

The Potential of Area I/O for Future Processor Systems

Etienne Hirt, Michael Scheffler, Jean-Pierre Wyss
Electronics Laboratory, ETH Zürich
CH-8092 Zürich, Switzerland

1 Introduction

“The next performance increase in microprocessor systems is waiting just behind the corner”. This statement seems to be proven every year when the new generation of microprocessors is launched. Although one cannot deny that progress is made primarily at chip level and that it is much slower on the system level: Whereas state-of-the-art processors already today show internal clock rates being predicted for the future as shown in the NTRS road map (Tab. 1), the external clock rate and especially memory bus width cannot keep up with this development. Even by adding several levels of cache hierarchies to overcome this discrepancy between off-chip bandwidth and on-chip speed, the maximum latency will enlarge more and more [1]. Today, even the I/O busses proposed in the road map are difficult to connect to the outside world.

Thus, performance figures as bandwidth, latency, system speed, and not to forget, size of future microprocessor systems are highly dependent on the interconnection technologies. In fact, when taking into account the expected feature size reduction due to semiconductor technology improvements, interconnection will be the performance show stopper.

Year	1998	2001	2004	2007	2010
	-	-	-	-	-
	2000	2003	2006	2009	2012
On-Chip Frequency [MHz]	200	300	400	500	625
Off-Chip Frequency [MHz]	66	100	100	125	150
I/O Bus width	64	128	128	256	256

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

Evolving from these shortcomings, in the late 80ies multi-chip modules (MCMs) were the first step to improve the off-chip interconnectivity[3]. With this technology, a block of functional components, the so-called “partition”, preferably with a large amount of high speed interconnects, is grouped together on a single substrate, thus forming a “new” component. This substrate is an additional interconnect level between system board and chip, on which unpackaged chips (so-called “bare dies”) can be mounted.

For the first MCMs, wire bonding was used to connect the bare dies to the substrate. But soon it became clear that wire bond interconnect is not the right choice for wide and fast I/O bus structures and components having large numbers of I/Os: It is a time consuming high precision process to bond several hundred wires with small pitches to the substrate. In addition, the high inductivity of wire bonds degrades overall signal speed.

Today, several other interconnect technologies are available on the market, of which CSPs (chip size/scale packages) are foreseen to have the strongest market growth. CSPs

re-distribute the peripheral I/O pads over the entire chip area to relieve the pitch constraints, using bump interconnections instead of wire bonds.

So, from the interconnect/package point of view, area I/O is already the method of the future. But on chip level, almost nothing has changed as the I/O pads of a chip are still on the periphery. Typical objections against area I/O are

- chip area will be bigger as no circuitry can be placed under the pad area,
- too complicated, no tools to do the routing on chip,
- no performance benefit, just adding cost,
- loss of flexibility, as area I/O has to use flip chip (FC) and cannot be wire bonded.

This paper exploits the use of area I/O on-chip for future microprocessor systems. Based on a case study for a Pentium® class system, we will show that with area I/O a significant performance gain can be obtained, together with a size reduction on system and chip level. Consequent use of area I/O in combination with high density interconnects opens the door towards novel system partitioning between package and chip. Hence, in the near future both chip designers and the packaging community can profit from a closer cooperation.

2 Peripheral I/O vs. Area I/O

A standard chip layout with peripheral pads as shown in Fig. 1a is the state-of-the-art used by the industry today. The core area is marked grey, surrounded by the chip I/O pads (white) with their pin electronic (dark grey). This pin electronic includes I/O buffers as well as ESD (electro static discharge) protection circuits.

