April 10, 2018

Massive On-Die Capacitor for Integrated Circuits

Woonseong Kim

Houle Gan

Teckgyu Kang

Massive On-Die Capacitor for Integrated Circuits

Abstract

This document proposes a solution to the problem of fluctuations in voltage supply voltages due to changing current demands of power-hungry integrated circuits (ICs) such as microprocessors, application-specific integrated circuits (ASICs), and field-programmable gate arrays (FPGAs). By fabricating the IC with additional layers including a massive capacitor, such as a deep-trench capacitor, di/dt voltage ripple can be reduced to a negligible level. Two configurations of massive on-die capacitors are proposed, including a wafer-to-wafer configuration, where a capacitor wafer substantially underlies the entire IC, and a chip-to-wafer configuration, where smaller capacitors are placed underneath regions of the IC with the highest current fluctuations, surrounded by dielectric.

Introduction

IC supply voltage fluctuations have emerged as a serious cause for concern in high performance processor design. These perturbations, sometimes known as “ground bounce”, occur when the IC demands rapid changes in load current over a relatively small time scale. Since the power delivery system has substantial parasitic inductance, this current variation produces voltage ripples on the IC’s supply lines. This is significant because if the supply voltage rises or drops below a specific tolerance range, the CPU may malfunction.

In some applications, the change in current draw (di) of an IC between quiescent operation and heavy load or max load can be 180 amps, 370 amps, 500 amps or more. Meanwhile, the interval (dt) over which this change in current draw occurs is getting shorter; for example, 700 MHz, 1,110 MHz, 1,500 MHz, and beyond. Thus, di/dt is increasing explosively with the advance of high-speed, high-power ICs.

This fundamental challenge is known as the di/dt problem, because the magnitude of the voltage ripples is affected by the change of current with respect to time. The di/dt problem can limit the performance scaling for power-hungry ICs; for example, machine learning ASICs.

Related Technology

Current technology uses capacitors at various levels within the die, package, and PCB. One such system uses metal-oxide semiconductors in the die. Another such system uses metal-insulator-metal capacitors in the die. These configurations can suffer from timing issues in the die. For example, on-chip variation is caused by voltage variation within the IC. The variations can be caused by resistance voltage drops (IR drops) or by di/dt. The voltage variations can lead to higher bit error rates. While these issues could be addressed by adding area—either in the form of additional capacitance, buffering, or redriving—IC real estate is limited.

Another way to mitigate di/dt is to fabricate a deep-trench capacitor in the IC die itself; however, integrating a deep-trench process with standard IC fabrication processes is difficult.

Solution

A massive on-die capacitor is able to mitigate the fundamental di/dt issue caused by faster switching and higher current draws. On-die capacitors, however, are limited in size due to IC real estate constraints. Therefore, adding a separate on-die capacitor layer onto the IC circuitry can maximize the on-die capacitor size.

This document proposes a solution to the problem of fluctuations in voltage supply voltages due to changing current demands of power-hungry integrated circuits (ICs) such as microprocessors, application-specific integrated circuits (ASICs), and field-programmable gate arrays (FPGAs). By combining the IC wafer with a massive capacitor wafer, such as a deep-trench capacitor, di/dt voltage ripple can be reduced to a negligible level. Two configurations of massive on-die capacitors are proposed, including a wafer-to-wafer configuration, where a capacitor wafer substantially underlies the entire IC, and a chip-to-wafer configuration, where smaller capacitors are placed underneath regions of the IC with the highest current fluctuations.

Construction

Figure 1 below shows four stages of fabrication of a massive on-die capacitor for ICs. In the first stage, the capacitor wafer and the IC wafer (in this case, an ASIC wafer), each with exposed bond pad copper (orange blocks), are stacked together. The light blue represents the capacitor wafer dielectric layer, and the dark blue represents the IC wafer dielectric layer. In the second stage, the two wafers are bonded together in a manner that forms a robust conductive path across the respective copper pads, and forms a secure mechanical bond. Bonding can be performed via a hybrid bonding process, an oxide bonding process, a copper bonding process, or some combination thereof.

