A Robust Docking Strategy for a Mobile Robot using Flow Field Divergence

Chris McCarthy, *Member, IEEE,* Nick Barnes, *Member, IEEE* and Robert Mahony, *Senior Member, IEEE*

***Abstract*— We present a robust strategy for docking a mobile robot in close proximity with an upright surface using optical flow field divergence and proportional feedback control. Unlike previous approaches, we achieve this without the need for explicit segmentation of features in the image, and using complete gradient-based optical flow estimation (*i.e.* no affine models) in the optical flow computation. A key contribution is the development of an algorithm to compute the flow field divergence, or time-to-contact, in a manner that is robust to small rotations of the robot during ego-motion. This is done by tracking the *focus of expansion* of the flow-field and using this to compensate for ego rotation of the image. The control law used is a simple proportional feedback, using the unfiltered flow field divergence as an input, for a dynamic vehicle model. Closed-loop stability analysis of docking under the proposed feedback is provided. Performance of the flow field divergence algorithm is demonstrated using off-board natural image sequences, and the performance of the closed-loop system is experimentally demonstrated by control of a mobile robot approaching a wall.**

***Index Terms*— Image motion analysis, optical flow, time-to-contact, focus of expansion, robot vision systems.**

I. INTRODUCTION

DOCKING is an essential capability for any mobile robot seeking to interact with objects in its environment. Tasks such as plugging into a re-charging station, pallet lifting or transporting goods on a factory floor are common tasks requiring some form of docking manoeuvre to be performed. Of particular importance is the control of the robot's deceleration to an eventual halt, close enough to the object that the interaction may take place while also avoiding collision. To achieve this, the robot must acquire a robust estimation of *time-to-contact* (τ), and control the robot's velocity accordingly. The accuracy and robustness of the τ estimate is therefore crucial to the stability, and safety of the robot in performing this task.

For a single, forward-facing camera approaching an upright surface, a common method of estimating τ is to measure the image expansion induced by the apparent motion of the surface towards the camera. This can be obtained from the optical flow field divergence. This image expansion, or *looming effect*, is characterised by flow vectors diverging from a single point in the image known as the focus of expansion (FOE), as shown in Figure 1. The use of visual motion to gauge τ is well supported by observations in biological vision. Srinivasan *et al.* [19] observe how honeybees use visual motion to decelerate and perform smooth graze landings. Lee [6] theorised that a human driver may visually control vehicle braking based on τ estimation obtained from image expansion.

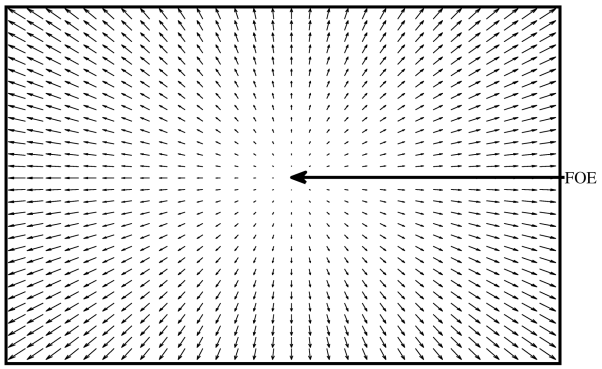

Fig. 1. Diverging optical flow vectors and the focus of expansion (FOE)

Optical flow and flow divergence are commonly used to estimate τ for obstacle avoidance [12], [1], [2]. Few, however, have applied optical flow to tasks requiring finer motion control such as docking. Cipolla and Blake [2], for example, measure τ using divergence computed from the temporal derivative of the moments of area for a closed-contour region of the image. The τ estimator is shown to be sufficiently robust for closed-loop control of collision avoidance on-board a camera-mounted robot arm. The authors note, however, that performance degrades significantly when in close proximity with the target surface due to a break down of the assumed affine motion. Examples where visual motion has been applied explicitly to docking include Santos-Victor and Sandini [15], who apply an affine model of image motion to obtain an approximation from normal flow vectors. Time-to-contact is measured from the inverse of an affine flow parameter and used to control forward velocity while approaching a planar docking surface. Questa *et al.* [13] also use an affine approximation of flow, from which they measure divergence and calculate τ.

An important drawback of these approaches is that all require the explicit segmentation of the surface, and directly estimate image motion from an assumed plane. Where closed-contour deformation is measured, there is also the problem of

Manuscript received October 23, 2006; This paper was presented in part at the IEEE International Conference on Intelligent Robots and Systems, Beijing, China, 2006.

C. McCarthy is with the Embedded Systems Theme, NICTA, and the Department of Information Engineering, Australian National University, Canberra ACT 0200, Australia. N. Barnes is with the Embeded Systems Theme, NICTA, Canberra ACT 2601, Australia. R. Mahony is with the Department of Engineering, Australian National University, Canberra ACT 0200, Australia (email: chris.mccarthy@anu.edu.au; nick.barnes@nicta.com.au; robert.mahony@anu.edu.au).

NICTA is funded by the Australian Government as represented by the Department of Broadband, Communications and the Digital Economy and the Australian Research Council through the ICT Centre of Excellence program.

reliably finding closed shapes when at close proximity with the surface [2].

An alternative approach is to compute τ from general optical flow. Methods for estimating general optical flow fields from local image regions, such as proposed by Lucas and Kanade [8], require no *a priori* knowledge of scene structure, and therefore, no segmentation. In general, for systems such as road vehicles, optical flow is often used for other functions, such as a general sensor for salience to detect moving hazards over the whole scene, as well as for particular functions such as obstacle detection. Affine approximations of image motion are not adequate for this type of general use, and having multiple methods for calculating flow is implausible on restricted embedded hardware. A key requirement for a robust docking control is a τ estimation algorithm based on general optical flow computation.

In much of the previous work with divergence-based τ estimation, divergence is measured at the same image location in each frame [1], [12], [3]. This, however, ignores the effect of FOE shifts on the divergence measure across the image. Mobile robot ego-motion is rarely precise, and even where only translational motion is intended, rotations will be present. Small directional control adjustments, fluctuations in direction due to steering control or differing motor outputs, bumps and undulations along the ground surface, and noisy optical flow estimation will all cause instantaneous, frame-to-frame rotations of the robot. As such, the optical axis will be subject to small rotations about the predominant direction of motion. As a result, the FOE is unlikely to be fixed with respect to the image centre. Given such rotations are likely to be small with respect to the robot's forward motion, the predominant direction of motion should remain constant. Therefore, to ensure consistency in τ estimates over time, we argue that divergence should be measured with respect to the FOE, and not the image centre.

Robustly estimating τ when the optical and translation axes are not physically aligned has been examined previously. Subbarao [20] considers τ with surfaces of arbitrary orientation, for a camera of arbitrary alignment with respect to the direction of motion. Subbarao, however, does not consider the effects of instantaneous rotations during ego-motion, and therefore assumes the point of interest lies along the camera's optical axis. While a fixation-based strategy such as that used by Questa *et al.* [13] can keep the target point centred, a mobile robot is unable to achieve this without additional hardware support. In many cases, such hardware is unavailable to facilitate high speed fixation.

An alternative approach is to account for instantaneous rotations in the image domain, by tracking the location of the FOE. Van Leeuwen and Groen [22], [21] consider the use of FOE tracking to correct for the physical misalignment of the optical and translational axes as a result of the camera-robot configuration. However, while accounting for the constant physical misalignment of these axes, they do not extend the use of FOE tracking explicitly to the removal of small frame-to-frame rotational effects during ego-motion, nor do they apply τ directly to control the vehicle's velocity. In general, while previous work has considered the use of FOE tracking for camera stabilisation during ego-motion, no one has applied such an approach to tasks requiring fine motion control (such as docking), nor provided a theoretical analysis supporting the advantages of such a strategy, and its potential use for control.

