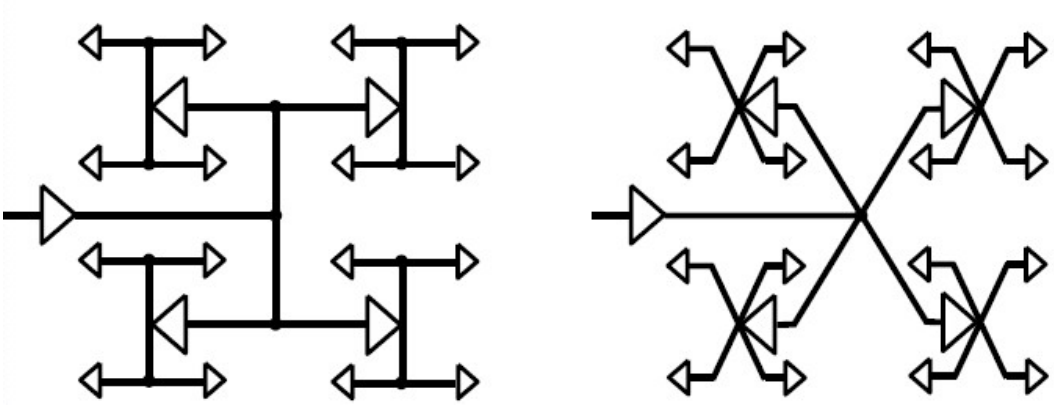

Figure 2.8: H-tree and X-tree.

In this approach, the clock driver is connected to the centre of the main “H” (or “X”). The clock signal is transmitted to the four corners of the main “H”. These four almost identical clock signals are provide to next level “H” structures. The process continues through several levels progressively. The H-tree end points are used to drive local registers or to provide the clock signal to the next hierarchical clock network level.

Here, each clock path, from the clock source to the local registers, has practically the same delay. The primary cause of different delays at the clock paths is the process parameter variations that affect interconnection impedances and, in particularly, any distributed buffer. The amount of clock skew in a H-tree clock distribution network is strongly dependent upon the physical size, semiconductor process control and the number of distributed buffers utilized in the network [FRI-00].

Planar H-tree structures place constraints on the physical layout of clock distribution networks as well in the methodology design. Furthermore, the interconnect capacitance and the power dissipated is much greater as compared with standard buffered clock trees since the wires total length tend to be greater. This capacitance growth exemplifies an important tradeoff between clock delay and clock skew when high-speed clock distribution networks are being designed [FRI-00].

Symmetric structures are utilized to minimize the clock skew. However, an increase in clock signal delay is incurred. Therefore, the increased clock delay must be considered when choosing between buffered tree and H-tree clock distribution networks. Also, since clock skew only affects sequentially adjacent registers, the obvious advantages of using highly symmetric structures to distribute clock signals are significantly degraded. There may, however, be certain sequentially adjacent registers distributed across the integrated circuit. For this situation, a symmetric H-tree structure is particularly appropriate to distribute the global portion of the clock network [FRI-00].

- Mesh

Occasionally, a mesh version of a clock tree structure is used. Nodes are connected in other metallization level by a mesh to minimize the interconnect resistance within the clock tree. A mesh structure places the branch resistances in parallel, minimizing the clock skew. In the figure 2.10, a 3D view of a mesh structure is illustrated.

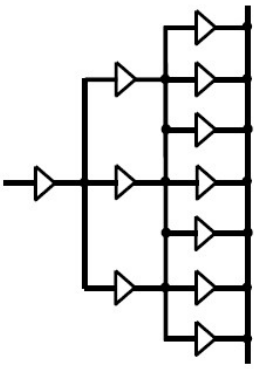

Figure 2.9: Mesh structure.

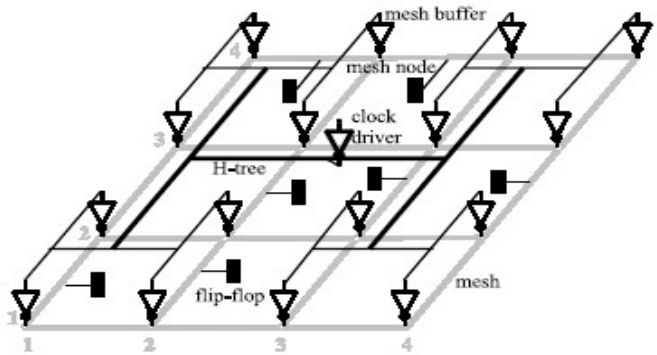

Figure 2.10: 3D view of a mesh structure [YEH-06].