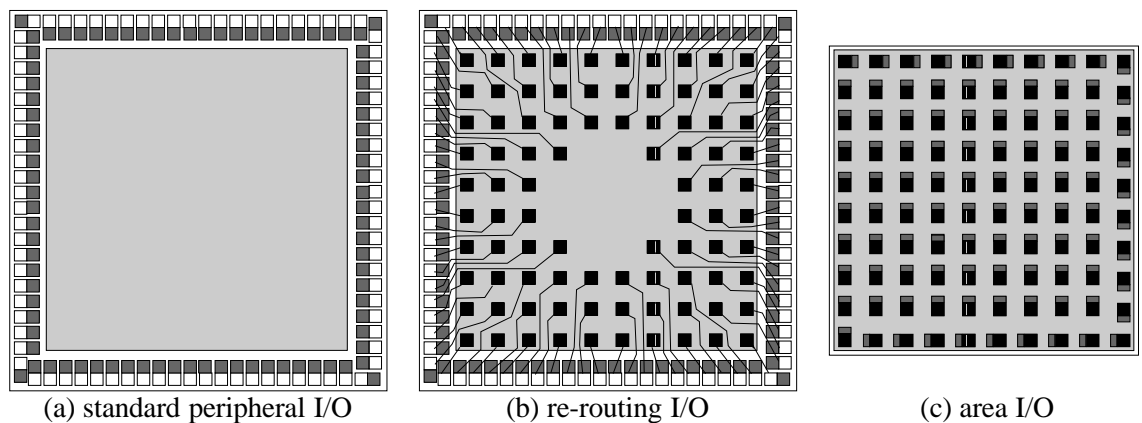


Figure 1: Pad Placement: white cells represent chip pads, black cells package pads, grey cells show the pin electronic

The significant disadvantage of this arrangement is: it shows a low ratio of core area to I/O area for high pin count, pad limited ICs, even at small pitches. Only the chip edges can be used to place the pads. Moreover, wire bonding causes parasitic effects due to its high inductivity. Using flip chip (FC) technology, solder bump connections instead of wire bonds, most of these parasitics can be eliminated. Also, unlike wire bonding, the time to flip chip an IC depends not on the number of I/Os leading to the same assembly time for high and low I/O count ICs. But as the I/O pad pitch of high pin-count ICs could be down

70 μm , this results in high manufacturing constraints making flip chip for these types of ICs almost impossible.

Modern CSPs (chip size/scale packages) on the other hand re-distribute the I/O pads from periphery again to the entire chip area to relieve these constraints (see Fig. 1b, [4]). Chip **SIZE** Packages, having a size of almost the die area, feature additional dielectric and metal layers deposited on wafer level to spread the pads over the surface. However, this additional process is expensive and causes yield loss due to wafer breakage risk.

Chip **SCALE** Packages, having about 1.2 times the size of the die area, connect the die to another substrate which re-distributes the pads. For both types of CSPs, the interconnect length between nearest neighbor chips is enlarged significantly: route from the core to the periphery, and back to center before leaving the chip level.

Area I/O (see Fig. 1c) spares this wiring from core to peripheral pad and back to the center.

Two major problems were encountered in the past when thinking about area I/O: placing the pin electronic near the pads consumes too much area, and no circuitry can be placed under I/O pads.

But for interconnects remaining inside the module and therefore not being exposed to human body touch, ESD restrictions are much lower. Thus, ESD cells could be reduced in size[5]. Regarding the attachment force, bump interconnection induces far less stress to the layers below pads, so circuitry can be placed under the pads. This allows to move all signal/power pads to an area arrangement with the reduced pin electronic near the pads as shown in Fig. 1c. The I/O pads are directly placed where they are connected to the core, reducing on-chip routing as well as the capacitive loads. Due to the distributed power pads, a good power supply is guaranteed. Thus, the advantages of area I/O are

- released pad pitch leading to higher manufacturing yield,
- increased I/O count allowing a better power distribution and wider busses,
- and reduced interconnect length due to smaller die sizes.

With these benefits, in the next step we will assess the performance gains for a Pentium® class system, realized on a high density substrate.

3 Improving the Form Factor: State-of-the-Art MCM Technology

As a benchmark for our case study, we use a Pentium® MCM on a high density substrate designed by ETH, Electronics Lab in 1996 [6].

The MCM technology [3] offers an intermediate layer, the high density substrate, with via and line pitch closely matched to the chip I/Os that have a pad pitch down to 70 μm . Up to 80% of the total number of interconnections remain completely on the substrate leading to a significant reduction of number of I/Os. Compared to a packaged IC, not only the component size itself shrinks due to the use of unpackaged “bare dies”, but also the fan out¹ area is significantly smaller.

