Trade-off Analysis of Robust Carrier Phase Tracking Techniques in Challenging Environments

Rui Sarnadas, Teresa Ferreira, *GMV*
Jordi Vilà-Valls, Gonzalo Seco-Granados, José A. López-Salcedo, *UAB*
Fernando D. Nunes, Fernando M. G. Sousa, *IT*
Paolo Crosta, Francesca Zanier, Roberto Prieto-Cerdeira, *ESA*

BIOGRAPHY

Rui Sarnadas received his degrees in Electrotechnical Engineering and Electronic Systems from Instituto Superior Técnico in Lisbon, Portugal. He joined GMV in 2008 and he is working on GNSS receiver technologies and signal processing.

Teresa Ferreira has been working in receiver and navigation-related technologies and applications since she joined GMV in 2004. Her degrees are in telecommunications and GNSS.

Jordi Vilà-Valls received the PhD degree in Electrical Engineering from the Grenoble INP (INPG), France, in 2010. Currently he holds a Postdoctoral Research Associate position at SPCOMNAV, UAB. His primary areas of interest include statistical signal processing, Bayesian estimation theory, and adaptive/robust filtering methods, with applications to positioning and tracking systems.

Gonzalo Seco-Granados received the M.Sc. and Ph.D. degrees in Telecommunication Engineering in 1996 and 2000, respectively, from Universitat Politècnica de Catalunya (UPC). He also received an MBA from IESE-University of Navarra, Barcelona, in 2002. From 2002 to 2005, he was member of the technical staff at the European Space Research and Technology Center (ESTEC), European Space Agency (ESA), Noordwijk, The Netherlands, involved in the Galileo project and leading the activities concerning indoor GNSS. Since 2006, he is Associate Professor at the Department of Telecomunications and Systems Engineering, Universitat Autònoma de Barcelona (UAB), and member of the SPCOMNAV group. He was a co-guest editor for a special issue of the IEEE Signal Processing Magazine. Since 2009, he is Director of the Chair of Knowledge and Technology Transfer "UAB Research Park–Santander".

José A. López-Salcedo received the M.Sc. and Ph.D. degrees in Telecommunication Engineering in 2001 and 2007, respectively, from Universitat Politècnica de Catalunya (UPC). From 2002 to 2006, he was a Research Assistant at UPC involved in R+D projects related with synchronization techniques for digital receivers, satellite communications and iterative decoding techniques for MIMO wireless systems, both for private industry and public administrations. Since 2006, he is Assistant Professor at the Department of Telecomunications and Systems Engineering, Universitat Autònoma de Barcelona (UAB), and member of the SPCOMNAV group.

Fernando D. Nunes received his E.E., M.Sc. and Ph.D. degrees in Electrical Engineering all from the Instituto Superior Técnico (IST), Technical University of Lisbon, Portugal. He is currently an Assistant Professor at the Department of Electrical and Computer Engineering, IST, and researcher with the Communication Theory group at Instituto de Telecomunicações. His main research interests include communication theory and signal processing in mobile radio communications and navigation systems.

Fernando M. G. Sousa received his B.Sc. in Electronics and Telecommunications Engineering from Instituto Superior de Engenharia de Lisboa (ISEL) and his E.E. and M.Sc. degrees in Electrical Engineering from the Instituto Superior Técnico, Technical University of Lisbon, Portugal. He is currently a coordinator professor at the Electronics, Telecommunications and Computer Engineering Department, ISEL, and researcher with the Communication Theory group at Instituto de Telecomunicações. His main research interests include communication theory and signal processing in mobile radio communications and navigation systems.

Paolo Crosta is a Radio Navigation System Engineer in the Radio Navigation Systems and Techniques Section at the European Space Agency Technical Directorate where

he provides support to the EGNOS and Galileo ESA programs. He received the M.S. degree in Telecommunications Engineering from the University of Pisa, Italy in 2002 and then he worked for several years as signal processing expert in Thales Alenia Space Italy.

Francesca Zanier received the M.E. (cum laude) in Telecommunications Engineering and the PhD degree in Information Engineering from the University of Pisa (Italy) in 2004 and 2009, respectively. Since 2009, she has been with the Radionavigation section at the European Space Agency (ESA/ESTEC), Noordwijk, The Netherlands working mainly on GNSS signal processing.

Roberto Prieto-Cerdeira received his Telecommunications Engineering degree from the University of Vigo, Spain and followed postgraduate studies on Space Science and Radioastronomy in Chalmers University of Technology (Gothenburg, Sweden). Since 2004, he has been with the Wave Interaction and Propagation Section in the European Space Agency, ESA/ESTEC, The Netherlands where he is responsible of the activities related to radiowave propagation in the ionosphere and local environment for Global Navigation Satellite Systems (GNSS) and Satellite Mobile Communications. He actively participates in ITU-R Study Group 3 and the SBAS-Iono group. He is a member of IEEE and the Institute of Navigation (ION).

ABSTRACT

A rather benign propagation environment is needed to reliably exploit the benefits of the carrier phase observables, since carrier phase tracking techniques are rather sensitive to fast fading phenomena (e.g. coming from ionospheric scintillations, high dynamics or fading) that can lead to the occurrence of cycle slips or even loss of lock.

Thus, the interest of techniques to provide robust carrier phase measurements in challenging environments becomes apparent since that would make it possible to extend advanced applications to a new set of scenarios, such as environments affected by severe scintillations or in high dynamics.

In this study, three different families of techniques are explored for robust carrier phase tracking: advanced closed loop and open loop scalar techniques, and vector tracking techniques.

The performance of these techniques was assessed under harsh environments, including fading, ionospheric scintillations and high dynamics. In this sense, both the underlying technique implementation details and the tailoring for different environments are discussed, as well as a preliminary trade-off analysis. The main driver is to analyze each family of techniques in different propagation environments, so as to enable the potential for robust and accurate carrier phase tracking under difficult conditions.

The results shown depict the different realms of applicability and performance achieved for the tracking techniques, taking into account different signal impairments and propagation environment. It can be seen that better carrier phase tracking performance can be achieved with such implementations, and the inherent trade-offs and tailoring are discussed.

1. INTRODUCTION

Carrier phase tracking techniques are typically used in GNSS applications for which the best solution performances are required. Such applications include precise ranging (e.g. using PPP to reach accuracies as low as a few centimeters), propagation characterization (e.g. in order to characterize atmospheric phenomena), data demodulation and robust tracking.

This study aims at assessing the performance of innovative carrier phase tracking techniques in challenging environments in order to understand their limits, analyze their performance and assess the improvement with respect to traditional techniques. For this purpose, three different families of techniques are explored for carrier phase tracking: closed loop, open loop and vector tracking techniques.

The advanced closed-loop techniques considered herein are mostly Kalman Filter (KF) based solutions, which provide the best performance and robustness under harsh environments since they are able to adapt to the current working conditions through a systematic formulation and a flexible architecture. On their side, open-loop techniques are able to cope with the presence of outliers that may cause loss of lock in traditional closed-loop receivers due to their batch-processing architecture. Vector architectures essentially exploit a deeper level of integration between the signal processing and the navigation units, allowing the receiver to continue tracking a signal with very low power - even if with degraded performance - avoiding loss of lock. Vector tracking techniques are expected to provide robust solutions in difficult environments.

The performance of the techniques is assessed via simulation focusing on high dynamic and ionospheric scintillation environments and the results are compared to that of classical techniques.

First the simulation environment is described and then the main findings are presented leading to the trade-off analysis among the techniques and the applicable environments/scenarios.

2. PROPAGATION ENVIRONMENT

In the scope of this study, harsh environments for carrier phase tracking are considered as environments where the signal suffers: ionospheric scintillations and receiver high dynamics. In addition to the effects described hereafter, Additive White Gaussian Noise (AWGN) is also added in order to push the techniques to work near the tracking threshold and to assess their sensitivity.