This is the traditional way to distribute clocks in high-performance microprocessors. Under any H-tree, there is usually placed a mesh. It achieves very low skew and jitter (i.e., robustness to variations), but sacrifices power and is difficult to analyze.

- Grid

A clock grid is probably the simplest clock distribution structure and is usually utilized for distributing the clock signal in the final stages of a hierarchical network. Essentially, it consists in a grid of wires driven by one or more buffers. In figure 2.11, this structure is depicted. The main advantage of a clock grid is that the distribution of the clock signal to each destination point is simple. The reason is that the regular structure arrives near to each place where clock signal is needed. Furthermore, the clock grid ensures good process variation tolerances and an easy design methodology.

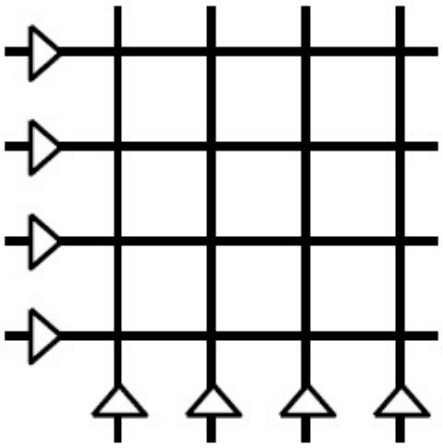

Figure 2.11: Grid structure.

There are some disadvantages with this structure. The main of them is that the grid wires increase the total capacitive load of the clock network, increasing the power consumption. Furthermore, the clock grid is inefficient with regard to used area.

2.2.3. Compensation techniques

It is usual to utilize compensation techniques when clock distribution networks for customized VLSI circuit are designed. They are necessary to minimize the variation of interconnect impedances and capacitive loads between different clock paths. Figure 2.12 shows this situation. Nodes i, j and k represent different loads. The paths to these nodes also have different buffering levels. Buffer locations are often chosen so that the active buffer output impedance is comparable to or greater than the interconnect resistance seen at the buffer output. This characteristic ensures that the locally distributed interconnect section can be accurately modelled as being mostly capacitive. However, in general, the interconnect impedance should be modelled as a distributed RC section of interconnect.

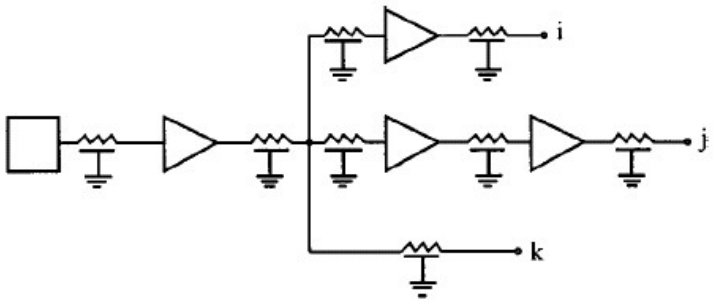

Figure 2.12: Different interconnect impedances and capacitive loads [FRI-00].

There are two important techniques to control and compensate the delay of each clock signal path and minimize the skew between them. Above all, they are effective in buffered tree clock distribution networks. First of them is to insert passive RC elements [IBM-85]. The second is to size geometrically transistor widths [FRI-86].

Clock buffers are located along the clock path such that the highly resistive long lines drive loads with low capacitance and vice versa. Placing a centralized module of parameterized clock buffers within the buffered tree (figure 2.13) it is possible to control the clock skew. Parameterizing the current of clock buffers in the module compensates the variation of the clock signal delay between each one of the functional elements.

