

The use of Area I/O or a Look on Future Architectures

Etienne Hirt, Michael Scheffler, Gerhard Tröster
Electronics Laboratory, ETH Zürich
CH-8092 Zürich, Switzerland

Abstract

To date designers seek to achieve ever smaller systems with ever more functionality, but more and more they face the interconnection technology as a show stopper. To overcome this bottleneck we propose a chip-package codesign approach; a close cooperation between chip and package designers exploiting the synergism. Our approach distributes the on-chip pads all over the IC area near the pads associated core area. This technique results into smaller ICs with more and faster I/Os being much easier to package. In this paper, a case study for a Pentium class system shows why other approaches such as wire bond, re-routing and chip size package (CSP) have shortcomings. Finally, we present an outlook to new system architectures that are enabled by area I/O: A processor system with first level cache on separate ICs instead of being integrated on the CPU itself.

Key words: Chip-Package Codesign, Area I/O, CSPs, Multichip Modules

Introduction

“The next performance increase in microprocessor systems is waiting just behind the corner”. This statement seems to be proven every year when a new generation of microprocessors is launched. Although it cannot be denied that progress is made primarily at chip level while it advances much slower on system level. For example, processors show ever increasing internal clock rates. On the other hand, external clock rates and memory bus width have not followed this development, as shown in table 1.

Even if several levels of cache hierarchies are added in order to overcome this discrepancy between off-chip bandwidth and on-chip speed, the maximum latency in future generations will increase considerably [1]. Today, even the I/O busses proposed in the NTRS 1997 road map [2] are difficult to connect to the outside world.

Thus, performance figures such as bandwidth, latency, system speed and also size of future microprocessor systems are highly dependent on the interconnection technologies. In fact, when the expected feature size reduction due to semiconductor technology improvements are taken into account, interconnection will be THE performance limitation.

In this paper we present an area I/O chip-package codesign approach featuring ICs designed for area connection instead of the actual

peripheral one. Using this technique, smaller ICs with more and faster interconnections are designed. The benefits of this approach are illustrated with a case study on the basis of a Pentium¹ class system. Finally, an outlook to a new processor system architecture is presented.

Real Area I/O: A Codesign Approach

A decade ago, designers made the first attempt to improve off-chip interconnectivity with multichip modules (MCMs) [3]. For those MCMs, wire bonding was used to connect the bare dies to the substrate to shorten signal path and reduce system size. Soon it became obvious that wire bond interconnect is not the final choice for wide and fast I/O bus structures and components having large numbers of I/Os. It is a time consuming, high precision manufacturing process to bond thousands of wires with small pitches to the substrate. In addition, the high inductivity of wire bonds degrades the overall signal speed.

Looking at today's packages, more appropriate interconnect technologies redistribute peripheral I/O over the entire back side of the package. The use of solder balls instead of leads relieves pitch constraints. So, from the interconnect/packaging point of view, area I/O is already

¹Pentium is a registered trademark of Intel Corporation

Table 1: NTRS Road Map for Performance of Microprocessors [2]

Year	1997	1999	2001	2003	2006	2009	2012
Technology Generation	250 nm	180 nm	150 nm	130 nm	100 nm	70 nm	50 nm
On-Chip Frequency [MHz]	350	526	727	928	1108	1468	1827
<i>Relative to 1997</i>	100%	150%	200%	270%	320%	420%	520%
Off-Chip Frequency [MHz]	75	100	100	125	125	150	150
<i>Relative to 1997</i>	100%	133%	133%	166%	166%	200%	200%
I/O Bus width	64	64	128	128	128	256	256
Chip Pad Count	800	976	1193	1458	1968	2656	3587

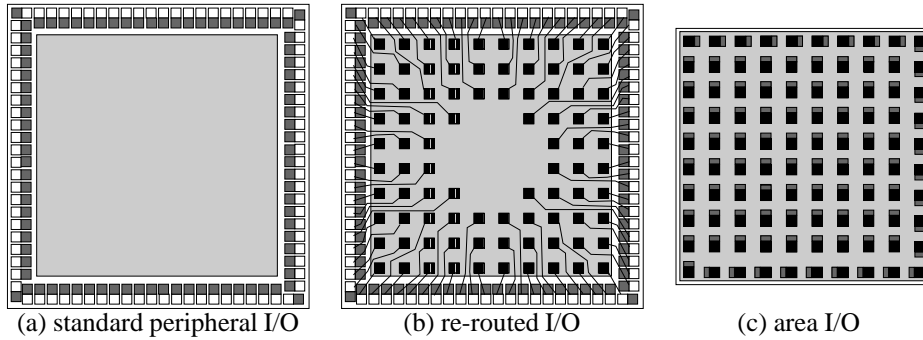


Figure 1: Pad Placement: white cells represent chip pads, black cells package pads, grey cells show the pin electronic

the method of the future. But instead of extending this concept to the chip level, designers typically place the chip's pads peripherally (figure 1a) because:

- wire bonding is standard,
- the IC core design is much easier
- and it is preferable to place ESD protection in pad ring.