A. Ionospheric Scintillations

Ionospheric scintillations are fast-fading phenomena observed on GNSS signals propagating through the ionosphere, typically occurring in equatorial regions during active periods of the solar maximum (it is also common in high latitudes but not considered herein). As ionospheric scintillations affect both the magnitude and the phase of the incoming signal, the phase may suffer disturbances, typically synchronized with a deep fade, that could lead to the occurrence of cycle slips (phase jumps), errors in the data demodulation and ultimately loss of lock.

This study considers simulated scintillations in the form of complex time-series generated using the Global Ionospheric Scintillation Model (GISM) as described in ITU-R-531-11 [ITU12]. The GISM is based on a multiple phase screen technique and it is driven by an electron density climatological model (NeQuick) underneath. The ionospheric scintillations are characterized by the intensity fluctuations given by the scintillation index, S4 and the standard deviation of the phase scintillations.

GISM is able to generate complex time-series of ionospheric scintillations for a given geometry, location, time, date and solar flux. For convenience of analysis, two time-series have been considered in this work: one for moderate scintillations (S4 = 0.5) and one for strong scintillations (S4 = 0.9) with an artificially long duration (3000 seconds each) and sampled at 10 ms (it is considered that scintillation time-series have relevant spectral components up to 10-25 Hz).

Note that such scintillations can be viewed as a complex phasor that travels in a complex plane at a speed that is the inverse of the channel decorrelation time. Figure 1 depicts the time series used in this study for both moderate and strong scintillations over time. While for moderate scintillations the phasor is always on the right hand plane and therefore there are no phase jumps, for the strong scintillations the phasor moves all over the plane. In fact, when a deep fading occurs, the complex vector is shortened and goes near the origin; this fact implies that small changes in the phase when it is near the origin may lead to phase jumps at the receiver output.

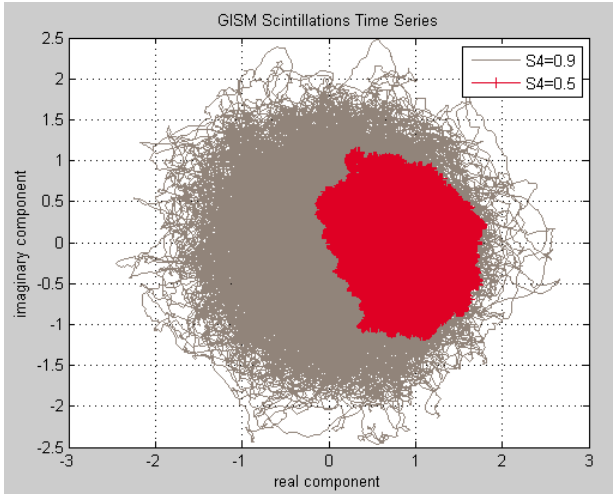

Figure 1: GISM time series for Moderate (S4=0.5) and Strong (S4=0.9) Scintillations.

B. High Dynamics

Carrier phase tracking loops are quite sensitive to high dynamics and their variations, which cause different problems to classical carrier phase tracking structures.

On a first level, the order of the receiver filters determines their capability to track input phase dynamics: whereas most mass-market GNSS receivers are equipped with first-order loop filters, which allow tracking of input phase values affected by constant frequency offsets, advanced receivers typically implement second order loop filters which allow tracking of frequency drifts. Also, the choice of the order of the filter drives not only its cost but also its bandwidth. If, on one hand, large bandwidths could allow tracking higher dynamics, on the other hand they also increase the carrier phase jitter.

High dynamics are simulated through velocity, acceleration and jerk user profiles, which is where carrier tracking techniques are more prone to loss of lock. Such profiles are then applied to the signals as line-of-sight (LoS) projections of the user's dynamics, resulting in a maximum Doppler frequency perceived by the receiver given by:

$$f_{Doppler} = v_{max} * f_0/c \qquad [1]$$

where v_{max} is the maximum relative velocity vector along the LoS of the satellite, f_0 is the carrier frequency of the signal (1176.45 MHz for L5 and 1575.42 MHz for L1), and c is the speed of light.

For a static receiver, v_{max} can be considered as 800 m/s [Kap06] leading to Doppler ranges in the order of ±4.2 kHz. It would take a receiver with a velocity of around 150 m/s in the direction of the satellite to increase the

Doppler range to about ±5 kHz, which would be the equivalent to the speed of an aeronautical user.

The maximum Doppler rate for a static user was derived in [Tsu00] and computed to be around 0.178 m/s^2. As a consequence, the main driver for the Doppler rate (and for the variation of the Doppler rate) is actually the user dynamics: in fact, a user with an acceleration of 1 g (gravitational acceleration of 9.8 m/s^2) in the LoS direction leads to a Doppler rate of about 52 Hz/s.

The rate of change of the Doppler rate will also be driven by the user dynamics, since the derivative of the Doppler rate caused by the satellite-only motion is very smooth, as shown in [Tsu00]. Therefore, a maximum jerk of around 52 Hz/s^2 can be considered for a maximum jerk of the user dynamics of 1 g/s. The expected Doppler-related orders of magnitude (worst case) for this study are depicted in Table 1. Note that velocity, acceleration and jerk are only LoS components, and that Doppler frequencies are for L1 center frequency.

Table 1: Dynamic profiles and impact on Doppler frequency

User Type	Velocity / Doppler	Acceleration / Doppler rate	Jerk / Doppler accel.
Vehicular	50 m/s 262,55 Hz	10 m/s^2 (1 g) 52,5 Hz/s	1 m/s^3 (0.1 g/s) 5,25 Hz/s^2
Aeronautical	500 m/s 2625,7 Hz	20 m/s^2 (2 g) 105 Hz/s	10 m/s^3 (1 g/s) 52,5 Hz/s^2

It is important to state that upon entering tracking, the Doppler frequency changes caused by the satellite's motion are considered negligible. Therefore, the Doppler frequency values shown in Table 1 do not take into account the 800 m/s of initial Doppler, nor the Doppler rate and acceleration from the satellite-only components, considered negligible when compared to the user's high dynamics variations which have an impact on the tracking techniques.

3. SCALAR TRACKING TECHNIQUES

This section provides a brief overview on the state-of-the-art for both closed-loop and open-loop scalar tracking techniques, together with a GNSS and state-space signal model definition.

A. GNSS Signal Model

The baseband analytic representation of a signal received from a generic GNSS satellite can be expressed as:

$$x(t) = P_x(t)\, d\big(t-\tau(t)\big)c\big(t-\tau(t)\big)e^{j(2\pi f_d(t)+\theta_e(t))} + w(t) \qquad [2]$$

where $P_x(t)$, $d(t)$, $c(t)$ and $w(t)$, stand respectively for the signal amplitude, the navigation message, the spreading code and the noise term, which may include the thermal noise, signals from the same or other satellites, replicas of the transmitted signal due to multipath, and any other interference. The synchronization parameters are the code delay $\tau(t)$, the carrier Doppler frequency shift $f_d(t)$, and the carrier phase $\theta_e(t)$. The digitized signal at the output of the radio frequency (RF) front-end feeds the digital receiver M channels. The goal of each channel is to acquire and track the signal of a single satellite.

Taking into account the problem at hand, which is to focus on the carrier phase tracking problem, a simplified (i.e. semi-analytic) signal model will be adopted herein by working with the complex samples of the prompt correlator output, which can be modeled as follows:

$$r_k = \alpha_k e^{j\theta_k} + n_k \qquad [3]$$

where each sample is taken at a time instant $t = kT_s$ with T_s the coherent correlation (i.e. pre-detection) integration time. The amplitude, α_k, may include possible amplitude fading, and the carrier phase, θ_k, includes both the phase variations due to the receiver's dynamics and phase variations due to propagation disturbances (e.g., scintillation). For simplicity, we will assume perfect code synchronization and that no data bits are present either because data wipe-off has already been implemented or because the received signal corresponds to a pilot component.

B. Advanced Closed-Loop Techniques

State-of-the-Art

The carrier phase tracking techniques implemented in conventional GNSS receivers rely on well-known phase-locked loop (PLL) architectures [Kap06]. The problem of standard PLLs is the existing tradeoff between noise reduction and dynamic range, which is driven by the bandwidth and order of the loop. These two parameters are the ultimate factors that control the phase tracking jitter and stress error: a small bandwidth is needed to be able to filter out the noise and track signals with low carrier to noise ratios (C/N0), and a large bandwidth has to be used to cope with high dynamics (i.e. fast variations on the parameters of interest). Traditional techniques have been shown to deliver poor estimates or even fail under harsh propagation conditions [Lia04] [Zha09].

