“Making of a Chip” Illustrations

22nm 3D/Trigate Transistors – Version

January 2012

The illustrations on the following foils are low resolution images that visually support the explanations of the individual steps.

For publishing purposes there are high resolution JPEG files posted to the Intel website: www.intel.com/pressroom/kits/chipmaking

Optionally same resolution (uncompressed) images are available as well. Please request them from markus.weingartner@intel.com

Sand / Ingot

Sand

Silicon is the second most abundant element in the earth's crust. Common sand has a high percentage of silicon. Silicon – the starting material for computer chips – is a semiconductor, meaning that it can be readily turned into an excellent conductor or an insulator of electricity, by the introduction of minor amounts of impurities.

Melted Silicon –

scale: wafer level (~300mm / 12 inch)

In order to be used for computer chips, silicon must be purified so there is less than one alien atom per billion. It is pulled from a melted state to form a solid which is a single, continuous and unbroken crystal lattice in the shape of a cylinder, known as an ingot.

Monocrystalline Silicon Ingot –

scale: wafer level (~300mm / 12 inch)

The ingot has a diameter of 300mm and weighs about 100 kg.

Ingot / Wafer

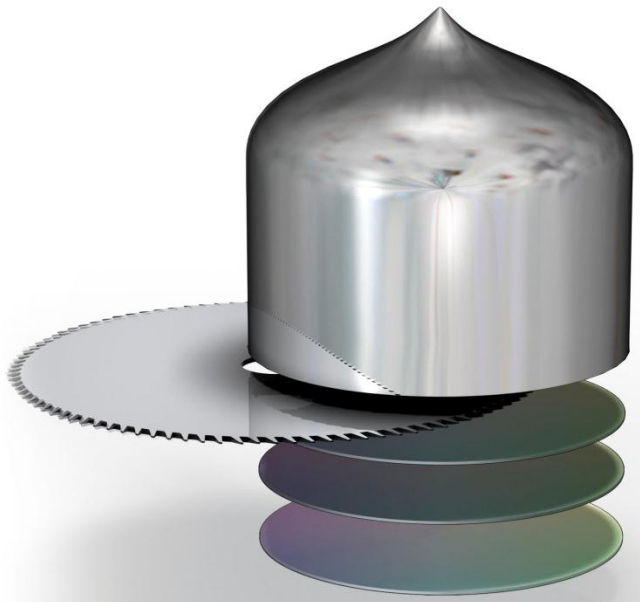

Ingot Slicing –

scale: wafer level (~300mm / 12 inch)

The ingot is cut into individual silicon discs called wafers. Each wafer has a diameter of 300mm and is about 1 mm thick.

Wafer –

scale: wafer level (~300mm / *12 inch*)

The wafers are polished until they have flawless, mirror-smooth surfaces. Intel buys manufacturing-ready wafers from its suppliers. Wafer sizes have increased over time, resulting in decreased costs per chip. when Intel began making chips, wafers were only 50mm in diameter. Today they are 300mm, and the industry has a plan to advance to 450mm.

Fabrication of chips on a wafer consists of hundreds of precisely controlled steps which result in a series of patterned layers of various materials one on top of another.

What follows is a sample of the most important steps in this complex process.

Photolithography

Applying Photoresist –

scale: wafer level (~300mm / 12 inch)

Photolithography is the process by which a specific pattern is imprinted on the wafer. It starts with the application of a liquid known as photoresist, which is evenly poured onto the wafer while it spins. It gets its name from the fact that it is sensitive to certain frequencies of light ("photo") and is resistant to certain chemicals that will be used later to remove portions of a layer of material ("resist").

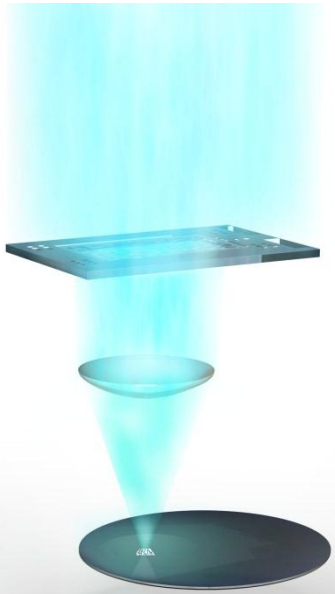

Exposure –

scale: wafer level (~300mm / 12 inch)

The photoresist is hardened, and portions of it are exposed to ultraviolet (UV) light, making it soluble. The exposure is done using masks that act like stencils, so only a specific pattern of photoresist becomes soluble. The mask has an image of the pattern that needs to go on the wafer; it is optically reduced by a lens, and the exposure tool steps and repeats across the wafer to form the same image a large number of times.