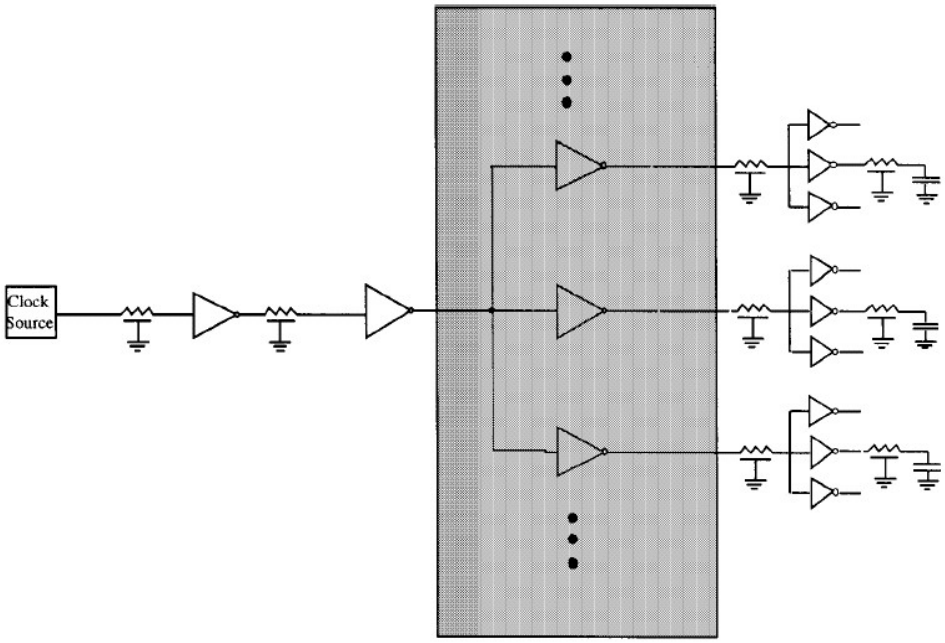

Figure 2.13: Module of parameterized buffers [FRI-00].

- Compensation advantages

The most important advantage of using compensation techniques is controlling and reducing the on-chip skew. Furthermore, the clock signal delay from the clock source to the clocked register is reduced. The reason is the improved partition of the RC loads. Since inverters located within each functional block often drive large capacitive loads, the interconnection impedance (resistance in particular) driven by any specific clock buffer is small compared to output buffer resistance. The fairly long distances of the clock signal paths are fairly resistive. However, these paths are isolated of the highly capacitive loads. In this way, RC time constants are decreased, reducing the total clock delay (RC delay was analysed in section 2.3.1).

Other advantage is that the partitioning of clock system becomes easier with the buffer insertion. The VLSI circuit design can be partitioned hierarchically in many different functional sub-blocks in which different clock distributions strategies can be implemented according different requirements. The utility of compensation techniques is very dependent upon the ability of characterize the interconnect devices and impedances within the clock distribution network. With accurate impedance

estimations, parameterized buffers can be designed to satisfy a specific clock skew schedule.

- Compensation disadvantages

Is important to note an important disadvantage of this compensation technique. Unlike interconnection impedances, transistor conductance tends to be highly sensitive to variations in supply voltage and process and environmental conditions (radiations and temperature for example). The delay of a clock signal path dominated by device impedances is more susceptible to the clock skew than a path dominated by interconnection impedances.

2.3. Clock distribution network parameters

The function of a clock distribution network is to provide a global time reference to all the clocked registers of a synchronous digital system. This way, it regulates the flow of information within the chip. In this section, main parameters involved in a clock distribution network are described. They are essential when the clock system is designed.

2.3.1. Propagation delay

One of the targets when a clock distribution system is designed is to minimize the propagation delay time of the clock signal from the source to the registers. It is required because the clock period time must be larger than the time for the clock signal to propagate from the clock source to the end of the longest clock path. The clock signal in the source cannot change before it arrives to every destination.

A clock signal path within a clock distribution network has a single input, the clock generator, and a single output, the register clock input. Branches are typically composed of distributed buffers and interconnect sections. In order to calculate the clock path

delay and skew, a simple model of a CMOS inverter driving another inverter with line resistance and capacitance between the two inverters is often used (figure 2.14).

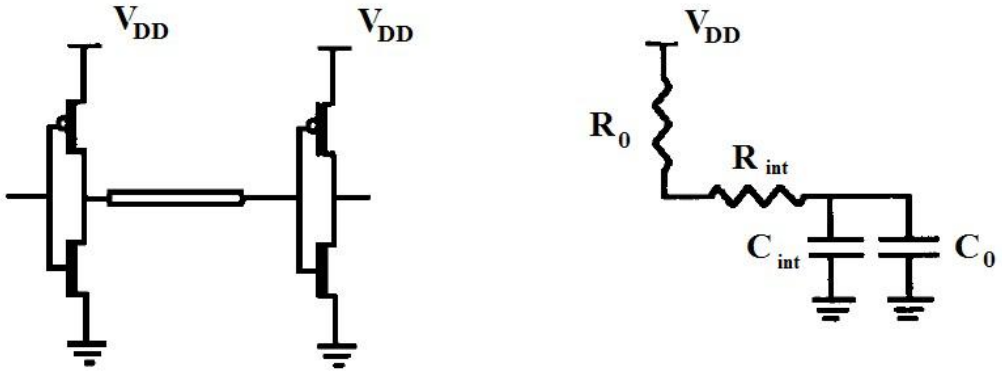

Figure 2.14: Interconnection delay model [AFG-89].