In this paper, we present a robust strategy for docking a mobile robot in close proximity with an upright planar surface using optical flow field divergence. Unlike previous approaches, we achieve this without the need for explicit segmentation of the surface in the image, and using complete gradient-based optical estimation (*i.e.* no affine models are used to estimate the optical flow field) in the control loop. In addition, we require only a simple proportional control law to regulate the vehicle's velocity, using only the unfiltered flow field divergence as an input. Central to the robustness of our approach is the derivation of a τ estimator that accounts for small rotations of the robot during ego-motion through tracking of the FOE. We provide a theoretical justification for the constant tracking of the FOE as a means of accounting for not just the physical misalignment of the optical and translational axes, but also frame-to-frame shifts of the optical axis due to instantaneous rotations during ego-motion. The proposed control is designed for the full dynamics of a vehicle, making the results applicable to a wide range of autonomous robotic vehicles. A simple proportional feedback, using the computed flow divergence error as the driving term is chosen. The control is simple to apply but leads to singular, non-linear closed-loop dynamics of the vehicle. A full theoretical analysis is undertaken that proves stability of the closed-loop system under ideal conditions. We present off-board and on-board experiments demonstrating the application of this strategy to the task of docking a mobile robot. Note that this paper extends preliminary results first presented in [10]. Here we provide additional experimental results, and a full analysis of the system's closed-loop stability.

The paper is structured as follows. Section II provides theoretical background, and the derivation of the proposed FOE-based τ estimator outlined above. Section III provides analysis of the technique's stability for the closed-loop control of a mobile robot during docking. Section IV describes all experiments conducted, and results achieved. Section V concludes the paper.

II. Theory

A. Background

Flow divergence is measured by examining the partial spatial derivatives of image velocity components in orthogonal directions at a given image location. This measure is commonly defined as:

$$D(x,y) = \frac{\partial u(x,y)}{\partial x} + \frac{\partial v(x,y)}{\partial y}, \tag{1}$$

where (x, y) is a point in the image, and u and v are the image velocity components in the x and y directions respectively. Time-to-contact (τ) to a point along the optical axis of the camera can be measured from flow divergence, and

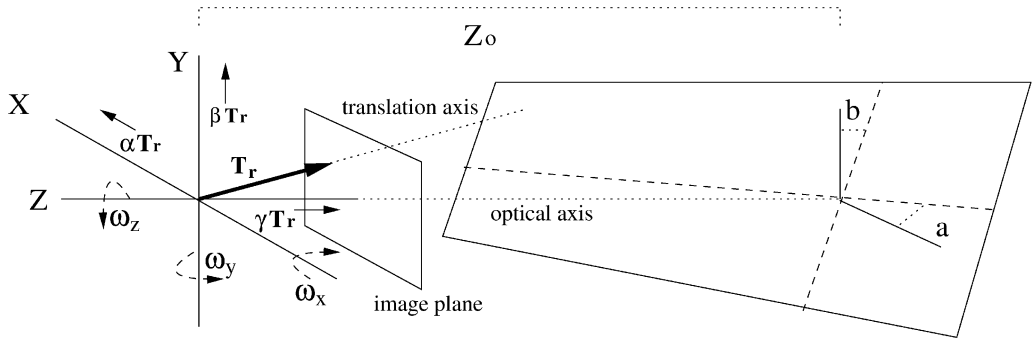

Fig. 2. Geometric configuration

is commonly defined as [3]:

$$\tau = \frac{-Z}{T_r} = \frac{-2}{D(x_0, y_0)}, \tag{2}$$

where Z is the distance to the object in the direction of heading, and T_r is the velocity in this direction. Note that for the typical scenario of a robot approaching a surface, we measure $Z > 0$, and $T_r < 0$, thereby decreasing the value of Z as the robot approaches. In particular, the divergence $D < 0$ is negative for a diverging flow field and τ is defined to be positive. The above relationship between D and τ, assumes the heading direction is aligned with the camera's optical axis, at the image centre: (x_0, y_0).

Flow divergence is constant across the image plane if the surface plane is perpendicular to the camera's optical axis (*i.e.* fronto-parallel with the image plane), and can therefore be calculated anywhere in the imaged area of the surface. If precise fronto-parallel alignment with the docking plane is not maintained then further image deformation is introduced, causing the measured divergence to vary across the projected surface [20].

Given instantaneous rotations during ego-motion, precise surface alignment is unlikely to exist. In the image domain, such effects are characterised by frame-to-frame shifts of the FOE, causing the divergence at any given image location to vary. As a result, (2) is unlikely to provide an accurate estimate of τ in the presence of such rotations. To improve τ estimates during ego-motion, accounting for rotational effects is essential.

B. Derivation of proposed time-to-contact estimator

The analysis presented here extends on the geometric modelling used by Santos-Victor and Sandini [15]. As in [15], we represent the docking surface as a plane in a camera centred coordinate system:

$$Z(X, Y) = Z_0 + aX + bY, \tag{3}$$

where Z_0 is the distance to the surface along the optical axis, X and Y represent points on the surface, and a and b give the slant and tilt with respect to the optical axis. By introducing the perspective projection equations into (3), the surface plane can be expressed as a function of the image coordinates, (x, y) [14]:

$$Z(x, y) = \frac{Z_0}{1 - a\frac{x}{f_\mathrm{x}} - b\frac{y}{f_\mathrm{y}}}, \tag{4}$$

where f_x and f_y are focal lengths expressed in pixels.

Given a fixed camera with respect to the robot's direction of motion, we represent the translational velocity of the camera, T_c, as proportions of the forward translational velocity, T_r, of the robot:

$$T_c = \left[\begin{array}{ccc} \alpha T_r & \beta T_r & \gamma T_r \end{array}\right]. \tag{5}$$

The camera's angular velocity (ω_c) is given by:

$$\omega_c = \left[\begin{array}{ccc} \omega_x & \omega_y & \omega_z \end{array}\right], \tag{6}$$

where each component represents rotation about the axis indicated by its subscript. Figure 2 shows the geometric configuration.

The optical flow induced by the apparent motion of the docking plane is defined by the well known equations [15]:

$$u(x,y) = f_\mathrm{x}\left[\frac{\gamma T_r(\frac{x}{f_\mathrm{x}} - \alpha)}{Z(x,y)} + \omega_x\frac{xy}{f_\mathrm{x}f_\mathrm{y}} - \omega_y(1 + \frac{x^2}{f_\mathrm{x}^2}) + \omega_z\frac{y}{f_\mathrm{y}}\right], \tag{7}$$

$$v(x,y) = f_\mathrm{y}\left[\frac{\gamma T_r(\frac{y}{f_\mathrm{y}} - \beta)}{Z(x,y)} + \omega_x(1 + \frac{y^2}{f_\mathrm{y}^2}) - \omega_y\frac{xy}{f_\mathrm{x}f_\mathrm{y}} - \omega_z\frac{x}{f_\mathrm{x}}\right], \tag{8}$$

where $u(x, y)$ and $v(x, y)$ are the horizontal and vertical components of motion.