On the basis of conventional PLL architectures, some improvements have been proposed in the literature: hybrid architectures coupling the PLL with a frequency-locked loop (FLL) to reduce the dynamics of the signal to be tracked [Mao10], then being able to use a smaller bandwidth; wavelet denoising techniques to reduce the noise affecting the system [Lia04], then being able to use large bandwidths to track high dynamics; adaptive

methods that sequentially adapt the bandwidth of the system according to the actual working conditions, based on the estimation of some performance metrics [Sko05] [Mao08], or Kalman filter (KF)-based tracking techniques [Psi07], where the filter is automatically adjusted so as to minimize the mean square error.

Standard KFs, which are formulated from an optimal filtering approach, have been shown to outperform all methods based on PLL architectures previously mentioned [Lia04] [Zha09] [Hum05] [Yu06], and that is why these methods are in the core of all the advanced carrier phase tracking techniques. A further improvement of KF-based methods is the so-called Adaptive KF (AKF) [Hu03], which sequentially adapts the filter parameters (e.g., the covariance matrix of the measurement noises) to the actual working conditions [Zha10] [Won12a]. In the present contribution, a performance assessment between standard PLLs and both standard and adaptive KFs is provided.

Kalman Filter State-Space Formulation

In the carrier phase tracking problem, considering that the input to the tracking block is given by equation [3], the parameter of interest is the phase θ_k, which includes the time-varying evolution caused by the receiver dynamics and the possible disturbance effects.

Concerning the phase evolution due to the receiver dynamics, the following model is considered:

$$\theta_k = \theta_0 + 2\pi\left(f_{d,k}kT_s + \frac{1}{2}\dot{f}_{d,k}k^2T_s^2\right) \quad [4]$$

where θ_0(rad) is a random constant phase value, $f_{d,k}$ (Hz) is the carrier Doppler frequency shift and $\dot{f}_{d,k}$ (Hz/s) the Doppler frequency rate (i.e., the Doppler frequency shift dynamics). Note that this model considers a second order approximation of the phase evolution, thus neglecting possible higher order terms. Considering this phase evolution the state to be tracked is defined as:

$$x_k = \left[\theta_k, f_{d,k}, \dot{f}_{d,k}\right]^T \quad [5]$$

with the phase expressed in cycles and the Doppler shift in cycles/sample. The process equation is:

$$x_k = \begin{pmatrix} 1 & 1 & 1/2 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{pmatrix} x_{k-1} + v_k \quad [6]$$

where $v_k \sim N(0, \boldsymbol{Q})$ is the process noise with covariance matrix $\boldsymbol{Q}$, representing possible uncertainties or errors on the state transition model. This process noise covariance matrix is designed according to the problem at hand and depending on the system working conditions. The general expression is:

$$\boldsymbol{Q} = \boldsymbol{G}\boldsymbol{G}^T\sigma_j^2 \quad [7]$$

where $\boldsymbol{G} = \left[\frac{1}{6}, \frac{1}{2}, 1\right]^T$ and σ_j^2 is the Doppler frequency rate error variance. The previous equations define the state-space formulation of the model. To complete the model definition, the Gaussian measurement noise variance is computed from the C/N0 (i.e., variance of the phase noise at the output of the discriminator) as:

$$\sigma_n^2 = \frac{1}{8\pi^2\, C/N0\, T_s}\left(1 + \frac{1}{2\, C/N0\, T_s}\right) \quad [8]$$

C. Open-Loop Techniques

The continuous-mode transmission of GNSS signals has a strong influence on the receiver implementation, which has typically been addressed through the adoption of closed-loop (i.e. feedback) architectures that keep track of the input signal by driving an error signal to zero. Closed-loop schemes are thus the natural choice for positioning receivers, where precise tracking of the carrier phase and code delay is of paramount importance. On the contrary, open-loop (i.e. feedforward) architectures have typically been relegated to communication receivers, particularly those based on burst-mode transmissions, where a batch of L signal samples (i.e. a snapshot) is processed at a time for estimating the required synchronization parameters.

So far, both approaches have remained apart due to their different fields of application, and open-loop techniques have remained outside the realm of GNSS systems. This trend, however, started to change with the advent of high-sensitivity GNSS (HS-GNSS) receivers, which are specifically tailored to allow operation in harsh working conditions where traditional GNSS receivers fail [Sec12]. The need for open-loop architectures is motivated by the tight constraints of traditional closed-loop schemes, where the noise filtering and the dynamics tracking capabilities are often at odds with each other, thus making very difficult to succeed in carrier tracking in non-nominal working conditions.

Most open-loop carrier tracking techniques are indeed adaptations of carrier frequency estimation methods. Thus, their output is not an error signal that must be driven to zero but the actual estimate of the carrier frequency (and higher order moments, if considered). If we are interested in the carrier phase observables, we need first to compensate the snapshot samples with the carrier dynamics that have been estimated for that snapshot. This is done by correlating the samples with a complex exponential local replica whose frequency content is the one that has been estimated. Then, the result of this correlation is a complex scalar whose argument provides the phase estimate $\hat{\theta}_n$ that corresponds to this n^{th} batch of samples.

In the sequel, we will classify open-loop techniques in two main categories, namely optimal maximum likelihood (ML) techniques and practical autocorrelation-based techniques. For the latter, we will use the unbiased autocorrelation of the snapshot samples, which is defined as $\hat{R}_r(m) = \frac{1}{L-m}\sum_{k=0}^{L-1} r_k r_{k-m}^*$ for $1 \le m \le L-1$.

Maximum Likelihood Carrier Tracking

Let us assume that we are interested in the estimation of the carrier frequency f_0, carrier frequency rate f_1 and carrier frequency acceleration f_2, in order to accurately track the carrier dynamics. In that case, ML open-loop carrier tracking involves a multi-dimensional grid search whereby the carrier estimates are obtained as:

$$\{\hat{f}_0, \hat{f}_1, \hat{f}_2\} = \arg\max_{f_0,f_1,f_2} \left|\sum_{k=0}^{L-1} z_k(f_1,f_2) e^{-j2\pi f_0 k T_s}\right|^2 \text{[9]}$$

where $z_k(f_1,f_2) \doteq r_k e^{-jh_k(f_1,f_2)}$ and the trial function defined as $h_k(f_1,f_2) \doteq 2\pi\left(\frac{f_1}{2}k^2T_s^2 + \frac{f_2}{6}k^3T_s^3\right)$ [Vil89]. Note that if one is interested on the carrier frequency only (i.e. forcing $f_1 = f_2 = 0$), the cost function in [9] becomes the well-known periodogram.

Autocorrelation-based Techniques

We can find several techniques within this category such as the Fitz method [Fit91], whereby the carrier frequency estimate is obtained as $\hat{f}_0 = \frac{T_s}{\pi N(N+1)}\sum_{m=1}^{N} \arg\{\hat{R}_r(m)\}$ and with the optimal performance being achieved for $N = L/2$. A similar carrier frequency estimator is the Luise&Reggiannini (L&R) method [Lui95], whose estimate is given by $\hat{f}_0 = \frac{T_s}{\pi(N+1)}\arg\{\sum_{m=1}^{N}\hat{R}_r(m)\}$. As can be seen, the complexity of these two estimators is quite low, compared to the one required by the ML solution. Nevertheless, this comes at the expense of a reduction in the visibility margin of these two algorithms, since they are only able to resolve carrier frequencies such that $|f_0| \le 1/(2NT_s)$ for the Fitz method and $|f_0| \le 1/(NT_s)$ for the L&R method. In practice, N may be a large number, and this restricts these techniques to low-dynamics or static applications. For instance, for a typical closed-loop bandwidth of $B_L = 5$ Hz, we would have an equivalent behavior with an open-loop snapshot of $L = 100$, according to the relationship $B_L \approx 1/(2LT_s)$ and using a pre-detection integration time of $T_s = 1$ ms. [Men97]. With this snapshot length, the frequency visibility margin would be of just 10 and 20 Hz for the Fitz and the L&R methods, respectively.

