Supersonic Corner Flows in Rectangular Channels

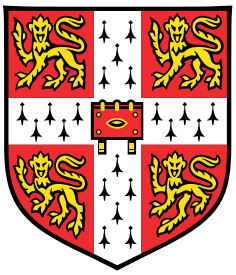

Kshitij Sabnis

Department of Engineering
University of Cambridge

This dissertation is submitted for the degree of
Doctor of Philosophy

Trinity College April 2020

Declaration

I hereby declare that except where specific reference is made to the work of others, the contents of this dissertation are original and have not been submitted in whole or in part for consideration for any other degree or qualification in this, or any other university. This dissertation is my own work and contains nothing which is the outcome of work done in collaboration with others, except as specified in the text and Acknowledgements. This dissertation contains fewer than $65,000$ words including appendices, bibliography, footnotes, tables and equations and has fewer than 150 figures.

Kshitij Sabnis
April 2020

Acknowledgements

First and foremost, I would like to thank my supervisor, Professor Holger Babinsky, for the opportunity to conduct exciting research and for his invaluable guidance and support throughout.

I would like to take this opportunity to thank our colleagues at the US Air Force Research Laboratory: Dr John Benek, Daniel Galbraith, and Dr Charles Tyler. The data from their simulations of the Cambridge wind tunnel has been used extensively in this thesis. The high-quality computational work and our close collaboration made a real difference to this project. I am also deeply grateful for the mentoring and guidance provided by Dr Philippe Spalart at Boeing Commercial Airplanes as I negotiated the world of turbulence modelling.

Of course, I must thank Dave Martin, Sam Flint, Tony Luckett, and Ciaran Costello, for maintaining, installing, and operating the supersonic wind tunnel. I would also like to thank everyone else in the research group for patiently teaching me and for sharing with me the highs and lows of wind tunnel testing: Dr Shunsuke Koike, Dr Andrea Coschignano, Dr Xue Xiang, Dr Charlotte Coles, Connor O'Pray, and Rhys Williams.

There are many others who have made my years in Cambridge so enjoyable: modern pentathletes, colleagues from the office, and many more. I will not attempt a list of names for fear of missing someone out, but I am eternally grateful to all of you for so many wonderful memories.

I would also like to thank my family, without whom I would not be here today. Finally, of course, many thanks to my wonderful girlfriend, Charlotte, for her unending love and for making my PhD years truly special.

This material is based upon work supported by the US Air Force Office of Scientific Research under award FA9550–16–1–0430.

Abstract

Rectangular channel geometries are widely encountered in supersonic flows, such as in wind tunnels and in aircraft inlets. Shock–boundary-layer interactions in these flows are known to exhibit significant three-dimensionality, due to the presence of sidewalls and associated corner boundary layers. The main effect is on the local separation of these corner regions, which then affects the wider flow field. Successful prediction of the overall flow therefore relies on the corner separation to be determined accurately. This, in turn, requires knowledge of the flow momentum distribution within the corner boundary layers. However, numerical methods struggle to reliably compute these flows and there is not much experimental data on supersonic corner boundary layers for comparison. This thesis addresses the outstanding gap in knowledge by performing validation-quality experiments on the corner regions of a Mach 2.5 channel flow, with a unit Reynolds number of approximately 40×10^6 m^{-1}. The experiments are conducted in the rectangular test section of a supersonic wind tunnel at the University of Cambridge.

An analysis of the wind tunnel experiments, alongside computational data provided by the US Air Force Research Laboratory, reveals that the corner boundary layers are strongly influenced by the geometry of the two-dimensional nozzles used to produce the supersonic flow. The dominant effect is related to bulk vertical velocities within the sidewall boundary layers, induced by vertical pressure gradients in the nozzle. For some very particular geometries, a second influence may be associated with a region of separated flow immediately ahead of the nozzle, which generates vortices within the sidewall boundary layer. Through these mechanisms, the nozzle geometry is seen to strongly influence both the thickness and the structure of the corner boundary layers.

High-quality experimental data in the corner regions are used to validate relevant numerical methods. Simple linear eddy-viscosity type turbulence models are found to compute these flows particularly poorly, with a 7% discrepancy in streamwise velocity. This is largely due to the fact that they do not capture known, stress-induced, corner vortices. However, the quadratic constitutive relation improves prediction of the corner boundary-layer structure, reducing experimental–computational differences by as much as half. This improvement is associated with vorticity generation in these corner regions, albeit with slightly different properties to the physical vortices. This production of vorticity depends only on the presence of a quadratic term in the eddy-viscosity model and not on which particular quadratic term is used. A more general form of the quadratic constitutive relation with one additional term is proposed, which appears to exhibit substantial improvements in the prediction of turbulent stress anisotropies.

The nozzle geometry effects are exploited to produce two otherwise-identical experimental setups with distinctly different momentum distributions in the corner boundary layers. A full-span wedge introduces an oblique shock with flow deflection angle, 8°, which impinges on the floor boundary layer. The two setups exhibit quite dissimilar separation behaviour, not only in the corner regions but also on the tunnel's centre span, with a difference in central separation length of as much as 35%. The observed behaviour is consistent with expectations based on local flow momentum affecting corner separation size, and on the displacement effect of this corner separation influencing the wider flow.

Table of contents

Nomenclature

Roman Symbols

A, B	constants in Stern–Volmer relation
A_V	cross-sectional area of LDV probe volume
B	intercept of log-law ($= 5$)
C^+	dimensions of corner separation associated with the positive-z sidewall
C^-	dimensions of corner separation associated with the negative-z sidewall
$\tilde{C}$	spatial correlation coefficient
$c_{cr1}, c_{cr2}, c_{cr3}, c_{cr4}$	constants in the quadratic constitutive relation
$\tilde{c}_{cr3}$	effective constant in the quadratic constitutive relation ($= c_{cr3} - c_{cr4}$)
c_f	skin friction coefficient ($= \tau_w / q_\infty$)
d	diameter of static pressure tap
d_p	mean diameter of seeding particles used for LDV
f_D	frequency of Doppler burst
f_s	frequency shift introduced to laser beam by Bragg cell
H	boundary-layer shape factor ($= \delta^* / \theta$)
h	half-height of channel
I	emitted light intensity
K	Stern–Volmer constant
k	turbulent kinetic energy
l	length of static pressure tap

L	dimensions of central separation
M	Mach number
N	number of measurements in data bin
n_p	number density of seeding particles
P	production of turbulent kinetic energy
p	pressure
q	dynamic pressure ($= \rho u^2/2$)
r	radius of pipe
R^2	coefficient of determination
Re	Reynolds number
s	constant in Musker model ($= 1.09 \times 10^{-3}$)
S_{ij}	symmetric stress tensor ($= \frac{1}{2}\left(\partial u_i/\partial x_j + \partial u_j/\partial x_i\right)$)
S_k	Stokes number
S^*_{ij}	traceless symmetric stress tensor ($= S_{ij} - 1/3\ \partial u_k/\partial x_k\ \delta_{ij}$)
T	temperature
t	time
$\tilde{t}$	characteristic timescale
$\vec{u} = (u, v, w)$	streamwise, floor-normal, and spanwise component of velocity
u^+	velocity (u) in inner units ($= u/u_\tau$)
u^*	velocity (u) transformed by the van Driest compressibility correction
u_τ	friction velocity ($= \sqrt{\tau_w/\rho_w}$)
W	width of wind tunnel
$\vec{x} = (x, y, z)$	streamwise, floor-normal, and spanwise coordinates
$\tilde{x}$	streamwise coordinate relative to inviscid shock location
y^+	wall-normal distance (y) in inner units ($= y u_\tau/\nu_w$)

Greek Symbols

α	angle between laser beams
β	pitch angle of laser emitting head from horizontal
Δ	error in measured quantity
δ	boundary-layer thickness
δ_{99}	boundary-layer thickness (99% of freestream velocity)
δ_f	distance between fringes in laser interference pattern
δ_{ref}	flow thickness
δ^*	boundary-layer displacement thickness
δ_{ij}	Kronecker delta tensor
ε	dissipation of turbulent kinetic energy
γ	ratio of specific heat capacities ($= 1.4$)
κ	Karman constant ($= 0.41$)
λ_1, λ_2	laser light wavelength for components 1, 2 of LDV
μ	dynamic viscosity
μ_t	eddy viscosity
ν	kinematic viscosity ($= \mu/\rho$)
ν_1, ν_2	velocity output from components 1, 2 of LDV
Ω_{ij}	strain tensor ($= \frac{1}{2}(\partial u_i/\partial x_j - \partial u_j/\partial x_i)$)
$\vec{\omega} = (\omega_x, \omega_y, \omega_z)$	streamwise, floor-normal, and spanwise component of vorticity ($= \nabla \times \vec{u}$)
ω_{RR}	rigid-body component of vorticity
ψ	rotation of laser emitting head from horizontal
ρ	density
$\tilde{\sigma}_{ij}$	normalised estimated Reynolds stress tensor ($= \tilde{\tau}_{ij} \,/\, \mu_t$)
τ	shear stress

τ_{ij}	Reynolds stress tensor ($= -\rho\overline{u_i' u_j'}$)
$\tilde{\tau}_{ij}$	estimated Reynolds stress tensor
θ	boundary-layer momentum thickness

Superscripts

$\bar{\cdot}$	ensemble-averaged quantity
$\cdot'$	fluctuations
$\tilde{\cdot}$	estimated value of quantity

Subscripts

$\cdot_0$	stagnation quantity
$\cdot_{01}$	pre-shock stagnation quantity
$\cdot_{02}$	post-shock stagnation quantity
$\cdot_b$	bulk quantity
$\cdot_{\text{comp}}$	quantity extracted from computation
$\cdot_e$	property evaluated at boundary-layer edge
$\cdot_{\text{exp}}$	quantity measured from experiment
$\cdot_f$	property of fluid
$\cdot_i$	incompressible definition of quantity
$\cdot_p$	property of seeding particles used for LDV
$\cdot_{\text{pit}}$	quantity measured using Pitot probe
$\cdot_{\text{max}}$	maximum possible value of quantity
$\cdot_{\text{ref}}$	quantity measured in reference condition
$\cdot_s$	property measured in settling chamber
$\cdot_\tau$	based on the friction velocity (u_τ)
$\cdot_w$	quantity evaluated at the wall
$\cdot_x$	dimension measured in streamwise direction

$\cdot_z$	dimension measured in spanwise direction
$\cdot_\infty$	freestream quantity

Acronyms / Abbreviations

CAD	computer-aided design
CFD	computational fluid dynamics
DNS	direct numerical simulation
DSLR	digital single-lens reflex
DTI	dial test indicator
HLLC	Harten–Lax–van Leer contact
LDV	laser Doppler velocimetry
LED	light-emitting diode
LES	large-eddy simulation
LEVM	linear eddy-viscosity model
NASA	National Aeronautics and Space Administration
PSP	pressure-sensitive paint
QCR	quadratic constitutive relation
RANS	Reynolds-averaged Navier–Stokes
SBLI	shock–boundary-layer interaction
SSOR	symmetric successive over-relaxation
SST	shear stress transport

Chapter 1

Research Motivations

The rising demand for fast air travel has stimulated a renewed interest in the development of civilian supersonic aircraft, which may enter service as early as the mid-2020s [1]. Several organisations are actively designing vehicles for supersonic air transport, ranging from business jets (Aerion Supersonic; Gulfstream Aerospace; Spike Aerospace) to commercial airliners (Boom Technology; Lockheed Martin Corporation) [2]. The difficulties associated with noise and emissions are relatively well known. Since fuel accounts for $20-35\%$ of typical airline operating costs, another key engineering challenge faced by designers is the efficiency of the inlet–engine system, shown in figure 1.1a [3]. The inlet design is particularly crucial at supersonic speeds, where the efficiency of this component affects the payload by 2.5 times as much as in subsonic aircraft [4].

The aircraft inlet slows down the incoming supersonic flow to subsonic speeds suitable for the engine, whilst minimising mechanical energy loss. Both military and civilian aircraft typically have inlets with a rectangular channel geometry. For internal-compression and mixed-compression inlets, the supersonic–subsonic transition takes place inside the inlet itself. The flow in the initial part of the channel is therefore supersonic. This flow is decelerated by a system of oblique shocks, shown in figure 1.1b. The shocks impinge on the boundary layers growing along the internal inlet surfaces. The resulting adverse pressure gradient causes the boundary layers to thicken or, perhaps, to separate.

The boundary layers along the streamwise corners of the channel are particularly problematic. Here, the sidewall boundary layer intersects with an equivalent boundary layer on the floor or ceiling, as shown in figure 1.1c. The corner boundary layer therefore contains extremely low-momentum flow, and is most susceptible to separation. As a result, the corner regions cause an increase in distortion and a significant loss of mechanical energy. This behaviour can, in turn, have a substantial impact on the overall aircraft. For example, the mechanical energy available to the engine is reduced by $8-10\%$, which decreases the range of supersonic fighter aircraft by as much as 9% [5].

Moreover, the corner boundary layers can also extend an influence on the wider flow field. The separation of the corner region results in the generation of pressure waves. These propagate into the flow, and so influence the pressure profile elsewhere. Through this mechanism, the corner flow can affect how the boundary layers, even outside the local region, respond to the incident shock [6].

a)

Photograph of the Concorde airliner, with inlet-engine system indicated, has been removed for copyright reasons. Copyright holder is David Parker/BWP Media/Getty Images.

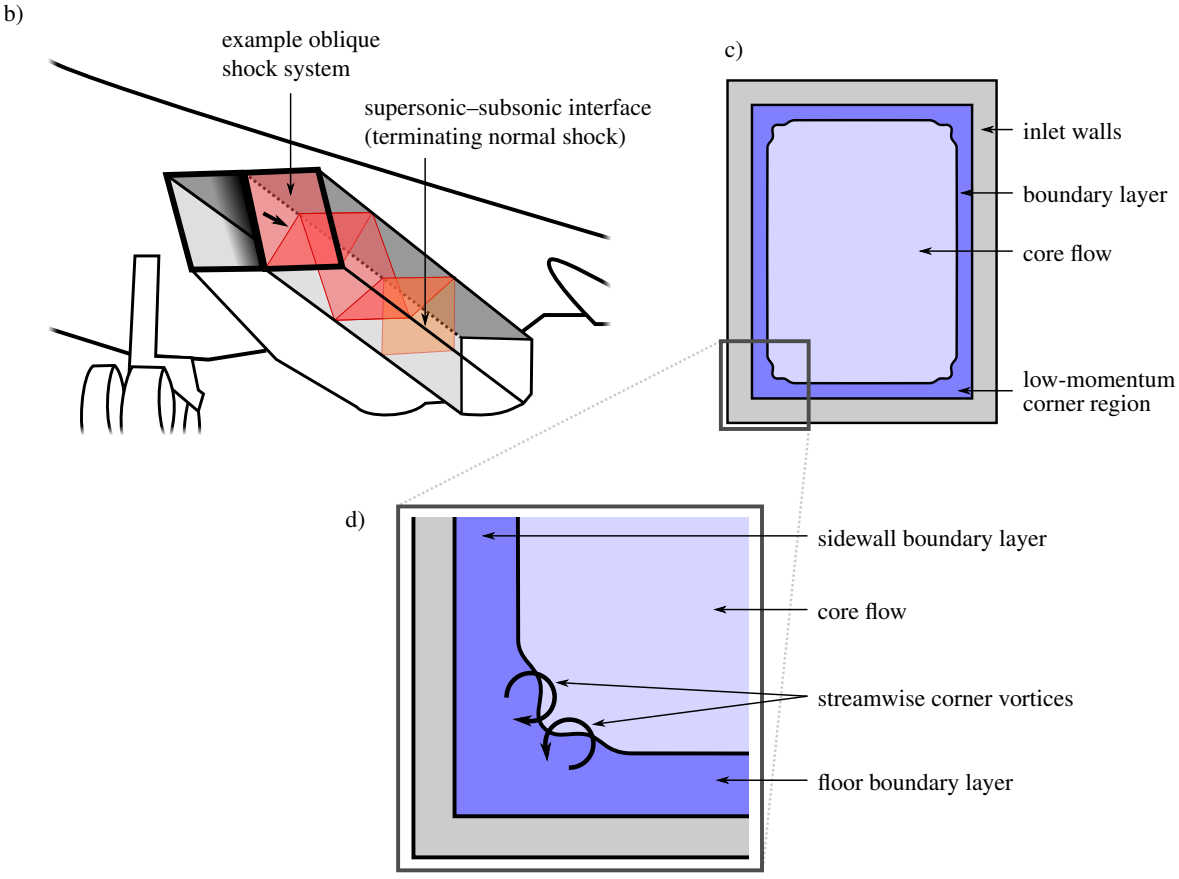

Figure 1.1: a) Photograph of the supersonic Concorde airliner. Image by David Parker / BWP Media / Getty Images. b) Schematic of inlet, with an example inviscid shock system within the rectangular channel. c) Cross-section of the inlet, showing boundary layers along the internal surfaces. d) Detailed view of the corner boundary layer, and the effect of corner vortices on momentum in this region.

The significant problems, both direct and indirect, associated with the corner boundary layers calls for the application of control techniques. However, despite this being an area of active research, there are still no consistently successful methods to mitigate corner separation, apart from (possibly) targeted surface bleed [7, 8]. The difficulties in alleviating these issues can arguably be attributed directly to a poor understanding of the detailed structure of the corner boundary layer.

Developing better insight into these corner flows is not, however, relevant just to aircraft inlets. Air ejectors and inlets for both supersonic missiles and scramjets are often designed with rectangular cross-sections [9–11]. The canonical supersonic rectangular-channel flow therefore applies to these systems too, as does the inherent importance of the boundary layers in the corner regions of the duct.

More immediately, supersonic wind tunnels also typically have a rectangular cross-section and therefore contain corner boundary layers. Experimental investigations on shock wave–boundary-layer interactions typically focus on the response of the floor boundary layer at the tunnel centre span, assuming the flow here to be two-dimensional. However, the significant influence that the corner regions exert on the overall flow field has been shown to impact separation even on the centre line [6]. In order to better understand how this class of experiments deviates from two-dimensionality, particularly when such data are quantitatively compared with computations, it is therefore necessary to better understand the flow in corner regions.

The main gap in knowledge is related to the streamwise vortices, generated by anisotropies in the local turbulent stresses, which are known to exist within the corner boundary layer [12]. These vortices, shown in figure 1.1d, transfer momentum between the core flow and the boundary layers, and so modify the momentum contained within the corner region. Since the corner momentum is expected to strongly influence the onset of separation in this region, a better understanding of the vortices is central to developing insight into the key flow physics.

One major barrier to developing this insight is that the corner vortices are not generally captured by standard numerical techniques commonly used in industry. This difficulty is exacerbated by a distinct lack of high-quality experimental data on the corner boundary layers in supersonic channels. The scarcity of relevant measurements hinders the validation of numerical techniques which seek to accurately capture the corner vortices. It is particularly problematic that the few experimental studies on these flows, conducted in different facilities, have reported quite inconsistent flow topologies, and the cause for these discrepancies has not yet been isolated [12–14].

The aim of this study is therefore to answer the question: “What are the key physical mechanisms governing the corner boundary layers of supersonic channel flows?” This is achieved through the collection of high-quality experimental data, which enables validation of relevant numerical methods, a particular focus being their ability to predict vortices in the corner regions.

Chapter 2

Shock wave–boundary-layer interactions in channels

This chapter presents a review of the literature relating to the interaction between shock waves and turbulent boundary layers in rectangular channel flows. The key physics of the well-understood two-dimensional flow fields are introduced. The mechanisms which modify these flow fields in channels are then explained, with a focus on the significant impact of the low-momentum corner regions. The chapter ends by summarising the current understanding of the corner boundary layer's structure – this is necessary to reliably compute the momentum of the corner flow and, in turn, the entire flow field.

2.1 Two-dimensional shock–boundary-layer interactions

In an inviscid fluid, a shock corresponds to a discontinuous increase in pressure. When the shock impinges on a viscous boundary layer, however, this pressure rise is significantly smeared. The mechanism governing this effect relies on the subsonic channel in the near-wall portion of the boundary layer. This channel permits upstream travel of pressure information, so that the flow is 'aware' of the upcoming pressure rise before encountering the shock itself.

The boundary layer responds to this pressure rise by growing thicker or, for sufficiently high adverse pressure gradient, by separating. In either case, the displacement effect of the boundary layer causes streamline curvature in the external flow. The pressure waves associated with this streamline curvature modify the shock structure, which in turn affects the pressure information transmitted upstream through the subsonic channel.

The feedback process described above sets up an equilibrium flow structure. This type of feedback system underlies every shock wave–boundary-layer interaction (SBLI). Normal shock interactions are, however, different in nature to oblique shock reflections. Similarly, there are fundamental differences in flow field depending on whether the boundary layer remains attached or separates. These different

cases (i.e. normal shock versus oblique shock, and attached versus separated boundary layers) are discussed in turn below.

2.1.1 Normal shock interactions

The current study is focused on the canonical case of the oblique shock reflection. However, it is useful to first consider the normal shock interaction in order to introduce some key flow features which can be found in all types of SBLIs.

Attached normal shock interactions

For sufficiently low adverse pressure gradient – typically corresponding to a normal shock at $M \lesssim 1.3$ – the incoming boundary layer is able to stay attached.

The equilibrium flow structure is depicted in figure 2.1b. This shows how the pressure rise transmitted through the subsonic channel causes the boundary layer to thicken upstream of the main shock. The associated concave streamline curvature generates a series of compression waves, which coalesce into the normal shock. This wave pattern effectively produces a smeared shock foot.

Figure 2.1c, shows significant stagnation pressure loss near the wall, due to viscous dissipation inside the boundary layer. Far away from the wall, the stagnation pressure loss corresponds to the inviscid case. In between these two limits, however, the stagnation pressure recovery exceeds the inviscid case. In this region, the flow encounters isentropic compression waves followed by a weaker shock.

The smearing of the static wall pressure can be identified in figure 2.1d, where the pressure rise starts upstream of the inviscid shock location. The pressure then gradually approaches that expected for an inviscid shock. The smearing effect can be significant – in fact, it can take up to 70 boundary-layer thicknesses for the pressure to reach the theoretical inviscid post-shock value [16].

Spreading the pressure rise over a greater streamwise distance also reduces the adverse pressure gradient, and so tends to help keep the boundary layer attached. Figure 2.1e presents the streamwise variation in the incompressible shape factor, H_i, over the SBLI.[1] This parameter peaks at the interaction itself, before the boundary layer begins to recover further downstream.

Separated normal shock interactions

As the shock strength is increased, there will be a critical point at which the boundary layer separates. This typically corresponds to an incompressible shape factor, H_i, at separation between 2.2 and 2.7 [17–19]. While the separation bubble itself may be very thin, it has a profound effect on the overall flow structure. The wave pattern corresponding to a separated normal shock interaction is shown in figure 2.2b.

The separation bubble contains a region of recirculating fluid. This tends to result in an enhanced feeding of pressure information upstream, and thus the interaction length is larger than for the attached

[1] In order to avoid Mach-number dependence, the incompressible definition of the shape factor is used as a measure of boundary-layer fullness. For more details, including the definition itself, the reader is referred to section 3.2.4.

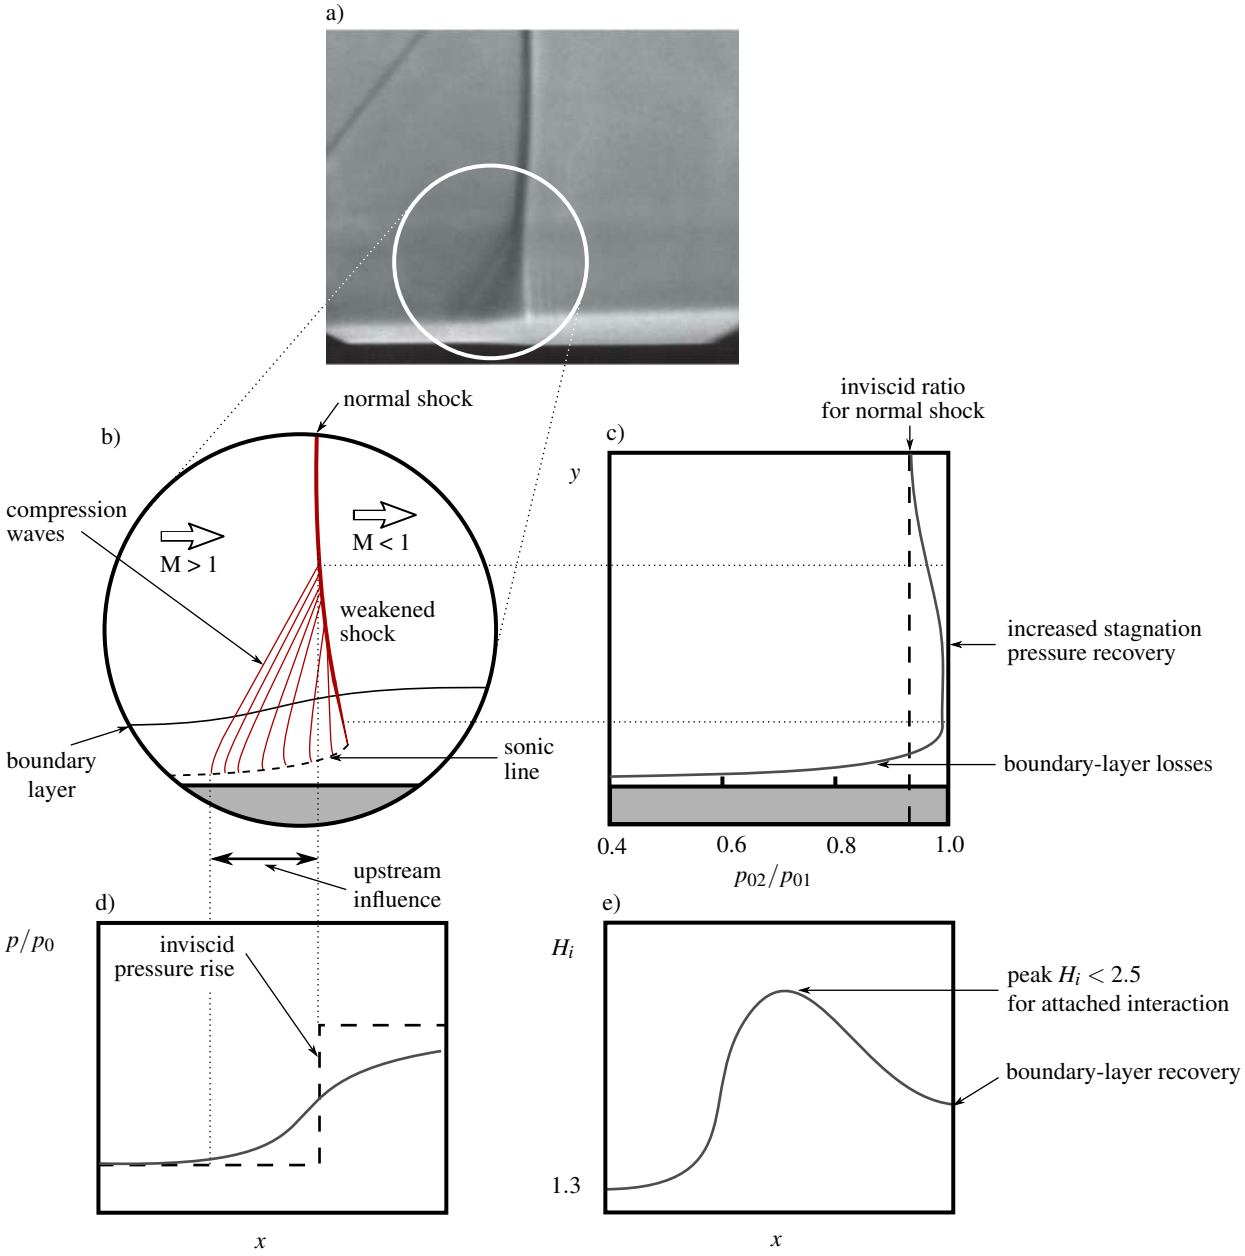

Figure 2.1: Attached normal shock interaction: a) schlieren image; b) schematic of equilibrium flow structure; c) recovered stagnation pressure profile; d) wall static pressure distribution; e) variation in incompressible shape factor, H_i. Adapted from [15].

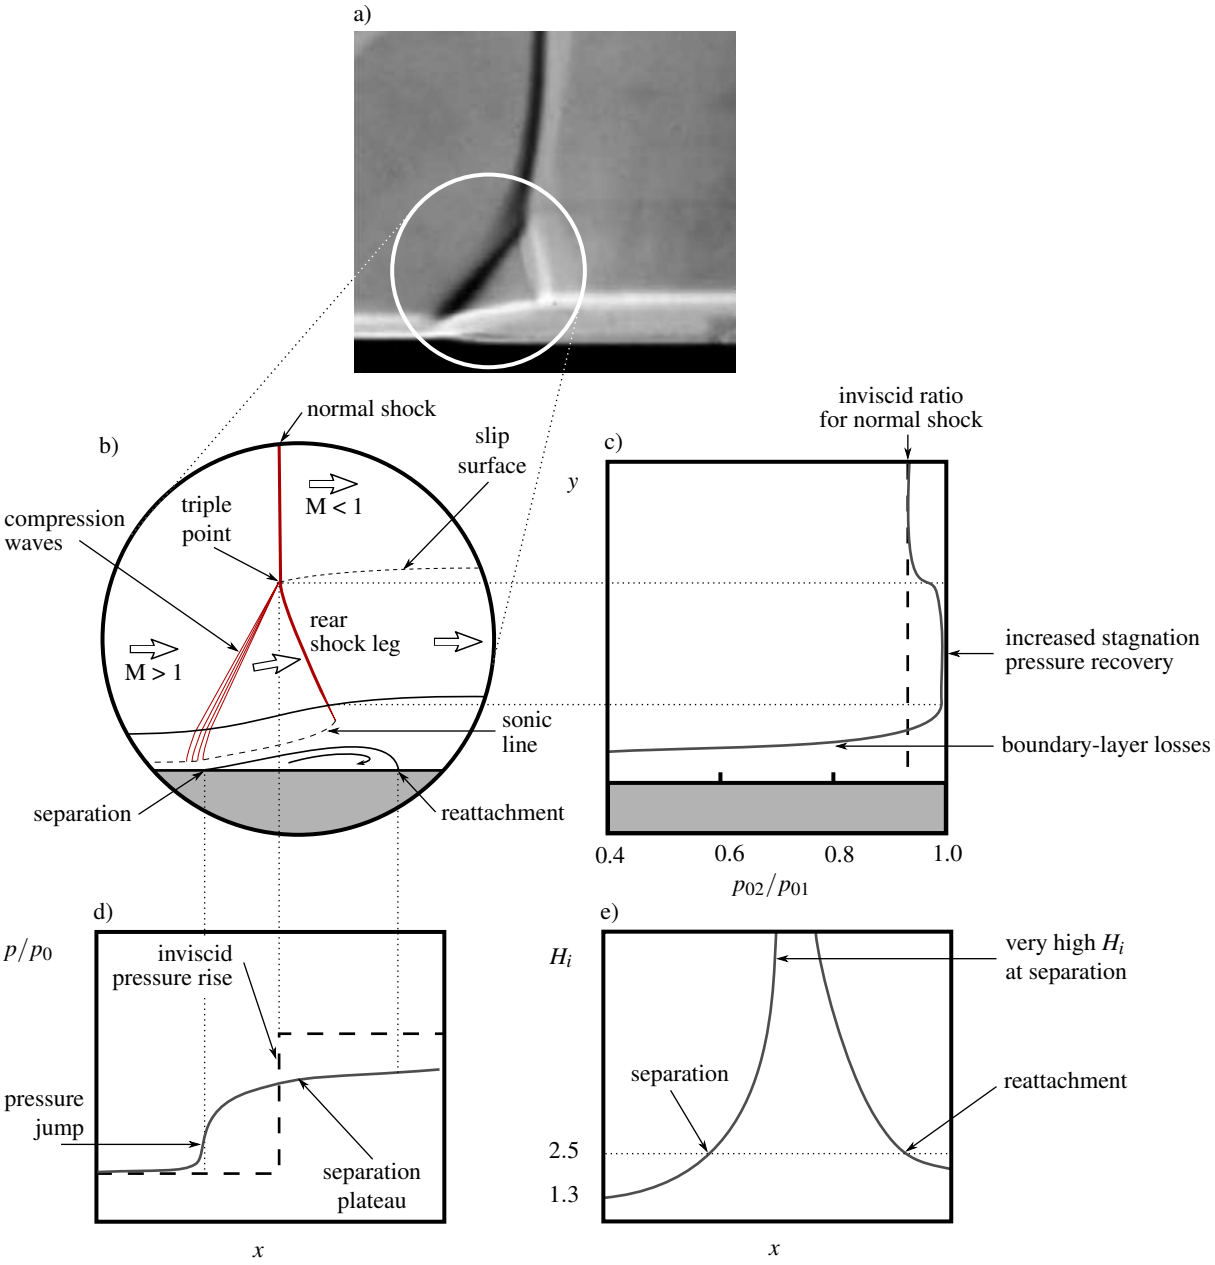

Figure 2.2: Separated normal shock interaction: a) schlieren image; b) schematic of equilibrium flow structure; c) recovered stagnation pressure profile; d) wall static pressure distribution; e) variation in incompressible shape factor, H_i. Adapted from [15].

case. Furthermore, the concave streamline curvature due to the presence of the separation bubble is associated with a series of compression waves. These waves coalesce into an oblique shock, which merges with the main normal shock. The pressure rise through oblique shock alone is less than that experienced by the outer flow passing through the normal shock, and so a rear shock leg is formed to create a 'lambda-shock' structure. The front and rear leg of the lambda-shock intersect at the 'triple point'.

A slip surface originates at the triple point [20]. This surface corresponds to a discontinuity in stagnation pressures and entropy, since these properties change differently through the two-shock lambda-foot compared to through the normal shock. In contrast, the static pressure and the flow direction either side of the slip surface are the same, since the velocity must be tangential to this streamsurface.

Figure 2.2c presents the stagnation pressure distribution downstream of the interaction. The viscous losses in the boundary layer are far greater than in the attached case, due to the presence of the separation bubble. For reasons similar to the attached case, the stagnation pressure recovery for the flow passing through the legs of the lambda-shock is greater than for the outer flow encountering the normal shock. The discontinuity in stagnation pressure across the slip line can also be observed.

The static streamwise pressure distribution (figure 2.2d) has a different structure to the attached case. At the separation point, upstream of the shock, there is a sharp increase in pressure – the magnitude of this increase is only weakly dependent on Mach number. There is then a pressure plateau, corresponding to a relatively uniform pressure distribution beneath the separation bubble. At the reattachment point, the wall pressure begins to increase asymptotically towards the inviscid value. In reality, the rear leg is locally smeared near the wall, and therefore a gradual pressure rise over the separation region is generally observed [21].

The shape factor, shown in figure 2.2e, increases dramatically in the separation region, corresponding to greatly distorted velocity profiles. On reattachment, H_i drops once more as the boundary layer recovers to an equilibrium profile.

Note the convex curvature of the boundary-layer edge downstream of the main shock or the rear shock leg in figures 2.1b and 2.2b, respectively. There are expansion waves associated with this curvature. The origin of these waves can also be attributed to the shock reflecting from the sonic line (a surface of constant pressure) as expansion waves. These waves have the potential to re-accelerate the flow to create a secondary pocket of supersonic flow immediately downstream of the main interaction. These secondary supersonic regions are not relevant to the current study, however, and so this feature is not described in detail. Instead, a more complete description is provided in reference 15.

2.1.2 The oblique shock reflection

The oblique shock reflection is the canonical flow field of interest in this thesis. Oblique shocks are inherently different from the normal shock case, since the core flow remains supersonic downstream of the interaction.

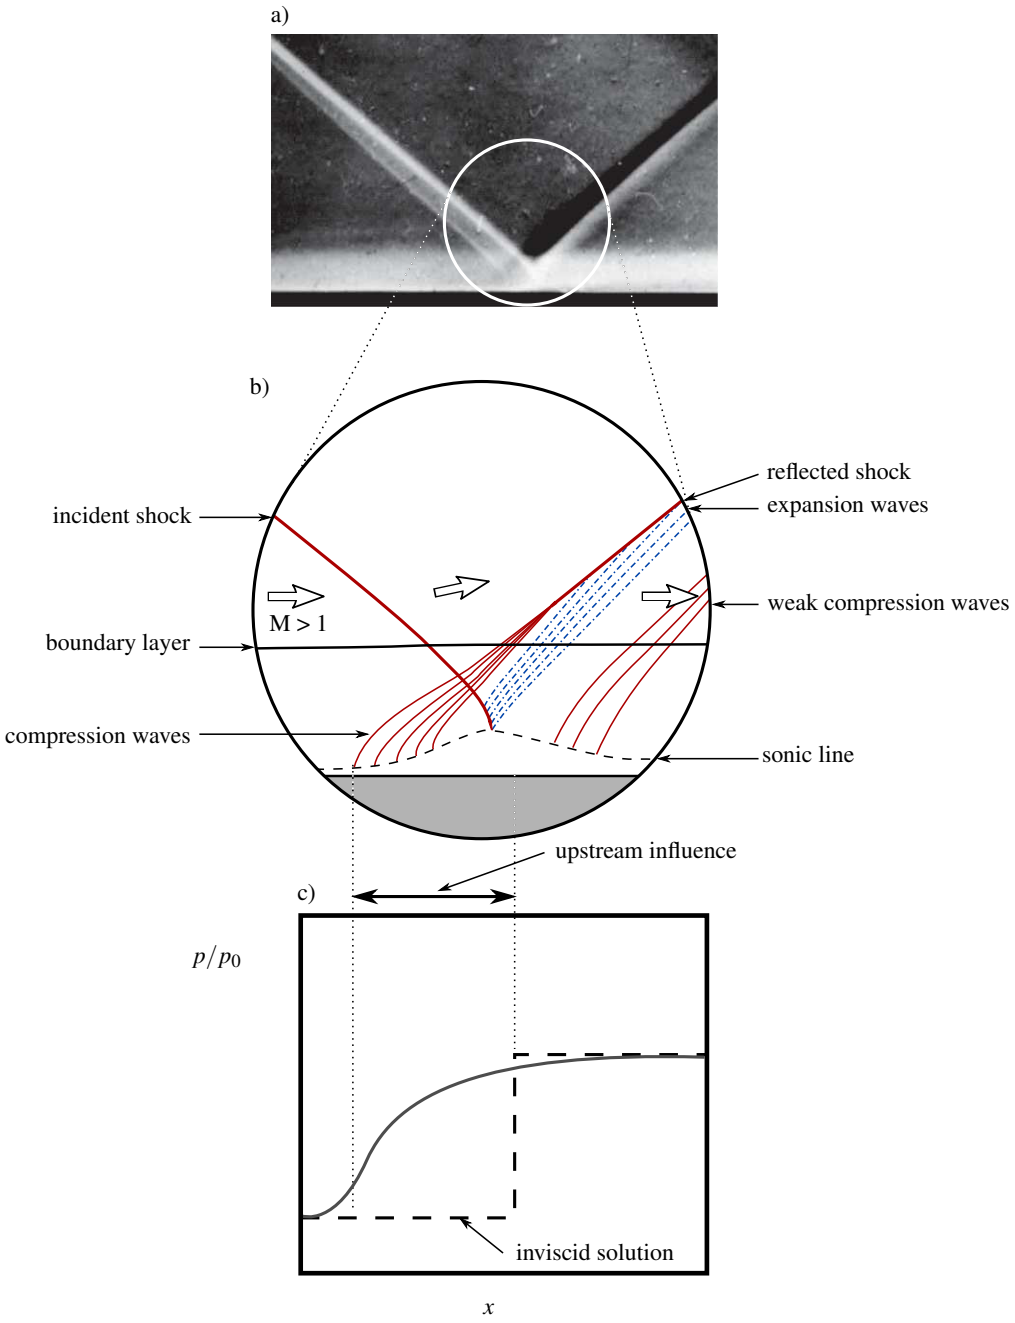

Figure 2.3: Attached oblique shock reflection: a) schlieren image; b) schematic of equilibrium flow structure; c) wall static pressure distribution. Adapted from [15, 22].

Attached oblique shock reflection

For a weak incident oblique shock, the flow will generally remain attached. A typical attached oblique shock–boundary-layer interaction is presented schematically in figure 2.3b. As in the normal shock interaction, the subsonic region within the boundary layer provides a channel to feed pressure information upstream. This leads to a streamwise smearing of the pressure increase (figure 2.3c). For an attached flow, the upstream influence is short, typically on the order of a boundary-layer thickness.

The boundary layer responds to the imposed pressure rise by thickening, upstream of the inviscid interaction location. The resulting concave flow curvature, visible in figure 2.3b, generates a series of compression waves, similar to the normal shock interaction. These waves coalesce to form the reflected shock.

Meanwhile, as the incident oblique shock enters the boundary layer, the local Mach number gradients cause it to bend progressively and to disperse into a series of compression waves within the boundary layer. When these compression waves encounter the constant pressure boundary formed by the sonic line, they are reflected as expansion waves. Downstream of the interaction, some weak compression waves are generated due to the streamline curvature as the boundary layer recovers and the flow returns to horizontal.

The complex wave pattern described above is embedded within the boundary layer, and so the outer flow is characterised simply by the incident and reflected shock. This corresponds to the flow field expected for the inviscid case. Viscous effects have only a small influence on the overall flow, and so this is often described as a ‘weak’ interaction process.

Separated oblique shock reflection

As the strength of the incident shock is increased, it imposes a larger adverse pressure gradient on the boundary layer, which eventually separates. The equilibrium flow structure is presented in figure 2.4b. The presence of a separation bubble, with its convex profile, defines the streamline curvature of the flow and thus the corresponding wave pattern.

At the separation point, there is concave streamline curvature outside the separation bubble. As a result, a series of compression waves is generated; these coalesce into an oblique ‘separation’ shock. When this separation shock intersects the incident shock, it is refracted to form the reflected shock. Simultaneously, the incident shock is refracted as the transmitted shock.

Similarly to the attached interaction, the transmitted shock is reflected from the sonic line as a series of expansion waves. These waves deflect the flow back towards to the wall at the top of the separation bubble. At the reattachment point, there is concave streamline curvature as the flow returns to horizontal. This generates a series of compression waves, which coalesce into the reattachment shock. In practice, this reattachment shock is typically smeared and reasonably weak in nature.

Note that the stagnation pressure of the flow that has passed through the incident/reflected shocks is different to the flow that has encountered the separation/transmitted shocks. This results in the generation of a slip line originating at the intersection point of these shocks.

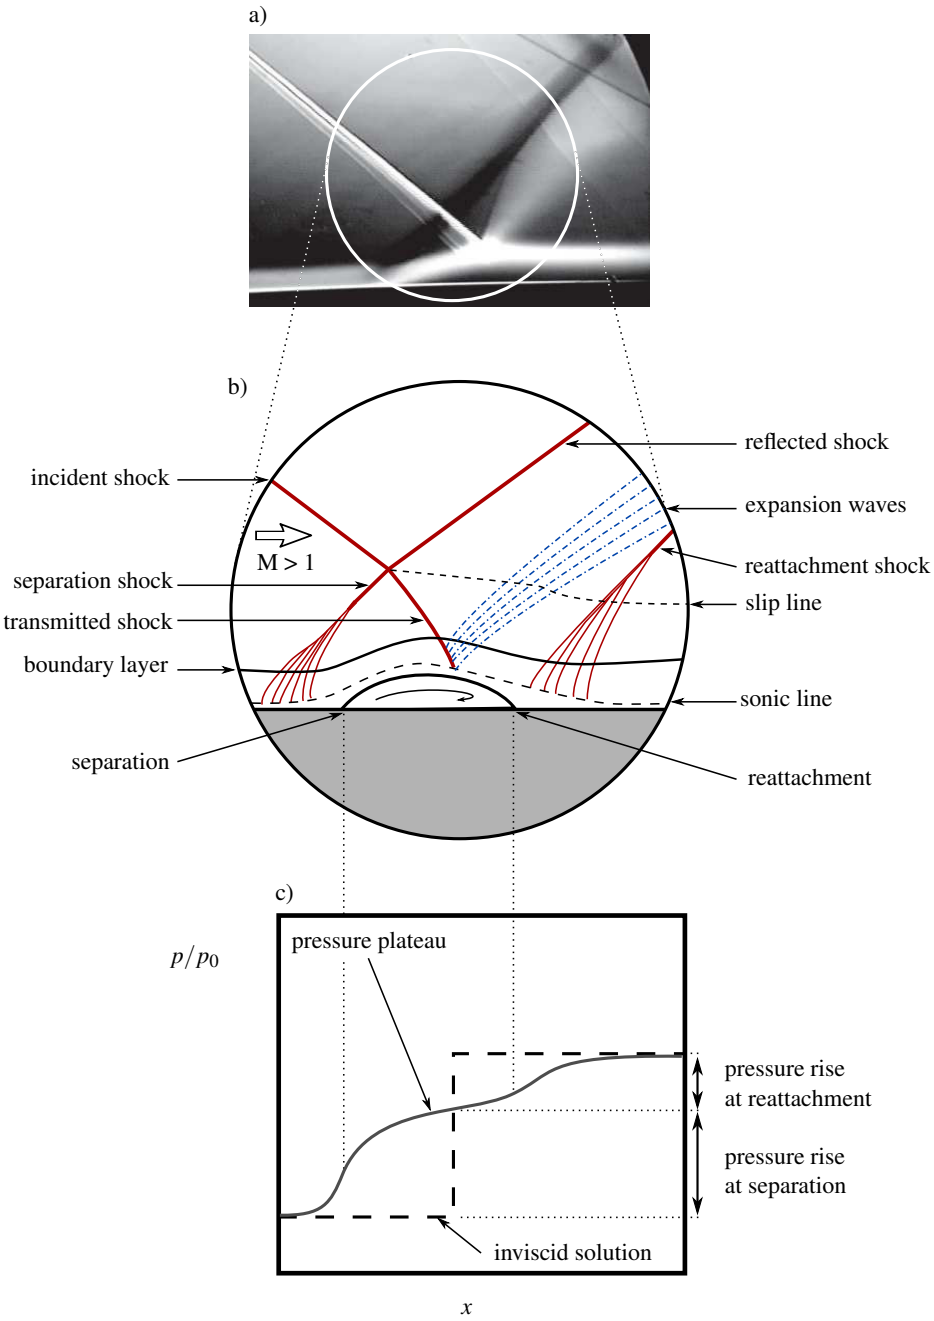

Figure 2.4: Separated oblique shock reflection: a) schlieren image; b) schematic of equilibrium flow structure; c) wall static pressure distribution. Adapted from [15].

