On the Averaging Correlation for Satellite Acquisition in Software Defined Radio Receivers

Michael Soudan, Bernhard C. Geiger
Signal Processing & Speech Communication Laboratory, Graz University of Technology
Graz, Austria

BIOGRAPHY

Michael Soudan received his MSc in Electrical Engineering from the University of Applied Sciences, Munich, in 2006. Currently, he is pursuing the PhD degree in Electrical Engineering at the Signal Processing and Speech Communication Laboratory, Graz University of Technology. His main research interests include digital compensation techniques for mixed-signal and time-varying systems.

Bernhard C. Geiger received his MSc in Electrical Engineering from Graz University of Technology in November 2009. In September 2009 he joined the Signal Processing and Speech Communication Laboratory, Graz University of Technology, where he is currently pursuing his PhD as a Research and Teaching Associate. His main research interests are information theory and wireless communications.

ABSTRACT

This paper presents theoretical background to the averaging correlation used for satellite acquisition in global navigation satellite systems. To this end, a correlation model is derived and multirate theory is employed to outline the equivalence of the averaging correlation as presented in the literature and the common circular correlation. Moreover, the enhanced acquisition performance that averaging correlation entails is explored and demonstrated employing simulation data.

INTRODUCTION

The Global Positioning System (GPS) is the most widespread global navigation satellite system, offering position information for both military and civil applications. All transmitted satellite signals convey Binary Phase Shift Keying (BPSK) modulated data in the L1 and L2 bands, with center frequencies of 1575.42 MHz and 1227.6 MHz, respectively. The modulation of these signals with near orthogonal Pseudo Random Noise (PRN) spreading sequences allows different satellites to allocate the same frequency band and enables the receiver to separate individual satellite signals. In particular, the signal transmitted on L1 is modulated with two PRN sequences, the civil Coarse Acquisition (C/A) and the military Precision Code (P(Y)). In the following, only the non-classified C/A code with a chipping rate of f_c = 1.023 MHz, a code period of T_{per} = 1 ms, and a code length of N_c = 1023 chips per code period is considered.
As the receiver starts its operation, the received composite signal is screened for all possible PRN sequences. In order to detect the presence of a particular satellite signal, the phase of the respective C/A code sequence and the effective Doppler shift that is associated with this satellite signal must be determined, which can be done by iterative demodulation and correlation. This results in a computationally intensive search over all PRN sequences, code phases and Doppler frequencies during satellite acquisition. To enable (circular) correlation with the received signal, the locally generated C/A code must be resampled to the sampling rate f_s of the RF front-end. This sampling rate is typically much higher than the code chipping rate in order to enhance signal fidelity. The high sampling rate, on the other hand, implies a higher computational complexity for the circular correlation of the received signal with the local code. The authors in [1] introduced averaging correlation (AC) as a method to reduce the computational complexity, and it was analyzed in terms of acquisition performance in [2]. Yet, the equivalence between the AC and a complete (circular) correlation could not be found in the literature to the best knowledge of the authors. This paper intends to fill this gap by employing filter bank theory. Moreover, as it was already anticipated in [1], AC can yield coarse code phase estimates even if not all possible downsampling phases are computed, thus reducing the computational complexity even further. This work tries to provide quantitative statements about the acquisition performance for such cases of incomplete averaging correlations.

ACQUISITION

During the acquisition process, the coarse values of carrier frequency and code phase are determined for all visible satellites. The latter parameter indicates the

23rd International Technical Meeting of the Satellite Division of The Institute of Navigation, Portland, OR, September 21-24, 2010

circular shift of the locally generated C/A code with respect to the code screened for in the received signal. If the frequency of the locally generated carrier and the code phase of the local C/A code match with the frequency and code phase of the individual satellite signal, the correlation between the local and the received signal is maximal. The two-dimensional search over code phase and carrier frequency may thus result in a correlation function as shown in Fig. 1.

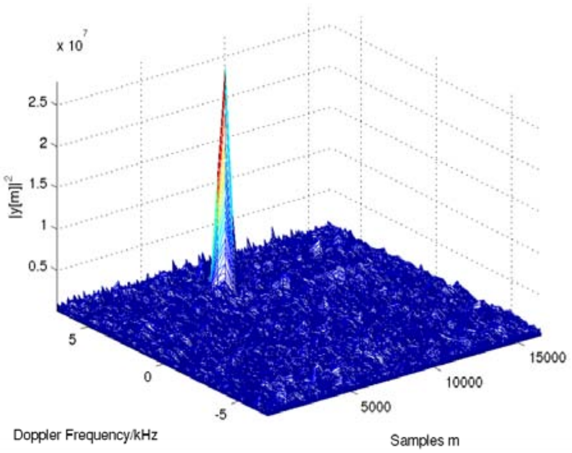

Fig. 1. Squared magnitude of the correlation function for 29 Doppler bins and N = 16368 code samples (corresponding to a sampling frequency of f_s = 16.3676 MHz).