Within this category of autocorrelation-based carrier frequency estimators, it is worthwhile to mention Kay's estimator, which has been widely adopted in the context of digital communication receivers, whose frequency estimate is given by $\hat{f}_0 = \frac{T_s}{2\pi}\sum_{k=1}^{L-1}\gamma_k \arg\{r_k r_{k-1}^*\}$ with $\gamma_k = \frac{3L}{2(L^2-1)}\left[1-\left(\frac{2k-L}{L}\right)^2\right]$ a window function [Kay89]. A generalization of this method was proposed in [Rib98], which allows the estimation of higher order terms of the carrier dynamics.

Linear Prediction Techniques

Finally, it is worthwhile to mention the use of linear prediction techniques for carrier tracking applications. These techniques have been reported to provide some advantages when estimating the frequency of a narrowband, rapidly time-varying signal, and they overcome the frequency resolution limitations of maximum likelihood methods. In this work, we have adopted the modified forward-backward linear prediction (MFBLP) method, whose frequency estimate is given by $\hat{f}_0 = \arg\max_f \left|\hat{H}_{LP}(e^{j2\pi f})\right|$ where $\hat{H}_{LP}(e^{j2\pi f})$ is the Fourier transform of the MFBLP linear prediction filter [Tuf82].

4. VECTOR TRACKING ARCHITECTURES

Among the primary advantages of vector tracking architectures we can mention the following: the effect of the noise is reduced in all channels making them less likely to enter nonlinear tracking regions; it can operate with momentary blockage of one or more satellites; and it can be better optimized than the scalar tracking approach. Vector tracking is also able to improve tracking in weak-signal or jamming environments, especially when integrated with inertial sensors. The primary drawback is that all satellites are intimately related, and any error in one channel can potentially adversely affect other channels [Pet08] [Las09].

For the vector architectures we consider two implementations: vector delay/frequency locked loop (VDFLL) and vector delay/phase locked loop (VDPLL). Both structures use an extended Kalman filter (EKF) to predict user's trajectory and clock error. In both implementations the state vector has usually eight components (position and velocity or PV model) but can be extended to eleven components if the estimation of acceleration is also required (position, velocity and acceleration or PVA model). In the former case the EKF estimates the X, Y and Z components in ECEF coordinates of the receiver's position and velocity, besides the clock bias and drift. In the latter case the three components of the acceleration are added.

The discrete-time dynamics model adopted for each user's motion component in the case of the PV model is:

$$\begin{bmatrix} x_1(k+1) \\ x_2(k+1) \end{bmatrix} = \begin{bmatrix} 1 & \Delta t \\ 0 & 1 \end{bmatrix} \begin{bmatrix} x_1(k) \\ x_2(k) \end{bmatrix} + \boldsymbol{w}(k) \qquad [10]$$

with Δt denoting the update interval. The noise covariance matrix is:

$$E\{\boldsymbol{w}(k)\boldsymbol{w}^T(k)\} = q_m \Delta t \begin{bmatrix} \frac{\Delta t^2}{3} & \frac{\Delta t}{2} \\ \frac{\Delta t}{2} & 1 \end{bmatrix} \quad [11]$$

with parameter q_m being selected according to the user's motion and channel characteristics. In both architectures we use the normalized non-coherent early-late power code delay discriminator and the arctangent carrier phase discriminator. In addition, the VDFLL uses non-coherent normalized frequency discriminators.

The basic structure of the code/frequency loops of our VDFLL is sketched in Figure 2 where the code and frequency loops are coupled through the navigation solution. This is a non-federated architecture where the code and frequency loops are updated at the same rate (typically 100 Hz).

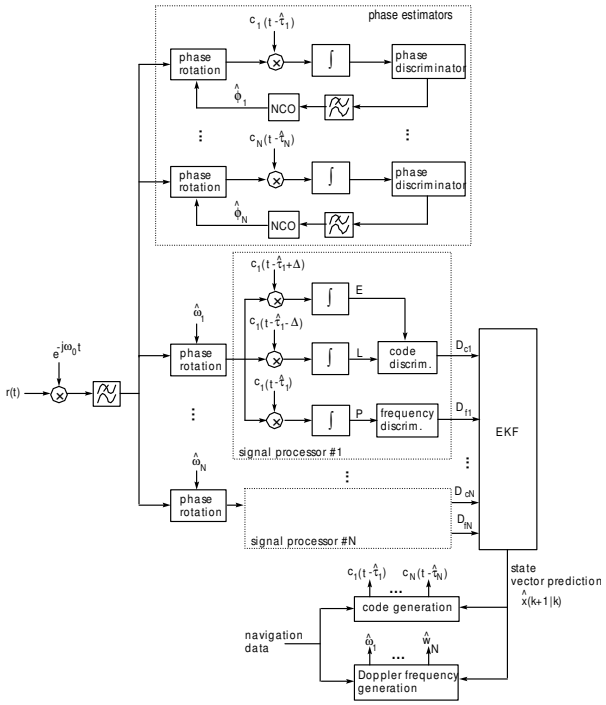

Figure 2: Block diagram of the VDFLL architecture.

Note that the carrier phase data is not used to update the navigation filter (EKF). The reason for that is because, in order to track the carrier phase using the position, the position would have to be known to the centimeter level. However, this level of accuracy is not attainable. Instead, carrier synchronization is achieved by propagating the Doppler frequency estimates generated by the navigation filter. Fine phase estimation can then be obtained using an additional bank of PLLs (located on the top of Figure 2) as in [Won12b]. However, this part does not affect the remainder of the receiver and may be omitted if the phase estimates are not required.

The alternative to the VDFLL is the proposed VDPLL depicted in Figure 3. It is constituted by two main parts: a vector delay locked loop (VDLL) based on the navigation filter (EKF) that receives samples from the bank of code discriminators and processes the residuals to predict the receiver's position, velocity and clock bias [Par96].

The VDLL is complemented with a carrier tracking structure based on the Co-Op architecture proposed in [Zho98]. The Co-Op algorithm includes a bank of PLLs (one per tracked signal) and a central block fed by the phase discriminators that assists the individual PLLs. This block is designed to track the receiver's motion and clock bias whereas the bank of PLLs is designed to track the satellites motion. Therefore, the characteristics of filters $F_1 \ldots F_N$ may be adjusted separately from the filters F_x, F_y, F_z, F_t. The relative amount of feedback from the central part is tuned through parameter γ, with $0 \leq \gamma \leq 1$. When $\gamma = 0$ the carrier phase synchronization is reduced to a bank of independent PLLs.

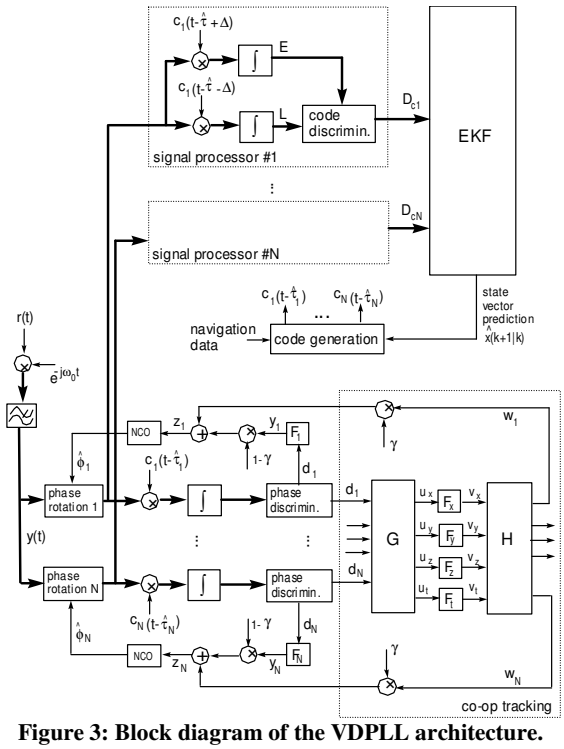

Figure 3: Block diagram of the VDPLL architecture.

