2008 IEEE/RSJ International Conference on Intelligent Robots and Systems
Acropolis Convention Center
Nice, France, Sept, 22-26, 2008

Avoiding Negative Depth in Inverse Depth Bearing-Only SLAM

Martin P. Parsley and Simon J. Julier

Abstract— **In this paper we consider ways to alleviate negative estimated depth for the inverse depth parameterisation of bearing-only SLAM. This problem, which can arise even if the beacons are far from the platform, can cause catastrophic failure of the filter. We consider three strategies to overcome this difficulty: applying inequality constraints, the use of truncated second order filters, and a reparameterisation using the negative logarithm of depth. We show that both a simple inequality method and the use of truncated second order filters are succesful. However, the most robust peformance is achieved using the negative log parameterisation.**

I. INTRODUCTION

The work in this paper is motivated by a subset of Simultaneous Localisation and Mapping (SLAM) called *bearing-only* SLAM. This is SLAM using a bearing-only sensor, as opposed to a range-bearing sensor, and is more difficult than the range-bearing case [1]–[3].

There has been a recent surge of interest in bearing-only SLAM because it allows the use of a single camera as a sensor. A camera has a number of advantages over range-bearing sensors (such as SICK scanners). It is a cheap and compact sensor, which is important in emerging areas such as wearable robotics, telepresence and television [3]–[5], where low cost and mobility are key. In addition, being a passive sensor, it will not interfere with any devices in its environment and is eye safe. A camera may also leverage a growing number of computer vision techniques when performing data association, and may recognise certain features of its environment [6]–[9]. An example is the use of image data to detect when loop closing has occured [10].

However, Bearing-only SLAM suffers from *limited observability* when attempting to initialise a new beacon (landmark) into the map [1], [9], [11]. Specifically, the observation model is non-invertible, so a single observation is insufficient to define a unique point in space [12]. Rather, it defines a ray emanating from the sensor along which a beacon may lie. The naïve approach of representing the initially unknown range by a large variance leads to linearisation issues, and difficulties in being processed correctly by the EKF [2], [3], [12].

Several approaches have been developed for performing beacon initialisation in bearing-only SLAM. *Delayed initialisation* delays initialising a beacon into the state until the depth can be estimated well [84]. However this discards valuable orientation information. *Range-parameterised* (RP) approaches treat the unknown depth as a multiple hypothesis problem, however the storage and computation of multiple hypotheses cannot easily be solved without introducing sub-optimality or consistency issues [9], [13], [14]. Recently, Davison proposed the *inverse depth parameterisation* (ID) [15]. This is based on the observation that the filter equations are less nonlinear if the filter parameterises the *inverse* of the distance to the beacon, instead of the distance itself. The method has the advantage of being able to cope with a large range of depths, from the camera position to infinity (i.e. sufficiently far away relative to the camera motion that no parallax is observed). In addition the measurement equation has low linearisation error at low parallax, allowing the inverse depth estimation uncertainty to be accurately modelled as a Gaussian. As a result, it has become a very popular means of performing bearing-only SLAM. However as will be shown later, it suffers from an issue with negative depth.

ID and RP approaches suffer from a beacon "pull-in effect", another serious problem which can also cause failure. This well known but seldom documented effect occurs as a result of apparent parallax that is actually the result of noise, and is more pronounced under certain platform movement [16]. While this paper does not address this problem, we will comment on the robustness of our algorithms to this effect.

The structure of the paper is as follows. Section II gives an overview of inverse depth parameterisation, and discusses the negative depth problem. Section III describes the use of second order filters. Section IV presents two methods to constrain the inverse depth to be positive; *truncation* of the inverse depth distribution, and *translation* of the inverse depth. Section V presents an alternative parameterisation using logarithms. Section VI compares the parameterisations and the effect of the second order filter on simulation data. Finally, Section VII presents a summary and conclusions.

II. INVERSE DEPTH PARAMETERISATION

Consider a vehicle performing EKF-based planar 2D bearing-only SLAM. At time k, the true state $\mathbf{x}(k)$ consists of the vehicle pose $\mathbf{x}_v(k)$ and the set of n static beacons $\mathbf{x}_{1\ldots n}$,

$$\mathbf{x}(k) = \begin{bmatrix} \mathbf{x}_v^T & \mathbf{x}_1^T & \ldots & \mathbf{x}_n^T \end{bmatrix}_k^T.$$

The mean and covariance of this estimate are

$$\hat{\mathbf{x}}(i|j) = \begin{bmatrix} \hat{\mathbf{x}}_v^T & \hat{\mathbf{x}}_1^T & \ldots & \hat{\mathbf{x}}_n^T \end{bmatrix}_{i|j}^T, \quad (1)$$

Martin Parsley was supported by a UK Research Council EPSRC Doctoral Training Award number EP/P502802/1.