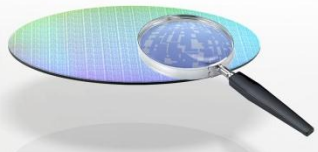

Resist Development–

scale: wafer level (~300mm / 12 inch)

The soluble photoresist is removed by a chemical process, leaving a photoresist pattern determined by what was on the mask.

Ion Implantation

Ion Implantation–

scale: wafer level (~300mm / 12 inch)

The wafer with patterned photoresist is bombarded with a beam of ions (positively or negatively charged atoms) which become embedded beneath the surface in the regions not covered by photoresist. This process is called doping, because impurities are introduced into the silicon. This alters the conductive properties of the silicon (making it conductive or insulating, depending on the type of ion used) in selected locations. Here we show the creation of wells, which are regions within which transistors will be formed.

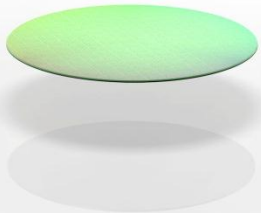

Removing Photoresist–

scale: wafer level (~300mm / 12 inch)

After ion implantation, the photoresist is removed and the resulting wafer has a pattern of doped regions in which transistors will be formed.

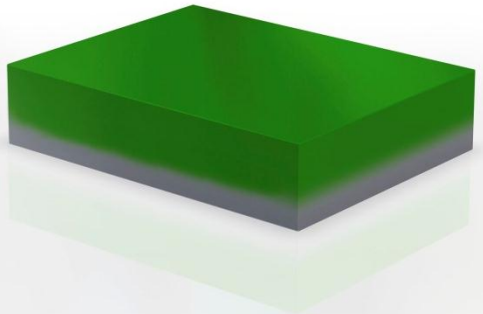

Begin Transistor Formation–

scale: transistor level (~50-200nm)

Here we zoom into a tiny part of the wafer, where a single transistor will be formed. The green region represents doped silicon. Today's wafers can have hundreds of billions of such regions which will house transistors.

Etching

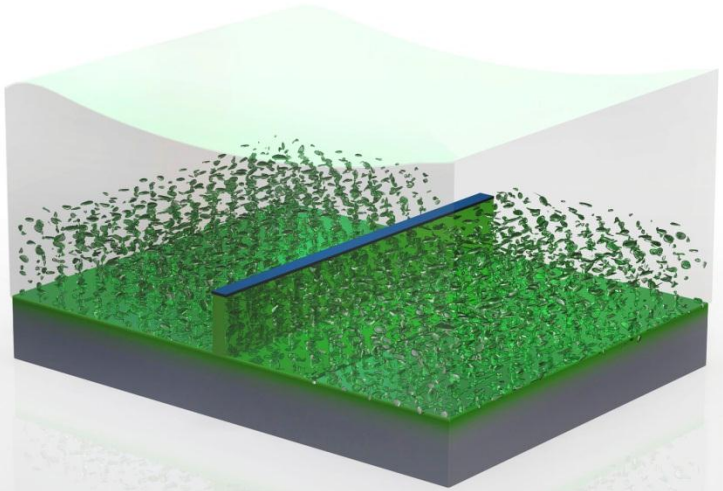

Etch–

scale: transistor level (~50-200nm)

In order to create a fin for a tri-gate transistor, a pattern of material called a hard mask (blue) is applied using the photolithography process just described. Then a chemical is applied to etch away unwanted silicon, leaving behind a fin with a layer of hard mask on top.

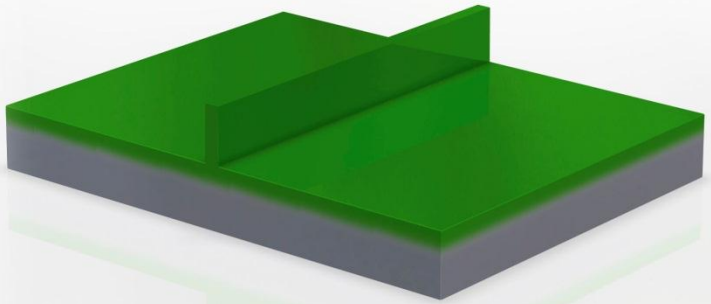

Removing Photoresist –

scale: transistor level (~50-200nm)

The hard mask is chemically removed, leaving a tall, thin silicon fin which will contain the channel of a transistor.