In short: while in the VDFLL the code and frequency discriminators feed the EKF through the filter residuals, in the VDPLL only the code discriminator outputs are processed by the EKF. In this case, the phase loops are closed using a hybrid solution: the individual loops track the satellite Doppler frequencies while the central loop follows mainly the frequency changes due to receiver's clock bias and motion.

5. COMPUTER SIMULATIONS

In this section, in order to provide illustrative numerical results, the performance obtained in challenging propagation scenarios using the three different architectures is discussed.

A. Scalar Architectures Tracking Results

Closed-Loop Tracking Results

In the sequel, the performance obtained with different closed-loop techniques is shown in a carrier phase tracking example where the signal of interest is corrupted by ionospheric scintillation, considering different user dynamics:

- *Moderate dynamics* – The signal considered a user starting from a stopped position and accelerating (Doppler rate = 52.5 Hz/s, Doppler accel. = 5.2 Hz/s^2) up to the maximum speed of 50 m/s (max. Doppler = 260 Hz), which is the case for a vehicular user, see Table 1.
- *High dynamics* – In this case the user starts from a stopped position and accelerates (Doppler rate = 50 Hz/s (1g), Doppler accel. = 25 Hz/s^2 (0.5g/s)) up to a maximum speed of 720km/h (max. Doppler = 1 kHz). This example represents the takeoff of a small aircraft (i.e., aeronautical user in-line with maximum dynamics provided in Table 1).

The following simulation considerations apply for all the results presented in the sequel:

1) Scintillation:
- Moderate scintillation – S4 = 0.5
- Severe scintillation – S4 = 0.9

2) Methods:
- FLL-assisted PLL (F-PLL) – used as a reference technique, with a FLL bandwidth Bf = 5 Hz and a PLL bandwidth Bp = 15 Hz.
- KF – Initialization: x_0 = [0,0,0], P_0 = [1/12,0,0;0,4e-4, 0;0,0,4e-12]. **R** tuned according to the actual C/N0 being simulated. **Q** with two different configurations:
  - KF 1, Q = G*G'*1e-12 (with G = [0.5;1.5;2]) being more reactive to estimate changes
  - KF 2 with a slightly lower Q = G*G'*5e-13 (with G = [0.2;0.3;0.5]) being less reactive but filtering more noise.
- AKF – Using the same configuration as the KF but adapting the measurement noise variance with a sequential C/N0 estimation.

Performance Analysis

The main problem when dealing with the carrier phase tracking under scintillation conditions is the fact that the desired phase due to the dynamics of the user and the other phase variations caused by the ionospheric scintillation are coupled in the composite phase, so they cannot be separated. A KF designed to correctly track the carrier phase under high dynamic conditions, i.e. including into the state variable the phase and the Doppler frequency shift components, will also track scintillation variations. Figure 4 shows this problem for the moderate dynamics case, where moderate scintillation is present from second 0.5 to 4.5, and severe scintillation is present from second 5.5 to 9.5.

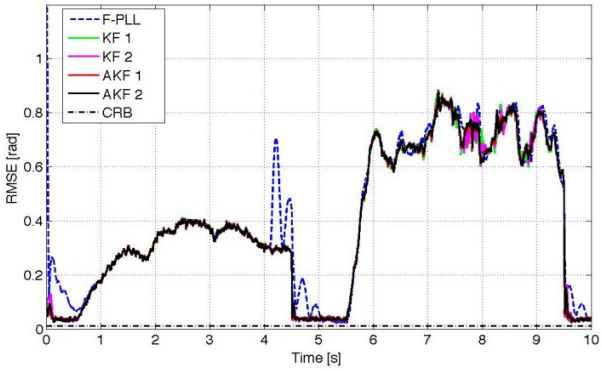

Figure 4: Phase estimation RMSE with moderate dynamics and scintillation (10 seconds of signal), C/N0=45 dB-Hz, and using the ATAN2 discriminator, vehicular user case.

The performance obtained with the different KFs is slightly better than with the F-PLL in terms of accuracy and much better in terms of adaptability. When using the ATAN discriminator the results obtained are not correct when having high dynamics (results not presented herein due to space limitations). There exists an ambiguity introduced by the discriminator (in addition to its reduced dynamic range) that cannot be compensated without an extra processing technique.

In the same scintillation scenario, Figure 5 provides the results for an aeronautical user. In this case the F-PLL is not able to track the carrier phase when the dynamics are too high and loses lock.

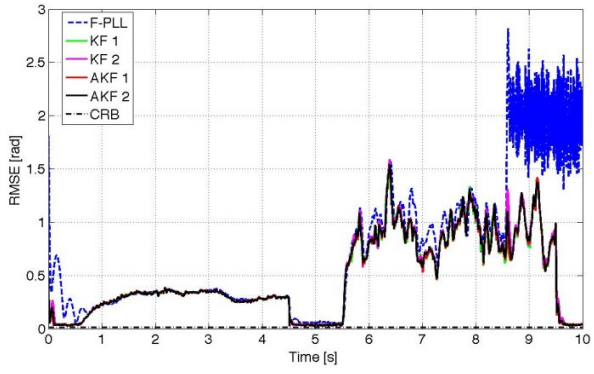

Figure 5: Phase estimation RMSE with high dynamics and scintillation (10 seconds of signal), C/N0=45 dB-Hz, and using the ATAN2 discriminator, aeronautical user case.

Hereafter, the noise and dynamic limits for the F-PLL and the KF-based solutions are assessed considering severe scintillation and moderate dynamics. First, the noise limit is computed, considering a C/N0 range from 26 to 50 dB-Hz and using the ATAN2 discriminator. The performance results are shown in Figure 6, where again the KF-based techniques are slightly better in the complete range of C/N0, but the error is large because of the severe scintillation affecting the signal.

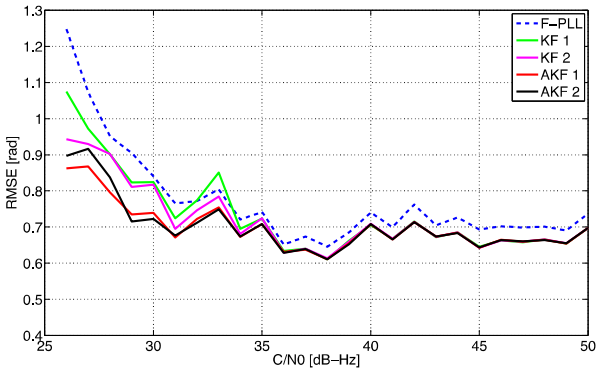

Figure 6: Phase estimation RMSE with moderate dynamics and severe scintillation, considering a C/N0 range in [26 - 50] dB-Hz, using the ATAN2 discriminator.

Note that at low to moderate C/N0, the AKFs are better than the standard KFs. This is an interesting result because it indicates that the effects of fadings caused by scintillation have a worse impact in the performance at low C/N0, and that in this case, having an estimation of the amplitude of such fadings may be helpful. The AKF used for this analysis considered a perfect C/N0 estimation of the received signal at each time epoch (i.e., desired signal + disturbances). The impact of a bad estimation of the C/N0 was found to be marginal. This is clear when comparing the performance of the KF w.r.t the AKF, where the former uses the nominal C/N0 and the latter the C/N0 estimation (i.e., the performance degradation when not using the correct C/N0 but a nominal fixed value is low).

Another important issue when dealing with high dynamics is the so-called dynamic limit, that is, up to which Doppler rate the methods tolerate. Figure 7 presents these results. Again, the F-PLL loses track for Doppler rates greater than 280 Hz/s, while the KF-based methods are able to follow the dynamics of the signal for a value of Doppler rate up to 1 kHz/s (20g). Note that the Doppler rate indicates the frequency increments at each sample in the FLL, thus the frequency jumps for Doppler rates > 200 Hz/s are too large for the FLL with a Bf = 5 Hz to correctly track the Doppler frequency shift. The problem is that the frequency dynamics are too high for a single FLL with a fixed and conservative configuration, to keep track of such an increasing dynamics for such a wide range of values.

