

Area-Efficient Advanced Multiuser WCDMA Receiver

Sébastien Jomphe¹, Karim Cheikhrouhou², Jean Belzile¹, Sofiène Affes² and Jean-Claude Thibault³
¹École de technologie supérieure (Montréal), ²Institut national de la recherche scientifique, énergie, matériaux et télécommunications (Montréal), ³ISR Technologies inc. E-mail: sjomphe@ele.etsmtl.ca

Abstract—We explore the benefits of employing STAR as a wideband CDMA receiver. We show that it inherently addresses data recovery and channel characterization through its analysis/synthesis paradigm. We present simulation results carried out on real-world channel measurements showing the robustness and quick adaptation rate of STAR in terms of time-delay tracking performance and carrier frequency offset recovery. We then give an overview of the preferred hardware embodiment along with hard values in terms of hardware resource utilization, showing that a 24-user STAR-enabled 3G base station receiver can be realized within a single XC2V6000 FPGA.

I. INTRODUCTION

Ever since the RAKE receiver was introduced by Price and Green in the 1950's [1], it has largely dominated the field of spread spectrum communications. It enjoys such a widespread acceptance that personal communication systems still use the same basic receiver structure today, albeit enhanced by various contributions in terms of power combining, DOA tracking, interference cancellation and the like [2].

We propose a different approach to the direct sequence spread spectrum receiver which employs antenna array processing and inherently characterizes the channel jointly in space and time [3]. For this purpose, we establish a post-correlation model (PCM) of the observation vector upon which we base all further processing. We thereby skip monitoring the channel structure in its spread form in favor of tracking its key parameters in its despread form. This considerably minimizes the effects of uncorrelated interference and additive noise, yet allows dispersive effects to fall through and be more easily tracked in despread space. Furthermore, the time-delay resolution is no longer tied to the oversampling factor, enabling direct chip-rate sampling, which in turn lowers the effect of clock (or phase) jitter on receiver performance.

The Spatio-Temporal Array-Receiver (STAR) described herein can be used both as a means to perform channel characterization and as a data receiver. As a part of its observation mechanism, key channel parameters such as multipath time-delays and channel fading coefficients are monitored. Those parameters are used to replace the observed space-time propagation matrix with a far more accurate synthesized version, but they could also be logged for later channel structure analysis.

Even though we aim at proving the concept using 3G specifications [4], ongoing work will show that this paradigm

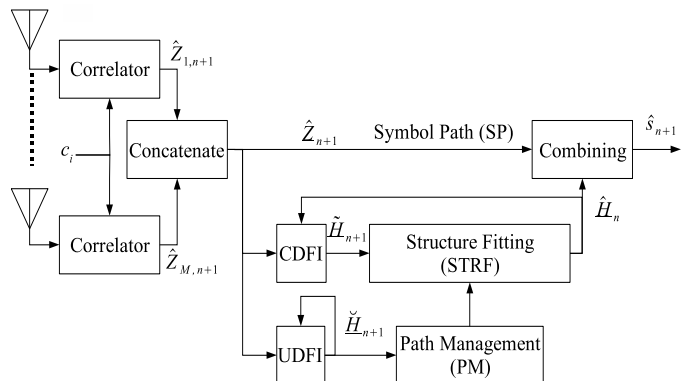


Fig. 1. Global topology of STAR with constrained and unconstrained decision feedback identifiers and structure fitting subsystem.

can be generalized to accommodate fourth generation systems which utilize multi-carrier modulation.

The structure of this paper is as follows. Section II will explain the STAR paradigm. Some of its key performances are outlined in section III while section IV shows the underlying hardware architecture of the system. Finally, we conclude with hardware resource utilization figures and expand those findings to a potential complete 3G base station receiver in Section V.

