Reducing dynamic power consumption in next generation DS-CDMA mobile communication receivers

V. Chandrasekhar

National Instruments, Austin, TX, USA
E-mail: Vikram.Chandrasekhar@ni.com

F. Livingston

Texas Instruments, Burlington, MA, USA
E-mail: Frank-Livingston@ti.com

J.R. Cavallaro*

Department of Electrical and Computer Engineering,
Rice University, Houston, TX, USA
E-mail: cavallar@rice.edu
*Corresponding author

Abstract: Reduction of the power consumption in portable wireless receivers is important for cellular systems, including UMTS and IMT2000. This paper explores the architectural design-space and methodologies for reducing the dynamic power dissipation in the Direct Sequence Code Division Multiple Access (DS-CDMA) downlink RAKE receiver. At the algorithm level, we investigate the tradeoffs of reduced precision and arithmetic complexity on the receiver performance. We then present and analyse two architectures for implementing the reference and reduced complexity receivers, with respect to dynamic power dissipation. The combined effect of reduced precision and complexity reduction leads to a 37.44% power savings.

Keywords: DS-CDMA RAKE receiver; VLSI architectures; mobile receiver; power reduction.

Biographical notes: Vikram Chandrasekhar received the MS Degree from Rice University in 2003 and the BTech Degree from the Indian Institute of Technology, Kharagpur in 2000. He is currently with National Instruments.

Frank Livingston received the MSEE Degree in 1995 and the BS Degree in 1992, both from the University of New Mexico. He is currently with Texas Instruments.

Joseph R. Cavallaro received the PhD Degree from Cornell University in 1988, the MS Degree from Princeton University in 1982, and the BS Degree from the University of Pennsylvania in 1981. He joined Rice University where he is currently a Professor in the Department of Electrical and Computer Engineering and Associate Director of the Centre for Multimedia Communication. He served as Program Director in the Prototyping Tools and Methodology program at NSF during 1996–1997, and has been a Visiting Professor at the University of Oulu, Finland during 2005. His research interests include computer arithmetic, VLSI and FPGA design, and VLSI/DSP architectures and algorithms for wireless communication systems.

1 Introduction

Achieving power-efficient architectures will be a major goal in the design of next-generation mobile communication receivers such as laptops, cell phones, PDA etc. Future portable receivers will need the ability to handle various multimedia data traffic irrespective of mobility, provide guaranteed Quality-of-Service (QoS) requirements, and integrate multiple functionality (GPS, World Wide Web, e-commerce etc.) simultaneously. The high bandwidth required by these applications implies that these functionality would come at the expense of a heavy drain on the available battery power. Shown in Table 1 are the

specifications of a next-generation wireless standard (IMT-2000) in Ojanperä and Prasad (2001). The high levels of expected performance, as well as the required data-rates will call for the implementation of advanced algorithms in the design of such receivers. With rapidly improving Integrated-Circuit (IC) technology as well as the decreasing cost of silicon area, there have been great advances in the ability to integrate the entire receiver chain on a single-chip (System-on-chip design). The point that has not been addressed in these designs is the system integration, with power minimisation as a key constraint. The design of such architectures forms a focal point in current research in communication systems.

Table 1 IMT-2000 Service requirements

Operating environment	Terminal speed (mph)	Peak bit-rate	Target BER
Rural outdoor	< 150	> 144 kbps	10^{-3}–10^{-7}
Urban/suburban outdoor	< 90	> 384 kbps	10^{-3}–10^{-7}
Indoor/low range outdoor	< 6	2 Mbps	10^{-3}–10^{-7}

1.1 Motivation

The RAKE receiver unit forms an important constituent of a DS-CDMA mobile receiver for performing single-user detection. The RAKE algorithm is a conceptually simple algorithm, however, its computational complexity increases linearly with the number of multi-path components being processed.

Even though there has been considerable research investigating techniques for improving the performance of DS-CDMA RAKE receivers in fading multi-path channels, there has been comparatively little research on investigating methodologies for minimising the power dissipation of the receiver architectures. A strength reduction technique has been described in Baghaie and Laakso (1991) for reducing the on line power dissipation in the complex RAKE multipliers by up to 25%. Power reduction techniques for a spread spectrum based correlator have been described in Garrett and Stan (1997) using a modified adder-tree structure and employing bus-invert coding. Low-power correlator architectures have been described in Sriram et al. (1999) that employ a partial correlation approach for reducing on line power dissipation during code acquisition in WCDMA based systems. To the best of our knowledge, there has been very little work on developing a framework which analyses the performance vs. power dissipation trade-offs in the context of mobile DS-CDMA RAKE receivers.

1.2 Contributions

The work presented in this paper has two principal aims. First, we analyse the impact of reduced precision and arithmetic complexity on the algorithm performance and power dissipation in the DS-CDMA mobile RAKE receiver. Next, we explore the architectural design-space for reducing the on line power dissipation. Starting with a conventional implementation of the RAKE receiver, we demonstrate design methodologies for achieving power reduction at the algorithm level and the architectural level. This 'proof of concept' architecture has been targeted towards a Xilinx Virtex-II FPGA and achieves the targeted data rate of 384 kbps. The resulting power-performance profiles have been obtained after passing synthesised complex receiver data simulating a urban three path fading channel through the targeted architectures.

- *Algorithm level.* We show that reduction of sampling rate of the input complex multi-path receiver data to the DS-CDMA RAKE correlator during de-spreading results in favourable trade-offs in power consumption vs. the corresponding receiver performance. Significant power savings are achieved through reduction in arithmetic complexity by decreasing the number of arithmetic operations during the RAKE correlation per symbol demodulation. For a 16 bit data-path, we have observed a 24.65% reduction in dynamic power dissipation in the reduced complexity RAKE receiver compared to the reference RAKE receiver implementation, with an acceptable performance loss of less than two dB.

- *Architectural level.* Starting with a 16 bit data path, and reducing precision till ten bits, we study the variation in the RAKE receiver performance with decreasing fixed-point precision. Word-length reduction alone results in power reduction of up to 25.6% in the reference RAKE receiver architecture, and 16.96% in the reduced complexity RAKE receiver architecture.

2 System description

We consider a K user DS-CDMA downlink system employing Binary Phase Shift Keying (BPSK) symbol modulation during transmission. The kth user's information sequence $b_k \in \{-1, 1\}$ is multiplied by a N chip Pseudo-Noise (PN) sequence whose bit duration equals $T_{bit} = NT_{chip}$. For purposes of estimating the complex channel coefficients (see Fantaccci and Galligani (1999), Viterbi (1995) and Rappaport (1986)), a common code-multiplexed pilot signal is broadcast by the base station to all mobile users. The sampled complex receiver data $r(n)$ at the DS-CDMA mobile receiver can be written in vector-matrix notation as in Latva-aho and Juntti (2000); Chandrasekhar (2002) as $\mathbf{r_i} = \mathbf{SH_iAb_i} + \mathbf{w_i}$ where

- $\mathbf{r}_i$ is the received sampled data (S samples/chip) corresponding to the ith information symbol represented by

$$\mathbf{r}_i = [r(iNST_s)\ r((iNS + 1)T_s)$$
$$\dots r([(i + 2)NS - 1]T_s)]^T$$
$$\in \mathbb{C}^{2NS \times 1}.$$

- **S** describes the signature matrix for all K active users and the pilot channel given by

$$\mathbf{S} = [\mathbf{s}_{1,1}, \ldots, \mathbf{s}_1 p, \ldots, \mathbf{s}_{K,1},$$
$$\ldots, \mathbf{s}_{K,P} \; \mathbf{s}_{\text{pilot},1}, \ldots, \mathbf{s}_{\text{pilot},P}]$$
$$\in \mathbb{R}^{2NS \times (K+1)P}.$$

Each of the columns $\mathbf{s}_{k,p}$, $1 \le k \le K+1$, $1 \le p \le P$ represents the appropriately delayed (by $\lceil NS\tau_p / T_b \rceil$ samples) signature waveform of the kth user and pth multipath. Therefore,

$$\mathbf{s}_{k,p} = \left[\mathbf{0}_{1 \times \lceil NS\tau_p / T_p \rceil} \mathbf{s}_k^T \quad \mathbf{0}_{1 \times (NS - \lceil NS\tau_p / T_b \rceil)} \right]^T$$
$$\in \mathbb{R}^{2NS \times 1}$$

and

$$\mathbf{s}_k = [s_k(T_s)\, s_k(2T_s) \ldots (NST_s)]^T$$
$$\in \mathbb{R}^{NS \times 1}.$$

where $s_k(t)$ represents the kth user's continuous-time spreading waveform given by the convolution of the user's spreading sequence $\{ck(n)\}$ and the transmitted chip-waveform $g_T(t)$.