M. Parsley and S. Julier are with the Department of Computer Science, University College London, Gower Street, London, WC1E 6BT, UK M.Parsley@cs.ucl.ac.uk; S.Julier@cs.ucl.ac.uk

$$\mathbf{P}(i|j) = \begin{bmatrix} \mathbf{P}_v & \mathbf{P}_{v1}^T & \dots & \mathbf{P}_{vn}^T \\ \mathbf{P}_{v1} & \mathbf{P}_{11} & \dots & \vdots \\ \vdots & \vdots & \ddots & \mathbf{P}_{1n}^T \\ \mathbf{P}_{vn} & \mathbf{P}_{1n} & \dots & \mathbf{P}_{nn} \end{bmatrix}_{i|j}, \quad (2)$$

where $(i|j)$ denotes the estimate at time i given observations up to and including time j [17]. $\mathbf{P}_v$ is the vehicle pose covariance, $\mathbf{P}_{nn}$ is the nth beacon covariance, $\mathbf{P}_{vj}$ is the cross correlation between the vehicle and the jth beacon, and $\mathbf{P}_{nj}$ is the cross correlation between the nth and jth beacons. The vehicle state $\hat{\mathbf{x}}_v$ is

$$\hat{\mathbf{x}}_v(i|j) = \begin{bmatrix} \hat{x}_v & \hat{y}_v & \hat{\theta}_v \end{bmatrix}_{i|j}^T,$$

with the associated covariance

$$\mathbf{P}_v(i|j) = \begin{bmatrix} P_x & P_{xy}^T & P_{x\theta}^T \\ P_{xy} & P_y & P_{y\theta}^T \\ P_{x\theta} & P_{y\theta} & P_\theta \end{bmatrix}_{i|j}.$$

Assume that at time m the vehicle first observes a new beacon j whose true Cartesian position is $\mathbf{x}_j^c$. The bearing from $\mathbf{x}_v(m)$ to $\mathbf{x}_j^c$ is θ_j. However, the distance d_j is not known. The estimated state $\hat{\mathbf{x}}(m|m)$ is augmented according to

$$\begin{aligned} \hat{\mathbf{x}}(m|m) &= \begin{bmatrix} \hat{\mathbf{x}}_{j-1}^T & \hat{\mathbf{x}}_j^T \end{bmatrix}_{m|m}^T \\ \hat{\mathbf{x}}_j(m|m) &= \begin{bmatrix} \hat{x}_v \\ \hat{y}_v \\ \hat{\rho} \\ \hat{\theta}_j \end{bmatrix}_{m|m}, \end{aligned} \quad (3)$$

where $\hat{\rho}$ is the inverse depth estimate, ϕ_j is the initial bearing observation with variance R_ϕ, and $\hat{\theta}_j = \hat{\theta}_v + \phi_j$. $\hat{\rho}$ is chosen to represent the mean of the inverse depth ranges we expect beacons to be at; for a beacon whose range lies between d_{min} and ∞, suitable heuristically-derived values from [15] are $\hat{\rho} = \frac{1}{2} d_{min}^{-1}$ and $R_\rho = \frac{1}{16} d_{min}^{-2}$ (for a 95% uncertainty region). The initial covariance is

$$\begin{aligned} \mathbf{P}_j(m|m) &= \begin{bmatrix} P_x & p_{xy}^T & 0 & P_{x\theta}^T \\ p_{xy} & P_y & 0 & P_{y\theta}^T \\ 0 & 0 & R_\rho & 0 \\ P_{x\theta} & P_{y\theta} & 0 & P_\theta + R_\phi \end{bmatrix}_{m|m}, \\ \mathbf{P}'(m|m) &= \begin{bmatrix} \mathbf{P} & \mathbf{P}\nabla^T \wp \\ \mathbf{P}\nabla\wp & \mathbf{P}_j \end{bmatrix}_{m|m}, \end{aligned} \quad (4)$$

where $\nabla\wp$ is the Jacobian of (3), $\mathbf{P}'(m|m)$ is the new covariance for the entire state, R_ϕ is the variance of the bearing sensor, and R_ρ is the initial inverse depth variance, chosen to encompass the inverse depth uncertainty region.

The observation function for the jth beacon is

$$\phi = \mathbf{h}_j\left[\mathbf{x}(k), v(k)\right],$$

where $v(k)$ is zero-mean observation noise with known variance R_ϕ.

