10.7 A 105GOPS 36mm² Heterogeneous SDR MPSoC with Energy-Aware Dynamic Scheduling and Iterative Detection-Decoding for 4G in 65nm CMOS

Benedikt Noethen, Oliver Arnold, Esther Perez Adeva, Tobias Seifert, Erik Fischer, Steffen Kunze, Emil Matus, Gerhard Fettweis, Holger Eisenreich, Georg Ellguth, Stephan Hartmann, Sebastian Höppner, Stefan Schiefer, Jens-Uwe Schlüßler, Stefan Scholze, Dennis Walter, René Schüffny

Technische Universität Dresden, Dresden, Germany

Modern mobile communication systems must face conflicting design constraints. On the one hand, the expanding variety of transmission modes calls for highly-flexible solutions supporting the ever-growing number and diversity of application requirements. On the other hand, stringent power restrictions (e.g., at femto base stations and terminals) must be considered while satisfying the increasingly-demanding performance requirements. In order to cope with these issues existing SDR platforms, e.g., [1] and [2], propose an MPSoC with a heterogeneous array of processing elements (PEs). MPSoC solutions provide programmability and parallelism for yielding flexibility, processing performance and power efficiency. To schedule the resources and to apply power gating a static approach is employed. In contrast, we present a heterogeneous MPSoC platform (*Tomahawk2*) with runtime scheduling and fine-grained hierarchical power management. This solution can fully adapt to the dynamically varying workload and semi-deterministic behavior in modern concurrent wireless applications. The proposed dynamic scheduler (*CoreManager*, CM) can be implemented either in software on a general-purpose processor or on a dedicated application-specific hardware unit. It is evident that the software approach offers the highest degree of flexibility; however it may become a performance-bottleneck for complex applications. A high-throughput ASIC was presented in [3], but this solution does not allow flexibly upgrading the scheduling algorithms. In this work, these limitations are overcome by implementing the CM on an ASIP.

The *Tomahawk2* MPSoC is composed of 20 heterogeneous cores connected by a hierarchical packet-switched star-mesh NoC, as depicted in Fig. 1. The NoC is clocked at 500MHz and provides a throughput of 80Gb/s per link, partly employing serial high-speed links [4]. This results in a compact top-level floorplan realization (Fig. 7). An ADPLL is attached to each unit, allowing individual adjustment of the clock frequency within the 83-666MHz range. The DDR2 interface connecting 2x128MB global memory at 400MHz provides a data rate of 12.8Gb/s. An FPGA I/O interface delivers 10Gb/s. The application processor (APP) is implemented as Tensilica 570T RISC core with 16kB data and 16kB instruction caches. It executes the applications' control code and sends task scheduling requests to the CM. The CM is based on a Tensilica LX4 core extended with a scheduling-specific instruction set [7], enabling efficient implementation of adaptive power management and dynamic task scheduling (including resource allocation, data-dependency checking and data management). For this purpose, the CM analyzes at runtime the scheduling requests and exploits the results to maximize data locality and to configure the dynamic voltage and frequency scaling (DVFS) performance levels of the PEs according to current system load, priorities and deadlines.

The Duo-PE is comprised of a vector DSP and a RISC core, connected to a shared local memory. This arrangement increases area efficiency and data locality. Each Duo-PE is equipped with a DMA unit, enabling concurrent data prefetching and task execution. The dual nature of these PEs allows high-performance 16-bit fixed-point signal processing on the 4-fold SIMD VDSP, as well as high-precision floating-point computing on the Tensilica LX4 RISC processor. To support fine-granular fast power management, each Duo-PE equips DVFS (Fig. 3). The core domain is connected to one of 3 global V_{DD} rails by a set of PMOS switches. A LUT contains a set of fractional frequency multipliers (N_1, N_2 and N_3) for the ADPLL, associated with the 3 supply levels. The power rails are controlled by an AVS scheme, which tracks temperature variations and finds the minimum V_{DD} guaranteeing error-free operation for the selected performance level. Therefore, each PE contains 3 ring oscillator HPMs replicating the critical timing of the design by its oscillation period T_{HPM}. The central AVS controller adjusts V_{DD} such that the oscillator period equals N times the reference period ($T_{HPM}=N*T_{REF}$). By configuring the HPMs with N_1, N_2 and N_3 multipliers, respectively, the DVFS target frequencies can be emulated. The voltages of all three V_{DD} rails can be hence slowly regulated by the AVS while the core is running at only one DVFS level, with the capability of fast change to another level.