- **H$_i$** denotes the complex channel impulse response coefficient matrix for the ith information symbol given by

$$\mathbf{H_i} = \begin{bmatrix} \mathbf{h}_i & 0 & \cdots & 0 \\ 0 & \mathbf{h}_i & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & & \cdots & \mathbf{h}_i \end{bmatrix} \in \mathbb{C}^{(K+1)P \times (K+1)}$$

and

$$\mathbf{h}_i = [h_{i,1} \quad h_{i,2} \quad \ldots \quad h_{i,P}]^T \in \mathbb{C}^{P \times 1}.$$

- **A** is the user/pilot amplitude matrix given by diag $\{A_1, A_2, \ldots, A_K, A_{\text{pilot}}\} \in \mathbb{R}^{(K+1) \times (K+1)}$

- **b$_i$** is the symbol vector for all K users and pilot corresponding to the ith transmission given by

$$\mathbf{b_i} = [b_1 \quad b_2 \quad \ldots \quad b_K \quad 1]^T \in \mathbb{R}^{(K+1) \times 1}.$$

2.1 DS-CDMA RAKE receiver

The DS-CDMA RAKE receiver attempts to collect the signal energy from all the received signal paths that fall within the delay line and carry the same information as described in Proakis (1995). Assuming that user one is the user of interest, we define the signature matrix,

$$\mathbf{S}_1 = [\mathbf{s}_{1,1} \quad \mathbf{s}_{1,2} \quad \ldots \quad \mathbf{s}_{1,P}]^T \in \mathbb{R}^{2NS \times P}$$

the RAKE receiver computes the decision statistic given by:

$$\hat{b}_{1,i} = \text{sgn}(\mathbf{S}_1 \hat{\mathbf{h}}_i A_1)^H \mathbf{r}_i$$
$$= \text{sgn}(\mathbf{S}_1 \hat{\mathbf{h}}_i A_1)^H (\mathbf{SH}_i \mathbf{Ab}_i + \mathbf{w}_i) \tag{1}$$

where $\hat{\mathbf{h}}_i \in \mathbb{C}^{P \times 1}$ is the complex channel coefficient estimate obtained from the output of a channel estimator. To estimate the complex channel coefficient for performing phase offset correction, a channel estimator is required. A L tap moving average filter performs channel estimation while demodulating the ith information symbol. An all-ones pilot symbol sequence (assumed to be known at the mobile receiver) is used for the purpose of channel estimation. For the pilot sequence, define

$$\mathbf{S}_{\text{pilot}} = [\mathbf{s}_{\text{pilot},1} \quad \mathbf{s}_{\text{pilot},2} \cdots \mathbf{s}_{\text{pilot},p}]$$
$$\in \mathbb{R}^{2NS \times P}$$

as the pilot code signature matrix. Then, the channel estimate $\hat{\mathbf{h}}_i$ is given by the expression

$$\hat{\mathbf{h}}_i = \sum_{i=n-L+1}^{n} b_{1,i}^* \mathbf{S}_{\text{pilot}}^H \mathbf{r}_i \tag{2}$$

where L is the length of the averaging filter.

3   Receiver architecture

Figure 1 shows the high-level description of the front-end in a wireless communication receiver. The architectures implemented in this paper are represented by the solid line blocks (corresponding to the RAKE receiver), while the dashed-line blocks are assumed to feed in the sampled wide-band signal and the estimated delays into the receiver. The sampled complex wide-band receiver data is input to the RAKE receiver for performing symbol-level demodulation, and the delay-tracker block for initial timing acquisition followed by fine synchronisation with a delay-locked loop. For the ith symbol interval, complex receiver data $\mathbf{r}_i$ is input to the RAKE receiver in a chip-serial fashion.

Figure 1 Front-end description of a wireless communication DS-CDMA receiver
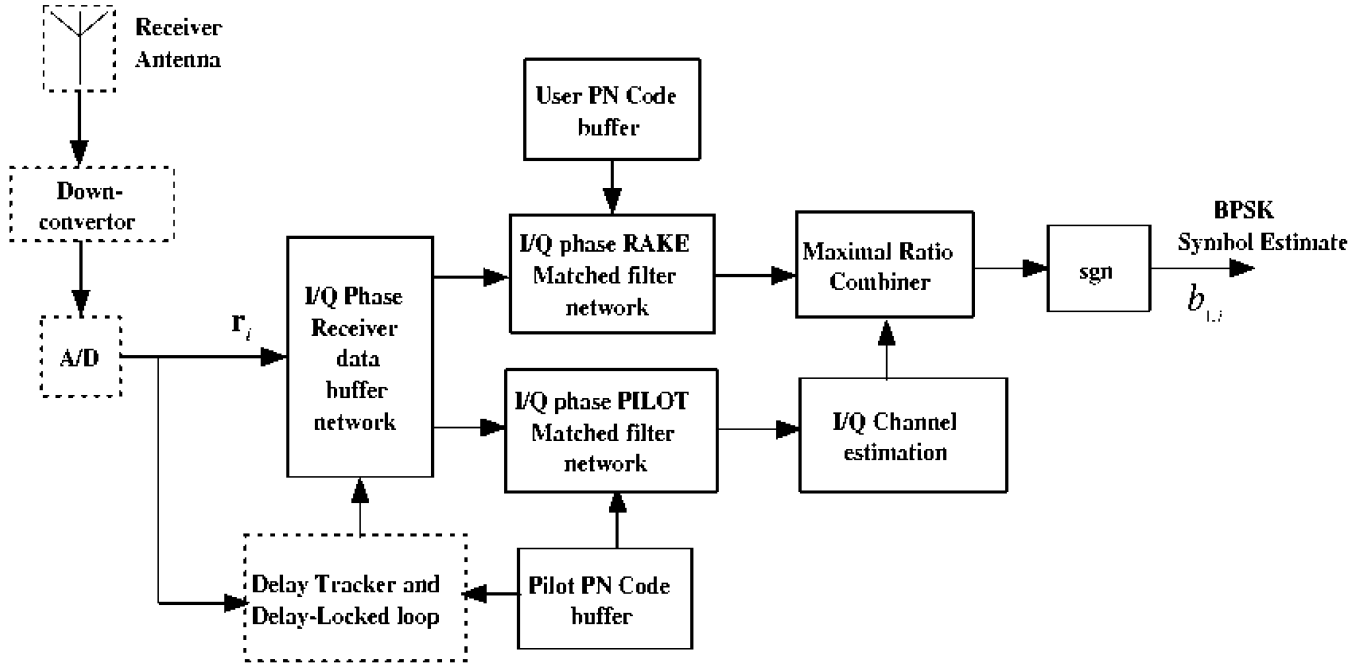
3.1 Front-end circular buffer

The front end circular buffer stores complex sampled receiver data $\mathbf{r}_i$ to be used for performing correlation during the channel estimation and detection operations, as shown in Figure 2. Denoting the maximum delay spread of the channel is a symbol duration D, N is the PN sequence length, S is the number of samples per chip, the minimum required buffer size is given by $B = NS \times [D/N + 1]$ words. Assuming a maximum delay spread of a symbol $D = T_b$, processing gain of $N = 32$ chips and $S = 2$ samples/chip, we obtain $B = 128$ words, requiring an address width of 7 bits. The buffer employs the following modes of operation:

- *Initialisation mode*. For the first $NS = 64$ cycles, the buffer is written into, till there is a symbol-duration worth of receiver data stored in the buffer. During this period, all the read-addresses are set to a value of 128 to ensure that there is no memory access conflict generated by the read and write addresses. The receiver data is written into the *recv_buff_wr_addr* specified from 0 to $NS - 1 = 63$.When *recv_buff_wr_addr = NS = 64*, the mode changes to the steady-state mode described below.