For the remainder of this work it is assumed, that the signal $s[m]$ has been successfully demodulated to the complex baseband. This assumption, although hard to fulfill in practice, does not influence the validity of the general statements made in this paper. The simplification is justified by the fact that this paper focuses on the correlation part of acquisition, and due to the fact that any residual carrier can be incorporated by reducing the signal-to-noise ratio accordingly.

Since both the received and the locally generated signal are periodic with the same period, the linear correlation can be substituted by a circular correlation. In software defined radio receivers, an efficient way to calculate the circular correlation $y[m]$ is to multiply the Fast Fourier Transform (FFT) $S(k)$ of the signal $s[m]$ with the conjugate FFT $C^*(k)$ of the C/A code $c[m]$ and to transform the result back to the time domain with the inverse FFT [3], which gives

$$Y(k) = C^*(k) S(k). \qquad (1)$$

Note that the conjugate FFT of the C/A code is equivalent to the FFT of the time-reversed C/A code $c[-m]$, which lets us interpret correlation as a filter operation. Using the FFT, the correlation function can be evaluated for all possible code phases simultaneously.

The FFT has a computational complexity of the order of $N\log(N)$ for N samples to be transformed. In particular, given a certain sampling frequency, the number of samples calculates to

$$N = f_s T_{per} = M f_c T_{per} = M N_c \qquad (2)$$

where for the sake of simplicity it is assumed that the oversampling factor M is integer. In cases where this assumption does not hold exactly, it might be necessary to resample the received signal digitally to obtain an integer M [4]. However, as it will be shown, in cases where the oversampling factor is sufficiently close to an integer, the error resulting from this simplification is negligible.

AVERAGING CORRELATION

The idea of [1] was to average M consecutive samples of the received signal and perform the circular correlation with N/M samples instead with N samples. Since there are M possibilities to average M consecutive samples (differing in the choice of the first sample), M sub-correlations are required for AC [1,5]. Still, this way, the computational complexity reduced to an order of $M(N/M)\log(N/M) = N\log(N/M)$. One way to perform averaging is by means of decimation, i.e., moving average (MA) filtering and downsampling [2]. In the following section it will be shown that the sub-correlation of each of the M possible downsampling phases corresponds to particular samples of the complete correlation. To be specific, it will be shown that the sub-correlation of the p-th downsampling phase of the input signal is identical to the complete correlation where the output is downsampled with phase p.

CORRELATION MODEL

The complete correlation $y[m]$ of the signal $s[m]$ and the resampled C/A code $c[m]$ can be split into M sub-correlations, whose result $y_p[n]$ corresponds to a downsampling of $y[m]$ with downsampling phase p. As shown in Fig. 2, by means of multirate signal processing, these sub-correlations can be recombined to form $y[m]$. This approach was also pursued in [2], although the authors did not mention the mathematical equivalence between AC and complete correlation.

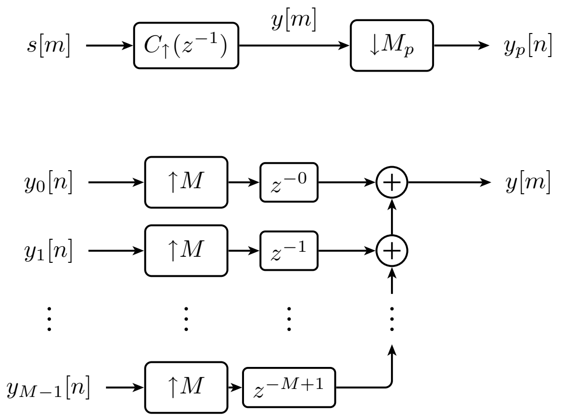

Fig. 2. Sub-correlation between resampled C/A code and received signal.