Let us now consider the effects of rotation, causing the FOE to shift with respect to the optical axis. Let (x', y') be an arbitrary point in the image representing the FOE. We define the depth of the surface, $Z(x, y)$, with respect to the FOE:

$$Z(x, y) = \frac{Z(x', y')}{1 - a\frac{(x-x')}{f_\mathrm{x}} - b\frac{(y-y')}{f_\mathrm{y}}}. \tag{9}$$

Substituting (9) into Equations 7 and 8, we obtain:

$$u(x,y) = \frac{\gamma T_r(x - f_{\text{x}}\alpha)}{Z(x',y')}\left[1 - \frac{a(x-x')}{f_{\text{x}}} - \frac{b(y-y')}{f_{\text{y}}}\right] + \omega_x\frac{xy}{f_{\text{y}}} - \omega_y(f_{\text{x}} + \frac{x^2}{f_{\text{x}}}) + \omega_z\frac{y}{f_{\text{x}}}, \quad (10)$$

$$v(x,y) = \frac{\gamma T_r(y - f_{\text{y}}\beta)}{Z(x',y')}\left[1 - \frac{a(x-x')}{f_{\text{x}}} - \frac{b(y-y')}{f_{\text{y}}}\right] + \omega_x(f_{\text{y}} + \frac{y^2}{f_{\text{y}}}) - \omega_y\frac{xy}{f_{\text{x}}} - \omega_z\frac{x}{f_{\text{x}}}. \quad (11)$$

Given the optical flow at the FOE is zero, substituting for $x = x'$ and $y = y'$ provides the following constraints on the optical flow at the FOE:

$$0 = \frac{\gamma T_r(x' - f_{\text{x}}\alpha)}{Z(x',y')} + \omega_x\frac{x'y'}{f_{\text{y}}} - \omega_y(f_{\text{x}} + \frac{x'^2}{f_{\text{x}}}) + \omega_z\frac{y'}{f_{\text{x}}}, \quad (12)$$

$$0 = \frac{\gamma T_r(y' - f_{\text{y}}\beta)}{Z(x',y')} + \omega_x(f_{\text{y}} + \frac{y'^2}{f_{\text{y}}}) - \omega_y\frac{x'y'}{f_{\text{x}}} - \omega_z\frac{x'}{f_{\text{x}}}. \quad (13)$$

Solving for ω_x and ω_y, we obtain:

$$\omega_x = \frac{f_{\text{y}}}{x'y'}\left[\frac{\gamma T_r}{Z(x',y')}(x' - f_{\text{x}}\alpha) + \omega_y(f_{\text{x}} + \frac{x'^2}{f_{\text{x}}}) + \omega_z\frac{y'}{f_{\text{y}}}\right], \quad (14)$$

$$\omega_y = \frac{1}{f_{\text{x}}f_{\text{y}}(1 + \frac{x'^2}{f_{\text{x}}^2} + \frac{y'^2}{f_{\text{y}}^2})}\left[\frac{T_r}{Z(x',' y')}\Big(x'y'\beta + f_{\text{y}}x' + f_{\text{x}}f_{\text{y}}\alpha + \frac{f_{\text{x}}\alpha y'^2}{f_{\text{y}}}\Big) - \omega_z\Big(y' + \frac{y'^3}{f_{\text{y}}^2} - \frac{x^2y'}{f_{\text{x}}f_{\text{y}}}\Big)\right]. \quad (15)$$

Taking the partial derivatives of Equations 10 and 11 in their respective directions, and again substituting for $x = x'$, $y = y'$, we obtain the partial derivatives at the FOE, defined as:

$$\frac{\partial u}{\partial x}\Big|_{\text{foe}} = \frac{\gamma T_r}{Z(x',y')}\left[1 - a\left(\frac{x'}{f_{\text{x}}} + \alpha\right)\right] + \omega_x\frac{y'}{f_{\text{y}}} - \omega_y\frac{2x'}{f_{\text{x}}}, \quad (16)$$

$$\frac{\partial v}{\partial y}\Big|_{\text{foe}} = \frac{\gamma T_r}{Z(x',y')}\left[1 - b\left(\frac{y'}{f_{\text{y}}} + \beta\right)\right] + \omega_x\frac{2y'}{f_{\text{y}}} - \omega_y\frac{x'}{f_{\text{x}}}. \quad (17)$$

Summing these, we obtain the flow field divergence at the FOE (D_{foe}):

$$D_{\text{foe}} = \frac{-\gamma T_r}{Z(x',y')}\left[a\left(\frac{x'}{f_{\text{x}}} + \alpha\right) + b\left(\frac{y'}{f_{\text{y}}} + \beta\right) - 2\right] + 3\left(\frac{\omega_x y'}{f_{\text{y}}} - \frac{\omega_y x'}{f_{\text{x}}}\right), \quad (18)$$

and from this we obtain an equation for the relative depth of the scene point projecting to the FOE:

$$\frac{Z(x',y')}{T_r} = \frac{\gamma}{D_{\text{foe}}}\left[a\left(\frac{x'}{f_{\text{x}}} + \alpha\right) + b\left(\frac{y'}{f_{\text{y}}} + \beta\right) - 2\right] - \frac{3Z(x',y')}{D_{\text{foe}}T_r(\theta)}\left(\frac{\omega_x y'}{f_{\text{y}}} - \frac{\omega_y x'}{f_{\text{x}}}\right). \quad (19)$$

Using Equations 14 and 15, we substitute for ω_x and ω_y in (19) and thus remove both rotations from (19) such that:

$$\frac{Z(x',y')}{T_r} = \frac{\gamma}{D_{\text{foe}}(\theta)}\left[1 + a\left(\frac{x'}{f_{\text{x}}} + \alpha\right) + b\left(\frac{y'}{f_{\text{y}}} + \beta\right) - \frac{3}{\gamma x'}\left(-f_{\text{x}}\alpha + \frac{(x'f_{\text{y}} + f_{\text{x}}f_{\text{y}}\alpha + x'y'\beta + \frac{y'^2f_{\text{x}}\alpha}{f_{\text{y}}})}{f_{\text{y}}(1 + \frac{x'^2}{f_{\text{x}}^2} + \frac{y'^2}{f_{\text{y}}^2})} + \frac{\omega_z y'T_r}{f_{\text{y}}Z(x',y')}\Big(f_{\text{y}} + \frac{y'^2}{f_{\text{y}}} - \frac{x'^2}{f_{\text{x}}} - 1\Big)\right)\right]. \quad (20)$$

Notably, the removal of ω_x and ω_y introduces a term involving camera roll (ω_z). If required, techniques for roll removal such as that of Hanada and Enjima [4] can also be applied without prior knowledge of the rotation.

If T_r is aligned with the FOE, then (20) gives a precise measure of τ. In the presence of rotations, this assumption is unlikely to hold. However, considering a docking scenario for a finite sized robot, the presence of small instantaneous rotations will also mean that the precise point of impact is unlikely to be known. Given that the FOE provides the only location in the flow field where rotation is accounted for, we can consider (20) to be a reasonable approximation of τ (referred to as τ_{foe}) under these conditions.

Time-to-contact for a ground-based mobile robot: Consider Equation 20 for the case of a mobile robot, moving on a ground plane towards a visible planar surface. Given a fixed, approximately forward facing camera, ω_z, will be negligible, and can therefore be set to zero. In addition, the camera orientation parameters with respect to the heading direction: α, β and γ, can also be set to known values ($\alpha = \beta = 0, \gamma = 1$). From these substitutions, (20) is reduced to:

$$\tau_{foe} = \frac{1}{D_{\text{foe}}}\left[1 + \frac{ax'}{f_{\text{x}}} + \frac{by'}{f_{\text{y}}} - \frac{3}{(\frac{x'^2}{f_{\text{x}}^2} + \frac{y'^2}{f_{y^2}} + 1)}\right]. \quad (21)$$

Note that the only potential unknowns in (21) are the surface orientation parameters: a and b. Given some directional control maintaining an angle of approach with the surface, θ should be known, and an upper and lower bound is likely to exist for the surface orientation parameters, and hence for τ_{foe}.

Constraints on rotation: The requirement for the FOE to lie within the image plane provides a natural constraint on the use of Equation 21. Perhaps most important is the limitation this imposes on the magnitude of rotation allowable. If too large, the FOE will no longer exist within the image plane.

Let us consider the maximum rotation to be that which causes the FOE to shift to the edge of the image plane. Let k be the 1-dimensional shift of the FOE from the image centre to the image edge. For a ground-based robot with forward facing camera, and rotation only about the Y axis (ω_y), we consider

only horizontal shifts of the FOE, and thus re-write (12) as:

$$\frac{T_r k}{Z(k)} = \omega_y (f_{\mathrm{x}} + \frac{k^2}{f_{\mathrm{x}}}). \tag{22}$$

From this, an upper and lower-bound on rotation is obtained:

$$-\frac{T_r k}{Z(k)(f_{\mathrm{x}} + \frac{k^2}{f_{\mathrm{x}}})} \leq \omega_y \leq \frac{T_r k}{Z(k)(f_{\mathrm{x}} + \frac{k^2}{f_{\mathrm{x}}})}. \tag{23}$$

It can be seen in (23) that as T_r increases, or $Z(k)$ decreases, the bound on rotation widens. Therefore, $\frac{T_r}{Z}$ must be kept sufficiently high to ensure the FOE lies within the image plane. Naturally, if ω_y is sufficiently large with respect to T_r, the FOE may not exist.