- *Steady state mode*. At the end of the Initialisation mode, there are $NS = 64$ words of receiver data (or 1 symbol worth of information) stored in the buffer. Conditioned on the current value of *recv_buff_wr_addr,* the read-port address *read_address(i)*, $i = 0 \dots P - 1$ is either initialised with the computed path delays *Delay(i)* (from the delay-tracker unit), or incremented by one. The read- port addresses are specified by,

$$read\_Address(i) = \begin{cases} Delay(i) \; if \\ recv\_buf\,f\_wr\_addr = 64 \\ \\ Delay(i) + 64 \; if \\ recv\_buf\,f\_wr\_addr = 0 \\ \\ read\_Address(i) + 1 \\ otherwise \end{cases}.$$

As the read-addresses get incremented, successive complex receiver data values get read from the buffer and are input to the RAKE/PILOT correlator units where the correlation of the receiver data with the user/pilot codes is carried out.

**Figure 2**   Buffer for storing complex receiver data
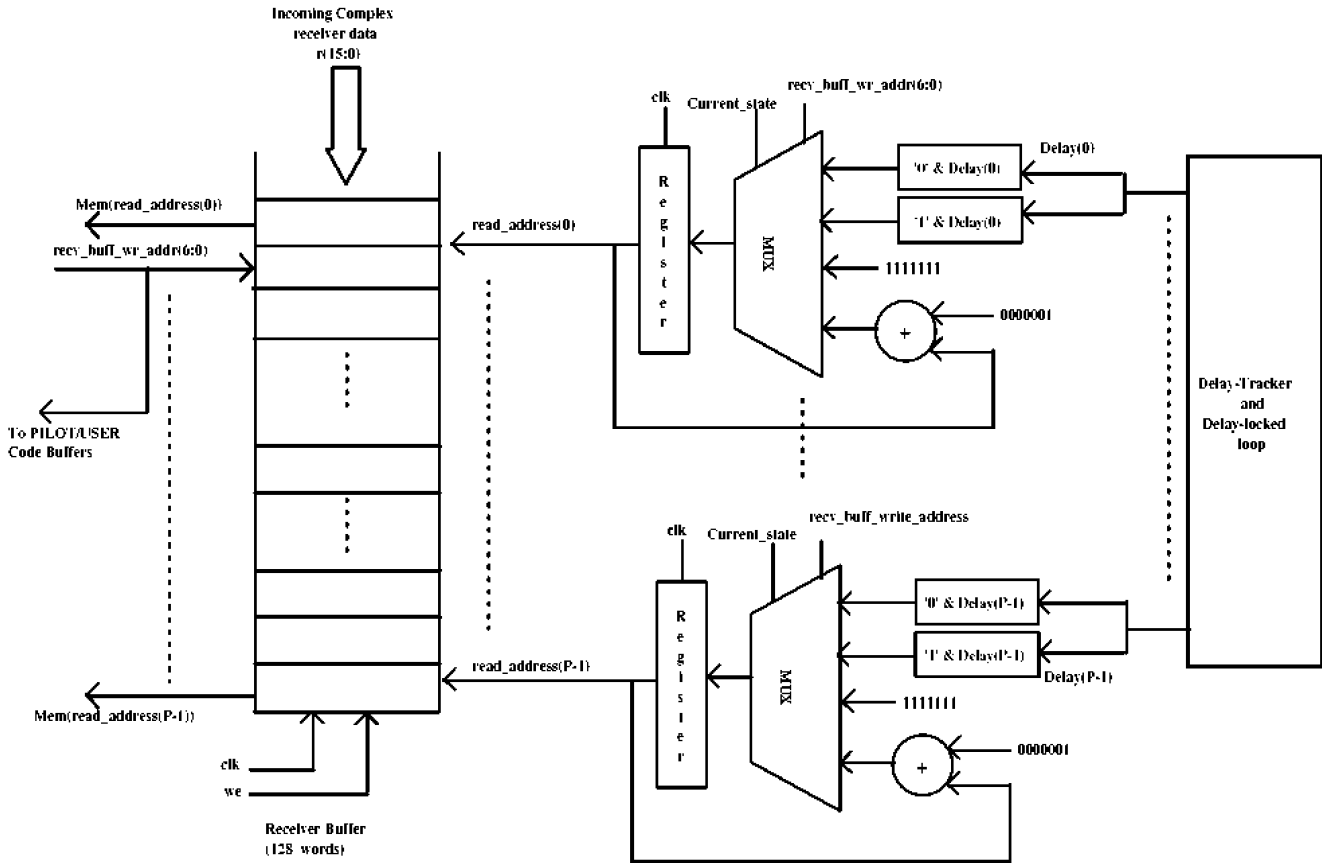
For a *P* finger RAKE receiver, the above description assumes a receiver buffer with *P* read-ports each for the real and imaginary parts of the complex receiver data. In a practical implementation however, truly multi-ported buffers are infeasible owing to the high output load capacitance which would dramatically increase the memory access time, and hence decrease the operating frequency of the design. An alternative is to use a serial shift register delay line implementation. For a *n* bit data-path, there are $\frac{n}{2}$ logic transitions (for real and imaginary data storage registers) potentially occurring at every node per clock cycle, due to shifting of data in the shift register unit as shown in Garrett and Stan (1997) and Sriram et al. (1999). For the shift register storage size of $2NS = 128$ words of *n* bit receiver data, this would amount to an average of $NSn = 64n$ logic transitions per clock cycle which is clearly power-inefficient. Consequently, a much simpler approach was adopted by instantiating *P* separate SRAM based dual-ported receiver data buffers, to store the real and imaginary components of the input sampled receiver data. This ensures a smaller output load capacitance at the data-bus compared to the register-file based approach. Moreover, the use of the pointer-based approach implies that the switching activity in the data-bus is reduced from $NSn = 64n$ logic transitions to just $Pn/2 = 1.5n$ logic transitions per clock cycle on an average.

3.2   User/Pilot PN code circular buffer

The User/Pilot code circular buffers (Figure 3) store the length $NS = 64$ PN sequences of the user $(\mathbf{s}_i \in \mathbb{R}^{64 \times 1})$ and pilot codes $(\mathbf{s}_{\text{pilot}} \in \mathbb{R}^{64 \times 1})$. We assume that the code coefficients are pre-determined at start-up and stored in the buffer. Since the symbol despreading operation begins when *recv_buff_wr_addr* = 1 or *recv_buff_wr_addr* = 65, the read-address pointer *code_read_addr* (6 bits wide) for the buffer is directly determined by the current write address *recv_buff_wr_addr* of the front end circular buffer, by the relation *code_read_addr* = *recv_buff_wr_addr* mod 64. While *recv_buff_wr_addr* counts from 0 to $2NS - 1 = 127$, *code_rd_addr* counts from 0 to $NS - 1 = 63$.

3.3   PILOT/RAKE matched-filtering block

The PILOT and RAKE correlator blocks take in the sampled wide-band signal, perform matched-filtering, and output a narrow-band signal at the symbol rate. The narrow band output and channel estimates are input to the Maximal Ratio combiner (MRC) which performs coherent demodulation. Figure 4 illustrates the architecture of the correlator network for computing the complex inner products, $pilot\_soft\_out\,(i) = \mathbf{S}_{\text{pilot}}^{\text{H}}\mathbf{r}_i$ and $rake\_soft\_out\,(i) = \mathbf{S}_i^{H}\mathbf{r}_i$ where i, $0 \le i \le P-1$ corresponds to the *i*th finger and $\mathbf{r}_i$ refers to the delayed multi-path data coming from the receiver circular buffer.