A well-known empirical estimate of the rise or fall time (90 %) of a single CMOS inverter driving an interconnect section with a capacitive load (representing the following CMOS inverter) is [SAK-93]:

$$T_{R/F} = 1.02 R_{int} C_{int} + 2.30\left(R_0 C_{int} + R_{int} C_0 + R_0 C_0\right) \qquad (2.4)$$

where R_{int} and C_{int} is the resistance and capacitance of the interconnect section respectively, and R_0 and C_0 is the output on-resistance of the driving buffer and the input load capacitance of the following buffer respectively. These four parameters are given by:

$$\begin{aligned}
&R_0 = \frac{1}{K \cdot \left(V_{DD} - V_T\right)}, \quad K = \frac{\mu \cdot C_{ox} \cdot W}{L_{eff}}, \quad C_{ox} = \frac{\varepsilon_{ox}}{t_{ox}} \\
&C_0 = C_{ox} \cdot W \cdot L_{eff} \\
&R_{\text{int}} = \frac{\rho}{W_{\text{int}} \cdot t_{\text{int}}} \cdot L_{\text{int}} \\
&C_{\text{int}} = \varepsilon_{ILD} \cdot L_{\text{int}} \cdot \left[\frac{W_{\text{int}}}{T_{ILD}} + 0.77 + 1.06 \cdot \left(\frac{W_{\text{int}}}{T_{ILD}} \right)^{0.25} + 1.06 \cdot \left(\frac{t_{\text{int}}}{T_{ILD}} \right)^{0.5} \right]
\end{aligned} \qquad (2.5)$$

where:

- W and L_{eff}: width and effective length of the transistor.
- C_{ox}: gate unit area capacitance.
- t_{ox}: gate oxide thickness.
- μ: charge carrier mobility.
- V_T: threshold voltage.
- ρ: metal resistivity.
- ε_{ox}: oxide dielectric constant.
- ε_{ILD}: interlevel dielectric constant.
- W_{int}, L_{int} and t_{int}: width, length and thickness of the interconnection line.
- T_{ILD}: Interlevel dielectric thickness.

C_{int} is derived from the empirical formula given in [WES-94] that includes the contribution of fringing fields. Occasionally, they are not taken into account.

The physical delay model represented by (2.3) is a fairly simple approximation of the delay of a CMOS inverter driving a distributed *RC* impedance. More complex and accurate delay models exist.

2.3.2. Clock skew

Clock skew is another important parameter to take into account when a clock distribution network is designed. It imposes timing constraint as propagation delay does. Clock skew can be defined as the difference in arrival times of clock signals to two sequentially adjacent clocked registers. In figure 2.15, clock skew is illustrated. There are two possibilities depending on in which register the clock signal arrives before.

Clock distribution systems are generally designed to achieve that clock signals arrive just at the same moment to every destination (zero skew design), but there are several factors that can cause differences in the propagation delay time between different clock paths (i.e. different line lengths, differences in delays of any active buffer, process and environmental parameter variations). In this case, the clock skew

can seriously affect to the system performance, reducing the maximum clock system frequency and causing catastrophic race conditions.

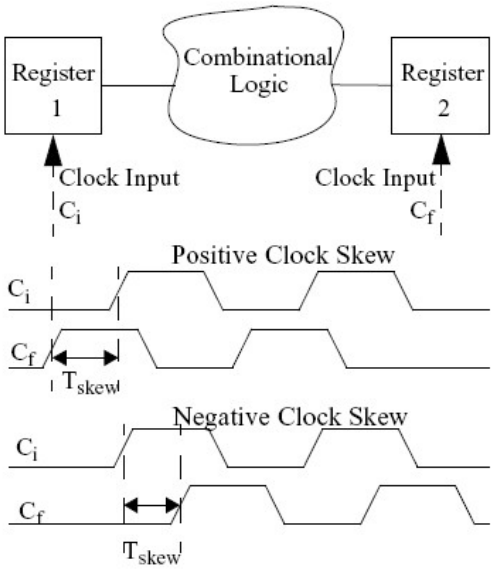

Figure 2.15: Clock Skew [STE-97].

A complete analysis of the clock skew is presented in section 3.1.

2.3.3. Clock jitter

Jitter represents the time varying behaviour of the clock signal. The main effect of jitter is that clock edge locations are shifted with regard to the ideal edge location. Therefore, the duty cycle of a clock signal could be different to the ideal value.