Constraints on angle of approach (surface orientation): As mentioned earlier, the surface alignment parameters, a and b in (21), may not be known. It is important to note, however, that the existence of the FOE within the projected surface does enforce some constraints on the range of possible angles of approach. At extreme approach angles, the FOE is unlikely to exist at all as the distance from the surface becomes infinite along the axis of motion. Therefore, ensuring the FOE always exists within the projected surface target area should maintain an approach angle that is within stability limits. This may also be used as a means of assessing the achievability of the task.

III. Closed-loop analysis

The goal of this section is to consider the closed-loop behaviour of a vehicle where the control input is generated by proportional feedback of the flow divergence. The analysis is undertaken for a full dynamic vehicle model to provide the most general results. Velocity controlled mobile robots can be dealt with in this framework by introducing a virtual dynamic state into the control law implementation, as is done for the experimental results that we detail in Section IV. The control is implemented as proportional feedback of the measured divergence and not τ to avoid possible ill-conditioning associated with inversion of a measured variable that may be close to zero. Consider the system

$$\dot{Z} = T_r, \qquad Z(0) = Z_0, \tag{24a}$$

$$m\dot{T}_r = F, \qquad T_r(0) = T_0 \tag{24b}$$

where $m > 0$ is the vehicle mass, F is the force input, Z is the distance to the wall (assumed to be positive), and T_r is the velocity of the robot orthogonal to the wall. For the purposes of the theoretical development, we assume a fronto-parallel approach angle. We discuss the more general case in remarks at the end of the section.

Let D_{ref} be a constant reference set-point for the flow divergence. Recalling (2), if the measured divergence

$$D(t) = \frac{2T_r(t)}{Z(t)} \equiv D_{\mathrm{ref}}$$

is exactly equal to the constant reference divergence at all times along the closed-loop trajectory, then substituting into (24a) one obtains

$$\dot{Z} = \frac{D_{\mathrm{ref}}}{2} Z,$$

and hence

$$Z(t) = \exp\left(\frac{D_{\mathrm{ref}}\, t}{2}\right) Z_0. \tag{25}$$

Choosing the reference divergence $D_{\mathrm{ref}} < 0$, corresponding to an expanding image as the robot approaches the wall, it is clear that $Z(t) \to 0$ exponentially. The velocity $T_r(t)$ is bounded and also converges to zero exponentially. This result was discussed in Srinivasan *et al.* [19] for honey bee landings and has been a key motivation for most of the TTC based docking and obstacle avoidance algorithms [13], [3], [1].

In practice, exact tracking of reference divergence is impossible and we propose a proportional feedback control

$$F = K\left(D_{\mathrm{ref}} - D(t)\right). \tag{26}$$

The feedback F will adjust the velocity T_r to force $D(t)$ to track the reference D_{ref}. For large gain $K \gg 0$ then the relative tracking error will be small and the closed-loop system trajectory should be close to (25). However, the actual closed-loop system dynamics are complicated by the non-linear dependence of $D(t)$ on the distance $Z(t)$ (2). Substituting (26) into (24), one obtains

$$\dot{Z} = T_r, \qquad Z(0) = Z_0 > 0, \tag{27a}$$

$$\dot{T}_r = \frac{K}{m}\left(D_{\mathrm{ref}} - \frac{2T_r}{Z}\right), \qquad T_r(0). \tag{27b}$$

Here we demand that $Z_0 > 0$ is positive in order that the underlying physical assumptions in the image model are valid. The authors have been unable to find an analytic solution to (27), however, the following theorem proves that its solutions have the desired qualitative behaviour. Note that the $1/Z$ singularity in (27b) complicates the analysis considerably at the limit point $Z \to 0$.

Theorem 3.1: Let $(Z(t), T_r(t))$ denote the solution of the closed-loop dynamics (27). Assume that $D_{\mathrm{ref}} < 0$. Then there exists a time $T > 0$, possibly infinite, such that $(Z(t), T_r(t))$ exist and are bounded, and $Z(t) > 0$, for all $t \in [0, T)$. Moreover, one has

$$\lim_{t \to T} Z(t) = 0, \qquad \lim_{t \to T} T_r(t) = 0.$$

Proof: The ODE (27) is smooth and non-singular on the domain $Z > 0$ and hence there exists a time $T > 0$, possibly infinite, such that the solution $(Z(t), T_r(t))$ is well defined on $t \in [0, T)$.

Firstly consider the case where $T_r(0) > 0$. In this case the vehicle is initially moving away from the wall. The control input (26) is negative for all $T_r(t) > 0$. Consequently, the velocity will decrease until $T_r = 0$ while $Z(t) > 0$ will increase during this period and (27) will not pass through a singularity. The negative driving reference D_{ref} will continue to drive the velocity negative and there must be a subsequent time $0 < t_0 < T$ such that $T_r(t_0) < 0$, $Z(t_r) > 0$. From this point on we will ignore this initial transient and assume without loss of generality that $Z_0 > 0$ and $T_r < 0$ in the remainder of the proof.

We continue by proving that if $T_r(0) < 0$ then $T_r(t) < 0$ for all $t \in [0, T)$. The proof is by contradiction. Assume the converse; that is, there exists a first time t_1 such that $T_r(t_1) =$

0 and $Z(t_1) > 0$ (the case $Z(t_1) = 0$ while $T_r(t_1) \leq 0$ is dealt with later). Since the solution of (27) is at least C^1 at t_1 ((27) is non-singular for $Z(t_1) > 0$), then for all $\epsilon > 0$ there exists $\delta > 0$ such that $-\epsilon < T_r(t) < 0$ for all $t \in [t_1 - \delta, t_1)$. Since $Z(t)$ is monotonic decreasing for $T_r < 0$, (27b) yields

$$\dot{T}_r \leq \frac{K}{m}\left(D_{\text{ref}} + \frac{2\epsilon}{Z(t_1)}\right)$$

Choosing $\epsilon < -Z(t_1)D_{\text{ref}}/2$ shows that $\dot{T}_r \leq 0$ on $[t_1 - \delta, t_1)$ and hence $T_r(t_1) < T_r(t_1 - \delta)$. This contradicts the assumption. It follows that $T_r(t) < 0$ on $[0, T)$ and $Z(t)$ is strictly monotonic decreasing on the whole interval $t \in [0, T)$.

A similar argument can be used to show that an asymptotic solution such that $\lim_{t\to T} Z(t) = Z(T) > 0$ and $\lim_{t\to T} T_r = 0$ is impossible. The proof is by contradiction. Assuming that such a solution exists, then for all $\epsilon > 0$ there exists a t_1, $0 < t_1 < T$, such that $-\epsilon < T_r(t) < 0$ for all $t \in [t_1, T)$. Once again (27b) yields

$$\dot{T}_r \leq \frac{K}{m}\left(D_{\text{ref}} + \frac{2\epsilon}{Z(T)}\right)$$

Choosing $\epsilon < -Z(T)D_{\text{ref}}/2$ shows that $T_r(T) < T_r(t_1)$ and contradicts the assumption that $\lim_{t\to T} T_r = 0$.

Next we prove that $T_r(t)$ cannot escape to infinity before $Z(t) \to 0$. The proof is by contradiction. Assume the converse; that is, there exists a first time t_1 such that $\lim_{t\to t_1} T_r = -\infty$ and $Z(t_1) > 0$. Thus, for all $B > 0$ there exists $\delta > 0$ such that $T_r(t) < -B$ for all $t \in [t_1 - \delta, t_1)$. Then (27b) yields

$$\dot{T}_r > \frac{K}{m}\left(D_{\text{ref}} + \frac{2B}{Z(t_1 - \delta)}\right)$$

Choosing $B > -Z(t_1 - \delta)D_{\text{ref}}/2$ ensures that $\dot{T}_r > 0$ on $[t_1 - \delta, t_1)$ and hence $T_r(t_1) > T_r(t_1 - \delta)$. This contradicts the assumption and it follows that $T_r(t)$ is well defined for all $Z(t) > 0$.

