ON THE LOCAL CONVERGENCE OF THE ITERATED EXTENDED KALMAN FILTER

P.J.G. Teunissen
Delft Geodetic Computing Centre (LGR)
Department of Geodesy, Delft University of Technology
Thijsseweg 11, 2629 JA Delft, The Netherlands

Abstract

It is shown in this paper that the iterated extended Kalman filter has a local linear rate of convergence. An expression for the rate of convergence is given. And some practical estimates of the rate of convergence for distance and azimuth observations are derived.

I. Introduction

Consider the minimization problem

$$(1) \qquad \min_{x_o \ldots x_k} \{\|y_o - A_o(x_o)\|^2_{y_o} + \sum_{i=1}^{k} \|y_i - A_i(x_i)\|^2_{y_i} + \|d_i + \Phi_{i,i-1}x_{i-1} - x_i\|^2_{d_i}\},$$

where: $y_i \in \mathcal{R}^{m_i}$ are given data vectors; $d_i \in \mathcal{R}^n$ are given system inputs; $x_i \in \mathcal{R}^n$ are the unknown state vectors of epoch i; $A_i : \mathcal{R}^n \rightarrow \mathcal{R}^{m_i}$ are (non)linear vector functions that map the state space $\mathcal{R}^n$ into the data space $\mathcal{R}^{m_i}$; $\Phi_{i,i-1} : \mathcal{R}^n \rightarrow \mathcal{R}^n$ are the transition matrices; $\|.\|^2_{y_i} = (.)^* Q^{-1}_{y_i}(.)$ with Q_{y_i} the variance matrix of $\underline{y}_i$; and $\|.\|^2_{d_i} = (.)^* Q^{-1}_{d_i}(.)$ with Q_{d_i} the variance matrix of $\underline{d}_i$.
Minimization problems of the type (1) typically occur in navigation applications. Consider for instance the problem of recovering the trajectory $(u(t), v(t))^*$ of an object that moves in the plane. If positioning is based on discrete-time sampling of distances and azimuths, then the entries of the maps A_i take the form

$$A_i(x_i) = ([u_i^2 + v_i^2]^{1/2}, arctan[u_i/v_i])^*.$$

And if in addition acceleration $(\ddot{u}(t), \ddot{v}(t))^*$ is observed as a white-noise process with spectral densities $q_{\ddot{u}}, q_{\ddot{v}}$, the state vectors take the form $x_i = (u_i, \dot{u}_i, v_i, \dot{v}_i)^*$ and the system input, its variance matrix and the transition matrix for the u-axis read

$$d_i = \int_{t_{i-1}}^{t_i} \begin{bmatrix} t_i - \tau \\ 1 \end{bmatrix} \ddot{u}(\tau) d\tau,$$

$$Q_{d_i} = q_{\tilde{u}} \begin{bmatrix} \frac{1}{3}(t_i - t_{i-1})^3 & \frac{1}{2}(t_i - t_{i-1})^2 \\ \frac{1}{2}(t_i - t_{i-1})^2 & (t_i - t_{i-1}) \end{bmatrix}, \Phi_{i,i-1} = \begin{bmatrix} 1 & (t_i - t_{i-1}) \\ 0 & 1 \end{bmatrix}.$$

In navigation applications the minimization problem (1) often has to be solved in real-time. This implies that (1) has to be solved recursively. If this is done when the maps $A_i : \mathcal{R}^n \rightarrow \mathcal{R}^{m_i}$ are *linear*, the well-known Kalman prediction- and filtering equations are obtained. They read

$$(2) \qquad \begin{cases} \hat{x}_{k|k-1} = \Phi_{k,k-1}\hat{x}_{k-1|k-1} + d_k, \\ \hat{x}_{k|k} = \hat{x}_{k|k-1} + K_k[y_k - A_k\hat{x}_{k|k-1}], \end{cases}$$

with corresponding variance matrices

$$\begin{aligned} Q_{k|k-1} &= \Phi_{k,k-1}Q_{k-1|k-1}\Phi_{k,k-1}^* + Q_{d_k}, \\ Q_{k|k} &= [I - K_kA_k]Q_{k|k-1}, \end{aligned}$$

where

$$K_k = Q_{k|k-1}A_k^*[Q_{v_k} + A_kQ_{k|k-1}A_k^*]^{-1},$$

is the so-called Kalman gain matrix.