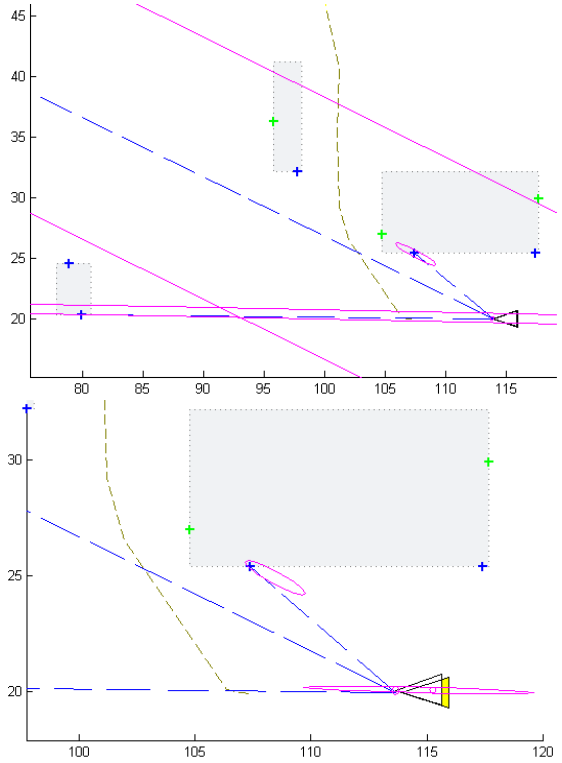

Fig. 1. Example showing failure due to the inverse depth of two beacons becoming negative. (Top) Close-up of the map just prior to failure. (Bottom) Following the update the inverse depths ρ have become negative, which maps to the positions shown behind the sensor in depth space (note the barely visible ellipse inside the estimated vehicle triangle).

For ID this is

$$\mathbf{h}\left[\mathbf{x}(k), v(k)\right] = \arctan\left(\frac{y_j - y_v + \sin(\theta_j)\hat{\rho}_j^{-1}}{x_j - x_v + \cos(\theta_j)\hat{\rho}_j^{-1}}\right)_k - \theta_v(k) + v(k). \quad (5)$$

Once linearity is achieved (using the test in [18]), the inverse depth is transformed into cartesian cartesian coordinates using

$$\hat{\mathbf{x}}_j^c = \mathbf{g}(\hat{\mathbf{x}}_j) = \begin{bmatrix} \hat{x}_v + \cos(\hat{\theta}_j)\hat{\rho}_j^{-1} \\ \hat{y}_v + \sin(\hat{\theta}_j)\hat{\rho}_j^{-1} \end{bmatrix}, \quad (6)$$

and the covariance transformed using

$$\mathbf{P}^c(k|k) = \nabla\mathbf{G}(k)\,\mathbf{P}(k|k)\,\nabla^T\mathbf{G}(k),$$

where $\nabla\mathbf{G}$ is the Jacobian of $\mathbf{g}(\cdot)$, and the c superscript denotes this as the cartesian estimate.

The ID parameterisation is only defined for positive depths. However, the Kalman filter update rule is linear and takes no account of nonlinear constraints. Therefore, the depth can become negative after an update, resulting in catastrophic failure. This is illustrated in Fig. 1, which shows a vehicle observing three beacons while performing 2D planar ID SLAM. After a single update the ranges of two

of the beacons become negative, and the covariance ellipses in depth space collapse behind the vehicle. In this paper we consider three strategies for mitigating this phenomenon. The first is the use of higher order filtering algorithms, the second explores methods for enforcing inequality constraints to filters, and the third approach involves reparameterising the depth using logarithms.

III. Higher order filters

Because the EKF performs a first-order linearization of the nonlinear system, the output distribution is often not adequately described by the mean and covariance alone. This can introduce large errors in the true posterior mean and covariance, which in the case of the parameterisations described in this paper may degrade their quality and lead to negative depth in ID.

Including second order terms in the Kalman update may reduce this effect, and increase the accuracy of the filter. Work has been done to include higher order terms in SLAM using the Unscented filter (UKF) [19]. Holmes [20] has obtained better results this way, suggesting this strategy could help to avoid the problem. However, computational issues means that the UKF does not enjoy many of the optimisation strategies used with SLAM [21], [22]. Various schemes [23] have been deployed to address this, however they lack theoretical rigour because correlation structure between the vehicle and beacons, which is very important, is not completely consistently maintained.