In order to accelerate computationally-intensive SDR baseband algorithms, two programmable application-specific cores are included: specifically, a sphere detection (SD) core and a multi-mode forward error correction (FEC) core for convolutional, Turbo, and LDPC codes [6]. In the particular case of coded systems, the communications performance can be significantly enhanced by means of iterative detection-decoding. For this purpose, the SD architecture presented in [6] has been largely extended by implementing the algorithms proposed in [5], allowing the SD to process a-priori information generated by the FEC.

The Tomahawk2 SDR MPSoC was fabricated in TSMC 65nm LP-CMOS technology (Fig. 7). It integrates 10.2M NAND gate equivalents and occupies 6mm x 6mm = 36mm² (Fig. 6). The MIMO iterative detection-decoding engine occupies 1.68mm², including 93kB of SRAM. Measurement results corresponding to SD and FEC modules can be found in Fig. 4. Each Duo-PE occupies 1.36mm², of which 0.8mm² is contributed by the two dual-port 32kB memories. The RISC core achieves a maximum frequency of 445MHz at 1.2V. For this configuration, 7.1GOPS are delivered for 8 cores. The VDSP reaches a maximum frequency of 500MHz at 1.2V, which yields a performance of 80GOPS for all 8 PEs. Executing a 2048-point complex FFT under these conditions, the VDSP and the RISC consume 98.1mW and 61mW, respectively. Due to its higher throughput for this application, the frequency and voltage of the VDSP can be downscaled, thus doubling the energy efficiency in comparison to the RISC.

The CM occupies 1.36mm², including 64KB memory for data and 32KB for instructions. It reaches a maximum frequency of 445MHz at 1.2V, resulting in scheduling-throughput of 1.1 Mtasks/s with a power dissipation of 69.2mW. To show the advantage of the hardware-accelerated scheduler, different CM implementations are compared in Fig. 2.

The presented heterogeneous SDR MPSoC integrates 8 Duo-PEs and an iterative detection-decoding engine. It features a hierarchical power management (at system- and PE-level) combined with flexible dynamic task scheduling. System-level power management is integrated in a programmable CM-ASIP which outperforms other scheduler implementations by a factor of 7 (ASIC) and 216 (SW) in terms of ATE product (Fig. 2). On PE-level the ultra-fast DVFS follows the dynamically adapting control of the CM, to further increase the energy efficiency. The flexible iterative multi-mode SD-FEC engine improves the area-performance relationship as well as the energy efficiency by a factor of 3 with regard to recent implementations (Fig. 4). This SDR MPSoC doubles the processing power compared to other SDR solutions (Fig. 5), enabling support of future communications standards. For a 4x4 MIMO 3GPP-LTE baseband application scenario the throughput is increased by nearly a factor of 6 at the same power consumption (Fig. 5).

Acknowledgements: This work has been supported by the German Ministry of Education and Research BMBF under grant number 13N10788 (CoolBaseStation), the state of Saxony under grant of ESF 100098198 and the German Research Foundation within cfaed. Furthermore, we would like to thank Synopsys and Tensilica for sponsoring Software and IP.

References

[1] F. Clermidy; et al., "A 477mW NoC-based digital baseband for MIMO 4G SDR," *ISSCC Dig. Tech. Papers,* pp. 278-279, 2010.

[2] D. Ilitzky; et al., "Architecture of the Scalable Communications Core's Network on Chip," *IEEE Micro*, vol.27, no.5, pp. 62-74, 2007.

[3] T. Limberg, et al., "A fully programmable 40 GOPS SDR single chip baseband for LTE/WiMAX terminals", *European Solid-State Circuits Conference*, pp. 466-469, 2008.

[4] D. Walter, et al., "A source-synchronous 90Gb/s capacitively driven serial on-chip link over 6mm in 65nm CMOS," *ISSCC Dig. Tech. Papers*, pp. 180-181, 2012.

[5] E. P. Adeva, et al., "VLSI Architecture for MIMO Soft-Input Soft-Output Sphere Detection", *Journal of Signal Processing Systems*, vol. 70, is. 2, pp. 125-143, 2013.