We have shown that $T_r(t)$ can only become unbounded at the point where $Z \to 0$, and if its limit at this point is well defined, then it follows that T_r is bounded on $[0, T)$. Moreover, since $T_r(t) < 0$ on the whole interval $[0, T)$ then the solution of the ODE is defined for all $Z > 0$. Since Z is strictly monotonically decreasing on $[0, T)$ and cannot have a positive limit, it follows that $\lim_{t\to T} Z(t) = 0$.

The final requirement of the proof is to show that $\lim_{t\to T} T_r(t) = 0$. To prove this we compute the first integral of (27). For $Z > 0$ and $T_r < 0$ one has

$$\frac{d}{dt}T_r = \frac{dT_r}{dZ}\frac{dZ}{dt} = \frac{dT_r}{dZ}T_r$$

thus (27b) may be rewritten as a differential equation in the variable Z

$$\frac{dT_r}{dZ} = \frac{K}{mT_r}\left(D_{\text{ref}} - \frac{2T_r}{Z}\right) \qquad T_r(Z_0) < 0. \tag{28}$$

Furthermore, since $Z(t)$ is monotonic decreasing on the time interval $[0, T)$ we know that $T_r(Z) < 0$ for all $0 < Z < Z_0$. We introduce a change of variables $s = -\log(Z)$ on $0 < Z < Z_0$ to get rid of the singularity in (28). Since $Z(t)$ is strictly monotonic decreasing on the interval $[0, T)$, $s := s(t) \to +\infty$ is strictly monotonic increasing as function of time, and we can think of s as a pseudo-time variable. Changing variables in (28) one obtains

$$\frac{dT_r}{ds} = \frac{2KT_r - Ke^{-s}D_{\text{ref}}}{mT_r}. \tag{29}$$

Consider the storage function $\mathcal{L} = m|T_r|^2/2$, then

$$\frac{d\mathcal{L}}{ds} = T_r\frac{dT_r}{ds} = 2KT_r - Ke^{-s}D_{\text{ref}} \leq -2K|T_r| + |u(s)|$$

where $u(s) = Ke^{-s}D_{\text{ref}}$ is viewed as an exogenous signal. From Theorem 1, Sontag and Wang [18] it follows that (29) is input-to-state-stable (ISS). Moreover, since $u(s)$ is a bounded, asymptotically stable input signal to an ISS system, then $\lim_{s\to+\infty} T_r(s) \to 0$ [17]. This proves that $\lim_{t\to T} T_r(t) = 0$ and completes the proof. ■

The complexity of this proof is associated with the difficult nature of analyzing singular differential equations. The singularity in (27b) that occurs at $Z = 0$ makes it impossible to apply standard stability arguments. The approach taken in the proof of Theorem 3.1 cannot preclude the possibility that convergence occurs in finite time, an outcome that appears unlikely given the exponential nature of (25), a limiting solution for the case when the gain $K \gg 0$. The authors believe that for any gain $K > 0$ the solution is defined on $[0, +\infty)$, however, we do not have a proof for this result.

Theorem 3.1 is proved for the case where the robot approaches the wall fronto-parallel. Let θ be the angle of approach with respect to the surface normal. In practice, Theorem 3.1 is valid also for arbitrary constant angle $\theta < \pi/2$. Define $T_z = \cos(\theta)T_r$ to be the component of the robot velocity orthogonal to the wall. Equation (27b) can be rewritten

$$\begin{aligned}\dot{T}_r &= \frac{K}{m}\left(D_{\text{ref}} - D(t)\right)\\ &= \frac{1}{\cos(\theta)}\frac{\cos(\theta)K}{m}\left(D_{\text{ref}} - \frac{2\cos(\theta)T_r}{Z}\right)\end{aligned}$$

Thus since θ is constant one has

$$\dot{T}_z = \frac{K'}{m}\left(D_{\text{ref}} - \frac{2T_z}{Z}\right)$$

with a new gain $K' = \cos(\theta)K$. The dynamics of this system are equivalent to those studied in Theorem 3.1.

The authors believe that for θ smoothly time-varying, with bounded derivative, and bounded away from $\pi/2$ the system will have the same qualitative behaviour shown in Theorem 3.1. The bulk of the above proof will hold in a straightforward manner, however, a full analysis is considered beyond the scope of the present paper.

The analysis in this section is undertaken in continuous-time although the real-world control signals will always be applied in discrete-time. For a sufficiently fast time-sampling then the discrete system should inherit the same stability as the continuous-time system. Characterising how fast the sampling must be to guarantee convergence is beyond the scope of the present work.

IV. Experimental Results

In this section, we present four sets of experiments demonstrating the performance of the proposed FOE-based τ estimator to the task of docking. We provide results from simulation, off-board image sequences, and from the technique's application to the closed-loop control of a mobile robot performing a docking manoeuvre. We first describe each experiment and discuss implementation issues relating to the application of the FOE-based τ strategy. We then present the results of these experiments.

A. Simulation experiment

To test the theory, a simulation modelling the motion of a ground-based mobile robot, with camera, towards a planar fronto-parallel surface, was conducted. A 2D motion model was used, allowing only forward velocity and a single rotation in the ground plane. As such, only the u component of flow across a single row of pixels was required to obtain τ estimates. From this, a set of sample flow fields were obtained.

For each consecutive sample, the distance to the surface was decremented by a constant amount. The robot was assumed to be initially aligned fronto-parallel with the surface before a constant translational velocity, and randomly selected instantaneous rotational velocity were applied to the scene with respect to the robot's location. The resulting motion vectors were then projected onto the robot's image plane, thereby generating the expected flow resulting from the robot's motion with respect to the scene. From this, the FOE (which shifts as a result of the rotation) was located, and τ computed using (21).

B. Off-board time-to-contact experiments

1) Indoor image sequence: A *looming wall sequence* was constructed to simulate the image expansion likely to be experienced when approaching a textured surface. Figure 3(a) shows sample frames from the sequence. In the construction of the image sequence, the camera was moved 3cm per frame towards a heavily textured, approximately fronto-parallel wall. Optical flow fields were estimated for each frame of the sequence, and from this, τ estimates obtained.

Flow divergence was estimated using four patches in the image, each centred on a distance of 12 pixels from the FOE, and each at 45 degrees from the horizontal and vertical axes that intersect at the FOE. Figure 3(a) shows this patch configuration. For comparison, τ was also estimated by placing the four patches about the image centre.

2) Outdoor image sequences: To test the technique's robustness under more natural conditions, two outdoor image sequences were constructed, depicting the motion of the camera towards different, more natural surfaces. Figures 3(b) and (c) show sample frames from both sequences: the *looming bush sequence*, and the *looming bricks sequence*. Both sequences depict the motion of a camera (attached to the front of a bicycle and walked) at an approximately constant velocity towards each respective surface. The camera's motion was subject to rotations induced by the uneven terrain (grass), and small adjustments of the bicycle's heading (including camera roll). The initial distance in both sequences was 9m. The average velocity of the camera depicted in the *looming bush sequence* is approximately 13cm per frame (6km$/$hr), and 20.5cm per frame (8.5km$/$hr) for the *looming bricks sequence*.

Flow divergence was estimated from optical flow vectors within a single 51×51 pixel patch centred on the estimated location of the FOE. Divergence was also measured at the image centre using the same patch size.

C. On-board docking experiment

To test the robustness of the FOE-based τ measure, the technique was integrated into a simple closed-loop docking behaviour for velocity control of a mobile robot. In the experiment, a robot with a single, fixed, forward facing camera approached a heavily textured, roughly fronto-parallel wall, attempting to decelerate and safely stop as close to the wall as possible without collision. Figure 4 shows the experimental workspace.