We use the alternative approach of truncated second order filters [24]. Although these require the Hessian to be computed, the sparseness of the observation Jacobian $\nabla\mathbf{H}$ means that the Hessian tensor is also sparse. The computational form of the Gaussian second-order update is

$$\begin{aligned}
\mathbf{S} &= \nabla\mathbf{H}(k)\mathbf{P}(k|k-1)\nabla\mathbf{H}^T(k)+\mathbf{R}(k)+\tfrac{1}{2}(\partial^2 h\mathbf{P}^2\partial^2 h)\\
\hat{\mathbf{x}}(k|k) &= \hat{\mathbf{x}}(k|k-1)+\mathbf{P}(k|k-1)\nabla\mathbf{H}^T(k)S^{-1}\\
&\quad \times[v-\tfrac{1}{2}(\mathbf{P}(k|k-1)\partial^2 h)]\,,\\
\mathbf{P}(k|k) &= \mathbf{P}(k|k-1)-\mathbf{P}(k|k-1)\nabla\mathbf{H}^T(k)\mathbf{S}^{-1}\nabla\mathbf{H}(k)\mathbf{P}(k|k-1)
\end{aligned}$$

where $\mathbf{S}$ is the innovation covariance and v is the innovation. The structure of the Hessian, showing its sparsity, for a beacon b is

$$\mathbf{M}_b=\begin{bmatrix}
\mathbf{M}_{vv} & \mathbf{0} & \dots & \mathbf{M}_{vb}^T & \dots & \mathbf{0}\\
\mathbf{0} & 0 & \dots & 0 & \dots & 0\\
\vdots & \vdots & \ddots & \vdots & \ddots & \vdots\\
\mathbf{M}_{vb} & 0 & \dots & \mathbf{M}_{bb} & \dots & 0\\
\vdots & \vdots & \ddots & \vdots & \ddots & \vdots\\
\mathbf{0} & 0 & \dots & 0 & \dots & 0
\end{bmatrix},$$

where $\mathbf{M}_{vv}$ and $\mathbf{M}_{bb}$ are the vehicle and beacon derivatives respectively, and $\mathbf{M}_{vb}$ are the mixed derivatives. Although the equations are inconvenient to implement, the overall computational costs remain $O(n^3)$.

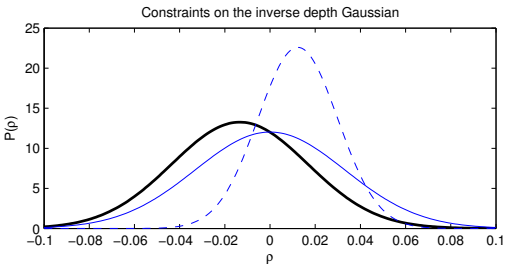

Fig. 2. Example from a run showing the application of the two constraint approaches. The negative ID Gaussian (thick line) is either truncated (dashed) or translated and inflated (solid).

Although the Second Order filter is theoretically more accurate, it does not in itself guarantee avoidance of negative inverse depth.

IV. Inequality Constraints

This section considers strategies that constrain negative inverse depths $\hat{\rho}$ to be positive.

A. Truncating the Inverse Depth Distribution

Shimada has developed a method to apply inequality constraints based on the observation that the probability outside of the constraints must be zero [25]. Assuming the unconstrained distribution is Gaussian, the first two moments of the constrained distribution can be readily calculated. The effect is illustrated in Fig. 2.

Following an update, for all beacons with ρ close to 0 or negative we apply a constraint of the form $\Phi^T\hat{\mathbf{x}} \leq \mathbf{0}$, where Φ is a vector the same size as $\hat{\mathbf{x}}$, with -1 for elements corresponding to negative inverse depths and zeros elsewhere.

B. Enforcing a Non-Negative Depth

The previous section considered truncation as an information gain. However, another way to consider the problem is that the EKF $\hat{\rho}$ estimate is the linear minimum mean squared error (MMSE), thus changing it increases the covariance.

Following an update, the following treatment is applied to the state and covariance to make all negative depths positive. An error term $\mathbf{n}(k)$ is computed, being the same size as the state $\hat{\mathbf{x}}$, with zeros in all elements except those that correspond to inverse depths $\hat{\rho} \leq \epsilon$,

$$\begin{aligned}
\mathbf{n}_i(k) &\in \begin{cases} \epsilon - \hat{\rho}_j(k|k) & \text{for } \hat{\rho}_j(k|k) < 0\\ \epsilon_i & \text{otherwise}\end{cases}\\
\hat{\mathbf{x}}'(k|k) &= \hat{\mathbf{x}}(k|k)+\mathbf{n}(k),\\
\mathbf{P}'(k|k) &= \mathbf{P}(k|k)+\mathbf{n}(k)\,\mathbf{n}(k)^T,
\end{aligned}$$

where ϵ is a small positive number (we used 10^{-6}) instead of 0 as (5) is not observable when $\hat{\rho}=0$.