Fig. 2 shows the realization of the high performance part of a microprocessor system on a 4-metal MCM-D (deposited) substrate with 20 μm line width, 30 μm spacing and 50 μm via land. The MCM is placed into a 320-pin plastic stud grid array (PSGA), containing a:

¹In the packaging community, fan out is the area overhead needed to mount and to connect a component to the next level, e.g. wire bond area, area needed for underfill, escape routing. Fan out plus component area lead to the overall footprint, the area where no other component can be placed.

- Pentium® 133 MHz processor,
- 512 kB second level cache, PBSRAM (Cache) and SRAM(Tag),
- system controller (MTSC),
- data path controller (MTDP),

thus fully encapsulating the high speed host bus.

Using MCM-D technology, it was possible to reduce the size of this functional block by 75% compared to a PCB (printed circuit board) solution (figure 4) and, at the same time, to improve signal integrity due to shorter lines with better high frequency characteristics.

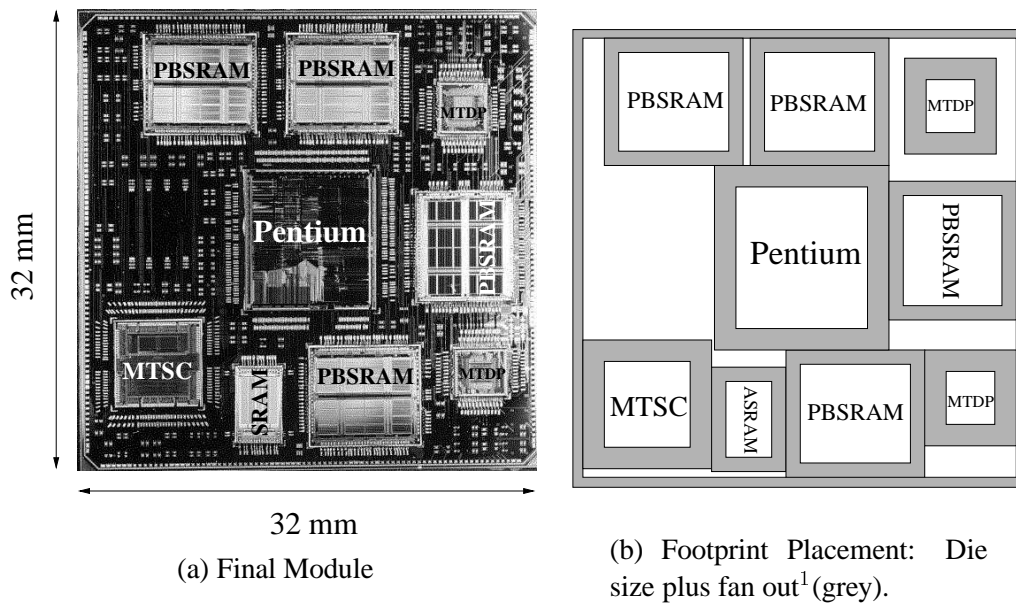


Figure 2: Module Placement

From Fig. 2, it is evident that the module size is limited not only by the component size, but also by the I/O fan out of the components on the module. This fan out overhead is caused only by the wire bond pads, not by additional escape routing, thus providing the smallest module available today.

But even with high density substrates used in today's MCM technology, the number of I/Os cannot be increased as demanded to achieve wider host busses. This could be changed by adapting IC design to the needs of the interconnection technology, using the advantages of area I/O.

4 Case Study: Pentium® Class System with Area I/O

In our case study, we will extend the Pentium® MCM with standard peripheral I/O ("Standard System S") shown to an area I/O system with FC interconnect featuring a wider host bus as proposed in the NTRS road map (see Tab. 1, "System A"). For already available 400MHz CPUs with 100MHz host bus, this leads to a 16Bytes wide host and DRAM data bus. Due to the lower FC interconnect inductivity, the power supply is improved, thus resulting in a signal-to-power pin ratio twice as high. Additionally, a common PCB solution

of System S featuring the same performance level is introduced for size and cost comparison purposes (“System P”).