[6] M. Winter; et al., "A 335Mb/s 3.9mm2 65nm CMOS flexible MIMO detection-decoding engine achieving 4G wireless data rates," *ISSCC Dig. Tech. Papers*, pp. 216-218, 2012.

[7] O. Arnold; et al., "Instruction Set Architecture Extensions for a Dynamic Task Scheduling Unit", *IEEE Symposium on VLSI*, pp. 249-254, 2012.

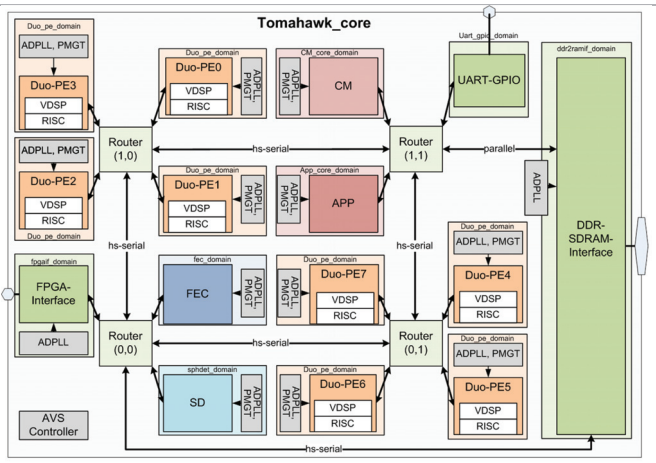

Figure 10.7.1: Block diagram of the Tomahawk2 MPSoC

	ASIC [3]	RISC[a]	ASIP (this work)
Scheduling Algorithm	As-soon-as-possible list-based		
Scheduling Configurability	Fixed	Flexible	Flexible
Task Queue size	16	16-256	16-256
Supported I/O transfers	8	12	12
Frequency [MHz]	175	445	445
Task scheduling [cycle]	60	7533	402
Task scheduling [us]	0.4	16.9	0.9
Technology [nm]	130	65	65
Supply Voltage [V]	1.3	1.2	1.2
Power Consumption [mW] @fmax	282	68	74.6
Energy per Task [nJ] @ fmax	113	1149	67
Area (logic) [mm²]	4.51	0.34[b]	0.49
ATE product [mm²*us*nJ]	204	6602	29

[a] Measured on the Tomahawk2 (this work) without application-specific instruction set
[b] Based on synthesis with Synopsys Design Compiler for 65nm LP TSMC process