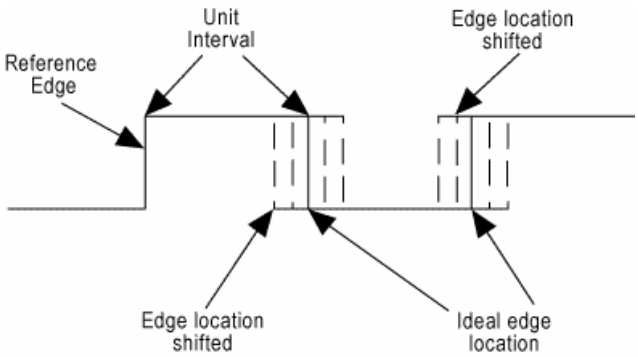

Figure 2.16: Clock jitter.

Clock source has always a specific clock jitter, but also clock distribution network introduces jitter in the clock signal. The noise sources that contribute to jitter most significantly are the following [STE-97]:

- Noise coupled through the circuits' power and ground connections.
- Noise coupled through adjacent or intersecting wires.
- Noise inherent to the circuits' transistors themselves.

2.3.4. Clock power dissipation

Another essential parameter in the design of a clock distribution network is the power dissipated by it. In a VLSI modern system, clock distribution networks have to drive thousands of registers, creating a great capacitive load that should be efficiently sourced. Therefore, each clock signal transition changes the state of each capacitive node within the network. Both high capacitive loads and the continuous demand of higher frequencies have led to an increasingly larger proportion of the total power of a system dissipated within the clock distribution network, in some applications much greater than 25% of the total power.

Total power consumption can be divided in three main sources:

$$P_{total} = P_{dyn} + P_{SC} + P_{lkg} \qquad (2.6)$$

where:

- P_{dyn}: It is the average dynamic power dissipation caused by the continuous charging and discharging of the capacitive nodes within the clock distribution network (parasitic interconnection capacitances and buffer input capacitances).

$$P_{dyn} = fC_{total}V_{DD}^2 \qquad (2.7)$$

- P_{SC}: It is the short-circuit power dissipation. The cause is that in a CMOS inverter, there is a small period in which both transistors, NMOS and PMOS, are simultaneously conducting. The power dissipated in this period is dependent on short-circuit current I_{SC}, input and output signal slopes, device parameters and supply voltages.

$$P_{SC} = I_{SC} V_{DD} \qquad (2.8)$$

- P_{lkg}: It is the power dissipated when there is no activity on the circuit. It is due to the leakage current I_{lkg}: reverse-bias diode leakage and the sub-threshold leakage. The first occurs when the drain-to-bulk voltage of either the PMOS or NMOS transistor is reverse biased. The second occurs when the gate-to-source voltage of the inverter is below threshold voltage, but it is enough to put the inverter into weak inversion.

$$P_{lkg} = I_{lkg} V_{DD} \qquad (2.9)$$

For CMOS VLSI circuits, the primary component of power dissipation is the dynamic power. The goal is to minimize the expression (2.6) without decrease system clock frequency. Therefore, for a given circuit implementation low dynamic power dissipation is best achieved by employing certain design techniques that minimize the power supply and/or the capacitive load.

2.4. Clock distribution networks examples

There are a lot of different examples to show different approaches for clock distribution networks in synchronous digital VLSI systems, from highly specialized customized circuits to commercial processors. These examples are going to be shown in chronological order to present a historical perspective of the general clock distribution design problem and related tradeoffs over the past two decades.

2.4.1. The DEC/Compaq 64-Bit Alpha Microprocessor Family

The DEC/Compaq 64-bit Alpha microprocessor family represent perfectly the evolution of VLSI circuits during the last decades and specifically their clock distribution networks. In the following table, some parameters of three DEC/Compaq microprocessors are shown:

Generation	**21064**	**21164**	**21264**
Technology (μm)	0.75	0.5	0.35
Transistors (million)	1.68	9.3	15.2
Frequency (MHz)	200	300	600
Die size (mm)	16.8 x 13.9	18.1 x 16.5	16.7 x 18.8
Capacitive clock load (nF)	3.25	3.75	----
Power supply (V)	3.3	3.3	2.2
Supply current (A)	< 10	~ 15	> 30
Total power (W)	~ 30	~ 50	~ 72
Power percentage dissipated within the clock distribution network	40 %	40%	44 %

Table 2.1: Characteristics of three Alpha Microprocessor generations [FRI-00].

In the Alpha 21064 microprocessor (year 1992), a single-phase clock signal is distributed globally on the higher level of the metal process (M3). The reason is that the resistivity per unit length of the third layer and the metal to substrate capacitance are smaller compared to the other layers.