Figure 3 Circular buffer for storing the user/pilot code coefficients
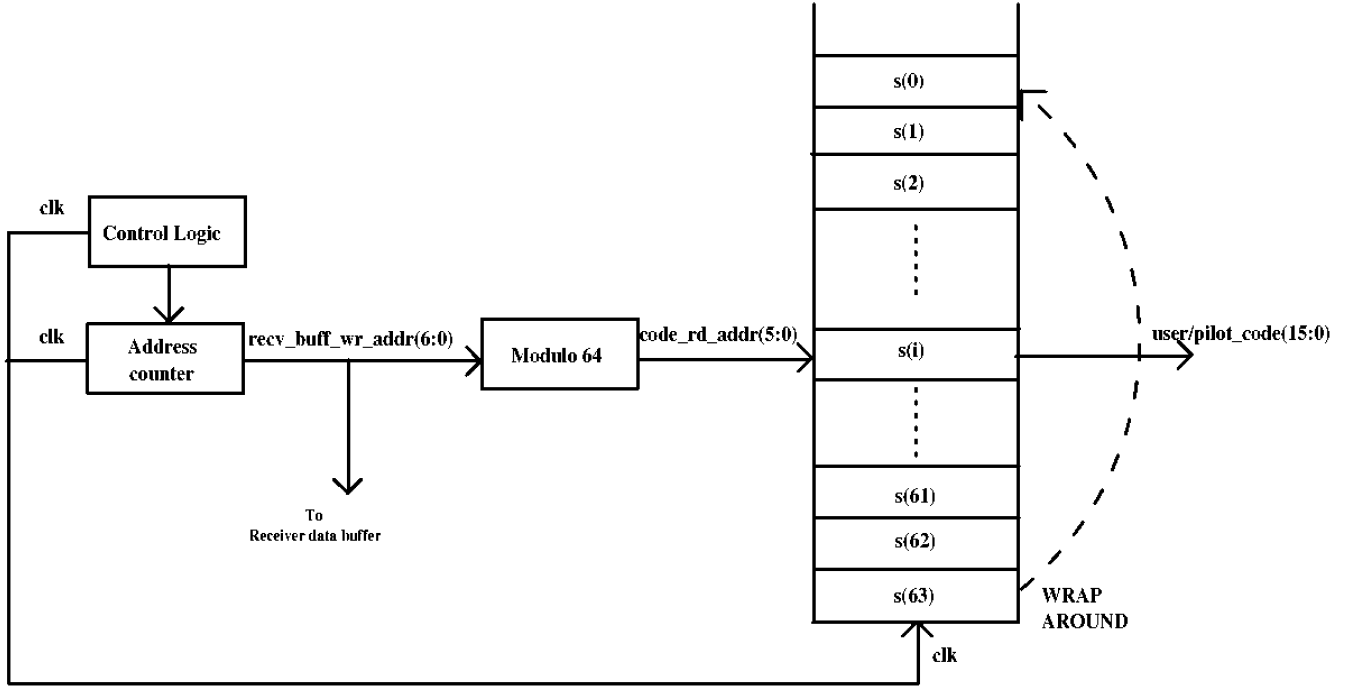
**Figure 4** Structure of the PILOT/RAKE finger network (see online version for colours)
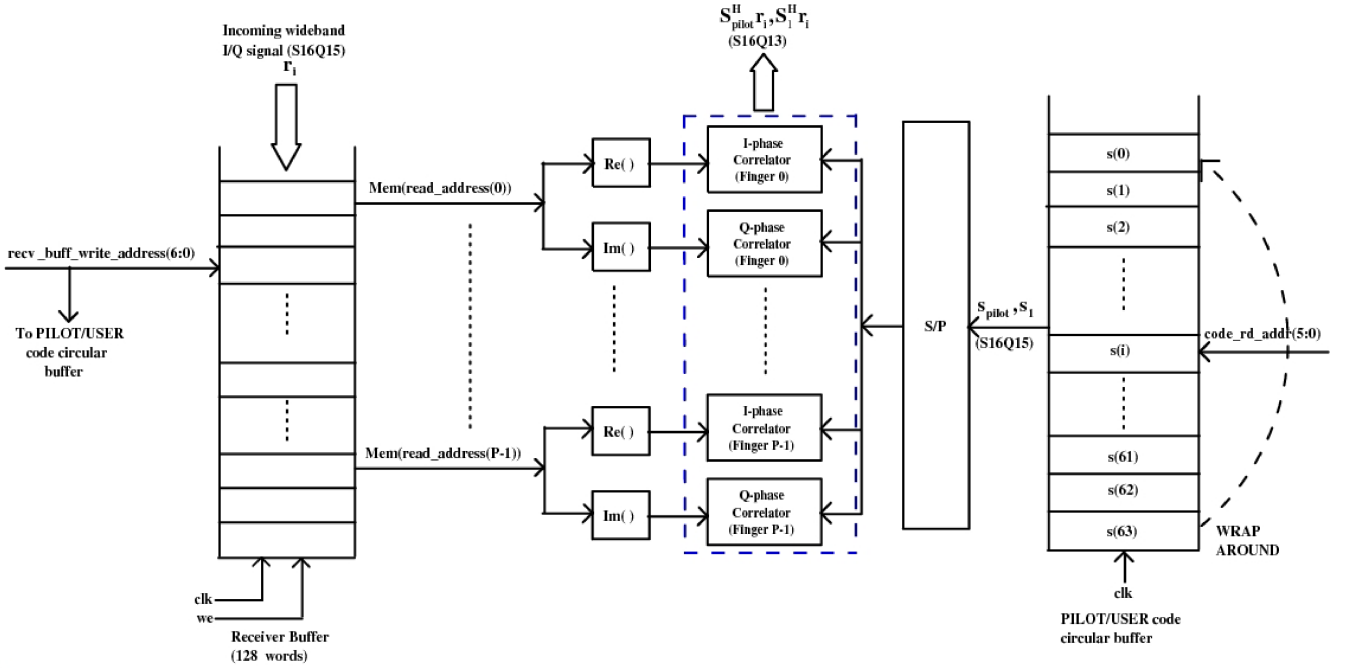
3.4 Channel estimation

The channel estimator block uses a simple moving averager filter to estimate the complex multi-path channel coefficients. For the ith symbol demodulation, the estimated channel coefficient $\hat{\mathbf{h}}_i \in \mathbb{C}^{p \times 1}$ is computed by,

$$\hat{\mathbf{h}}_i = \frac{1}{L} \sum_{k=n-L+1}^{n} b_{pilot,k}^* \mathbf{S}_{pilot}^H \mathbf{r}_k. \qquad (3)$$

Since the filtering operation (see Figure 5) computes the channel estimates based on the results of the previous L pilot correlations, a L word circular buffer based implementation is employed. In the practical implementation, the pilot sequence is assumed to be an all ones sequence, therefore, the computation of $\hat{\mathbf{h}}_i$ is simplified as shown,

$$\hat{\mathbf{h}}_i = \hat{\mathbf{h}}_{i-1} + \{ \mathbf{S}_{pilot}^H (\mathbf{r}_i - \mathbf{r}_i - L) \} \gg [\log_2 L]. \qquad (4)$$

Figure 5   Moving average based channel estimation
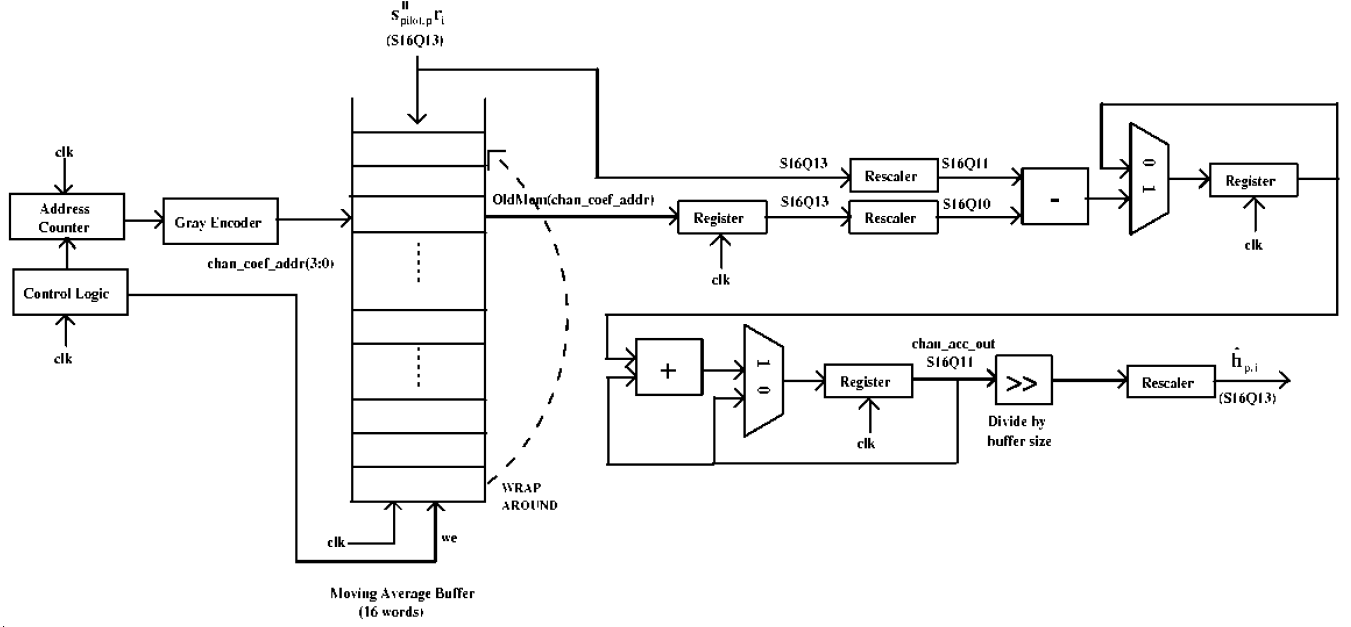
Corresponding to symbol *i*, the despread pilot correlation output $\mathbf{S}_{\text{pilot}}^{H}\mathbf{r}_i$ is written into the circular buffer address *i* mod *L* where it replaces the oldest pilot correlation output $\mathbf{S}_{\text{pilot}}^{H}\mathbf{r}_{i-L}$. The difference between these two values is used to compute the *i*th channel estimate as indicated in equation (4). A 'read before write' type of circular buffer is chosen for implementation, in order that the oldest pilot correlation value is read out before being overwritten by the new pilot correlation output. For performing the scaling, *L* is chosen to be a power of 2, in order to replace the division operation by right shifting by *L* bits.