In the present contribution we will assume the maps $A_i : \mathcal{R}^n \rightarrow \mathcal{R}^{m_i}$ to be *nonlinear*. This implies that in general no direct method exists for solving (1) either in batch or recursive form. One therefore has to take recourse to computational techniques that are iterative in nature. The necessity to iterate, however, puts a strain on the requirement of being able to perform the computations in real-time. It is therefore of importance for a particular application to know in advance the expected number of iterations that are required to obtain a given numerical precision in the computed results. In other words, one needs to know the expected *rate of convergence* of the iteration method applied. This will therefore be the subject of the present contribution.

Although various iterative techniques exist for solving nonlinear optimization problems like (1), we will restrict ourselves in the present contribution to the Gauss-Newton method. The reason is that the Gauss-Newton method takes advantage of the sums of squares structure of the objective function. The method is therefore especially suited for solving minimization problems like (1).

II. Nonlinear Least Squares

Since (1) is essentially a least-squares problem, a brief review is given in this section of some of the geometrical and numerical characteristics of general nonlinear least-squares problems. The main results are taken from [Teunissen, 1984, 1985]. The general nonlinear least-squares problem reads

$$(3) \qquad \min_x \|y - A(x)\|^2,$$

where: $\|.\|^2 = (.)^* Q_y^{-1}(.)$; Q_y is positive definite; y is an m-dimensional data vector; and $A(.)$ is a nonlinear vector function or map from $\mathcal{R}^n$ into $\mathcal{R}^m$.
For varying values of $x \in \mathcal{R}^n$, $A(x)$ traces locally an n-dimensional manifold embedded in $\mathcal{R}^m$. If the metric of $\mathcal{R}^m$ is described by the positive definite matrix Q_y^{-1}, the scalar $\|y - A(x)\|$ equals the distance from point y to the point $A(x)$ on the manifold. Hence the minimization problem (3) corresponds to the problem of finding that point on the manifold, say $\hat{y} = A(\hat{x})$, which has least distance to y. This geometry of the nonlinear least-squares problem is shown in figure 1.

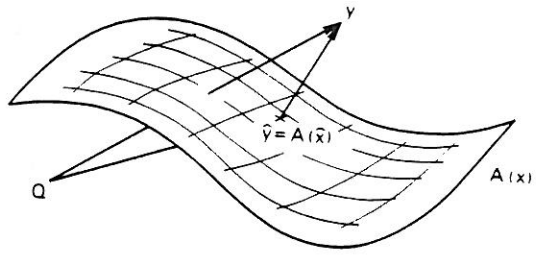

Figure 1: Geometry of nonlinear least-squares

The minimizer $\hat{x}$ of (3) can in principle be located by any one of the existing iterative descent methods. The Gauss-Newton method, however, is especially suited for solving nonlinear least-squares problems, since it takes advantage of the sums of squares structure of the objective function. With the Gauss-Newton method the minimizer $\hat{x}$ of (3) is located according to the following iterative scheme:

$$x^{\alpha+1} = x^\alpha + t^\alpha [\partial_x A(x^\alpha)^* Q_y^{-1} \partial_x A(x^\alpha)]^{-1} \partial_x A(x^\alpha)^* Q_y^{-1} [y - A(x^\alpha)], \quad \alpha = 0.1..... \tag{4}$$

The positive scalar t^α is used to control the line search strategy. It is determined such that $\|y - A(x^{\alpha+1})\| \leq \|y - A(x^\alpha)\|$ holds in each iteration step. In many applications it suffices to take $t^\alpha = 1$. The iterative scheme (4) is initialized through an externally provided initial parameter vector x^o. And the iteration is terminated once a stop criterion is met. The stop criterion is based on the convergency test