Figure 10.7.2: Comparison of CoreManager realizations

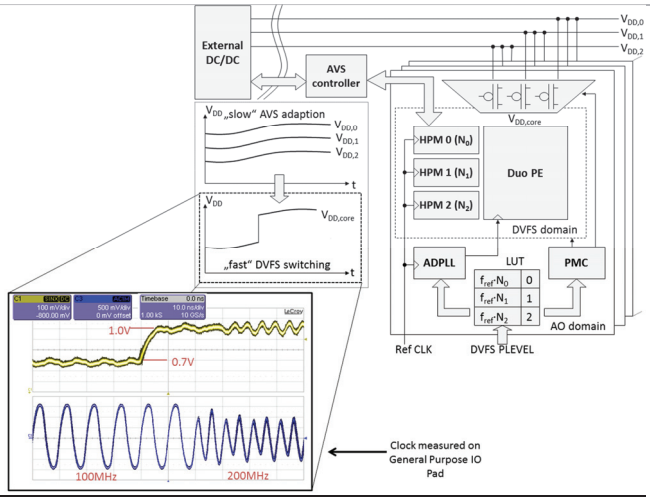

Figure 10.7.3: Combined power management, DVFS and AVS

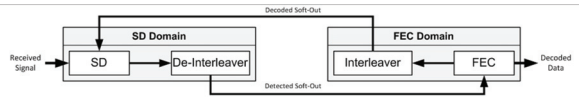

	SD		FEC		
	This work	[Borlenghi][a]	This work	[Borlenghi][a]	
Approach	SISO SD (up to 4x4, 64-QAM)	SISO SD (up to 4x4, 64-QAM)	Multimode decoding (LDPC, turbo, convolutional)	LDPC decoding only	Approach
Technology f_{max} (MHz) @ 1.2V Area (kGE)	65 nm 445 383[b]	65 nm 135 872[c]	65 nm 500 1.54[b]	65 nm 299 0.78[c]	Technology f_{max} (MHz) @ 1.2V Area (mm²)
Configuration	SO[d]	SISO[d] (2 it.)	SO[e]	SISO (2 it.)[e]	LDPC[f]	Turbo[g]	LDPC[h]	Configuration
Uncoded thrp. [Mb/s] @ f_{max}	396	162	194	66	155	45	586	Coded thrp. [Mb/s] @ f_{max}
AT-product (Mb/s/kGE)	1.03	0.42	0.22	0.08	100.92	29.41	751.00	AT-product (Mb/s/mm²)
Energy (pJ/b)	220	538	920	2690	232	1060	21	Energy (pJ/b/it)

[a] F. Borlenghi, et al., "A 2.78 mm2 65 nm CMOS gigabit MIMO iterative detection and decoding receiver", European Solid-State Circuits Conference, pp. 65- 68, 2012
[b] Memories for data exchange between SD and FEC included. [c] Memories for data exchange between SD and FEC excluded.
[d] 4x4 MIMO, 64-QAM, at 10^{-5} BER. [e] 4x4 MIMO, 64-QAM, at 1% BLER. [f] 768b code block, 3/4 rate, 10 iterations. [g] 1028b code block, 1/3 rate, 6 iterations. [h] 1944b code block, 5/6 rate, 10 iterations.

Figure 10.7.4: Comparison of detection-decoding engines

	Tomahawk2	Magali [1]	Tomahawk [3]	Intel SCC [2]
Platform Scope	MIMO 3GPP-LTE, WiMAX, 802.11n, SDR	MIMO 3GPP-LTE, WiMAX, 802.11n, SDR, Cognitive Radio	MIMO 3GPP-LTE, WiMAX, H264, SDR	SISO WiMAX, 802.11-04, SDR
Clocking and Power Management	GALS, local DVFS and AVS, power-gating	GALS, local DFS	Global frequency	Global frequency
Scheduling	Dynamic (flexible, energy adaptable algorithm)	Static	Dynamic (fixed algorithm)	Static
Memory Organization	Distributed and shared	Distributed	Distributed and shared	Local
Peak Performance	105 GOPS (3.6 GFLOPS)	37 GOPS	40 GOPS	-
Application for power measurements	4x4 MIMO 3GPP-LTE Rx baseband, 60 Mbit/s	4x2 MIMO 3GPP-LTE Rx, 2x2 MIMO Tx, MAC, 10.8 Mbit/s	LTE/WiMAX	-
Power consumption	480 mW @1.15V[a]	477 mW @ 1.2V	1.2W @ 1.3V	-
NoC Throughput (per link)	80 Gbit/s	17 Gbit/s	5.47 Gbit/s	8Gbit/s
Die size	36mm²	29.6 mm²	100 mm²	25mm²
Technology	65nm	65nm	130nm	65nm

[a] See application scenario power measurement results in Fig. 6

Figure 10.7.5: Comparison of state-of-the-art SDR chips

		Area [mm2] *total*	Area [mm2] *mem*	Mem size [bit]	f_{max} [MHz] @VDD =1.2 V	Throughput @f_{max}	P [mW] @f_{max}, VDD=1.2 V	$P_{Application\text{-}Scenario}$ [mW]
APP		0.582	0.245	274432	445	890 MOPS	49.7	off
CM		1.360	0.870	786432	445	1.1 MTasks/s	74.6	14.1 @200MHz, 0.9 V
Duo -PE[a]	**Xtensa**	1.357	0.800	532480	445	890 MOPS	61.5	off
	VDSP				500	10 GOPS	98.1	35.0 @282 MHz, 0.9 V
SD		0.522	0.260	292864	445	396 Mb/s	87.0	36.5 @200 MHz, 1.15 V
FEC		1.154	0.618	479232	500	155 Mb/s	360.0	132.2 @200 MHz, 1.15 V
FPGA-IF		0.602	-	-	500	10 Gb/s	-	-
DDR-IF		4.552	-	-	400	12.8 Gb/s	-	-
NoC		3.417	-	-	500	80 Gb/s[b]	32.0[c]	18.0 @286 MHz, 1.15 V

[a] Per instance. [b] Per link, max throughput. [c] Averaged over several test cases, NoC not fully utilized

Figure 10.7.6: Tomahawk2 components performance summary