Figure 2.4c presents the wall pressure distribution in the streamwise direction, and shows a two-step increase in pressure. The first step corresponds to the compression waves associated with the separation shock. This is followed by a plateau, with the separation bubble itself exhibiting a uniform pressure distribution. The second, more gentle and often smeared, pressure rise occurs on reattachment.

In contrast to the attached case, the features of the separated flow field demonstrate significant departure from the inviscid solution. This is therefore described as a 'strong' interaction process.

2.1.3 Onset of separation

Since the attached and separated flow fields are so different, it is important to be able to accurately predict the onset of separation. Intuitively, there are two factors which could affect when separation occurs – the strength of the shock and the properties of the incoming boundary layer.

A stronger shock imposes a larger pressure rise, which promotes separation of the incoming boundary layer. Note, however, that the key parameter governing whether the flow separates is not the magnitude of the pressure rise itself. Instead, a better measure of separation onset is the adverse pressure gradient experienced by the boundary layer.

The importance of adverse pressure gradient was elegantly demonstrated by Smith and Smits [23]. A two-dimensional compression-corner geometry with pressure jump ratio of 3.7 exhibits separation. However, when the sharp corner is replaced by a more rounded geometry, the flow remains attached. The pressure rise is identical for both setups. Therefore, the attached flow with the rounded corner is due to the greater streamwise distance over which the pressure increases, resulting in a lower adverse pressure gradient.

With regards to the boundary-layer properties, one might expect a boundary layer with lower momentum (i.e. higher shape factor) to separate more easily. However, in reality, the shape factor of the incoming boundary layer has a relatively minor effect. Indeed, Délery found the separation behaviour for the normal shock interaction to be almost independent of boundary-layer shape factor, after plotting the experimental data shown in figure 2.5a. The onset of separation instead appears to be governed mainly by the freestream Mach number, which determines the shock strength.

The insensitivity of separation to shape factor is somewhat counter-intuitive. However, this phenomenon can be explained as a result of two competing effects. A low shape factor does, indeed, increase the intrinsic resistance of the boundary layer to separation. On the other hand, this low shape factor boundary layer has a smaller subsonic channel to feed pressure information upstream, which results in a shorter interaction length. The pressure rise is therefore less smeared, and so the boundary layer experiences a stronger adverse pressure gradient, making it more prone to separate. These two competing effects tend to approximately compensate for one another, as illustrated by the near-horizontal line in figure 2.5a. Indeed, there are multiple studies which show that separation tends to occur, regardless of boundary-layer shape factor, for Mach numbers ranging from 1.30 to 1.35 [22, 16, 25].

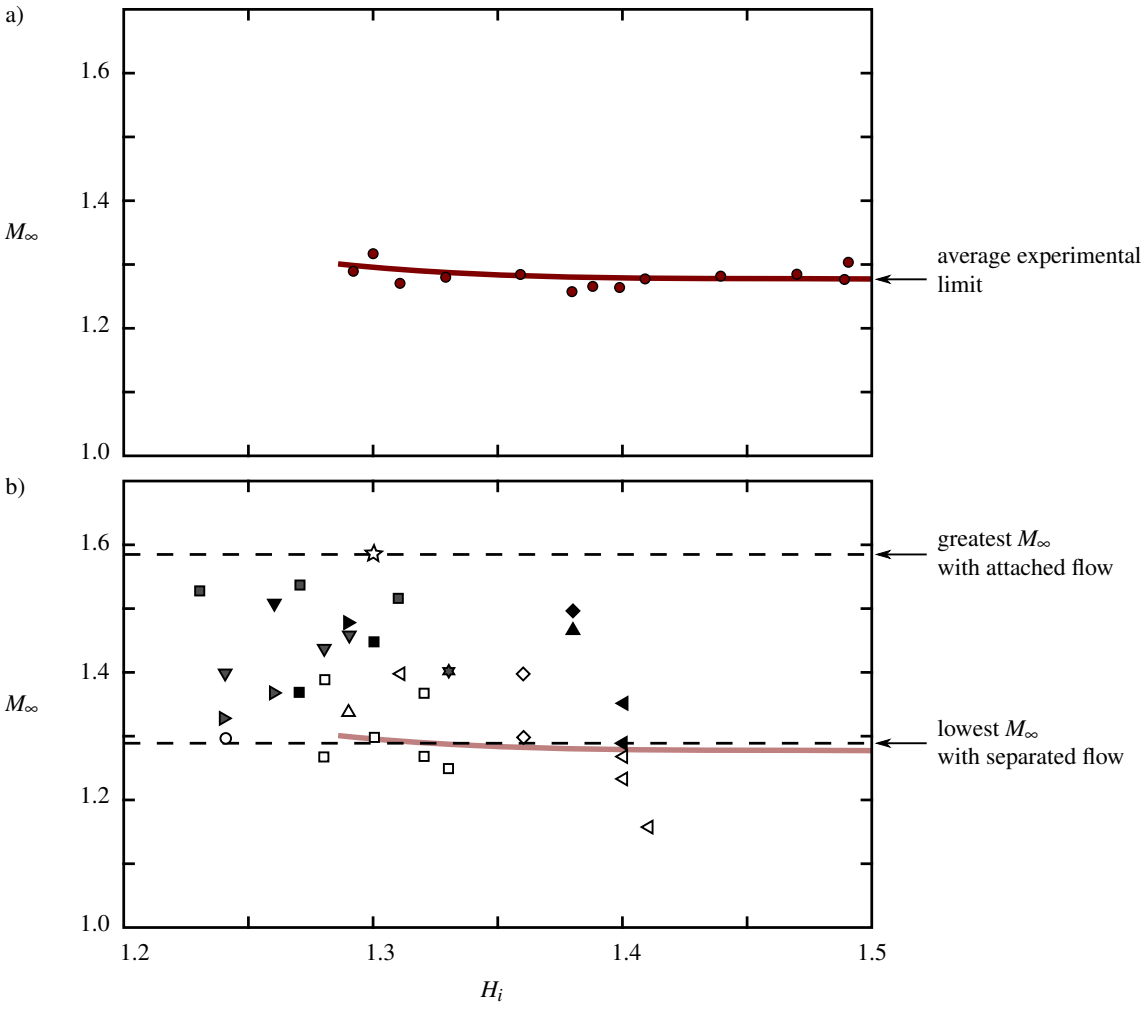

Figure 2.5: Experimental data for normal SBLIs, plotting freestream Mach number (M_∞) against incompressible shape factor of incoming boundary layer (H_i). a) Circles mark the freestream Mach number at which separation occurs for a given incoming boundary-layer state. Adapted from [22]. b) Nature of normal shock interaction obtained from a comprehensive survey by Bruce: open symbols correspond to attached flow and filled symbols represent separation. Each symbol shape represents a different experimental study, with definitions provided in reference 24. The average experimental limit from a) is indicated with a red line for reference. Adapted from [24].

A more complex picture emerges, however, when one considers the more comprehensive survey conducted by Bruce [24]. Figure 2.5b shows the boundary-layer state observed in studies on normal shock interactions conducted at many different experimental facilities [26]. The preceding discussion appears to suggest that all cases above the red line should display separation, whilst those below the line should correspond to attached flow. Instead, however, there is a significant variation in the critical Mach number required for separation – separated flows can be found at Mach numbers as low as 1.29 while there are cases of attached flow at Mach numbers up to 1.59. No clear trend can be observed with the shape factor of the incoming boundary layer, and the scatter cannot be explained by two-dimensional separation theories.

The description of shock–boundary-layer interactions presented thus far is entirely two-dimensional. In contrast, experimental studies (such as those from figure 2.5b) are carried out in wind tunnels with rectangular cross-section. Interactions studied in these facilities are fundamentally three-dimensional. A closer examination of the flow in this type of geometry helps to explain the apparent inconsistencies in predicting separation onset, and the relevant mechanisms are outlined in the following section.

2.2 Shock–boundary-layer interactions in a channel

The two-dimensional theory of shock–boundary-layer interactions is useful to introduce key physical concepts. However, there is substantial evidence [27–30] that this theory cannot be directly applied to three-dimensional geometries like the rectangular channel, which exhibit significant departures from the corresponding two-dimensional behaviour.

Chapter 1 has already highlighted the relevance of this type of channel flow, which appears in industrial applications such as the inlets of supersonic aircraft. More immediately, our understanding of shock–boundary-layer interactions is derived from experimental studies, which typically examine the response of the boundary layer growing along the wind tunnel floor to an incident shock. Since supersonic wind tunnels often have a rectangular cross-section, it is necessary to understand the impact of the three-dimensional channel geometry to develop further insight into this type of flow.

The geometry of a rectangular channel differs from the two-dimensional flow field in two main ways. Firstly, the flow does not extend to infinity, but is confined. The confinement imposes straight streamsurfaces on the symmetry planes, rather than at infinity for the unconfined case. This enhances any wall-normal variations and can strongly influence the flow field, particularly for normal shock interactions [15]. This effect is enhanced by the fact that there are boundary layers growing on all four channel walls, and these will all respond to the adverse pressure gradient imposed by the shock by thickening. As a result, the effective area of the channel is substantially reduced. Note, however, that most wind tunnels have boundary layers which occupy only a small proportion of the total cross-sectional area. For these facilities, the impact of flow confinement is generally assumed to be minor.

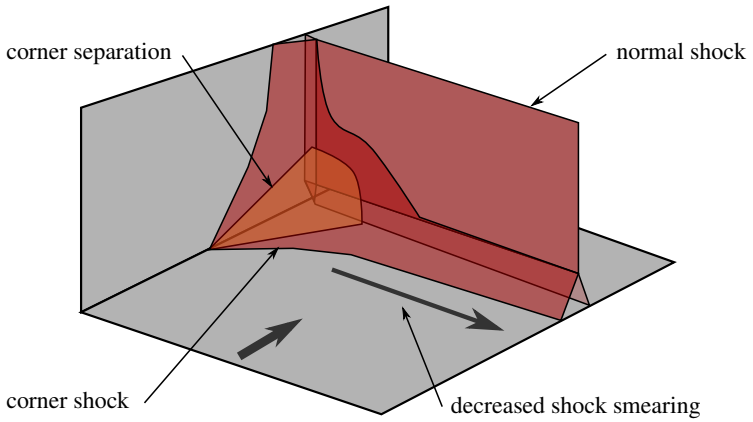

Figure 2.6: Schematic showing effects of corner separation on flow field for a normal shock interaction. Adapted from [31].

The second, dominant influence of the rectangular channel geometry is the presence of the corner regions, where the sidewalls meet the channel floor (and ceiling) at a right-angled, internal corner. The flow in this region can be treated as the intersection of the boundary layers growing on the floor and on the sidewall. Since both constituent boundary layers contain viscous flow, the corner region corresponds to an area with particularly low momentum flow.

The corner boundary layer does, in fact, have a more complex flow structure which can influence the local momentum, and this is detailed in section 2.3. It is necessary to account for this aspect of the corner flow to accurately compute the flow numerically. However, the resultant change in momentum is relatively small. As a result, the key physical concepts of shock–boundary-layer interactions in a channel can be understood by treating the corner flow simply as a region with low momentum.

In response to the adverse pressure associated with an impinging shock, the local low momentum causes the corner region to separate furthest upstream. The effect of this early separation on the flow field depends on whether the incident shock is normal or oblique in nature. The normal shock interaction will be summarised first, since it provides useful background of the key physical mechanisms. The current state of understanding how sidewalls influence the oblique shock reflection will then be discussed.

2.2.1 Sidewall effects on normal shock interactions

The early separation of the corner boundary layer, due to its low momentum, causes a displacement effect which requires the external flow to deflect. This streamline curvature is associated with a series of compression waves, which propagate into the flow field. These corner waves influence the pressure profile experienced by the rest of the flow. For simplicity, the compression waves are treated as a single, weak, conical shock wave, and this is depicted schematically in figure 2.6.

Figure 2.6 also shows the impinging normal shock, which has a lambda-shock structure near the floor if it is strong enough to cause separation. The corner shock merges with this lambda-shock

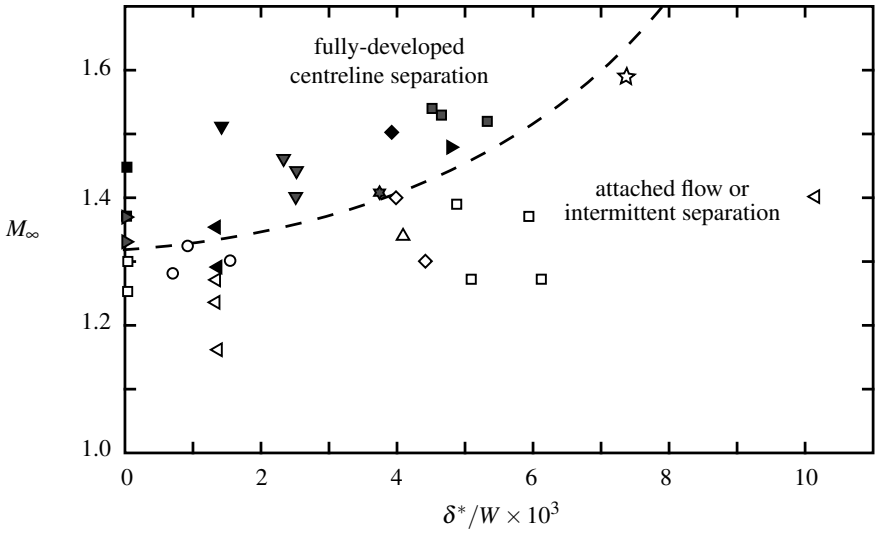

Figure 2.7: Experimental data for normal SBLIs plotted against viscous aspect ratio. Open symbols correspond to attached flow and filled symbols represent separation. Adapted from [24, 31]. Each symbol shape represents a different experimental study, with definitions provided in reference 24.

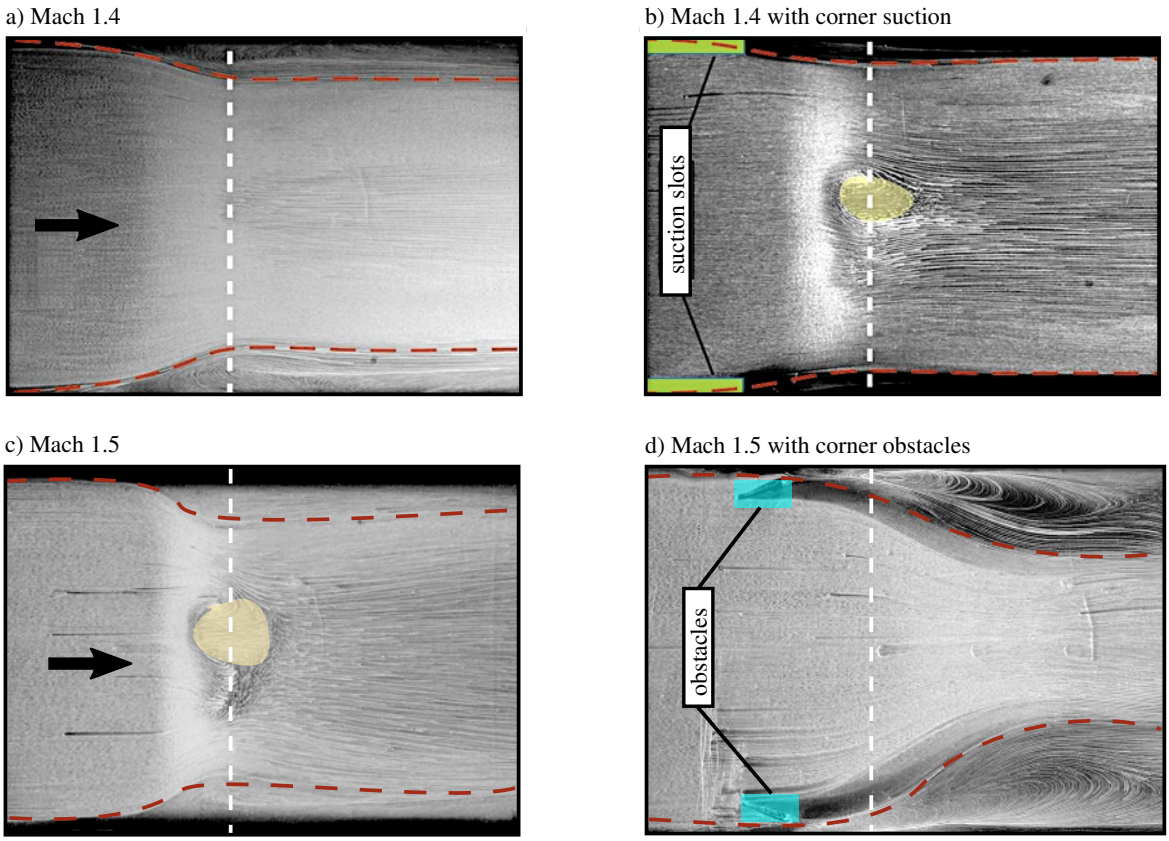

Figure 2.8: Surface oil-flow visualisation: a) attached flow with Mach 1.4 normal shock; b) use of suction to reduce corner separation; c) separated flow with Mach 1.5 shock; d) use of corner obstacles to increase corner separation. Any central separation is highlighted in yellow, and the boundary of the corner separation is marked with a dashed red line. Adapted from [24, 31].

foot. As a result, the distance between front and rear legs of the lambda-shock increases close to the sidewall. The increased streamwise interaction length reduces the adverse pressure gradient, and so delays or decreases the observed shock-induced separation. Further away from the sidewall, the corner effects become weaker and so separation is more likely.

It is reasonable to expect that the most two-dimensional flow field in the channel would be on the centre span, where the corner effects are weakest. The extent of shock smearing on this central plane depends on the relative size of the corner separation to the tunnel width. As the relative size of the corner separation is increased, the smearing is more pronounced so a stronger shock is required to separate the flow on the centre line.

In order to test this trend across many different facilities, it is necessary to know the size of corner separation for each physical experiment. However, this quantity is not commonly measured or reported. An approximate alternative parameter is based on the assumption that corner separation scales with the displacement thickness, δ^*, of the incoming boundary layer. This leads to the definition of the 'viscous aspect ratio', δ^*/W, where W is the tunnel width. When the data from figure 2.5b are replotted as a function of the viscous aspect ratio (figure 2.7), it is clear that this new parameter provides a reasonable measure by which to predict the effects of confinement [24].

Therefore, a larger corner separation should result in a more benign interaction on the centre span, and vice versa. Physical experiments have been conducted by Bruce *et al.* to directly test this proposed relationship [24]. Figure 2.8a shows a Mach 1.4 flow, which exhibits separation only in the corners – the centre-span flow remains attached. In figure 2.8b, suction slots are used in the corner regions to remove the lowest-momentum part of the flow, which achieves a reduction in the size of corner separation. At the same time, the oil-flow visualisation shows separation of the previously-attached centre-span flow.

Equivalently, the Mach 1.5 case shown in figure 2.8c exhibits separation on the centre line. Obstacles, in the form of vortex generators, placed in the corners are observed to increase the size of corner separation (figure 2.8d). Moreover, the centre-span flow is transformed from separated to attached, in agreement with the proposed trend. These results are consistent with other studies performed at a range of Mach numbers and in various facilities [27, 32–34].

2.2.2 Sidewall effects on oblique shock reflections

The flow field for the oblique shock reflection in a rectangular channel also tends to exhibit significant three-dimensionality [35–38]. As for normal shock interaction, waves generated by the corner separation provide the fundamental link between the corner regions and the centre-span separation [8, 39].

Analogous to the normal shock case, the early separation of the low-momentum corner region generates a series of compression waves. These waves can be simplified as a conical corner shock, as shown in figure 2.9. However, the core flow downstream of an oblique shock reflection remains

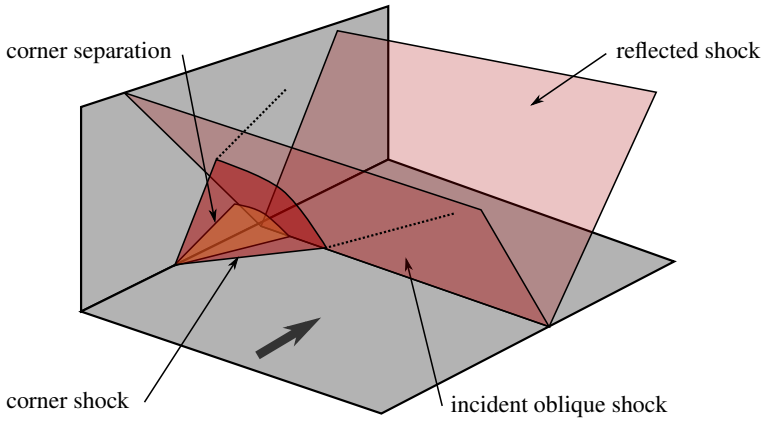

Figure 2.9: Schematic showing effects of corner separation on flow field for an oblique shock reflection. The detailed two-dimensional wave structure is represented by an inviscid oblique shock reflection for simplicity. Adapted from [31].

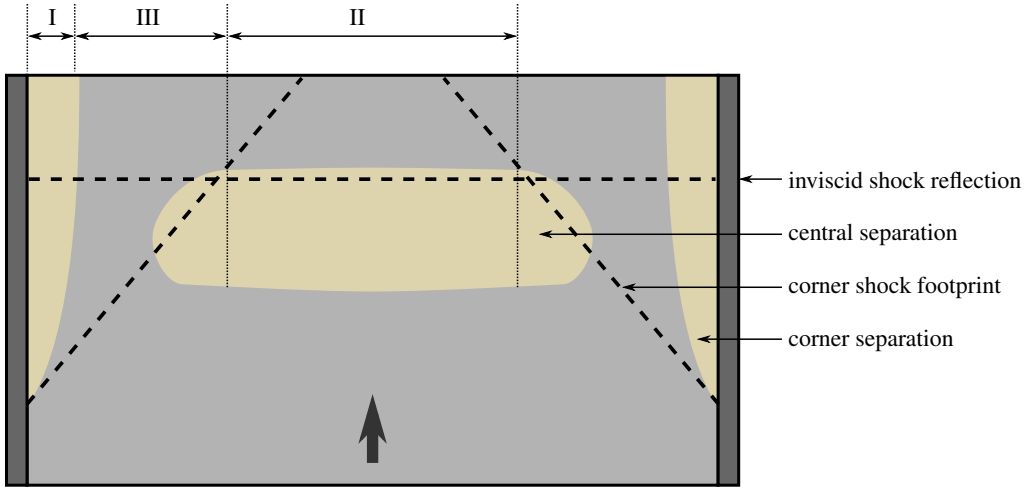

Figure 2.10: Flow field along floor for an oblique reflection in a channel, divided into regions: (I) corner separation region; (II) central separation region; (III) intermediate region. Adapted from [31].

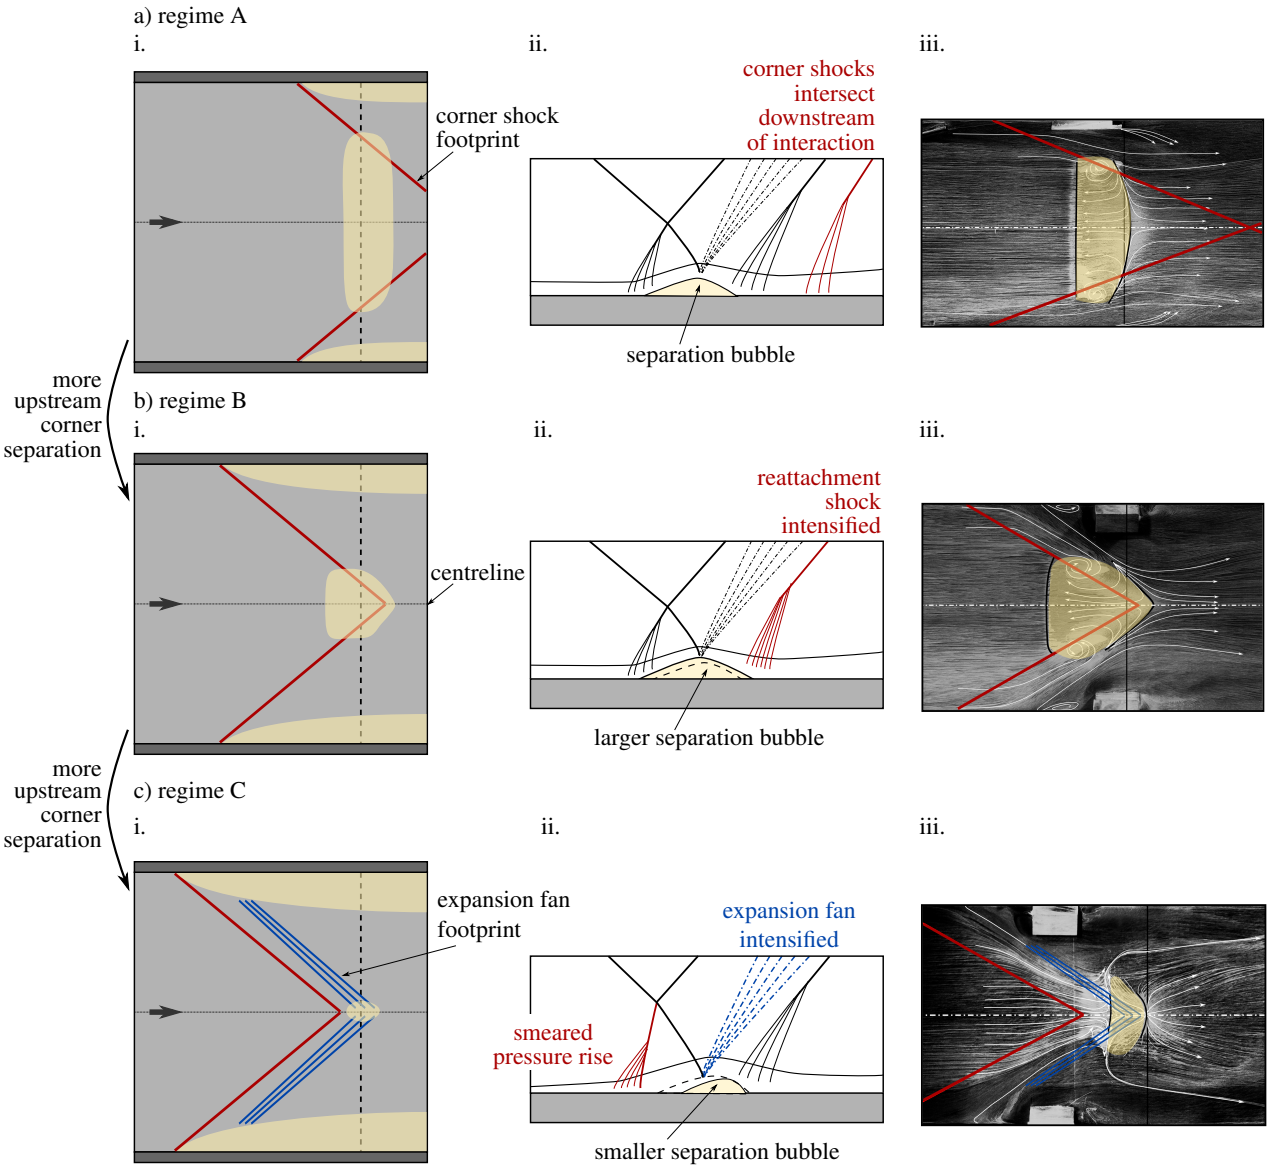

Figure 2.11: Corner effects on centre-span separation for the oblique shock reflection: a) regime A (small corner separation), b) regime B (medium corner separation), c) regime C (large corner separation). For each flow field: i. a schematic footprint of the flow field on the tunnel floor; ii. a shock structure schematic on the tunnel centre line; iii. oil-flow visualisation on the tunnel floor. Adapted from [6].

supersonic. As a result, the corner shock is able to pass through the incident and reflected oblique shocks. This is fundamentally different from the merging behaviour observed for normal shocks.

Figure 2.10 presents the footprint of the typical flow on the floor. This shows the corner shocks passing through the incident oblique shock. The flow field can be divided into three separate regions: (I) the corner separation region, which originates furthest upstream; (II) the central separation region, which is almost two-dimensional; (III) a third intermediate region, in which the corner shock smears the interaction, lowering the adverse pressure gradient to reduce or eliminate separation [8]. The analysis of this flow field is more involved, however, since different regimes are observed – these depend on how far upstream the corner regions separate.

Effect of corner separation location

The effects of corner separation can, perhaps, be best understood in the context of a comprehensive set of experiments performed by Xiang and Babinsky [40, 6]. This study focuses on the response of the floor boundary layer to an oblique shock reflection at Mach 2.5, with an 8° flow deflection by the incident wave. The onset of corner separation is artificially varied and the resulting effect on the entire flow field is measured.

The corner separation point is moved in both the upstream and downstream directions. In order to move this point further upstream, small rectangular obstacles are placed ahead of the interaction, at the junction between the wind tunnel floor and sidewall. The corner boundary layer separates around these obstacles and thus the shock-induced corner separation is enhanced. On the other hand, a combination of two techniques are used to move the corner separation point further downstream. Vortex generators placed near the floor-sidewall junction serve to energise the corner boundary layer, while suction slots placed underneath the interaction in the corners removes the lowest-momentum flow in this region. The combined use of both methods is successful in delaying the onset of corner separation.

Figure 2.11a shows the flow field for the smallest corner separation (regime A), where vortex generators and suction are used in combination. As a result, the corner separation point has shifted downstream. The corner shocks therefore intersect, on the centre line, far downstream of the main interaction. This means that the waves have a negligible impact on the central separation. Indeed, as the corner separation point is moved further upstream, the streamwise length of the central separation remains roughly constant as long as the corner shocks intersect downstream of the interaction. The only difference between these cases is the width of the central separation. This parameter depends on the spanwise extent of the region where pressure smearing due to the corner waves enables attached flow. This causes the central separation to become narrower as the corner separation point moves further upstream.

As the corner separation point moves further upstream (figure 2.11b), the flow field shifts to regime B. Here, the corner shocks intersect within the interaction region. These waves impose an additional pressure rise on the flow before it can reattach, and so extend the central separation.

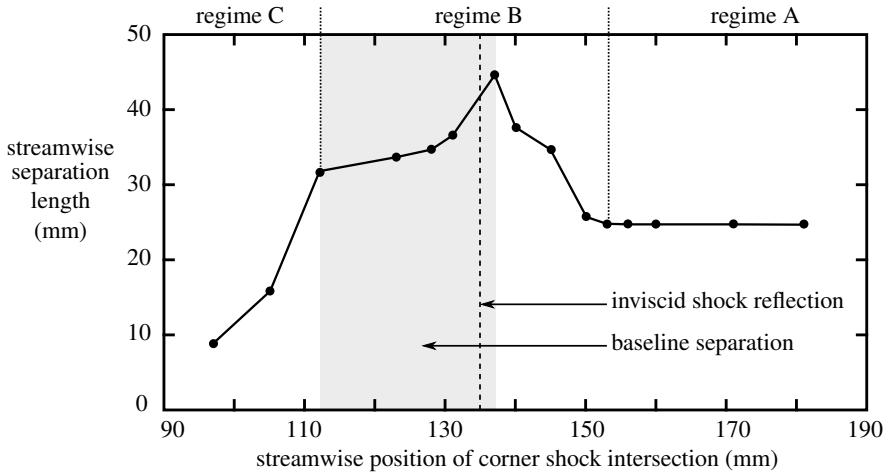

Figure 2.12: Effect of crossing point of corner shocks on centre-span separation. Adapted from [6].

Even more severe corner separation results in regime C, where the corner shocks originate far enough upstream that they intersect ahead of the main interaction. This case is shown schematically in figure 2.11c. Here, the corner separation point is so far upstream that another set of waves need to be taken into account. These are the expansion waves generated by the convex streamline curvature, where the corner separation begins to turn back towards the sidewall. Note that, in figures 2.11a and 2.11b, these corner expansion waves follow a trajectory downstream of the main interaction, and so do not influence the central separation. However, for the flow shown in figure 2.11c, the corner expansion waves intersect within the interaction region. As a result, it is necessary to consider the influence of these waves on the centre-span pressure profile.

In figure 2.11c, the centre-span flow sees a small pressure rise from the corner shocks ahead of the interaction. This gentle 'pre-compression' smears the overall pressure rise over a greater streamwise extent, reducing the adverse pressure gradient. In addition, the expansion waves reduce the pressure rise over the main interaction. The combination of these two effects results in a more benign interaction, with a shorter central separation.

Figure 2.12 collates the results from the entire experimental campaign. Instead of the corner separation point, the streamwise position of corner shock intersection is used as a more robust parameter. The variation in central separation size, simply by modifying the corner separation point, is significant. In fact, for very upstream corner separations (regime C), the streamwise extent of separation on the centre line can be even smaller than quasi-two-dimensional case (regime A) from figure 2.11a.

Similar trends can also be observed in a computational study performed by Benek *et al.* [41, 42]. RANS calculations are used to study flows at Mach 2.5, 2.7, and 2.9, with incident wave angles between 8° and 13°. Rather than changing the corner separation point artificially, the width of the tunnel is varied, and the resulting effects studied. Whilst the results have not been validated experimentally so should be treated with caution, the observed behaviour agrees well with the

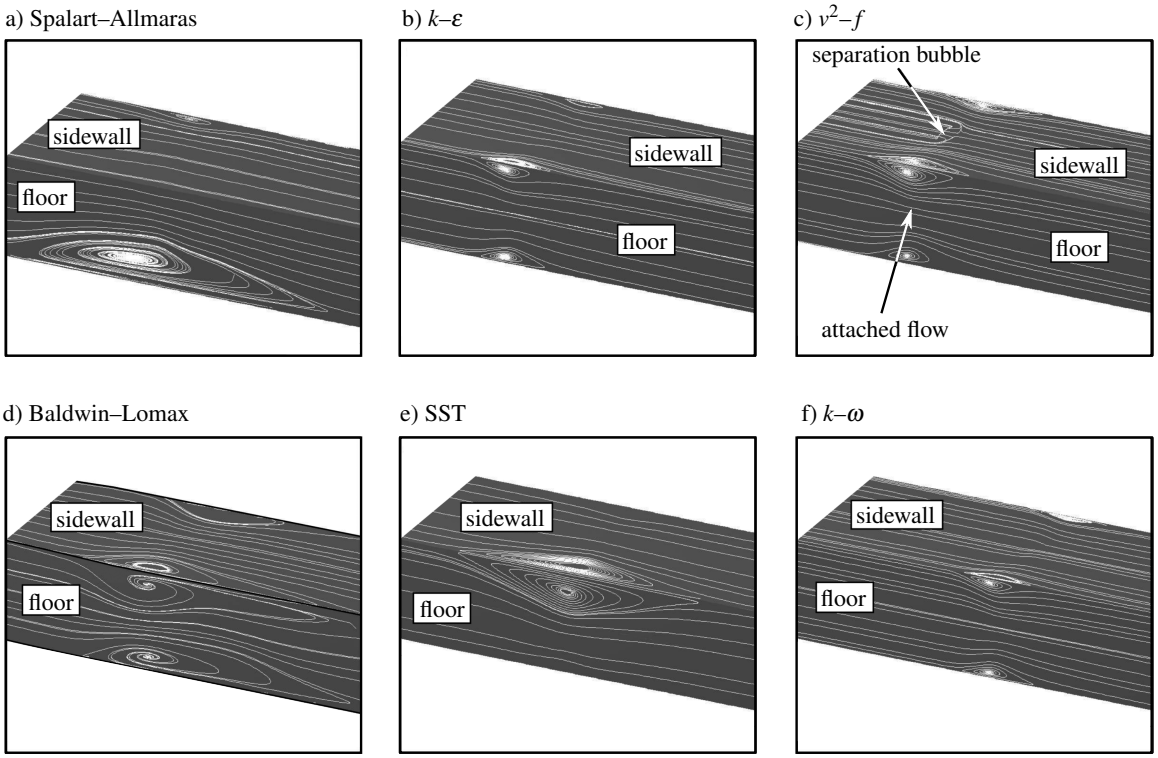

Figure 2.13: RANS computations for a normal shock at M = 1.4. The corner separation structure varies with turbulence model used: a) Spalart–Allmaras, b) k–ε, c) v^2–f, d) Baldwin–Lomax, e) SST, f) k–ω. Adapted from [44].

mechanisms already discussed. Higher fidelity large-eddy simulations have also been performed for different tunnel widths in a Mach 2.7 flow, by Wang *et al.* [43]. The waves generated by corner separation are captured by the simulations, and the trends of the central separation are consistent with the proposed effects of these waves.

The physical mechanisms linking corner separation to the rest of the flow field therefore appear to be relatively well understood. Despite this insight, computational methods commonly used in industry struggle to reliably predict flows with shock-induced separation in channels [44, 5]. This is partially due to a poor prediction of the corner waves, since the corner separation is typically computed inaccurately. The onset of this separation is a function of the local momentum, which is determined by the flow structure within the corner boundary layer. The complex flow in these corner regions is examined in the following section.

2.3 Structure of the corner boundary layer

The corner boundary layer is fundamentally the intersection of the sidewall with the floor boundary layer, and therefore has low momentum. However, embedded within this region is a pair of streamwise counter-rotating vortices [45].

The corner vortices transfer momentum between the core flow and the boundary layers, and so can modify the momentum of the corner flow. Difficulties in accurately capturing these vortices in simulations can help explain the poor prediction of corner separation, and thus the entire flow field, in channel flows with shocks.

The struggle to reliably capture corner separation is illustrated by Bruce *et al.* in figure 2.13. Here, a normal shock interaction in a Mach 1.4 channel is calculated using RANS methods with six different turbulence models. The various models produce quite distinct flow topologies, and exhibit significant differences in the existence and size of separation on the floor. Unsurprisingly, the corner separations also vary considerably between the six cases [44]. This can, at least in part, be attributed to the accuracy of corner momentum prediction and, in turn, the ability of different modelling approaches to capture corner vortices.

2.3.1 The origin of corner vortices

The origin of the corner vortex pair has been studied extensively through analytical treatments [46, 47], as well as numerical [48–51] and experimental [52–54] investigations. The production of the corner vortices can be understood by first considering the incompressible case [55, 56]. For a steady, fully-developed flow, the mean streamwise vorticity equation is [57]:

$$v\frac{\partial \omega_x}{\partial y} + w\frac{\partial \omega_x}{\partial z} = \underbrace{\omega_y\frac{\partial u}{\partial y} + \omega_z\frac{\partial u}{\partial z}}_{\text{A}} + \underbrace{\nu\left(\frac{\partial^2}{\partial y^2} + \frac{\partial^2}{\partial z^2}\right)\omega_x}_{\text{B}} + \underbrace{\left(\frac{\partial^2}{\partial y^2} - \frac{\partial^2}{\partial z^2}\right)\left(-\overline{v'w'}\right)}_{\text{C}} + \underbrace{\frac{\partial^2}{\partial y \partial z}\left(\overline{v'^2} - \overline{w'^2}\right)}_{\text{D}}, \tag{2.1}$$

where (x, y, z) denote the streamwise, wall-normal and spanwise coordinates. (u, v, w) are the corresponding components of the mean velocity, (u', v', w') are the turbulent fluctuations in velocity. The vorticity components are given by ω_x, ω_y and ω_z, while ν is the kinematic viscosity of the flow.

This equation shows the balance of the advection of vorticity (the left-hand side) with production and diffusion on the right-hand side. Term A on the right-hand side corresponds to the generation of Prandtl vortices of the first kind, corresponding to a bending of vortex lines by the mean shear. The second term (B) represents the viscous diffusion of vorticity. The third and fourth terms (C and D) correspond to the production of Prandtl vortices of the second kind. This vorticity production mechanism is related to anisotropies of the turbulent Reynolds stresses in the flow [58, 59].

The streamwise vortices which exist in corner flows are believed to be generated by mechanisms relating to Prandtl vortices of the second kind, corresponding to terms C and D [45, 51]. The relative

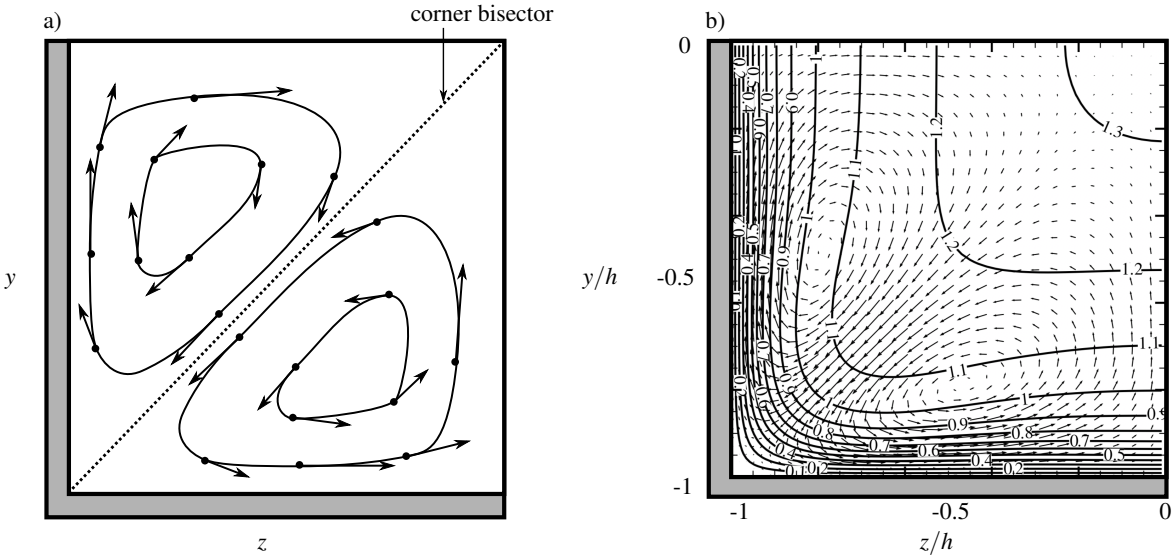

Figure 2.14: Secondary flow structure in a subsonic corner boundary layer. a) Experimentally-determined transverse components of velocity in the $y-z$ plane. Adapted from [60]. b) Direct numerical simulation of a square duct flow, showing transverse velocity (arrows) and streamwise velocity contours. Adapted from [51].

contributions of the two terms has not been fully established, but the net effect is the generation of two counter-rotating vortices in the corner region [60, 61]. In the absence of pressure gradients or bulk transverse flows, this vortex pair is symmetric about the corner bisector, as shown in figure 2.14a.

The transverse velocity of the vortices is directed along the corner bisector, from the core flow into the corner. This corresponds to an entrainment of the core flow into this region, and so the momentum of the corner boundary layer is increased. Indeed, the distortion of the edge of the corner boundary layer is evident from direct numerical simulations of the flow, shown in figure 2.14b.