The distribution of the loads is asymmetric, so it is necessary a specialized strategy for the clock network. The single 200-MHz clock signal is distributed through five levels of buffering, as shown in figure 2.17, where the total network consists of 145 separate elements. Each of the elements contains four levels of buffering with a final output stage locally driving the clocked registers. Vertical straps are placed on the second level of metal (M2), configuring a clock mesh structure, to minimize the skew that could appear within the initial four-stage portion of the buffer tree. This strategy is accomplished by centrally locating the clock generation circuitry within the integrated circuit [HOR-92].

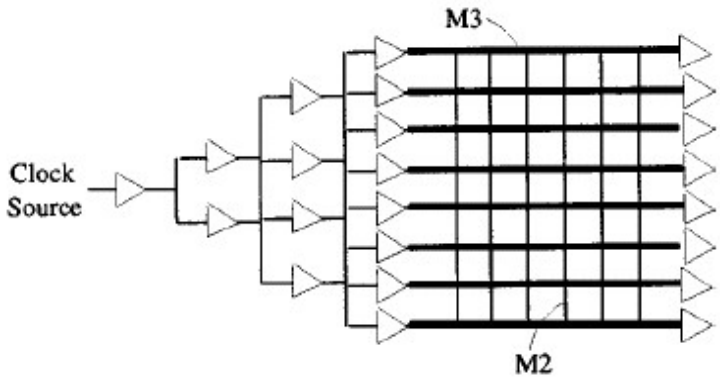

Figure 2.17: Clock distribution network of Alpha 21064 microprocessor [HOR-92].

The Alpha 21264 microprocessor (year 1998) is the first version of the Alpha family that utilizes a hierarchy of clocks. That hierarchy includes a gridded global clock, six gridded major clocks and many local and conditioned local clocks [BAI-98], [GRO-98].

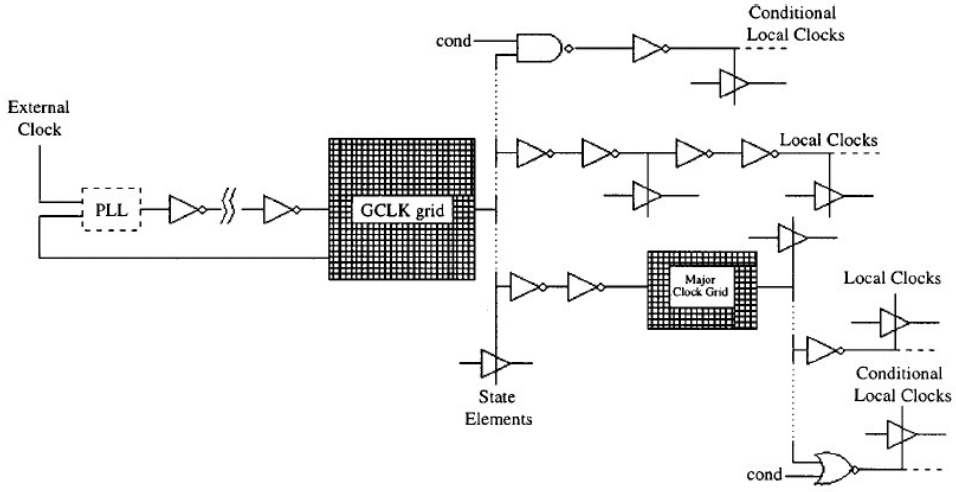

Figure 2.18: Hierarchy in Alpha 21264 microprocessor clock network [GRO-98].

An on-chip PLL is used to generate the clock signal. As shown in figure 2.19, the clock signal is routed from the center of the die and distributed by X and H-trees to 16 distributed clock drivers. The final portion of the clock distribution network uses a tree configuration as compared to an H-tree to save power and area. All interconnections within the GCLK grid are shielded (both laterally and vertically by power/ground lines).