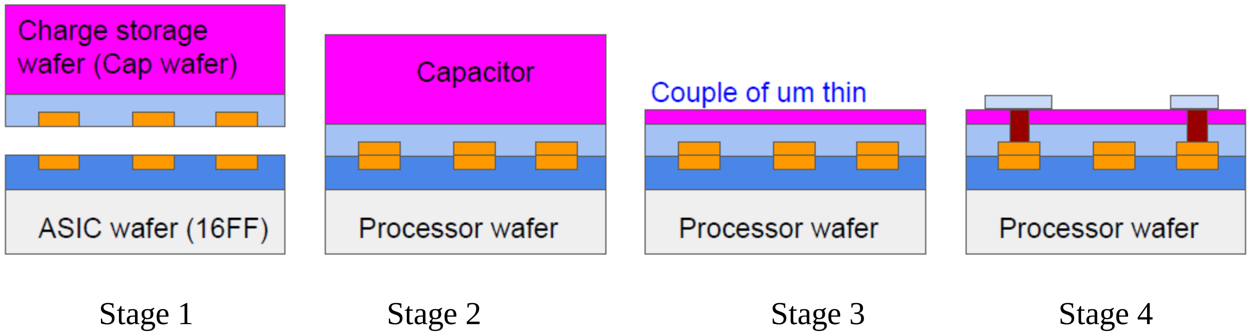

Figure 1: Example fabrication steps of a massive on-die capacitor for ICs

In the third stage, the capacitor wafer is planed down in a bulk silicon removal process. In some cases, the silicon can be removed in a process of H doping plus selective silicon etching in a process of chemical mechanical polishing. The resulting capacitor will be several microns thick; for example, 1 to 100 microns thick. In some cases, the capacitor can be between 1 and 10 microns thick. In some cases, the capacitor can be between 4 and 6 microns thick. In some cases, the capacitor can be 5 microns thick. In the fourth stage, silicon vias can be formed through the capacitor wafer and dielectric to the copper pads. The vias can be filled with copper. External copper pads can be deposited on the outer surface such that they electrically couple to the vias.

The capacitor wafer can be a deep-trench capacitor. The capacitor wafer can include a high-K dielectric such as halnium oxide fill. The top and bottom surfaces of the capacitor wafer can form the capacitor electrodes; however, the internal structure of the capacitor wafer can include a pattern of deep and narrow trenches of first conductors separated by thin dielectric layers surrounding fins of second conductors. The capacitor wafer can ultimately provide a capacitance of between 400 and 1,500 nanofarads/mm^2. Effectively applying this capacitance to decouple the power supply voltages requires a dense pitch of vias. For example, the vias and copper pads may be fabricated at a pitch of 5 or 6 microns. The vias connect the copper pads to the power supply mesh layers in the IC wafer (the darker blue regions in Figures 1 and 2A and 2B).

Figures 2A and 2B show examples of massive on-die capacitor configurations. Figure 2A shows an example wafer-to-wafer configuration, where a capacitor wafer substantially underlies the entire IC. This arrangement can provide for maximum decoupling from the power supply and highest protection against di/dt voltage variation.

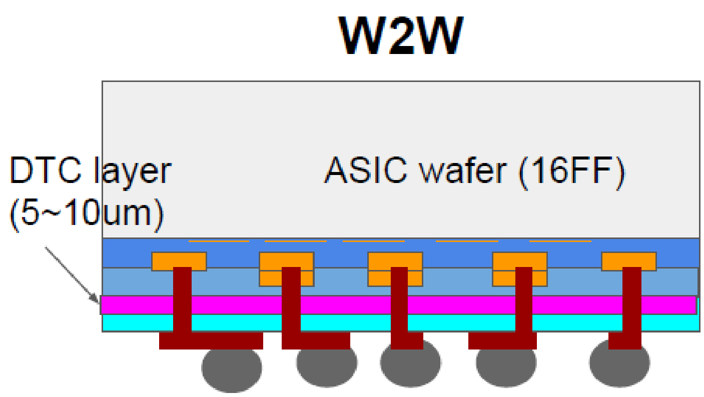

Figure 2A: Example wafer-to-wafer on-die capacitor configuration

Figure 2B shows a chip-to-wafer configuration, where smaller capacitors are placed underneath regions of the IC with the highest current fluctuations. The individual capacitor or capacitors can be surrounded by dielectric. For example, the capacitors can be arranged underneath a vector processing unit so as to provide for a relatively short path to the VPU in order to minimize resistive voltage drop and inductive reactance to current changes. The chip-to-wafer configuration can be advantageous because the capacitor wafer may be most helpful in the immediate vicinity of the most power-hungry regions of the IC.

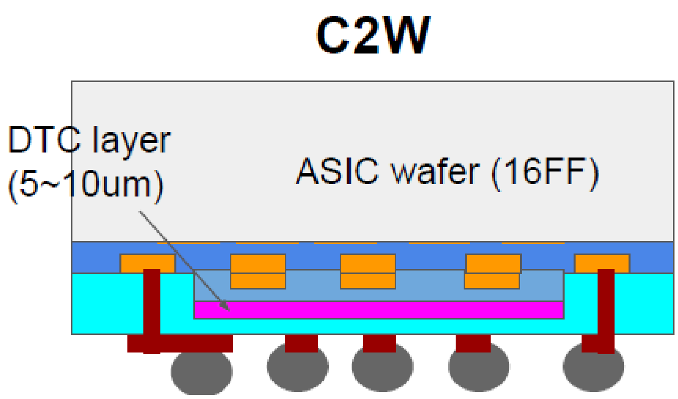

Figure 2B: Example chip-to-wafer on-die capacitor configuration

Results

The massive on-die capacitor configurations described above can be effective in reducing voltage variation from a di/dt of up to and over 7 amps per nanosecond; however, higher values of di/dt may be effectively addressed with this technology.