Thus the mean of negative inverse depths is translated to ϵ, and the covariance is inflated to compensate, as shown by the thin solid line in Fig. 2. Note that unlike truncation,

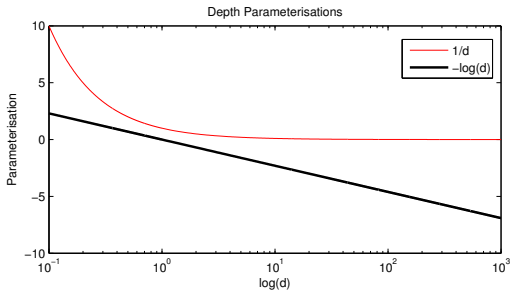

Fig. 3. Plot showing the inverse depth and negative log parameterisations.

negative depth is not assumed to provide depth distribution information in this case. Theoretically, it is possible for the covariance of a beacon to grow without bound if its inverse depth repeatedly becomes negative after each update, such that the covariance increases faster than the Kalman update can reduce it. However, we have not seen evidence of this occuring in practice.

Rather than correct the estimate when negative depths occur, an alternative approach is to parameterise the depth such that negative values cannot occur. In the next section we discuss one such parameterisation, the negative log parameterisation.

V. Negative Log Parameterisation

The negative log parameterisation is the same as in (3)-(5), except instead of the inverse depth ρ, we parameterise the depth as $d = e^{-l}$. Thus the equivalent estimated state (3) for the beacon j, first seen at time m is

$$\hat{\mathbf{x}}_j(m|m) = \begin{bmatrix} \hat{x}_v \\ \hat{y}_v \\ \hat{l} \\ \hat{\theta}_j \end{bmatrix}_{m|m} . \tag{7}$$

The parameterisation is shown in Fig. 3. Unlike inverse depth, the negative log curve is not asymptotic as depth increases and thus cannot represent the full $(0 \ldots \infty)$ range of depths that inverse depth can. The parameterisation is conceptually straightforward, as its components are the same as for inverse depth, with the exception that the depth d is given by $d = e^{-l}$ instead of ρ^{-1}. Thus the initialisation of a new beacon into the state follows that in (3) and (4), with $\hat{\rho}$ replaced by $\hat{l}$. The initial inverse depth estimates $\{\hat{\rho}_i, R_\rho\}$ are replaced by their equivalents $\{\hat{l}_i, R_l\}$, which likewise are computed heuristically to encompass the mean $\hat{l}_i$ and variance R_l of the uncertainty region in l space.

Beacon updates and conversion to cartesian form are the same as in (5) and (6) respectively, using e^{-l} instead of ρ^{-1}.

The disadvantage of the negative log parameterisation is that it is not asymptotic with increasing depth. Therefore, the initial depth and covariance estimate cannot encompass as high a range as inverse depth can, i.e. $d = (0, n]$ rather than $(0, \infty)$, where n is a large number; for our experiments with $d_{max} \gtrsim 10^3$ the incidence of failures increases and the accuracy decreases compared to inverse depth.

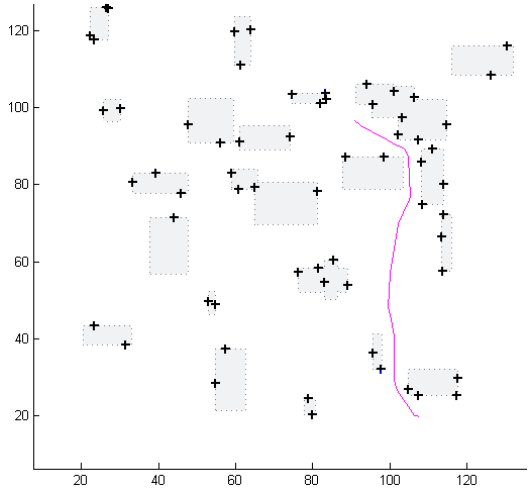

Fig. 4. Map of the run, in metres. The solid line represents the trajectory of the vehicle, which starts at the bottom. The crosses are beacons. The dashed lines represent building outlines that cause occlusions. These ensure that observations are made at a variety of ranges.

VI. Results

Because simulation studies provide ground truth, and a controlled environment for comparing the algorithms, we performed 2D simulations of a vehicle performing bearing-only SLAM with known data association in the environment shown in Fig. 4. The trajectory covers 106m at a constant speed of 14ms^{-1}. The true mean range at which beacons are first initialised has a mean and variance of 45.6m and 646m^2 respectively. The Monte Carlo simulations were run until $N = 1000$ successful runs were obtained for each filter. A run was unsuccessful (i.e. failed) if the (Gaussian-assumed) observation likelihood [26] became less than 10^{-100} (effectively zero) at any time step; failed runs are not represented in the results. The average of the N successful runs was then computed from

$$\begin{aligned} \hat{\mathbf{x}}(k|k) &= \tfrac{1}{N}\sum_{h=1}^{N} \hat{\mathbf{x}}^h(k|k)\,, \\ \mathbf{P}(k|k) &= \tfrac{1}{N}\sum_{h=1}^{N} \mathbf{P}^h(k|k)\,. \end{aligned}$$

We computed the normalised estimation error squared (NEES) at each time step, and averaged them over the N successful runs according to [27].