Guided by the knowledge that traditional chip-to-chip interconnections are slower and more inhomogeneous than on-chip interconnections (about factor ten), designers favor a single chip solution designing ever larger ICs. These ICs are difficult to package as the pad pitch is decreasing and the number of I/Os as well as its speed is increasing.

Much better designs can be found by exploiting the synergy of ICs and packaging through their concurrent and matched design to meet system-level objectives. This chip-package codesign approach [4] features distribution of the functionality between IC and package. To do so, we propose a distribution of the on-chip pads all over the IC area.

The first approaches to connect the ICs on an area connection were re-routing the peripheral pads to an area I/O arrangement (figure 1b). This method is available for existing ICs but it

needs additional steps on wafer level to form this redistribution. Furthermore, this redistribution increases significantly the propagation delay as compared in table 2. Thus, we propose real on-chip area I/Os (figure 1c). With this method, the pin electronic (buffer, receiver and optionally ESD protection) and the I/O pad are placed near the associated core area. This approach allows to place more I/Os at a larger pitch on a smaller IC. The benefits are:

- Flip chip is used instead of wire bonding
- smaller, faster and cheaper than CSPs
- uses same assembly technology as SMT

Chip-to-Chip Speed Improvements

To compare the speed of the different first-level interconnect, a point-to-point connection between two adjacent ICs of 8 mm side length is modeled. The driver is modeled as a voltage source having an resistance of 21Ω and a rise time of 55 ps (0 to 2V). It drives one load. The simulation results comparing wire bonding, area I/O and thin film (TF) and IC metal layer re-routing are shown in figure 4. For re-routing a short (2 mm) and long (4mm) redistribution

Table 2: Simulation Parameters and Results

	Wire bonding	Area IO	Re-routing IC short	Re-routing TF short	Re-routing IC long	Re-routing TF long
Chip size [mm ²]	8 * 8	7 * 7	8 * 8	8 * 8	8 * 8	8 * 8
Max. substrate path length [mm]	10	8	9	9	9	9
Redistribution length [mm]	n/a	n/a	2	2	4	4
Propagation Delay at 2V [ps]	265	165	245	195	365	220
Skew at switching level(2V) [ps]	100	0	80	30	200	55
Rise time (0.2 to 2V) [ps]	115	85	115	85	195	90

Table 3: Interconnection Parameters

Parameter	Line width [μm]	Line thickness [μm]	Relative dielectric constant (ϵ_r)	Dielectric thickness [μm]
Thin Film	20	3	2.65	7
IC	3	1	3.9	1
PCB	125	30	4.7	200

length was simulated. As summarized in table 2, area I/O has the lowest propagation delay and rise time. Wire bonding or short re-routing on either a IC metal layer (curve FCreceiver11 of figure 4) or a thin film layer (curve FCreceiver12) can be feasible for high speed interconnection. But long re-routing on a metal layer (labeled FCreceiver12) delays the system with up to 200ps compared to area I/O. The simulation parameters as well as detailed results can be found in tables 2 and 3.

System Improvements

A case study for a Pentium class system (figure 3) exploits the use of area I/O on-chip versus wafer level re-routing (figure 1b) and CSPs for future microprocessor systems. This system consists of the CPU, the second level Cache (4 Pipelined Burst SRAMS 'PBSRAM', 1 Asynchronous 'ASRAM') and the system controller split into two data path chips 'MTDP' and one controller 'MTSC', thus being the computing core. Additional components are the main memory and any peripherals connected to the peripheral bus (PCI).

Comparing a wire bonded MCM-D solution[5] using off-the-shelf components to an area I/O solution, a significant performance gain as well as a size reduction on system and chip level for area I/O can be found.

Size Comparison

The Area I/O system (figure 2(b)) is 40% smaller than the wire bond version and only a fifth of a

conventional PCB implementation. Figure 2(a) shows the wire bonded version. The white rectangles show the chip size, and the grey border marks the overhead needed for interconnection. It can be seen that the CPU as well as the chip set (marked as MTDP, MTSC) have a larger overhead than the memories as they need two rows of wire bonding. The size reduction for area I/O is partly due to the smaller overhead for flip chip attach saving about 20%. Additionally, for area I/O the silicon area is reduced by 10% for the CPU and 30% for the chip set because area consuming routing to the padding and the padding itself can be skipped. The core area is only slightly enlarged as active area can be placed under area pads and big power rails are removed.