3.5   *Maximal ratio combining*

The MRC weights the narrow band despread outputs of the RAKE finger network, by the corresponding complex conjugated channel coefficient estimates. Figure 6 shows the implementation of the MRC unit. The five stage pipelined multipliers implement the phase rotation operation $\hat{\mathbf{h}}_i^{H}\mathbf{S}_1^{H}\mathbf{r}_i$. The [$\log_2 P$] stage deep adder tree network combines the phase rotated outputs and produces a single soft symbol estimate. Finally, the hard symbol estimate $b_{1,i}$ corresponding to the *i*th transmitted symbol of user one is computed by taking the sign of the MRC output as shown in equation (5).

$$b_{1,i} = \text{sgn}(\text{Re}(\hat{\mathbf{h}}_i^{H}\mathbf{S}_1^{H}\mathbf{r}_i)). \tag{5}$$

Figure 6   Maximal ratio combiner unit
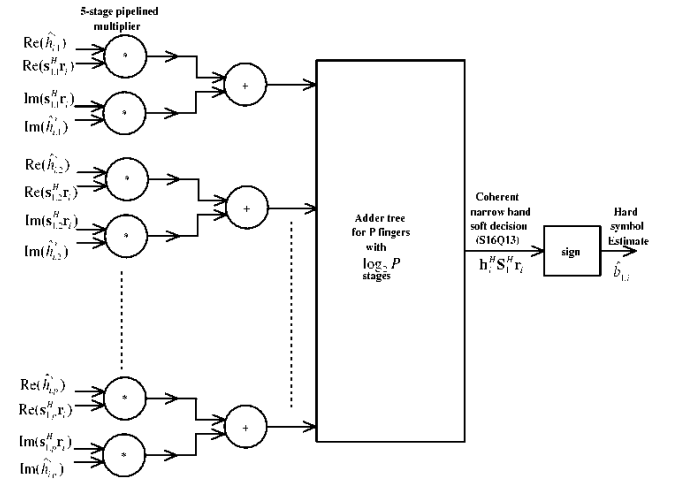
4   Power-efficient architectures

Dynamic power dissipation is usually the dominant source of power dissipation in CMOS VLSI circuits. The dynamic power consumption P_{dyn} at any node in a CMOS-based design is a function of the node capacitance *C*, the switching activity α of the node (defined as the average number of node transitions per clock cycle), the clocking frequency f_{clock}, and the supply-voltage V_{cc} employed in the design, given by equation (6)

$$P_{dyn} = \frac{1}{2}\alpha C V_{cc}^2 f_{clock}. \tag{6}$$

Since P_{dyn} is quadratically related to V_{cc}, voltage reduction yields the biggest savings in power consumption. As voltage reduction results in increased combinational logic delays as shown in Chandrakasan et al. (1995), techniques such as pipelining and parallelism are employed for maintaining a constant throughput of the design. In addition, optimisations such as reduced algorithmic complexity, re-ordering of arithmetic expressions, word-length reduction can markedly reduce the overall capacitance and node switching activity in the design, thereby reducing the power-dissipation (detailed description is provided in Chandrakasan et al. (1995) and Rabaey and Pedram (1996)). At the circuit level, clock-disabling techniques that turn off idle functional units can be exploited to extract further power savings.

4.1 Reduction in arithmetic complexity

The computationally most intensive operation involved in the RAKE receiver is the correlation operation where the sampled complex multi-path receiver data is correlated with the spreading waveform vector for the user and pilot channels. For the pth finger, the correlation output $X_{cor}^p(i)$ corresponding to the ith signaling interval can be represented by,

$$X_{cor}^p(i) = \int_{iT+\tau_p}^{(i+1)T+\tau_p} r(t)s_1(t-iT-\tau_p)dt \tag{7}$$

where

$$s_1(t) = \sum_{n=0}^{N-1} c_k(n)g_T(t-nT_c)$$
$$= \left(\sum_{n=0}^{N-1} c_k(n)\delta(t-nT_c)\right) \otimes g_T(t).$$

When implementing the correlation operation as a digital matched filter, the complexity of the correlation operation is governed by the length of the signature waveform vector N_{corr} and the number of active fingers P. The signature waveform vector $\mathbf{s}_{1,p}$ is represented by the discrete-time convolution of the length N spreading sequence $\{c_1(n)\}$ and the M tap raised cosine filter with impulse response $\{g_T(n)\}$. The square-root raised cosine filter is given by $M = 2DS + 1$ taps (being linear phase) where D is the group delay of the filter and S is the upsampling rate at the filter input. The length of the convolution output is given by
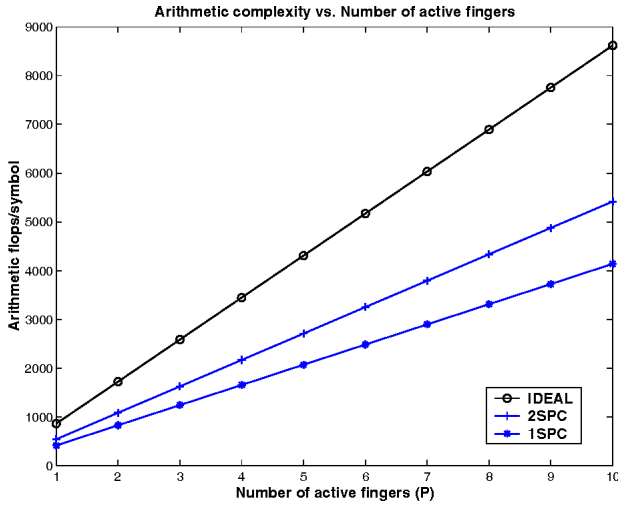
$N_{conv} = M + NS - 1$ samples. Assuming values of $D = 10$ samples, $S = 2$ samples/chip, we obtain $M = 41$, $N_{conv} = 2N + 40$ samples, hence the overall correlator length is specified by $N_{corr} = N_{conv}$. For typical values such as a spreading code of length $N = 32$, $P = 3$ path channel, $L = 16$ tap channel estimator , the arithmetic complexity of the RAKE receiver with ideal correlation equals $16NP + 318P + 2LP - 1 = 2585$ flops/symbol. We explore the reduction in the correlator length as a means for achieving reduction in arithmetic complexity in Table 2 and Figure 7. We consider the following two schemes:

- *Sampling at 2 samples/chip*. The starting and ending $DS = 20$ samples of the spreading waveform at the convolution output occur due to the group delay of the filter $g_T(n)$. By discarding these $2DS = 40$ samples and retaining the steady state response, the correlator length reduces to $N_{corr} = N_{conv} - 40 = 2N$ samples/symbol, which translates into savings in arithmetic complexity. Thus the number of correlation operations involved in the pilot correlators (for channel estimation) and rake correlators (for despreading and detection) are reduced by $320P = 960$ flops/symbol to 1625 flops/symbol. In the results , the performance of the resulting receiver (with truncated correlation waveform) is shown to be almost identical with that obtained with perfect correlation. We call this receiver as the *reference RAKE receiver*.