The ideas of vorticity production in corner geometries can also be applied to high-speed boundary layers. Compressibility effects only tend to change the thickness of flat-plate boundary layers, but do not generally affect their structure [62, 63]. This concept also extends to corner boundary layers, as Modesti *et al.* showed using direct numerical simulations of compressible flow in a square duct [64]. Therefore, the vorticity production mechanisms for incompressible flow also apply to the corner boundary layers of supersonic channel flow.

2.3.2 Numerical simulation of corner boundary layers

Section 2.2 concluded that successful prediction of flow fields with shock-induced separation in a channel requires accurate calculation of the momentum in the corner boundary layer. The preceding section has shown that this corner momentum is strongly affected by the vortex pair which exists in this region. Therefore, the ability of numerical methods to capture corner vortices is a major factor in accurately computing supersonic channel flows which feature shocks.

Modesti *et al.* performed direct numerical simulations of a square duct flow at Mach 1.5 and Mach 3 [64]. As one might expect, these high-fidelity simulations do indeed capture the counter-rotating vortex pair in the corner region. However, this type of calculation is too computationally expensive to use in industrial applications, which have high Reynolds numbers and complex geometries. Instead, simulations of this type are usually restricted to the much cheaper approach provided by Reynolds-averaged Navier–Stokes (RANS) methods. It is therefore somewhat problematic that simple RANS computations do not tend to capture the vortices in corner boundary layers.

This poor prediction of corner flows is due to most simple RANS approaches using linear eddy-viscosity models (LEVMs), which are based on the Boussinesq assumption [65]. This imposes a linear relationship between the Reynolds stresses and mean velocity gradients, as defined by:

$$\tilde{\tau}_{ij} = 2\mu_t S^*_{ij} + \frac{2}{3}\rho k \delta_{ij} \,, \tag{2.2}$$

where $\tilde{\tau}_{ij} = -\rho\overline{u'_i u'_j}$ is the estimated Reynolds stress tensor and μ_t is the eddy viscosity. In this equation, the traceless stress tensor $S^*_{ij} = S_{ij} - \frac{1}{3}\frac{\partial u_k}{\partial x_k}\delta_{ij}$, with δ_{ij} denoting the Kronecker delta, and $S_{ij} = \frac{1}{2}\left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i}\right)$. The mean velocity components $u_i = \{u, v, w\}$ correspond to the spatial coordinates $x_i = \{x, y, z\}$ in the streamwise, wall-normal and spanwise directions, respectively. The density is given by ρ and k corresponds to the turbulent kinetic energy. Note that the second term, $\frac{2}{3}\rho k\delta_{ij}$, is neglected in one-equation turbulence models. These have instead:

$$\tilde{\tau}_{ij} = 2\mu_t S^*_{ij} \,. \tag{2.3}$$

Evaluation of the Reynolds stresses for these linear eddy-viscosity models (equations 2.2 and 2.3) gives $\tau_{yz} = 0$ and $\tau_{yy} = \tau_{zz}$. In turn, it is possible to calculate the production terms, C and D, from the streamwise vorticity equation (equation 2.1):

$$\text{term C}: \quad \frac{1}{\rho}\left(\frac{\partial^2}{\partial y^2} - \frac{\partial^2}{\partial z^2}\right)\tau_{yz} = 0 \,, \tag{2.4}$$

$$\text{term D}: \quad \frac{1}{\rho}\frac{\partial^2}{\partial y \partial z}\left(\tau_{zz} - \tau_{yy}\right) = 0 \,. \tag{2.5}$$

Since the vorticity production terms equate to zero, RANS simulations based on LEVMs are unable to generate the quasi-streamwise vortices which exist in corners, resulting in a poor prediction of these flows.

Therefore, more sophisticated Reynolds stress closures are required in order to improve the predictive capabilities of corner flows. There are a range of approaches, such as algebraic Reynolds stress models [66, 67] and differential Reynolds stress models [68]. However, these methods require additional assumptions to enable closure, which limits their scope of application, and they involve higher-order correlations, which are difficult to measure. Furthermore, the non-linearity introduced by

these methods can result in a lack of numerical robustness, especially when dealing with complex flows [59].

A promising approach to this problem is the use of non-linear eddy-viscosity models. These are typically simple, cheap extensions to LEVMs which better predict the anisotropy of Reynolds stresses. Gatski devised a family of these non-linear models by expressing the turbulent stress matrix using a basis of various combinations of the stress and strain tensors [69]. One model from this family was developed by Spalart which, for only a mild increase in complexity, has shown some success in predicting corner flows [70]. The quadratic constitutive relation (QCR) considers terms that are quadratic in the mean vorticity and strain tensors, when evaluating the Reynolds stress tensor.

These extensions to the Boussinesq assumption still relate the eddy viscosity with the turbulent stresses using only properties of the mean flow, rather than any variables specific to the turbulence model in use. There have been two iterations of the quadratic constitutive relation, developed by Spalart. The original statement of the relation, QCR-2000 [70], applies a modification to the Boussinesq assumption, and is defined by:

$$\tilde{\tau}_{ij} = 2\mu_t S^*_{ij} - \frac{4c_{cr1}\mu_t}{\sqrt{S_{kl}S_{kl} + \Omega_{kl}\Omega_{kl}}} \left[\Omega_{ik}S^*_{kj} - S^*_{ik}\Omega_{kj}\right] , \tag{2.6}$$

where the strain tensor $\Omega_{ij} = \frac{1}{2}\left(\frac{\partial u_i}{\partial x_j} - \frac{\partial u_j}{\partial x_i}\right)$. Note that since this modification uses only mean-flow properties, it can be applied to any single-equation or multi-equation turbulence model that is based on a linear eddy-viscosity model. The additional term $\frac{2}{3}\rho k\delta_{ij}$ can be included for turbulence models which provide the turbulent kinetic energy, k.

A more recent version of this relation, QCR-2013 [71], was proposed for turbulence models which do not provide k and includes an additional term:

$$\begin{aligned}\tilde{\tau}_{ij} =& 2\mu_t S^*_{ij} - \frac{4c_{cr1}\mu_t}{\sqrt{S_{kl}S_{kl} + \Omega_{kl}\Omega_{kl}}} \left[\Omega_{ik}S^*_{kj} - S^*_{ik}\Omega_{kj}\right] \\ & - c_{cr2}\mu_t \left[\sqrt{2S^*_{kl}S^*_{kl}}\right] \delta_{ij} .\end{aligned} \tag{2.7}$$

This additional term $c_{cr2}\mu_t \left[\sqrt{2S^*_{kl}S^*_{kl}}\right] \delta_{ij}$ approximately accounts for the $\frac{2}{3}\rho k\delta_{ij}$ term in the Boussinesq equation (equation 2.2), which had been omitted in equations 2.3 and 2.6. The constants c_{cr1} and c_{cr2} were proposed as 0.3 and 2.5, respectively, after calibration in the outer part of an equilibrium turbulent boundary layer [71].

The use of QCR results in improved prediction of the counter-rotating vortex pair in corner flows [71, 72]. The improvement in flow prediction results in a better estimate of corner separation [73], including shock-induced separation in transonic and supersonic flows [74].

Whilst QCR does appear to improve the prediction of corner flow, even these computations are limited in their accuracy. The extent to which the the corner vortices in simulation reflect those in physical flows has not yet been established. Furthermore, when the constant c_{cr1} (which, in practice,

is treated like a tuning parameter) exceeds the recommended value of 0.3, the appearance of additional non-physical vortices has been noted by Leger *et al.* [75].

2.3.3 Corner boundary layers in supersonic flow

It is therefore necessary to better understand these limitations and, indeed, develop deeper insight into the flow physics involved. This requires a comparison of the results from RANS and LES simulations [77, 78] with experimental data. However, apart from a few isolated studies [79, 13, 80], there is a distinct lack of validation-quality measurements available for this purpose. Addressing this absence of data is particularly important since, without a good quantitative understanding of the structure of corner boundary layers, predictions of separation due to an adverse pressure gradient will inevitably be inaccurate.

The problem of having insufficient high-quality experimental data of supersonic corner flows is made worse by the fact that much of this data appears to be somewhat inconsistent. Whilst all measurements of the corner boundary layer do identify quasi-streamwise vortices, the reported configuration of these vortices are not always compatible between studies.

Figure 2.15a shows measurements, performed at the University of Washington by Davis *et al.*. Pressure measurements are used to study the corner region of a square duct suspended within a Mach 3.9 tunnel core flow (figure 2.16a) [12, 76, 81]. The study was focused on the boundary layers growing along the internal duct walls. The vortex pair shown in figure 2.15a is symmetric about the corner bisector, analogous to the subsonic case [60].

An independent survey of the corner boundary layer was conducted by Rice *et al.* in the Mach 2 flow of the high speed wind tunnel at the University of Tennessee Space Institute. This supersonic flow is generated using a symmetric 'full' converging-diverging nozzle setup, where the two contoured nozzle surfaces are symmetric about the tunnel centre height (figure 2.16b). These particle image velocimetry measurements, shown in figure 2.15b, suggest that the corners of this channel contain only a single primary vortex [80, 14]. This corresponds to a fundamentally different flow structure to that observed by Davis or in subsonic flow.

Yet another corner flow topology has been observed in a Mach 2.75 flow at the University of Michigan's Glass Wind Tunnel facility. Here, a 'half' nozzle configuration, which features a curved ceiling nozzle surface and a straight horizontal floor, is used to generate the supersonic flow (figure 2.16c). The corner region of this channel was interrogated using particle image velocimetry (figure 2.15c), and shows a vortex pair sitting on the tunnel floor [13]. These vortices are not symmetric about the corner bisector, and so this constitutes a third class of flow structure.

The three separate vortex structures in figure 2.15 implies that there is no single 'typical' corner flow in supersonic channels. Moreover, it is not clear which factors influence the flow topology in this region. It has been suggested that the flow features can be related to characteristics of the nozzle geometry [13, 80] or the aspect ratio of the tunnel [13], but no coherent mechanism has been proposed to explain them. Furthermore, the distinct corner flow structures have not been directly compared with

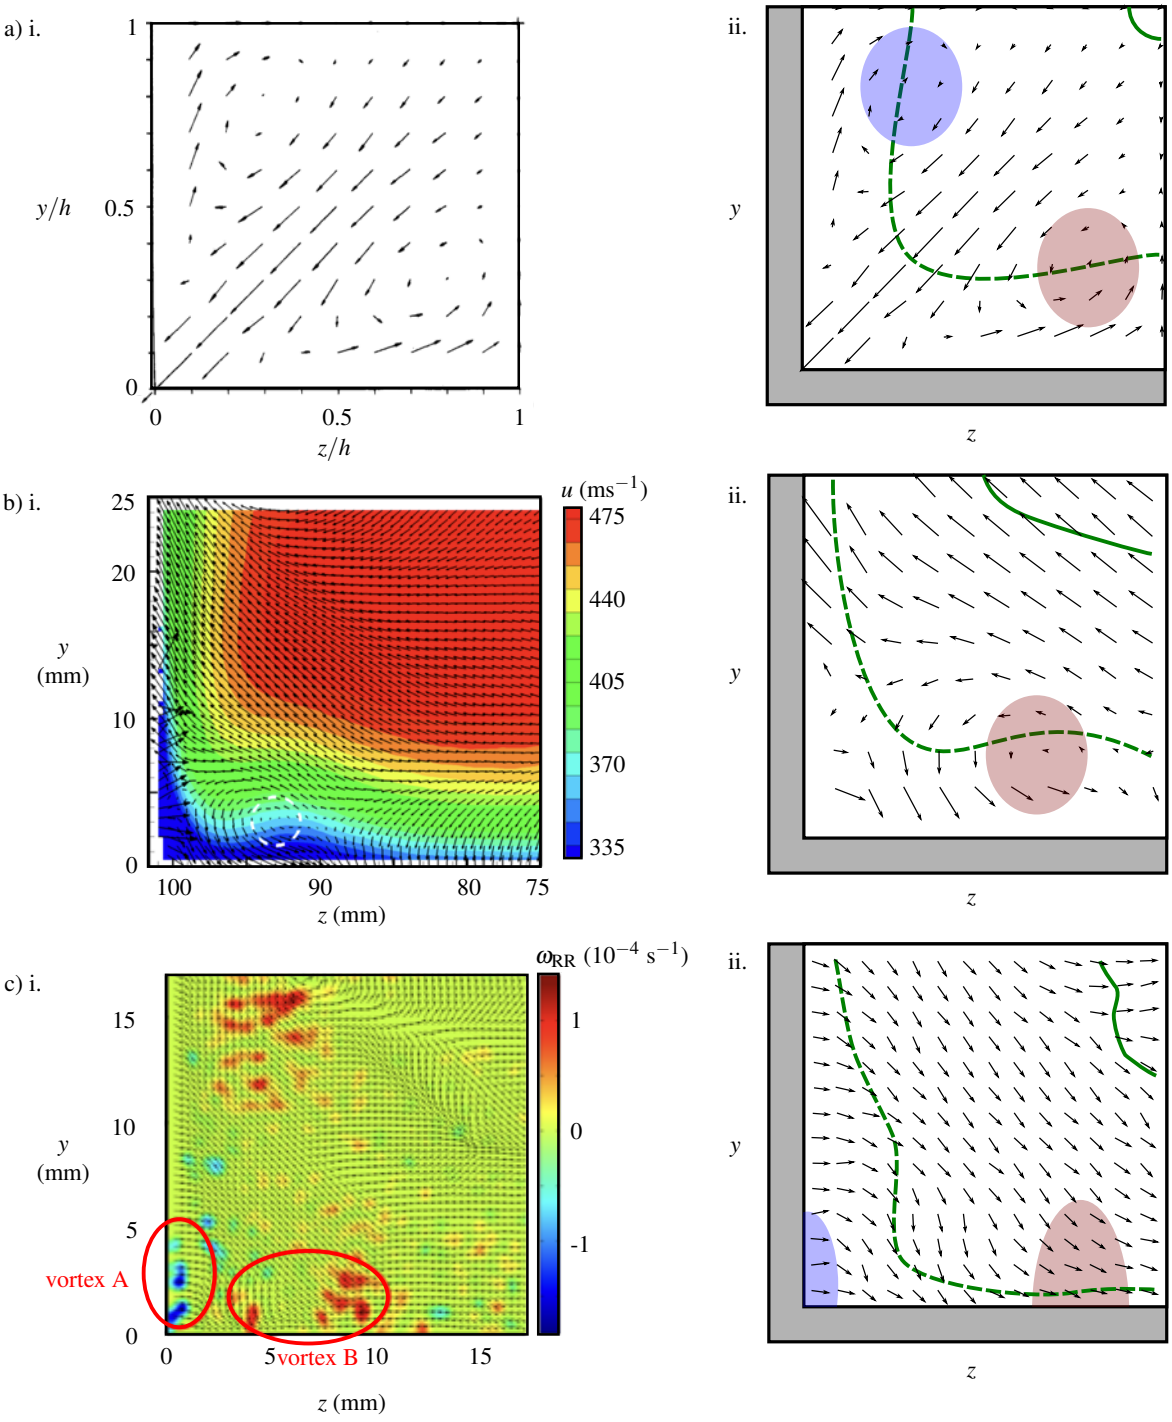

Figure 2.15: Measured corner flows from the literature: a) Davis and Gessner (1989) [76] in a square duct; b) Peltier *et al.* (2018) [14] with a symmetric nozzle; c) Morajkar *et al.* (2015) [13] using an asymmetric nozzle setup. For each study, the original reported flow field is presented, along with an adapted schematic flow field. The schematic flow fields are scaled such that the corner boundary layers each cover approximately the same area. The solid green line shows the boundary-layer edge, the dashed green line is a representative contour within the boundary layer, and red and blue areas correspond to suggested regions of negative and positive vorticity, respectively.

Figure 2.15a: Reprinted by permission of the American Institute of Aeronautics and Astronautics, Inc. from D.O. Davis and F.B. Gessner. Further experiments on supersonic turbulent flow development in a square duct. AIAA Journal, 27(8):1023–1030, 1989. Permission conveyed through Copyright Clearance Center, Inc.

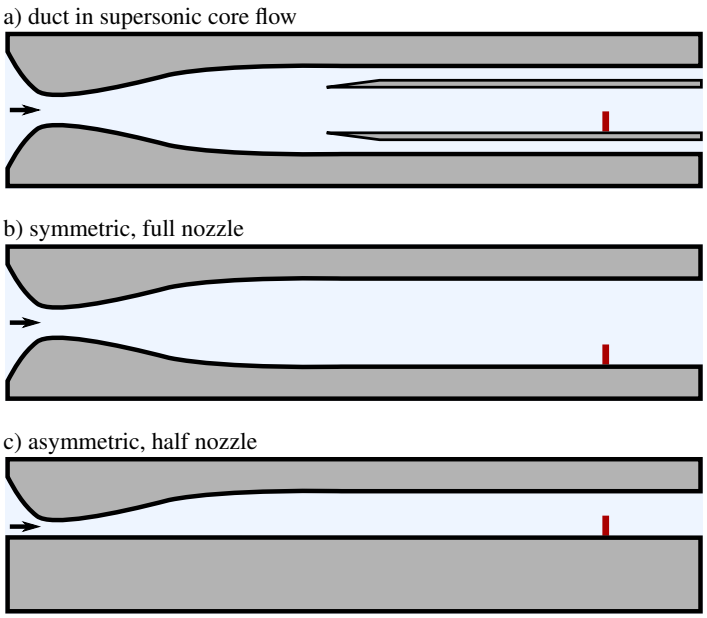

Figure 2.16: Schematic nozzle configurations used for the measured corner flows in figure 2.15: a) duct suspended in supersonic core flow, Davis and Gessner (1989) [76]; b) symmetric, full nozzle setup, Peltier *et al.* (2018) [14]; c) asymmetric, half nozzle setup, Morajkar *et al.* (2015) [13]. The red line is indicative of where the corner flow measurements were performed.

one another till now and the reasons for the differences between them are not well understood. This is compounded by an inability to easily isolate possible causes for the observed variations since previous investigations were performed under a variety of different conditions (facilities, Mach numbers, nozzle configurations, tunnel aspect ratios, etc.).

The apparent inconsistencies between the structures of experimentally-measured corner flows is particularly problematic for the validation of the quadratic constitutive relation or, indeed, any other numerical models which aim to better predict the flow in these regions. If simulations compute the flow in a rectangular channel geometry, it is not clear which of the three flow structures presented in figure 2.15 should be considered as a benchmark for comparison with computational data. Moreover, even if a numerical method correctly predicts one of the three vortex topologies, the 'missing flow physics' preventing accurate computation of the other two flow structures would remain unknown. As a result, the numerical method could not be applied with confidence to a new scenario with unknown vortex topology. Understanding the physical reasons for the differences between the corner boundary layers measured in various experiments is therefore of paramount importance for the validation of relevant numerical models.

2.4 Summary

A review of the current literature presented in this chapter shows that the two-dimensional response of a boundary layer to an incident shock is well understood. However, two-dimensional flow fields of this type do not generally exist in physical scenarios. The simplest extension to the two-dimensional case is, perhaps, the supersonic channel flow. This kind of geometry features in the inlets of supersonic aircraft, and in many high-speed wind tunnels where shock–boundary-layer interactions are studied.

In order to design aircraft inlets or to correctly interpret wind tunnel experiments, it is therefore necessary to understand the behaviour of shock–boundary-layer interactions in channels. These sidewall effects have been extensively studied in transonic and, more recently, supersonic flows. The low-momentum in the corner boundary layers of the channel have a profound influence on the flow field. The dominant mechanism of this effect is based on the corner waves, which are generated by the displacement effect of the separated corner regions and which influence the pressure profile elsewhere [24, 6].

A comprehensive set of experiments performed by Xiang and Babinsky illustrate the effect of the corner regions on an oblique shock reflection in a Mach 2.5 channel flow [6]. The corner separation point is artificially varied using a combination of obstacles, vortex generators and suction at the junction between the floor and sidewall. An analysis of the floor boundary layer in each case shows that the separation flow field depends on where the corner waves intersect at the tunnel centre line.

In order to accurately predict the point of corner separation, it is necessary to understand the momentum within the corner boundary layer. This requires knowledge of the flow structure and, in particular, the streamwise vortices which exist in this region. It is therefore particularly problematic that the few experimental studies on the corner regions of supersonic channel flows have reported quite inconsistent vortex topologies. Moreover, the possible causes for the observed variations have not, as yet, been identified.

This acute gap in our knowledge is exacerbated by the inability of numerical methods commonly used in industry to accurately predict these corner flows. The quadratic constitutive relation modification to simple RANS methods has shown some promise in this field, but its capabilities and limitations remain unknown. Importantly, the validation of these methods is inhibited by the aforementioned scarcity of relevant experimental data on supersonic corner flows.

2.4.1 Research objectives

The survey of the literature shows that there is a pressing need for a high-quality experimental study, focused on the corner boundary layers of supersonic channel flows. This provides an opportunity to understand the inconsistencies in previous studies and enables validation of relevant numerical methods. This study therefore aims to improve the physical understanding of supersonic corner flows. To achieve this aim, the following questions need to be addressed:

- what determines the flow structure in the corner boundary layers of supersonic channel flows – can the different vortex topologies in the literature be explained?
- how do differences in the structure of the corner boundary layer influence shock-induced corner separation?
- how do differences in the corner boundary layer affect separation elsewhere, and are these effects consistent with mechanisms in the literature?
- what are the capabilities and limitations of modelling methods, such as the quadratic constitutive relation?

In order to answer these questions, the corner regions of the supersonic wind tunnel at Cambridge University Engineering Department will be interrogated. This tunnel has the option of operating in both full and half nozzle configurations, analogous to the arrangements in figures 2.16b and 2.16c, respectively. This provides a rare opportunity to test the influence of nozzle configuration on corner boundary layers in a single facility.

These experiments will be conducted in close collaboration with RANS computations. A detailed characterisation process will enable the simulations to capture the real flow in the physical wind tunnel. By modelling the entire wind tunnel, including the nozzle region, it is hoped to capture the mechanisms governing the differences in corner flow structure shown in figure 2.15. As a result, computational methods, such as the quadratic constitutive relation, can be critically tested. Furthermore, the relevant flow physics can be better understood by analysing both experimental and computational data simultaneously, since simulations are able to provide information which cannot be accessed by experiments alone.

By gaining deeper insight into the corner boundary layers, it is hoped that shock–boundary-layer interactions in supersonic channel flows will be better understood. This might enable more appropriate interpretation of wind tunnel data, alongside more informed design of aircraft inlets. Ideally, this work will also help to improve the numerical modelling capability for corner flows and, in turn, flow fields featuring shock waves in supersonic channels.

Chapter 3

Research methodology

This chapter presents the methods employed in order to investigate the supersonic corner boundary layers of the Cambridge wind tunnel. The experimental facility is described in detail and the techniques used to interrogate the flow are introduced, along with an assessment of the associated uncertainties. The chapter ends by summarising RANS computations of the wind tunnel, performed by the US Air Force Research Laboratory, whose data are presented throughout this thesis.

3.1 Experimental facility

Experiments are performed in Supersonic Wind Tunnel No. 1 at the Cambridge University Engineering Department. This is an intermittent blow-down wind tunnel, driven from a high-pressure reservoir of dried air by a manually operated valve. This sets the stagnation pressure to an accuracy of 0.5%. The stagnation pressure, chosen to avoid tunnel unstart, is 308 ± 1 kPa for empty tunnel tests. Upon entering the settling chamber, shown in figure 3.1, the flow passes through a series of eight flow-straightening and turbulence grids.

The operating stagnation temperature is measured using a set of four T-type thermocouples located in the settling chamber. These thermocouples do not tend to produce consistent temperature values due to thermal stratification within the settling chamber, a phenomenon typical of facilities with a high contraction ratio [83]. In addition, the stagnation temperature tends to increase by about 2 – 3 K over the course of a wind tunnel run, and can vary slightly from one day to another. Due to these various factors, the measured stagnation temperature does not take a single value but instead lies in a band from 280 – 290 K. Therefore, the nominal stagnation temperature is quoted as 285 ± 5 K.

The flow is accelerated to high subsonic speeds by a 37:1 contraction, which also incorporates a round-to-rectangular geometry transition. Two-dimensional nozzle blocks then further accelerate the flow to the required Mach number. Depending on which nozzle blocks are installed, Mach numbers in the range 0.7 to 3.5 can be produced in this facility. For the current study, the nominal freestream Mach number is fixed at $M_\infty = 2.5$. The structure of the corner boundary layer is thought to be affected by whether the floor boundary layer develops along a straight wall from the contraction

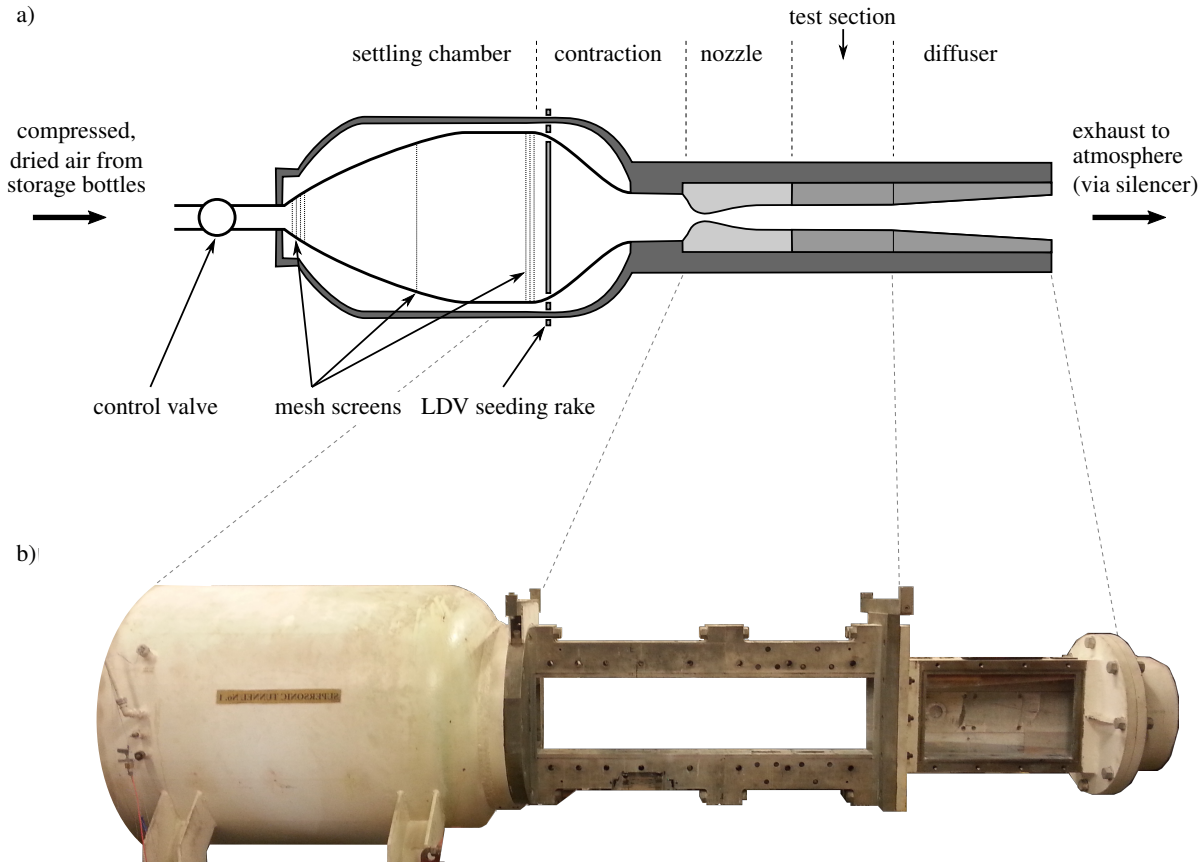

Figure 3.1: The Supersonic Wind Tunnel No. 1 facility at the Cambridge University Engineering Department: a) schematic of wind tunnel; b) photograph, adapted from [82]. The light grey blocks in the schematic represent liners which can be interchanged; the setup shown is the empty wind tunnel with a full nozzle configuration.

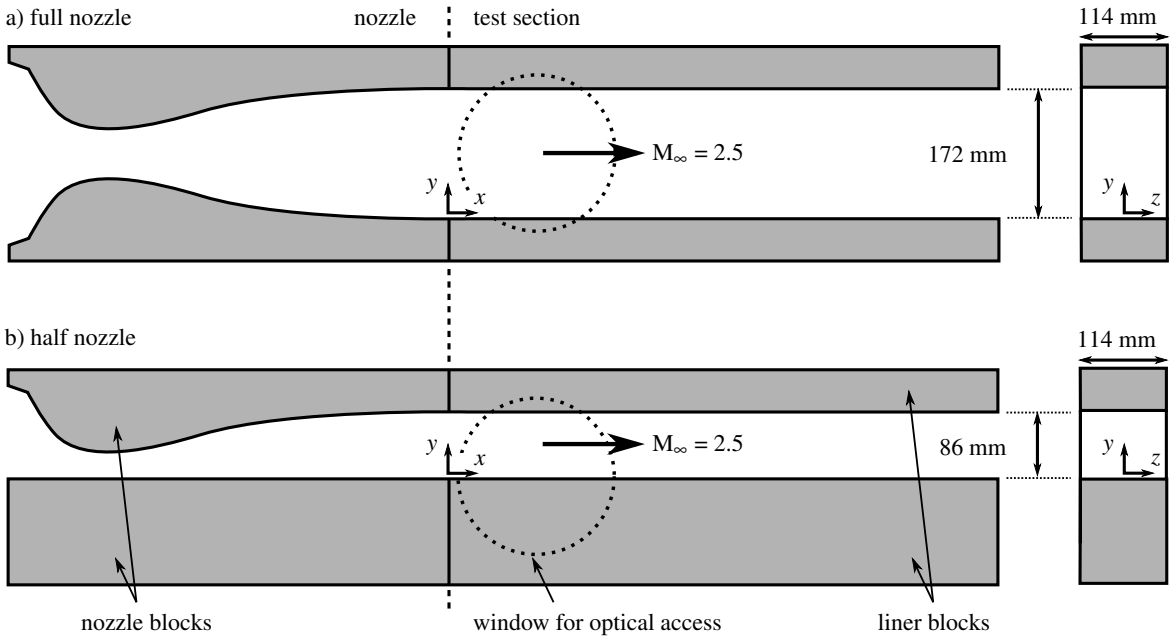

Figure 3.2: Wind tunnel setup for a) the full and b) the half nozzle configurations. The dashed circle corresponds to a window in the tunnel sidewall, which provides optical access.

or whether it follows the curve of a contoured surface. Therefore, the supersonic flow is generated using both the full (symmetric) and half (asymmetric) nozzle configurations, which are depicted in figure 3.2. This offers the particular benefit for this study that two corner boundary layers with quite different 'histories' can be compared within the same facility.

The test section has a rectangular cross-section with width 114 mm. The height of the channel is 172 mm for the full setup and 86 mm mm for the half setup. A particular focus of this investigation is the turbulent, naturally-grown boundary layers on the tunnel surfaces within the test section. These are approximately 7 mm to 8 mm thick, and have a Reynolds number based on incompressible displacement thickness of around $\mathrm{Re}_{\delta_i^*} = 30,000$. The floor and the ceiling are designed with a small divergence angle of 0.2° from horizontal, in order to provide boundary-layer relief. This ensures that, as the boundary layers on these surfaces grow naturally, the effective area for the core flow remains constant, reducing any Mach number variations in the streamwise direction.

The sidewalls of the test section take the form of removable doors, which are bolted onto the tunnel structure. These sidewalls incorporate circular windows of 203 mm diameter, which provide optical access. The static pressure in the tunnel's test section is below atmospheric pressure. Seal strips are used between the tunnel structure and sidewalls to prevent air from the surrounding environment being sucked into the tunnel due to the difference in pressures. This assists in maintaining good flow quality within the tunnel.

The coordinate system in use is shown in figure 3.2: x represents the streamwise direction, as measured from the end of the nozzle; y indicates the floor-normal direction, with $y = 0$ mm set at

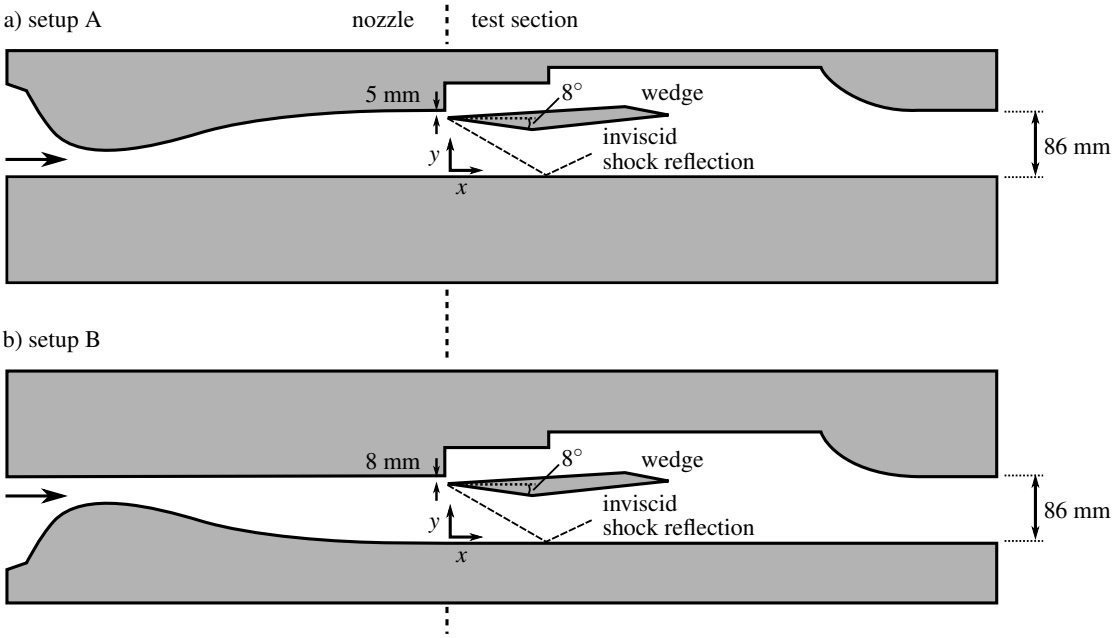

Figure 3.3: Wind tunnel setup for the two configurations with an oblique shock reflection: a) setup A, and b) setup B.

the tunnel floor; z is the spanwise coordinate measured from the centre span, such that $z = \pm 57$ mm correspond to the tunnel sidewalls.

3.1.1 Setup to investigate the oblique shock reflection

Figure 3.3 shows the setup used to test the response of the floor boundary layer to an oblique shock reflection, for distinct corner boundary layers. Since the corner regions are thought to be influenced by the nozzle geometry, the Mach 2.5 flow is generated using two different types of asymmetric half nozzle setup. Setup A corresponds to a 'conventional' half nozzle configuration, with a contoured ceiling and flat floor. On the other hand, this nozzle is inverted vertically to produce setup B. For these tests a higher stagnation pressure of 380 ± 1 kPa is chosen to prevent tunnel unstart.

A full-span wedge on the tunnel ceiling imposes a flow deflection angle of 8°, and thus generates an oblique shock. This choice of flow deflection angle corresponds to substantial separation of the floor boundary layer. Importantly, however, the sidewall boundary layer sees two distinct pressure rises from the incident and reflected shocks, each with approximately half the pressure rise experienced by the floor boundary layer. The chosen incident wave angle enables the sidewall boundary layer to remain attached, which simplifies the flow field to be investigated.

In order to avoid tunnel unstart, the wedge is retracted during startup and is then deployed once supersonic flow has been established. The wedge is placed below the upper surface of the tunnel; this

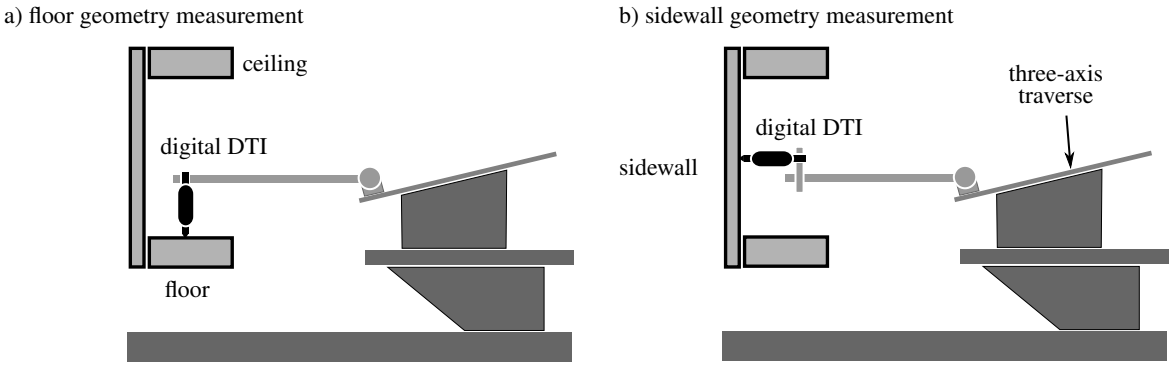

Figure 3.4: Schematic setup for measuring the tunnel geometry using a dial test indicator, configured to measure a) the tunnel floor, and b) the tunnel sidewall. To scan these surfaces, the three-axis traverse moves in the x-direction, i.e. normal to the plane of the page.

enables the ceiling boundary layer to disappear into the gap and ensures that the shock is generated in clean flow.[1]

The oblique shock wave then impinges on the turbulent boundary layer growing along the tunnel floor. At the inviscid shock reflection location, this floor boundary layer is approximately 6 mm thick, and has a Reynolds number based on incompressible displacement thickness of around $\mathrm{Re}_{\delta_i^*} = 24,000$.

3.2 Experimental techniques

3.2.1 Tunnel geometry measurement

It is important for the tunnel geometry to be accurately specified in a validation experiment, so that computational model best represents the physical tunnel. For this purpose, the mesh for CFD simulations is generated using a CAD model of the complete flow path from the settling chamber to the diffuser.

However, it is also important to accurately measure the as-installed, as opposed to design, geometry. This process enables simulations to accurately represent the physical wind tunnel. Conducting these measurements is particularly important given that the tunnel liners are removed and reinstalled between each set of experiments, and so there may be variations between different installations.

The precise, as-installed, tunnel geometry is measured by securing a dial test indicator to a three-axis traverse, and scanning along the target surface. This is performed for both the tunnel floor (figure 3.4a) and the sidewall (figure 3.4b), so that the profile of these surfaces can be mapped to within 0.02 mm.

Note, however, that there is also measurement uncertainty introduced by any angular misalignment between the traverse mechanism and the wind tunnel itself (which defines the co-ordinate system). This angular alignment is typically accurate to approximately 0.1°.

[1] Note that setups A and B have two different gap sizes, as explained further in section 7.1.

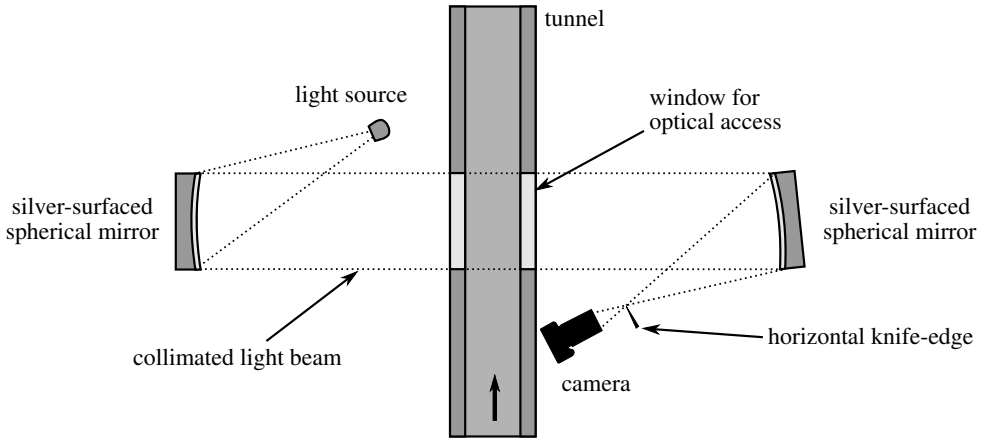

Figure 3.5: Schematic plan view of setup for flow visualisation using z-type schlieren imaging. Adapted from [84].

3.2.2 Flow visualisation

Schlieren imaging

Schlieren photography is an extensively used and well-established method to identify features of high-speed flows. It is based on the physical principle that the refractive index of a gas depends upon its density. Regions where there is a variation in density, such as shock waves or boundary layers, refract a collimated light beam, rendering them visible with this technique [85].

For the current study, a two-mirror z-type optical setup is used for schlieren imaging, as depicted in figure 3.5. A collimated light beam is produced using a 4 W light source, located at the focus of a spherical concave mirror. This beam then passes through the test section, where it may encounter flow with density gradients. Here, the associated variations in refractive index cause the light rays to bend towards regions of higher density. These light rays are then brought to a focus by a second spherical concave mirror.

A camera placed at this focal point would record a 'shadowgraph'. Such an image has light and dark regions corresponding to spanwise-averaged variations in density, and can be shown to be sensitive to the second derivative of density, $\nabla^2\rho$.

By placing a knife-edge at the focal point ahead of the camera lens, part of the light is removed. This increases the sensitivity of the resulting 'schlieren' image to density gradients. More specifically, this image is a measure of the first derivative of density. A horizontal knife-edge is most commonly used in this investigation, and so the resulting images visualise the spanwise-integrated vertical density gradient, $\frac{\partial \rho}{\partial y}$. This enables identification of the floor and ceiling boundary layers, as well as any oblique or streamwise-running flow feature, such as waves.

Two different cameras are used, depending on the specific application. When a short exposure time is not required, a Nikon D7000 DSLR camera is used to capture high resolution images. For high-speed applications, however, schlieren visualisation is performed using a Photron Fastcam Nova

S6 camera, which features shutter speeds as fast as 0.2 μs [86]. However, the limited intensity of the point light-source used for this technique sets a minimum exposure time of 1.1 μs.

Oil-flow visualisation

Further information about the studied flow fields can also be obtained by analysing the time-averaged skin-friction lines from surface oil-flow visualisation. This technique involves coating the target surface (e.g. the wind tunnel floor) with a thin layer of oil, formed from a powdered pigment in a carrier medium. When the wind tunnel is operating, the shear stress imposed by the flow sweeps the oil in the direction of the local skin friction. Over the course of the wind tunnel run, the carrier medium evaporates, depositing the pigment along skin-friction lines.

The oil mixture is formed by a pigment of finely-powdered titanium dioxide (TiO_2) in a paraffin carrier medium. These are combined in approximately equal amounts, and a small amount ($\approx$ one drop per 10 g of mixture) of oleic acid is added. This emulsifying agent enhances mixing of the oil, and inhibits coagulation during wind tunnel operation. Lubricating oil is also added in small quantities ($\approx$ 1 g per 10 g of mixture), in order to increase viscosity and thus inhibit smearing of skin-friction lines on tunnel shutdown [87, 6].

The composition of the oil does require fine-tuning for the specific application, depending on the available run time and the local skin friction. Stanbrook has proposed an empirical relation for the time, t, required for the oil to dry [88]:

$$t = 36{,}000 \times \frac{\mu_{\text{oil}}}{q c_f} \,, \tag{3.1}$$

with a variation about this value of approximately 30%. In equation 3.1, q denotes the dynamic pressure, c_f is the local skin friction coefficient and μ_{oil} is the viscosity of the oil. For the flow field with a shock, there is a considerable range of local skin friction. This coefficient is approximately 0.002 in the Mach 2.5 flow upstream of the interaction, but drops to zero at the point of separation on the tunnel centre line. This corresponds to a wide range of drying times for an oil mixture of given viscosity.

The intermittent blowdown wind tunnel, with the configuration required for the current study, has a maximum run time of approximately sixty seconds. It is therefore somewhat difficult to ensure that the oil in all parts of the flow dries after the desired flow condition has been established, but before tunnel shutdown. In order to achieve an appropriate drying time, an iterative procedure is used to refine the oil mixture. More paraffin or more titanium dioxide is added if the mixture dries too early or too late, respectively, and the amount of lubricating oil is increased if the pattern smears on shutdown.

The target surface is coated with gloss black paint to maximise contrast with the white titanium dioxide pigment. During the course of a wind tunnel run, the oil flow is monitored *in situ* to ensure that the oil does dry once the desired flow condition has been established, but before tunnel shutdown. The final surface oil-flow pattern is illuminated by an LED array, and photographed using a Nikon D7000 DSLR camera. The interpretation of the skin-friction lines is based on the work by Délery [89],

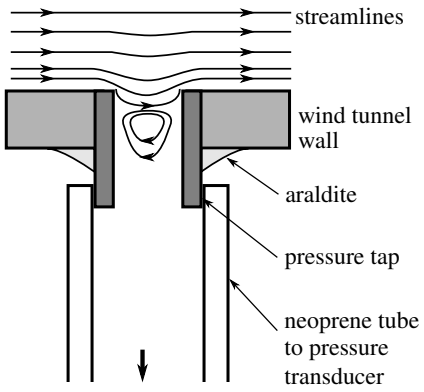

Figure 3.6: Schematic of static pressure taps in wind tunnel surface. Representative streamlines extracted from [91].

who has developed methods to infer the topology of three-dimensional separation based on surface flow patterns.