Let us first focus on the upper part of Fig. 2. Resampling the C/A code can be accomplished using the filter bank shown in Fig. 3. The time-reversed C/A code here given by the signal $c[-n]$ is upsampled by M in each branch k of the block, where $k = \{0,1, ..., M-1\}$. Each of these upsampled signals is delayed by a branch-specific delay and the resulting signals are summed. The equation in (3) outlines the structure of these signals, where each signal x_k is represented by a row of the matrix. Summing along the columns yields the resampled C/A code $c[-m]$ (the variable m indicates that the received signal is oversampled, contrary to the code $c[-n]$ which is clocked at the chipping rate). The symbol $S(z)$ represents a filter whose impulse response is the received signal $s[m]$. In other words, in the proposed model instead of filtering the received signal with a time-reversed version of the resampled code, the time-reversed code is filtered by the received signal. This exchange between filter impulse response and input signal can be done without loss of generality and will simplify the forthcoming derivations significantly. The output of the filter is downsampled by a factor of M and a downsampling phase p, in order to obtain a chipping rate version of the correlation function $y_p[n]$.

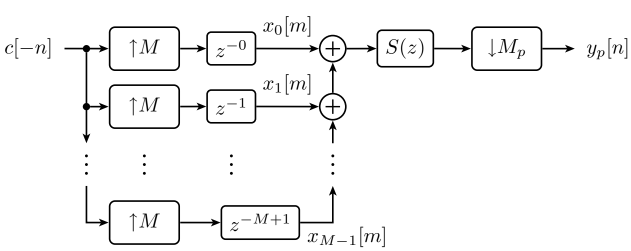

Fig. 3. Multirate system for resampling the C/A code.

$$
\begin{bmatrix} x_0 \\ x_1 \\ \vdots \\ x_{M-1} \end{bmatrix} = \begin{bmatrix} c[0] & 0 & \cdots & 0 & \cdots & c[N_c-1] & 0 & \cdots & 0 \\ 0 & c[0] & \cdots & 0 & \cdots & 0 & c[N_c-1] & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & c[0] & \cdots & 0 & 0 & \cdots & c[N_c-1] \end{bmatrix} \tag{3}
$$

By linearity, an equivalent structure is shown in Fig. 4 where the upsampled C/A code is processed by $S(z)$ before being delayed.

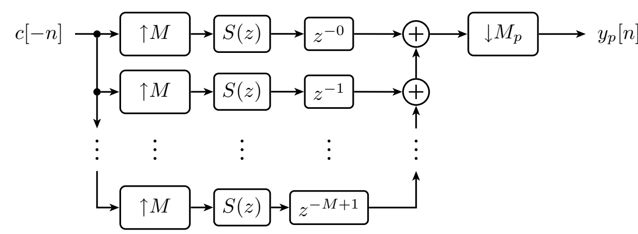

Fig. 4. Incorporating the signal filter in the multirate structure.

The upsampler followed by $S(z)$ occurs multiple times in the structure while providing identical outputs, so by combining all M paths the structure can be reduced as shown in Fig. 5.

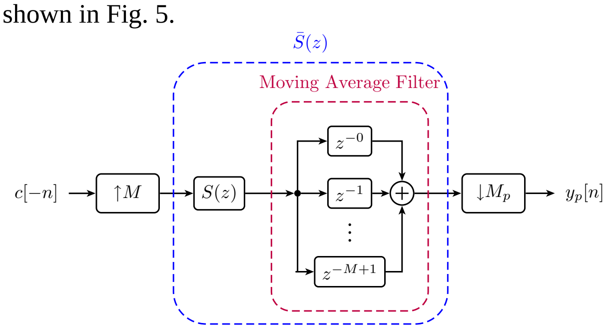

Fig. 5. Simplified system model for computing sub-correlations.

Now, the signal filter $S(z)$ and the moving average (MA) filter can be combined to a single filter $\overline{S}(z)$. This corresponds to filtering the received signal $s[m]$ by a MA filter and using the output, $\overline{s}[m]$, as an impulse response. By introducing the polyphase representation [6] of the upsampler followed by $\overline{S}(z)$, i.e.,

$$
\overline{S}(z) = \sum_{k=0}^{M-1} \overline{S}_k\left(z^M\right) z^{-(M-1-k)} \tag{4}
$$

the polyphase components of $\overline{S}(z)$ can be moved in front of the upsamplers and therefore, operate on the lower data rate (see Fig. 6). The impulse response of the filter $\overline{S}(z)$ is derived by MA-filtering the received signal $s[m]$ and by downsampling the output with a factor M and a downsampling phase k.