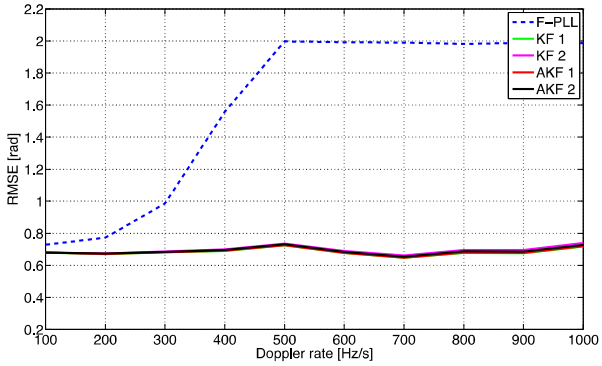

Figure 7: Phase estimation RMSE under severe scintillation, considering a Doppler frequency rate in the range [0.1 – 1] kHz/s (no Doppler acceleration) for C/N0 = 45 dB, using the ATAN2 discriminator.

To conclude, note that when dealing with high dynamic scenarios, the KF-based methods are more accurate, reliable and robust than the F-PLLs. The AKFs will be the choice for low C/N0 values, as it has been shown that the impact of fades is worse in this situation. It is also important to mention that, under these conditions, only the ATAN2 discriminator delivers good estimates without adding extra processing.

Open-Loop Tracking Results

For the open-loop tracking techniques that have already been discussed, we present herein some performance results for a low dynamics scenario with the same type of scintillation effects already considered for closed-loop tracking techniques. The results are shown in Figure 8 for C/N0 = 30 dB-Hz in terms of the probability density function of the carrier frequency estimation errors.

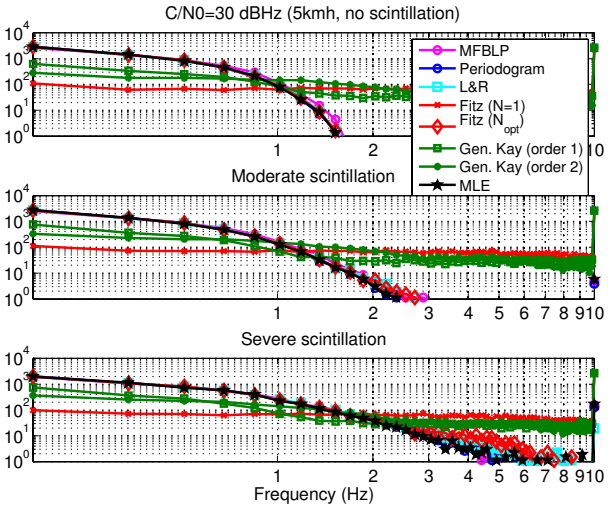

Figure 8: Probability density function of carrier frequency errors provided by different candidate open-loop techniques.

The first observation is that Kay's method provides a kind of uniform error distribution, which means that the algorithm is unable to operate at such a low C/N0 value (which is consistent with the fact that the linear phase approximation that it assumes is only valid at medium to high C/N0). The second observation is that all the remaining open-loop techniques have a pretty similar performance in the absence of scintillation, but as this degradation becomes more relevant, the ML technique is the one that exhibits a slightly narrower error distribution, thus, providing a slightly better performance than the rest of techniques. A comparison with closed-loop techniques is shown in Figure 9, where for severe scintillation and C/N0 = 30 dB-Hz, open-loop techniques seem to provide a better performance. That is, they exhibit a narrower phase error distribution compared to their closed-loop counterparts.

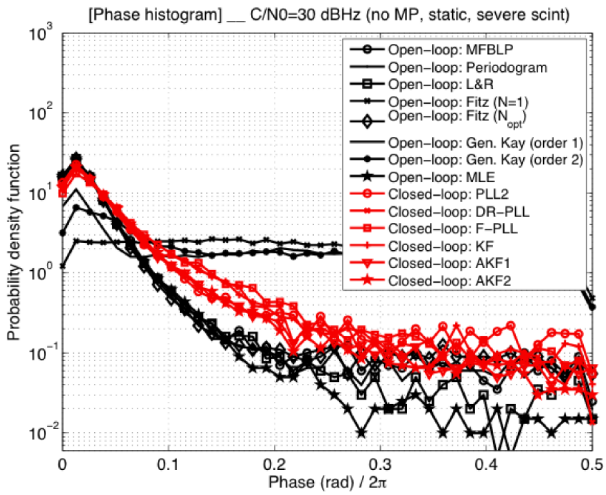

Figure 9: Probability density function of carrier frequency errors provided by different open-loop and closed-loop techniques.

B. Vector Architectures Tracking Results

The simulations are carried out at longitude 14.409° W, latitude 7.975° S, and altitude 920 meters, corresponding to a location at Ascension Island. Two scenarios are considered: fixed user in scintillation channel and moving user with circular motion in AWGN channel. The date, 27 March 2000, was selected as it is associated with a period of large ionospheric activity and experimental data exist for that location and time. Over the Ascension Island, on March 27th, 23:30 UT, 12 GPS satellites are visible: one not healthy (PRN 16) and nine with elevation angle greater than 5 degrees.

The simulations use 4 and 8 satellites. The former is the minimum number that ensures the solution of the navigation equation whereas the latter is nearly the maximum number of satellites visible above a reasonable mask angle. Each satellite set is selected using the minimum GDOP criterion. Simulations with 4 satellites use PRN 4, 7, 19, and 26 (GDOP = 2.67); simulations with 8 satellites use PRN 2, 4, 7, 8, 13, 19, 24, and 26 (GDOP = 2.25).

Modulation BOC(1,1) in a pilot channel is considered. The EKF assumes that all received signals exhibit the same C/N0 with 15 dB-Hz below the true level. Simulations were performed with the PV model and the parameter q_m of the dynamics noise covariance matrix [11] was made equal to 0.01 for the fixed user scenario and 10 for the circular motion.

The bandwidths of the scalar loop filters and the central filters were made equal to 2 Hz and 20 Hz, respectively. In the fixed user scenario all the filters are first-order (second-order PLL). Second-order filters were used in the circular motion scenario (third-order PLL). The weight of the central filters γ was made equal to 0.5.

The algorithms were simulated over a range of C/N0 ratios (25 to 50 dB-Hz) using a semi-analytic approach, i.e., modeling the correlator outputs and simulating only the non-linear parts of the system. At each C/N0 ratio, one hundred simulation runs of 300 second duration were carried out. Each run uses different clock noise and thermal noise samples. The pre-detection integration interval is made equal to 10 ms (which is the EKF update interval), a temperature compensated xtal oscillator model is considered, and the early-late space was made equal to 0.1 T_c, where T_c is the chip period. In the Figures 10 and 11 the (mod 2π) rms phase error is defined as:

$$\frac{1}{R}\sum_{r=1}^{R}\sqrt{\frac{1}{SN}\sum_{s=1}^{S}\sum_{n=1}^{N}\left[\cos^{-1}\left(\cos\left[\hat{\phi}_s^{(r)}(n)-\phi_s^{(r)}(n)\right]\right)\right]^2}$$

where $\phi_s^{(r)}(n)$ and $\hat{\phi}_s^{(r)}(n)$ denote, respectively, the true value and the estimated carrier phase of satellite s (with $s = 1, \ldots, S$), at epoch n (with $n = 1, \ldots, N$), and run r (with $r = 1, \ldots, R$).

Figure 10 shows the vector architectures performance versus the C/N0 for a fixed user in a scintillation channel using four satellites affected by scintillation with S4 = 0.5 and S4 = 0.9 (see section Propagation Environment). The carrier phase performances of both architectures are equivalent and increase as the scintillation parameter S4 decreases. In presence of severe scintillation, the mean rms phase errors are almost independent of the C/N0 which means that the ionospheric scintillation effect is predominant over the thermal noise. Note that both architectures try to track the overall phase which includes a significant part due to the scintillation effect. As all the four tracked satellites are affected by scintillation with equal intensity, the asymptotic value of the mean rms

phase error gives approximately the phase standard deviation corresponding to each scintillation time-series. For comparison in Fig. 10 we present also the architectures performance for AWGN channel.