II. A NEW RECEIVER PARADIGM

Fig. 1 shows the global topology of STAR. M antennas each feed an observation vector to an array of M continuous despanders from which we get the PCM, defined as:

$$\underline{z}_n = \underline{\hat{H}}_n s_n + \underline{N}_n \quad (1)$$

where $\underline{\hat{H}}_n$ is the spatio-temporal propagation vector of the channel, s_n is the transmitted symbol and \underline{N}_n is an additive noise term. The underlined notation borrowed from [3,5,6] denotes vector-reshaping of the corresponding matrix otherwise noted in bold face. Along the symbol path (SP), the PCM reaches a combiner from which we extract the symbol:

$$\hat{s}_n = \frac{1}{M} \hat{H}_n^H \hat{z}_n, \quad (2)$$

where M is the number of receiving antennas and \hat{H} has a norm inherently constrained to \sqrt{M} . Aside from the isolation of the spread observation vector, note that multiple spatio-temporal channel estimates \underline{H} lie within the structure. $\underline{\hat{H}}$ and $\underline{\tilde{H}}$ are both from LMS-type space-time tracking decision feedback identifiers (DFI) [3] which adaptively monitor the

past state of the channel structure, symbol recovery statistics and the PCM of incoming observation vectors. They are defined as:

$$\check{\mathbf{H}}_{n+1} = \check{\mathbf{H}}_n + \zeta (\mathbf{Z}_n - \check{\mathbf{H}}_n \hat{\mathbf{b}}_n) \hat{\mathbf{b}}_n^* \quad (3a)$$

$$\check{\mathbf{H}}_{n+1} = \hat{\mathbf{H}}_n + \mu (\mathbf{Z}_n - \hat{\mathbf{H}}_n \hat{\mathbf{s}}_n) \hat{\mathbf{s}}_n^* \quad (3b)$$

where μ and ζ are adaptation step-sizes and $\hat{\mathbf{b}}$ is a hard-quantized version of $\hat{\mathbf{s}}$. $\check{\mathbf{H}}_n$ is a synthetic, noiseless version of $\hat{\mathbf{H}}_n$ from the structure fitting subsystem (STRF) which prevents the CDFI from going astray.

A second point of interest is the path management unit (PM). Its only role is to monitor the power level of known paths as well as unit time-delays and assesses the necessity to either lock onto emerging paths or release fading paths. This decision relies on a hysteresis mechanism to prevent false detections.

Crossing into the STRF, we first extract the spatial dimension from $\check{\mathbf{H}}$:

$$\hat{\mathbf{J}}_{n+1} = \check{\mathbf{H}}_{n+1} \hat{\mathbf{D}}_n^T, \quad (4)$$

where $\hat{\mathbf{J}}_{n+1}$ is an M -by- P spatial propagation matrix comprising the MP channel fading coefficients arising from the P paths and M antennas, and $\hat{\mathbf{D}}_n$ is a synthesized P -by- L temporal support matrix for each of the P paths from the previous structure-fitting iteration. An updated time estimate is then extracted by LMS-fitting $\hat{\mathbf{D}}_n$ onto the new spatio-temporal observation $\check{\mathbf{H}}_{n+1}$:

$$\hat{\mathbf{D}}_{n+1} = \hat{\mathbf{D}}_n + \frac{\xi}{M} (\check{\mathbf{H}}_{n+1}^T - \hat{\mathbf{D}}_n \hat{\mathbf{J}}_{n+1}^T) \hat{\mathbf{J}}_{n+1}^* \quad (5)$$

where ξ is an adaptation step-size and $\hat{\mathbf{D}}_{n+1}$ is the new temporal support estimate. We then determine the new optimal center-position of the chip impulse response for each of the P paths. This procedure is omitted for lack of space but can be found in [3]. The resulting temporal delays $\hat{\tau}_p$ have a precision of $0.001 T_c$ [3], far better than is possible with oversampling in RAKE-type receivers. The new synthetic temporal support matrix $\hat{\mathbf{D}}_{n+1}$ is then populated with P replicas of a delayed chip impulse response. The STRF then recombines the processed spatio-temporal components as:

$$\hat{\mathbf{H}}_{n+1} = \hat{\mathbf{J}}_{n+1} \hat{\mathbf{D}}_{n+1}^T. \quad (6)$$

We name this receiver paradigm the ‘‘analysis/synthesis’’ approach [5]. Key channel parameters such as $\hat{\tau}$ and $\hat{\mathbf{J}}$ are computed prior to the synthesis operation, and can be used for channel characterization. Further enhancements such as carrier frequency offset recovery (CFOR) [5] integrate seamlessly within this algorithm with only minor processing of $\hat{\mathbf{J}}$.