Note, however, that whilst this method is intrusive in nature, Squire has shown that the effect of the oil layer on the boundary-layer flow can generally be neglected [90]. The notable exception, relevant in the current study, is near separation. Here, the oil tends to accumulate upstream of the separation line, since the mixture cannot cross the separation boundary. Moreover, the forces in the large adverse pressure gradient associated with shock-induced separation are large enough to affect the behaviour of the oil layer. As a result, this technique exhibits a small error in determining the separation line. This uncertainty is approximately 0.2 boundary-layer thicknesses, which corresponds to 1.2 mm for the present study [90]. Despite these limitations, oil-flow visualisation allows the flow topology to be determined, and comparisons of the size of separation regions between different setups are considered to be reliable.

3.2.3 Pressure measurements

Static pressure measurements

Steady-state surface pressure measurements are undertaken using 0.15 mm diameter static pressure taps, shown schematically in figure 3.6. The cylindrical holes are drilled in the wind tunnel floor and sidewalls, allowing the pressure distributions over these surfaces to be measured. The pressure taps are connected via 1 mm diameter neoprene tubes, which are approximately one metre long, to a NetScanner 9116 pressure transducer. This 16-channel differential transducer uses atmospheric pressure as a reference. The required conversion from gauge to absolute pressure is performed using atmospheric pressure measurements from a mercury barometer. The static pressures are typically non-dimensionalised by the settling chamber stagnation pressure, p_{0s}.

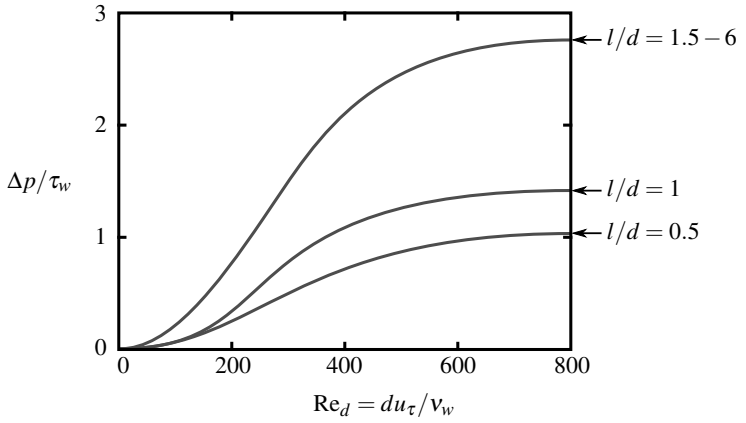

Figure 3.7: Error due to static tap geometry on pressure measurements: effect of Reynolds number based on tap diameter (Re_d), and length-diameter ratio of the tap (l/d). Adapted from [93].

There are some limitations associated with these measurements, however. A combination of the small internal diameter of the taps and the long neoprene tubes results in a long response time of around 30 ms [92]. This technique is therefore only used to take time-averaged measurements.

The tap geometry results in an overprediction of the static pressure, due to streamline deflection into the tap and possible secondary flows within it [91]. These effects are illustrated by the representative streamlines in figure 3.6. The error due to different tap geometries has been shown by Shaw [93] and Rayle [94] to depend on four key parameters: the tap edge profile; the local wall shear stress, τ_w; the length-diameter ratio of the tap; the Reynolds number based on tap diameter, Re_d. Figure 3.7 shows that, for the sharp-edged taps in use, the uncertainty reaches asymptotic limits for length-diameter ratio greater than 1.5 and Re_d larger than 800. The experimental setup in the current study satisfies both these conditions, and so the overestimate in pressure due to tap geometry is $2.8\tau_w$. This corresponds to an error of 0.5%.

The manufacturer-rated uncertainty of the NetScanner 9116 transducer is 0.05%. This estimate was verified directly, by comparison with mercury barometer pressure readings. The atmospheric pressure, used to convert from gauge to absolute pressures, is measured using a mercury barometer which typically yields errors of around 30 Pa.

The estimate of the settling chamber stagnation pressure, p_{0s}, used to normalise all pressure measurements, is obtained by using a static tap in the wall of the settling chamber. This is, therefore, really a measurement of the local static pressure. However, the flow velocity in the settling chamber is relatively low (approximately $8\,\mathrm{m\,s^{-1}}$). Therefore, the overestimate of the true stagnation pressure is on the order of 0.02% [95], and so a reasonable approximation of this quantity is obtained. The above discussion regarding transducer error and uncertainty in measuring atmospheric pressure also applies to this measurement, and so the overall uncertainty in p_{0s} is 0.3%.

The sources of uncertainty relating to tap geometry, transducer error, atmospheric pressure measurement, and stagnation pressure measurement need to be combined. This leads to an overall uncertainty in static pressure ratio, p/p_0, of around 1%.

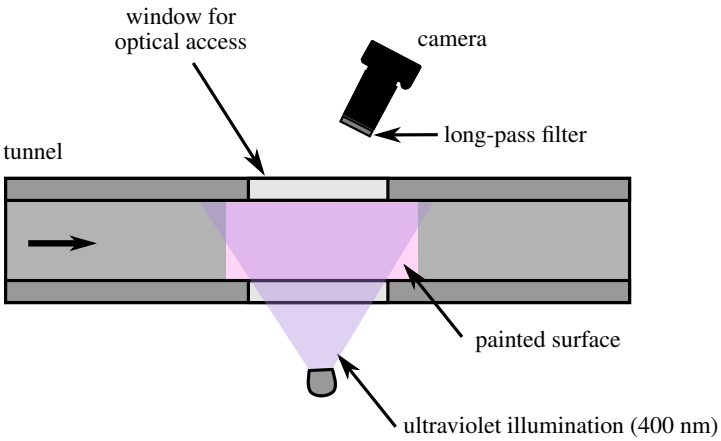

Figure 3.8: Schematic plan view of setup for pressure-sensitive paint measurements. Adapted from [84].

Pressure-sensitive paint

The spatial resolution of pressure measurements using static taps is limited by the distribution of taps and the number of transducer channels. To overcome this limitation, steady-state surface pressure measurements are performed using pressure-sensitive paint (PSP).

The surface of interest is sprayed with a special polymer binder seeded with luminescent molecules. When these molecules are irradiated by ultraviolet light, they are excited to a higher energy state. The subsequent return to ground state can take place via two separate processes: the emission of red-shifted light, or the transfer of energy to the vibrational mode of free oxygen molecules nearby. The relative frequency of the two processes depends on the partial pressure of oxygen – fewer oxygen molecules in the vicinity leads to more light emission [96, 97]. Importantly, the partial pressure of oxygen at a given point is proportional to the local pressure.

As a result, it is possible to relate the intensity of emitted light, I, to the local wall pressure, p, using the Stern–Volmer relation:

$$\frac{I_{\max}(T)}{I(p)} = 1 + K(T)p \,, \tag{3.2}$$

where K is the Stern–Volmer constant. $I_{\max}$ is the maximum possible emitted light intensity, corresponding to the case of no nearby free oxygen molecules [97]. Instead of trying to determine the temperature-dependent parameters, $I_{\max}$ and K, in the laboratory, an alternative approach is used.

Alongside measurements when the wind tunnel is operating, the intensity is also measured in a reference wind-off condition, in the absence of airflow. Using equation 3.2, the ratio of intensities between the wind-on and wind-off conditions is given by:

$$\frac{I_{\text{ref}}}{I(p,T)} = A(T) + B(T)\frac{p}{p_{\text{ref}}} \,, \tag{3.3}$$

where I_{ref} corresponds to the light intensity in the wind-off condition, and p_{ref} is atmospheric pressure (101 kPa) everywhere [97].

Therefore, the pressure distribution, p, is obtained by calculating the intensity ratio between two separate images, the wind-on and the wind-off images. Calculating this ratio on a pixel-by-pixel basis also accounts for any non-uniformities in ultraviolet illumination and any variations in the thickness of the paint layer. The constants, $A(T)$ and $B(T)$, are determined by *in situ* calibration using static pressure taps on the surface of interest. This procedure is detailed later in this section, after a description of the experimental setup.

A schematic of the setup used for pressure-sensitive paint measurements is shown in figure 3.8. An array of light-emitting diodes, at wavelength 400 nm, ensures uniform illumination of the target surface. This surface is coated with UniCoat pressure-sensitive paint, supplied by Innovative Scientific Solution Incorporated. At least five coats of paint are used to ensure sufficient intensity response [98]. The pressure in the Mach 2.5 flow field with shock-induced separation varies from 29-57 kPa. This range of pressures is sufficiently large to provide reliable measurements [99].

A Nikon D7100 DSLR camera is focused on the coated target surface to measure the emitted light. A 550 nm long-pass filter is used in front of the camera lens. This filter permits the light emitted from the paint to pass through, but removes any light with wavelengths corresponding to the incident illumination. Koike has shown that the DSLR camera features acceptable linearity and that the 16-bit intensity resolution is sufficient for these measurements [100]. A focal ratio of $f/8$ is chosen to ensure that the depth of field covers the wind tunnel span.

The choice of exposure time is slightly more involved. The regions of supersonic flow during wind tunnel operation correspond to surface pressures lower than atmospheric pressure. As a result, the light emitted from the pressure-sensitive paint is brighter than the reference wind-off case. With *a priori* knowledge of how much brighter this image is likely to be, an exposure time can be selected in the wind-off setup to maximise the range of measured intensities, without exceeding the maximum possible intensity value for any pixel.

The reference wind-off image is captured immediately prior to wind tunnel startup. Over the course of the tunnel run, ten images are captured. An average of these images is calculated in order to reduce random noise. The emission from the paint is confined to wavelengths between 500 nm and 600 nm, so the red channel of each pixel in RGB colour-space is extracted for analysis.

Equation 3.3 is used to convert the measured light intensity to pressure measurements. For this, the intensity values from wind-off image are divided by those from the wind-on image, on a pixel-by-pixel basis. It is still necessary, however, to determine the values of the constants, A and B.

This calibration procedure is performed *in situ*, using measurements from at least five static pressure taps in the target surface. This number of pressure taps is generally considered to be sufficient for robust calibration [97]. Figure 3.9a shows that the resulting calibration curve exhibits a linear relationship. This enables the values of A and B in equation 3.3 to be extracted, so that conversion

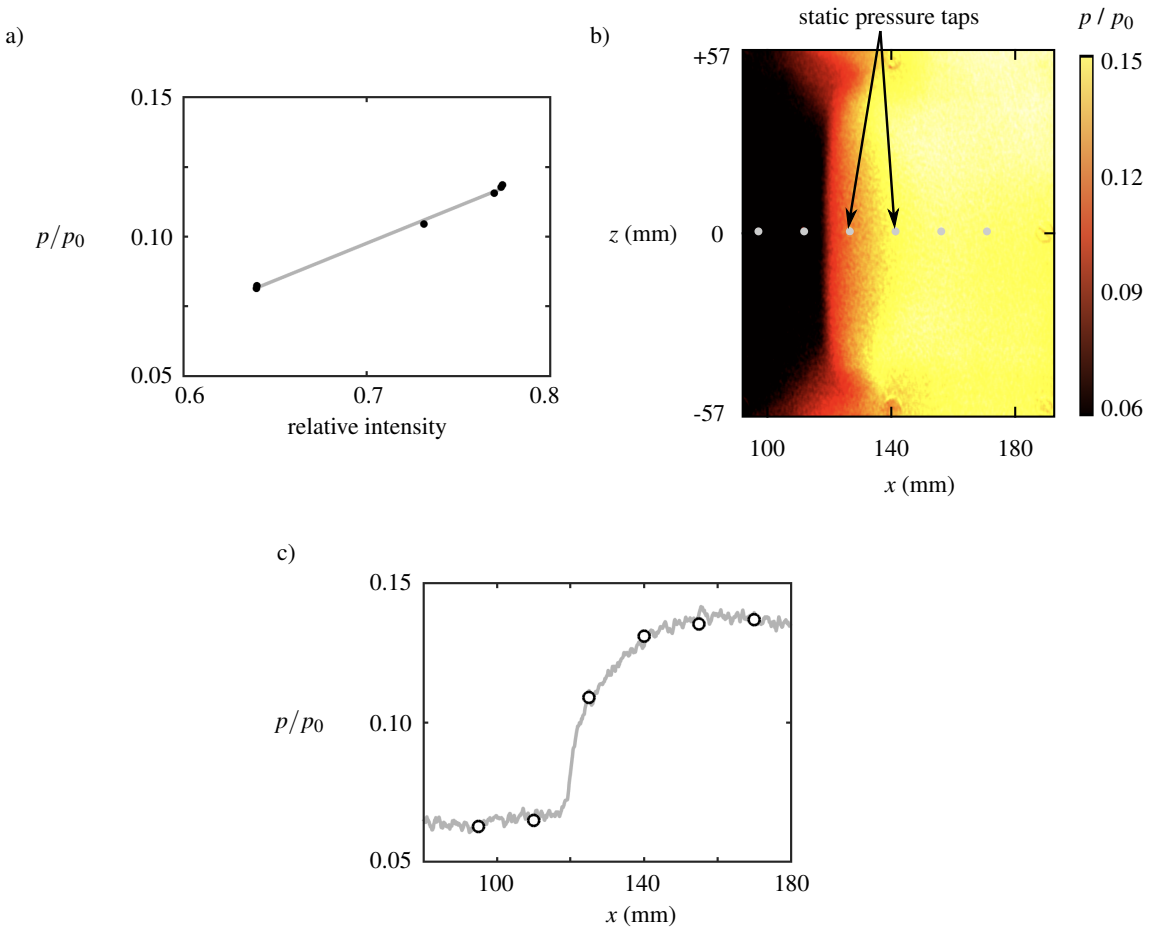

Figure 3.9: a) Calibration curve for a typical run, relating light intensity to static tap pressure; b) calibrated pressure field for this run, measured using pressure-sensitive paint; c) direct comparison between pressure-sensitive paint and static tap measurements along the tunnel centre span.

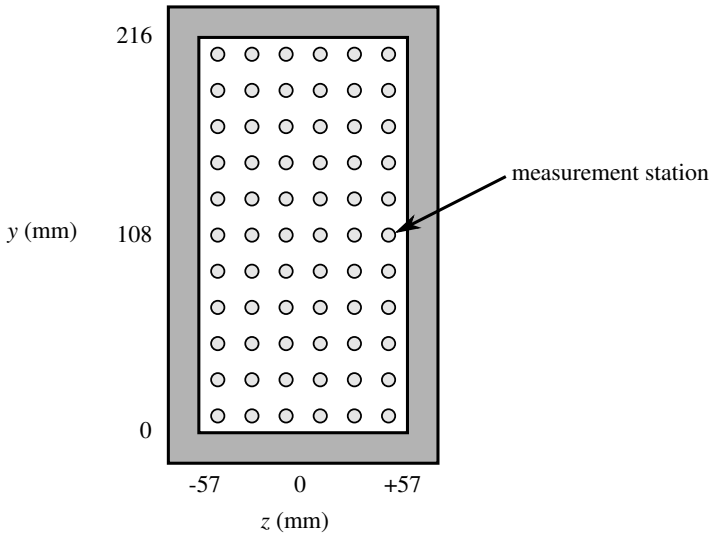

Figure 3.10: Location of measurement stations for Pitot probe measurements, immediately upstream of the nozzle. Pitot probe size not to scale.

from intensity ratio to absolute pressures is possible. The global pressure field, such as that shown in figure 3.9b, is thus obtained.

A comparison between static tap pressures and the calibrated PSP data, on the tunnel centre line, is presented in figure 3.9c. The noise in the pressure-sensitive paint data is reduced by averaging over a 10 pixel $\times$ 10 pixel region, corresponding to a square of approximate dimension 0.5 mm$\times$0.5 mm. The comparison in figure 3.9c places an error bound of 3% on these measurements.

The temperature-dependence of the constants, A and B, also needs to be considered in this blowdown facility, where run times are typically not long enough for surface temperatures to reach equilibrium. The light intensity varies considerably (approximately 10%) between wind-off images captured immediately before and after wind tunnel runs. This difference can be attributed to the wind tunnel surfaces being cooled over the course of the run. However, the intensity difference is uniform spatially, and the variations between different parts of the target surface are less than 1%.

Since the global effect of temperature variations are accounted for by *in situ* calibration and the spatial variations are small, the error due to the temperature-dependence of this technique is generally considered to be negligible when compared to the random noise. However, in regions where the thermal properties of the surface change (e.g. filler material) the calibration from figure 3.9a is no longer valid, and a much greater error is observed.

Pitot tube measurements

Stagnation pressure measurements by Pitot tubes are used to assess the inflow pressure uniformity, and to check the accuracy of direct velocity measurements in boundary layers (section 3.2.4).

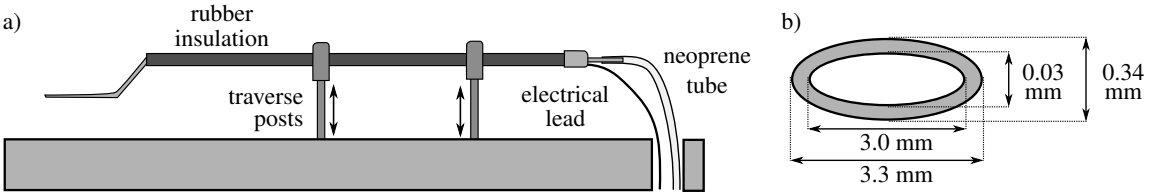

Figure 3.11: a) Schematic of side-on view of Pitot tube traverse setup. b) Dimensions for cross-section of flattened Pitot probe. Adapted from [101].

The stagnation pressure distribution at the exit of the settling chamber, upstream of the nozzle is measured using a rake. This consists of Pitot probes arranged at 20 mm intervals in a square grid across the tunnel cross-section. This grid, shown in figure 3.10, consists of six stations and eleven stations in the spanwise and vertical directions, respectively. The Pitot tubes each have internal diameter of 0.7 mm, and are connected to the NetScanner 9116 differential pressure transducer. The corresponding uncertainty has contributions from intrinsic transducer error, atmospheric pressure measurement and the settling chamber stagnation pressure. These are all discussed in relation to static tap pressure measurements, and result in an overall error of 0.7%.

In order to measure floor boundary-layer profiles within the test section, a Pitot tube is connected to a stepper-motor controlled traverse, shown schematically in figure 3.11a. The probe used for this purpose is formed of a flattened tube of internal hole dimension 3.0 mm $\times$ 0.03 mm in the spanwise and vertical directions, respectively (figure 3.11b). This probe shape enables high-resolution measurements in the wall-normal direction.

In supersonic flow, a shock is formed ahead of the Pitot tube – this is a normal shock directly upstream of the probe head. Having crossed this shock, the flow is then decelerated isentropically. The measured pressure therefore corresponds to the post-shock stagnation pressure, averaged over the corresponding open cross-sectional area of the tube. The measured pressure will be referred to as the 'Pitot pressure', denoted by p_{pit}. The Pitot pressure is generally non-dimensionalised by the settling chamber using the approach outlined for wall static tap measurements.

Within the floor boundary layer, there are significant gradients of Pitot pressure in the vertical direction. Pitot probe measurements are prone to errors caused by such gradients within the probe area [102]. This can generally be treated as a bias error. The effect of this phenomenon is reduced by using a flattened tube shape, as shown in figure 3.11b. Such a geometry minimises the vertical probe height, and any residual error can be taken into account by adjusting the effective probe position by 0.15 times the outer diameter of the tube [102]. The uncertainties due to transducer error and atmospheric pressure measurement, discussed in detail for wall static pressures, lead to a combined error of 0.3%.

However, there is a much larger contribution to the overall error from uncertainty in the wall-normal position of the probe. To avoid positional errors caused by constant probe bending under aerodynamic load, wall-normal positions are calibrated by the contact between the Pitot tube head and the tunnel floor in the wind-on condition. This contact is detected electrically. There is rubber

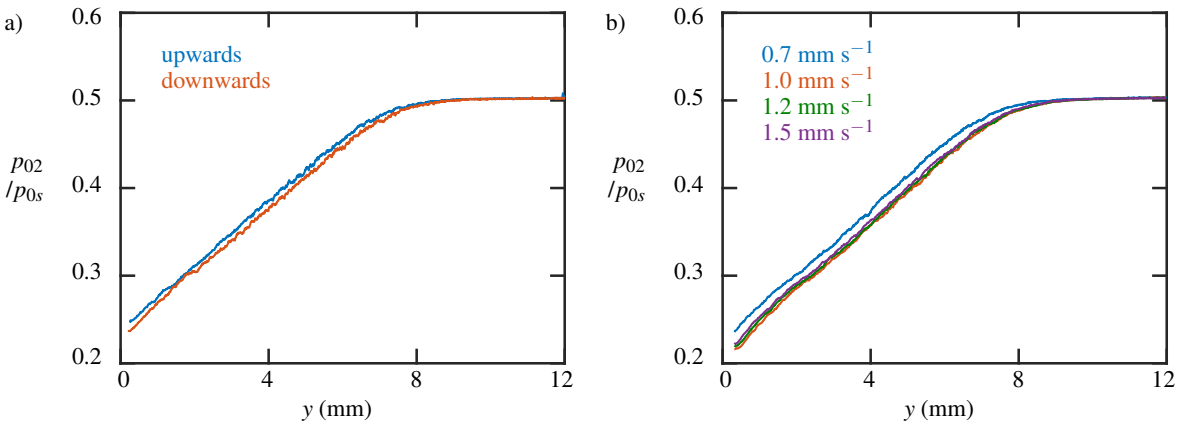

Figure 3.12: Pitot tube boundary-layer measurements carried out at $x = 60$ mm and $z = 0$ mm. These correspond to traverses performed a) both towards and away from the tunnel floor, and b) with varying traverse speed, conducted in the downwards direction.

insulation between the probe supports and the Pitot probe itself, so an electric circuit is completed when the tip of the probe makes physical contact with the tunnel floor. This acts as a relay, to stop the stepper motor as the probe tip touches the floor. A pulse counter records the number of motor steps during the traverse, and so the position where the traverse begins can be determined.

It is possible to check that probe errors due to slow response time do not affect the measurements by comparing data from traverses performed in opposite directions (figure 3.12a) and at varying speeds (figure 3.12b). The former displays the effects of temporal lag, and allows the resultant uncertainty in these measurements to be estimated as $\pm 3\%$. Meanwhile, figure 3.12b shows that at speeds lower than 0.7 mm s^{-1}, the y position is inaccurate due to the motor skipping steps, while it stalls at speeds faster than 1.5 mm s^{-1}. Thus, a suitable traverse speed of 1.2 mm s^{-1} is selected for this study.

3.2.4 Velocity measurements: Laser Doppler velocimetry

Two-component laser Doppler velocimetry (LDV) is used to measure the streamwise and floor-normal velocities, u and v, respectively. This is a non-invasive, optical technique which gives direct measurements of flow velocities with high spatial resolution.

Physical concept

The underlying concept behind laser Doppler velocimetry is the detection of light scattered by small, illuminated particles suspended within the fluid of interest. The seeding particles, which are discussed in more detail later in this section, are tiny droplets of paraffin, small enough that they faithfully represent the fluid motion [103].

The LDV system used in the current study measures two velocity components, and each of these components consists of a pair of laser beams, shown in figure 3.13a. Each beam has a Gaussian beam profile, and so they intersect in an ellipsoidal region, termed the probe volume. Within the probe

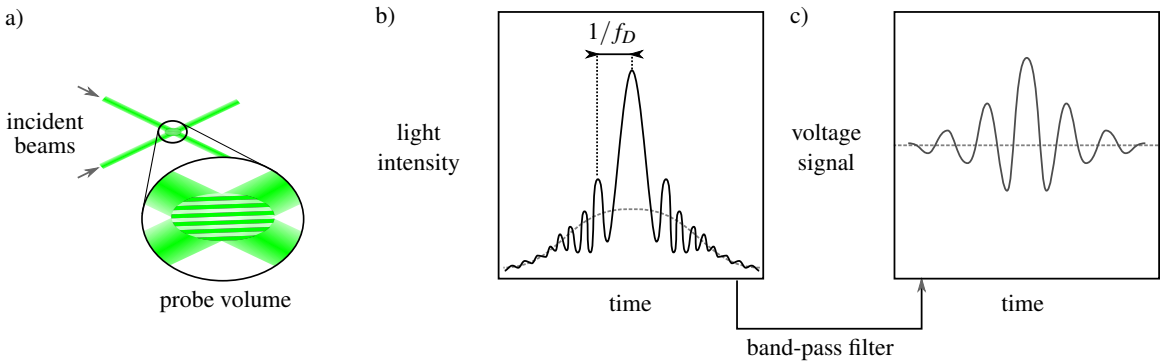

Figure 3.13: a) Intersection of incident laser beams to form an interference pattern in the probe volume. b) Scattered light intensity measured by the receiving optics, known as a Doppler burst. This consists of a predominant frequency, f_D, modulated by a Gaussian envelope (dashed). c) Sinusoidal voltage signal extracted using a band-pass filter. Adapted from [84].

volume, an interference pattern, consisting of bright and dark fringes, is formed [104]. The distance between fringes, δ_f, depends only on the laser wavelength, λ_1, and the angle between the two beams, α:

$$\delta_f = \frac{\lambda_1}{2\sin\frac{\alpha}{2}} \,. \tag{3.4}$$

As the seeding particles traverse the probe volume, they scatter light. These particles are only illuminated as they pass through the bright fringes, and so the scattered light intensity varies sinusoidally (figure 3.15b). The predominant 'burst frequency', f_D, is the inverse of the time taken for the particle to traverse one fringe spacing:

$$f_D = \frac{v}{\delta_f} \,, \tag{3.5}$$

where v is the component of velocity perpendicular to the fringe direction. Thus, a faster-moving particle produces a higher burst frequency. The resulting scattered signal, known as a Doppler burst, is detected by the receiving optics, which is focussed on the probe volume. Figure 3.15b shows that this signal contains a sinusoidal oscillation modulated by a Gaussian envelope. This envelope is an artefact of the laser beam profile, and is removed using a band-pass filter (figure 3.13c).

Note that the setup described above cannot distinguish between particles moving forward versus backward through the probe volume. In order to overcome this problem, the frequency of one of the two laser beams is shifted by $f_s = 40$ Hz, using a Bragg cell. This results in the fringe pattern to propagate through the probe volume at a constant, known velocity. The direction of the fringe propagation depends on which beam the frequency-shift is applied to. In practice, this beam is chosen such that the fringes move in an upstream direction, relative to the mean flow. Figure 3.14 shows that the resultant effect is a bias in the frequency of the Doppler burst, enabling the speed of particles travelling up to some maximum negative velocity to be unambiguously determined. With this arrangement, negative velocities correspond to burst frequencies less than f_s [104].

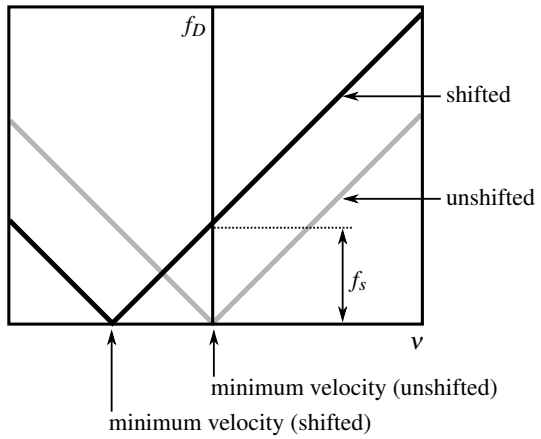

Figure 3.14: Relationship between measured velocity, v, and Doppler frequency, f_D, with and without the 40 Hz Bragg cell frequency shift (f_s). Adapted from [84].

The velocity of the tracing particle can then be determined by combining equations 3.4 and 3.5, and accounting for the frequency shift, f_s:

$$v_1 = \frac{(f_D - f_s)\,\lambda_1}{2\sin\left(\frac{\alpha}{2}\right)}. \tag{3.6}$$

The second component of velocity, v_2, is measured simultaneously using a second pair of laser beams, at a different wavelength, λ_2:

$$v_2 = \frac{(f_D - f_s)\,\lambda_2}{2\sin\left(\frac{\alpha}{2}\right)}. \tag{3.7}$$

The two pairs of laser beams are perpendicular to one another and, for the current study, aligned at $\psi = 45^\circ$ to the freestream direction. This arrangement, shown in figure 3.15, ensures that the measured components, v_1 and v_2, are similar in magnitude. Therefore, the maximum velocity in any one channel is minimised, and so reduces the fringe spacing required to capture the fastest particles (equations 3.4 and 3.6). A reduced fringe spacing corresponds to a probe volume with smaller spatial extent. This, in turn, means that velocity measurements are averaged over a smaller region and so the spatial resolution of the measurements is increased.

The mean streamwise and wall-normal velocities can then be extracted from the measured components using a simple rotation of frame:

$$u = v_1 \cos\psi + v_2 \sin\psi\,, \tag{3.8}$$

$$v = v_1 \sin\psi - v_2 \cos\psi\,. \tag{3.9}$$

Experimental setup

The laser Doppler velocimetry measurements are conducted using a Powersight diode-pumped solid-

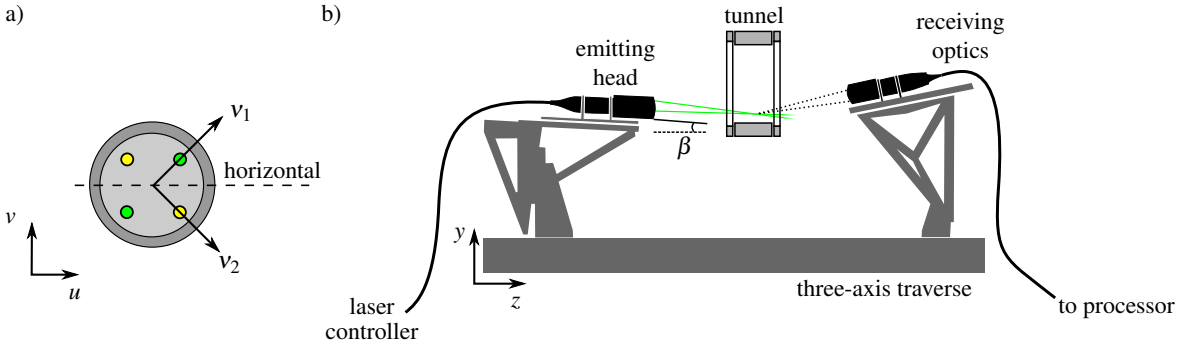

Figure 3.15: a) Arrangement of beams on face of the emitting head, showing rotation of laser beam pairs by $\psi = 45°$ from horizontal. b) Schematic side-on view of setup for laser Doppler velocimetry. Adapted from [84].

Table 3.1: Physical parameters of the laser Doppler velocimetry setup.

property	component 1	component 2
wavelength (nm)	561	532
focal length (mm)	362.6	
laser beam diameter (mm)	2.10	
fringe spacing (μm)	4.08	3.87
probe waist (μm)	123.3	117.0
probe span (mm)	1.43	1.35
band-pass filter (MHz)	$20-175$	
maximum velocity (m/s)	550	520
minimum velocity (m/s)	-80	-80

state laser system supplied by TSI [105]. Figure 3.15b shows that the emitter is inclined from the horizontal at an angle, $\beta = 8°$. This prevents the bottom pair of beams (i.e. those below the dashed line in figure 3.15a) from being clipped when the probe volume is located very close to the wind tunnel floor.

The two pairs of lasers have different wavelengths so that the receiving optics can detect the channel of any scattered light, using band pass filters. The laser pairs have wavelengths of 561 nm and 532 nm, both in the green region of the visible spectrum. These laser beams are focused onto the probe volume using a lens with focal length 363 mm.

This lens focal length and the initial beam separation (50 mm) determine the angle between incident beams, $\alpha = 7.88°$. Each laser beam has a diameter of 2.10 mm, and so the probe volume spans approximately 0.1 mm in the streamwise and vertical directions, but 1.4 mm in the spanwise direction. A more comprehensive definition of the parameters of the probe volume is provided in table 3.1.

Optical alignment is required to ensure that the four beams converge on the same measurement volume. This is verified by passing the beams through a 35 μm pin hole. When all four laser beams

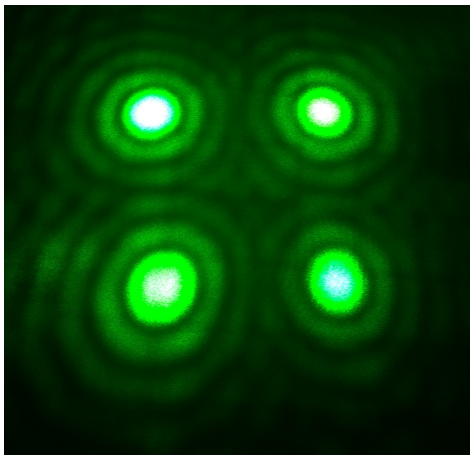

Figure 3.16: Diffraction pattern projected onto screen when all four laser beams pass through the 35 μm pinhole.

converge at the same point, and when this probe volume corresponds to the pin-hole location, the interference pattern shown in figure 3.16 is observed on a screen positioned at the far side of the pin hole.

The light scattered by seeding particles in the moving fluid is captured by a TSI RV 70 optical receiver, which is focused on the measurement volume. Figure 3.16 shows that this detector is placed on the far side of the wind tunnel, with respect to the laser emitter. Since the receiver detects light for a scattering angle (much) less than 90°, this setup is referred to as a forward-scatter configuration. This type of setup maximises the intensity of the received light, and therefore enhances the signal-to-noise ratio [106].

The Doppler burst detected by the receiving optics is converted into a voltage signal using photomultipliers in a photo-detector module (TSI PDM 1000). This signal is then analysed using a digital processor (TSI FSA 4000), which converts the dominant frequencies to a flow velocity through Fourier analysis. During this process, a minimum voltage threshold is imposed to distinguish between signals generated from seeding particles and from noise. In addition, a band pass filter restricts the frequency range of any detected signal, such that $f_D - f_s$ must lie between 20 MHz and 175 MHz. As a result only velocities between around $-80\,\mathrm{m\,s^{-1}}$ and $520\,\mathrm{m\,s^{-1}}$ can be detected.

The emitter-receiver assembly is mounted on a three-axis traverse, shown in figure 3.15b, which is controlled using stepper motors. Data can be collected continuously as the traverse is conducted. The maximum traverse speed is set to $1\,\mathrm{mm\,s^{-1}}$, in order to limit the load on the motors and to prevent vibration of the apparatus.

The location of the probe volume relative to the wind tunnel co-ordinate system needs to be calibrated. The spanwise calibration is achieved during verification of the beam alignment using the 35 μm pin hole, which is placed on the tunnel centre span in a 3D-printed holder. The accuracy of this process is limited by the resolution of the 3D printing process, i.e. $\Delta z \approx 0.1$ mm. Note that this is much smaller than the probe volume dimension in this spanwise direction.

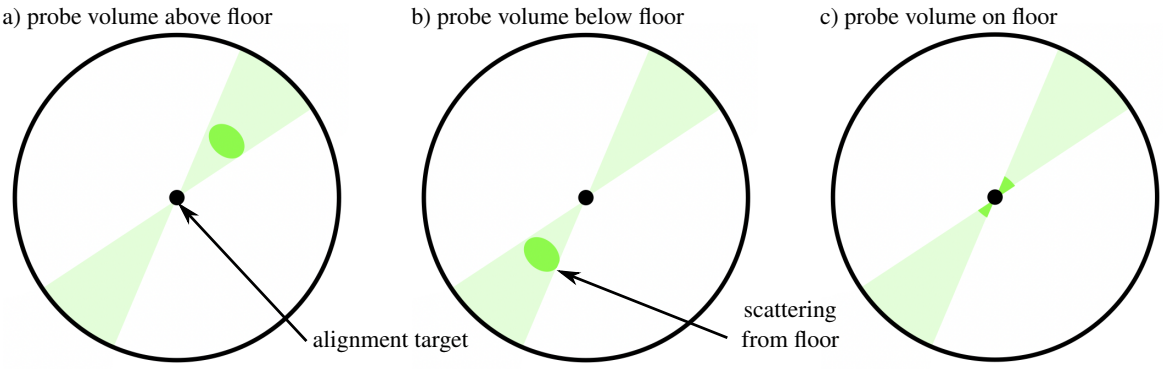

Figure 3.17: Scattering from floor, viewed through the receiving optics, when the laser beams intersect: a) above the floor; b) below the floor; c) on the floor surface. Adapted from [84].

The vertical position is calibrated by observing the probe volume through the receiver optics, as it is moved onto the wind tunnel floor. When the measurement volume is located on the floor, there is significant scattering from the surface. Figures 3.17a and 3.17b shows that if the probe volume is slightly misaligned with the floor, the scattered light forms a bright spot either above or below the alignment target in the receiving optics. However, when the laser beams are focused precisely on the floor, as in figure 3.17c, this bright spot lies on top of the alignment target. The corresponding positional error is given by the size of the traverse steps in the vertical direction, and so $\Delta y \approx 0.005$ mm. Note that, as discussed in section 3.1, the tunnel floor and ceiling diverge to provide boundary-layer relief – as a result, the calibration in the vertical direction needs to be performed at each measurement station.

The probe volume's streamwise position is determined by visually placing it on top of a feature with known x-coordinate. This might be a pressure tap, or a junction between the liner blocks which make up the tunnel floor. The corresponding positional accuracy is determined by the feature size and so, typically, $\Delta x \approx 0.2$ mm.

Flow seeding

The seeding particles in the flow are generated by introducing high-pressure air (≈ 750 kPa) to a paraffin reservoir through a double Laskin nozzle, as shown in figure 3.18. This results in atomisation of the paraffin. Whilst heavier droplets fall back into the reservoir, the smallest, lightest particles are forced by the air pressure into a secondary chamber, called the impactor. This chamber contains a series of perforated plates, which trap any large particles. Therefore, only the smallest paraffin droplets remain.

Colliss conducted a thorough investigation of the seeding particle size distribution, in the same facility as the current study [84]. The particle lag through a normal shock was measured using LDV, and compared with a theoretical model based on Stokes flow. Figure 3.19 shows this comparison for theoretical models with different particle sizes, which suggests that the mean particle diameter, d_p, is approximately 0.5 μm.

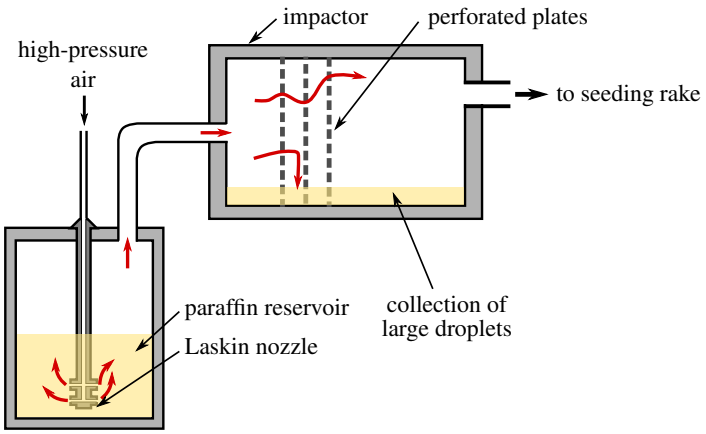

Figure 3.18: Schematic of seeding particle generation, with the path of paraffin droplets indicated by red lines. Adapted from [84].

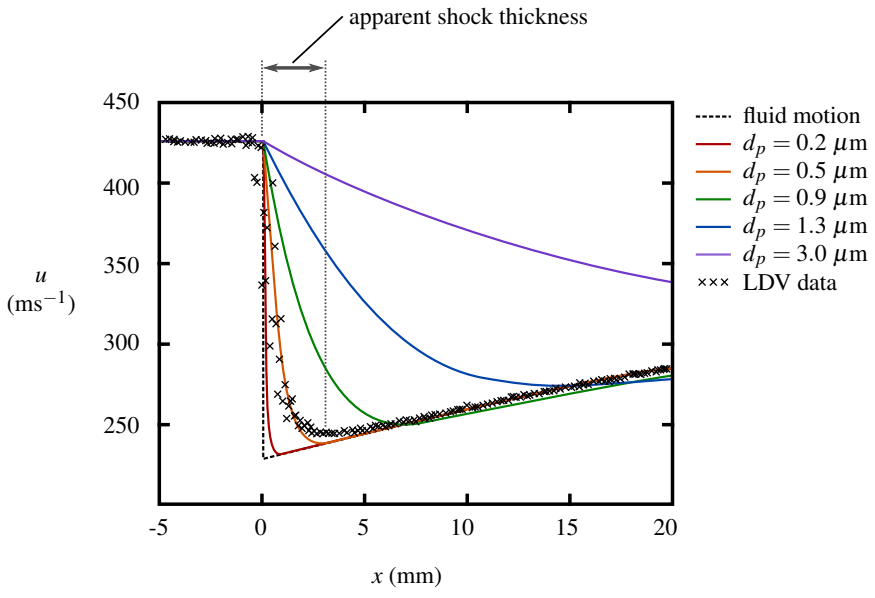

Figure 3.19: Size of seeding particles, determined by comparing the measured particle lag through a normal shock with a theoretical model. Adapted from [84].

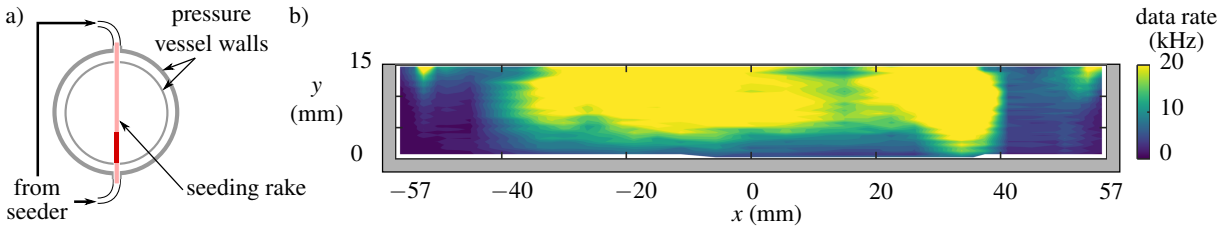

Figure 3.20: a) Schematic cross-sectional view of seeding rake (light red), in the tunnel's settling chamber. The dark red section of the rake indicates where seeding particles are introduced to the flow. b) Seeding rates in the floor boundary layer. Measurements conducted in the test section, across the tunnel span at $x = 120$ mm.

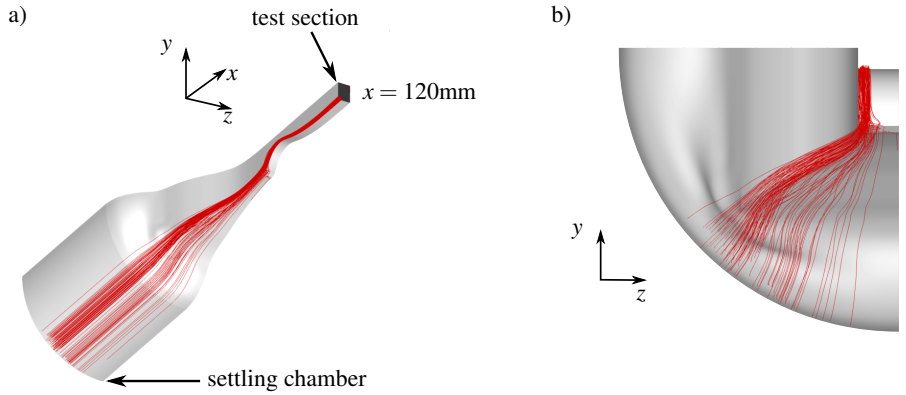

Figure 3.21: The origin of streamlines, from RANS simulations, which end up in the corner region of the test section: a) isometric view; b) cross-sectional view.

Knowledge of the seeding particle size is required to estimate their ability to accurately follow the flow. It is necessary for the particles to respond to the significant velocity gradients present in the studied flow field with minimal lag. The Stokes number, S_k, is estimated using [103]:

$$S_k = \frac{\tilde{t}_p}{\tilde{t}_f} = \frac{\rho_p d_p^2 u_\infty}{18\mu\delta} \,, \tag{3.10}$$

where the characteristic timescale of the particle, $\tilde{t}_p = \rho_p d_p^2 / 18\mu$. For shock–boundary-layer interactions, the timescale of the fluid is typically defined using outer-flow properties, i.e. $\tilde{t}_f = \delta / u_\infty$ [107]. In these expressions ρ_p denotes the particle density, μ gives the dynamic viscosity of the fluid, u_∞ is the freestream velocity, and δ is the boundary-layer thickness. For $u_\infty \approx 570$ ms^{-1} and $\delta \approx 6$mm, equation 3.10 evaluates to $S_k \approx 0.06$. Since $S_k < 0.1$, the seeding particles are expected to faithfully capture the motion of the carrier fluid, with errors not expected to exceed 1% [103].