$$\|x^{\alpha+1} - x^\alpha\|_\alpha < \epsilon,$$

with ϵ an a priori given tolerance level and $\|.\|_\alpha^2 = (.)^* Q(x^\alpha)^{-1}(.), Q(x^\alpha) = [\partial_x A(x^\alpha)^* Q_y^{-1} \partial_x A(x^\alpha)]^{-1}$. Once the stop criterion is met, $x^{\alpha+1}$ is accepted as the solution of (3).
Each iteration step of the Gauss-Newton method provides a solution of a linear(ized) least-squares problem. The geometry of the Gauss-Newton method is therefore one

of orthogonal projection. That is, the vector $\partial_x A(x^\alpha)(x^{\alpha+1} - x^\alpha)$, wich lies in the tangent space of the manifold $A(x)$ at $A(x^\alpha)$, is the orthogonal projection of the residual vector $e(x^\alpha) = y - A(x^\alpha)$ onto this tangent space. See figure 2. This indicates that the geometry of the manifold $A(x)$ must play an important role in the local behaviour of the Gauss-Newton method. And indeed it turns out that the local convergence of the Gauss-Newton method depends on the *normal curvatures* of the manifold $A(x)$ at $A(\hat{x})$.

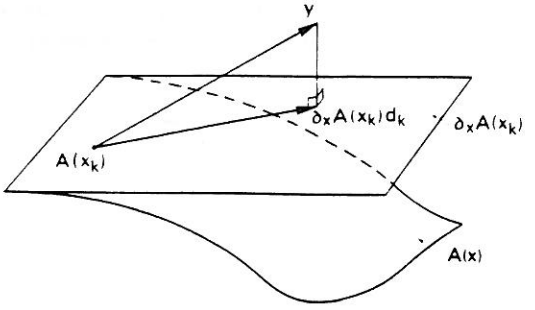

Figure 2: Orthogonal projection onto tangent space of $A(x)$ at $A(x^\alpha)$.

Since [see Teunissen, 1985, section IV.4],

$$(5)\quad x^{\alpha+1} - \hat{x} = [(1 - t^\alpha)I + t^\alpha Q(\hat{x})[e(\hat{x})^* Q_y^{-1} \partial_{xx}^2 A(\hat{x})]](x^\alpha - \hat{x}) + 0(\|x^\alpha - \hat{x}\|^2),$$

it follows that the Gauss-Newton method has a *linear* rate of convergence for points x^α sufficiently close to solution $\hat{x}$. If we define the normal curvature of manifold $A(x)$ for $v \in \mathcal{R}^n$ and $n \in R(\partial_x A(x))^\perp$, with $n^* Q_y^{-1} n = 1$ as

$$(6)\qquad k_n(v) = \frac{v^*[n^* Q_y^{-1} \partial_{xx}^2 A(x)]v}{v^* Q(x)^{-1} v},$$

the $Q(\hat{x})^{-1}$-weighted norm of (5) follows for $t^\alpha = 1$ as

$$(7)\qquad \|x^{\alpha+1} - \hat{x}\| \leq \max_v |k_{\hat{n}}(v)| \|e(\hat{x})\| \|x^\alpha - \hat{x}\|,$$

where $\hat{n} = e(\hat{x})/\|e(\hat{x})\|$. This important result shows that local convergence of the Gauss-Newton method is *guaranteed* if the observation point y lies within a hypersphere with centre $A(\hat{x})$ and radius $1/\max_v |k_{\hat{n}}(v)|$.

III. Iterated Filtering

In this section we will apply the results of the previous section to the recursive solution of the minimization problem (1). For $k = 0$, the recursion first has to be initialized. This is done by solving with the Gauss-Newton method the nonlinear least-squares problem

$$\min_{x_o} \|y_o - A_o(x_o)\|^2_{y_o}. \tag{8}$$

The next step after initialization is the propagation in time of the minimizer $\hat{x}_{1|0}$ with corresponding variance matrix $Q_{1|0}$. Prediction is accomplished with the time-update equations

$$\hat{x}_{k|k-1} = \Phi_{k,k-1}\hat{x}_{k-1|k-1} + d_k, \tag{9}$$

and

$$Q_{k|k-1} = \Phi_{k,k-1}Q_{k-1|k-1}\Phi^*_{k,k-1} + Q_{d_k}. \tag{10}$$

It should be noted that although both equations (9) and (10) are exact, the matrix $Q_{k-1|k-1}$ and therefore also the matrix $Q_{k|k-1}$, will usually be only an approximation of the actual variance matrix. This is due to the nonlinearity of the maps $A_i, i = 0, 1, ..., (k-1)$. If $Q_{k-1|k-1}$ is taken as the inverse of the normal matrix of the corresponding linearized least-squares problem, then higher order terms of the variance of unit weight are neglected. See [Teunissen, 1989] for an expression of the variance matrix which includes the second order terms of the variance of unit weight.