III. PERFORMANCE ANALYSIS AND VERIFICATION

We proceed to show simulation results which outline the robustness of STAR in terms of time-delay tracking and




Fig. 2. Performance analysis of STAR along test route 2 in Laval near Montreal. (a) shows impulse response contour and corresponding extracted time-delay values. (b) shows power density spectrum and corresponding extracted CFO (Δf) and Doppler spread ($\pm f_D$).

CFOR, namely. We used a recording of a wideband CDMA channel from test route 2 [5]. These recordings employ a base rate spreading factor of 256, a carrier frequency of 1.9825 GHz, a chip rate of 4.096 Mcps and a power control of ± 0.25 dB at a rate of 1600 Hz. Further details can be found in [5].

Time-delay synchronization is the most crucial of all aspects of the receiver. Ref. [6] shows that time drifts significantly degrade the performance of enhanced WCDMA receivers. Relying on its analysis/synthesis paradigm, STAR can monitor the precise time-evolution of multipath components in a realistic manner. Fig. 2-a shows the time-delay impulse response contour measured from the recording alongside the corresponding time-delays extracted by STAR. The algorithm was able to maintain tracking for 100%, 97% and 60% of the entire recording time, respectively, for paths of power levels 0 dB, -4.3 dB and -8.0 dB.

CFOR also plays a major role. Ref. [4] allows a 0.10 ppm mismatch between transmitter and receiver carrier frequencies on the uplink but significant losses in SNR have been shown to happen for even smaller discrepancies [3].

The CFOR algorithm of STAR instantaneously estimates and compensates such imperfections and reduces the SNR losses accordingly. To verify this, we show in Fig. 2-b the power spectral density of the first tracked multipath from test route 2 along with the extracted carrier frequency offset. Also depicted is the maximum Doppler spread. CFO extraction relies entirely on the PCM and compensation requires no explicit hardware in the RF chain.

IV. DATAFLOW AND BUILDING BLOCKS

We have previously published a framework from which STAR could be implemented [7]. Having recognized the duality of repetitive and logical (high branch count) operations, we have chosen a codesign approach. To this end, we have split STAR into three computational domains and kept cross-boundary bandwidth requirements at a minimum.

A) Algorithm Partitioning

Because a live receiver must handle incoming symbols at a fixed rate, the time allotted to (1) and (2) is finite. Intuitively, each received Z_{n+1} should be combined with a matched \hat{H}_{n+1} , but the amount of computations involved in synthesizing an optimal \hat{H}_{n+1} from \tilde{H}_{n+1} requires more time than is available.

Meanwhile, [6] shows that relaxing this one-to-one constraint has very little impact in terms of time-synchronization and received BER. It is suggested that updating \hat{H} every $n_{ID}=10$ symbols is acceptable when $L=32$. This amounts to structure-fitting the channel once every 10 symbols, or every 83 μ s. This unties the timing requirements of the STRF from those of the SP and allows us to confine them to separate clock domains, exchanging \tilde{H}_{n+1} and \hat{H}_{n+1} periodically.

To implement the PM subsystem, power level monitoring is required. By monitoring \hat{J} and \tilde{H}_{n+1} , the PM can assess the need to drop vanishing paths or lock onto emerging ones. Preliminary benchmarks suggest that the PM should run at a rate of approximately $10 \cdot n_{ID}$. Because this processing occurs at a comparatively low rate and is mainly composed of comparisons and branching, the PM is better suited to the software realm. Fig. 3 shows the three domains and how they interconnect.

B) Resource Reuse

STAR is ultimately meant as a multi-user receiver. It follows, then, that every spreading code in use will require its own SP, or more precisely its own despreader, in order to keep up with the incoming symbol rate. Sharing of other resources, such as DFIs, combiner, power estimator and STRF is, however, possible.

Structure-fitting \tilde{H}_{n+1} into \hat{H}_{n+1} is a sequential process that involves multiple intermediary variables. As such, this task can be segmented and carried out by specialized nanoprocessors separated by distributed memory resources acting as data conduits. The structure fitting is basically dependent on a feedback of \hat{D} between iterations n and $n+1$ (from m to f in Fig. 3). This translates into a lower bound on the value of n_{ID} that the hardware can offer for any one user ($n_{ID}=3$).