The seeding particles are introduced to the flow through a rake in the settling chamber, shown in figure 3.20a. The standard rake design consists of vertical rod with approximately ten holes of diameter 2 mm, directed in the downstream direction. Several rakes are available with different distributions of holes along the rake height.

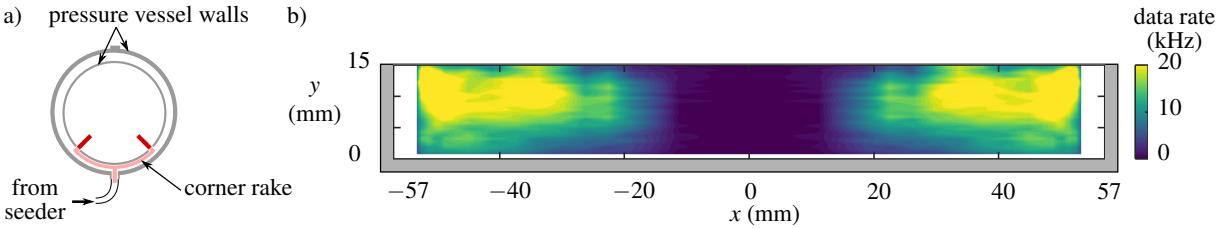

Figure 3.22: a) Schematic cross-sectional view of corner seeding system (light red), in the tunnel's settling chamber. The dark red section of the rake indicates where seeding particles are introduced to the flow. b) Seeding rates in the floor boundary layer when using the corner seeding system. Measurements conducted in the test section, across the tunnel span at $x = 120$ mm.

The seeding particles are introduced on the tunnel centre span. However, as they are advected downstream, entrainment into the floor boundary layer and turbulent mixing spreads the seeding particles in the spanwise direction. The resultant seeding particle distribution, shown in figure 3.20b, is parametrised by the 'data rate', i.e. the number of particles that pass through the LDV probe volume in one second. Typically data rates in excess of 10 kHz are required to obtain validation-quality LDV measurements.

Figure 3.20b shows that the corner regions are not seeded sufficiently when the vertical rake is used. This is particularly problematic for an investigation focused on the corner boundary layers of the wind tunnel channel. In order to probe these regions, modifications to the seeding system were required to introduce seeding particles into the corner boundary layers.

This modified seeding system, located at the same streamwise location as the vertical rake, is designed with the help of RANS computations of the wind tunnel, described in section 3.3. In order to establish the optimum rake location, streamlines within the corner boundary layer of the test section were isolated. These streamlines were then traced upstream to the settling chamber, and so their origin in the seeding plane could be determined (figure 3.21).

Placing the 'corner rake' at the location shown in figure 3.22a therefore ensures that seeding particles will end up in the corner boundary layer within the test section. The seeding rates measured after the installation of the new rake are shown in figure 3.22b. This particle distribution demonstrates that the modification was successful, and so validation-quality reference data can be collected in the corner regions.

Boundary-layer profiles

A common use of laser Doppler velocimetry is to take measurements of the wind tunnel boundary layers. The method used to conduct these tests is therefore presented below in some detail. Whilst LDV data have also been obtained in some other parts of the flow — for example, streamwise traverses through an oblique shock reflection — the analysis procedure is analogous to that described in this section.

The velocities of seeding particles are measured continuously during a traverse normal to the surface, such as the wind tunnel floor. The data are then split into bins with spatial resolution,

$\Delta y = 0.15$ mm, on the same order as the vertical dimension of the probe volume. There are typically at least 2000 particles within each bin, except in the immediate vicinity of the wall, where there may be fewer than 1000 particles per bin. Within each bin, any data points outside three standard deviations of the mean velocity are considered to be outliers and discarded.

In order to accurately calculate the mean velocity, it is necessary to account for a phenomenon called velocity bias. This occurs in both shear flows and turbulent flows, where seeding particles at a range of velocities pass through the probe volume [108]. The number of particles traversing the probe volume per second is given by $A_V \times n_p \times u$, where A_V is the cross-sectional area of the probe volume (normal to the flow) and n_p is the number density of particles in the flow. Therefore, assuming uniform seeding, in every second there are more high velocity particles that pass through the measurement volume than low velocity. As a result, a simple average of all data points biases the mean towards higher velocity. In order to account for this effect, every data point is weighted by inverse velocity, such that lower-velocity particles are assigned higher weights. The mean velocity within each bin, u, is therefore

$$u = \frac{N}{\sum_{i=1}^{N} u_i^{-1}} , \tag{3.11}$$

where u_i is the velocity of particle i, and the sum is evaluated over the N particles in each bin [108].

From these mean velocities, it is possible to generate a profile of the traversed boundary layer. In order to gain further insight, several parameters which define the state of the local boundary layer are calculated. Perhaps the simplest parameter is δ_{99}, which corresponds to the wall-normal distance where the velocity reaches 99% of the freestream value. More robust integral parameters for boundary-layer thickness include the displacement thickness, δ_i^*, which quantifies the deflection of streamlines due to the boundary layer, and the momentum thickness, θ_i, which is a measure of momentum loss.

The i subscript indicates that these are definitions associated with incompressible flow, assuming uniform density, which is clearly not valid for the flows in this study. However, the corresponding compressible definitions, δ^* and θ, are functions of Mach number, rendering them inappropriate as universal parameters [22]. It is therefore common practice in the field to quote equivalent incompressible boundary-layer parameters. This approach is further justified by Morkovin's hypothesis that compressibility effects on the boundary layer are negligible for Mach numbers lower than 5 [109]. The boundary-layer parameters are defined by:

$$\delta_i^* = \int_0^{\delta_{99}} \left(1 - \frac{u}{u_e}\right) \mathrm{d}y , \tag{3.12}$$

$$\theta_i = \int_0^{\delta_{99}} \frac{u}{u_e}\left(1 - \frac{u}{u_e}\right) \mathrm{d}y , \tag{3.13}$$

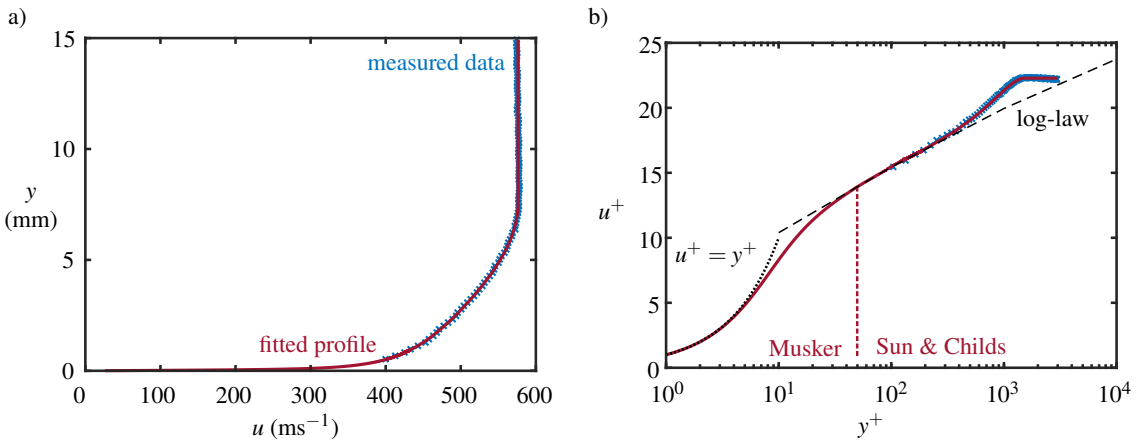

Figure 3.23: Boundary-layer LDV data and fitted profile in: a) dimensional, and b) non-dimensional units. Measurements performed at $x = 60$ mm, $z = 0$ mm. The non-dimensional data are presented in wall units, $u^+ = u/u_\tau$ and $y^+ = yu_\tau/\nu_w$. The error bars are contained within the symbol size with $\Delta u = 5$ ms^{-1} and $\Delta y = 0.1$ mm.

where u_e is the streamwise velocity at the boundary-layer edge. In order to characterise the fullness of the boundary-layer velocity distribution, the shape factor H_i is used, where:

$$H_i = \frac{\delta_i^*}{\theta_i} \,. \tag{3.14}$$

These incompressible definitions are used to determine the boundary-layer properties. Alongside the insensitivity to variations in Mach number, this approach has the additional benefit that fewer assumptions are required to calculate the parameters from raw velocity data, compared to the corresponding compressible definitions [15].

In order to minimise any errors due to discretisation of the velocity data [110], the measured velocities are fitted to theoretical profiles, as shown in figure 3.23. This fitting procedure has been developed in the Cambridge facility — in particular by Colliss [84], Oorebeek [111], and Davidson [82] — over many years.

A Sun & Childs (1973) fit [112] is used for the outer layer. This is based on a linear combination of the log-law of the wall and Coles' wave function [63], but with two key modifications. First, compressibility effects are accounted for using a van Driest correction [113]. Second, the shear stress is assumed not to be uniform, but instead to vary linearly across the boundary layer – this ensures that the wall-normal velocity gradient is zero at the boundary-layer edge.

The assumption of adiabatic flow is valid in the current study due to the small difference between recovery temperature (approximately the stagnation temperature at moderate Mach number) and the wall temperature (approximately room temperature). For such an adiabatic flow, the mathematical

formulation of this model is given by:

$$\frac{u}{u_e} = \frac{1}{\sqrt{\sigma}} \sin\left\{ \left(\sin^{-1}\sqrt{\sigma}\right) \left[1 + \frac{1}{\kappa}\frac{u_\tau}{u_e^*} \left(\ln\left(\frac{y}{\delta}\right) + 2\sqrt{1-\frac{y}{\delta}} - 2\ln\left(1+\sqrt{1-\frac{y}{\delta}}\right)\right) - \frac{\Pi}{\kappa}\frac{u_\tau}{u_e^*}\left(1+\cos\left(\frac{\pi y}{\delta}\right)\right)\right]\right\} , \quad (3.15)$$

where:

$$\sigma = \frac{\frac{\gamma-1}{2}M_e^2}{1+\frac{\gamma-1}{2}M_e^2} , \quad (3.16)$$

$$\frac{u_\tau}{u_e^*} = \frac{1}{\sin^{-1}\sqrt{\sigma}}\sqrt{\frac{\sigma c_f}{2(1-\sigma)}} , \quad (3.17)$$

$$\frac{\Pi}{\kappa} = \frac{1}{2}\left[\frac{u_e^*}{u_\tau} - \frac{1}{\kappa}\ln\left(\frac{\delta u_\tau}{\nu_w} - \mathrm{B}\right) + \frac{0.614}{\kappa}\right] . \quad (3.18)$$

In these expressions, the Karman constant, $\kappa = 0.41$, and the log-law intercept, $\mathrm{B} = 5.0$. In addition: u^* is the velocity (u) transformed by the van Driest compressibility correction; ν_w is the kinematic viscosity at the wall; γ is the ratio of specific heats; subscript e denotes a quantity at the boundary-layer edge [112]. The friction velocity, $u_\tau = \sqrt{\tau_w/\rho_w}$, where τ_w and ρ_w are the shear stress and the density at the wall, respectively.

The three remaining independent parameters define the boundary-layer properties: the skin-friction coefficient (c_f), the boundary-layer thickness (δ), and the velocity at the boundary-layer edge (u_e). These three quantities are used as fitting parameters to minimise the residual between the velocity data and the theoretical model in equation 3.15. Even for boundary-layer profiles which are far from equilibrium in adverse pressure gradients, this model fits the data well and the estimated R^2 value typically exceeds 0.99.

However, the Sun & Childs model only applies to the outer layer of the boundary layer. In order to resolve the viscous sub-layer and the buffer zone, an expression for the velocity profile is determined using the method proposed by Musker (1979) [114]. Scaling arguments are used to determine a cubic variation in eddy kinematic viscosity with wall-normal distance, which leads to the experimentally-verified expression:

$$\frac{\mathrm{d}u_i^+}{\mathrm{d}y^+} = \frac{\frac{y^{+2}}{\kappa}+\frac{1}{s}}{y^{+3}+\frac{y^{+2}}{\kappa}+\frac{1}{s}} , \quad (3.19)$$

where the constant intrinsic to the model, $s = 1.09\times10^{-3}$ [114]. This model is developed for an incompressible boundary layer, but the equivalent compressible form is found, as for the outer layer, by applying a van Driest correction [113]. Equation 3.19 uses wall (or inner) units for the parameters, $u^+ = u/u_\tau$ and $y^+ = yu_\tau/\nu_w$. In these expressions, the compressible friction velocity, u_τ, is determined using equation 3.17 with the skin-friction coefficient, c_f, of the fitted Sun & Childs profile. This

process ensures that, when the Musker and the Sun & Childs profiles are matched at $y^+ = 100$, the velocity (u) and the shear ($\mathrm{d}u/\mathrm{d}y$) are continuous.

A continuous boundary-layer profile which extends to the wall, and which closely matches the discrete velocity measurements, is therefore constructed. From this fitted profile, the characteristic boundary-layer integral parameters — displacement thickness (δ_i^*), momentum thickness (θ_i), and shape factor (H_i) — are calculated by numerical integration.

Uncertainty in LDV measurements

It is important to carefully quantify the uncertainties in velocity measurements, since they propagate through the data-processing procedure and influence the calculated boundary-layer parameters.

The error intrinsic to the LDV system is, in fact, relatively small since the data processor dynamically selects an appropriate sampling frequency on detection of a Doppler burst. This maximises the accuracy of signal frequency calculation, and so the manufacturer-quoted uncertainty is 0.1% of the maximum measurable velocity [105]. Table 3.1 indicates that this maximum velocity is $550\,\mathrm{m\,s^{-1}}$, and so the corresponding error is $\pm 0.55\,\mathrm{m\,s^{-1}}$.

A more dominant source of random error is related to the number of seeding particles detected by the receiving optics. This is determined firstly by the alignment of the optics, in particular the accuracy to which the four beams converge upon the same probe volume. A second factor is the intrinsic number of seeding particles passing through the region occupied by the probe volume. These two factors influence the number of data points within every measurement bin, which affects the corresponding statistical error in the mean velocity. Most bins contain at least 2000 particles. With a nominal 10% maximum turbulence intensity within the boundary layer, a conservative error estimate of 2% is obtained. However, this uncertainty goes up to 3% for bins with only 1000 data points.

There is a further contribution to the overall uncertainty from the fact that the laser emitting head is rotated about its axis (figure 3.15a) and inclined to the horizontal (figure 3.15b). This means that the fringes within the probe volume are not aligned with the co-ordinate axes. For emitting head rotation angle, ψ, and inclination angle, β, the measured velocity components in the two LDV channels are:

$$v_1 = u\cos\psi + v\sin\psi\cos\beta + w\sin\psi\sin\beta\ , \tag{3.20}$$

$$v_2 = u\sin\psi - v\cos\psi\cos\beta - w\cos\psi\sin\beta\ , \tag{3.21}$$

where u, v and w are the true velocity components in the streamwise, vertical, and spanwise directions, respectively. To recover estimates for u and v from our two-component LDV data, the spanwise component, w, is typically assumed to be zero. This gives:

$$\tilde{u} = v_1\cos\psi + v_2\sin\psi\ , \tag{3.22}$$

$$\tilde{v} = v_1\sin\psi - v_2\cos\psi\ . \tag{3.23}$$

Substituting the above expressions for v_1 and v_2:

$$\tilde{u} = u \,, \tag{3.24}$$

$$\tilde{v} = v\cos\beta + w\sin\beta \,. \tag{3.25}$$

These extracted velocities are independent of the emitting head rotation angle, ψ. In addition, the extracted value of u corresponds exactly to the true value, though there is a small associated error of 0.1% due to uncertainties in the values of ψ and β themselves.

The tilt of the emitting head produces an inaccurate estimate for the vertical velocity, v. This error can be as small as 1% when $w = 0$. However, the spanwise and vertical velocities are generally expected to be of the same order, resulting in a more appropriate error estimate of 12%.

An additional source of error is the particle lag due to the finite size of particles [103]. Figure 3.19 shows that, despite the particle diameter being only 0.5 μm, there is still some lag through a normal shock – the infinitely-thin shock wave has an apparent thickness of 3 mm. However, this source of error will only be significant in the immediate vicinity of shock waves, where the velocity gradients are significant.

Collating these contributions, the uncertainties in u and v are 2% and 14% respectively. Note, however, that the error can be higher in near-wall regions, where the seeding rates are generally lower, or in the vicinity of shock waves, where particle lag is significant.

Whilst LDV measurements are typically reported as point values, they are really an average velocity across the finite probe volume. The assumption of point measurement cannot be assessed using experiments alone, since the true 'un-averaged' velocities are not known *a priori*. Therefore, the RANS computations described in section 3.3 are used to provide representative boundary-layer profiles. These simulations are then interrogated in an analogous manner to the LDV measurement process, by averaging over an orthogonal mesh with dimensions $0.1 \times 0.1 \times 1.4$ mm. This mesh, shown in figure 3.24, represents the probe volume. The 'long' grid dimension of 1.4 mm is aligned in the spanwise direction.

The representative LDV volume is centred at one point at a time from the extracted profile, and the RANS solution is interpolated onto the LDV volume grid. The velocity values in this region are then averaged. By averaging uniformly across the simulated LDV volume, an estimate is obtained for the averaging effect in real LDV measurements. This estimate is conservative because the real probe volume is ellipsoidal, and because the measured data are biased towards particles passing through the centre of this region.

Figure 3.25a shows the results for a floor boundary-layer profile on the tunnel centre line, while figure 3.25b shows the equivalent results for a sidewall boundary layer. Selected data points from these figures are listed in table 3.2. The floor boundary layer does not show any noticeable difference between the true and the averaged profiles, except within 0.1 mm of the wall. However, the equivalent analysis of the sidewall boundary layer is noticeably different. In this case, the velocity values normal

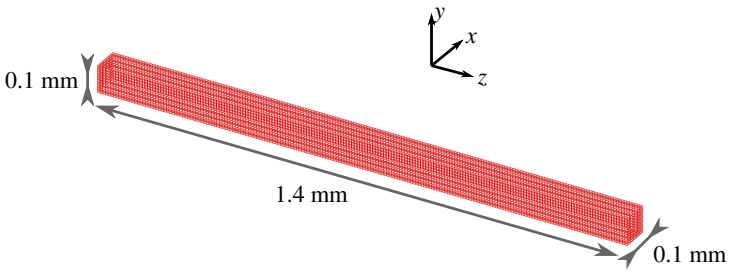

Figure 3.24: Mesh representing LDV probe volume, over which measured velocities are averaged.

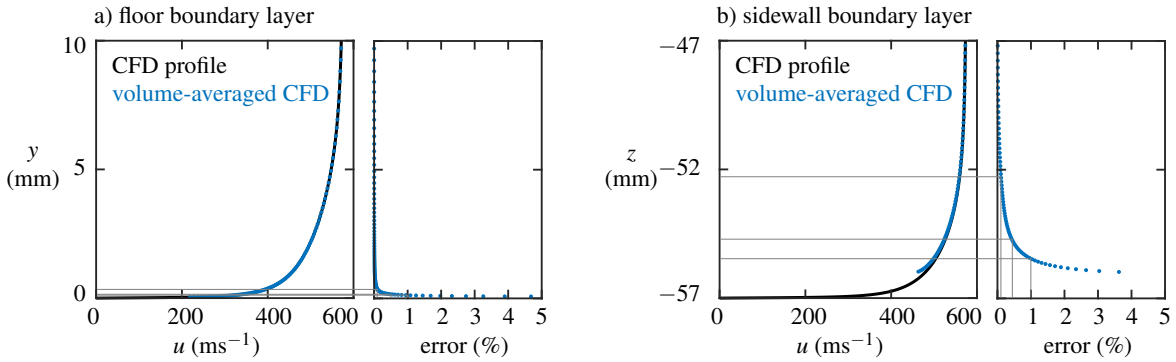

Figure 3.25: Comparison of streamwise velocities extracted directly from the RANS solution with velocities averaged over a representative measurement volume: a) the floor boundary layer ($x = 60$ mm, $z = 0$ mm), and b) the sidewall boundary layer ($x = 60$ mm, $y = 15$ mm).

Table 3.2: Error due to averaging over a finite probe volume for the floor boundary layer (figure 3.25a), and the sidewall boundary layer (figure 3.25b).

floor boundary layer		sidewall boundary layer	
error (%)	$y - y_{\text{wall}}$ (mm)	error (%)	$z - z_{\text{wall}}$ (mm)
0.1	0.62	0.1	4.89
0.5	0.35	0.5	2.13
1.0	0.16	1.0	1.54

to the wall are averaged over 1.4 mm rather than 0.1 mm. As a result, the averaging region now forms a significant proportion of the boundary-layer thickness. Therefore, averaged velocities within 1.4 mm of the wall show a discrepancy from the true values. However, within 1 mm of the wall, no data can be obtained because the probe volume intersects the sidewall surface, and the resulting bright scattering inhibits any measurements.

There is a further contribution to the overall velocity uncertainty, due to the fact that the probe volume is constantly traversing through the boundary layer, rather than stopping to collect data at each station. Therefore, each measurement 'point' is in fact an average over the region traversed during data collection for that bin . Typically, the traverse length per measurement point is 0.15 mm, and so extends the 'effective probe volume' by this amount in the direction of the traverse. As a result, the effective probe volume during a floor boundary-layer traverse is in fact 0.27 mm in the vertical direction, and for a sidewall boundary-layer traverse it has a spanwise extent of 1.5 mm. In theory, these values should have been used for the representative LDV probe volume in figure 3.24, but the existing computations can still provide useful information. The estimated error in figures 3.25a and 3.25b is less than 1%, except within 1.1–1.3 times the probe volume extent in the traverse direction. Using this rough approximation, we might expect to reach 1% error further away from the wall than table 3.2 suggests, i.e. at around 0.35 mm for floor boundary-layer traverses and 1.64 mm for sidewall boundary-layer traverses.

Therefore, the LDV measurements can be treated as point values in most regions of the flow except within 0.35 mm of the wind tunnel floor and within 1.64 mm of the sidewalls, where the finite size of the probe volume needs to be taken into account.

Further uncertainty is introduced when calculating the boundary-layer parameters (δ_i^*, θ_i, and H_i) from the discrete velocity data. A comprehensive assessment of errors introduced by this type of method was conducted by Titchener *et al.* [110]. Direct numerical simulations of boundary layers provided the true values of relevant parameters. These were compared with the values obtained by processing the discrete velocity data extracted from the DNS profiles.

This study highlights the benefits of fitting theoretical models, as compared to linear interpolation between the discrete data points. The fitting method invoked typically reduces the error in boundary-layer parameters by a factor of two. The most important factors affecting the parameter uncertainty are the spatial resolution of data points and, in particular, the location of the closest data point to the wall. The LDV measurements obtained in this study typically have around 40 measurement points within the boundary layer and the closest data point to the wall is at around $y^+ = 80$. This corresponds to an uncertainty in integral parameters of around 5% for an equilibrium turbulent boundary layer [110].

3.2.5 Summary of experimental errors

A validation study requires an informed comparison between measured data and the results of computations. This demands a comprehensive assessment of the experimental uncertainties. The

relevant sources of error have been discussed for each technique in this section, and are summarised in table 3.3 for reference.

Table 3.3: Summary of experimental uncertainties.

property	technique	source	uncertainty
wind tunnel geometry x, y, z	dial-test indicator	instrument error angular alignment	± 0.02 mm $\pm 0.1^\circ$
static pressure (point) p/p_0	pressure taps	tap geometry transducer error atmospheric pressure stagnation pressure	$\pm 0.5\%$ $\pm 0.05\%$ $\pm 0.03\%$ $\pm 0.3\%$
static pressure (surface) p/p_0	pressure-sensitive paint	intrinsic noise tap error (calibration)	$\pm 2\%$ $\pm 1\%$
Pitot pressure p_{pit}/p_0	Pitot probe traverse	transducer error atmospheric pressure stagnation pressure temporal lag & position	$\pm 0.05\%$ $\pm 0.03\%$ $\pm 0.3\%$ $\pm 3\%$
flow velocity u, v	laser Doppler velocimetry	processor error statistical error angled laser emitter	$\pm 0.1\%$ $\pm 2\%$ $\pm 0.1\%$ in u $\pm 12\%$ in v
probe volume position x, y, z	laser Doppler velocimetry	feature size traverse step size 3D print resolution	± 0.2 mm in x ± 0.005 mm in y ± 0.1 mm in z
boundary-layer parameters δ_i^*, θ_i, H_i	laser Doppler velocimetry	velocity error spatial resolution	$\pm 2\%$ $\pm 5\%$

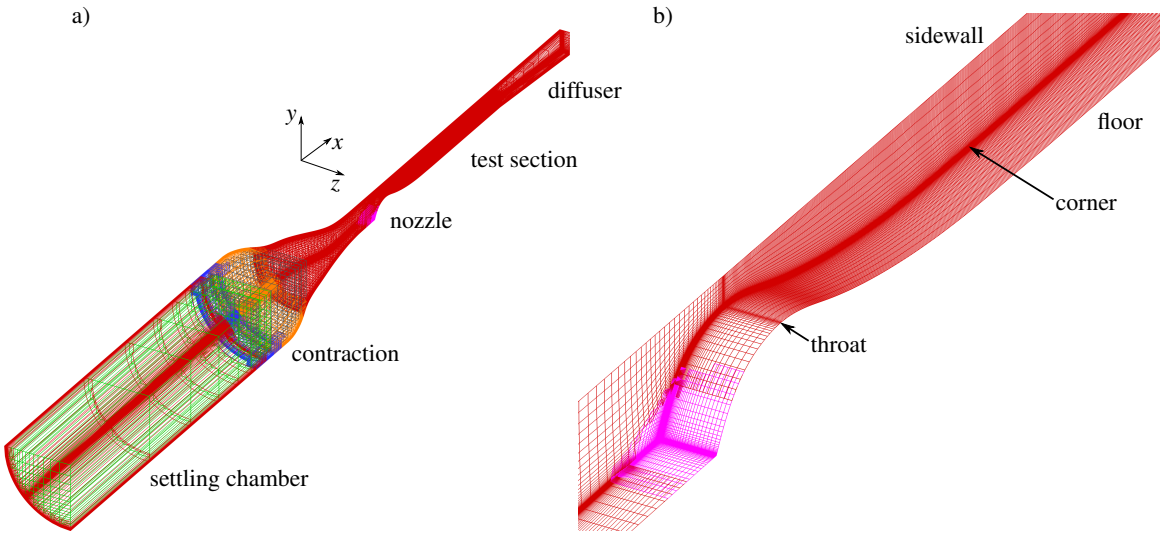

Figure 3.26: Representation of mesh used in computations to simulate physical tunnel: a) isometric view of entire mesh; b) surface grid in nozzle and start of test section.

3.3 Computational method

Alongside the experiments carried out as part of this study, numerical RANS simulations of the flow in the tunnel are performed by Daniel Galbraith and John Benek at the US Air Force Research Laboratory. Some of their data is included in this thesis to illustrate the use of flow characterisation information for calibration, to directly compare computational models against experimental data, and to provide experimentally-inaccessible information about the flow field.[2] Therefore, a brief description of the method for these computations is provided below.

The solver OVERFLOW 2.2l [115] is used to solve the RANS equations using the third-order accurate upwind finite difference HLLC scheme [116] combined with the Koren limiter [117]. The time integration uses an unfactored SSOR implicit solution algorithm [118]. The computations are performed using the Spalart–Allmaras turbulence model [119, 120] with and without the quadratic constitutive relation [70, 71], and the resulting flow fields are compared.

To accelerate convergence to a steady state, both grid sequencing and multi-grid are implemented. Grid sequencing involves computing an initial solution on a coarse grid, then successively finer grids. Meanwhile, multi-grid is a process in which, for every solution step, the solution vector is updated based on corrections calculated on a coarse grid; this has the effect of damping out high frequency errors in the fine grid solution.

Simulations of empty wind tunnel

In order to generate the mesh, a CAD model of the wind tunnel is used, supported by measurements of the physical tunnel geometry. In contrast to many computations of experimental facilities, the entire

[2]Whilst the computations were conducted by engineers at the US Air Force Research Laboratory, analysis of the data was performed by the author.

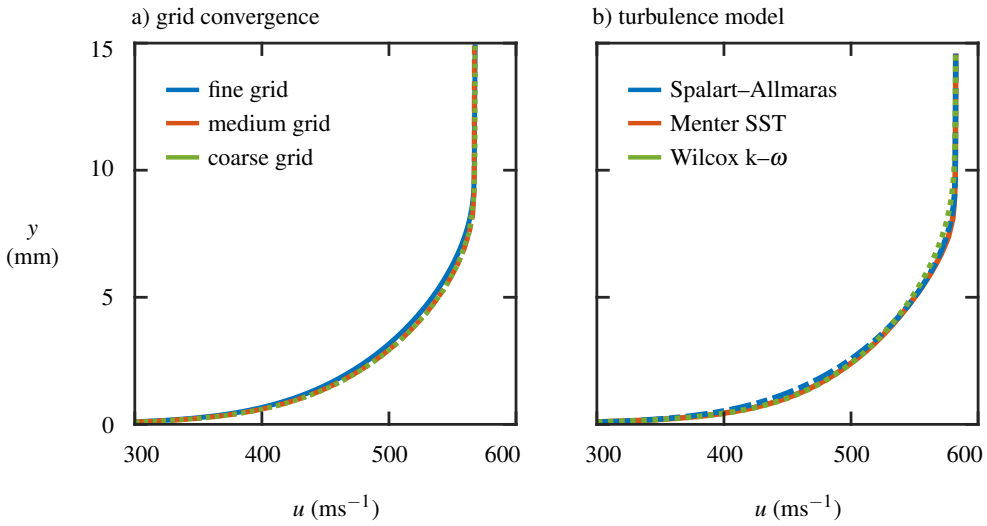

Figure 3.27: Floor boundary-layer profiles computed on the centre span at $x = 60$ mm: a) comparison of profiles for the coarse, medium and fine grids; b) comparison of profiles for the Spalart–Allmaras, Menter SST and Wilcox k–ω turbulence models.

wind tunnel flow is simulated. The mesh starts in the settling chamber and ends in the diffuser, well downstream of the test section.

For the full nozzle, only one quarter of the tunnel is modelled since it is, in theory, top–bottom and left–right symmetric. Similarly, only one half of the tunnel with the half nozzle configuration is simulated. The Chimera overset grid technique [121] is used to create a smooth mesh in the contraction and at the sharp corner upstream of the nozzle. The final grid system, created using the mesh generation software Pointwise [122], contains 181.7M points across seven grids (figure 3.26). A viscous wall spacing of 1.5×10^{-7} m is used with a growth rate of 5% – this corresponds to $y^+ < 1$ at the first point from the wall.

At the inflow plane, the stagnation pressure and the stagnation temperature are set at 308.0 kPa and 288.0 K, respectively, to match the conditions of the physical wind tunnel. The flow angle in this plane is defined to be perpendicular to the boundary. These boundary conditions impose a uniform velocity profile at the inflow [113]. Therefore, the inflow plane cannot contain a boundary layer along the tunnel surfaces, and so a no-slip wall cannot be used at its interface with the tunnel walls. To overcome this, the first few cells along the walls that meet with the inflow plane are modelled as adiabatic slip-walls.

Meanwhile, an extrapolated boundary condition is used for the outflow plane, where the flow is assumed to be supersonic. This type of boundary condition specifies that solution on the outflow plane is equal to the solution one cell upstream of this plane. The walls are modelled with a no-slip condition and are isothermal at a room temperature of 291 K. A cheap solution is first generated using one-dimensional nozzle theory – this is used as an initial solution for the complete mesh to accelerate convergence.

A grid resolution study is conducted by computing the flow on three meshes with different spatial resolutions. To maintain the point distribution of the original grid, every other point is removed from this fine mesh to create a medium grid (22.9M points), and every other point is removed from the medium grid to produce a coarse grid (2.9M points). Figure 3.27a shows the centre-span boundary layer on the test section floor for the three grid levels. The three velocity profiles are coincident, indicating that the centre-span profiles have reached grid convergence.

The Spalart–Allmaras turbulent model is chosen for its simplicity, numerical stability, and established accuracy for a range of similar flow conditions to the experiment. A further advantage of the Spalart–Allmaras model is its ease of integration with the quadratic constitutive relation. In order to justify this choice, centre-span floor boundary-layer profiles using Spalart–Allmaras [119, 120], Menter SST [123, 124], and Wilcox k–ω [125, 126] turbulence models are compared in figure 3.27b. The minimal differences between the profiles shows that computations of this flow appear to be insensitive to the turbulence model used. This suggests that the use of the Spalart–Allmaras model in this flow field is, indeed, valid.

3.4 Summary

This chapter describes the different nozzle geometries used in the Cambridge facility to investigate supersonic corner flows, as well as the response of the floor boundary layer to an oblique shock reflection. A range of techniques are used to comprehensively interrogate the flow field. Both schlieren imaging and surface oil flow enable the flow to be visualised. Pressure measurements are obtained using a fixed Pitot rake (error: $\pm 0.7\%$), Pitot probe traverse (error: $\pm 3\%$), static pressure taps (error: $\pm 1\%$), and pressure-sensitive paint (error: $\pm 3\%$). In addition, laser Doppler velocimetry provides measurements of both the streamwise velocity (error: $\pm 2\%$) and the vertical velocity (error: $\pm 14\%$).

The chapter also introduces RANS simulations conducted by the US Air Force Research Laboratory, from which data are presented in the remainder of this thesis. The flow is computed for the entire tunnel, from the settling chamber to the diffuser. The simulations use the Spalart–Allmaras turbulence model, and the computed velocities are shown to be grid-converged and independent of the chosen turbulence model.

The methodology for these computations is based on the wind tunnel parameters, but a comparison with experimental data has not yet been presented. For a validation study, which aims to determine the accuracy and limitations of numerical models, it is important to characterise the physical wind tunnel and to carefully calibrate simulations. It is now accepted that a perfect wind tunnel is near-impossible to build and operate. Instead, the emphasis is on careful measurement of the 'imperfections' of the physical wind tunnel, so that computations can accurately reproduce the true flow [127]. The process of characterising the wind tunnel flow to calibrate the RANS simulations forms the subject of Chapter 4.

Chapter 4

Flow characterisation for validation

An important part of this validation study is the characterisation of the wind tunnel flow to calibrate relevant computations. Compared to most phenomenological investigations, this characterisation process demands a more detailed focus on experimental accuracy and on mitigating systematic errors. This chapter details the comprehensive tests performed to obtain suitable calibration data, as well as the synergetic benefits of close experimental–computational collaboration during this process.

4.1 Flow characterisation for a routine experimental study

If the current study had simply been an experimental investigation of flow physics, the characterisation process would have been much simpler. For these 'phenomenological' studies, the focus is often on the identification of key flow features or on determining the response of a flow field to various changes. As a result, it is generally adequate to perform basic characterisation tests in order to ensure that the flow quality is 'sufficient' to capture the features of interest.

One important measure of the flow quality of a wind tunnel is the uniformity of pressure at the entrance to the test section. The stagnation pressure in this region, measured using a rake of Pitot probes, is presented in figure 4.1. This shows a highly uniform flow over the tunnel cross-section, with a maximum variation in stagnation pressure across the entire inflow of 0.1%.

Within the test section itself, schlieren images are captured for both nozzle configurations used in the current study. Figures 4.2a and 4.2b show that the flow in the tunnel is established. It is also possible to identify the boundary layers along the tunnel's floor and ceiling, as well as waves emanating from imperfections in these surfaces.

A final measure of flow quality in a routine experimental study is the thickness of the wind tunnel boundary layers relative to the core flow. This provides, for example, an estimate of the tunnel's viscous aspect ratio, introduced in section 2.2.1, which indicates whether confinement and three-dimensional effects are likely to strongly influence the flow. A representative boundary-layer

Section 4.2.3 of this chapter presents data from RANS computations performed by Daniel Galbraith and John Benek at the US Air Force Research Laboratory. Analysis of the data was performed by the author.

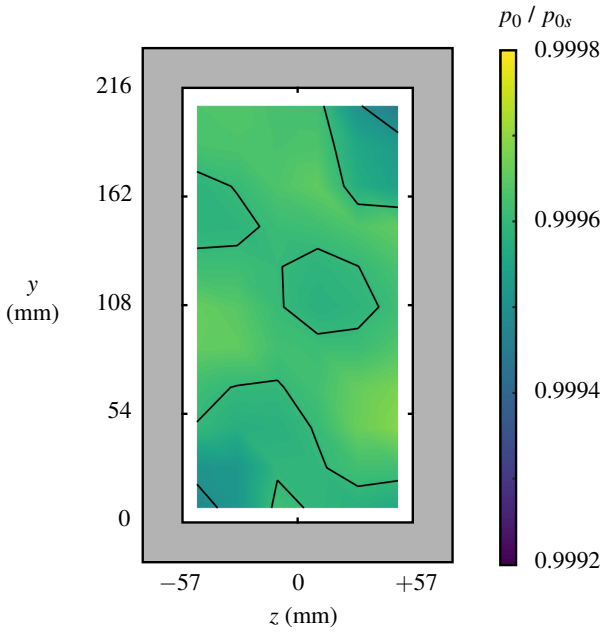

Figure 4.1: Stagnation pressure measured by Pitot rake at exit of the settling chamber, upstream of the nozzle. The Pitot probes are located in a square grid, at 20 mm intervals, between $x = -50$ and 50 mm and from $y = 8$ to 208 mm.

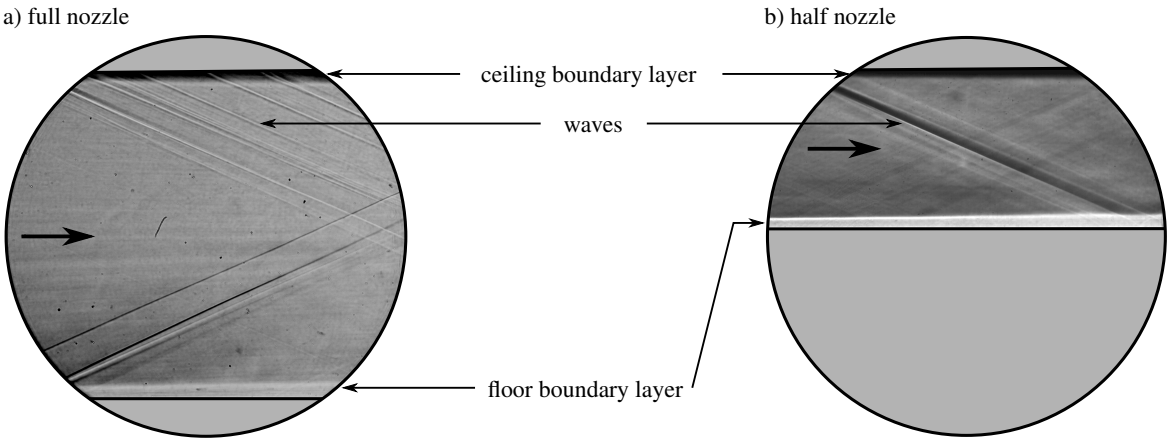

Figure 4.2: Schlieren image for the empty tunnel with flow established for a) the full nozzle setup, and b) the half nozzle setup.

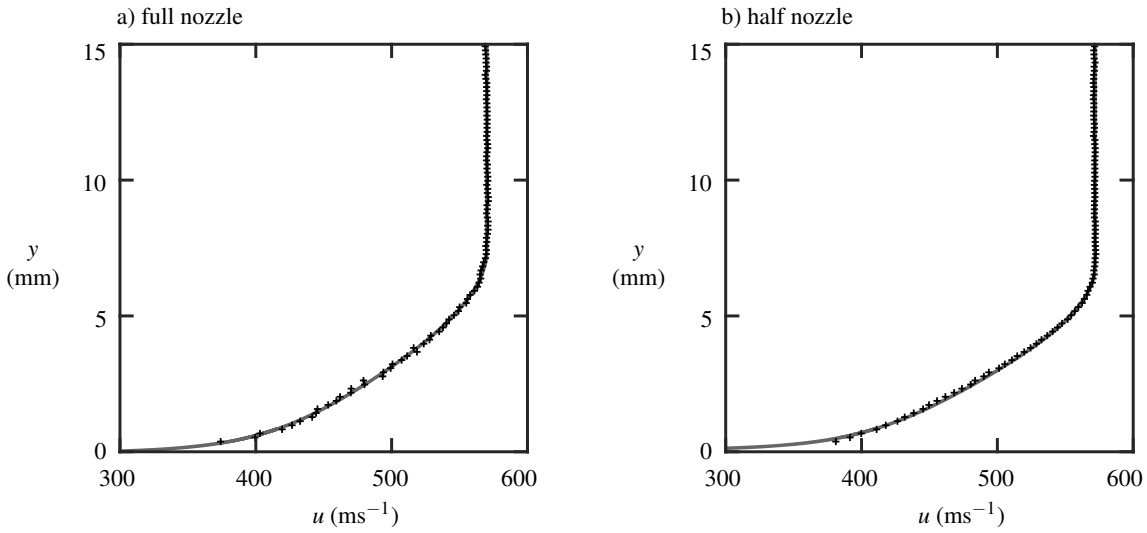

Figure 4.3: Laser Doppler velocimetry measurements of the floor boundary layer, performed at $x = 60$ mm and $z = 0$ mm for a) the full nozzle setup, and b) the half nozzle setup.

Table 4.1: Boundary-layer parameters for profiles measured on the tunnel centre span at $x = 60$ mm. These correspond to the profiles presented in figure 4.3.

	freestream velocity, u_∞ (ms^{-1})	boundary-layer thickness, δ (mm)	displacement thickness, δ_i^* (mm)	momentum thickness, θ_i (mm)	shape factor, H_i
full nozzle	570	7.5	0.97	0.73	1.34
half nozzle	572	6.8	0.92	0.68	1.36

thickness is typically estimated by measuring the velocity profile of the floor boundary layer at a single location on the tunnel's centre line.

These measurements are performed at $x = 60$ mm for both the full nozzle (figure 4.3a) and half nozzle (figure 4.3b) configurations. Informative incompressible boundary-layer parameters, calculated from these velocity profiles, are listed in table 4.1. The boundary layers have a shape factor of 1.34 and 1.36 for the full nozzle and half nozzle, respectively. These values are typical of an equilibrium turbulent profile. The boundary layers for both nozzle configurations are about 7–8 mm thick. For a wind tunnel of 114 mm width, these boundary layers are relatively thin and so the effects of confinement and three-dimensionality are expected to be fairly benign.

4.2 Additional tests for a validation study

A more involved characterisation process than the previous section is required for validation studies, however. The physical wind tunnel needs to be measured well enough for computations to accurately represent the true flow. Frameworks by which this might be consistently achieved [128–131] place an

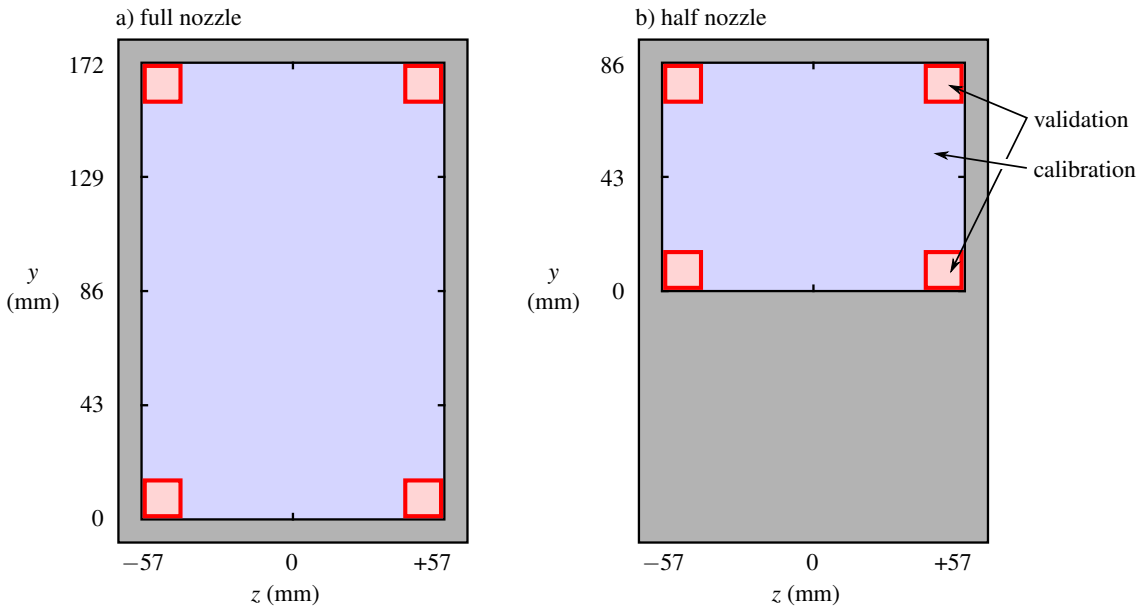

Figure 4.4: Division of regions for the purpose of validation (red) and of flow characterisation i.e. calibration of computations (blue) for a) the full nozzle setup, and b) the half nozzle setup.

emphasis on detailed measurements of the inflow conditions, the 'as-tested' geometry, and the tunnel boundary layers.