- *Sampling at 1 sample/chip*: To achieve a reduction in the arithmetic complexity, we reduce the sampling rate for the despreading operation in the RAKE correlators to one sample/chip, and investigate the resulting complexity vs. performance trade-offs. This halves the length of the correlator for the RAKE de-spreading operation to $N_{corr} = N$ samples/symbol and a corresponding reduction in the overall flop count by $4NP = 384$ flops/symbol to 1241 flops/symbol. As the performance of detection is heavily influenced by the accuracy of channel estimates, the pilot channel correlation is still performed at two samples/chip. The complexity reduction comes at the tradeoff of reduced correlator output energy owing to the halved correlation length. The results demonstrate a significant power reduction with acceptable detection performance due to this optimisation. We call this receiver as the *reduced complexity RAKE receiver*.

Table 2 Arithmetic complexity per symbol detection in reference and reduced complexity receivers

Operation	Multiplications	Additions
$\mathbf{S}_1^H \mathbf{r}_i \in \mathbb{C}^{P\times 1}$	$4NP/2NP$	$2P(2N-1)/2P(N-1)$
$\mathbf{S}_{pilot}^H \mathbf{r}_i \in \mathbb{C}^{P\times 1}$	$4NP$	$2P(2N-1)$
$\hat{\mathbf{h}}_i = \frac{1}{L}\sum_{k=i-L+1}^{i} b_k \times \mathbf{S}_{pilot}^H \mathbf{r}_k \in \mathbb{C}^{P\times 1}$	–	$2P(L-1)$
$Re(\hat{\mathbf{h}}_i^H \mathbf{S}_1^H r_i) \in \mathbb{R}^{1\times 1}$	$2P$	$2P-1$
RAKE receiver (2 samples/chip)	$(16NP + 2LP - 2P - 1)$ flops	
RAKE receiver (1 sample/chip)	$(12NP + 2LP - 2P - 1)$ flops	

Figure 7   Arithmetic complexity in flops per symbol (see online version for colours)

parameterised, so that the precision requirements of the entire design could be changed 'off-line' with minimal modifications.

**Table 3**   Fixed-point precision requirements for the RAKE receiver

Detector variable	Description	Integer bits
$\mathbf{r}_i \in \mathbb{C}^{2NS \times 1}$	Complex receiver input data	1
$\mathbf{S}_1 \in \mathbb{R}^{2NS \times P}$	User signature matrix	1
$\mathbf{S}_{\text{pilot}} \in \mathbb{R}^{2NS \times P}$	Pilot signature matrix	1
$\mathbf{S}_1^H \mathbf{r}_i \in \mathbb{C}^{P \times 1}$	Soft rake correlator output	3
$\mathbf{S}_{\text{pilot}}^H \mathbf{r}_i \in \mathbb{C}^{P \times 1}$	Soft pilot correlator output	3
$\sum_{k=i-L+1}^{i} b_{1,k}^* \mathbf{S}_{\text{pilot}}^H \mathbf{r}_k \in \mathbb{C}^{P \times 1}$	Moving average accumulator	5
$\hat{\mathbf{h}}_i = E[b_{1,i}^* \mathbf{S}_{\text{pilot}}^H \mathbf{r}_i] \in \mathbb{C}^{P \times 1}$	Channel coefficient estimate	3
$(\mathbf{S}_1 \hat{\mathbf{h}}_i A_1)^H \mathbf{r}_i \in \mathbb{C}^{1 \times 1}$	Maximal ratio combiner output	6

In a practical implementation of the reduced complexity DS-CDMA RAKE receiver, the halving in the input sampling rate to the RAKE despreading unit would imply that the RAKE correlation would complete twice as fast as the PILOT correlation. This means that the RAKE correlator would remain idle for half the symbol duration, and still be clocked by the sample-rate clock, resulting in wasteful dissipation of idle clocking power. Therefore, the clock input for the RAKE correlation network is derived from the global clock at half the input sampling rate. Note that the reduced clocking rate *does not* reduce the effective symbol rate of the system.

4.2   Reduction in fixed-point precision in the DS-CDMA RAKE receiver

All the DS-CDMA architectures presented in this paper are based on a fixed-point implementation. A quantisation analysis tool developed at the University of Texas, Dallas in Linebarger et al. (2000) was used for determining the dynamic range and precision requirements of the RAKE receiver. This paper assumes that all the fixed-point variables are quantised with a uniform width and only differ in their integer bit requirements.

Table 3 shows the fixed-point integer requirements of the individual RAKE receiver variables after quantisation analysis. The corresponding fractional bit-width requirements were determined from the difference of the overall precision and the number of integer bits. From the obtained fixed-point formats, extensive simulations were carried out using MATLAB/C with C++ classes in SystemC providing the fixed-point arithmetic support.

A minimum word-length of 10 bits was required for the RAKE receiver to achieve acceptable performance (within 1 dB) of the equivalent floating point version of the algorithm (this will be discussed further in the next section highlighting the results). The implementations of the reference and reduced complexity RAKE receivers (performed on the Virtex-II FPGA) were made

4.3   Architecture description

To quantify the effect of the aforementioned methodologies, architectures incorporating the power saving techniques were implemented on a Virtex-II FPGA. This paper describes two distinct architectures based upon which the results are reported. They are enumerated below:

- *Reference architecture*. Figure 8 shows the reference architecture of the RAKE receiver. This implementation employs a uniform input sampling rate of two samples/chip for both the PILOT and RAKE correlator matched filtering operations. The external clock is passed through a delay-locked loop to derive the global clock buffer *CLK* running at the input sample frequency of $f_{\text{samp}} = 24.576$ MHz.

- *Reduced Complexity architecture*. To explore the effects of reduced arithmetic complexity on the resulting power consumption of the RAKE receiver, the wide-band signal was input at the rate of two samples/chip to the PILOT correlator and 1 sample/chip to the RAKE correlator. Figure 9 shows the architecture of the resulting reduced complexity RAKE receiver with two separate clocking domains namely *CLK* (shown by the solid box) and *CLKDV* (shown by the dashed box) running at $f_{\text{samp}} = 24.576$ MHz and $\frac{f_{\text{samp}}}{2} = 12.288$ MHz respectively. While the global clock buffer distribution *CLK* was used to clock the PILOT matched filtering operation, the second clock buffer *CLKDV* was used to clock the RAKE matched filtering, channel estimation and Maximal Ratio Combining blocks. The presence of two independent clocking domains required the use of additional synchronising logic to transfer signals (such as the pilot soft matched filter output) from the *CLK* domain to *CLKDV* domain. Further, separate state machines were encoded in order to describe the control logic for operation of each of these domains.