To investigate the contribution that including second order terms can make, we performed all the experiments both with and without second order updates. The vehicle was modelled as a steered bicycle [28] receiving control inputs $u_v(k) = 14\text{m/s}$ and steer angle $u_s(k)$ at each time step, with associated errors v_v and v_s. The linearised process model $\nabla\mathbf{F}(k)$ and noise $\mathbf{Q}(k)$ for the vehicle are the same as in [28].

$\triangle t$, the time elapsed between $k-1$ and k was 0.02s for our simulations and the vehicle wheelbase $B = 2\text{m}$. The sensor range of the vehicle was 100m and the sensor sweep $\pm 60°$. The velocity and steer noise standard deviations were 0.5m and $1°$ for the vehicle, and $1.15°$ for the bearing sensor.

The initial parameterised depth distribution was computed heuristically by parameterising 100 depths uniformly distributed between 1 and 100m, then computing the mean and covariance of the distribution. This gives better results for ID than using the distribution in [15], and makes the ID and negative log parameterisations directly comparable on this trajectory.

A. Inverse depth results

The Gaussian $\{\hat{\rho}_i, R_\rho\}$ values for inverse depth are $\{0.0519, 0.0138\}$.

The environment causes the standard inverse depth parameterisation to consistently fail (i.e. 100% failure rate) due to negative inverse depth at the beginning of the run, as illustrated in Fig. 1. This occured regardless of whether the Second Order filter was used. The truncated Gaussian constraint (Section IV-A) did not prevent failure, though it generally delayed it for up to a third of the run. The truncated Gaussian constraint fails because the decrease in the inverse depth covariance as a result of truncation causes convergence to an incorrect depth, suggesting that the concept of truncation (that negative depth is informative) is inappropriate. The non-negative inverse depth constraint described in Section IV-B was the only method successful at preventing negative depth failure, and always converged in our study. Thus, all the quantitative inverse depth results in this paper use this method. When we tried the initial depth distribution from [15], we came to the same conclusions. Table II summarises the main strategies used, and for the effective ones ranks their ability to increase robustness to failure by beacon pull-in, the dominant failure mode after negative depth.

B. Negative log results

The Gaussian $\left\{\hat{l}_i, R_l\right\}$ values for negative log are $\{-3.64, 0.861\}$. The ID depth bounds above could have been chosen to encompass ∞, however as currently posed, the negative log distribution cannot cope with an equivalently large depth space; for $d_{max} \gtrsim 10^3$ the performance would become worse than inverse depth. The negative log parameterisation appears to be sensitive to the choice of $\{\hat{l}_i, R_l\}$ (compared to the actual beacon depth distribution) to a greater extent than inverse depth. This is likely to be due to the the Gaussianity assumptions of the parameterisation breaking down due to the shallower gradient of the parameterisation over ID. For environments with a greater depth range, the negative log parameterisation would need to be modified, either by remapping the depth space, or by using an alternative distribution; a possibility is $d = ae^{-bl}$, where $a, b > 1$.

Fig. 5 shows that both parameterisations return consistent estimates with similar accuracy for the vehicle pose. The state estimate of the negative log parameterisation shows greater consistency than inverse depth; Fig. 6 shows that the NEES estimate is an order of magnitude lower. It also shows

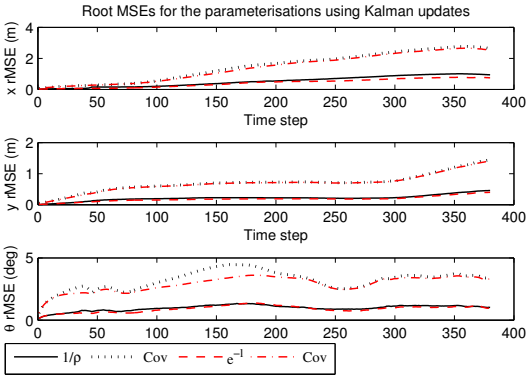

Fig. 5. Comparison between the inverse depth and negative log parameterisations using a Kalman update. Both estimates have similar accuracy and uncertainty estimates (2σ bounds are shown as dashed lines).

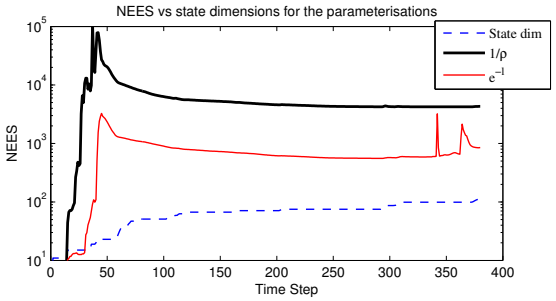

Fig. 6. NEES compared to the state size for the parameterisations.

that the period between 0 and 50 time steps is particularly difficult for inverse depth, as there are few observations with low parallax.