A summary of the data used for the different configurations is given in Tab. 2 and illustrated in Fig. 3.

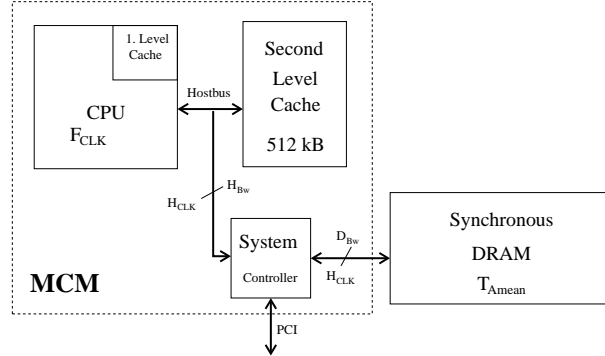


Figure 3: Pentium® MCM System

Table 2: Specifications used for System Comparison (following [2, 7])

		Systems S & P Peripheral I/O MCM/PCB Small Bus	System A Area I/O MCM Wide Bus
CPU speed	F_{CLK}	200 MHz	400 MHz
Host Bus speed	H_{CLK}	66 MHz	100 MHz
DRAM mean cycle time	$T_{A,mean}$	30ns	15 ns
DRAM bus width [Bytes]	D_{Bw}	8	16
Host bus width (Data) [Bytes]	H_{Bw}	8	16
Signal to Power Ratio	SPR	4	8

4.1 Implications on size

On chip level the use of area I/O with the pin electronics near the pads reduces the die size as shown in Fig. 1c. Since for the processor the I/O pads can be placed on top of the core area, the space reserved for the peripheral pads can be saved, leading to a size reduction for this IC of about 10%.

For the system controller (MTSC) even more area can be saved. Due to its two rows of pads on-chip and a lot of wiring needed to escape from the IC core to the pads, chip area can be reduced by about 30% using area I/O. For the data path controller the savings are about the same.

The SRAMs on the other hand do not need to be adopted to area I/O, since a peripheral pad pitch with more than $150 \mu m$ is not critical for flip chip.

On substrate level, the use of flip chip makes systems smaller compared to wire bonding as FC dies only need about 0.5mm fan out instead of 1.5mm for two rows of bond pads. Together with the size reduction for ICs, the MCM substrate for System A shrinks by 44%, compared to Substrate S (see Fig. 4).