The robot used is velocity controlled, that is, the control signal is passed to a servo motor that controls the rolling speed of the drive wheels. Initial experimental tests showed that direct proportional feedback of the drive wheels lead to highly aggressive control action due to the noise in the divergence measure. By incorporating a virtual model of robot dynamics in the control design, the closed-loop behaviour of the vehicle was smooth and well conditioned. The discrete time realisation of the proposed control law is

$$v_t = \Delta v_{t-1} + \frac{\Delta K_p}{m}(D_{\text{ref}} - D_t), \tag{30}$$

where $v(t)$ is the velocity control input at time t, Δ is the discretisation time, m is a virtual vehicle mass, K_p is a proportional gain, D_t is the most recent flow divergence estimate, and D_{ref} is the reference set-point for flow divergence ($\frac{\Delta K_p}{m} = 0.0325$ and $D_{\text{ref}} = 0.022$ for these trials). Along with the discrete-time kinematics

$$z_t = \Delta v_{t-1}. \tag{31}$$

Flow divergence was estimated using two 40×40 pixel image patches, each placed at 45 degrees on either side of the vertical axis passing through the FOE, and each centred on a distance of 25 pixels from the FOE. The patches were placed only above the FOE to avoid measuring divergence on the imaged ground plane. Reasons for the variation of patch size and configuration used in the off-board experiment were based on empirical observations of performance on-board. Due to the noisier conditions on-board, larger patch sizes were used to obtain a more robust estimate of flow divergence during ego-motion. In general, a range of patch sizes and configurations were found to obtain strong results.

D. Optical flow and FOE estimation

Throughout the experiments, Lucas and Kanade's [8] gradient-based method was applied. This technique was chosen based on strong performances in a recent comparison of optical flow techniques for robot navigation tasks [9]. For the indoor *looming wall* sequence, a standard implementation of

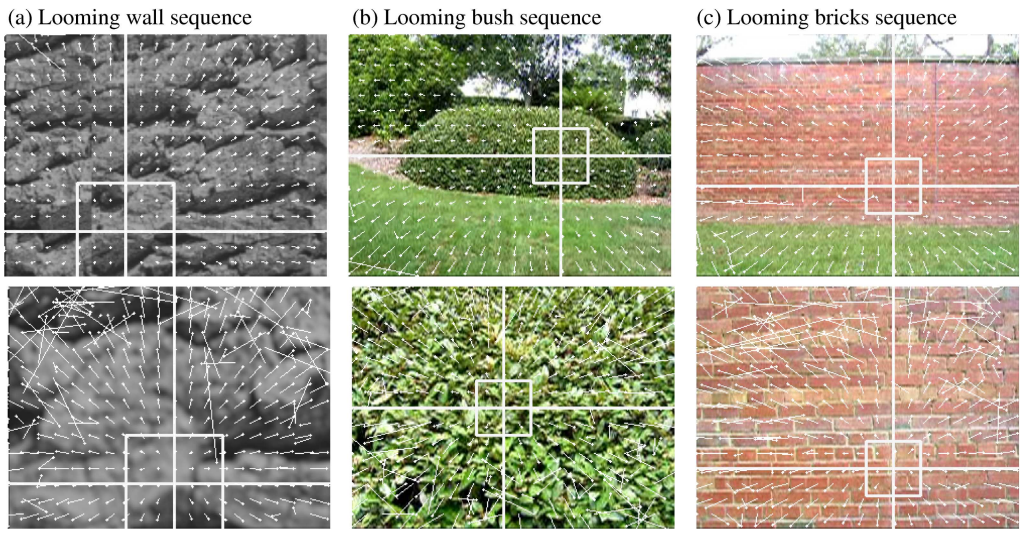

Fig. 3. Sample frames and flow fields from each image sequence used in off-board experiments: (a) *looming wall*, (b) *looming bush*, and (c) *looming bricks*. Line intersections show estimated FOE for frame, and boxes indicate the divergence patch configurations used for FOE-based τ estimation.

Lucas and Kanade's algorithm was applied, and flow vectors were obtained for all image points. Due to significantly larger flow experienced in both outdoor sequences, a pyramidal implementation of Lucas and Kanade's technique was applied. To offset the increased computation load of this approach, flow vectors were only estimated for every fifth pixel.

In all experiments, the FOE was calculated using a simple algorithm that requires the imaged surface to occupy the entire viewing field (or at least, the section of the viewing field for which the FOE is expected to lie within). To obtain x', each row in the image was used to count the number of positive and negative horizontal flow components, which were then differenced, and averaged over all rows to locate the overall zero point for x. The algorithm was applied similarly to obtain y', using the signs of vertical components of flow. While more sophisticated algorithms for locating the FOE do exist, it is important to note that in many cases, pure (or close to pure) translational motion is assumed (*e.g.*[16], [11], [5]). In contrast, the technique applied here provides a relatively high tolerance to rotation, such that the FOE will always be located so long as it lies within the imaged area, and other local minima in the flow field do not exist. Given only the sign of flow vectors are used to estimate the FOE, the computation associated with its estimation is negligible in comparison with the flow estimation itself. It should be noted that other suitable techniques do exist, such as [7], that do not require the segmentation of the object surface area. The algorithm employed here was chosen primarily for its efficiency in achieving reasonably accurate FOE estimates.

E. Results

1) Simulation results: Figure 5 gives the simulation results, showing a direct comparison of τ obtained using the FOE-based estimator defined by (21), and estimates obtained from the measured divergence at the image centre (using (2)). Ground truth τ is also provided, computed from the robot's

Fig. 4. Setup for on-board docking tests.

distance from the surface and its known constant forward velocity towards the surface. It can be seen that the FOE-based τ measure closely reflects ground truth. Small discrepancies between the FOE-based measure and ground truth are the result of unavoidable quantisation errors in the image, disallowing the precise location of the FOE.

In contrast, τ estimates taken along the optical axis exhibit significant fluctuation compared with that obtained at the FOE. It is also evident that the image centre always provides an over estimate of τ, a result of the optical axis deviating from its fronto-parallel alignment with the surface. While errors in τ are reduced as the distance to the surface approaches zero, it is important to note that this is due to the robot's constant velocity towards the surface. As the surface draws near, the translational flow increases, thereby diminishing the effects of the robot's rotation in the flow field.

2) Indoor image sequence results: Figure 6(a) shows τ estimates for each frame of the indoor looming-wall sequence for the FOE-based, and image-centre-based strategies. Ground truth τ is also shown, obtained from the camera's known velocity, and a best linear fit over τ measures obtained from ground truth flow fields constructed from camera calibration.

From these results, a significant improvement in the consistency of τ estimates is achieved when divergence is calculated with respect to the FOE. Of particular note, the FOE-based

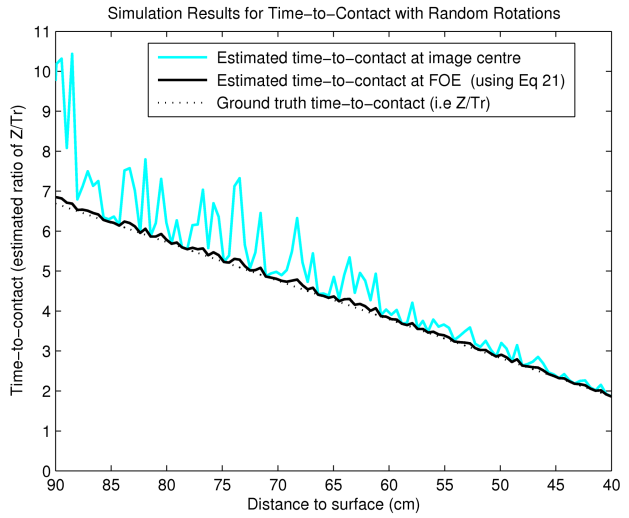

Fig. 5. Simulation results compare our FOE-based τ estimator (Equation 21), with τ estimates obtained at the image centre using (2) for the simulated 2D motion of a ground-based mobile robot translating at constant speed towards a fronto-parallel, planer surface. For each sample, the robot's forward speed, and randomly chosen instantaneous rotational velocity ($-0.1 \leq \omega_y \leq 0.1$) were used to compute the corresponding horizontal flow. From this, τ estimates were obtained. Ground truth shows the exact τ for each sample, given the robot's forward velocity and distance from the surface. For all samples, the camera's focal length is set to 188px.

strategy achieves a close match with ground truth from the fifteenth frame onward. In contrast, the image centre-based method consistently over-estimates τ, and exhibits larger fluctuations across the sequence.