It is also important to bear in mind that calibration and validation are two separate stages in this type of study [128]. Calibration involves setting up simulations to represent the physical wind tunnel (using flow characterisation data) while validation tests the accuracy of physical models and numerical methods. The same data cannot be used for both purposes. The primary focus of the investigation are the corner boundary layers of the channel. Therefore, the data from a 15 mm $\times$15 mm region around each corner, marked in red on figure 4.4 is used for validation. Meanwhile, the remainder of the flow field, marked in blue in the figure, is treated as largely independent of the corner flows. These regions are measured for tunnel characterisation and for model calibration.

4.2.1 Quantitative characterisation data

As-installed tunnel geometry

The mesh generated by the US Air Force Research Laboratory for their RANS computations is based on a CAD model of the tunnel's design geometry. However, this may not correspond exactly to the true geometry, since the tunnel liners are removed and reinstalled between each set of experiments. Detailed measurements of the as-tested geometry, conducted using a dial test indicator, therefore provide an estimate of the CAD model's accuracy.

An early measurement of the tunnel floor with the full nozzle setup is presented in figure 4.5a. The floor and ceiling are designed with a 0.2° linear divergence in order to provide boundary-layer relief. However, the measured floor profile shows deviations of up to 0.3 mm (4% of a boundary-layer

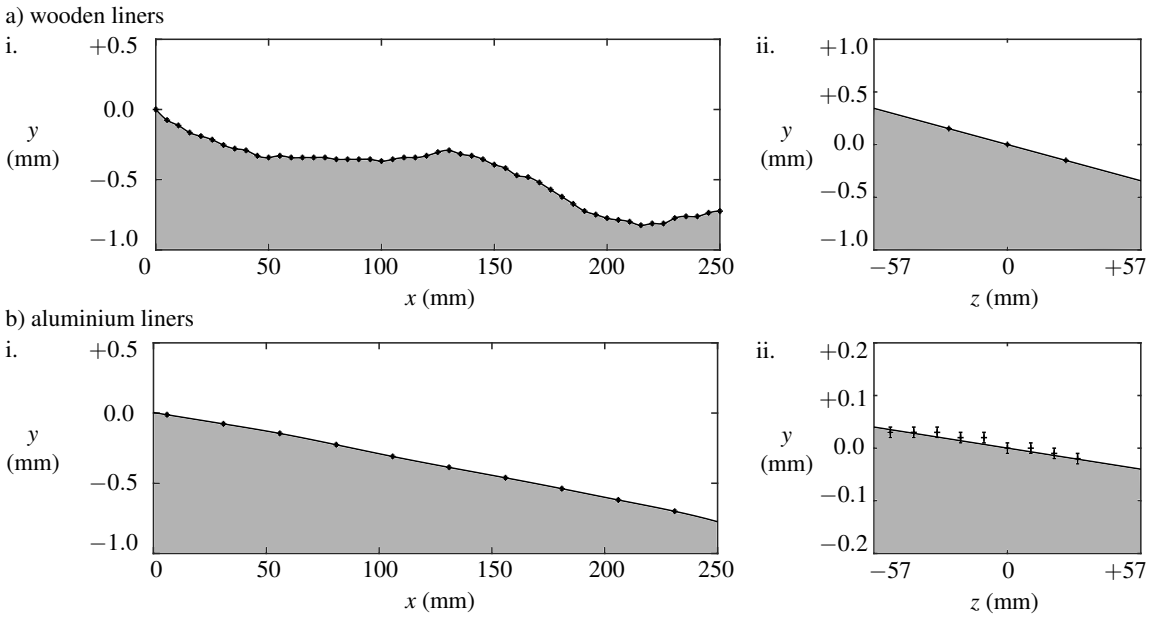

Figure 4.5: Measurement of tunnel floor geometry for the full nozzle setup, with a) wooden liners and b) aluminium liners. The measurements are performed i. in the streamwise direction at $z = 0$ mm, and ii. in the spanwise direction at $x = 0$ mm. For these measurements y is defined to be zero at $x = z = 0$ mm. The error bars are contained within the symbol size with $\Delta y = 0.01$ mm.

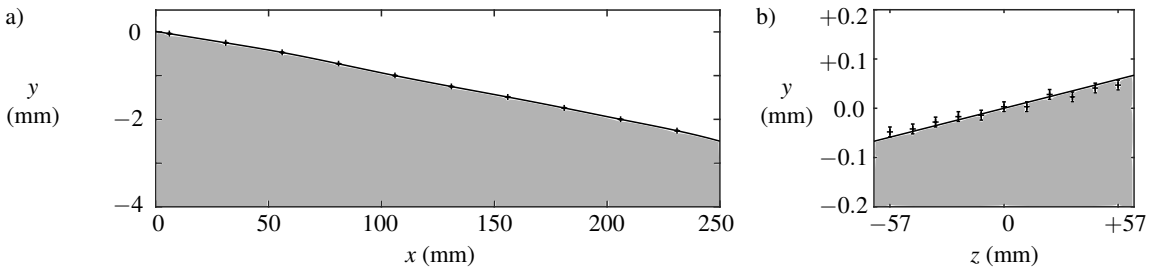

Figure 4.6: Measurement of tunnel floor geometry for the half nozzle setup. The measurements are performed a) in the streamwise direction at $z = 0$ mm, and b) in the spanwise direction at $x = 0$ mm. For these measurements y is defined to be zero at $x = z = 0$ mm. The error bars are contained within the symbol size with $\Delta y = 0.01$ mm.

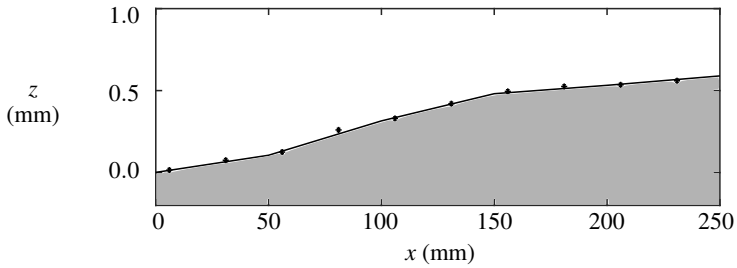

Figure 4.7: Measurement of tunnel sidewall geometry. The measurements are performed in the streamwise direction at the tunnel centre height ($y = 86$ mm). For these measurements z is defined to be zero at $x = 0$ mm. The error bars are contained within the symbol size with $\Delta z = 0.01$ mm.

thickness) from this linear design. Whilst such small variations are not expected to significantly impact the flow quality, the departures from the design geometry could affect any quantitative comparisons between experimental data and corresponding simulations. The observed 'bumpiness' was attributed to warping of the wooden liner material, and grooves burned into the surface by the LDV laser.

In order to reduce the size of these small imperfections, the liners were therefore replaced with aluminium equivalents. Figure 4.5b shows the equivalent floor profile with aluminium liners. The variations from the design geometry are now on the order of 0.05 mm, or 0.8% of a boundary-layer thickness. The tilt in the spanwise direction is $0.06 \pm 0.10^\circ$. The surface can therefore be considered flat in this direction, within experimental error.

Figure 4.6 shows the floor geometry with the half nozzle configuration, for which aluminium liners were used. The floor is designed to have a larger linear divergence angle of 0.6° for this setup, since the sidewalls do not provide boundary-layer relief and this channel has half the area of the full nozzle configuration. The deviations from the design geometry are approximately 0.07 mm, or 1% of a boundary-layer thickness. In the spanwise direction, the measured tilt is $0.07 \pm 0.10^\circ$. Therefore, within measurement accuracy, the floor can be treated as horizontal in this direction.

The measured tunnel sidewall geometry, presented in figure 4.7, has an approximately linear profile. Since there is no boundary-layer relief built into the sidewalls, this surface is expected to be parallel to the x-direction. The observed deviation from this design is $0.18 \pm 0.10^\circ$. The difference is attributed to the variations in torque applied when the sidewalls are bolted to the tunnel structure. The manual nature of this process implies that the sidewall angle changes each time these tunnel doors are closed. As a result, the measured sidewall angle is treated as an order-of-magnitude estimate for typical run-to-run variations.

The as-installed geometry measurements therefore complement the CAD model to generate an accurate mesh for simulations. These measurements also provide estimates for the size of typical deviations from the design geometry. The estimates can, in turn, be used to perform a sensitivity analysis on how the computed flow responds to such deviations.

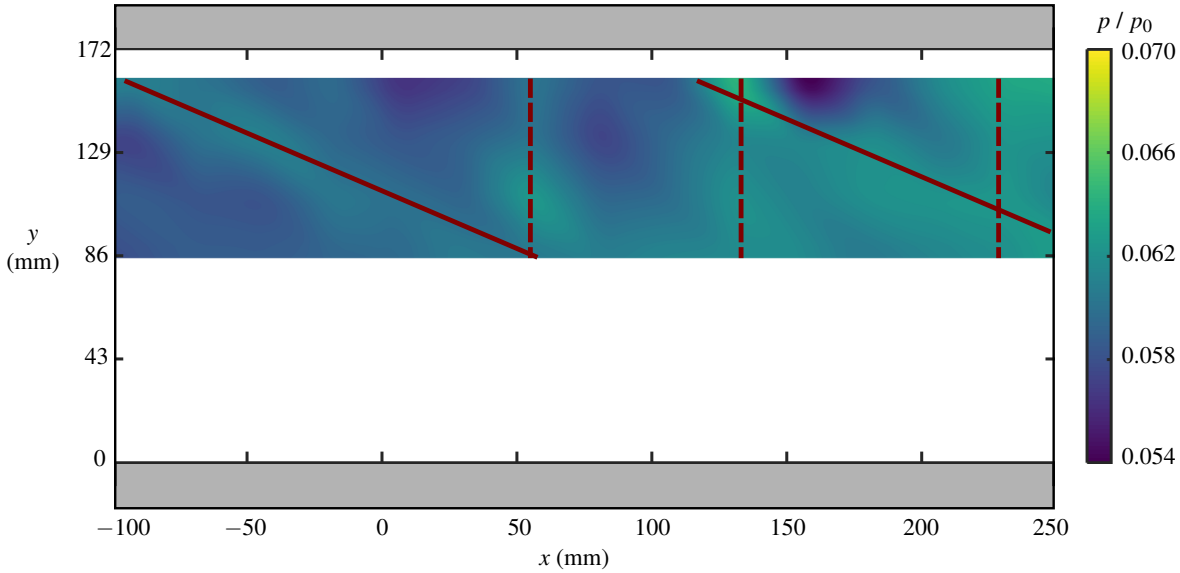

Figure 4.8: Static pressure distribution measured using taps located in the tunnel sidewall for the full nozzle setup. The red lines correspond to high-pressure areas, which take the form of oblique (solid) and vertical (dashed) regions.

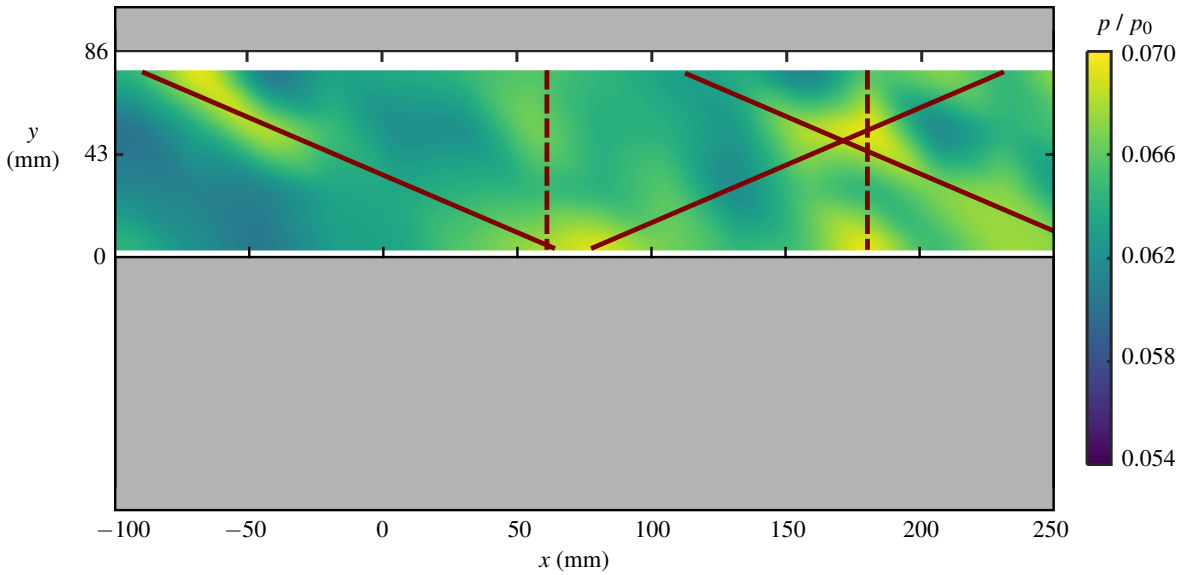

Figure 4.9: Static pressure distribution measured using taps located in the tunnel sidewall for the half nozzle setup. The red lines correspond to high-pressure areas, which take the form of oblique (solid) and vertical (dashed) regions.

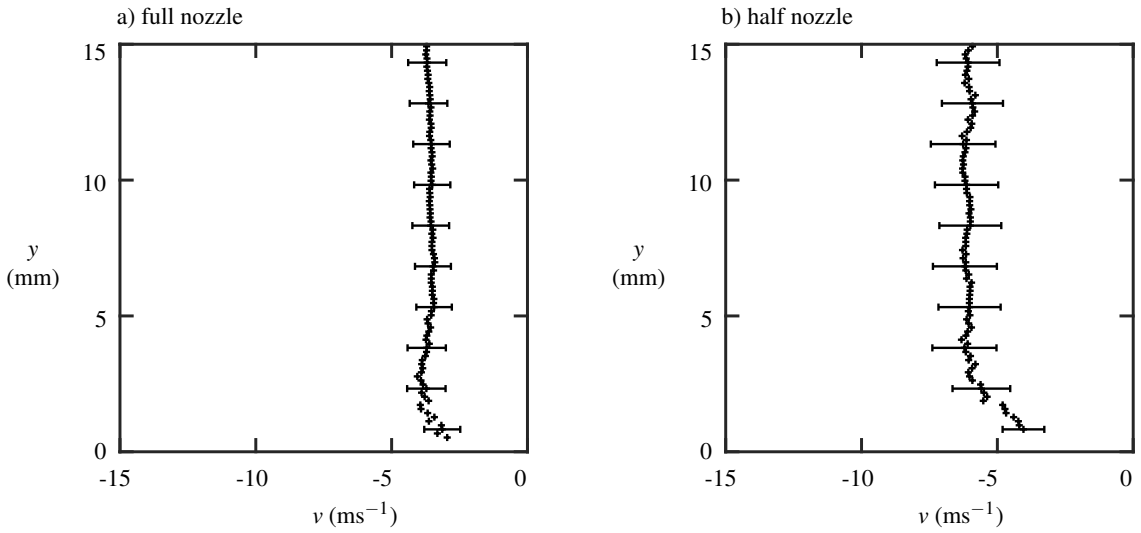

Figure 4.10: Measurements of the vertical velocity, v, in the floor boundary layer, performed at $x = 60$ mm and $z = 0$ mm for a) the full nozzle setup, and b) the half nozzle setup.

Sidewall static pressures

Waves are visible in the schlieren images of the flow (figures 4.2a and 4.2b). The strength of these waves is quantified using steady-state surface pressure measurements on the tunnel sidewall.

For the full nozzle setup, figure 4.8 shows that the wall static pressure varies by 4% throughout the test section. These deviations correspond to departures in Mach number of 0.02 from a mean value of 2.48. The oblique features in the plots correspond to weak waves generated from the tunnel floor and ceiling. The wave angle is approximately $24 \pm 2^\circ$ which, for a Mach 2.5 flow, matches closely to the Mach angle (23.6°). These waves can therefore be treated as isentropic.

Meanwhile, the vertical columns are likely due to spanwise-travelling waves generated at the sidewalls. The angle of these waves cannot be extracted from the data in figure 4.8. However, the fact that the Mach number varies by less than 1% in the test section does imply that these waves are also weak.

The equivalent data for the half nozzle configuration are presented in figure 4.9. Here, the static pressure variation across the test section is approximately 7%. The wall pressures correspond to a mean Mach number of 2.44, and associated deviations from this value of 0.03. As for the full nozzle, the pressure distribution exhibits both oblique and vertical features. These correspond to waves generated by imperfections in the tunnel floor and ceiling (oblique), and in the sidewalls (vertical). A comparison of the static pressures and the schlieren images indicates that these waves are stronger than for the full setup, though they are still weak in nature.

Velocities in the test section

In order to calibrate computational models, it is not sufficient to use only the one set of boundary-layer profile data shown in figure 4.3. One example of useful additional data is the vertical component of

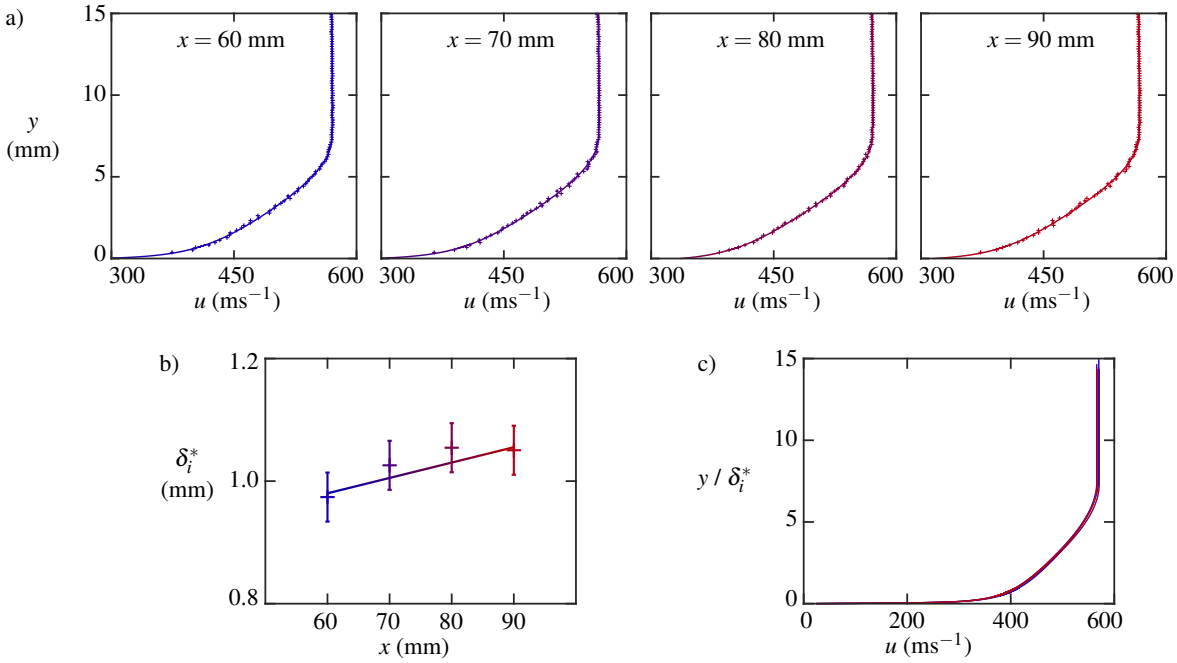

Figure 4.11: a) Floor boundary-layer profiles measured using LDV for the full nozzle setup. These measurements are performed at different streamwise locations on the centre span. b) Variation in displacement thickness with streamwise direction. c) Collapse of profile shape when y is non-dimensionalised by incompressible displacement thickness, δ_i^*.

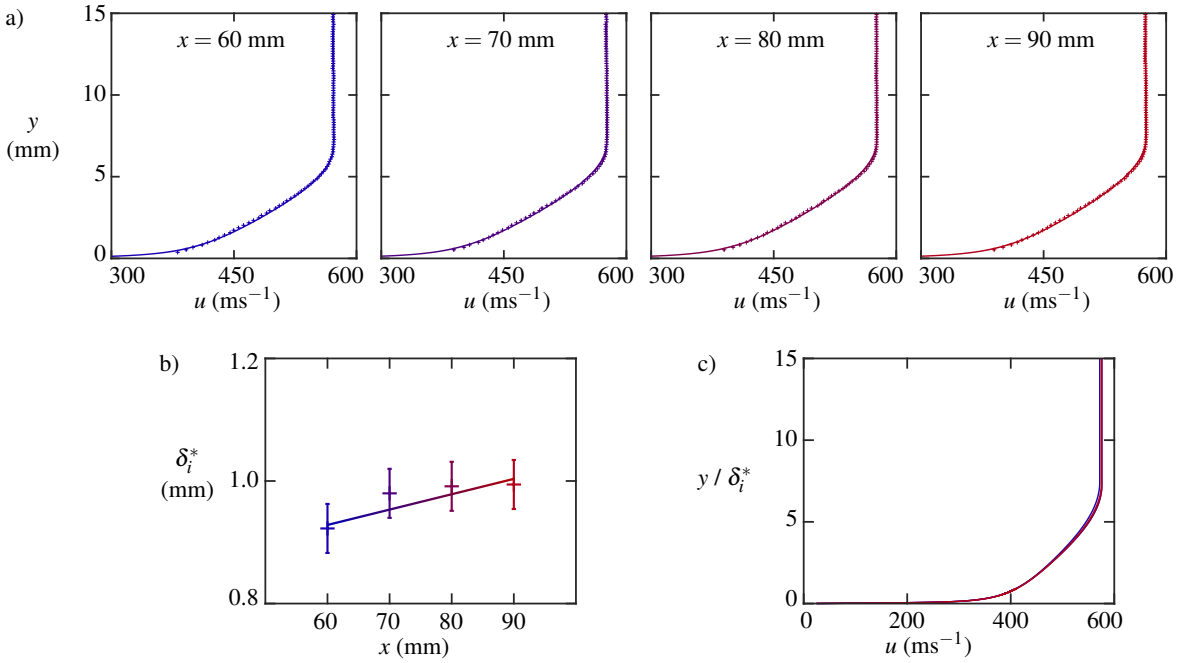

Figure 4.12: a) Floor boundary-layer profiles measured using LDV for the half nozzle setup. These measurements are performed at different streamwise locations on the centre span. b) Variation in displacement thickness with streamwise direction. c) Collapse of profile shape when y is non-dimensionalised by incompressible displacement thickness, δ_i^*.

velocity, which is shown in figure 4.10 for the centre-span floor boundary layer at $x = 60$ mm. These measurements have a much larger intrinsic experimental uncertainty ($\approx 14\%$) than the streamwise velocity (2%) but still provide quantitative flow characterisation data. The vertical velocity in the freestream is -4 ms^{-1} for the full nozzle and -6 ms^{-1} for the half nozzle setup, which corresponds to approximately 1% of the equivalent streamwise velocity. This non-zero downwards velocity is due to the slight divergence of the floor and ceiling liner blocks for boundary-layer relief, which is discussed in section 3.1. The divergence angle is larger for the half nozzle setup than for the full nozzle, which results in a more negative vertical freestream velocity. Figure 4.10 also shows that the magnitude of vertical velocity decreases to zero within the boundary layer, as required to satisfy the no-slip condition at the floor.

In considering the vertical velocity component, it is important to bear in mind that, due to the low velocity magnitudes typically measured, this quantity is expected to be particularly sensitive to the divergence angle of the liner blocks and to the precise alignment of the LDV apparatus relative to the tunnel. As a result, such measurements are used only for flow characterisation purposes and are not used for direct comparison with corner flow computations in the validation stage of the process.

In addition, the streamwise velocity profiles for the centre-span floor boundary layer are measured at several other streamwise locations to characterise its growth. These measurements are shown in figures 4.11a and 4.12a for the full nozzle and the half nozzle setups, respectively. The boundary-layer parameters corresponding to these profiles are listed in table 4.2. The profile shapes are typical for a fully-developed turbulent equilibrium boundary layer, with an incompressible shape factor $H_i \approx 1.3 - 1.4$.

The growth of this boundary layer with streamwise co-ordinate is shown by the increasing displacement thickness in this direction (figures 4.11b and 4.12b). When normalised for boundary-layer thickness, as performed in figures 4.11c and 4.12c, there is a collapse of the profiles. The constant profile shape in this region suggests that the waves identified in figures 4.8 and 4.9 are not strong enough to disturb the boundary layers growing along the wind tunnel walls.

In addition to the centre-span measurements, velocity data is also obtained for the boundary layer away from the symmetry plane, and on the sidewalls. The measured velocity distribution is shown in figure 4.13, for the two nozzle configurations. The core flow represents a large fraction of the channel area for both setups. The floor and ceiling boundary layers have approximately constant thickness across the tunnel span, but there is a significant variation in sidewall boundary-layer thickness. These variations are related to the nozzle geometry, and form the subject of Chapter 5.

The additional velocity measurements represent a more complete characterisation of the overall flow field. These measurements consist of 128 boundary-layer profiles for the full nozzle setup and 81 profiles with the half nozzle, forming a valuable data set for quantitative comparison with computations of the wind tunnel flow.

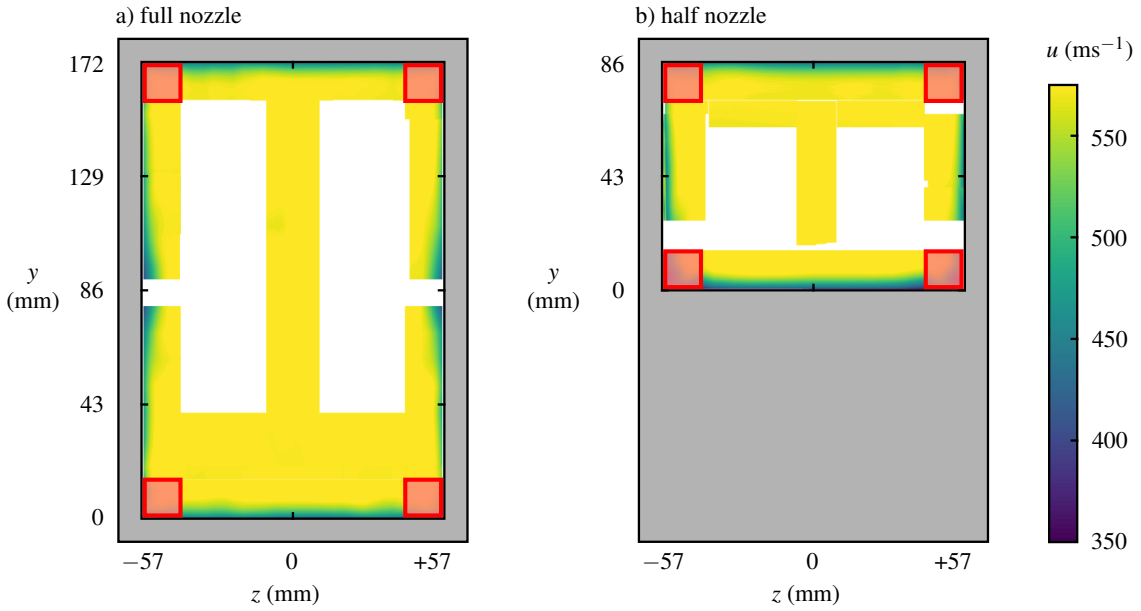

Figure 4.13: Streamwise flow velocities measured across the tunnel cross-section at $x = 120$ mm. Corner regions, marked by red boxes, are excluded for calibration purposes. Regions with an insufficient number density of seeding particles for high-quality velocity measurements are in white.

Table 4.2: Incompressible boundary-layer parameters, measured along the tunnel centre line. These correspond to experimental profiles presented in figure 4.11a (full nozzle) and figure 4.12a (half nozzle).

	full nozzle				half nozzle			
x (mm)	δ (mm)	δ_i^* (mm)	θ_i (mm)	H_i	δ (mm)	δ_i^* (mm)	θ_i (mm)	H_i
60	7.49	0.97	0.73	1.34	6.80	0.92	0.68	1.36
70	7.59	1.03	0.76	1.35	7.03	0.98	0.72	1.37
80	7.70	1.05	0.78	1.36	7.14	0.99	0.73	1.36
90	7.71	1.05	0.78	1.35	7.26	0.99	0.73	1.36

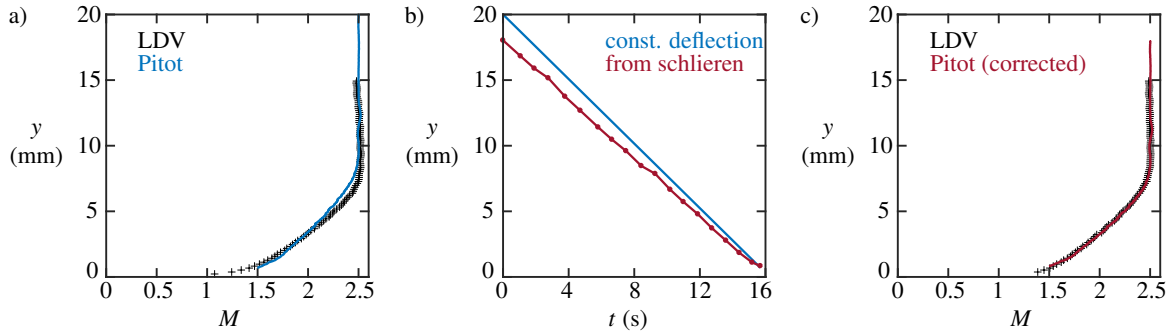

Figure 4.14: a) Initial comparison of floor boundary-layer profiles measured at $x = 120$ mm and $z = 0$ mm between Pitot tube and LDV measurements. b) Probe position over course of traverse assuming constant deflection under aerodynamic load and measured from schlieren images. c) Comparison of the floor boundary-layer profile between Pitot tube and LDV measurements with corrected probe position.

4.2.2 Efforts to identify systematic errors

During the flow characterisation process, random errors in the data can be minimised by performing a greater number of relevant measurements. However, in order to use data for model calibration, it is also essential to identify and account for any systematic errors. These types of error tend to bias the measurements, and are generally not simple to detect or resolve. An identification of potential systematic errors is performed by implementing redundant tests, assessing the reliability of methods, and evaluating the repeatability of data.

Note that the following tests are performed only on the full nozzle setup, for two main reasons. First, these procedures are time-consuming in nature. Secondly, any systematic errors are expected to affect both configurations similarly, and should therefore be captured by these tests.

Tests of LDV data using Pitot probe measurements

The validity of LDV flow velocity data is tested through direct comparison with equivalent measurements using a Pitot probe traverse introduced in section 3.2.3. The Pitot probe experiments are conducted with no seeding rake (or seeding particles), and so this comparison also identifies any effects of the seeding apparatus on the flow.

An initial comparison of Pitot probe data with LDV measurements, presented in figure 4.14a, shows reasonable agreement. However, there is a discrepancy in Mach number towards the outer edge of the flow. This small disagreement is not problematic for purely experimental studies where, often, only the difference between measured profiles is of importance. However, errors of the magnitude seen here would hinder calibration of computational models.

To avoid any unwanted interference the supporting structure of the Pitot probe is kept slender (figure 3.11a). However, this makes the setup mechanically fragile. When the wind tunnel is turned on, the probe therefore bends slightly in the y-direction under aerodynamic load. As described in section 3.2.3, any constant probe deflection is usually accounted for by detecting electrical contact

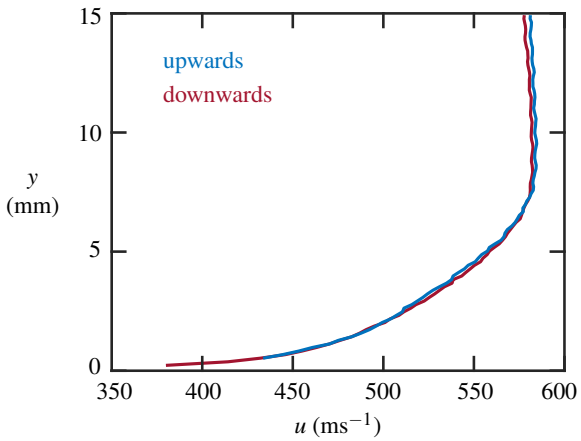

Figure 4.15: Floor boundary-layer profiles measured using LDV at $x = 120$ mm and $z = 0$ mm, comparing velocity measurements between traverses performed in the upwards and downwards directions.

between the probe and the wind tunnel floor at the end of the traverse. However, rather than assuming the deflection to be constant throughout a traverse across the boundary layer, the discrepancies in figure 4.14a demand that a more careful approach is taken instead. The position of the probe is accurately determined using schlieren images of the tunnel captured throughout the wind tunnel run. Figure 4.14b shows the variation with time of the probe position assuming constant deflection and of the probe position calibrated from schlieren images. The different trajectories of the two lines in the figure indicates that the Pitot probe deflection varies over the course of the run.

Figure 4.14c presents the revised boundary-layer profile with wall-normal co-ordinate, y, now calibrated using the schlieren images. This demonstrates better agreement with the LDV data, and any discrepancies are within intrinsic experimental error. Therefore the effect of the seeding rake in the settling chamber and any influence due to the seeding particles is smaller than measurement resolution, and so LDV data can be used for model calibration with greater confidence.

Tests of the traverse system

In order to identify systematic errors with the LDV–traverse measurement process, floor boundary-layer traverses are performed in both the upwards and downwards direction at the same measurement station. A comparison of the resulting velocity data is presented in figure 4.15. The profiles collapse well and calculated boundary-layer parameters show excellent agreement. The differences are all within the estimated uncertainty in these parameters (5%).

In addition, the velocity distribution across the tunnel cross-section is typically measured by interpolating between several floor-normal traverses in the y-direction. An example of this process is shown in figure 4.16a-i. In order to ensure that these measurements are valid, traverses are also performed over the same region in the z-direction (figure 4.16a-ii). A visual representation of the two interpolated flow fields appears to be identical.

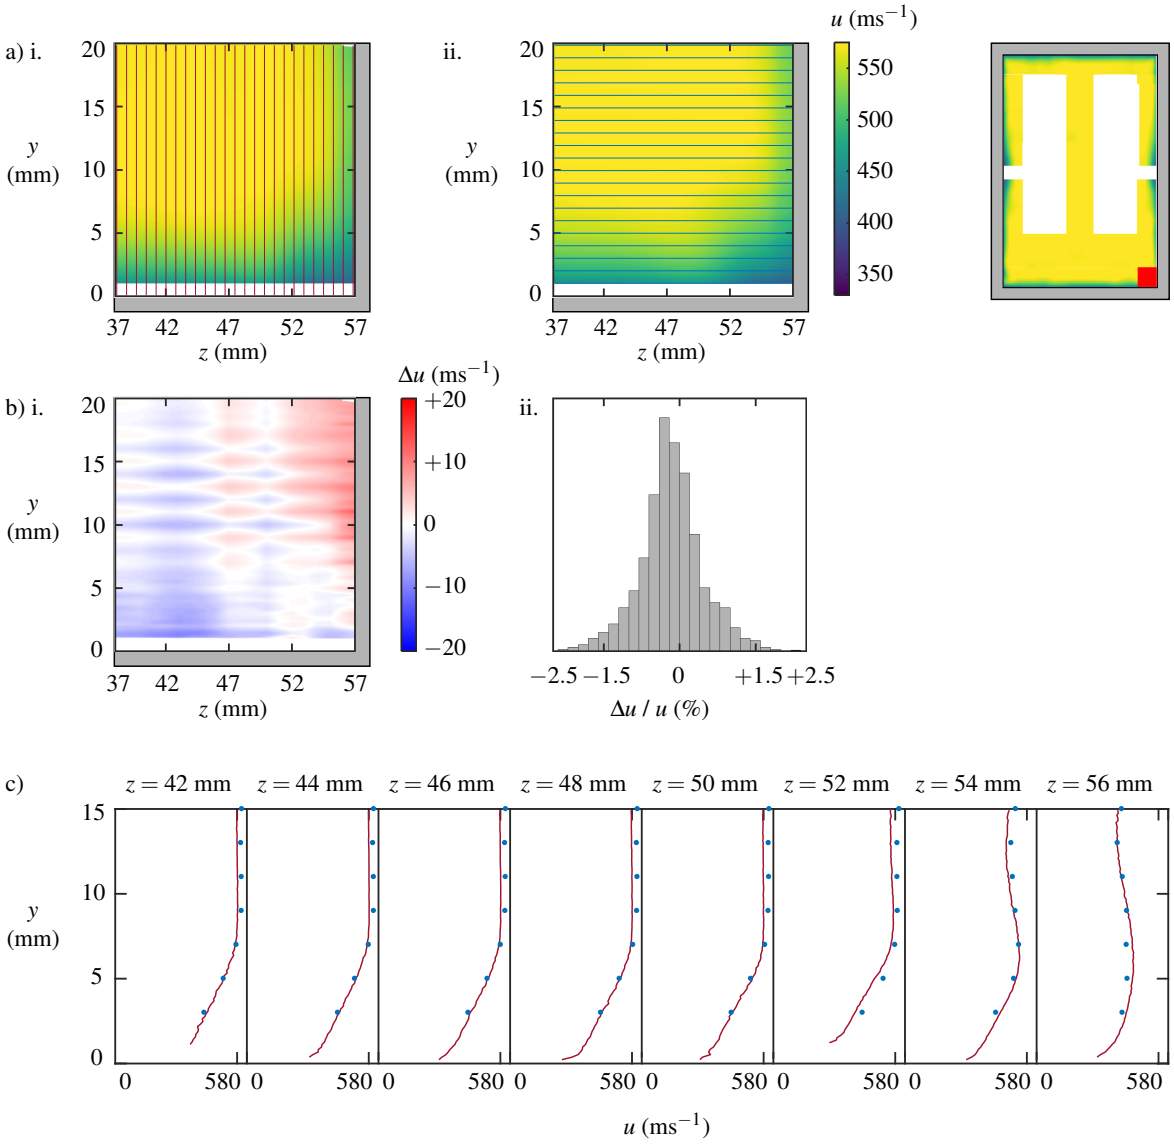

Figure 4.16: Measurement of the streamwise velocity distribution in the corner region at $x = 120$ mm. Interpolation of velocity field from data obtained using a) i. vertical traverses (red lines) and ii. horizontal traverses (blue lines). b) i. Difference between interpolated velocity distributions using both data sets, and ii. corresponding histogram of percentage velocity differences. c) Boundary-layer profiles measured using selected vertical traverses from a-i (red lines). The blue dots mark the velocities measured using selected horizontal traverses from a-ii, at the intersection points with the vertical traverse under consideration. The solid red box in the top-right inset marks the measurement region.

Table 4.3: Boundary-layer parameters for LDV measurements on the tunnel centre span at $x = 120$ mm, comparing profiles between traverses performed in the upwards and downwards directions. These correspond to the velocity data presented in figure 4.15.

	δ (mm)	δ_i^* (mm)	θ_i (mm)	H_i
upwards	7.91	0.91	0.69	1.32
downwards	8.23	0.95	0.73	1.31

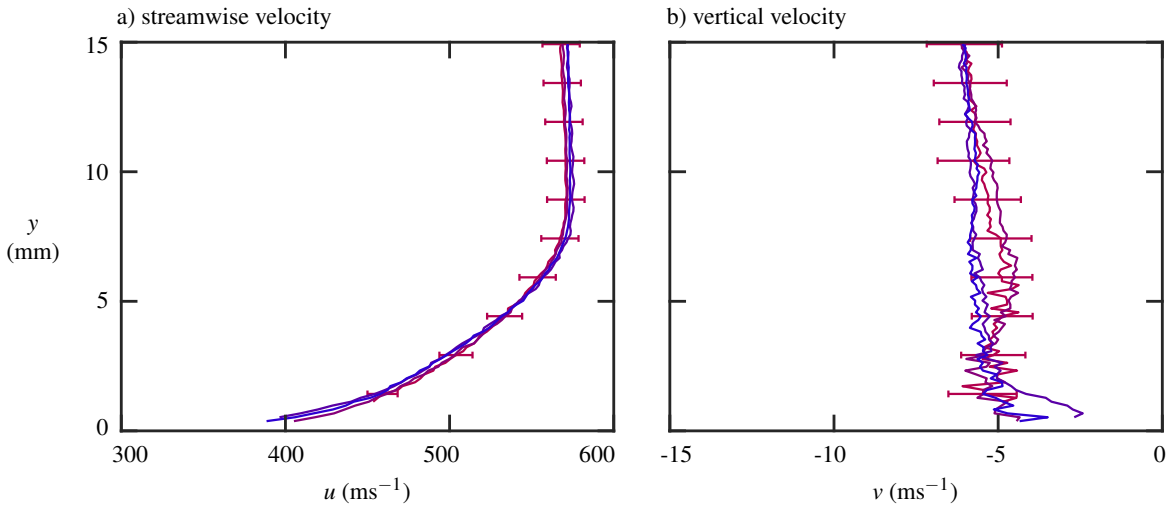

Figure 4.17: Floor boundary-layer profiles measured using LDV at $x = 120$ mm and $z = 0$ mm, comparing measurements between four different tunnel installations: a) streamwise velocity, u, and b) vertical velocity, v. Error bars are marked on every ten data points for a single profile for clarity.

A more quantitative assessment is performed by evaluating the difference between interpolated velocities, as shown in figure 4.16b. Similarly, the measured velocities are compared at the intersection points between the two sets of traverses (figure 4.16c). These figures both show the difference between interpolated velocity distributions is less than 10 ms^{-1}, i.e. 1.5% of the freestream velocity.

It is therefore possible to conclude from these tests that any systematic error associated with the LDV–traverse system are no larger than the intrinsic experimental error associated with measurements.

Tests of repeatability

As described in section 3.1, the wind tunnel setup consists of liners within the tunnel structure, which are removed between each set of experiments. Phenomenological investigations are able to gain useful information by comparing data obtained within a single set of experiments. In order to use the data to calibrate simulations, however, it is important to assess the repeatability of the flow field across different installations.

Figure 4.17a presents a floor boundary-layer traverse at the same location conducted during four separate tunnel installations several months apart. The excellent agreement between the profiles demonstrates good repeatability of the flow field. More quantitative indications of the repeatability

Table 4.4: Boundary-layer parameters for LDV measurements on the tunnel centre span at $x = 120$ mm, comparing profiles between four different tunnel installations. These correspond to the data presented in figure 4.17a.

	u_∞ (ms^{-1})	δ (mm)	δ_i^* (mm)	θ_i (mm)	H_i
October 2017	572.8	8.33	0.94	0.72	1.30
February 2018	573.9	8.23	0.95	0.73	1.31
May 2018	570.7	8.50	0.89	0.69	1.29
June 2018	569.7	8.32	0.90	0.70	1.29

are the calculated boundary-layer parameters from these measurements, listed in table 4.4. These quantities are consistent with a high level of repeatability, exhibiting small variation in the freestream velocity (570 ms^{-1} to 574 ms^{-1}) and a narrow range of measured boundary-layer thicknesses (8.2 mm to 8.5 mm). The relatively small variations which do exist are thought to be primarily due to differences in the stagnation temperature between different tunnel installations, discussed in section 3.1.

The validation process is conducted primarily using the streamwise velocity component. Nevertheless, the equivalent boundary-layer measurements of the vertical velocity, v, are shown in figure 4.17b for reference. The variations between installations appear to be proportionally larger than for the streamwise velocity. However, the profiles do agree with one another to within experimental uncertainty. This is a direct consequence of the significantly larger intrinsic error in measurements of v (14%) than for u (2%). Note further that, due to the lower velocity magnitude, the vertical component of velocity is expected to be particularly sensitive to any variations in the divergence angle of the liner blocks and in the alignment of the LDV apparatus.

Tests of flow symmetry

A common assumption of wind tunnel testing is that the flow matches the symmetry of the design geometry. In other words, the flow field in this full nozzle setup might be expected to exhibit top–bottom and left–right symmetry. Typical scientific studies are insensitive to any subtle asymmetries, but quantifying how symmetric the flow is forms an essential part of the characterisation process.