A CSP implementation is 20% larger than the wire bonded reference. Whereas the memory dies can be packaged on the same size as their die, the CPU as well as the controller packages are much larger. The overhead from the chip area to the CSP area is mainly caused by the first level interconnect (wire bonding from die to CSP interposer).

A close look at the First Level Interconnect

The CPU with its 75 μm peripheral pad pitch as well as the system controller with its two pad rows (resulting pitch 60 μm) are very difficult to assemble and package. The interconnection overhead is so large that CSPs can be built with 0.8mm ball pitch and still the size is defined by the on-CSP wire bonding. Area I/O or re-routing on the other side relieves the CPU pitch to more

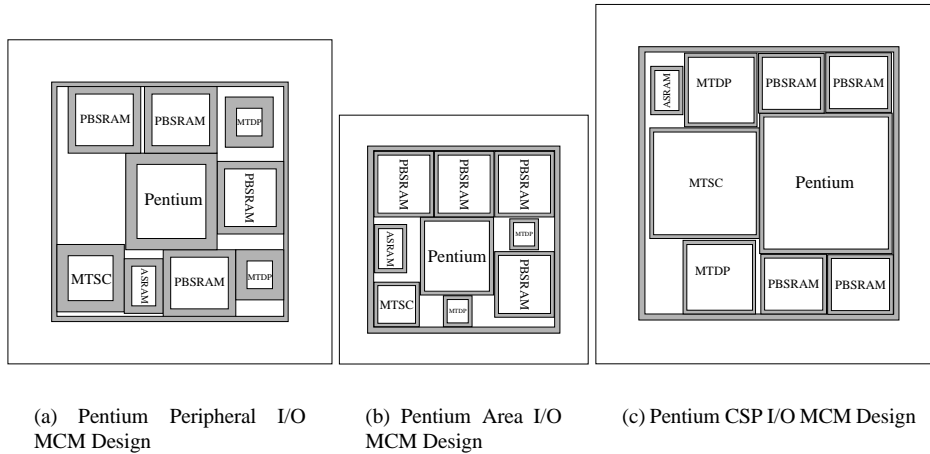


Figure 2: Size Comparison of Different Configurations (80% to scale): White rectangulars represent IC or CSP size, grey borders show the footprint and the substrate pad ring and the white frames show the module package size

than $350 \mu m$. The dies could be packaged into CSPs with $0.5mm$ pitch using flip chip (FC) on a very simple interposer as no redistribution is needed. Thus, the assembly cost is reduced.

Speed Comparison

Extending the chip-to-chip skew simulations a interconnection speed comparison was done for the different implementations in order to show the impact for this system. The simulation covers the address bus interconnecting the CPU with the system controller (TSC) and the four cache ICs (PBSRAM).

As the summary in table 4 shows, the full PCB system has a very large propagation delay ($4.2 ns$). To date, busses and drivers are designed to accomplish this skew. But, for faster interconnections as needed for high speed applications this is no longer acceptable. The area I/O configuration shows the best performance figures as the propagation delay is smaller than $1 ns$. Re-routing introduces a small speed penalty but other alternatives such as wire bonding or CSP introduce too much skew (up to $500ps$). Thus, in terms of speed area I/O is the best solution, and re-routing can be a viable alternative.

Cost comparison

A cost comparison of the presented implementation alternatives was done with the MOE tool (Modular Optimization Environment [6]). It features a process oriented cost representation and

includes direct cost, non-recurring expenditure (NRE), test and yield.

The detailed comparison showed that an optimized system, featuring less than fifth the size of the PCB implementation, causes slightly higher cost. But, it is less expensive than the wire bonded MCM solution. For Area I/O the direct cost are lower due to smaller ICs and the assembly yield is better because of the much larger pitch. Further details can be found in [7].

Outlook

Consequent use of area I/O in combination with high density interconnects opens the door towards novel system partitioning between package and chip. The system shown in figure 5 is a proposal: The first level cache is accessed at CPU speed even when it is not on the CPU die. Thus, the caches can be built with dedicated memory technology and therefore need only half the silicon area. The CPU die size decreases and can accommodate several main memory interfaces. So, the possibility to have more I/Os is used to improve the bandwidth from main memory by factor four. In this example, all busses (cache, memory and peripheral) are independent and therefore parallel accesses are possible.

With this architecture, designers could freely distribute a systems functionality on an MCM to achieve optimal performance and yield. Thus, by designing systems in a package they directly profit from the chip-package codesign approach.