The step following prediction is filtering. This implies solving the nonlinear least-squares problem

$$\min_{x_k} \|y_k - A_k(x_k)\|^2_{y_k} + \|\hat{x}_{k|k-1} - x_k\|^2_{k|k-1}, \tag{11}$$

where $\|.\|^2_{k|k-1} = (.)^* Q^{-1}_{k|k-1}(.)$. The filtered estimate $\hat{x}_{k|k}$ follows then as the minimizer of (11). Since the map A_k is assumed to be nonlinear, filtering has to be iterated in principle. When the Gauss-Newton method with line search is applied to (11), the following *iterative filtering scheme* is obtained:

$$\boxed{x^{\alpha+1}_{k|k} = [1 - t^\alpha_k]x^\alpha_{k|k} + t^\alpha_k[\hat{x}_{k|k-1} + K_k(x^\alpha_{k|k})[y_k - A_k(x^\alpha_{k|k}) - \partial_x A_k(x^\alpha_{k|k})(\hat{x}_{k|k-1} - x^\alpha_{k|k})]]} \tag{12}$$

with

$$K_k(x^\alpha_{k|k}) = Q_{k|k-1}\partial_x A_k(x^\alpha_{k|k})[Q_{y_k} + \partial_x A_k(x^\alpha_{k|k})Q_{k|k-1}\partial_x A_z(x^\alpha_{k|k})^*]^{-1}$$

or

$$K_k(x^\alpha_{k|k}) = [Q^{-1}_{k|k-1} + \partial_x A_k(x^\alpha_{k|k-1})^* Q^{-1}_{y_k}\partial_x A_k(x^\alpha_{k|k})]^{-1}\partial_x A_k(x^\alpha_{k|k})^* Q^{-1}_{y_k}.$$

The positive line search scalar t^α_k is used to enforce that the value of the objective function of (11) decreases in each iteration step. Although the iterative filtering scheme (12) can be initialized with an *externally* provided initial state vector $x^o_{k|k}$, it is in the present situation more expedient to take as the initial state vector the predicted estimate $\hat{x}_{k|k-1}$:

(13) $$x^0_{k|k} := \hat{x}_{k|k-1}.$$

If with (13) and for $t^\alpha_k = 1$, $x^1_{k|k}$ is accepted as the filtered estimate $\hat{x}_{k|k}$, it follows from (12) that

(14) $$\boxed{\hat{x}_{k|k} = \hat{x}_{k|k-1} + K_k(\hat{x}_{k|k-1})[y_k - A_k(\hat{x}_{k|k-1})]}$$

This is the *extended Kalman filter* [see e.g. Gelb, 1974]. This shows that the extended Kalman filter is identical to the *first* iteration step of the Gauss-Newton method when applied to (11). In a similar way we obtain the *iterated extended Kalman filter* [see Gelb e.g., 1974] from (12) by excluding the line search strategy:

$$\boxed{x^{\alpha+1}_{k|k} = \hat{x}_{k|k-1} + K_k(x^\alpha_{k|k})[y_k - A_k(x^\alpha_{k|k}) - \partial_x A_k(x^\alpha_{k|k})(\hat{x}_{k|k-1} - x^\alpha_{k|k})] \quad \alpha = 0, 1, ...}$$

(15)

This shows that the iterated extended Kalman filter is *identical* to the Gauss-Newton method when applied to (11). The implication is therefore that the convergency characteristics of the iterated extended Kalman filter can be found from those of the general Gauss-Newton method. This idea will be pursued in the next section.