Figure 8 Architecture of the reference DS-CDMA downlink RAKE receiver (see online version for colours)
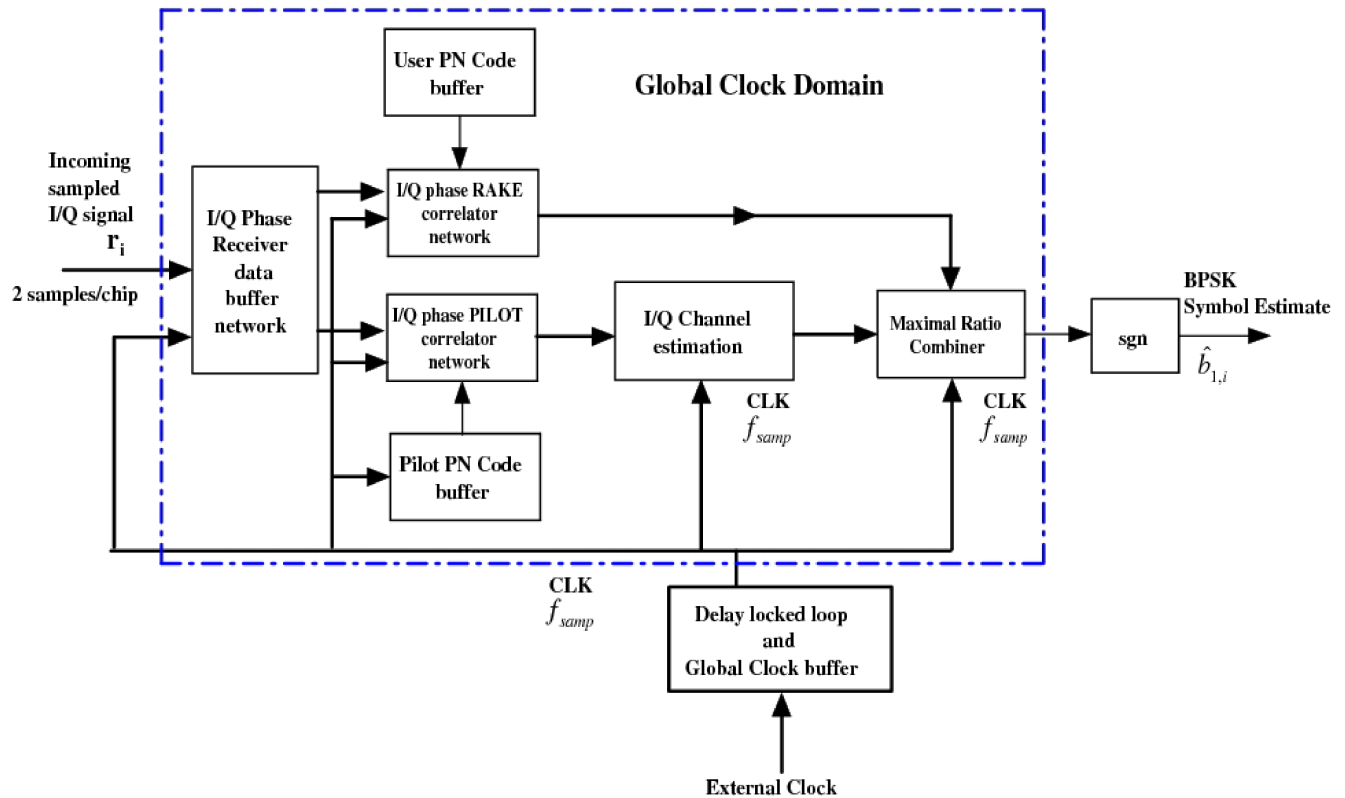
Figure 9 Architecture of the DS-CDMA downlink RAKE receiver with reduced complexity (see online version for colours)
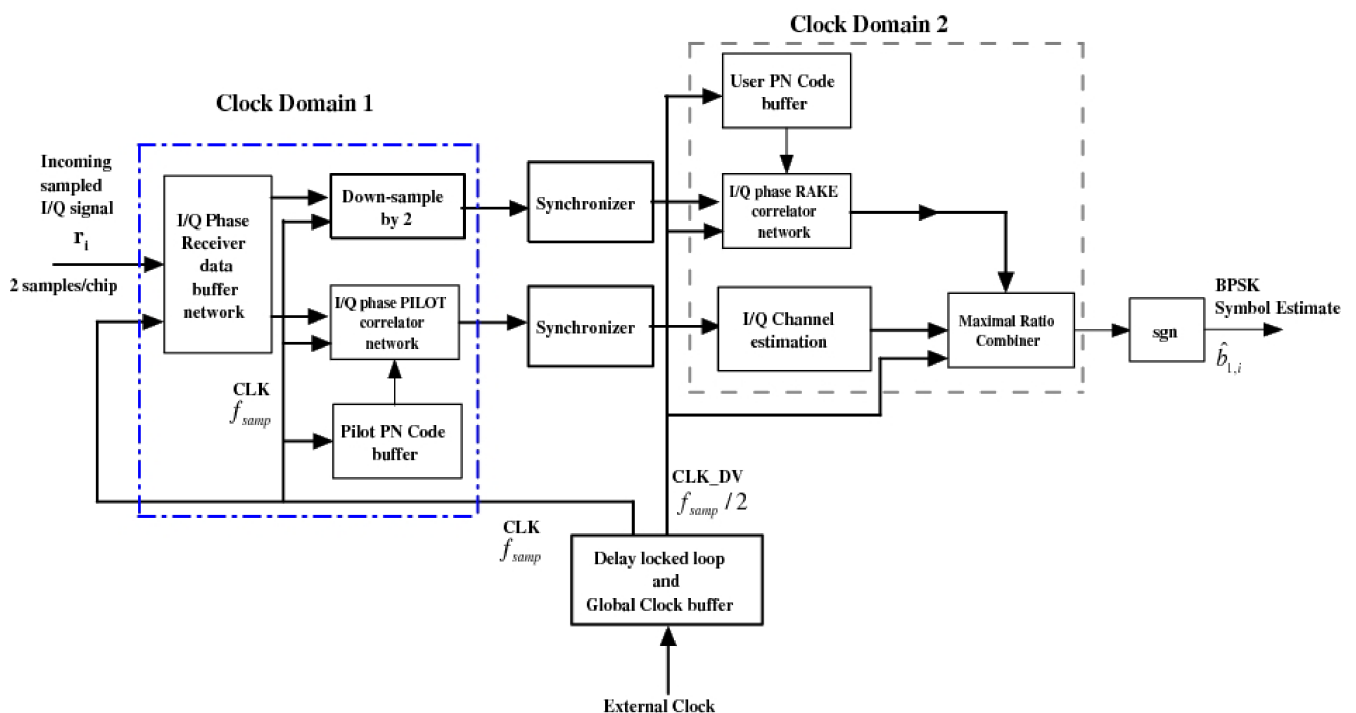
5 Results

For studying the impact of precision reduction on the resulting algorithm performance, the mobile receivers were simulated based on 10, 12, 14, 16 bit fixed-point word-length and compared with a floating point implementation. For each word-length format, the average received $SNR = 10\log_{10}(E_b/N_o)$ was varied to study the effect on the bit-error rate performance of the algorithm. In the computer simulations, five equal power users employing length 32 extended Gold sequences were considered. The scenario in consideration was a 5 user, 3 path correlated Rayleigh fading channel based on the Jakes mobility model. For each data-point, 40 random test cases of 5000 transmitted bits were tested. The multi-path delays were fixed for each simulation and varied from one simulation to the next. All the users were assigned unit transmit amplitudes. An additional code-multiplexed pilot channel with a three dB higher power was employed for channel estimation at the mobile receiver. The over-sampling rate at the transmitter and receiver front end was chosen to be two samples/chip in order to account for fractional multi-path delays. The A/D converter at the receiver front end was chosen to have an 8 bit width ($S8Q7$ format). We consider the performance of the following DS-CDMA RAKE receivers:

- Reference RAKE receiver performing truncated correlation sampled at 2 samples/ chip (Complexity 16NP – 2LP – 2P – 1 flops/symbol).
- Reduced arithmetic complexity RAKE receiver performing truncated correlation sampled at 1 sample/chip for detection and 2 samples/chip for channel estimation (Complexity = 12NP – 2LP – 2P –1 flops/symbol).

The performance of these receivers were compared against a DS-CDMA RAKE receiver employing perfect correlation (highest complexity of 16NP + 318P + 2LP – 1 flops/symbol.

5.1 Multi-user, multi-path fading channel

We describe the performance of the reference and reduced complexity RAKE receivers for a multi-path channel in the presence of interferers.

Figure 10 shows the performance of the reference DS-CDMA RAKE receiver for the above scenario. We notice that the receiver performance in fixed-point is close to the ideal floating point performance, with negligible performance degradation for the 10 bit precision (less than 1 dB loss) upto an SNR of ten dB.

Figure 10 Error probability reference RAKE (see online version for colours)
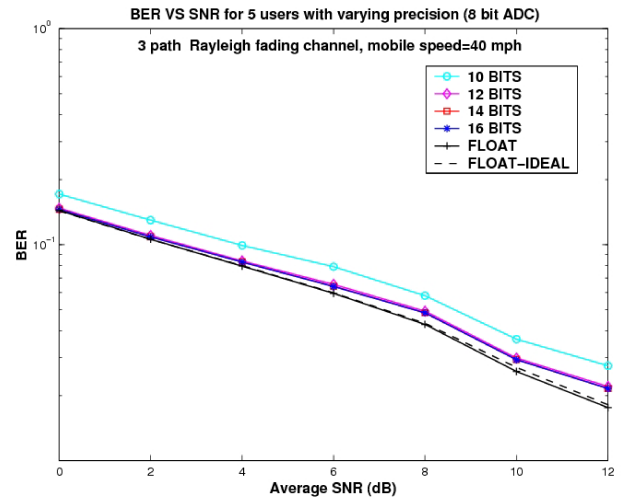

Figure 11 shows the performance of the reduced complexity DS-CDMA RAKE receiver. The reduction in complexity for reducing the dynamic power consumption, causes a performance degradation of two dB compared to the ideal DS-CDMA RAKE receiver employing ideal correlation (shown by the dashed line in black), owing to the reduced energy at the output of the RAKE correlator. We note that the receiver performance in fixed-point is almost identical with the floating-point performance up to a 10 bit precision.