Temporary Gate Formation

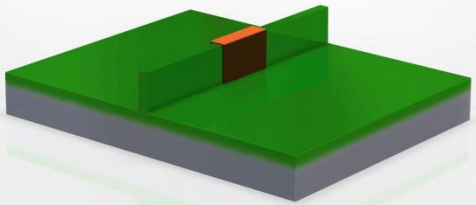

Silicon Dioxide Gate Dielectric–

scale: transistor level (~50-200nm)

Using a photolithography step, portions of the transistor are covered with photoresist and a thin silicon dioxide layer (red) is created by inserting the wafer in an oxygen-filled tube-furnace. This becomes a temporary gate dielectric.

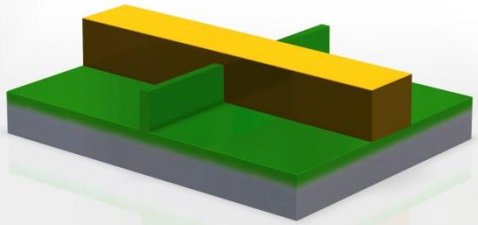

Polysilicon Gate Electrode–

scale: transistor level (~50-200nm)

Again using a photolithography step, a temporary layer of polycrystalline silicon (yellow) is created. This becomes a temporary gate electrode.

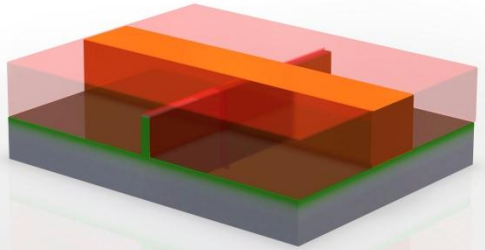

Insulator–

scale: transistor level (~50-200nm)

In another oxidation step, a silicon dioxide layer is created over the entire wafer (red/transparent layer) to insulate this transistor from other elements.

Intel uses a "gate last" (also known as "replacement metal gate") technique for creating transistor metal gates. This is done in order to avoid transistor stability problems which otherwise might arise as a result of some subsequent high temperature process steps.

"Gate-Last" High-k/Metal Gate Formation

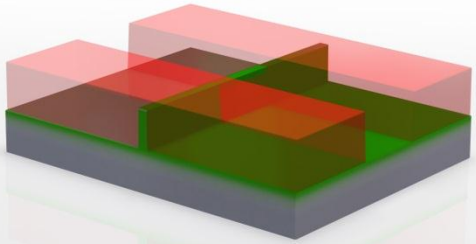

Removal of Sacrificial Gate–

scale: transistor level (~50-200nm)

Using a masking step, the temporary (sacrificial) gate electrode and gate dielectric are etched away. The actual gate will now be formed; because the first gate was removed, this procedure is known as "gate last".

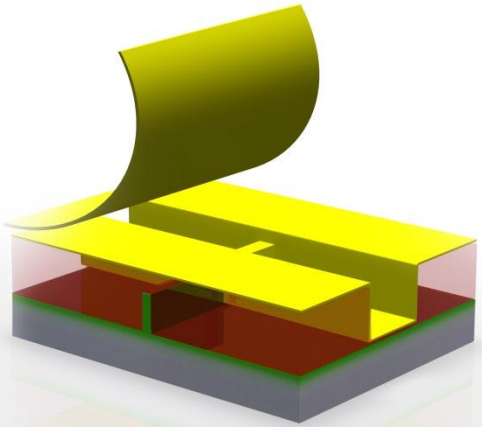

Applying High-k Dielectric –

scale: transistor level (~50-200nm)

Individual molecular layers are applied to the surface of the wafer in a process called "atomic layer deposition". The yellow layers shown here represent two of these. Using a photolithography step, the high-k material is etched away from the undesired areas such as above the transparent silicon dioxide.

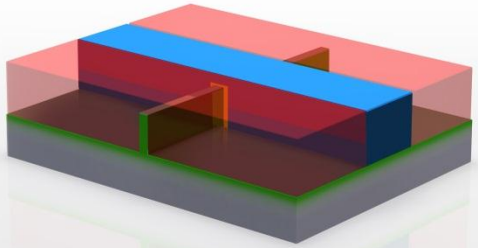

Metal Gate–

scale: transistor level (~50-200nm)

A metal gate electrode (blue) is formed over the wafer and, using a lithography step, removed from regions other than where the gate electrode is desired. The combination of this and the high-k material (thin yellow layer) gives the transistor much better performance and reduced leakage than would be possible with a traditional silicon dioxide/polysilicon gate.