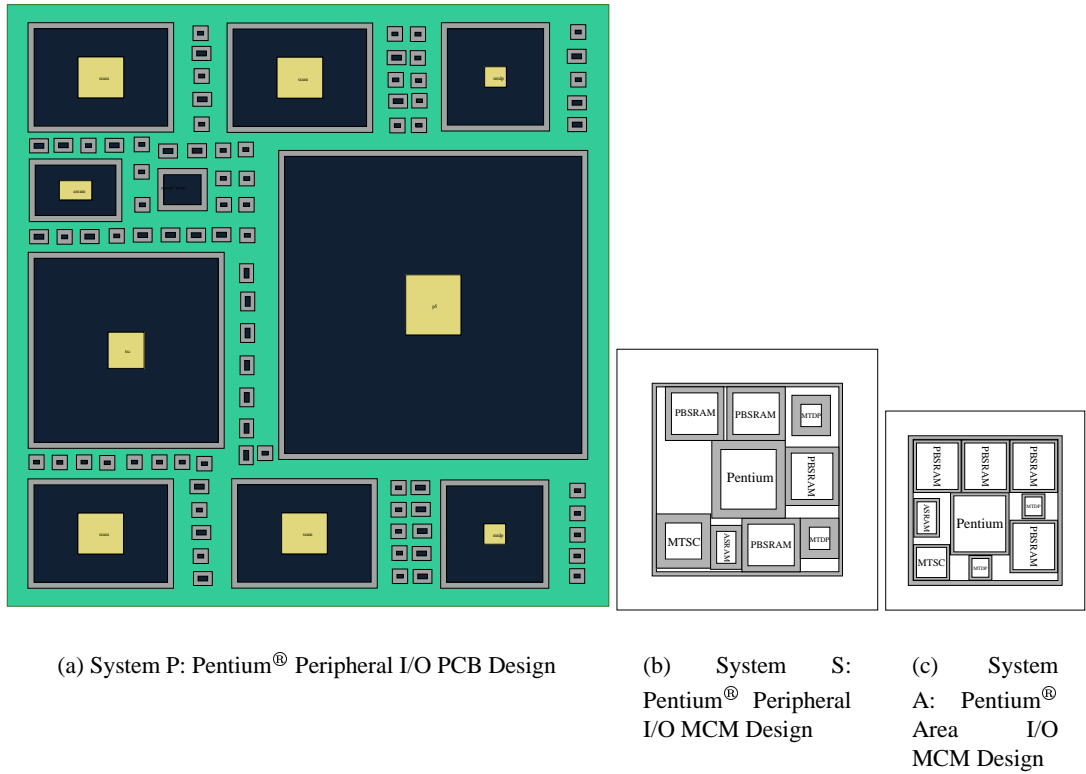


Figure 4: Size Comparison of Different Configurations (80% to scale)

4.2 Implications on performance

As we defined the available bandwidth as the crucial feature, we will use this figures to illustrate the performance gain. Calculations are based on the assumptions in Tab. 3 and the data in Tab. 2. The formulas (1) - (4) are considered to be a sufficiently detailed model to provide a rough estimation of the required and available bandwidth as well as the pin count.

Table 3: Additional Data for all Models, based on Tillamook Design[7, 8]

Bytes/command	B_C	2
CPU cycles/command	C_C	3
Data bytes/command	B_D	4
Host address	H_A	29
General purpose pins	GP	60
Core power pins	C_p	74

The results shown in Tab. 4 illustrate the bandwidth limitation of the standard System S and the improvement by using area I/O. Comparing the available to the needed bandwidth, it can be seen that the Systems P & S are bandwidth limited, whereas for System A preload of data² becomes possible. System A using an 8Byte host bus would result in an available bandwidth of 528MB/s, thus being still bandwidth limited. This shows that only with area I/O the full computing performance available can be used.

²The CPU may load speculatively data from the main memory into the cache.

$$\text{Available Bandwidth from DRAM} = D_{Bw}/T_{A,mean} \quad (1)$$

$$\text{Needed Bandwidth by CPU} = \frac{f_{CPU}}{C_C} * (B_C + B_D) \quad (2)$$

$$\text{Signal Pins} = (H_{Bw} * 10) + H_A + GP \quad (3)$$

$$\text{Power Pins} = \left(\frac{\text{Signal Pins}}{SPR} \right) * 2 + C_p \quad (4)$$

Table 4: Performance Comparison: whereas the peripheral Systems P & S are bandwidth limited, for the area I/O System A speculative cache preload becomes possible.

	System S & P Peripheral I/O MCM/PCB Small Bus	System A Area I/O MCM Wide Bus
Available Bandwidth from DRAM	264 MB/s	1056 MB/s
Needed Bandwidth by CPU	400 MB/s	800 MB/s
Signal Pins CPU	169	249
Power Pins CPU	158	132
Total Pins	327	381
Minimum Pad Pitch CPU calculated	114 μm	455 μm
Minimum Pad Pitch CPU measured	75 μm	300 μm^3

Tab. 4 shows also the pad pitch required for the CPU IC, illustrating the manufacturing constraints. System A realized with peripheral I/O pads would result in a CPU pad pitch of 85 μm (calculated) that might be going below 60 μm ³. Whereas the peripheral I/O MCM with a small bus still could be wire bonded, this is very difficult for the peripheral configuration with the wider bus, and FC mounting is even impossible. The area I/O technology used in System A, on the other hand, allows a very comfortable pad pitch.