This pipeline structure is an obvious overdesign for a single user, but seeing as the algorithm is well segmented and globally sequential, interleaved processing of multiple disjoint data sets is possible, and can recover the otherwise idle processor time and thus increase data throughput. We can further exploit this, since only $n_{ID}=10$ is required, by simply interleaving 10 different users in the STRF while only

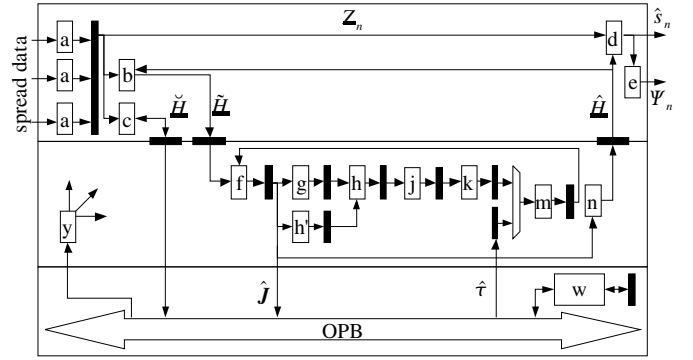


Fig. 3. Three-domain architecture of STAR. The top section shows the symbol path (SP), the middle section is the structure fitting pipeline (STRF) and the bottom section shows the path management subsystem (PM). Black rectangles are block RAM resources. Refer to ID column in Table I for legend.

incurring a 3-symbol latency to each one. Should the upper bound on n_{ID} in [6] be proved too strict, we could potentially relax it further and handle every user in a single STRF

The required nanoprocessors belong to five categories: despreader, matrix multiplier, FFT processor, linear regression fitter and norm. The different computations require these resources to be programmable both in terms of data bus bit widths, and operand sizes to suit dynamic operand sizes. The STAR Control Unit ("y" in Fig. 3) dynamically handles these configuration steps through a unidirectional setup bus.

D) Codesign

The PM unit is shown in the the lower part of Fig. 3 and runs within a Microblaze 32-bit soft microprocessor connected to an industry-standard on-chip peripheral bus (OPB). Sampling power levels is done by passive monitoring of STRF memory locations. Once it triggers a path removal or arrival for a user, the PM raises a flag to the STAR control unit so that the STRF pipeline will discard data from the time-delay update processor (Fig. 3-k) for this user and fetch new time-delays ($\hat{\tau}'$) from the PM to the fractional-delay impulse mapper (Fig. 3-m). The fetching operation only takes $8(MP+P)$ clock cycles to complete, much less than any other nanoprocessor would.

The software nature of the path management algorithms also presents the possibility of swapping them in a live system depending on the transmission environment, further enhancing the flexibility of STAR.

E) Data Passing

Two problems arise from inter-domain and intra-pipeline data-passing. First, we must consider the need for the input data from stage q of the pipeline to remain stable (and available) while stage $q-1$ produces other results. This is handled by using dual-port distributed block RAM (BRAM) resources (two independent read ports, one write port) and allocating a different memory address space to each possible data set (each user). The STAR Control Unit reprograms this information into each nanoprocessor at each pipeline hop.

TABLE I
HARDWARE RESOURCE USAGE FOR STAR ($M=1$)

Resource Name	ID	LUTs	MULTs	BRAMs
Symbol Path				
Despreader	a	456	-	-
Constrained DFI	b	566	2	-
Unconstrained DFI	c	566	2	-
Combiner	d	365	4	-
Power Estimator	e	62	3	1
Channel Structure Fitting				
Space-Time Separation	f	382	4	-
Conjugate and Multiply	g	69	2	-
Time Matrix Update (1/2)	h'	417	4	-
Time Matrix Update (2/2)	h	401	4	-
Fast Fourier Transform	j	1039	8	-
Time-delay Update	k	338	5	-
Fractional Impulse Mapper	m	101	-	3
Reconstruction	n	376	4	1
Path Management				
Microblaze and glue logic	w	5929	3	32
Misc				
STAR Control Unit	y	896	1	-
Pipeline framework	-	528	2	36
3G PHY				
Viterbi decoder		~18000	-	4
Cyclic Redundancy Checker		~200	-	-
Interleavers		~100	-	1
Total		~30791	48	78

Inter-domain data-passing as shown in Fig. 3 draws from the same conclusions. A further concern arises from the asynchronous nature of the STRF pipeline with regards to the SP. The exact time required for structure-fitting varies on a number of dynamic factors and cannot be determined with complete accuracy beforehand, so simultaneous access to the same \hat{H}_n can occur from both STRF and SP.

