INTEL TECHNOLOGY AND MANUFACTURING DAY GLOSSARY

- **cell height** – The logic cells in a cell library are rectangular and designed with a fixed height so they can be placed side by side in parallel rows, and readily interconnected. The cell height is the thickness of these rows.

- **cell library** – A library (collection) of hundreds of logic cells such as the NAND2. A designer chooses specific cells and connects them in a specific manner to form a logic circuit – an adder is an example – part of a chip design.

- **COAG (contact over active gate)** – A process feature whereby the gate contact is stacked on top of the transistor gate rather than at its side. Intel has implemented COAG in its 10nm process – an industry first – thereby improving transistor density.

- **CPT (cost per transistor)** – The cost of building a single transistor, calculated as the number of transistors that can be manufactured on a wafer divided by the cost of processing a single wafer. Perhaps the most important benefit of Moore's Law is the steady reduction in CPT, thereby delivering ever more value to the end user.

- **contact** – A conductor that connects one of the terminals of a transistor to a metal layer. Contacts are usually created from tungsten.

- **double patterning or LELE** – The use of two steps (LELE = litho-etch-litho-etch) to create the desired pattern on one layer. Double patterning is required when the wavelength of the lithography tool (scanner) doesn't provide sufficient resolution to create the desired feature sizes. The LELE process has yield and performance risk from misalignment between patterns.

- **dummy gate** – A gate that isn't part of a transistor, put on the edge of a logic cell, needed to isolate one cell from another. Traditional processes used two dummy gates per cell; Intel's 10 nm process requires only a single dummy gate, thereby improving transistor density.

- **FinFET (aka tri-gate transistor)** – Introduced by Intel in 2011, it is a transistor structure whereby the gate wraps around the channel (the region where current flows from source to drain), which is in the shape of one or more vertical fins. This improves performance and reduces power over the traditional planar transistor in which the gate controls channel flow only from above.

- **high-k metal gate** – A technique adopted by Intel in 2007 for reducing transistor leakage (wasted power) while increasing transistor performance, allowing transistor scaling to continue in accordance with Moore's Law. The solution was to replace the transistor's silicon dioxide gate dielectric (a thin layer below the gate) with a hafnium-based "high-k" material.

- **logic cell** – Logic cells contain a small number of transistors connected together to form a simple binary logic or storage function. A simple example is a NAND2 cell, which is built with four transistors. It has two inputs and one output. Each input can be a zero or one (zero is represented by low voltage, one by high voltage). If both inputs are one, then the output is zero. If one or both inputs is zero, the output is one. Logic circuits are built by combining logic cells.

- **metal, interconnect** – An integrated circuit consists of a large number of transistors whose three terminals (source, drain and gate) are connected together in a very specific fashion by conducting interconnects (electrical wires). These interconnects are made of metal and sometimes referred to simply as "metal." The most common metal used for this purpose is copper.

- **Mtr/mm^2** – The unit for transistor density. Measures how many million transistors can be packed into an area 1 mm by 1 mm.

- **nm** – Nanometer, one-billionth of a meter. It takes only four silicon atoms in a crystal lattice to make one nanometer.

- **patterning** – The process of creating a layer of material on a wafer in a specific pattern. The most common way to do so is to cover the wafer with a photoresist (light-sensitive material), shine light through a patterned mask onto the wafer to expose desired portions, then etch away the exposed resist, leaving a particular pattern on the wafer.

- **pitch** – A measure of how closely spaced a linear feature can be in a particular process. For instance, in Intel's 14 nm process, interconnects placed together side by side have a 52 nm pitch, meaning the sum of the minimum width and the minimum spacing is 52 nm.

- **SADP (self-aligned double patterning)** – A patterning technique that uses a spacer material to create features with half the original lithographic pitch. It allows better pitch scaling than LELE while adjacent lines are kept self-aligned to each other.

- **SRAM cell** – A high-speed memory cell, built into arrays, used in logic chips for registers and caches, for temporarily storing data on a processor. It typically contains six transistors. As today's processors tend to have very large caches, an important consideration in the design of new processes is to minimize the SRAM cell size.

- **strained silicon** – A technique adopted by Intel in 2003 for speeding up transistors by creating a compressive strain for PMOS transistors and a tensile strain for NMOS transistors, increasing current flow when a transistor is in the "on" state.

- **transistor** – A tiny switch that controls the flow of electricity. It is the fundamental building block of logic circuits. There are two types, NMOS and PMOS (negative and positive metal oxide semiconductors). The transistor has three terminals: source, drain and gate. In an NMOS transistor, current flows from the source to the drain only when the gate is at high voltage; in a PMOS transistor, current flows only when the gate is at low voltage. NMOS and PMOS are complementary, making today's CMOS processes.

- **transistor performance, drive current** – The performance of a transistor – how fast it can switch from off to on, or vice versa – is proportional to the amount of current that it drives from its drain. An important consideration in the design of a new process is to increase drive current and to reduce power per transistor. Drive current is a factor in determining the maximum frequency at which a chip can run.