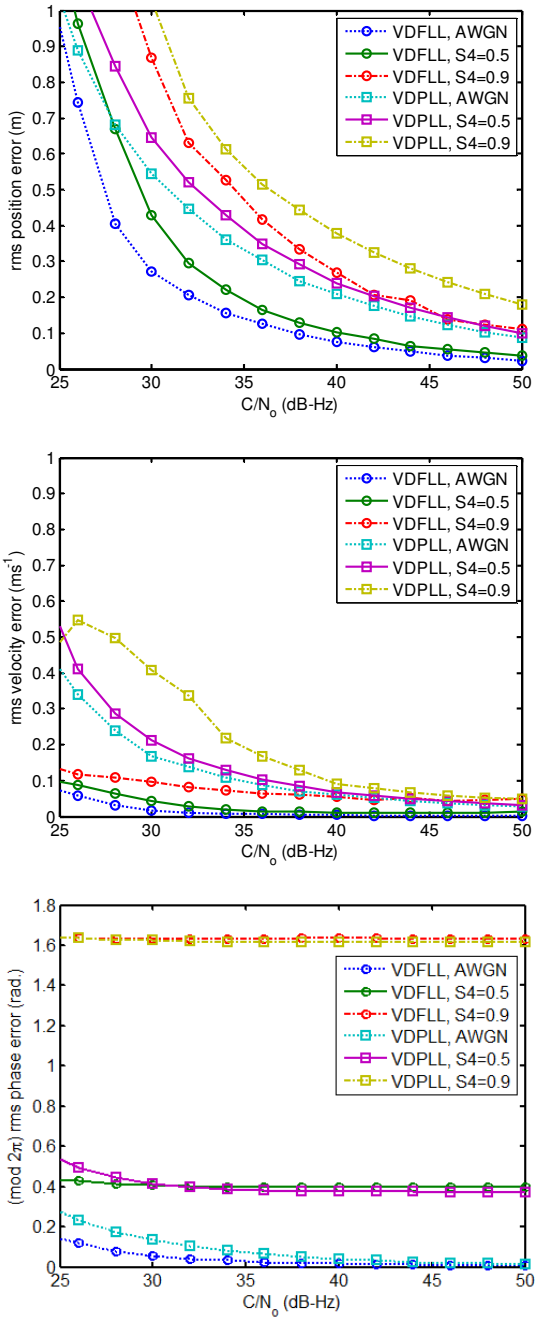

Figure 10: Vector architectures performance: fixed user in scintillation channel (GISM time-series) using 4 satellites.

Figure 11 shows the vector architectures performance versus the C/N0 in AWGN channel for a moving user along a circular path: uniform motion in a circle of radius 250 m in the Y-Z plane with 50 m/s velocity and 10 m/s^2 acceleration. Four and eight satellites were considered.

The sudden growth of the rms position error for the VDFLL with four satellites evidences the fact that the architecture becomes more stable when the number of tracked satellites increases. In both architectures a C/N0 greater than 30 dB-Hz ensures a small probability of divergence of the central EKF.

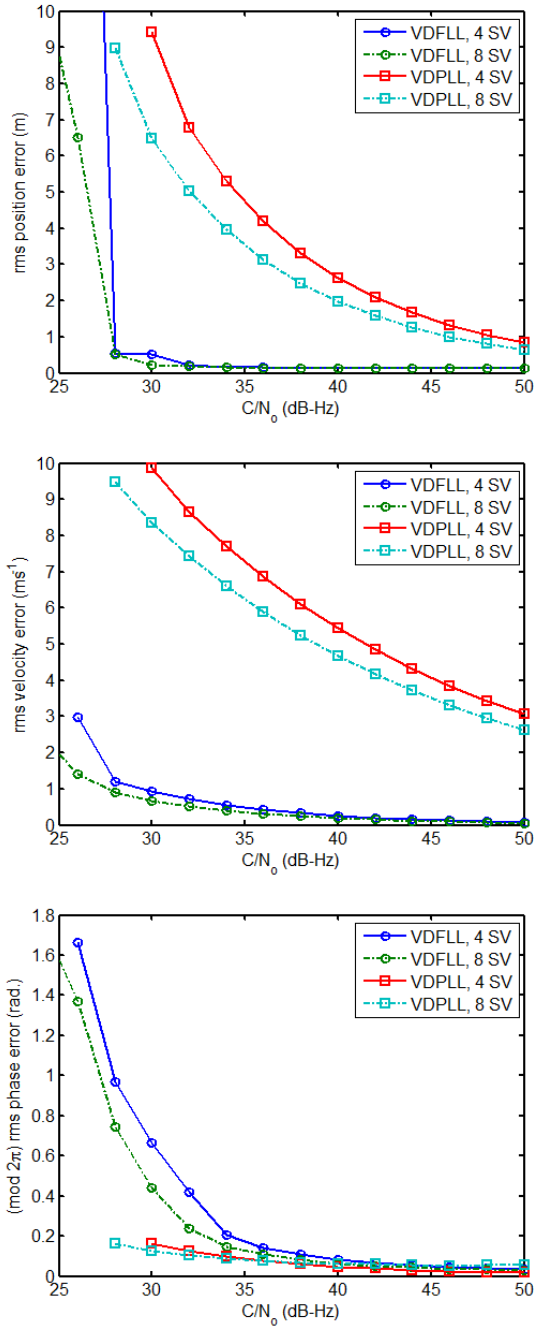

Figure 11: Vector architectures performance: moving user with circular motion in AWGN channel using 4 and 8 satellites.

The circular motion in Figure 11 is characterized by a constant acceleration. In order to avoid steady-state phase errors, it is necessary to employ second-order filters in the individual and the common phase lock loops of the VDPLL as mentioned, for instance, in [Kap06]. This corresponds to use third-order PLLs. In contrast, the VDFLL feeds the discriminators outputs (residuals) directly to the EKF, thus avoiding this drawback. In fact, the VDFLL performance does not degrade significantly if first-order loop filters are used.

Figure 12 depicts the position and clock errors along the first 120 seconds of a run in the VDFLL architecture using four satellites, with C/N0=40 dB-Hz, for a fixed user in AWGN channel. In the interval 50<t<60 s the occlusion of one satellite is simulated. The results of the figure illustrate the ability of the vector architectures to withstand satellite occlusions without loss of tracking. When the number of satellites passes from 4 to 3 there will be a gradual increase of the tracking errors but this situation is rapidly reversed after the occlusion finishes. This contrasts with the scalar architectures where the same occlusion leads to the loss of tracking and the need to re-acquire code and phase.

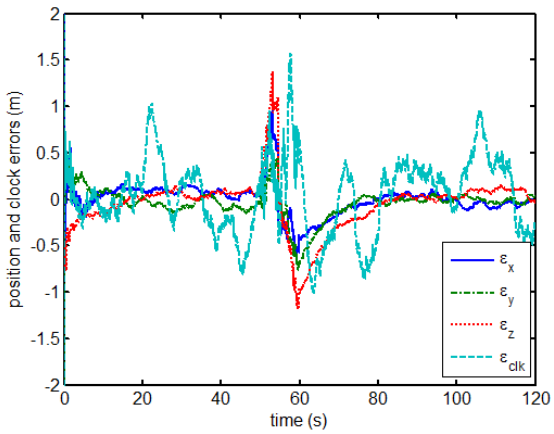

Figure 12: VDFLL position and clock errors: fixed user in AWGN channel with PRN #19 occlusion for 50 < t < 60 s.

The VDFLL architecture exhibits better results than the VDPLL in terms of position and velocity errors. The good performance of the VDFLL can be explained by the fact that the EKF is fed with two different types of inputs: the code and the frequency discriminators, as shown in Figure 2. This leads to improved architecture robustness which is crucial in scenarios where the number of received signals with good C/N0 varies rapidly, such as in scintillation and fading environments.

6. CONCLUSIONS AND FUTURE WORK

In this work, different techniques for robust carrier phase tracking in GNSS receivers are described. These techniques range from advanced closed-loop scalar techniques, to open-loop implementations of tracking techniques, up to vector tracking approaches.