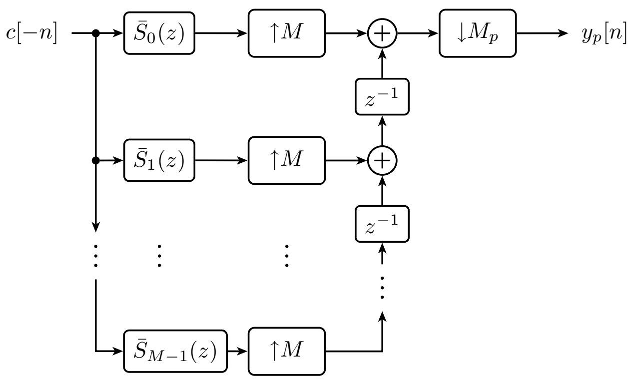

Fig. 6. Polyphase representation for computing sub-correlations.

The downsampler at the output with a phase $p = \{0,1,...,M\text{-}1\}$ deletes the influence of all filter paths except for the p-th. Consequently, upsampler and downsampler cancel and only a single filter remains, whose impulse response is identical to the received signal $s[m]$, which has been MA-filtered and downsampled by M with phase p (see Fig. 7).

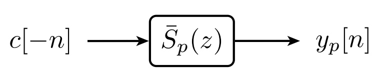

Fig. 7. Combining the downsampler with the polyphase representation.

Swapping now again input signal and impulse response as shown in Fig. 8, and representing the input signal by the operations required to obtain it, the proof of the equivalence is completed.

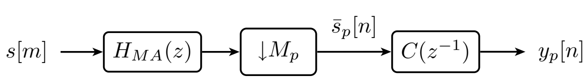

Fig. 8. Equivalent system model for computing sub-correlations.

Here, $H_{MA}(z)$ denotes the z-transform of the MA filter, which is given by

$$H_{MA}(z) = \sum_{k=0}^{M-1} z^{-k} \tag{5}$$

Comparing now Fig. 2 with Fig. 8, one can see that it is equivalent to downsample the complete correlation or to downsample the received signal and perform AC. Varying the downsampling phases p from 0 to $M\text{-}1$ and combining the outcomes of all paths as shown in the lower part Fig. 2 of yields the oversampled circular correlation as anticipated by [1], while reducing the computational complexity by $N\log(M)$.

INCOMPLETE AVERAGING CORRELATION

In [1], the authors briefly mention that, in case a coarse (chip rate) code phase estimate suffices, not all M possible downsampling phases have to be computed. To be precise, for a factor M = 5 it is stated that "one or two" phases may be sufficient to locate the correlation peak and, thus, the code phase. On the one hand, omitting the computation of a few phases leads to missing values of the complete correlation - consequently, there is a non-zero chance that the correlation peak is missing, which translates to a reduced acquisition performance. On the other hand, instead of M only K sub-correlations have to be computed. The complexity of the incomplete averaging correlation (IAC) can be reduced to N $(K/M)\log(N/M)$ with $K < M$.

In particular, let K be such that $M/K = L$ is an integer. In that case one obtains a downsampled version of the complete correlation by a factor of L. Fig. 9 illustrates this for $M = 16$ and $L = 4$ (the ideal correlation main lobe is downsampled by a factor of 4) where the peak does not show up in the IAC. In fact, by reducing the number of paths computed from M to K, the probability that the peak is preserved is only $K/M = 1/L$.

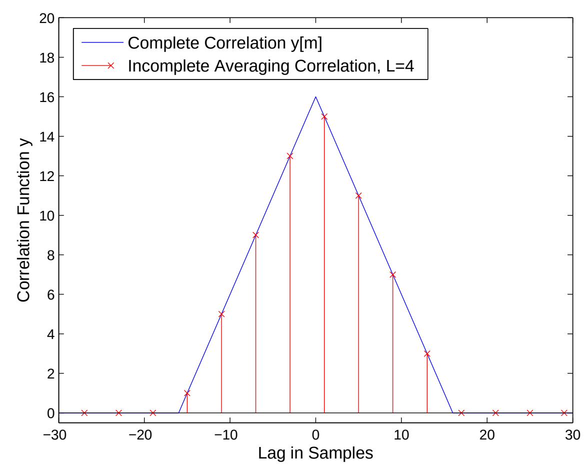

Fig. 9. Correlation main lobe for complete correlation and incomplete averaging correlation (M = 16, L = 4, K = 4).