Three factors make this irrelevant. First, the BRAM is programmed with a vendor-specific "read after write" attribute which ensures that no erroneous data will be read should both ports access the same location simultaneously. Second, by architectural choice, the rate at which \hat{H}_n is read by the SP always exceeds that at which it can be written by the STRF, so that there can be no continuous read/write contentions. Third, by the nature of the algorithm, the overall shape of \hat{H}_{n+1} will never evolve by more than the adaptation step size μ (3b) compared to that of \hat{H}_n from one iteration to the next. An overlap could not catastrophically affect the system.

V. HARDWARE RESOURCE UTILIZATION

Our prototype uses a Xilinx Virtex2 6000 FPGA containing in excess of 67,000 look-up tables (LUTs), 144 dedicated block multipliers (MULTs) and 144 18-kilobit block RAM (BRAMs). Table I gives a breakdown of the elements in STAR

in terms of FPGA resources. Considering the amount of resources needed by the despreader unit compared to those of the STRF pipeline, the multi-user, single-STRF scenario seems promising. The multiple antenna case only incurs supplementary resources for added despreaders, and slightly increases the STRF cycle time for each user.

Using a Viterbi decoder from the Xilinx LogiCore™ library, parametrized to suit requirements in [4], we find that a dual-rate, single-channel unit would require 18,000 LUTs and 4 BRAMs and could be reused by each user. We further estimate that the remainder of the PHY layer which is comprised of interleavers and a CRC decoder would easily fit in the remaining resources (in excess of 35,000 LUTs). The global resource requirement for a 24-user base station as defined in [4], with a 10% overhead in glue logic would still fit in the afore-mentioned FPGA with a reasonable packing factor of 60%.

VI. CONCLUSION

We have presented the Spatio-Temporal Array-Receiver and shown how it exploits a new receiver paradigm which translates into increased channel identification performance compared to existing array receivers. The global analysis/synthesis approach results in a far more accurate spatio-temporal identification of the channel, thereby dramatically increasing the performance and inherently making STAR a live channel characterization tool. Of particular interest are time-delay extraction and carrier frequency offset recovery which have been addressed in Section IV with real-world channel measurements.

We then outlined the hardware framework needed to realize STAR, and have discussed specifics of its implementation such as partitioning, pipelining, resource reuse and codesign approach. Finally, we have given tangible resource utilization figures along with a breakdown of a potential STAR-enabled 3G base station receiver contained within a single FPGA.

REFERENCES

- [1] R. Price and P.E. Green, "A Communication Technique for Multipath Channels", *Proc. IRE*, vol. 46, 1958, pp. 555-570.
- [2] G.E. Bottomley, T. Ottoson, and Y.E. Wang, "A Generalized RAKE Receiver for Interference Suppression", *IEEE J. Select. Areas Commun.*, vol. 18, no. 8, 2000, pp.1536-1545.
- [3] S. Affes, and P. Mermelstein, "A New Receiver Structure for Asynchronous CDMA: STAR-- The Spatio-Temporal Array-Receiver", *IEEE J. Select. Areas Commun.*, vol. 16, no. 8, pp.1411-1422, Oct. 1998
- [4] 3rd Generation Partnership Project (3GPP), Technical Specification Group (TSG), Radio Access Network (RAN), and Working Group (WG4), "UE Radio Transmission and Reception (FDD)", TS 25.101, V3.3.0, 2000.
- [5] K. Cheikhrouhou et al., "Design Verification and Performance Evaluation of an Enhanced Wideband CDMA Receiver using Channel Measurements", to appear in EURASIP-JASP, 2nd Quarter 2005.
- [6] K. Cheikhrouhou, S. Affes, and P. Mermelstein, "Impact of Synchronization on Performance of Enhanced Array-Receivers in Wideband CDMA Networks", *IEEE J. Select. Areas Commun.*, vol. 19, no. 12, December 2001, pp. 2462-2476.
- [7] S. Jomphe, J. Belzile, S. Affes, and K. Cheikhrouhou, "Codesign Implementation of a 3G CWCDMA Base Station Receiver", in *Proc. IEEE CCECE'04*, Niagara Falls, Canada, May 2004, pp. 1191-1194.