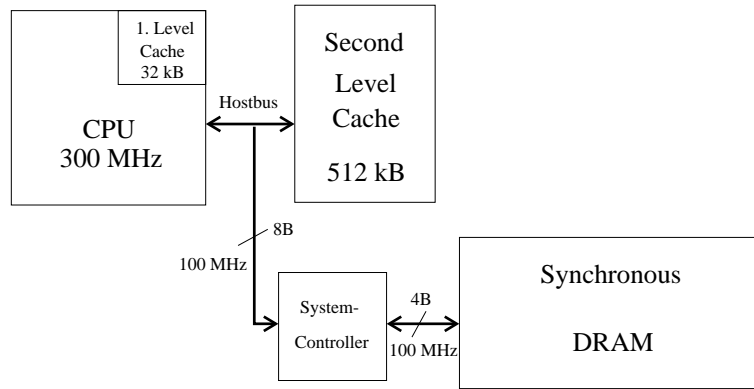


Figure 3: Pentium MCM system architecture underlying all case study implementations.

Table 4: Comparison Results

	Full PCB System	Wire bonded	Re-routing TF	Area I/O	CSP
System Size [cm^2]	81	20	16	12	24
Propagation Delay [ps]	4200	1150	980	870	1390
Skew [ps]	3330	280	110	0	520

Summary

Interconnection is the key to closing the gap between on- and off-chip bus speed. Area I/O in particular promises to meet the increasing pin count and speed as foreseen by the NTRS. Bump bonding and smaller ICs improves signal speed and quality due to shorter interconnect length and lower parasitics. The relaxed pitch also improves manufacturability as well as reliability. The CSP community can also profit from Area I/O as the interposer can be kept very simple and in the presented examples even smaller. Finally, the speed degradation introduced by CSPs becomes marginal. But meeting the challenging predictions of the NTRS road map necessitates a closer cooperation between chip and package designers. Then, we can be assured that a faster microprocessor system is indeed "just around the corner."

References

- [1] D. Patterson, T. Anderson, N. Cardwell, R. Fromm, K. Keeton, C. Kozyrakis, R. Thomas, and K. Yelick, "A case for intelligent ram," *IEEE Micro*, pp. 34–44, Apr. 1997.
- [2] "The national technology roadmap for semi-conductors," 1997.
- [3] P. D. Franzon, *Multichip Module - Technologies and Alternatives*. New York: Van Nostrand Reinhold, 1993.
- [4] G. Troester, "New Route in System Integration: Chip-Package Codesign," *IEEE Micro*, vol. 18, pp. 7–9, July/August 1998.
- [5] E. Hirt, J.-P. Wyss, C. Habiger, and G. Troester, "A pentium based mcm for embedded computing," in *11th European Microelectronics Conference*, (Venice, Italy), pp. 516–523, International Society for Hybrid Microelectronics, 1997.
- [6] M. Scheffler, D. Ammann, A. Thiel, C. Habiger, and G. Troester, "Moe - a modular optimization environment for concurrent cost reduction," in *Proc. of the 5th European Concurrent Engineering Conference (ECEC)*, (Nuremberg, Germany), pp. 115–119, Apr.26-29 1998.
- [7] E. Hirt, J.-P. Wyss, and M. Scheffler, "Area i/os potential for future processor systems," *IEEE Micro*, pp. 2–9, Aug. 1998.




Figure 4: Point-to-Point Skew Comparison of Area I/O vs On-Chip Re-routing: Curve X10.H1 shows the voltage source, FCRECEIVER11 shows short re-routing on IC metal layer, FCRECEIVER12 is short re-routing on a thin film layer and FCRECEIVER2x show the signal behavior for long re-routing

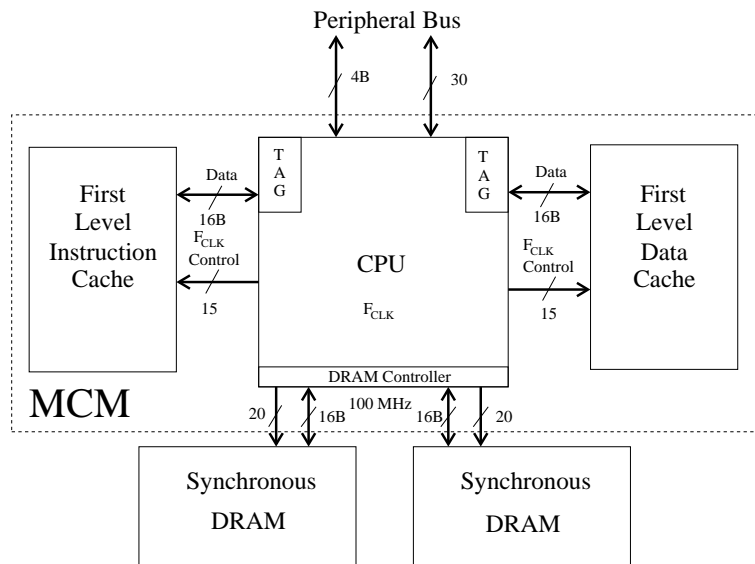


Figure 5: Future System Architecture: Although the first level cache is moved from the processor IC, it is accessible with CPU speed.