Each family of techniques is shown to have both strengths and weaknesses, which are in turn related to the channel characteristics, propagation environment, and system dynamics. Therefore, the trade-offs presented are of paramount importance, leading to a careful tailoring of the design, parameterization of the technique, and applicability scenario.

In the presence of fading or scintillation phenomena, it can be seen that the scalar closed-loop carrier phase tracking techniques tend to show similar degradation, although the performance itself is comparable in term of behavior. For the case of the closed-loop techniques, for instance, the Kalman-based approaches tend to show less phase estimation error with respect to the more "conventional" F-PLL. When high dynamics are observed, however, the differences in performance are more notorious: in fact, Kalman-based approaches seem more robust, whereas the F-PLL shows lower dynamic limits.

The applicability of open-loop scalar tracking techniques is also discussed, where it can be seen that some benefit is achieved for lower C/N0 ranges. This is the case in the presence of fading or even indoor situations. The MLE carrier tracking technique is found to provide the most robust performance, and when compared to closed-loop techniques, a gain is observed in terms of a narrower error distribution, at the cost of additional complexity.

Vector tracking techniques traditionally target a better position-based solution, i.e. to reduce the position error by taking advantage of the channel combination processes. However, such implementations can also be tailored for carrier phase tracking, as in the case of the VDFLL with the additional bank of coupled PLL.

By extending the parameterization and tuning of vector tracking techniques, the resolution of the position solution achieved by the vector implementation can be re-introduced in steering the tracking of each signal's carrier phase.

In conclusion, the different approaches presented for robust carrier phase tracking show promising results for harsh environments. In particular, the impact of propagation phenomena such as fading, ionosphere scintillations and high dynamics in carrier phase tracking can be reduced by proper trade-off analysis and technique tailoring.

In today's GNSS applications, an increasing number of external information sources, like aiding or inertial

sensors, can be coupled with such techniques to enable rapid (i.e. real-time) configurability and adaptability to the changing propagation environment. Hence, one direction to explore is the tailoring of receiver architectures to specific applications: in other words, a receiver architecture may use a combination of these robust techniques, together with any information from external sources, to take the most of their strengths, whether in terms of low C/N0 tracking capability, carrier phase tracking accuracy or improved dynamic limits.

7. ACKNOWLEDGMENTS

The content of the present article reflects solely the authors' view and by no means represents the official ESA or GALILEO views.

8. REFERENCES

[ITU12] "Ionospheric propagation data and prediction methods required for the design of satellite services and systems", Recommendation ITUR P.531-11, January 2012

[Kap06] "Understanding GPS – Principles and Applications", E. Kaplan, C. Hegarty, Artech House, Boston, Second Edition, 2006

[Tsu00] "Fundamentals of Global Positioning Systems – A Software Approach", Bao-Yen Tsui, J. John, Willey & Sons Inc., 2000

[Lia04] P. Lian, "Improving tracking performance of PLL in high dynamic applications, Ph.D. thesis, University of Calgary, Calgary, Canada, 2004

[Zha09] L. Zhang and Y.T. Morton, "Tracking GPS signals under ionosphere scintillation conditions," in Proc. ION GNSS ITM, Sept. 2009, pp. 227– 234

[Mao10] X. Mao, Y. T. Morton, L. Zhang, and Y. Kou, "GPS carrier signal parameters estimation under ionosphere scintillation," in Proc. ION GNSS, 2010

[Sko05] G. Skone, G. Lachapelle, D. Yao, W. Yu, and R. Watson, "Investigating the impact of ionospheric scintillation using a GPS software receiver," in Proc. of the ION GNSS, Long Beach, CA, Sept. 13-16, 2005

[Mao08] W.-L. Mao, A.-B. Chen, "Mobile GPS carrier phase tracking using a novel intelligent dual–loop receiver," Intl. Journal of Satellite Comm. and Networking, vol. 26, pp. 119– 139, 2008

[Psi07] M.L. Psiaki, T. Humphreys, A. Cerruti, S. Powell, and P. Kintner, "Tracking L1 C/A and L2C signals through ionospheric scitillations," in Proc. ION GNSS ITM 2007, pp. 246–268

[Hum05] T.E. Humphreys, M.L. Psiaki, P.M. Kintner, and B.M. Ledvina, "GPS carrier tracking loop performance in the presence of ionospheric scintillations," in Proc. ION GNSS, 2005

[Yu06] W. Yu, G. Lachapelle, and S. Skone, "PLL performance for signals in the presence of thermal noise, phase noise, and ionospheric scintillation," in Proc. of the ION GNSS, Fort Worth, TX, Sept. 2006

[Hu03] C. W. Hu, W. Chen, Y. Chen, and D. Liu, "Adaptive Kalman filtering for vehicle navigation," Journal of Global Positioning Systems, vol. 2, no. 1, pp. 42–47, 2003

[Zha10] L. Zhang, Y. T. Morton, and M. M. Miller, "A variable gain adaptive Kalman filter-based GPS carrier tracking algorithms for ionosphere scintillation signals," in Proc. ION GNSS ITM, 2010

[Won12a] J.-H. Won, B. Eissfeller, T. Pany, and J. Winkel, "Advanced signal processing scheme for GNSS receivers under ionospheric scintillation," in Proc. of the IEEE/ION PLANS, 2012, pp. 44–49

[Sec12] G. Seco-Granados, J. A. López-Salcedo, D. Jiménez-Baños, and G. López-Risueño, "Challenges in indoor global navigation satellite systems," IEEE Signal Processing Mag., vol. 29, pp. 108–131, March 2012

[Vil89] V. A. Vilnrotter, S. Hinedi, R. Kumar, "Frequency estimation techniques for high dynamic trajectories", IEEE Trans. on Aerospace and Electronic Systems, vol. 20, no. 4, pp. 559-576, July 1989

[Fit91] M. P. Fitz, "Planar filtered techniques for burst mode carrier synchronization", Proc. IEEE GLOBECOM, vol. 1, pp. 365-369, 1991

[Lui95] M. Luise, R. Reggiannini, "Carrier frequency recovery in all digital modems for burst mode transmissions", IEEE Trans. on Communications, vol. 43, no. 2/3/4, pp. 1169-1178, Feb./Mar./Apr. 1995

[Men97] U. Mengali and A. N. D'Andrea, "Synchronization techniques for digital receivers", Plenum Press, 1997

[Kay89] S. Kay, "A fast and accurate single frequency estimator", IEEE Trans. on Acoustics, Speech and Signal Processing, vol. 37, no. 12, pp. 1987-1990, December 1989

[Rib98] J. Riba, M. A. Lagunas, “Instantaneous open-loop frequency estimation methods for navigation receivers”, Proc. 6th Workshop on Digital Signal Processing Techniques for Space Applications (DSP), pp. 1-8, 1998

[Tuf82] D. W. Tufts, R. Kumaresan, “Estimation of frequencies of multiple sinusoids: making linear prediction perform like maximum likelihood”, Proc. of the IEEE, vol. 70, no. 9, pp. 975-989, September 1982

[Pet08] M. Petovello, C. O’Driscoll, G. Lachapelle, “Carrier Phase Tracking of Weak Signals Using Different Receiver Architectures,” ION NTM-2008, San Diego, CA, pp. 1-11, Jan. 2008

[Las09] M. Lashley, Modeling and Performance Analysis of GPS Vector Tracking Algorithms, PhD dissertation, Auburn University, Al, 2009

[Won12b] J.-H. Won, T. Pany and B. Eissfeller, “Characteristics of Kalman filters for GNSS tracking loops”, IEEE Trans. on Aerospace and Electronic Systems, vol. 48, no. 4, pp. 3671-3681. October 2012

[Par96] B. Parkinson, J. Spilker, Global Positioning System: Theory and Applications, vol. I, chapter 7, AIAA, 1996

[Zho98] M. Zhodzishsky, S. Yudanov, V. Veitsel, J. Ashjaee, “Co-Op Tracking for Carrier Phase,” in Proc. ION GPS 1998, Nashville, TN, September 1998, pp. 653-664