Table I shows the failure rate of the runs. As negative depth has been eliminated, the failures are primarily caused by beacon pull-in, mainly between 0 and 50 time steps. The low failure rate of the negative log parameterisation shows that it is far more robust to failures than inverse depth.

The uncertainty estimate and accuracy of the vehicle poses using the second order filter are almost identical to the first order case, however the state estimates become vastly more consistent, as shown in Fig. 7. The negative log parameterisation remains more consistent than inverse depth. Neither parameterisation failed at all when using the second order filter.

VII. Conclusion

In this paper we have considered several strategies for overcoming the problem of negative depth in the inverse

TABLE I

FAILURE RATE OF THE PARAMETERISATIONS, MEASURED IN THE COURSE OF OBTAINING 1000 SUCCESSFUL RUNS.

Parameterisation	Failures	Failure rate (%)
ρ^{-1} Kalman update	65	6.1
e^{-l} Kalman update	7	0.7
ρ^{-1} Second Order	0	0
e^{-l} Second Order	0	0

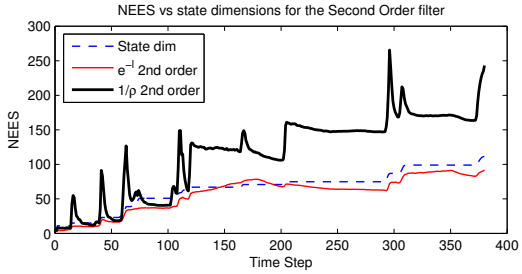

Fig. 7. NEES compared to the state size for the parameterisations using a Second Order filter

TABLE II

SUMMARY OF THE METHODS USED AND THEIR EFFECTIVENESS

Method	Effectiveness
Unconstrained Second Order filter	ineffective
Inequality constraints	ineffective
Ineq. constraints + Second Order filter	ineffective
Translating the depth	moderate
Translating the depth + Second Order filter	high
Negative log parameterisation	high

depth parameterisation of bearing-only SLAM. We examined three different strategies – second order filters, implementations of inequalities and reparameterisations. Their effectiveness is summarised in Table II. We found that it is worth exploring alternative parameterisations to inverse depth. The benefits include an increase in consistency and robustness to pull-in failure, in addition to the elimination of the negative depth failure mode. We have used a basic e^{-l} parameterisation to remove the discontinuity, and found that it is as accurate as inverse depth, with improved consistency in our simulations. This parameterisation has the shortcoming of being unable to cope with very large variations in depth observations (as may be encountered when running outdoors). However, other similar parameterisations, for instance $d = ae^{-bl}$, where $a, b > 1$, may be tailored to avoid this.

In further work, the cause of negative depth, and the convergence of inverse depth with the positive depth constraint need to be investigated. The linearity of the negative log parameterisation also needs to be investigated, and the parameterisation tested on an environment with a greater range of beacon depths. A multiple Gaussian approach, in the same manner as Range Parameterisation, could be used instead of a single Gaussian, as the parameterisation used to represent each depth interval could be tailored to reduce nonlinearity and provide optimum performance.

REFERENCES

[1] T. Bailey, "Constrained initialisation for bearing-only slam," in *Robotics and Automation, 2003. Proc. ICRA '03. IEEE International Conference on*, vol. 2, 2003, pp. 1966–1971 vol.2.

[2] K. Bekris, M. Glick, and L. Kavraki, "Evaluation of algorithms for bearing-only slam," in *Robotics and Automation, 2006. ICRA 2006. Proc. 2006 IEEE International Conference on*, 2006, pp. 1937–1943.

[3] N. Kwok and G. Dissanayake, "An efficient multiple hypothesis filter for bearing-only slam," in *Intelligent Robots and Systems, 2004. (IROS 2004). Proc. 2004 IEEE/RSJ International Conference on*, vol. 1, 2004, pp. 736–741 vol.1.

[4] A. Davison, "Real-time simultaneous localisation and mapping with a single camera," in *Computer Vision, 2003. Proc. Ninth IEEE International Conference on*, 2003, pp. 1403–1410 vol.2.

[5] A. Davison, Y. G. Cid, and N. Kita, "Real-time 3D SLAM with wide-angle vision," in *Proc. IFAC Symposium on Intelligent Autonomous Vehicles, Lisbon*, Jul. 2004.

[6] V. Lepetit and P. Fua, "Keypoint recognition using randomized trees," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 28, no. 9, pp. 1465–1479, 2006.