Figure 11 Error probability: reduced complexity RAKE (see online version for colours)
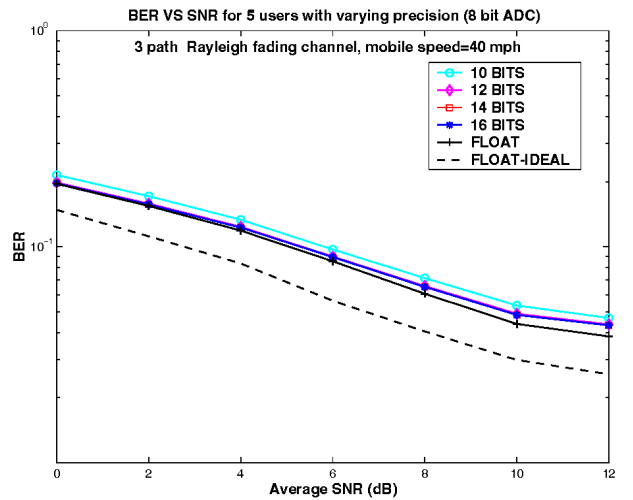
5.2 Results of FPGA implementation

Two different architectures for the RAKE receiver were targeted for a 2 million gate Virtex-II (XC2V2000 series)

FPGA, which employs a supply voltage of $V_{cc} = 1.5$ V. Synthesised complex receiver data for an urban three path Rayleigh multi-path channel was passed through each receiver implementation, and symbol detection was carried out. The results of each simulation were corroborated with the corresponding SystemC/MATLAB simulation to verify correctness of performance.

5.2.1 Timing simulation

For finding the dynamic power consumption in the design, the synthesised receiver data was run through the receiver. An external clock running at 50 MHz was produced to clock the receiver. The analysis was carried out following the synthesis, translation, mapping, netlist extraction, and the post-placement and routing phase. Extensive timing simulations were carried out in the Modelsim simulator to model true-device behaviour. All internal node transitions occurring during the course of the simulations were dumped into a '.vcd' (Value-Change-Dump) file format. The .vcd files were then analysed by the power analysis tool XPower in Xilinx, Inc. (2005a) provided for Xilinx FPGAs described in Xilinx, Inc. (2005b). A power report was generated as a result of the analysis that contained the overall power consumption, as well as a summary of the dominant power consumption among the individual blocks of the design. Finally, the dynamic power consumption was obtained after calculating the difference of the overall design power consumption and the quiescent power (225 mW) of the FPGA.[1]

In Table 4, the results of implementation of the reference and reduced complexity architectures for the DS-CDMA downlink RAKE receiver have been reported. The area shown in the table is represented in FGPA slices as well as the percentage occupancy in the FPGA, with the available area being 10752 slices in a Virtex-II FPGA. Considering *only* the effect of reduced precision, the reference architecture shows a power reduction of 25.6% for the 10 bit data-path compared to the 16 bit data-path. For the reduced complexity architecture, we observe power savings of 16.96% for the 10 bit data-path. These power savings are quite significant considering that the 10 bit data-path achieves almost close to the equivalent floating point performance for both the reference and reduced complexity receivers (performance loss being less than 1 dB).

Next, we consider the effect of complexity reduction on the resulting power savings. The 16 bit reduced complexity RAKE receiver achieves a power saving of 24.65% compared to the 16 bit reference RAKE receiver implementation. The combined effect of reduced precision and arithmetic complexity results in 37.4% reduction in dynamic power consumption for the 10 bit RAKE receiver, with a three dB degradation in performance (Figure 11). The tradeoff of dynamic

baseband power consumption with receiver performance is important for battery operated mobile wireless terminals. In scenarios where there is a strong received signal, then adaptive methods to reduce the dynamic digital baseband processing as proposed in this paper will greatly increase battery life.

Table 4 FPGA implementation complexity

Type	Bits	Area (slices)	P_{dyn} (mW)	Savings (%)
Reference architecture	16	3572 (33%)	109.5	–
	14	3000 (28%)	97.5	10.95
	12	2341 (22%)	93	15.06
	10	1844 (17%)	81.5	25.6
Reduced complexity architecture	16	3724 (35%)	82.5	24.65
	14	3134 (29%)	73	33.33
	12	2457 (23%)	68.5	37.44
	10	1942 (18%)	68.5	37.44

6 Conclusion

We have examined design methodologies and performance trade-offs for reducing the online power dissipation in a DS-CDMA mobile RAKE receiver. At the algorithm level, reduction in arithmetic complexity has been investigated for obtaining savings in the dynamic power dissipation. At the architectural level, precision reduction and activity rate reduction have been exploited for additional savings.

Reduction in precision shows that a 10 bit data-path achieves near floating point performance with minimal performance degradation for the reference RAKE receiver. Power-efficient architectures based on a Xilinx Virtex-II FPGA have been proposed for implementing both the conventional and reduced complexity DS-CDMA RAKE receiver. For a 16 bit data-path, we have observed a 24.65% reduction in dynamic power dissipation in the reduced complexity RAKE receiver compared to the reference RAKE receiver implementation, with an performance loss of less than 2 dB. The combined effect of reduced precision and complexity reduction leads to a 37.44% savings in digital baseband power consumption which will extend the operation of mobile wireless terminals.

References

Baghaie, R. and Laakso, T. (1991) 'Implementation of low power CDMA RAKE receivers using strength reduction transformation', *Proceedings of IEEE Vehicular Technology Conference*, Saint Louis, MO, pp.543–548.

Chandrakasan, A., Potkonjak, M., Mehra, R., Rabaey, J. and Brodersen, R. (1995) 'Optimizing power using transformations', *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, Vol. 14, No. 1, January, pp.12–31.

Chandrasekhar, V. (2002) *Reducing Dynamic Power Consumption in Next Generation DS-CDMA Mobile Communication Receivers*, Master's Thesis, Rice University, Available from www.ece.rice.edu/~ cvikram

Fantaccci, R. and Galligani, A. (1999) 'An efficient RAKE receiver architecture with pilot signal cancellation for downlink communications in DS-CDMA indoor wireless networks', *IEEE Transactions on Communications*, Vol. 47, No. 6, pp.823–827.

Garrett, D. and Stan, M. (1997) 'Power reduction techniques for a spread spectrum based correlator', *Proceedings of IEEE International Symposium on Low Power Electronics and Design*, Monterey, CA, 18–20 August, pp.225–230.

Latva-aho, M. and Juntti, M. (2000) 'LMMSE detection for DS-CDMA systems in fading channels', *IEEE Transactions on Communications*, Vol. 48, No. 2, pp.194–199.

Linebarger, D., Zeid, F.A. and Shrivastava, A. (2000) *Dynamic Range Tool, Signal Processing Lab, Engineering and Computer Science Department*, University of Texas, Dallas.

Ojanperä, T. and Prasad, R. (2001) *WCDMA: Towards IP Mobility and Mobile Internet*, Artech House Publications, Boston, MA.

Proakis, J. (1995) *Digital Communications*, McGraw-Hill, New York.

Rabaey, J. and Pedram, M. (1996) *Low Power Design Methodology*, Kluwer Academic Publishers, Boston, MA.

Rappaport, T. (1986) *Wireless Communications*, McGraw-Hill, New York.

Sriram, S., Brown, K. and Dabak, A. (1999) 'Low-power correlator architectures for wideband CDMA code acquisition', *Proceedings of IEEE 33rd Asilomar Conference Signals, Systems and Computers*, Pacific Grove, CA, Vol. 1, 24–27 October, pp.125–129.

Viterbi, A. (1995) *CDMA Principles of Spread Spectrum Communication*, Addison Wesley, Reading, MA.

Xilinx, Inc. (2005a) *FPGA Xpower Tutorial*, Available from http://support.xilinx.com

Xilinx, Inc. (2005b) *Xilinx FPGA Products*, Available from http://www.xilinx.com

Note

[1]The quiescent power (Q-Power) of a FPGA is fixed by the FPGA area, internal operating voltage and independent of the size of the design. The Virtex-II FPGA has a Q-Power specification of 225 mW at an operating voltage of $V_{ccint} = 1.5$ V.