PERFORMANCE ANALYSIS

To analyze the performance of the proposed correlation methods, a series of simulations was performed. To this end, 10 s of recorded satellite signals (available from [3]) where used for acquisition. The signals where recorded using an intermediate frequency of 4.1304 MHz and a sampling rate of 16.3676 MHz, resulting in a non-integer oversampling factor of M = 15.9996. For simplicity, the

23rd International Technical Meeting of the Satellite Division of
The Institute of Navigation, Portland, OR, September 21-24, 2010

digital signal was not resampled to an oversampling factor of $M = 16$, but the signal was used as is.
In order to exclude any influence of imperfect demodulations, all signals where first downconverted to baseband using a slightly adopted version of the tracking program from [3]. The baseband signals where then correlated according to the introduced methods.

The acquisition strategy adopted in this paper relies upon comparing the ratio between the first and the second largest correlation energy peak against a threshold [7]

$$R = \frac{\max|y[n]|^2}{\text{2nd}\max|y[n]|^2}\,. \tag{6}$$

If the ratio R exceeds the threshold, the satellite under consideration is assumed to be visible and the Doppler estimate used to demodulate the received signal is assumed correct. The code phase estimate, in that case, is retrieved from the position of the maximum of the correlation function.
For AC, M correlations of smaller length are computed, each corresponding to a certain downsampling phase. To determine the correct downsampling phase, i.e., the phase whose sub-correlation contains the correlation peak of the complete correlation, the authors in [1] propose to find the strongest among all sub-correlation peaks. For this particular sub-correlation, the ratio between its first (the complete correlation peak) and second maximum (a random noise sample) is computed.

It turns out (see Fig. 10), that the thus obtained ratio is never smaller than the ratio computed by dividing the first by the second maximum of the complete correlation function [1]. This is understood intuitively, since while the value of the largest peak is preserved, there is a high probability that the value of the second largest peak, which is due to noise, will be lower for shorter correlation functions. In the worst case (which occurs with probability $1/M$), both the correlation peak and the largest noise sample of the complete correlation function are downsampled with the same phase p, and the ratio remains the same. Thus, one can assume that the detection probability will increase significantly by employing AC. Although the false alarm probability will increase for a smaller number of samples in the correlation function, it has been shown recently (see [8]) that the increased detection performance outweighs this drawback when it comes to receiver operating characteristics (ROC; detection probability as a function of the false alarm probability).

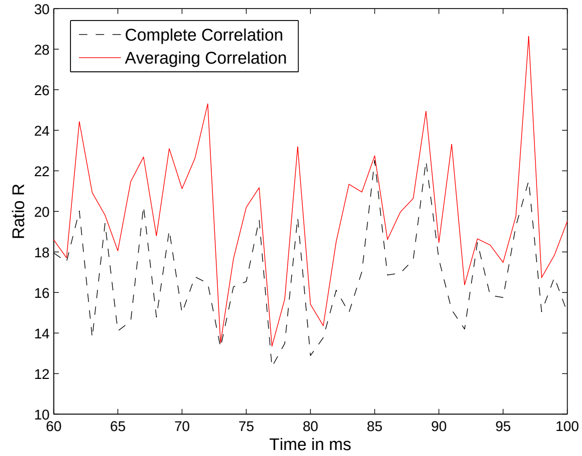

Fig. 10. Ratios R for complete and averaging correlation ($M = 16$). The strong satellite was considered.

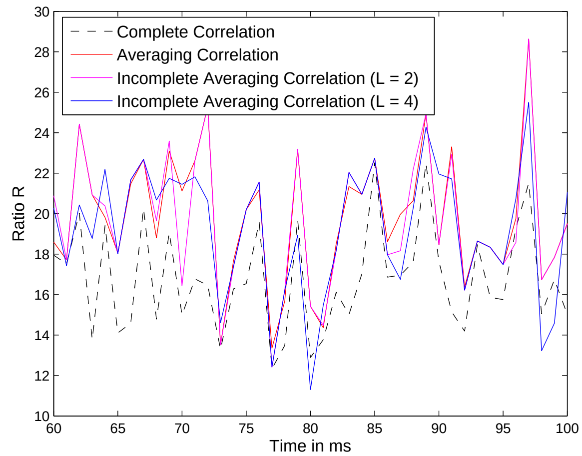

Fig. 11. Ratios R for complete, averaging, and incomplete averaging correlation ($M = 16$). The strong satellite was considered.

Surprisingly, by omitting certain downsampling phases, i.e., by computing only $K < M$ sub-correlations, the ratio between the largest and the second largest peak still very often exceeds the ratio obtained from the oversampled correlation function. Fig. 11 for example shows the ratios obtained for $L = 1$, 2, and 4. It can be seen, that there are cases where the ratio from IAC exceeds the ratio from AC. This can be explained by the fact that while the largest correlation peak still contains significant signal energy, the second largest correlation peak, which is due to noise, may be substantially smaller. As expected, the ratio is reduced for a decreasing number of computed downsampling phases, as it can be seen in Fig. 12.