[7] S. Thrun, *Robotic mapping: a survey*. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 2003.

[8] P. Smith, I. Reid, and A. Davison, "Real-Time Monocular SLAM with Straight Lines," in *British Machine Vision Conference*, vol. 1, September 2006, pp. 17–26.

[9] T. Lemaire, S. Lacroix, and J. Sola, "A practical 3d bearing-only slam algorithm," in *Intelligent Robots and Systems, 2005. (IROS 2005). 2005 IEEE/RSJ International Conference on*, 2005, pp. 2449–2454.

[10] M. Cummins and P. Newman, "Fab-map: Probabilistic localization and mapping in the space of appearance," *The International J. Robotics Research*, vol. 27, no. 6, pp. 647–665, 2008.

[11] S. S. A. S. T. Vidal-Calleja, M. Bryson and J. Andrade-Cetto, "On the observability of bearing only slam," *Proc. the IEEE International Conference on Robotics and Automation*, 2007.

[12] H. Z. J Klippenstein and X. Wang, "Feature initialization for bearing only visual slam using triangulation and the unscented transform," in *Mechatronics and Automation, 2007. Proc. 2007 IEEE International Conference on*. Harbin, China, August 5-8, 2007, 2007, pp. 157–164.

[13] M. Parsley and S. Julier, "The common state filter for slam," February 2008, accepted for publication in IROS 2008.

[14] M. Bryson and S. Sukkarieh, "Building a robust implementation of bearing-only inertial slam for a uav," *J. Field Robotics*, vol. 24, no. 1-2, pp. 113–143, 2007.

[15] J. Montiel, J. Civera, and A. Davison, "Unified Inverse Depth Parametrization for Monocular SLAM," in *Robotics: Science and Systems*, August 2006.

[16] A. J. D. Javier Civera and J. M. M. Montiel, "Interacting multiple model monocular slam," in *Robotics and Automation, 2008. Proc. ICRA '08. IEEE International Conference on*, 2008.

[17] X.-R. Li and Y. Bar-Shalom, "Multiple-model estimation with variable structure," *Automatic Control, IEEE Trans.*, vol. 41, no. 4, pp. 478–493, Apr 1996.

[18] J. Civera, A. Davison, and J. Montiel, "Inverse Depth to Depth Conversion for Monocular SLAM," in *IEEE International Conference on Robotics and Automation*, April 2007.

[19] S. Julier and J. Uhlmann, "A new extension of the kalman filter to nonlinear systems," in *Int. Symp. Aerospace/Defense Sensing, Simul. and Controls*, 1997.

[20] S. Holmes, G. Klein, and D. W. Murray, "A square root unscented kalman filter for visual monoslam," in *Robotics and Automation, 2008. Proc. ICRA '08. IEEE International Conference on*, 2008.

[21] S. Thrun, Y. Liu, D. Koller, A. Y. Ng, Z. Ghahramani, and H. Durrant-Whyte, "Simultaneous Localization and Mapping with Sparse Extended Information Filters," *The International J. Robotics Research*, vol. 23, no. 7-8, pp. 693–716, 2004.

[22] J. Knight, A. Davison, and I. Reid, "Towards constant time slam using postponement," *Intelligent Robots and Systems, 2001. Proceedings. 2001 IEEE/RSJ International Conference on*, vol. 1, pp. 405–413 vol.1, 2001.

[23] S. Challa, R. J. Evans, and X. Wang, "A bayesian solution and its approximations to out-of-sequence measurement problems," *Information Fusion*, vol. 4, no. 3, pp. 185–199, Sep. 2003.

[24] A. H. Jazwinski, *Stochastic Processes and filtering theory*. Academic Press, 1970.

[25] N. Shimada, Y. Shirai, Y. Kuno, J. Miura, "Hand Gesture Estimation and Model Refinement Using Monocular Camera — Ambiguity Limitation by Inequality Constraints," in *Proc. The 3rd International Conference on Automatic Face and Gesture Recogni tion*, 1998, pp. 268–273.

[26] S. B. Williams, H. Durrant-Whyte, and G. Dissanayake, "Constrained Initialization of the Simultaneous Localization and Mapping Algorithm," *The International J. Robotics Research*, vol. 22, no. 7-8, pp. 541–564, 2003.

[27] Y. Bar-Shalom and X.-R. Li, *Estimation with Applications to Tracking and Navigation*. New York, NY, USA: John Wiley & Sons, Inc., 2001.

[28] S. Clark and H. Durrant-Whyte, "Autonomous land vehicle navigation using millimeter wave radar," in *Robotics and Automation, 1998. Proc. 1998 IEEE International Conference on*, vol. 4, 1998, pp. 3697–3702 vol.4.