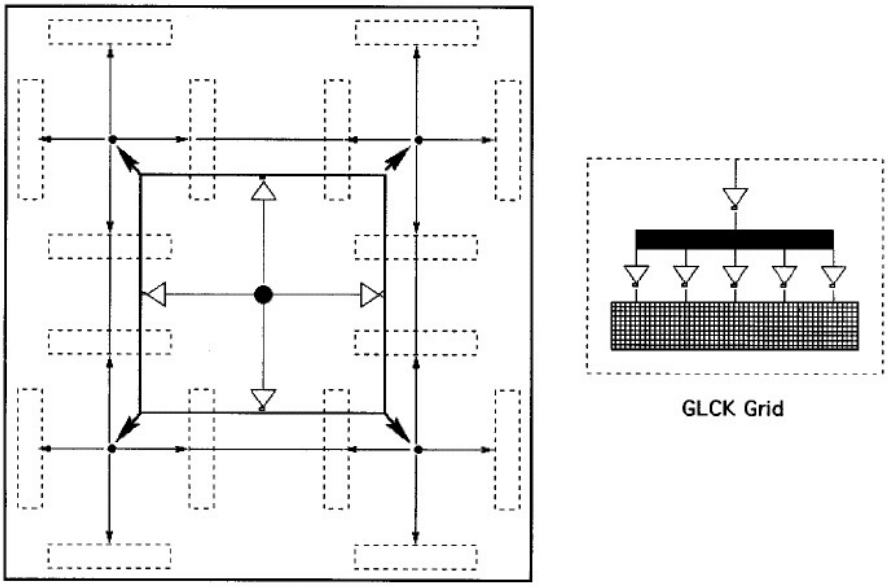

Figure 2.19: Global clock distribution network in the Alpha 21264 [BAI-98].

2.4.2. The Intel IA-64 Microprocessor

The Intel IA-64 is a recent example of a microprocessor designed to operate beyond the gigahertz frequency level [RUS-00]. The general clock distribution design strategy applied in this microprocessor is to minimize the clock skew through the use of distributed programmable deskew circuits while supporting local optimization of the clock distribution network.

The architecture of the IA-64 clock distribution network consists of three components: a balanced tree structure to globally distribute the clock signal, multiple deskew circuits distributed regionally with balanced clock trees, and multiple local clock buffers that drive the individual circuits and registers. In addition, a separate reference clock signal is distributed along with the global clock signal for use within the deskew circuit.

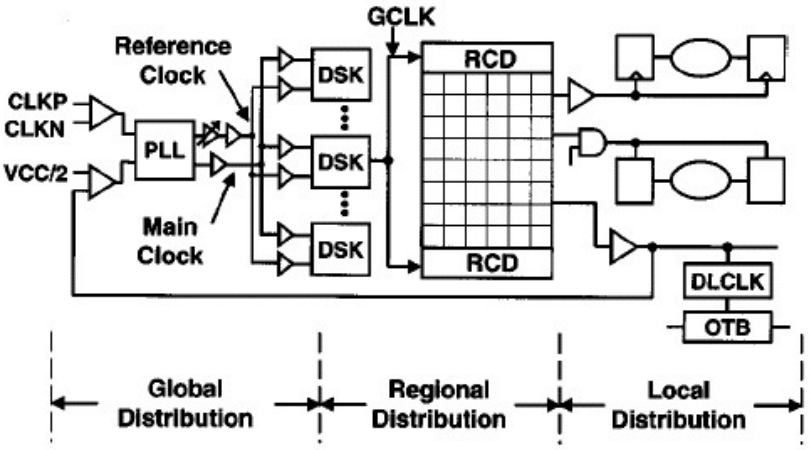

Figure 2.20: IA-64 clock distribution topology [RUS-00].

The global portion of the clock distribution network is originated from the clock source (on-chip PLL), and distributes the clock signal through a network structured as an H-tree until it reaches one of eight deskew clusters, each containing up to four deskew circuits. Clock lines are shielded by ground and power lines in order to minimize the effects of any capacitive and inductive coupling between the clock lines and any adjacent signal lines.

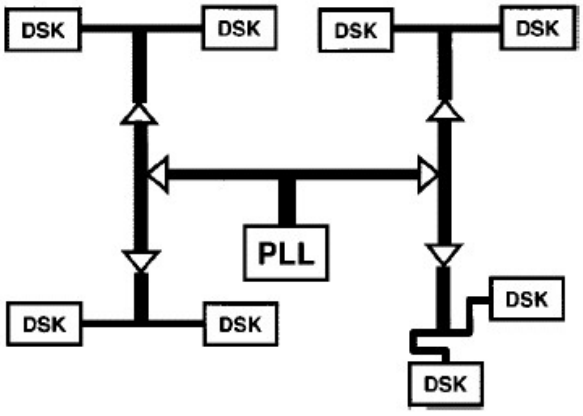

Figure 2.21: IA-64 global clock distribution [RUS-00].

The deskew buffer consists of a phase detector and a digitally controlled analog delay line. The buffer provides local delay compensation by comparing the reference clock signal to a local feedback signal, permitting the delay of the clock signal to be adjusted through the analog delay line and the phase detector.