Metal Deposition

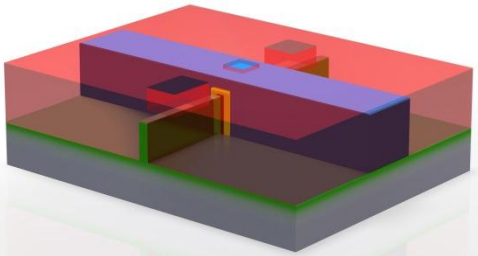

Ready Transistor –

scale: transistor level (~50-200nm)

This transistor is close to being finished. Three holes have been etched into the insulation layer (red color) above the transistor. These three holes will be filled with copper or other material which will make up the connections to other transistors.

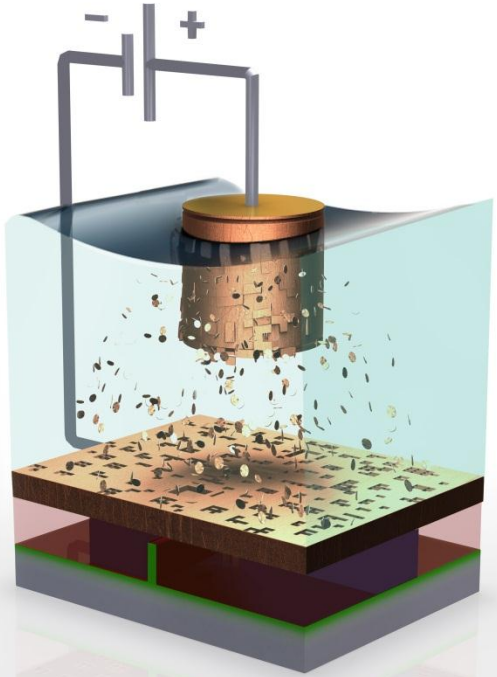

Electroplating –

scale: transistor level (~50-200nm)

The wafers are put into a copper sulphate solution at this stage. The copper ions are deposited onto the transistor thru a process called electroplating. The copper ions travel from the positive terminal (anode) to the negative terminal (cathode) which is represented by the wafer.

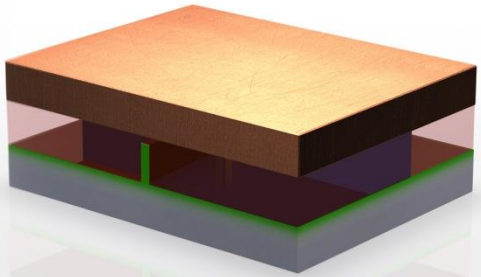

After Electroplating –

scale: transistor level (~50-200nm)

On the wafer surface the copper ions settle as a thin layer of copper.

Metal Layers

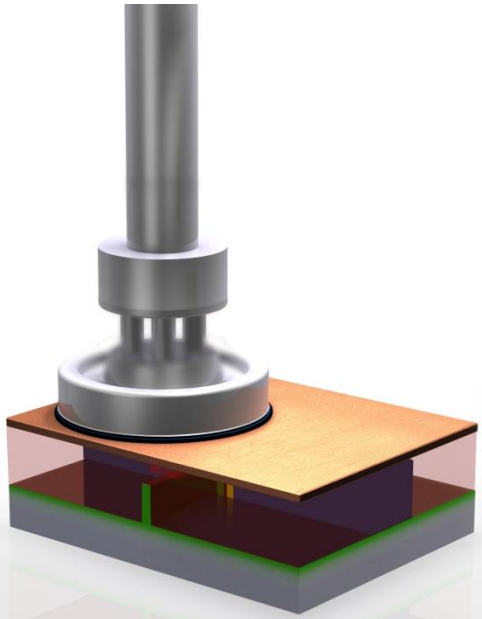

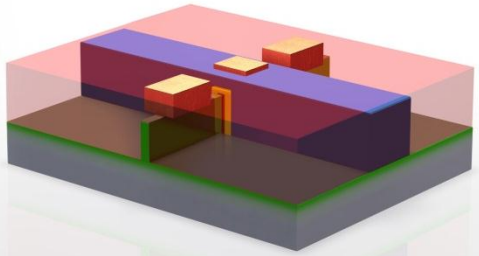

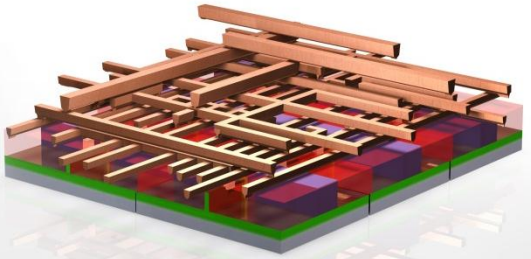

Polishing –

scale: transistor level (~50-200nm)

The excess material is mechanically polished away to reveal a specific pattern of copper.