IV. The Rate of Convergence of the Iterated Extended Kalman Filter

It will be clear that at each recursion step of the filter, iteration is needed in principle to take care of existing nonlinearities. But iteration contributes to the computation time and therefore puts a strain on the requirement of real-time filtering. It would therefore be ideal if one can show for a particular application that it suffices to base the filtering on the extended Kalman filter. In that case iteration is avoided and filtering is based on simple recursive formulas. But in order to be able to show this, one first needs to know the convergency behaviour of the iterated extended Kalman filter. That is, the decision whether to iterate or not should be based on the number of significant digits that are gained in each iteration step. And this implies that knowledge of the rate of convergence of the iterated extended Kalman filter is needed. The local convergency behaviour of the iterated extended Kalman filter can be derived from the results of section II. By excluding the line search strategy, we obtain analogous to (5) for the iterated extended Kalman filter:

$$x^{\alpha+1}_{k|k} - \hat{x}_{k|k} = [Q_{k|k}(\hat{x}_{k|k})[e_k(\hat{x}_{k|k})^* Q^{-1}_{y_k} \partial^2_{xx} A_k(\hat{x}_{k|k})]](x^\alpha_{k|k} - \hat{x}_{k|k}) + O(\|x^\alpha_{k|k} - \hat{x}\|^2),$$

(16)

where: $Q_{k|k}(\hat{x}_{k|k}) = [Q^{-1}_{k|k-1} + \partial_x A_k(\hat{x}_{k|k})^* Q^{-1}_{y_k} \partial_x A_k(\hat{x}_{k|k})]^{-1}$ and $e_k(\hat{x}_{k|k}) = y_k - A_k(\hat{x}_{k|k})$. The residual $e_k(\hat{x}_{k|k})$ should not be confused with the predicted residual.

Equation (16) shows that the iterated extended Kalman filter has a *linear* rate of convergence for points $x^a_{k|k}$ sufficiently close to the filtered estimate $\hat{x}_{k|k}$. In practical applications this will often be the case if the iteration is initialized with (13).
From taking the $Q_{k|k}(\hat{x}_{k|k})^{-1}$-weighted norm of (16), the linear convergency factor (CF) of the iterated extended Kalman filter follows as

$$CF \leq \max_{v} \left| \frac{v^*[\hat{u}^* Q_{y_k}^{-1} \partial^2_{xx} A(\hat{x}_{k|k})]v}{v^* Q_{k|k}(\hat{x}_{k|k})^{-1} v} \right| \| y_k - A_k(\hat{x}_{k|k}) \|_{y_k} \tag{17}$$

with $\hat{u} = e_k(\hat{x}_{k|k}) / \| e_k(\hat{x}_{k|k}) \|_{y_k}$. Note that the unit vector $\hat{u}$ not necessarily has to be orthogonal to the range space of $\partial_x A_k(\hat{x}_{k|k})$. The above result shows that the rate of convergence of the iterated extended Kalman filter is governed by:

1. The consistency of the data: $\| y_k - A_k(\hat{x}_{k|k}) \|_{y_k}$
2. The nonlinearity of the map A_k : $\partial^2_{xx} A_k(\hat{x}_{k|k})$.
3. The a priori and a posteriori precision: $Q_{y_k}, Q_{k|k}(\hat{x}_{k|k})$.

Hence, an increasing precision of the predicted state vector for instance, will result in a higher rate of convergence. This implies that a higher discrete-time data sampling rate will also increase the rate of convergence.

V. Estimates of the Rate of Convergence for Distance and Azimuth Observables

In this section we will derive some *practical estimates* for the rate of convergence of the iterated extended Kalman filter. The estimates will be given for both distance and azimuth observables. The estimates will be based on the following upperbound of (17):

$$CF \leq [\sum_{i=1}^{m_k} \lambda_i^2 / \sigma^2_{y_i}]^{1/2}, \tag{18}$$

where the λ_i are the in *absolute value largest* eigenvalues of the eigenvalue problems

$$|\partial^2_{xx} A^i_k(\hat{x}_{k|k}) Q_{k|k}(\hat{x}_{k|k}) - \lambda_i I| = 0, \quad i = 1, ..., m_k. \tag{19}$$