The output signal of each deskew circuit is distributed through a balanced tree network to a set of local buffers. These buffers drive a regional grid of clock ports. Thus, by applying this intermediate compensation technique, the global clock distribution design problem can be treated as a local design problem, making the entire chip integration process significantly easier and more manageable. Localized clock skew scheduling and clock gating for power management are also supported at this level of the clock distribution network.

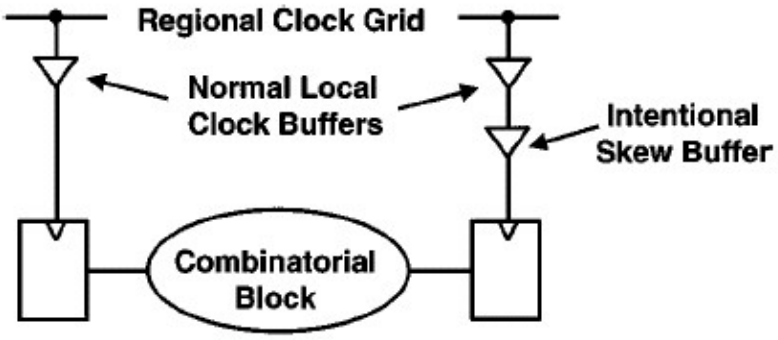

Figure 2.22: IA-64 local clock distribution [RUS-00].

2.4.3. The Intel Itanium 2 Microprocessor

The Itanium 2 microprocessor [TAM-04] is one of the newest developed by Intel in the last years. It is fabricated on the 130-nm CMOS process with six layers of copper interconnects. The processor has a total of 410 million transistors and operates at 1.5 GHz at 1.3 V. The clock distribution network design employs a tree-based differential global clock network. To address the clock skew issue and to increase the frequency of operation, a fuse-based de-skew circuit is implemented, reducing the clock skew caused by on-die process variations and clock network design mismatches.

The clock distribution network is implemented using a multilevel tree architecture. Figure 2.23 shows the clock distribution hierarchy.

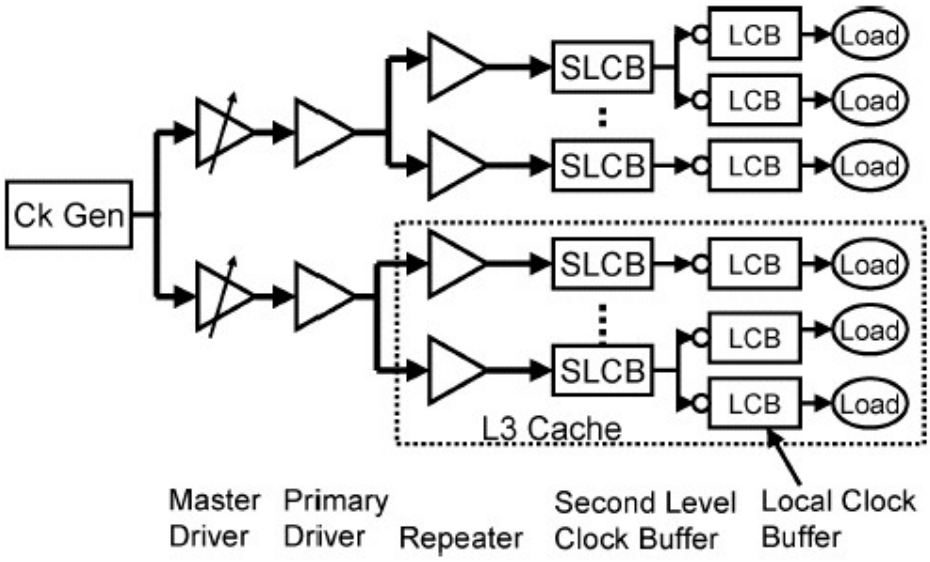

Figure 2.23: Clock distribution hierarchy [TAM-04].

The global clock distribution is a differential two-level tree structure that originates at the clock generator (Ck Gen) and terminates at the second-level clock buffer (SLCB). The master drivers, the primary drivers, and the repeaters are used along the distribution to increase the signal strength.

The zonal clock distribution consists of the SLCB and the local clock buffer (LCB). The SLCB converts the differential global clock to a single-ended zonal clock, SLCBO. There are 23 zonal clocks in the core, strategically partitioned to support the various functional units. Deskew circuits are implemented in this part of the clock distribution network to minimize clock skew effects.

To minimize the inductance effects in the clock networks, the clock lines are shielded in layers to provide low-impedance current return paths. To minimize skin effects, wires wider than about 4 µm are not used.