3) Outdoor image sequence results: Figures 6(b) and (c) show τ estimates for both outdoor image sequences, again comparing the FOE-based, and image-centre-based strategies.

As with the indoor *looming wall sequence*, improvements in τ estimation are achieved by the FOE-based strategy as the surface approaches. This is evident from frame 40 onward for the *looming bush* sequence, and from frame 20 in the *looming bricks* sequence.

Across all sequences, larger fluctuations are evident in early frames for both strategies. This is unsurprising given the flow due to camera translation is unlikely to be large enough to be reliably measured at this distance from the wall. It is also likely that the FOE is poorly defined at this distance. In early frames of both outdoor sequences, the FOE's location was observed to shift significantly, and in some cases (particularly for the *looming bush* sequence), fall outside the imaged area of the surface. As divergence increases, however, the FOE-based strategy quickly stabilises, and begins to outperform the image-centre-based estimator.

In addition to rotational effects, the FOE-based strategy was observed to provide increased robustness to flow exceeding measurable levels in each sequence. This effect is evident in flow fields shown in the bottom row of Figure 3, where peripheral flow vectors become noisy and unreliable. While generally only in the periphery, this region of flow becomes larger as $\frac{Tr}{Z}$ increases (*i.e.* $Z \rightarrow 0$). As a result, any shifting of the FOE when in close proximity to the surface may

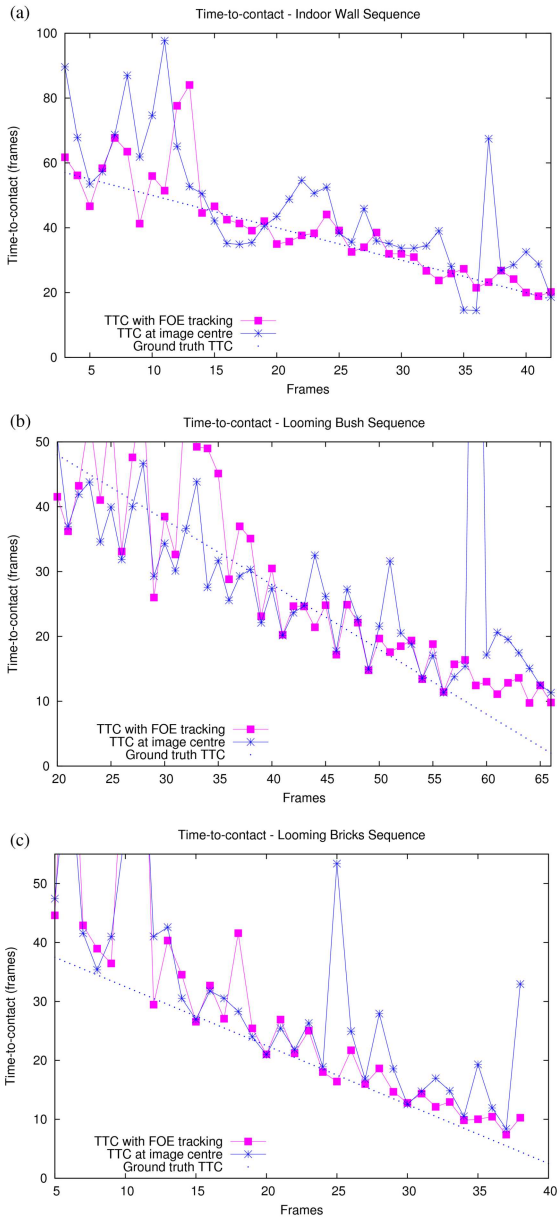

Fig. 6. τ estimates for: (a) indoor looming wall sequence, (b) looming bush sequence, and (c), looming bricks sequence.

cause this region to inhabit image-centre-based divergence patches. This is the likely cause of larger fluctuations in image-centre-based τ estimates in the later frames of each sequence (particularly for the *looming bricks* sequence, where forward velocity was significantly faster). In contrast, τ estimates taken with respect to the FOE remain stable under these conditions, and in accordance with simulation results, appear to improve in consistency as Z decreases. This improvement appears also to result from the FOE itself being more clearly defined, and therefore more accurately located when $\frac{T_r}{Z}$ is large.

4) On-board docking results: Six trials of the FOE-based docking strategy were conducted, and data recorded. Figure 7 shows the velocity-distance profiles and the plotted approach of the robot towards the surface for each trial. Also shown is the theoretically expected velocity-distance profile based on

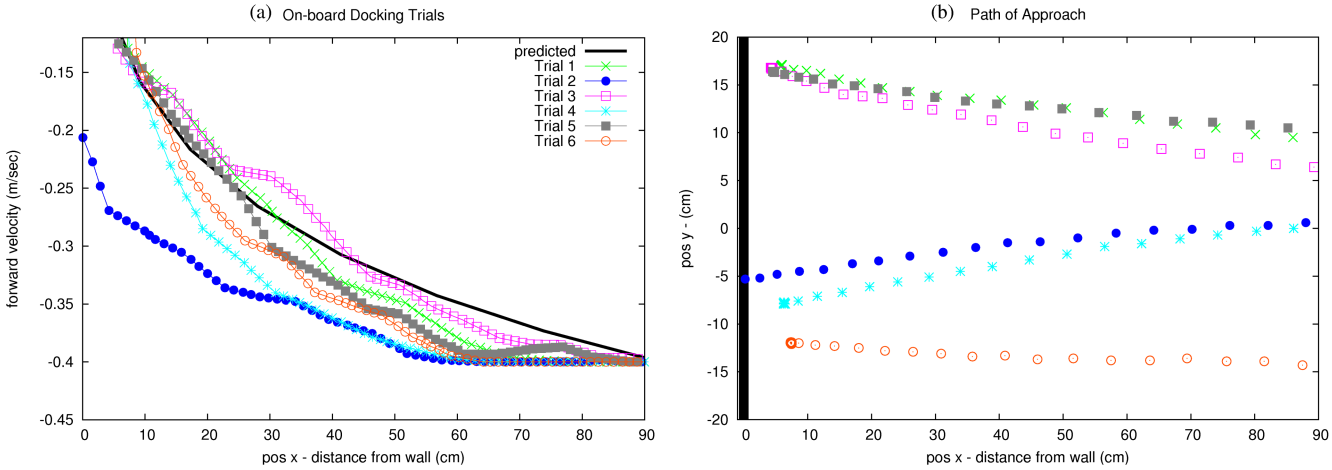

Fig. 7. On-board docking results showing (a) velocity-distance profiles, and (b), the plotted paths of the robot for each trial.

the integrating of (27) in discrete time for the initial velocity, distance and tuning parameter values used in the trials. Of the six trials conducted, the FOE-based strategy docked in close proximity to the surface five times without collision. Only one collision, Trial 2, was observed. Results shown in Figure 7 suggest this was most likely due to noise effected divergence estimates obtained around 30cm from the surface.

Among the successful trials, close proximity stopping distances were achieved with surprisingly high consistency. Recorded velocity-distance profiles, and stopping distances also appear consistent with theoretical expectation. Notably, results show an early lack of response compared with the predicted deceleration. This is a likely result of divergence being too small to measure at such distances. As the robot approaches, the measured divergence increases, and the velocity-distance profiles begin to resemble theoretical expectations. The average stopping distance achieved over the successful trials was 6cm, with the furthest distance recorded being just 7cm. This consistency in stopping distance is encouraging when considering the simple control law used, and significant differences in the plotted approach path of the robot during each trial. Figure 7(b) shows considerable variation in both the robot's initial starting position, and the extent (and direction) of the lateral drift experienced during each approach.