It is not possible to take measurements using laser Doppler velocimetry in the upper half of the tunnel, due to the distribution of seeding particles and the tilt of laser optics. The scope to directly test top–bottom symmetry is therefore somewhat limited. As an alternative, the nozzle blocks and liners are swapped from the top to the bottom half of the tunnel, and vice versa. This test enables any top–bottom asymmetries introduced downstream of the contraction to be identified.

In order to ensure that reversing these liners does not affect the flow itself, velocities at equivalent locations before and after the swap are compared. Due to the seeding particle distribution, this could only be performed for a small sub-region of the core flow. Figure 4.18a displays very similar (uniform) velocity distributions in this area before and after the swap. The difference between these

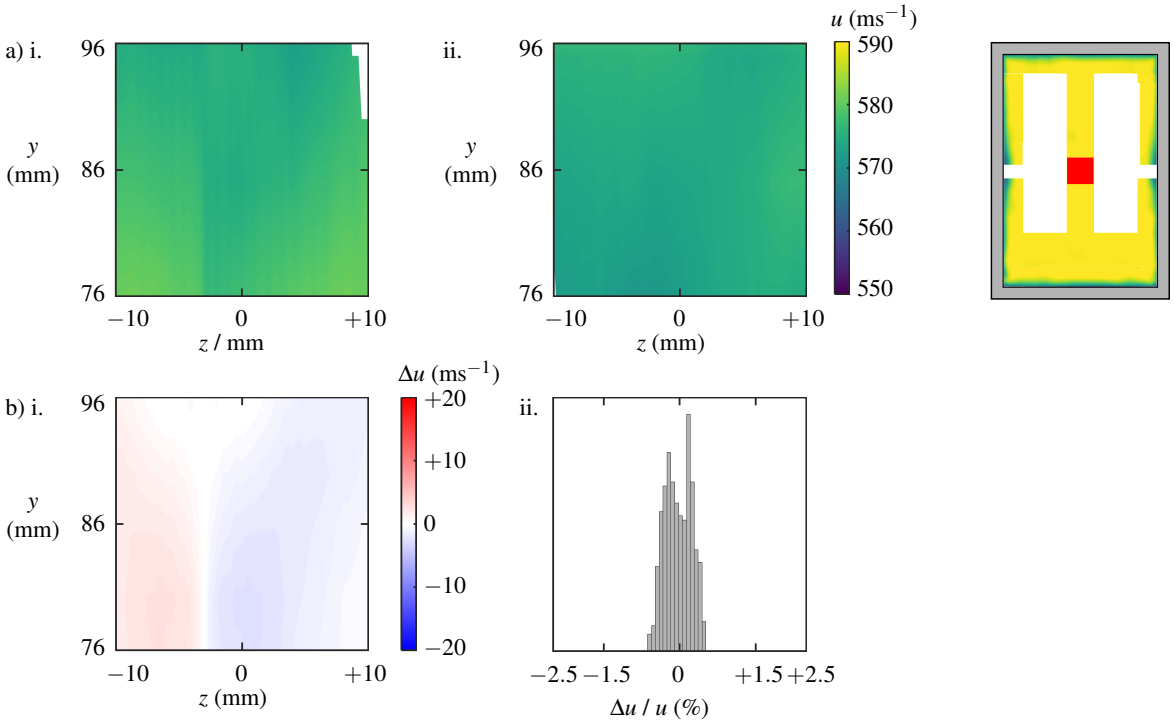

Figure 4.18: Core flow in overlapping region at $x = 120$ mm a) i. before and ii. after nozzle blocks and liners have been top–bottom swapped. b) i. Difference between equivalent velocity measurements, and ii. histogram of percentage velocity differences. The solid red box marks the measurement region.

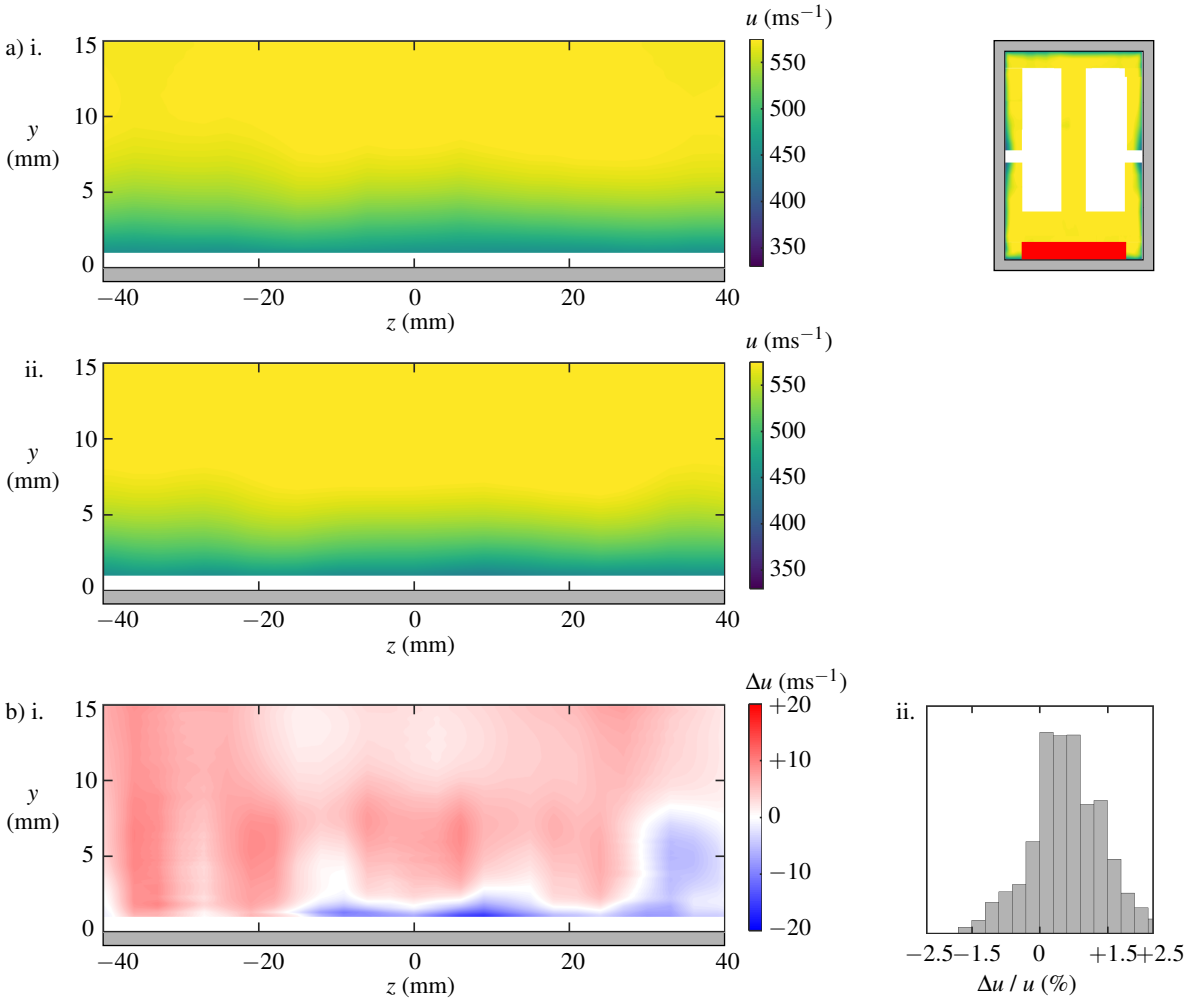

Figure 4.19: Floor boundary layer measured across tunnel span at $x = 120$ mm a) i. before and ii. after nozzle blocks and liners have been top–bottom swapped. b) i. Difference between velocity distributions, and ii. histogram of percentage velocity differences. The solid red box marks the measurement region.

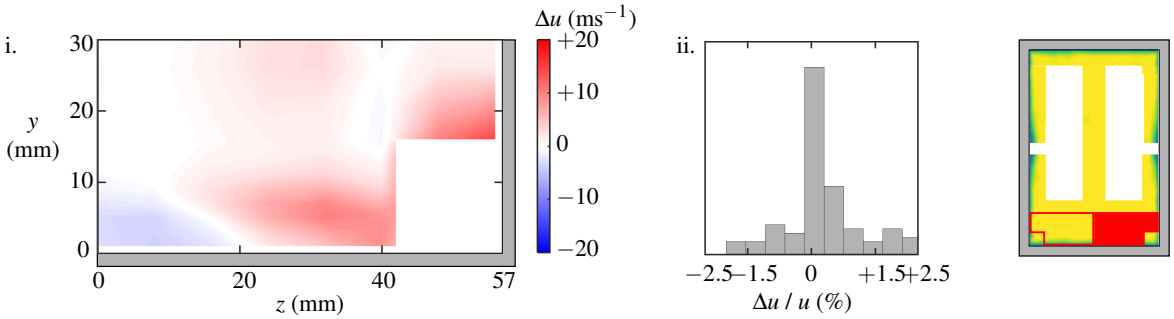

Figure 4.20: Difference in streamwise velocity between the left- and right-hand halves of tunnel cross-section measurements at $x = 120$ mm (from figure 4.13), presented in i. spatial and ii. histogram form. The difference in velocity is calculated between each point in the solid red box and the corresponding point in the empty red box. The bottom corner regions are used for validation and so excluded from this analysis, as per figure 4.4.

two distributions is presented in figure 4.18b. This figure shows that the velocities at equivalent locations in the core flow change by less than 0.7% on swapping the liners.

Figure 4.19a presents the velocity distribution in the floor boundary layer before and after the swap is performed. Qualitatively, the velocities do not appear to have changed substantially. The difference between the two distributions is plotted in figure 4.19b. In this figure, velocity measurements at equivalent locations differ by at most 5 ms^{-1}, or 0.8% of the freestream velocity.

In order to assess the left–right symmetry of the tunnel, all measurements are performed on both sides of the tunnel centre line. Figure 4.13, for example, presents the velocity field across the entire tunnel span. The velocity difference between equivalent locations in the left- and right-halves of the tunnel is shown in figure 4.20. Departures from a perfect left–right symmetry are no larger than 10 ms^{-1}, or 1.6% of the freestream velocity.

Incidentally, these checks of left–right symmetry have been useful in identifying seal failures, when air is sucked into the tunnel from the surroundings. In this scenario, the bolting mechanism of the tunnel doors is such that air is not sucked into the test section symmetrically. Seal failures can therefore be identified by checks of left–right symmetry. As an example, figure 4.21a-i presents an initial measurement of the floor boundary layer across the span of the tunnel. It is clear that there is a very different flow field between the left- and right-halves of the tunnel. This difference is highlighted by the marked contrast between individual boundary-layer profiles measured at equivalent locations, $z = -52$ mm and $+52$ mm, and shown in figure 4.21a-ii. After replacing the wind tunnel seals, the flow appears to be left–right symmetric (figure 4.21b-i) and equivalent boundary-layer traverses, shown in figure 4.21b-ii, are approximately coincident. Thus this test of symmetry can be used to identify leaks, which are often difficult to detect but which can impact large areas of the flow field. To avoid problems associated with leaks, the seals are also inspected visually every time the tunnel doors are unbolted, i.e. at least every five wind tunnel runs.

Note that the equivalent velocities in both sets of symmetry test data are very similar but not identical. This is a characteristic feature of an ‘imperfect’ physical wind tunnel. The corresponding

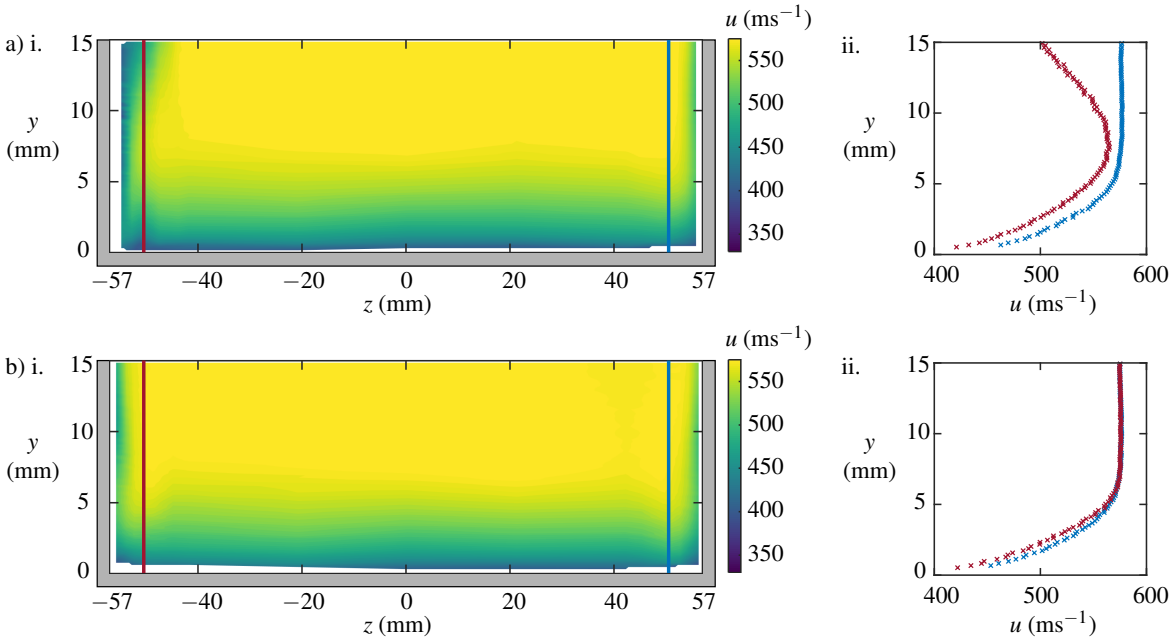

Figure 4.21: Floor boundary layer at $x = 120$ mm with the tunnel a) poorly sealed and b) correctly sealed: i. velocity distribution across tunnel span and ii. individual boundary-layer profiles at $z = -52$ mm (red) and $z = +52$ mm (blue).

differences can be used to more reliably calibrate computations, and to estimate the error associated with these measurements.

4.2.3 Comparison between computation and experiment

Presentation of data

In science discovery investigations, measurements are often presented in such a way as to highlight features of interest, perhaps at the expense of other aspects of the flow. Examples of this include the choice of colour map for a contour plot or the knife-edge orientation for schlieren techniques. The presentation of, and comparison between, data sets in a validation experiment needs to be more objective than this, however. The colour map for data presentation is chosen to be perceptually uniform with appropriate ranges to accurately represent the information. Schlieren images of the flow are captured using various knife-edge orientations (figures 4.22a and 4.22b), in combination with a shadowgraph (figure 4.22c), to obtain a more complete picture of the flow field.

Comparison of equivalent quantities

An important aspect of calibrating a computational model is the comparison with experimental data. During this process, it is important to compare measurements with computations in a like-for-like manner.

One example of this is the processing of sidewall static pressures. This quantity, extracted from the computational data provided by the US Air Force Research Laboratory, is shown for the full

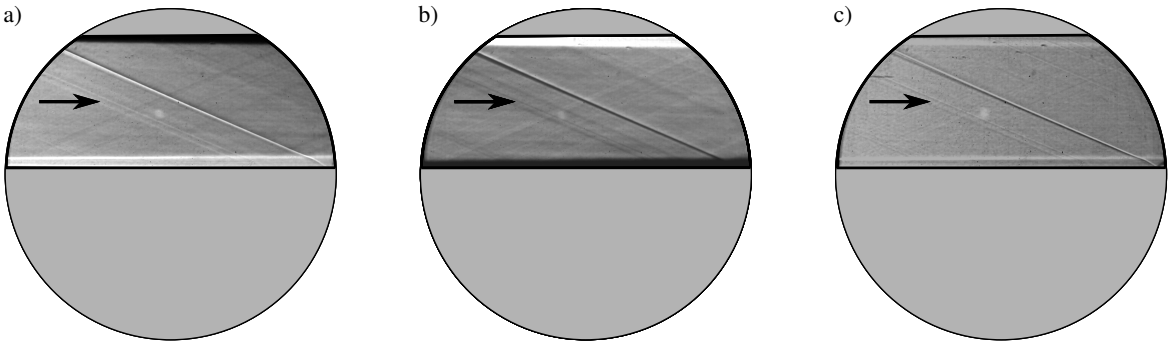

Figure 4.22: Visualisation of density gradients in the flow with the half nozzle setup: a) schlieren with horizontal knife-edge on bottom; b) schlieren with horizontal knife-edge on top; c) shadowgraph, i.e. with no knife-edge.

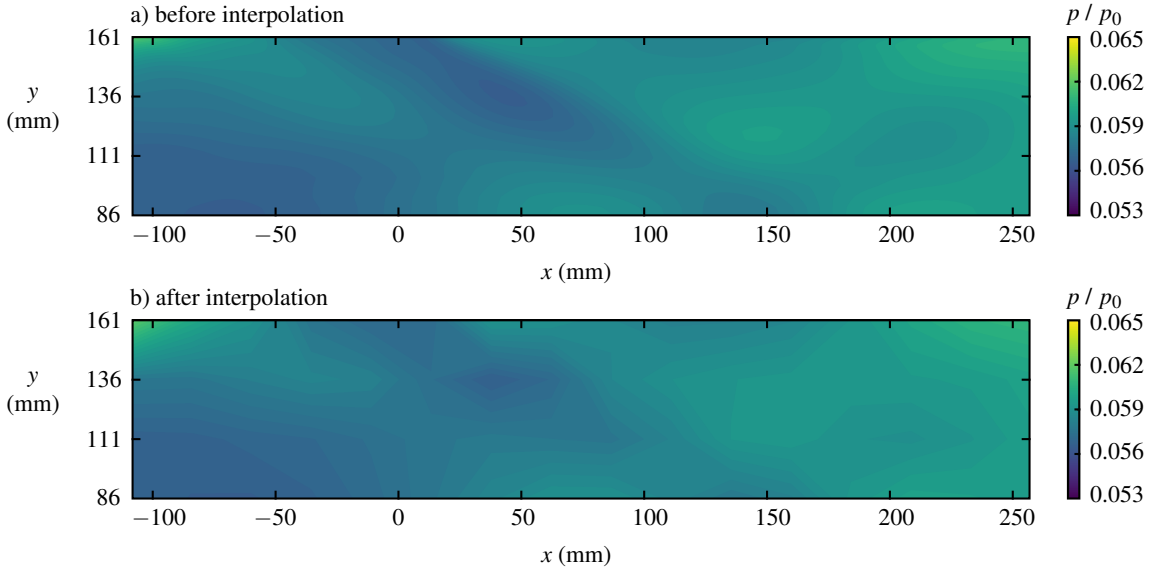

Figure 4.23: Computed static pressure for the full nozzle setup on the tunnel sidewall a) prior to, and b) after interpolating onto measurement grid from experiment.

nozzle in figure 4.23a. The computational solutions are interpolated onto the measurement locations from experiment before any comparisons are made. The interpolated pressures, given in figure 4.23b, exhibit some small differences to the more highly-resolved original distribution. Therefore, the interpolation process prevents any perceived differences caused by the difference in spatial resolution, which could in turn affect interpretation of physical results.

Similarly, boundary-layer profiles from computations also need to be treated carefully when they are compared with experimental data. LDV measurements are typically reported as point values, but they are, in fact, an average velocity of particles which pass through the finite probe volume. During the direct comparison of velocity data, the computational solutions are averaged over a representative probe volume (figure 3.24). Following the discussion in section 3.2.4, this is likely to be particularly relevant at locations within 1.6 mm of the sidewalls and within 0.35 mm of the tunnel floor or ceiling.

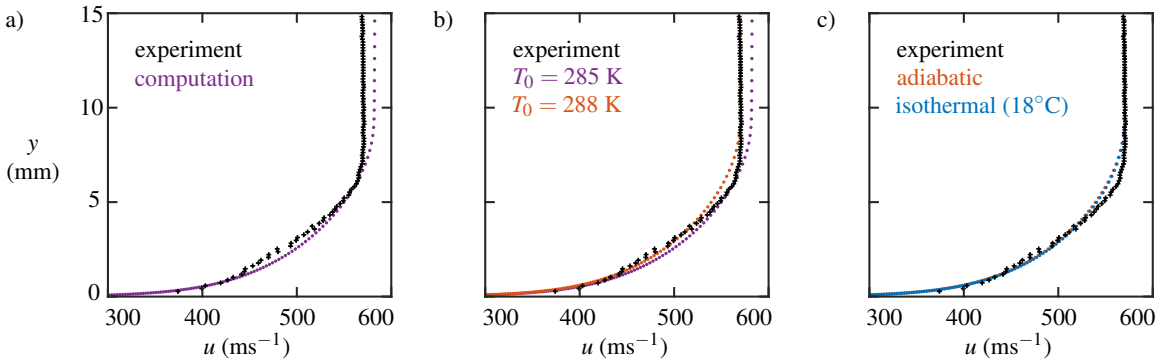

Figure 4.24: Floor boundary-layer profiles computed with the two-dimensional grid on the tunnel centre span at $x = 60$ mm, showing the calibration process: a) an initial comparison, b) the effect of stagnation temperature, and c) the effect of wall heat-transfer condition.

This averaging procedure prevents introducing artificial error, which could be as large as 3% and which could affect the interpretation of any direct comparisons between experimental and computational data.

Calibration process (two-dimensional grid)

In order to calibrate the RANS simulations conducted at the US Air Force Research Laboratory, preliminary solutions were first computed on a two-dimensional grid representing the centre span of the tunnel.[1] These solutions are much cheaper to compute than their three-dimensional counterparts, but are expected to capture much of the key physics required for calibration.

The first simulations were conducted on this two-dimensional grid with adiabatic walls and nominal inflow stagnation conditions of 308 kPa and 285 K. The boundary-layer profile is compared to corresponding LDV measurements in figure 4.24a, which shows that the freestream velocity at $x = 60$ mm is not accurately reproduced by the computations.

The size of this discrepancy cannot be accounted for purely by three-dimensional effects, such as the growth of the sidewall boundary layers. It was therefore necessary to run computations at a different stagnation temperature in order to match the measured freestream velocity. The best agreement was found for inflow stagnation temperature set to 288 K, as shown in figure 4.24b. Note that this value falls within the uncertainty bands for stagnation temperature, i.e. 280 – 290 K. Thus, it seems as though a simple arithmetic mean of the typical temperature range is insufficient to estimate a representative stagnation temperature and so a value of 288 K is used in all future computations.

Next, the heat-transfer condition at the walls was calibrated. When conducting measurements, the run time of the tunnel is short enough that the tunnel walls are generally assumed to be isothermal. A wall temperature of 291.2 K is chosen to represent the room temperature in which the test section resides. A comparison between boundary-layer profiles using adiabatic walls and isothermal walls at 291.2 K is shown in figure 4.24c. The resulting effect on the velocity profile appears to be negligible.

[1] The calibration procedure was performed by Daniel Galbraith and John Benek at the US Air Force Research Laboratory.

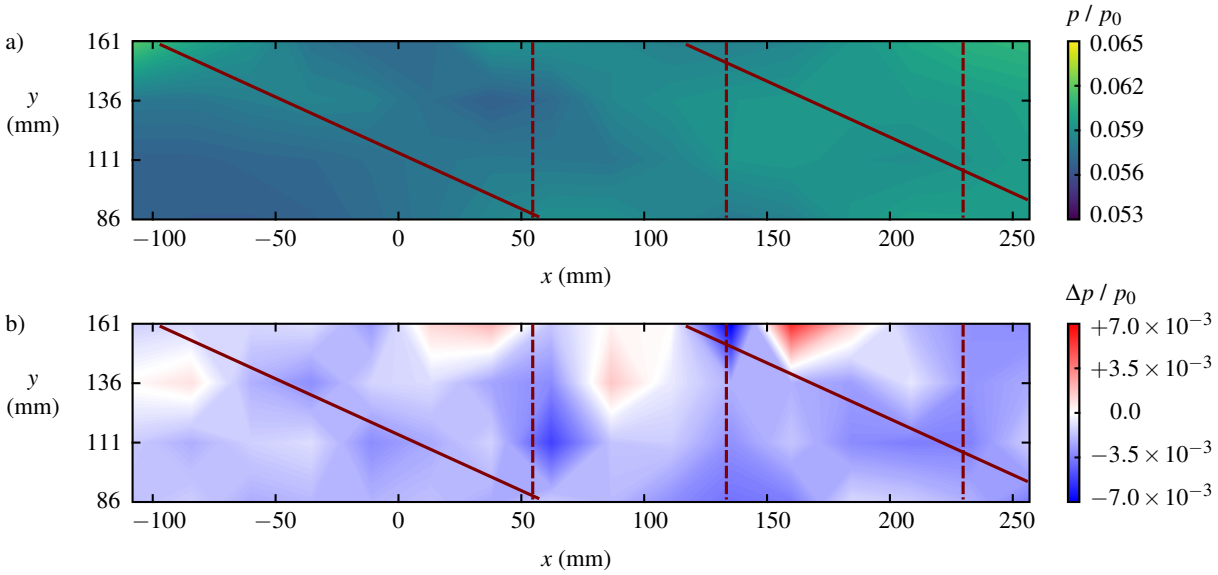

Figure 4.25: a) Computed static pressure for the full nozzle setup on the tunnel sidewall. b) The difference between the computed pressure distribution and equivalent experimental data (figure 4.8). The red lines correspond to high-pressure areas; these take the form of oblique (solid) and vertical (dashed) regions.

Therefore, since isothermal walls form a more accurate representation of the actual wind tunnel, the three-dimensional simulations are modelled with isothermal walls.

Evaluation of computations on the three-dimensional grid

The computations performed by the US Air Force Research Laboratory are evaluated using a range of characterisation data, in order to test how well they capture the physical wind tunnel flow.[2] In particular, it is important to assess: the pressure profile through the test section; the agreement of velocities across the channel's cross-section; and the profiles of the boundary layers growing along the walls.

For the full nozzle setup, the difference between the computational (figure 4.25a) and experimental (figure 4.8) sidewall pressure distributions is plotted in figure 4.25b. The difference in the static pressure ratio p/p_0 is around 3×10^{-3}, which corresponds to a Mach number discrepancy of about 1.5%.

The equivalent computational data for the half nozzle setup is given in figure 4.26a. The difference compared to the measurements from figure 4.9 are plotted in figure 4.26b. The discrepancy in static pressure ratio is roughly 6×10^{-3}, larger than the equivalent full nozzle comparison. Note, however, that this still corresponds to a relatively small difference in the corresponding Mach number of approximately 2.4%.

[2]This analysis was performed by the author.

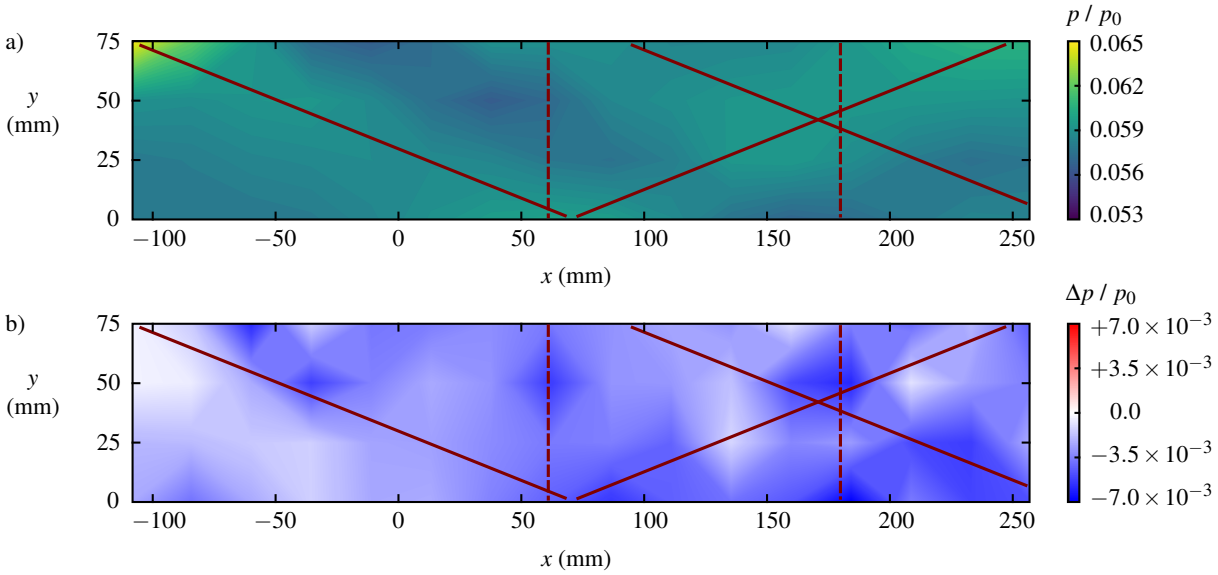

Figure 4.26: a) Computed static pressure for the half nozzle setup on the tunnel sidewall. b) The difference between the computed pressure distribution and equivalent experimental data (figure 4.9). The red lines correspond to high-pressure areas; these take the form of oblique (solid) and vertical (dashed) regions.

In both nozzle setups, the computations appear to under-predict the pressure, with the flow expanding to a lower pressure and a higher Mach number than the experiment. The top-left corner shows a higher pressure region that is just completing its expansion through the nozzle. The simulations appear to capture the weak waves generated by the tunnel floor and ceiling, which indicates that these correspond to uncancelled waves from the nozzle. However, the computations do not seem to reproduce the vertical features, which are believed to be waves from the tunnel sidewalls.

A second metric which can be used to evaluate how well the computations capture the physical flow is the streamwise velocity across the tunnel cross-section. This was measured for both nozzle setups using laser Doppler velocimetry (figure 4.13), and equivalent computational data are presented in figure 4.27a-i and 4.27b-i. These figures appear, on first glance, to be consistent with the experimental data – there is a large core region, the floor/ceiling boundary layers are spanwise-uniform with similar thicknesses, and there are significant variations in the thickness of the sidewall boundary layer.

A more quantitative measure of agreement can be obtained by calculating the difference between the computed and the experimental velocity distributions. This difference is shown in figures 4.27a-ii and 4.27b-ii for the two nozzle setups. In these figures, the velocities at most locations differ by no more than 10 ms^{-1}, or 1.7% of the freestream velocity. The notable exception is the sidewall boundary layers, where the discrepancy with experiment can be twice this large. Even though the calculation of the sidewall boundary layers is not fully accurate, 95% of data points have no more than 3.2% error for the full nozzle setup (4.27a-iii) and 3.5% error for the half nozzle setup (4.27b-iii).

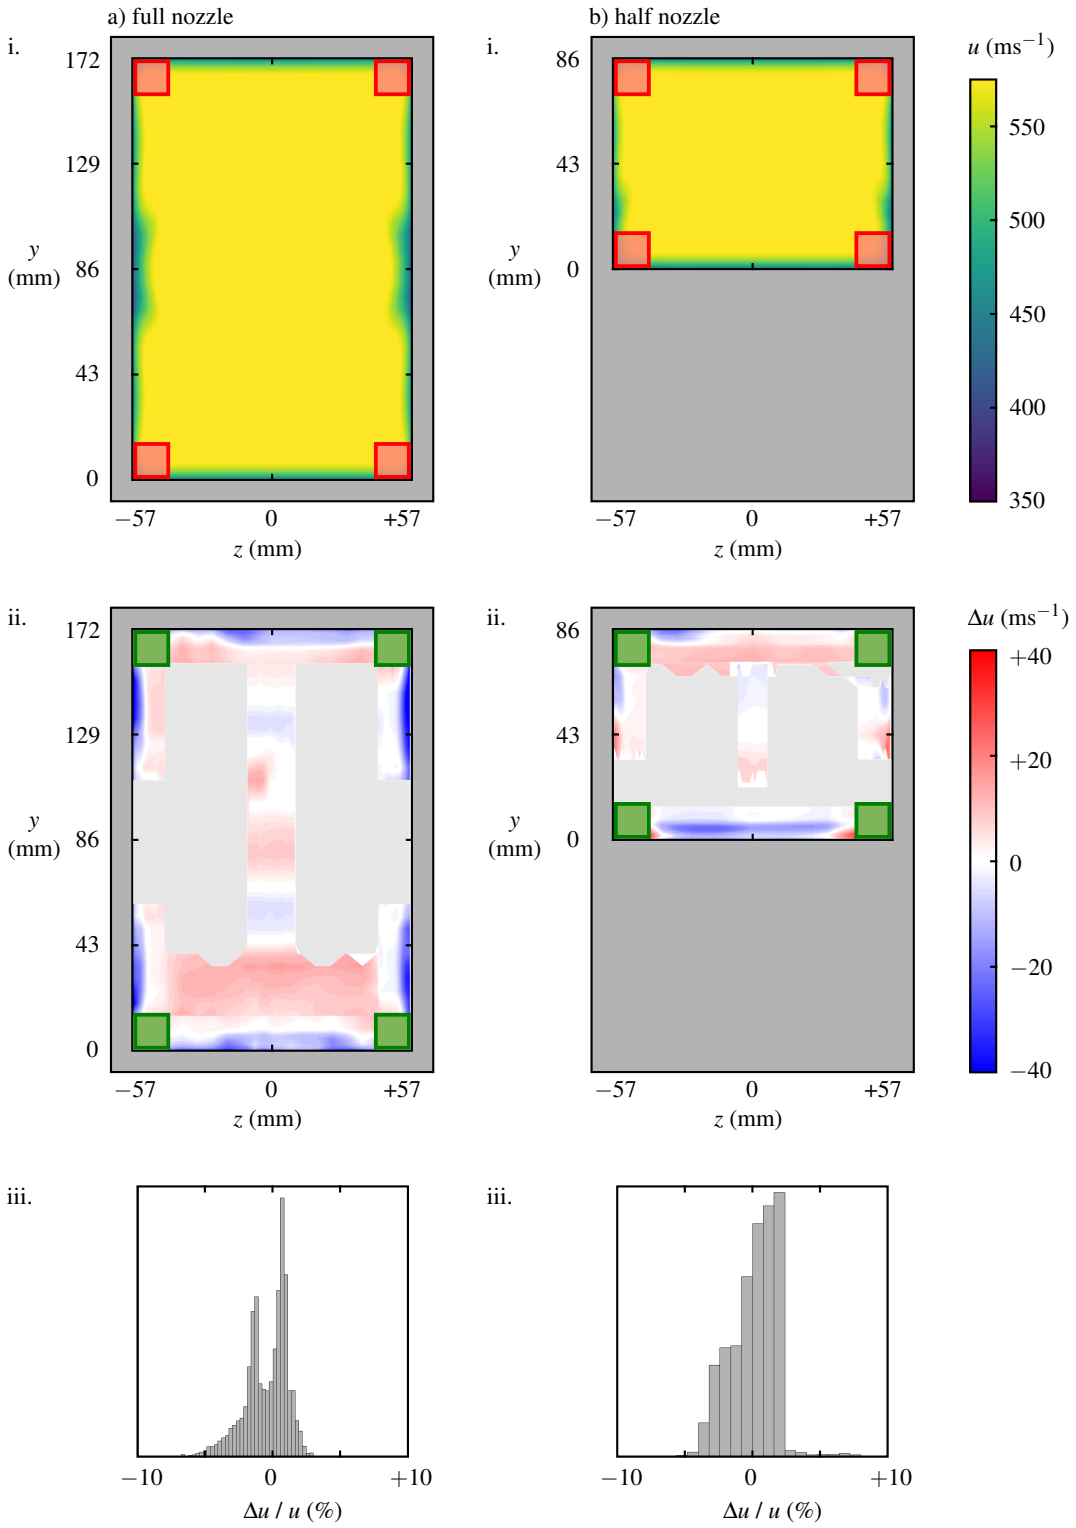

Figure 4.27: Comparison between computations and experiment of the streamwise velocity at $x = 120$ mm for a) the full nozzle setup, and b) the half nozzle setup: i. the computed streamwise velocity; ii. difference between the computed velocity and equivalent experimental data (figure 4.13); iii. histogram of percentage velocity differences between computations and experiment.

This is slightly larger than the estimated experimental error, indicating that the computations do not quite compute the true wind tunnel flow. However, a 3.5% error is considered to be small enough to show that the simulations do capture the flow sufficiently well to be used for validation purposes. Moreover, the differences between the two velocity distributions in this relatively well-understood part of the flow field can be used as a baseline 'expected minimum error' during validation of the corner regions.

The accuracy of computations is also evaluated by comparing the boundary-layer profiles to experiment. This is performed at various streamwise stations for the floor boundary layer on the centre span, assumed to be representative of the boundary layers growing along the wind tunnel walls.

A comparison between computational and experimental floor boundary-layer profiles is presented for the full nozzle in figure 4.28 and for the half nozzle in figure 4.29. In both cases the freestream velocity and the general profile shape appear to be in good agreement. The profiles do, however, differ slightly at the edge of the boundary layer for both nozzle setups.

A more quantitative assessment of the agreement between computations and experiment is performed by calculating the relevant boundary-layer parameters for the profiles in figures 4.28 and 4.29. These are listed in table 4.5, and can be compared to the equivalent experimental quantities from table 4.2. There is a good agreement of the integral parameters (δ_i^*, θ_i, and H_i). The deviations from experimental values are generally on the order of 4%. With an experimental uncertainty in these quantities of 5%, the computations do appear to predict these boundary-layer parameters with good accuracy.

Note however, that the deviations from experimental values of the boundary-layer thickness, δ, are on the order of 15%. This parameter therefore does not appear to be predicted particularly well by the computations. Whilst the boundary-layer thickness is known to be less robust than the corresponding integral parameters, the differences are believed to be largely due to the discrepancies in velocity near the boundary-layer edge. These discrepancies, which are particularly noticeable at the $x = 60$ mm measurement location, can perhaps be attributed to weak waves emitted from the junction between the nozzle block and the liner block. The differences are slightly larger than the estimated experimental uncertainties but are in fact relatively minor – the maximum velocity difference is 3%. This level of discrepancy is consistent with the typical differences between computed and experimental velocities across the tunnel cross-section, shown in figure 4.27. Therefore, whilst the precise wind tunnel flow might not be computed exactly, the simulations do appear to capture much of the key physics and can be used for validation purposes.

4.3 Benefits of close experimental–computational collaboration

It is relatively common for the experimental and computational stages of validation studies to be conducted separately, often with a significant time gap before the flow from an experiment is simulated. A particular advantage of the current study is that the experiments and computations are performed

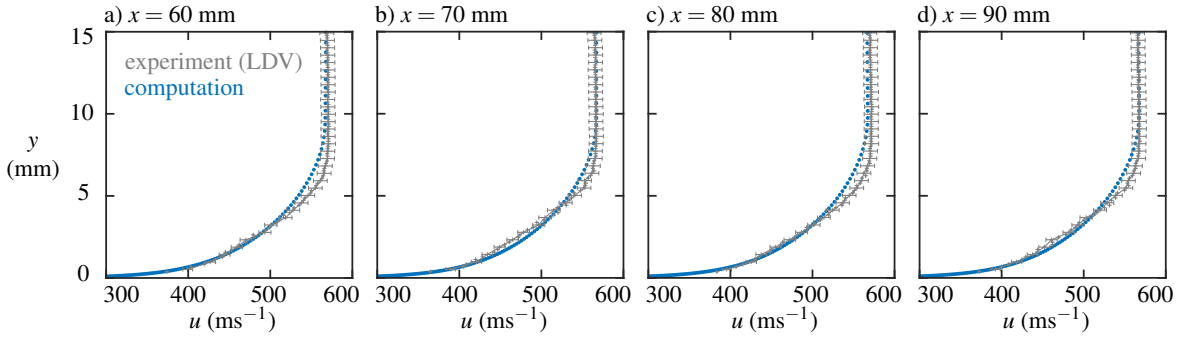

Figure 4.28: Comparison between computational and experimental floor boundary-layer profiles, with the full nozzle setup. CFD profiles are calculated using the fine mesh. The profiles are shown for the centre span at: a) $x = 60$ mm, b) $x = 70$ mm, c) $x = 80$ mm, and d) $x = 90$ mm.

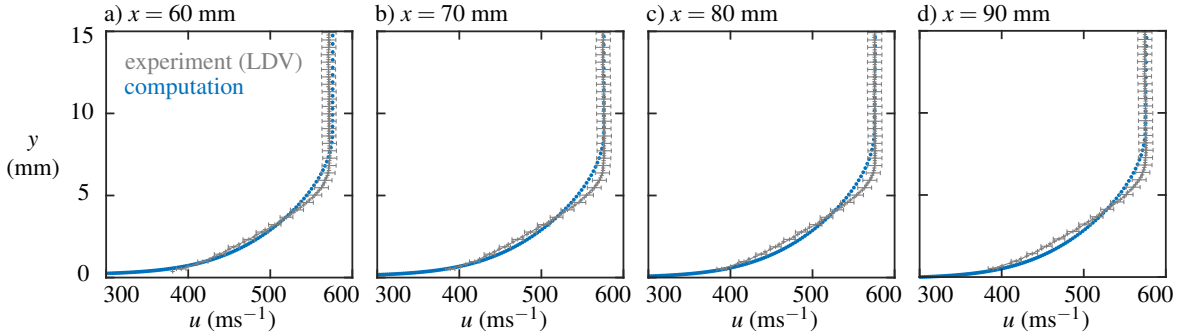

Figure 4.29: Comparison between computational and experimental floor boundary-layer profiles, with the half nozzle setup. CFD profiles are calculated using the fine mesh. The profiles are shown for the centre span at: a) $x = 60$ mm, b) $x = 70$ mm, c) $x = 80$ mm, and d) $x = 90$ mm.

Table 4.5: Incompressible floor boundary-layer parameters, along the tunnel centre line, extracted from the RANS computations. These correspond to CFD profiles presented in figure 4.28 (full nozzle) and figure 4.29 (half nozzle).

	full nozzle				half nozzle			
x (mm)	δ (mm)	δ_i^* (mm)	θ_i (mm)	H_i	δ (mm)	δ_i^* (mm)	θ_i (mm)	H_i
60	8.43	0.96	0.74	1.30	7.89	0.96	0.73	1.32
70	8.49	0.96	0.74	1.29	8.02	0.96	0.73	1.31
80	8.63	0.97	0.75	1.29	8.38	0.97	0.74	1.30
90	8.77	0.99	0.77	1.29	8.69	0.97	0.75	1.30

simultaneously, and close communication between the two approaches is maintained throughout the process.

On a fundamental level, this is beneficial to both parties – flow characterisation data required by computations can be collected when relevant wind tunnel configurations are still installed, whilst experiments benefit from the otherwise-inaccessible data that simulations can provide. More specifically, the flow characterisation process in the current study has highlighted a range of important synergetic benefits of close experimental–computational collaboration.

The collaborative approach leads to a better understanding of the experimental techniques in use, and their associated uncertainties. For example, section 3.2.4 details the analysis of computational profiles in order to better understand the effects of a finite probe volume in measurements using laser Doppler velocimetry. The departure from the 'ideal' point-measurement velocities are very small, except within 1.4 mm of the tunnel sidewalls, where there can be an effect as large as 3%. Importantly, it would not have been possible to extract this information from experiment alone, where the 'true' point-measurement velocities are unknown. Close collaboration is thus necessary to gain this kind of insight into experimental techniques.

The joint experimental–computational approach has also supported substantial improvements in the wind tunnel infrastructure. In particular, section 3.2.4 describes the problems associated with introducing LDV seeding particles to the flow through commonly-used rakes. Too few of these particles end up in the corner boundary layers of the test section for validation-quality measurements in these regions. As shown in figure 3.21, streamlines in the computations were traced from the areas of interest upstream to the location of the seeding rakes. Using this information, a new system could be configured, and these rakes successfully seeded the corner boundary layers (figure 3.22).

This corner seeding system enables high-quality measurements to be performed not only in the corner regions but also in the tunnel's sidewall boundary layers. There is substantial motivation to conduct these sidewall boundary-layer measurements, in order to fully characterise the physical wind tunnel flow field for calibration of computations. The resulting measurements, which form the subject of section 5.1, provide physical insights which are not only surprising but can have far-reaching consequences that affect the corner boundary layers themselves.

4.4 Summary

This chapter details the procedure to characterise the flow for a validation study. The process includes careful measurement of the as-installed tunnel geometry, the sidewall static pressure distribution, and the velocity field across the tunnel cross-section. In performing these measurements, the validity of common assumptions about both techniques and the wind tunnel flow are assessed. Examples of such assumptions are that measurement techniques are accurate, wind tunnel experiments are repeatable, and that any symmetry of the geometry is reflected in symmetry of the flow. The tests described in this

chapter determine the extent to which these assumptions are valid, which also enables identification of potential systematic errors.

The collected data serve two key purposes. Firstly, the uniformity of the flow (for example, Mach number variations of less than 2% over the test section) shows that the flow quality is sufficiently good for a validation study. Second, the quantitative characterisation measurements are used to set up computations and to evaluate how well these simulations capture the physical wind tunnel flow. The computations are shown to accurately predict the sidewall static pressure distribution (discrepancy: 2%), the floor boundary-layer profile (integral parameter discrepancy: 4%), and the cross-sectional streamwise velocity distribution (discrepancy: 3%). The data from the simulations can therefore be used to validate the capabilities of different numerical methods in the corner regions.