Metal Layers – *scale: transistor level (six transistors combined ~500nm)*

Multiple metal layers are created to interconnect (think: wires) all the transistors on the chip in a specific configuration. How these connections have to be "wired" is determined by the architecture and design teams that develop the functionality of the respective processor (e.g. 2nd Generation Intel® Core™ i5 Processor). While computer chips look extremely flat, they may actually have over 30 layers to form complex circuitry. A magnified view of a chip will show an intricate network of circuit lines and transistors that look like a futuristic, multi-layered highway system.

After all the interconnect layers are formed, an array of solder bumps is put on each die. These are the electrical connections with which the chip will communicate with the outside world, through the package in which it is later inserted. (these bumps are not shown in the illustrations)

When wafer processing is complete, the wafers are transferred from the fab to an assembly/test facility.

There, the individual die are tested while still on the wafer, then separated, and the ones that pass are packaged. Finally, a thorough test of the packaged part is conducted before the finished product is shipped.

Wafer Sort / Singulation

Wafer Sort –

scale: die level (~10mm / ~0.5 inch)

This portion of a ready wafer is being put through a test. A tester steps across the wafer; leads from its head make contact on specific points on the top of the wafer and an electrical test is performed. Test patterns are fed into every single chip and the response from the chip is monitored and compared to "the right answer".

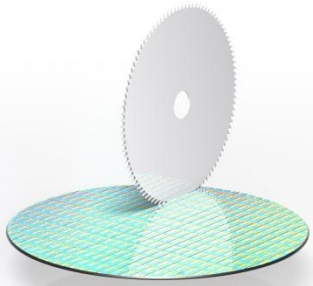

Wafer Slicing –

scale: wafer level (~300mm / 12 inch)

The wafer is cut into pieces (called die). The above wafer contains future Intel processors codenamed Ivy Bridge.

Selecting Die for Packaging –

scale: wafer level (~300mm / 12 inch)

The die that responded with the right answer to the test patterns will be packaged.

Packaging

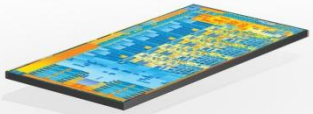

Individual Die –

scale: die level (~10mm / ~0.5 inch)

These are individual die which have been cut out in the previous step (singulation). The die shown here is Intel's first 22nm microprocessor codenamed Ivy Bridge.

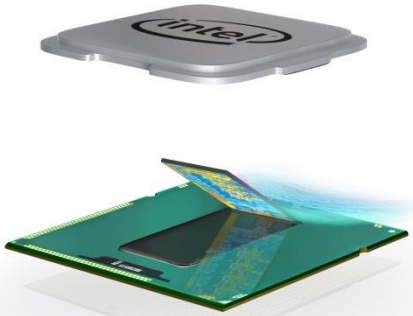

Packaging –

scale: package level (~20mm / ~1 inch)

The package substrate, the die and the heat spreader are put together to form a completed processor. The green substrate builds the electrical and mechanical interface for the processor to interact with the rest of the PC system. The silver heat spreader is a thermal interface which helps dissipate heat.

Processor –

scale: package level (~20mm / ~1 inch)

Completed processor (Ivy Bridge in this case). A microprocessor has been called the most complex manufactured product made by man. In fact, it takes hundreds of steps – only the most important ones have been included in this picture story - in the world's cleanest environment (a microprocessor fab).

Class Testing / Completed Processor

Class Testing –

scale: package level (~20mm / ~1 inch)

During this final test the processor is thoroughly tested for functionality, performance and power.

Binning –

scale: package level (~20mm / ~1 inch)

Based on the test result of class testing, processors with equal capabilities are binned together in trays, ready for shipment to customers.

Retail Package –

scale: package level (~20mm / ~1 inch)

The readily manufactured and tested processors either go to system manufacturers in trays (see binning image) or into retail stores in a box such as the one shown here.