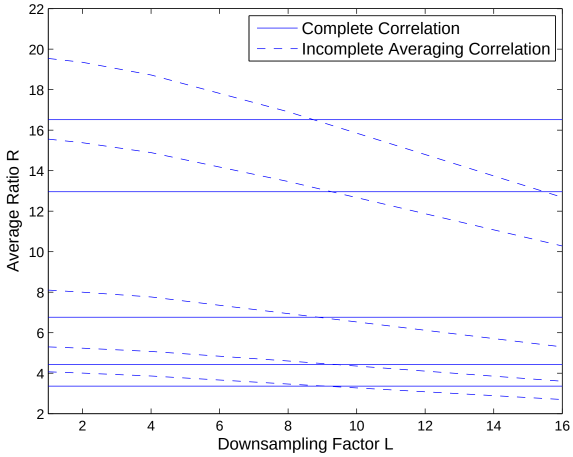

Fig. 12. Average ratio R depending on L. Different signal strengths are shown.

From Fig. 12, which shows the ratios R depending on different values of L for different signal strengths, one can see that on average the ratio from IAC will be higher than the ratio from complete correlation even when only $K = 2$ of 16 downsampling phases are computed (corresponding to $L = 8$). Thus, one can assume that the detection probability can be maintained while reducing the computational complexity from $N\log(N)$ to $2(N/M)\log(N/M)$.

By reducing the number of samples in the correlation function, the false alarm probability increases [8]. Thus, although the detection probability for IAC may remain unchanged compared to complete correlation, acquisition performance in terms of ROC may drop. Since AC outperforms complete correlation, it can be assumed that there will be a certain value of K required sub-correlations to obtain the same ROCs. Thus, there is a trade-off between computational complexity and acquisition performance. Quantifying this trade-off, e.g., by finding the appropriate value of K where the ROCs are identical, is within the scope of future work.

CONCLUSION

In this work the equivalence between averaging correlation and complete correlation is proved by means of multirate theory. In particular, it is shown that averaging the received signal and downsampling it with a certain phase prior to correlation is equivalent to performing the correlation at a high rate prior to downsampling with the same phase. While averaging correlation already allows a substantial reduction of computational complexity compared to complete correlation, this complexity can be decreased even further by omitting the correlation for certain downsampling phases.

A performance analysis of averaging correlation by means of the ratio between the largest and the second largest correlation energy peaks for each sub-correlation reveals that, on average, incomplete averaging correlation, which omits a significant amount of sub-correlations, outperforms complete correlation.

ACKNOWLEDGMENTS

This work was partially funded by the Austrian Research Promotion Agency under the project "SoftGNSS 2", project number 819682.

REFERENCES

[1] J. A. Starzyk, Z. Zhu, "Averaging Correlation for C/A Code Acquisition and Tracking in Frequency Domain", IEEE MWSCS, Fairborn, Ohio, 2001.

[2] M. Fantino, M. Pini, P. Mulassano, M. Nicola, and A. Nordio, "Signal Compression for an Efficient and Simplified GNSS Signal Parallel Acquisition", ION GNSS, Savannah, Georgia, 2008.

[3] K. Borre, D. M. Akos, N. Bertelsen, P. Rinder, and S.H. Jensen, "A Software-Defined GPS and Galileo Receiver: A Single Frequency Approach", Birkhäuser, 2007.

[4] G. Hamza, A. Zekry, and I. Motawie, "Implementation of a Complete GPS Receiver Using Simulink", IEEE Circuits and Systems Magazine, 9, 2009.

[5] M. Fantino, A. Molino, and M. Nicola, "An Acquisition Strategy Suitable for Software GNSS Receivers", ENC-GNSS, Naples, Italy, 2009.

[6] P. Vaidyanathan, "Multirate Systems and Filter Banks", Prentice Hall, Englewood Cliffs, 1993.

[7] J. Jung, "Receiver having a ratio-based signal acquisition method", U.S. Patent 7,161,977, Jan. 2007.

[8] B. C. Geiger, M. Soudan, and C. Vogel, "On the Detection Probability of Parallel Code Phase Search Algorithms in GPS Receivers", IEEE PIMRC, Istanbul, Turkey, 2010.