4.3 Implications on cost

In order to make a cost comparison of the three different solutions, also the area consumed on PCB motherboard has to be considered. Moreover, all the cost data are referring to a high volume production per year (> 1'000'000). So, changes in NRE due to different design methodologies do not have to be considered. The assumptions for die/chip cost are based on two considerations. First, several KGD (known good die) programs offer the same cost and the same test level for packaged as well as for bare dies. Secondly, in the end the die cost is driven only by the die size: decreasing the die size will reduce also die cost (see Tab. 5, System A).

On chip level, no additional metal layer is needed for area I/O, as pads for bump interconnection can be placed over active area. Also, routing density is much lower on the two IC top metal layers because pin electronic is placed close to the I/O pads and the "output" of the core. Moreover, connecting the power locally, we can remove the big power rails.

On substrate level, the size reduction of the MCM substrate by 44% for System A requires an additional metal layer to provide sufficient wiring space.

On PCB main board level, again the area could be reduced by 20%.

The cost calculations have been carried out with the cost modeling tool MOE (Modular Optimization Environment[9]).

³The ratio between calculated and measured pitch for the standard System S, peripheral I/O MCM system with small bus, can be expected to be similar for the other configurations.

Table 5: Cost implications of different solutions: Aiming for high volume production, NRE cost is not included

		System P full PCB system small bus	System S Peripheral I/O MCM small bus	System A Area I/O MCM wide bus
System PCB	size	81 cm^2	20 cm^2	16 cm^2
	metal layers	6	4	4
	cost	\$ 0.12/ cm^2	\$ 0.1/ cm^2	\$ 0.1/ cm^2
MCM substrate	type	n/a	4-layer MCM-D	5-layer MCM-D
	cost		\$ 2/ cm^2	\$ 3/ cm^2
	size		10.2 cm^2	6.76 cm^2
MCM package		n/a	\$ 5	\$ 5
Die Cost	Processor	\$ 150	\$ 150	\$ 135
	System Controller	\$ 21	\$ 21	\$ 14
	Data Path Controller	2 * \$ 7	2 * \$ 7	2 * \$ 4
	Burst SRAM	4 * \$ 20	4 * \$ 20	4 * \$ 20
Number of I/Os	sum of all dies	1200	1200	1420
Assembly Cost		\$ 0.05 per pin	\$ 0.05 per bond	\$ 0.05 per bump
Yield Dies		0.999 all dies	all KGDs (0.999) except SRAM (0.99)	all KGDs (0.999) except SRAM (0.99)
Yield Substrate		n/a	0.99	0.99
Yield Assembly		0.999 per chip	0.9999 per bond wire	0.99 per FC attach

MOE features a process oriented cost structure representation, so the yield of the different manufacturing configurations can be taken into account easily. Inspired by Monte Carlo simulation, MOE calculates the cost for a virtual process lines, including direct cost, NRE, test and yield aspects. In this virtual process line, a functional test is assumed before the system is shipped, so the units containing errors are sorted out. This yield loss has to be added to the direct cost of every shipped system. Detailed results can be found in Tab. 6.

Table 6: Cost Modeling Results

	PCB system	System S	System A
direct cost	\$ 281	\$ 299	\$ 267
yield loss to be added per unit	\$ 6	\$ 63	\$ 39
overall cost	\$ 283	\$ 362	\$ 306

Whereas the cost penalty for the MCM system compared to the PCB system is caused mainly by the high MCM substrate cost and the yield loss due to the wire bonding, the direct cost of the area I/O solution are lowest of all. Additionally, the higher substrate cost is expected to decrease due to large area panel production of MCM-D substrate [10]. The flip chip assembly of the components is almost as reliable as standard soldering technology.

5 Outlook to Future System Architectures

The previous section showed that area I/O can already improve system performance as well as system size based on state-of-the-art architectures. But also different bus architectures as Intel®'s Dual Independent Bus and the AGP bus (Advanced Graphics Port) as well as completely new approaches [5, 11] can benefit from higher I/O count.