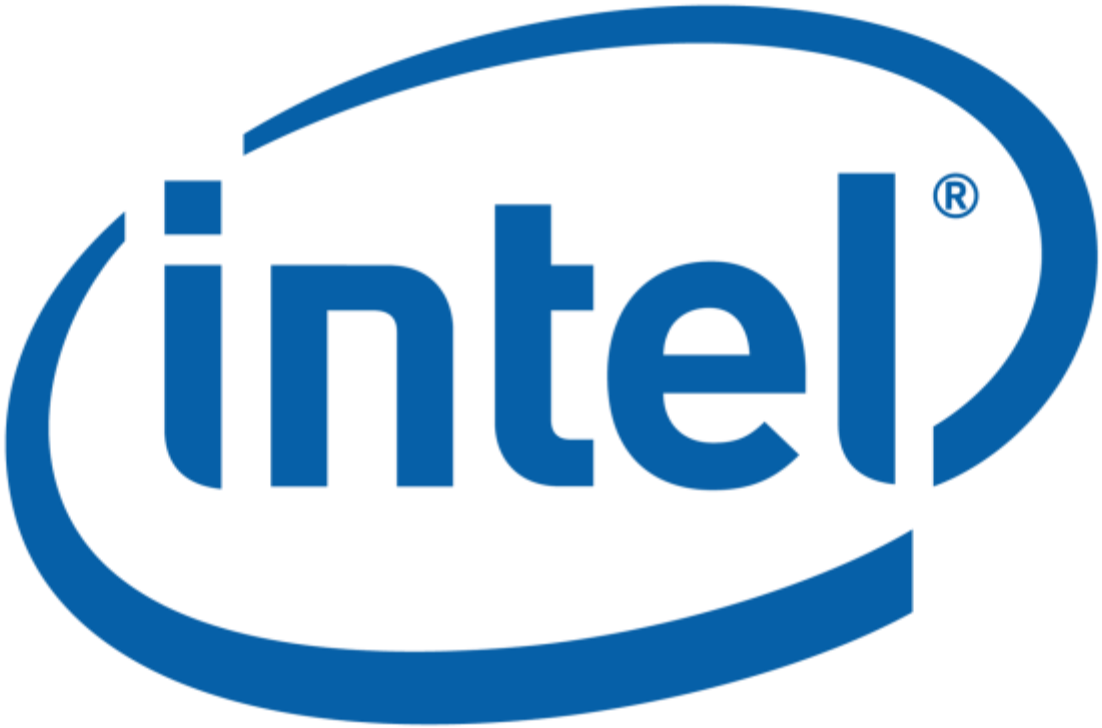

Wafer Sort / Singulation

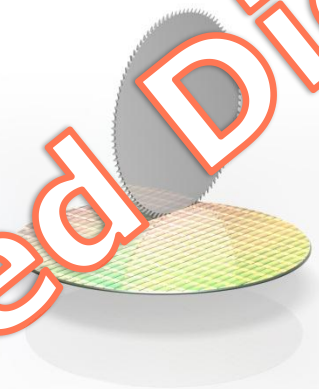

Wafer Sort –

scale: die level (~10mm / ~0.5 in)

This portion of a ready wafer is being put through a test. A tester stepper across the wafer; leads from its head make contact with each of the dies on the top of the wafer and an electrical test is performed. Test patterns are fed into every single chip and the response from the chip is monitored and compared to "the right answer".

Wafer Slicing –

scale: wafer level (~300mm / 12 inch)

The wafer is cut into pieces (called die). The above wafer contains future Intel processors codenamed Ivy Bridge.

Selecting Die for Packaging –

scale: wafer level (~300mm / 12 inch)

The die that responded with the right answer to the test patterns will be packaged.

Packaging

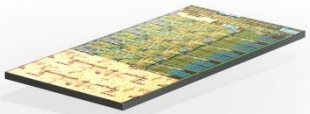

Individual Die –

scale: die level (~10mm / ~0.5 inch)

These are individual die which have been cut out in the previous step. The die shown here is Intel's [illegible] microprocessor codename Ivy Bridge.

Packaging –

scale: package level (~20mm / ~1 inch)

The package substrate, the die and the heat spreader are put together to form a completed processor. The green substrate builds the electrical and mechanical interface for the processor to interact with the rest of the PC system. The silver heat spreader is a thermal interface which helps dissipate heat.

Processor –

scale: package level (~20mm / ~1 inch)

Completed processor (Ivy Bridge in this case). A microprocessor has been called the most complex manufactured product made by man. In fact, it takes hundreds of steps – only the most important ones have been included in this picture story - in the world's cleanest environment (a microprocessor fab).