An attempt was made to compare the FOE-based on-board control scheme with the same control scheme using an image-centre based divergence measure. The raw divergence estimates obtained at the image-centre, however, were found to be unworkable for the simple proportional control scheme used. A large range of tuning parameter values were explored.

The FOE-based docking strategy compares well with previous work in flow-based docking. The final stopping distances achieved are a significant improvement on Questa *et al.* [13] (approximately 15cm), and comparable with Santos-Victor and Sandini [15]. Unlike previous work, we report highly consistent results over a set of trials. In addition, we obtain these results using general optical flow estimation (no affine approximations), and without filtering of the divergence estimates. We acknowledge, however, that we are using newer and faster computers than in previous work, thus allowing faster estimation of the optical flow.

V. Conclusion

This paper has presented a mobile robot docking strategy that utilises a time-to-contact (τ) estimation that is robust to noisy, instantaneous rotations induced by robot ego-motion. We have shown that through tracking the focus of expansion in the optical flow field, small rotations of the camera and misalignments of the optical and translational axes can be accounted for by calculating flow divergence with respect to the FOE. In this way, the effects of the rotation are effectively cancelled out, and improved accuracy and stability is achieved. Based on this, we have proposed a divergence-based control law for docking a robot with near fronto-parallel surfaces with closed-loop analysis proving its stability under ideal conditions, verified also through experimental trials. These results show a significant improvement in τ estimates when compared with common strategies that take no account of the shifting FOE during robot ego-motion. The accuracy and stability achieved using the FOE-based τ estimator was demonstrated to be sufficient for fine motion control of a mobile robot when in close proximity with the docking surface.

Acknowledgement

The authors thank the Dept. of Computer Science and Software Engineering, The University of Melbourne, for use of the robot and tracking software used in docking experiments.

References

[1] N. Ancona and T. Poggio, "Optical flow from 1d correlation: Application to a simple time-to-crash detector," in *Proceedings of the International Conference on Computer Vision*, 1993, pp. 209–14.

[2] R. Cipolla and A. Blake, "Image divergence and deformation from closed curves," *Int. J. Rob. Res.*, vol. 16, no. 1, pp. 77–96, 1997.

[3] D. Coombs, M. Herman, T. Hong, and M. Nashman, "Real-time obstacle avoidance using central flow divergence,and peripheral flow," *IEEE Transactions on Robotics and Automation*, vol. 14, no. 1, pp. 49–59, 1998.

[4] M. Hanada and Y. Ejima, "A model of human heading judgement in forward motion," *Vision Research*, vol. 40, no. 2, pp. 243–63, Jan. 2000.

[5] R. Jain, "Direct computation of the focus of expansion," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 5, no. 1, pp. 58–64, Jan. 1983.
[6] D. N. Lee, "A theory of visual control of braking based on information about time to collision," *Perception*, vol. 5, no. 4, pp. 437–59, 1976.
[7] H. Li, "Global interpretation of optical flow field: A least-squares approach," in *International Conference on Pattern Recognition*, 1992, pp. 668–71.
[8] B. Lucas and T. Kanade, "An iterative image registration technique with an application to stereo vision." in *Proceedings of DARPA Image Understanding Workshop*, 1984, pp. 121–30.
[9] C. McCarthy and N. Barnes, "Performance of optical flow techniques for indoor navigation with a mobile robot," in *Proceedings of the 2004 IEEE International Conference on Robotics and Automation*, 2004, pp. 5093–8.
[10] ——, "A robust docking strategy for a mobile robot using flow field divergence," in *2006 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2006, pp. 5564–5569.
[11] S. Negahdaripour., "Direct computation of the foe with confidence measures," *Computer Vision and Image Understanding*, vol. 64, no. 3, pp. 323–350, November 1996.
[12] R. C. Nelson and J. Y. Alloimonos, "Obstacle avoidance using flow field divergence," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 11, no. 10, pp. 1102–6, 1989.
[13] P. Questa, E. Grossmann, and G. Sandini, "Camera self orientation and docking maneuver using normal flow," in *Proceedings of Spie - the International Society for Optical Engineering*, vol. 2488, Orlando, USA, 1995, pp. 274–83.
[14] J. Santos-Victor and G. Sandini, "Visual behaviors for docking, Tech. Rep. LIRA-TR 2/94, June 1994.
[15] ——, "Visual behaviors for docking," *Computer Vision and Image Understanding: CVIU*, vol. 67, no. 3, pp. 223–38, 1997.
[16] D. Sazbon, H. Rotstein, and E. Rivlin, "Finding the focus of expansion and estimating range using optical flow images and a matched filter," *Machine Vision and Applications*, vol. 15, pp. 229–36, 2004.
[17] E. Sontag, "Smooth stabilization implies coprime factorization," *IEEE Transactions on Automatic Control*, vol. 34, no. 4, pp. 435–443, 1989.
[18] E. D. Sontag and Y. Wang, "On characterization of the input-to-state stability property," *Systems and Control Letters*, vol. 24, no. 5, pp. 351–359, 1995.
[19] M. V. Srinivasan, S. W. Zhang, J. S. Chahl, E. Barth, and S. Venkatesh, "How honeybees make grazing landings on flat surfaces," *Biological Cybernetics*, vol. 83, pp. 171–83, 2000.
[20] M. Subbarao, "Bounds on time-to-collision and rotational component from first-order derivatives of image flow," in *Computer Vision, Graphics and Image Processing*, vol. 50, 1990, pp. 329 – 41.
[21] M. B. van Leeuwen and F. C. A. Groen, "Motion estimation wth a mobile camera for traffic applications," in *Proceedings of the IEEE Intelligent Vehicles Symposium 2000*, 2000, pp. 58–63.
[22] ——, "Motion interpretation for in-car vision systems," in *Proceedings of the 2002 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2002, pp. 135–40.

PLACE PHOTO HERE

Chris McCarthy received the B.Sc degree with honours in 1999, and a Master of Computer Science degree in 2005 from the University of Melbourne. In 2000 he worked as an International Fellow in the Centre for Computer Studies, at Ngee Ann Polytechnic, Singapore. From 2001-2005, he was an Associate Lecturer with the Department of Computer Science and Software Engineering at The University of Melbourne. Since 2005, he has been working towards a PhD degree with the Australian National University and NICTA's Canberra Research Laboratory. His research interests are in vision for vehicle guidance and computational models of biological vision.

Mr. McCarthy received an International Foundation of Robotics Research Student Fellowship Award for his paper presented at the 2004 International Symposium on Experimental Robotics (ISER).

PLACE PHOTO HERE

Nick Barnes received the B.Sc. degree with honours in 1992, and a Ph.D. in computer vision for robot guidance in 1999 from the University of Melbourne. From 1992-1994 he worked for an IT consulting firm. In 1999 he was a visiting research fellow at the LIRA-Lab at the University of Genoa, Italy, supported by an Achiever Award from the Queens' Trust for Young Australians. From 2000 to 2003, he was a lecturer with the Department of Computer Science and Software Engineering, The University of Melbourne. Since 2003 he has been with NICTA's Canberra Research Laboratory. He is currently a principle researcher and research group manager in computer vision. His research interests include visual dynamic scene analysis, computational models of biological vision, feature detection, vision for vehicle guidance and medical image analysis.

PLACE PHOTO HERE

Robert Mahony is currently a reader in the Department of Engineering at the Australian National University. He received a PhD in 1995 (systems engineering) and a BSc in 1989 (applied mathematics and geology) both from the Australian National University. He worked as a marine seismic geophysicist and an industrial research scientist before completing a postdoctoral fellowship in France and a Logan Fellowship at Monash University in Australia. He has held his post at ANU since 2001. His research interests are in non-linear control theory with applications in robotics, geometric optimisation techniques and learning theory.