Therefore, area I/O can not only be used to make existing ICs smaller but it opens up complete new partitioning options. Functional blocks can be moved on/off chip as shown in Fig. 5 [12]. Using the much more efficient technology of a SRAM process instead of integrating SRAM on a logic process, the area used is more than two times smaller. Therefore, the first level cache can be moved off-chip which improves CPU die yield without performance loss. Additionally, the DRAM controller is implemented on the CPU chip ensuring low latency and allowing for two independent DRAM banks. With this architecture, the functionality of a CPU could be distributed freely on an MCM to achieve optimal performance and yield.

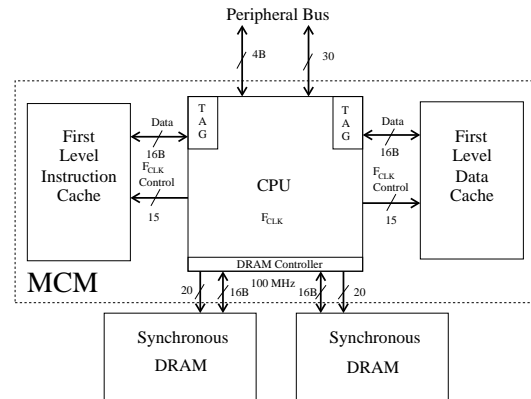


Figure 5: Future System Architecture: Although the first level cache is moved from the processor IC, it is accessible with CPU speed.

6 Conclusion

Interconnection is the important factor responsible for closing the gap between the diverging trend of on-chip and off-chip bus speed.

To overcome this divergence, area I/O interconnect is a promising technology. Major benefits of this technology are an increase in bandwidth due to higher I/O count and a relief of manufacturing constraints on the substrate due to larger pad pitch without additional cost. Moreover, area I/O profits from the consequent use of flip chip technology leading to lower interconnect parasitics.

As shown in the case study for a processor system,

- the ICs cost can be reduced due to their reduced size,
- the overall system PCB size is again cut by 20% compared to a peripheral MCM system,
- the available bandwidth exceeds the bandwidth required by the CPU, removing today's limitations,
- the system cost currently has a penalty of about 10 % compared to a simple PCB solution, caused by the high density substrate cost.

These results show that, based on chip level area I/O, the matched design of chip, package, and system level is a promising methodology to ensure future progress in computing performance. But in order to meet this challenge, a closer cooperation among chip-, package- and system designers is mandatory.

Acknowledgement

The authors wish to thank their colleagues D. Ammann, A. Thiel and C. Habiger for the valuable discussions and helpful comments.

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 semiconductors," 1997.
- [3] P. D. Franzon, *Multichip Module - Technologies and Alternatives*. New York: Van Nostrand Reinhold, 1993.
- [4] R. Fillion, R. Wojnarowski, R. Saia, and D. Kuk, "Demonstration of a chip scale chip-on-flex technology," in *International Conference on Multichip Modules*, (Denver, Colorado, USA), pp. 351–356, Apr.17-19 1996.
- [5] Q. Zhu and S. Tam, "Package clock distribution design optimization for high-speed and low-power vlsis," in *IEEE Transactions on CPMT/Part B*, Feb. 1997.
- [6] 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.
- [7] Intel Corporation, *Mobile Pentium[®] Processor with MMXTM Technology on .25 Micron*, 1997.
- [8] Intel Corporation, *Pentium[®] Processor Family Developer's Manual*, 1995.
- [9] 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.
- [10] N. Ammann, "Lap-the key to low cost multichip packaging," *Future Circuits International*, vol. 2, pp. 25–28, 1997.
- [11] L. Schaper, "Seamless high off-chip connectivity (shocc)," in *IEEE International Workshop on Chip Package Co-Design*, (Zurich, Switzerland), pp. 39–45, Mar.24-26 1998.
- [12] E. Hirt, J.-P. Wyss, M. Scheffler, D. Ammann, A. Thiel, C. Habiger, and G. Troester, "On the impact of area i/o on partitioning: A new perspective," in *IEEE International Workshop on Chip Package Co-Design*, (Zurich, Switzerland), pp. 33–38, Mar.24-26 1998.