In (18) it is assumed that Q_y is a diagonal matrix.
We will first consider the distance observable. The corresponding Hessian reads

$$\begin{bmatrix} \partial^2_{uu} l & \partial^2_{uv} l \\ \partial^2_{vu} l & \partial^2_{vv} l \end{bmatrix} = \frac{1}{l} \begin{bmatrix} v/l \\ -u/l \end{bmatrix} \begin{bmatrix} v/l \\ -u/l \end{bmatrix}^*,$$

and its singular value decomposition is given as

$$\begin{bmatrix} \partial^2_{uu} l & \partial^2_{uv} l \\ \partial^2_{vu} l & \partial^2_{vv} l \end{bmatrix} = \begin{bmatrix} cosa & sina \\ -sina & cosa \end{bmatrix} \begin{bmatrix} 1/l & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} cosa & -sina \\ sina & cosa \end{bmatrix} \tag{20}$$

where $sina = u/l$ and $cosa = v/l$. Note that the distance-Hessian is of rank one. If we substitute (20) into (19), it can be shown that the in absolute value largest eigenvalue is given as $\lambda = \hat{l}_{k|k}\sigma^2_{\hat{a}_{k|k}}$. From this follows that the convergency factor of the iterated extended Kalman filter for a single distance observable is bounded from above as

$$CF \leq \left[\frac{\hat{l}_{k|k}\sigma_{\hat{a}_{k|k}}}{\sigma_{l_k}}\right]^2 \left|\frac{l_k}{\hat{l}_{k|k}} - 1\right| \tag{21}$$

We will now consider the azimuth observable. The corresponding Hessian reads

$$\begin{bmatrix} \partial^2_{uu}a & \partial^2_{uv}a \\ \partial^2_{vu}a & \partial^2_{vv}a \end{bmatrix} = \frac{1}{l^4}\begin{bmatrix} -2uv & u^2 - v^2 \\ u^2 - v^2 & 2uv \end{bmatrix},$$

and its singular value decomposition is given as

$$\begin{bmatrix} \partial^2_{uu}a & \partial^2_{uv}a \\ \partial^2_{vu}a & \partial^2_{vv}a \end{bmatrix} = \tag{22}$$

$$\begin{bmatrix} -cos(\frac{1}{4}\pi + a) & sin(\frac{1}{4}\pi + a) \\ sin(\frac{1}{4}\pi + a) & cos(\frac{1}{4}\pi + a) \end{bmatrix} \begin{bmatrix} 1/l^2 & 0 \\ 0 & -1/l^2 \end{bmatrix} \begin{bmatrix} -cos(\frac{1}{4}\pi + a) & sin(\frac{1}{4}\pi + a) \\ sin(\frac{1}{4}\pi + a) & cos(\frac{1}{4}\pi + a) \end{bmatrix}$$

Note that the azimuth-Hessian is of rank two.
If we substitute (22) into (19), it can be shown after a considerable derivation that the in absolute value largest eigenvalue is bounded from above as $\lambda \leq [\hat{\mu}/\hat{l}_{k|k}]^2$ where $\hat{\mu}$ is the largest singular value of $Q_{k|k}(\hat{x}_{k|k})$. Hence, for the convergency factor of the iterated extended Kalman filter for a single azimuth observable follows that

$$CF \leq \left[\frac{\hat{\mu}}{\hat{l}_{k|k}}\right]^2 \left|\frac{a_k - \hat{a}_{k|k}}{\sigma_{a_k}}\right| \tag{23}$$

VI. References

Gelb, A. (Ed.) (1974): Applied Optimal Estimation. MIT Press, Cambridge MA.

Teunissen, P.J.G. (1984): A Note on Gauss' Formulas in Nonlinear Geodesic Adjustment. Statistics & Decisions, Suppl. Issue No. 2, p. 455-466.

Teunissen P.J.G. (1985): The Geometry of Geodetic Inverse Linear Mapping and Nonlinear Adjustment, Netherlands Geodetic Commission, Publications on Geodesy, New Series, Vol. 8, No. 1, Delft.

Teunissen, P.J.G. (1989): First and Second Moments of Nonlinear Least-Squares Estimatros, Bulletin Geodesique, 63, p. 253-262.