The flow characterisation process highlights the importance of close experimental–computational collaboration in improving the understanding of experimental techniques and in guiding improvements to the tunnel infrastructure, such as the corner seeding system. Chapter 5 provides further examples of such benefits by illustrating how a collaborative approach can lead to the discovery of new flow physics, both in the measurements of the sidewall boundary layers and in the study of previously-unknown features observed in schlieren images of the empty wind tunnel flow.

Chapter 5

Nozzle geometry and the corner boundary layers

This chapter presents the discovery of two different features of the flow in supersonic wind tunnels. These are both related to the geometry of the two-dimensional nozzles, which are typically used in facilities with a rectangular cross-section. The effects of these flow features on the structure of the corner boundary layers are then studied, and the numerical methods used to compute these flows are validated against experimental data. This chapter presents data from RANS computations performed by Daniel Galbraith and John Benek at the US Air Force Research Laboratory. Analysis of the data was performed by the author.

5.1 Secondary flow in the sidewall boundary layers

Figure 5.1 shows the streamwise velocity data for both nozzle setups from a series of traverses at $x = 120$ mm. For the reasons described in section 4.2, the corner boundary layers are not used for flow characterisation, and so these regions are omitted from the figure. The tunnel has relatively thin boundary layers with respect to the total cross-sectional area. For each setup, the floor and ceiling boundary layers have approximately the same thickness. This boundary-layer thickness, around 8 mm, is roughly equivalent for the two setups, as shown in figure 5.2. The figure also shows that these boundary layers exhibit small variations of around 4% across the tunnel span, and so are considered to have approximately constant thickness in this direction.

On the other hand, figure 5.3 reveals that the sidewall boundary-layer thickness is not uniform. Unfortunately, high-quality data cannot be obtained over the entire tunnel height due to insufficient seeding particles in some regions. There are therefore gaps in data from $y = 60 - 110$ mm for the full nozzle (figure 5.3a), and from $y = 15 - 30$ mm for the half nozzle (figure 5.3b). Nevertheless, within the measurement region variations in sidewall boundary-layer thickness of 44% in the full nozzle setup and 43% with the half nozzle can be observed. Note, however, that the two setups exhibit quite different behaviours. In the full nozzle setup (figure 5.3a), the sidewall boundary layers increase in

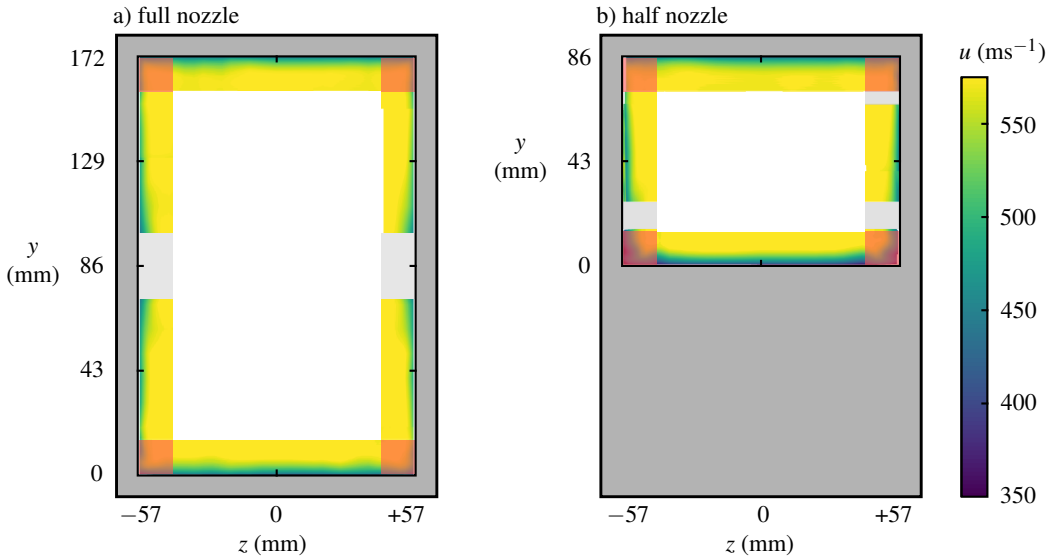

Figure 5.1: The streamwise velocity (u) in the tunnel boundary layers, measured using LDV at $x = 120$ mm for a) the full nozzle setup, and b) the half nozzle setup. Corner regions, marked in red, are excluded from this flow characterisation data. Regions with an insufficient number density of seeding particles for high-quality velocity measurements are marked in light grey.

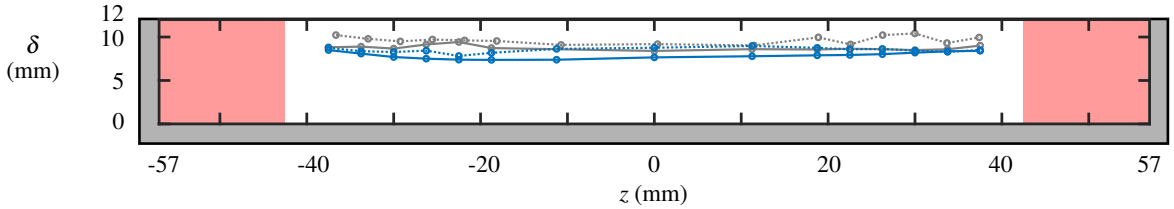

Figure 5.2: Thickness of the boundary layers at $x = 120$ mm across the tunnel span. Measurements shown for the ceiling (dashed) and floor (solid) boundary layers, for the full nozzle (grey) and the half nozzle (blue). Corner regions, marked in red, are excluded.

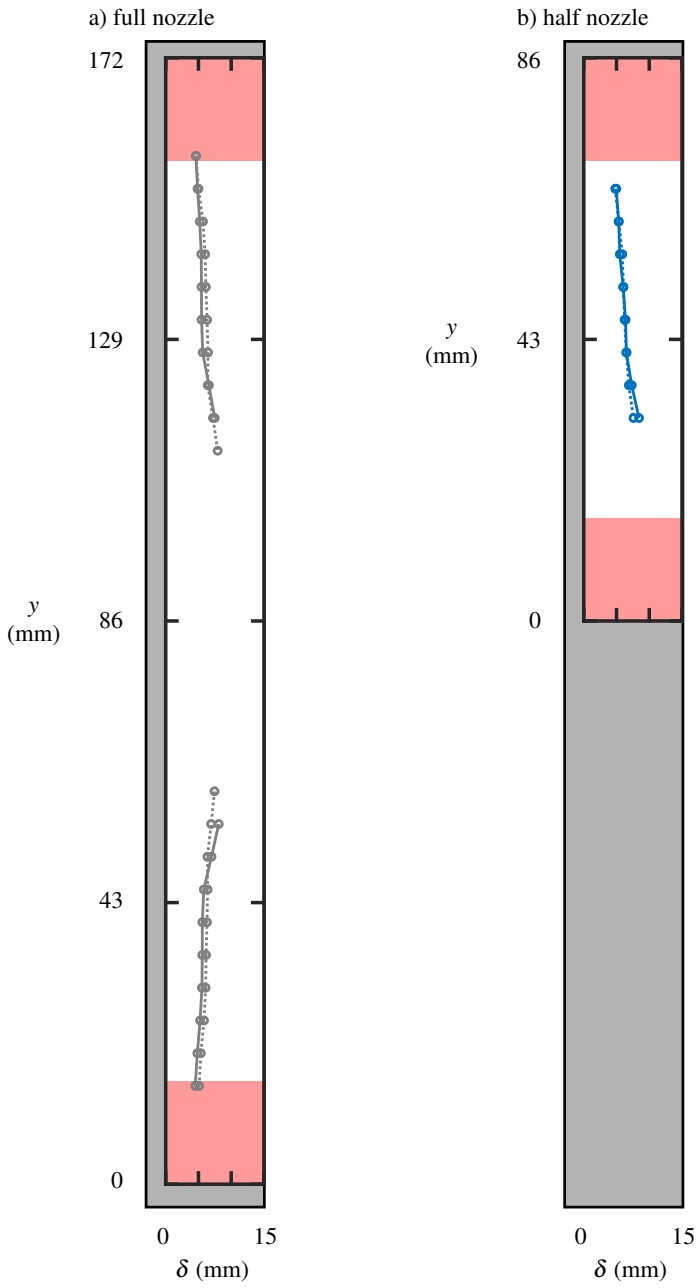

Figure 5.3: Thickness of the sidewall boundary layers at $x = 120$ mm for a) the full nozzle setup, and b) the half nozzle setup. In each case, measurements are shown for the positive-z sidewall boundary layer (dashed) and the negative-z sidewall boundary layer (solid). Corner regions, marked in red, are excluded.

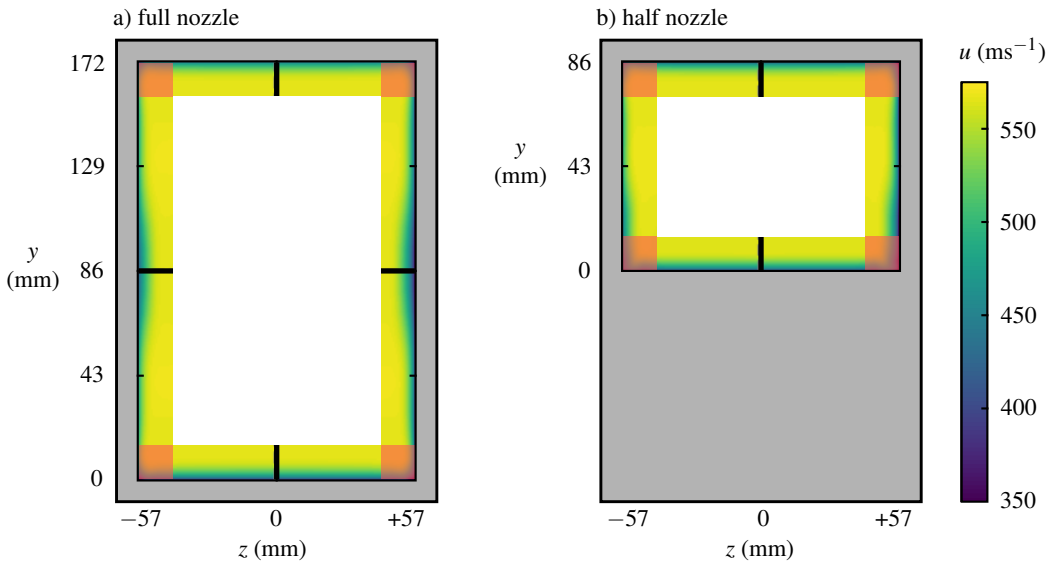

Figure 5.4: The streamwise velocity (u) in the tunnel boundary layers, extracted from computations at $x = 120$ mm for a) the full nozzle setup, and b) the half nozzle setup. Corner regions, marked in red, are excluded from this flow characterisation data.

thickness from each corner towards the centre height. With the half nozzle setup, however, figure 5.3b shows that these boundary layers get continuously thicker from top to bottom.

The large variation in sidewall boundary-layer thickness is somewhat unexpected, and is certainly not desirable. The available computational data, provided by the US Air Force Research Laboratory, are analysed to see if they too display similar behaviour.[1] The streamwise velocity distribution across the tunnel cross-section is therefore plotted in figure 5.4 for the two nozzle setups. This figure shows many of the same features as observed in experiment. The floor and ceiling boundary layers are all approximately 8 mm thick, and uniform across much of the tunnel span. Crucially, the computations reflect the same non-uniformities in the sidewall boundary layer as the experiments. The full nozzle setup in figure 5.4a exhibits a thickening boundary layer from the corners to the centre height. In contrast, the sidewall boundary-layer thickness in the half nozzle setup (figure 5.4b) continuously increases in the downwards direction. Therefore, the computations appear to successfully capture the physics governing the variations in sidewall boundary-layer thickness, and so these are used together with experiments to identify the causes for this behaviour.

5.1.1 Physical explanation

Figure 5.5 shows the pressure distribution in the nozzle region, extracted along the centre span from computations for both nozzle configurations. The expansion of the flow corresponds closely to the inviscid wave pattern, which is illustrated schematically in figures 5.6a and 5.7a for the two setups.

[1] This analysis was performed by the author.

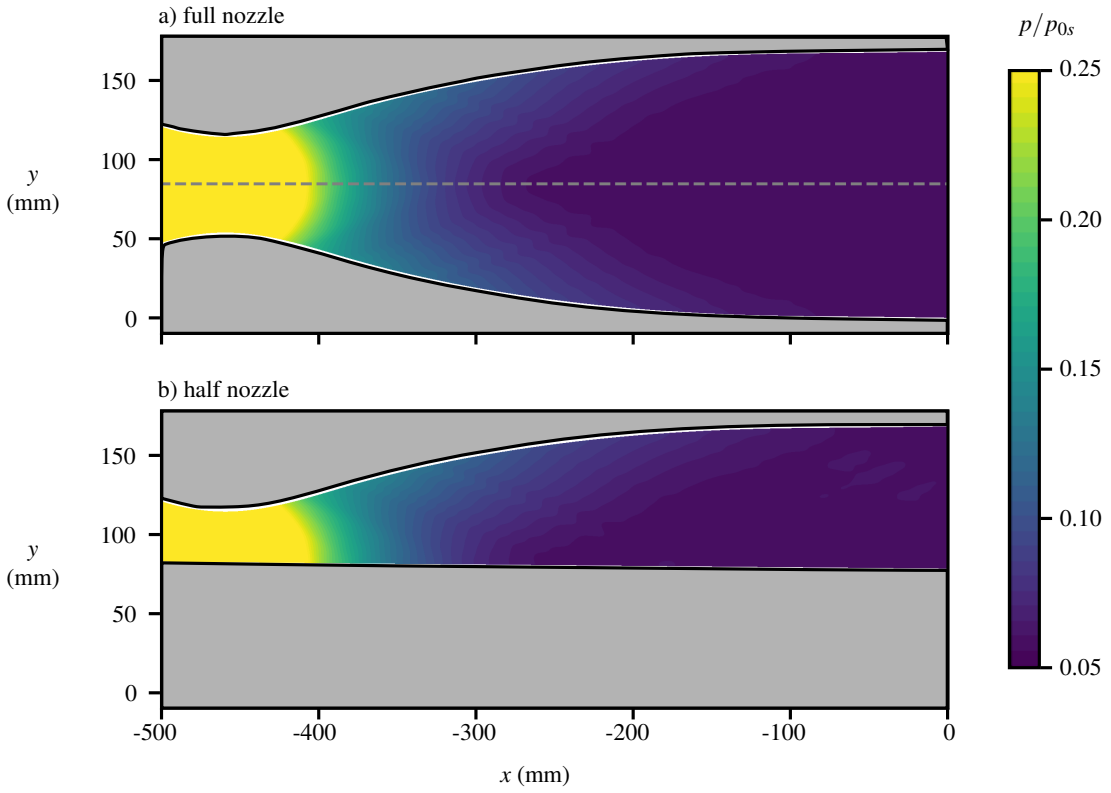

Figure 5.5: Pressure contours in the nozzle on the centre span, for a) the full nozzle setup and b) the half nozzle setup.

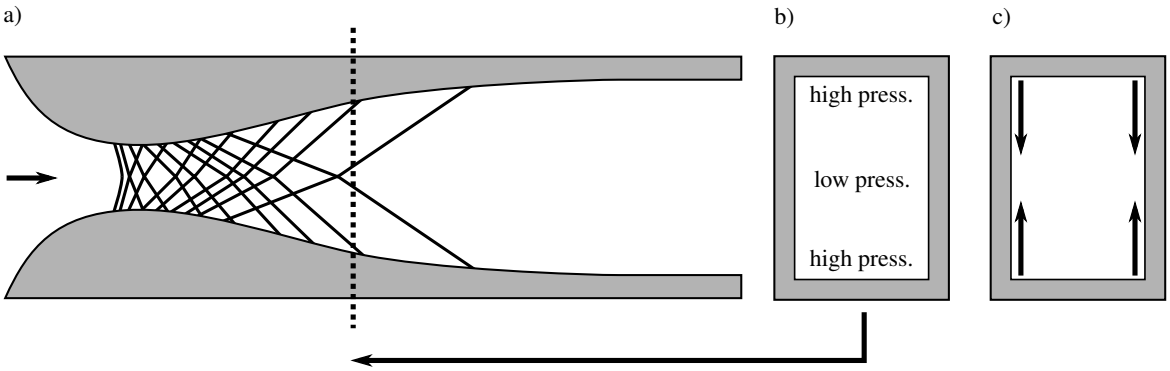

Figure 5.6: Schematic of the physical mechanism for the full nozzle setup: a) inviscid expansion wave pattern through the nozzle; b) the cross-sectional pressure distribution upstream of the nozzle exit; c) induced secondary flows in the sidewall boundary layers.

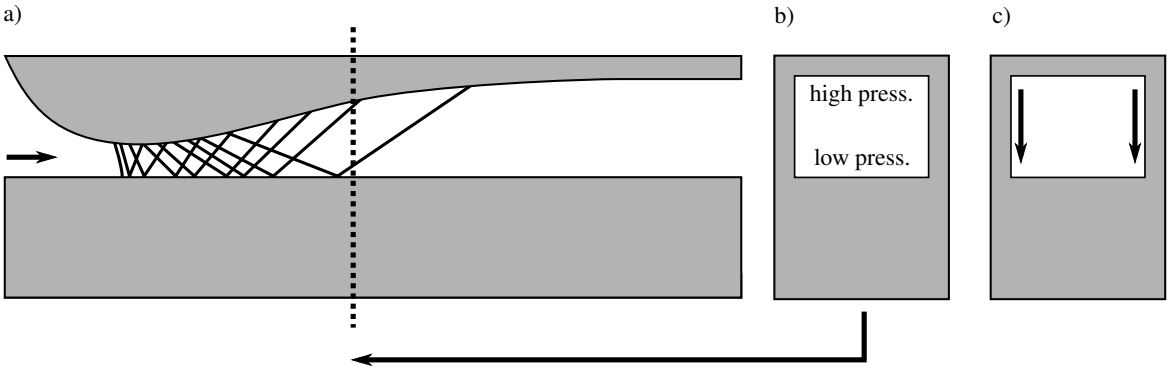

Figure 5.7: Schematic of the physical mechanism for the half nozzle setup: a) inviscid expansion wave pattern through the nozzle; b) the cross-sectional pressure distribution upstream of the nozzle exit; c) induced secondary flows in the sidewall boundary layers.

The full nozzle setup in figure 5.6 will be considered first. The dashed line in figure 5.6a lies approximately half-way between the nozzle throat and exit. Here, the flow at the centre height has expanded to the test section pressure. Meanwhile, the pressure towards the top and bottom of the tunnel is, comparatively, higher because here the flow expansion is not yet complete. Similar arguments hold elsewhere in the nozzle, with the pressure drop on the curved surfaces always lagging behind that at the centre height. This type of expansion therefore sets up a pressure gradient in the vertical direction throughout the nozzle (figure 5.6b). The sidewall boundary layers, with low momentum flow, are most susceptible to this pressure gradient and a secondary flow forms in response. This takes the form of a vertical bulk flow from the channel corners towards the centre height (figure 5.6c).

The mechanism for the half nozzle setup, illustrated in figure 5.7, is equivalent. Here, the vertical pressure gradient is from the top to the bottom of the tunnel (figure 5.7b). As a result, the bulk flow in the sidewall boundary layers, shown in figure 5.7c, is purely in the downwards direction.

In order to test the hypothesised mechanism, the vertical velocity distribution is measured for both nozzle setups. A lack of optical access necessitates these measurements are performed downstream

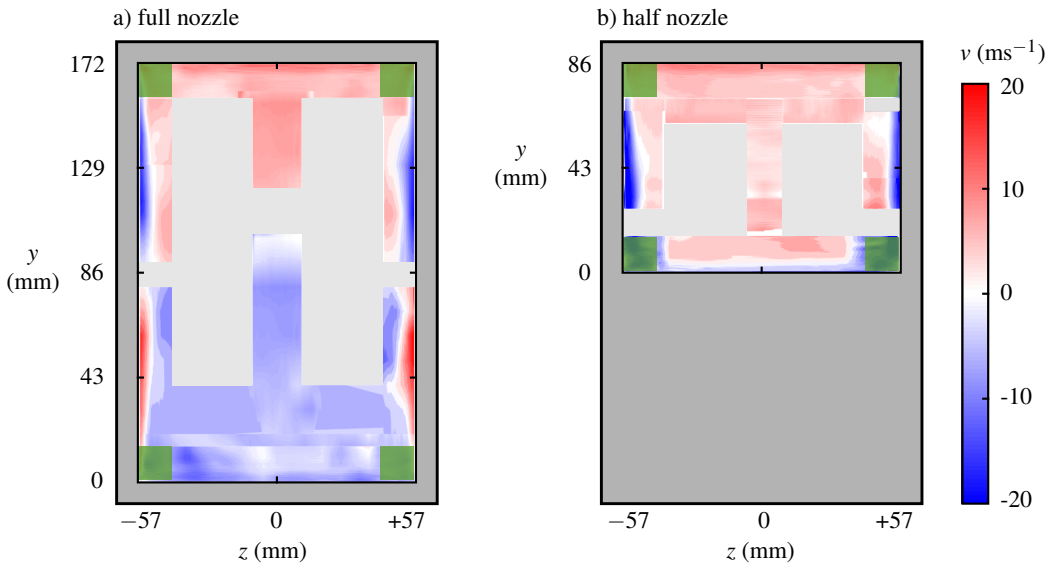

Figure 5.8: The vertical velocity (v) in the tunnel boundary layers, measured using LDV at $x = 120$ mm for a) the full nozzle setup, and b) the half nozzle setup. Corner regions, marked in green, are excluded from this flow characterisation data. Regions with an insufficient number density of seeding particles for high-quality velocity measurements are marked in light grey.

of the nozzle exit, at $x = 120$ mm. Even though there is no driving force in the form of a vertical pressure gradient within the test section itself, the sidewall secondary flows might be expected to persist for some distance downstream of the end of the nozzle. As a result, it is hoped that evidence of the secondary flows can be found at this downstream location.

Figure 5.8 presents the vertical velocity component obtained from experiments, and the equivalent data from computations are given in figure 5.9. For each nozzle setup, the vertical velocity distribution determined by both approaches displays the same general patterns. Within the core flow, the vertical velocities are very small, not exceeding 1% of the local freestream velocity. In contrast, there is a significant vertical velocity component (up to 5% of u) within the sidewall boundary layers, indicative of a secondary flow. The directions of these vertical velocities are consistent with the suggested mechanism – towards the tunnel centre height in the full nozzle setup and downwards with the half nozzle.

The variations in sidewall boundary-layer thickness observed in figure 5.1 can now be explained in the context of these sidewall secondary flows. The vertical velocities transport the low momentum fluid in these regions, so that the boundary layer is thinnest at the ‘upstream’ end of the cross-flow and thickens in the direction of the secondary flow. The resulting thickness distributions of the sidewall boundary layers are consistent with those measured in experiment.

The effects of these secondary flows are not restricted to the sidewall boundary layers. The sidewall boundary layers are closely coupled with the corner regions, which form the subject of this thesis. Therefore, the influence of the nozzle geometry, through the induced vertical velocities in the sidewall boundary layers, might also be expected to extend to the corner regions. This idea is

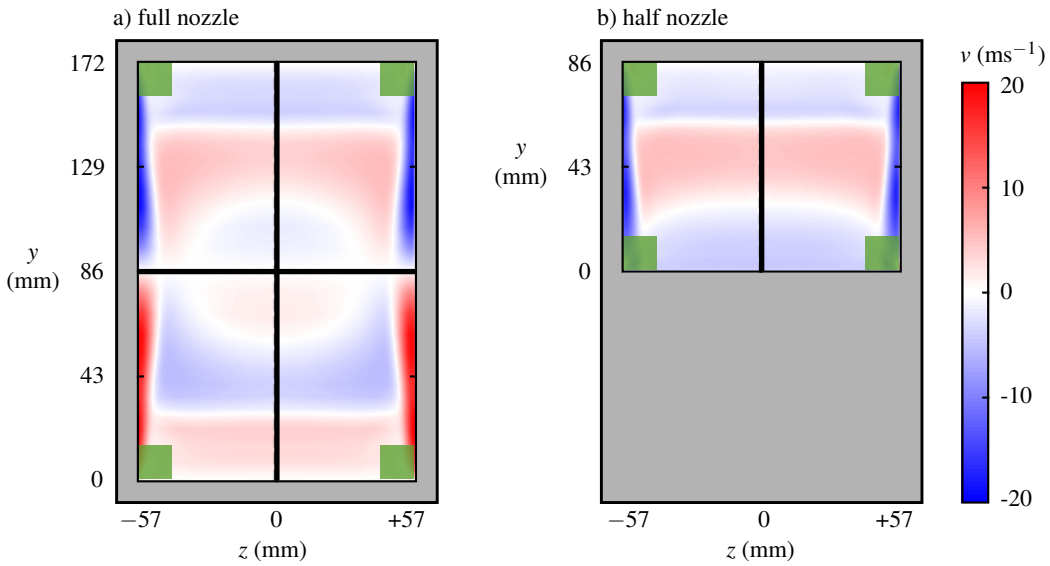

Figure 5.9: The vertical velocity (v) in the tunnel boundary layers, extracted from computations at $x = 120$ mm for a) the full nozzle setup, and b) the half nozzle setup. Corner regions, marked in green, are excluded from this flow characterisation data.

explored further in section 5.3, where the factors which influence the structure of the corner boundary layer are examined in detail.

5.2 Vortices induced by the nozzle geometry

The detailed characterisation process also identified another notable feature in the wind tunnel flow. Figure 4.2a shows a schlieren image, captured with an exposure time of 0.1 ms. The boundary layers on the tunnel floor and ceiling are visible, along with some weak Mach waves generated by surface imperfections in the tunnel walls. When the exposure time is reduced to 1.1 μs, as shown in figure 5.10, turbulent eddies become visible. These are thought to be inside the sidewall boundary layers. In addition, two subtle streamwise-coherent flow features, indicated by red arrows in the figure, can be resolved. These flow structures lie slightly above and below the tunnel centre height.

In order to try to identify these features in the computations, numerical schlieren images are generated by integrating the vertical density gradient $\partial\rho/\partial y$ across the tunnel span. This visualisation, shown in figure 5.11, captures the floor and ceiling boundary layers along with with weak waves, originating from the tunnel floor and ceiling. Most importantly, however, there appear to be two faint, streamwise-coherent structures that look similar to those observed in the experimental schlieren image (figure 5.10).

The features identified in figure 5.10 appear similar to artefacts observed in schlieren images from a range of facilities, shown in figure 5.12 [132–134]. However, these features have not been

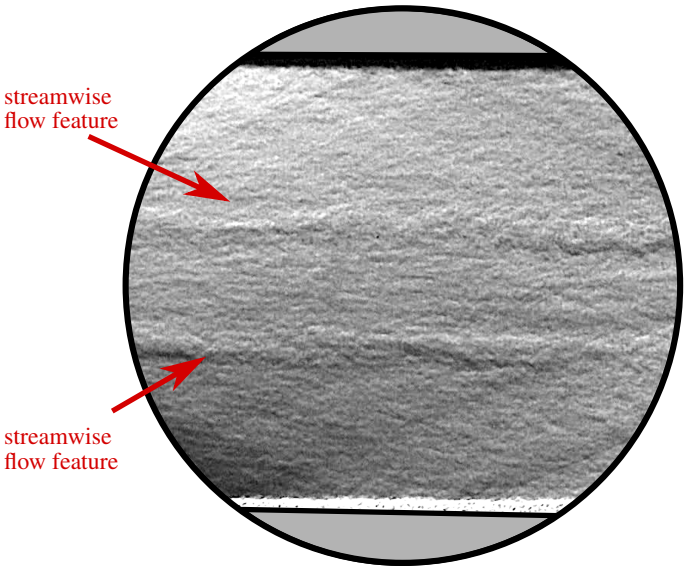

Figure 5.10: Schlieren image of the empty wind tunnel with a horizontal knife edge and an exposure time of 1.1 μs. Flow direction is from left to right.

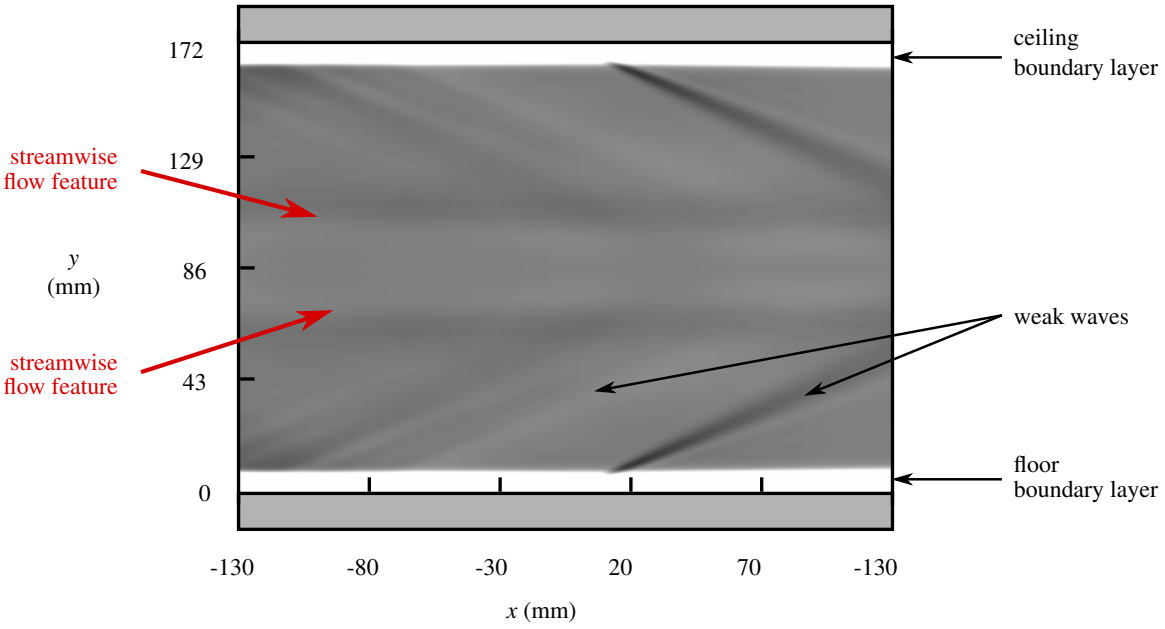

Figure 5.11: Numerical, spanwise-averaged schlieren image of the empty wind tunnel extracted from computations. The figure shows density gradients in the y-direction, analogous to a horizontal knife-edge. Flow direction is from left to right.

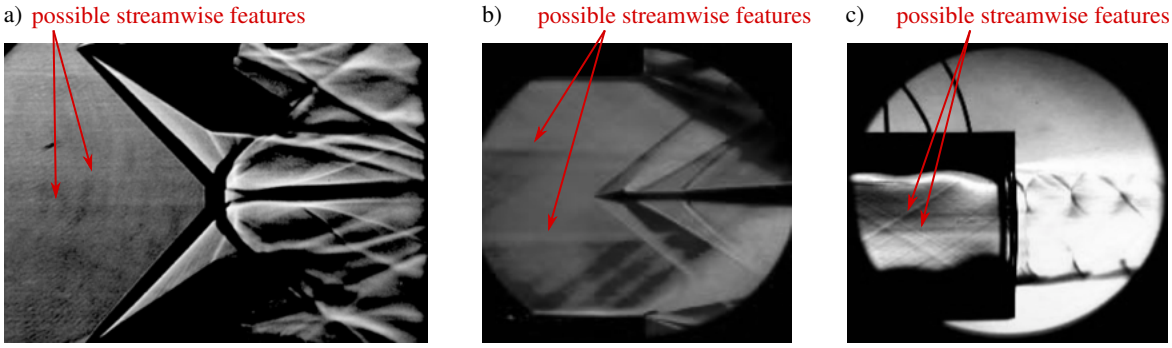

Figure 5.12: Suggestions of similar streamwise-coherent flow features in enhanced schlieren images from various facilities: a) Mach 3 flow at the University of Witwatersrand [132]; b) Mach 1.56 flow at the Military Technical Institute, Serbia [133]; c) Mach 2 flow at the University of Texas in Arlington [134]. In each image, flow direction is from left to right.

commented on in the literature, and their nature remains unknown. Nevertheless, the apparent presence of similar structures in other wind tunnels suggests that there may be a common underlying cause.

Origin of unknown features

Figure 5.13a shows the cross-sectional distribution of streamwise rigid-body vorticity, ω_{RR}, extracted at $x = 120$ mm from computations performed by the Air Force Research Laboratory. This quantity, defined by Kolář, is calculated by removing the effects of mean shear and of irrotational strain from the vorticity field [135]. The researchers who performed the computations noticed a region of strong vorticity inside the sidewall boundary layer, 18 mm away from the tunnel centre height. The vertical location of this 'sidewall vortex' matches the location of the unknown flow features observed in experiment (figure 5.10b) and in computations (figure 5.11). Note that this vortex is separate from the stress-induced counter-rotating vortex pair within the corner boundary layer, labelled as 'corner vortices' in figure 5.13a.

Figure 5.13 show that the sidewall vortices appear to originate from a separated flow region immediately upstream of the nozzle block. Here, there is a sharp geometry change at the start of the nozzle, which is shown more clearly in figure 5.14a. The computations predict that the floor boundary layer separates ahead of this sharp corner geometry (figure 5.13b). In order to confirm the physical presence of such a separation, oil-flow visualisation is performed in this region. The resulting skin-friction line topology, shown in figure 5.14b, is indeed consistent with a separated flow field. The separated region has a streamwise extent of 6 ± 1 mm on the tunnel centre span. This agrees well with the computations, which predict a separation length of 5.2 mm.

The existence of the separation in itself is not entirely surprising. Despite the favourable pressure gradient which exists upstream of the nozzle, the sharp geometry change is topologically similar to a forward-facing step and thus separation of the incoming boundary layer might be expected. However, the fact that the separation has an effect (in this case, a vortex) which persists through the nozzle into the test section is, perhaps, less obvious.

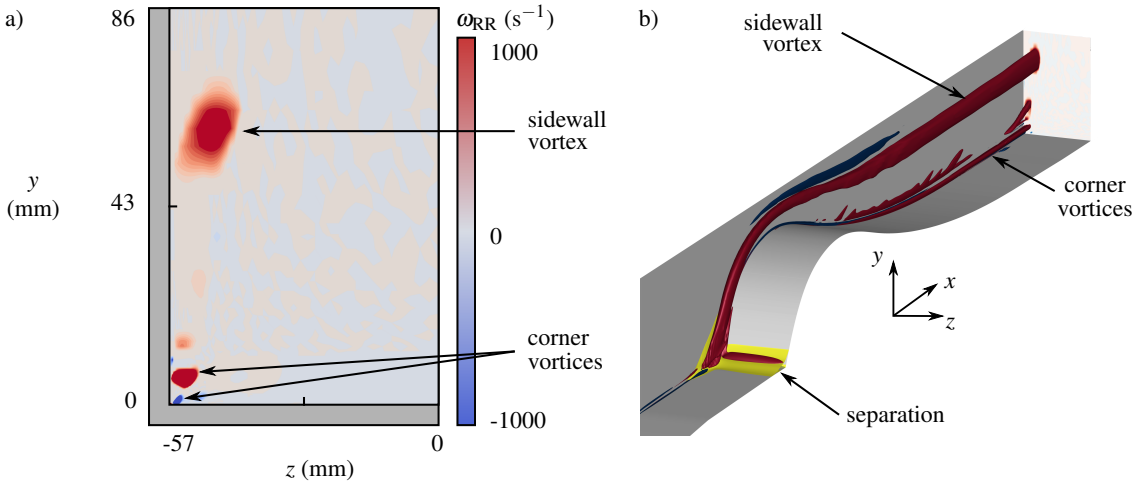

Figure 5.13: Rigid-body component of vorticity, ω_{RR}, from RANS computations: a) cross-sectional distribution at $x = 120$ mm; b) isometric view, showing isosurfaces of ω_{RR} in red and blue, with a region of separation highlighted in yellow. One quarter of the wind tunnel is shown.

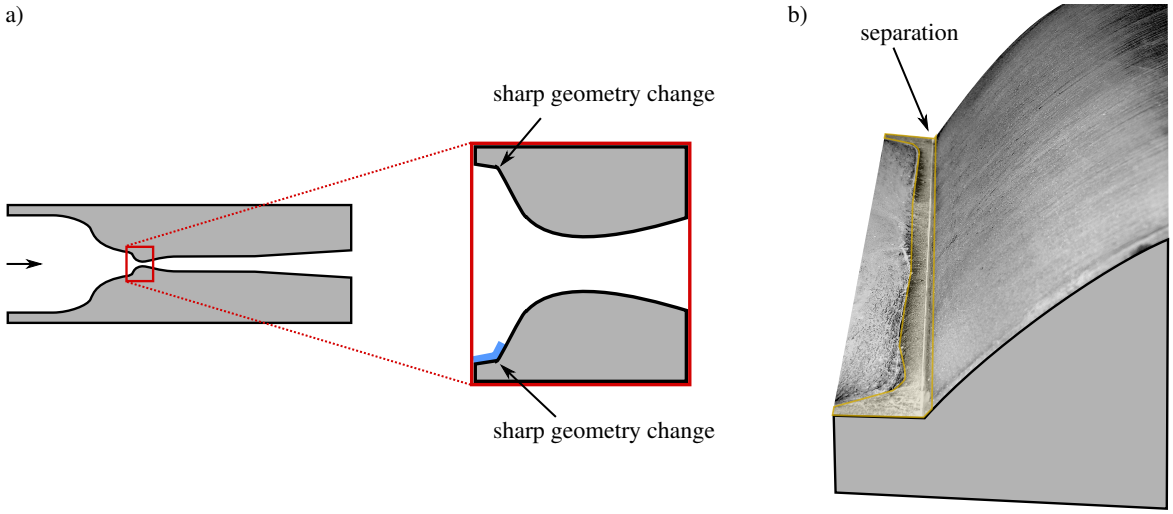

Figure 5.14: a) Magnification of the sharp geometry change upstream of the nozzle. The surface for oil-flow visualisation is marked in blue. b) Surface oil-flow pattern, with the separated region highlighted in yellow.

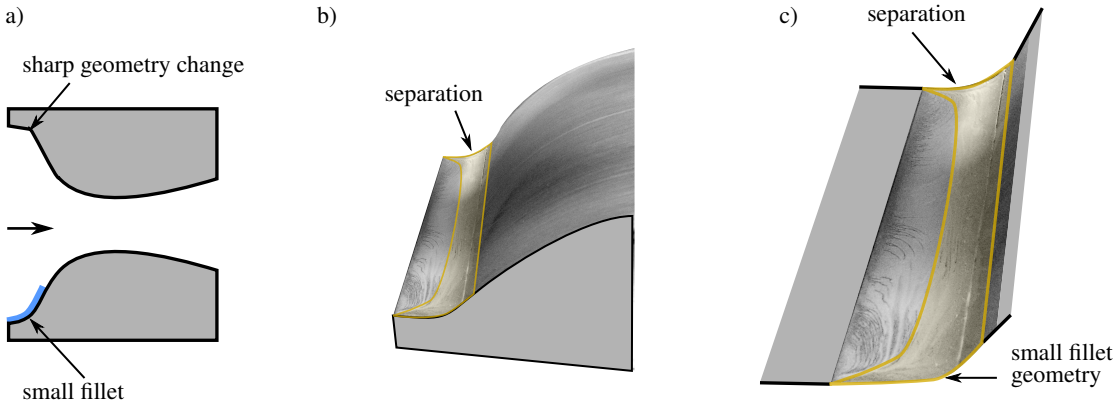

Figure 5.15: Experimental tests on the small fillet geometry on the tunnel floor, showing: a) the installed geometry, b) oil-flow visualisation (separated region highlighted in yellow) on the surface marked blue in a, and c) magnified view of separated region from b.

5.2.1 Elimination of the nozzle-induced vortices

The RANS computations suggest that the sidewall vortex originates in the separated region upstream of the nozzle. Therefore, removal of this separation should alter, or perhaps eliminate, the vortex. Removing the separation can be achieved by replacing the sharp geometry with a more gentle contour. In order to investigate this hypothesis further, the geometry upstream of the nozzle is redesigned by introducing a rounded fillet. Two different fillet radii are investigated: a 'small' fillet radius of 20 mm and a 'large' fillet radius of 100 mm. According to the numerical simulations, these radii correspond to 5 and 25 local boundary-layer thickness, respectively.

The streamwise flow features in figure 5.10 are quite weak and have only been observed recently, suggesting that a well set-up, sensitive schlieren system is required to capture them. If fillets are introduced to both the floor and ceiling geometries simultaneously, any observed differences in the flow field could not definitively be attributed to the installed geometry as opposed to an inadvertent change in the sensitivity of the schlieren system. Therefore, the filleted geometry is instead installed along the tunnel floor whilst the original sharp-cornered (or unfilleted) setup is left along the ceiling, as shown in figures 5.15a and 5.18a. With such a setup, a comparison between the top and bottom halves of the schlieren images directly tests the effect of the installed filleted geometry against the original baseline case.

Small fillet geometry

The small fillet geometry is investigated first. Surface oil-flow visualisation (figures 5.15b and 5.15c) shows that this setup still exhibits separation upstream of the nozzle, albeit with a slightly smaller streamwise extent (5 ± 1 mm). The equivalent computation for the small fillet geometry, shown in figure 5.16a, also predicts separation of the incoming flow. In addition, the computations show that the sidewall vortices remain in the test section and are more or less unchanged, as evident in figure 5.16a.

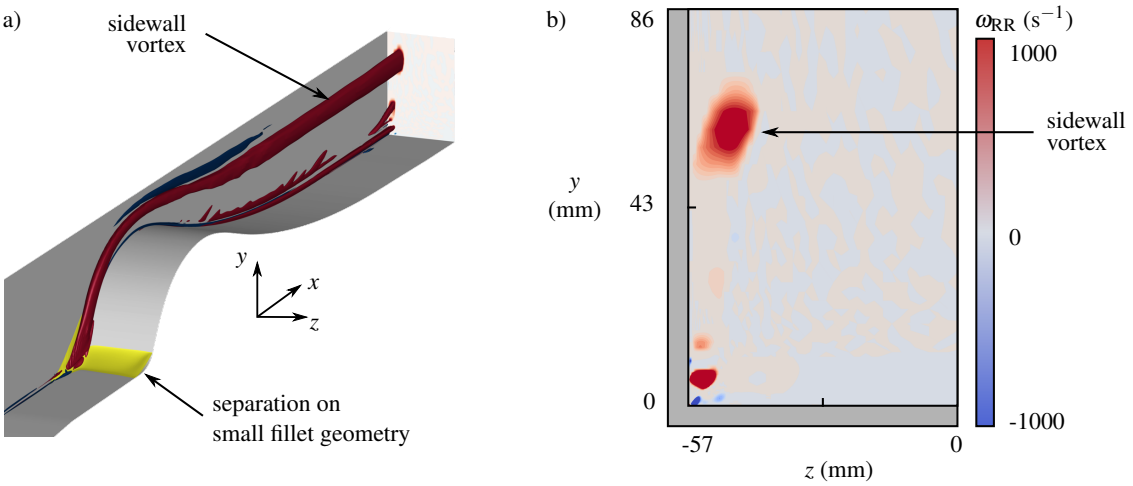

Figure 5.16: Rigid-body component of vorticity, ω_{RR}, from computations of the small fillet geometry: a) isometric view, showing isosurfaces of ω_{RR} in red and blue, with a region of separation highlighted in yellow; b) cross-sectional distribution at $x = 120$ mm. One quarter of the wind tunnel is shown.

The sidewall vortex marked in figure 5.16b has a rigid-body vorticity distribution almost identical to that determined for the unfilleted geometry (figure 5.13a).

Figure 5.17a is a numerical schlieren image based on the computations for this small fillet setup. The streamwise-coherent density gradients associated with the sidewall vortices are, unsurprisingly, still visible. Schlieren visualisation of the physical wind tunnel flow is presented in figure 5.17b. Weak waves generated by the tunnel's floor and ceiling at the junctions between liner blocks can be discerned. More importantly, the image shows the presence of streamwise-coherent structures with similar strength in both the top and bottom halves of the flow. Both experiments and computations therefore indicate that the small fillet geometry does not remove the separation ahead of the nozzle, and that the sidewall vortices continue to exist in the test section. These findings are consistent with the hypothesised relationship between upstream separation and the sidewall vortices.

Large fillet geometry

In order to fit the large fillet, geometric constraints require the nozzle to be shifted downstream by 150 mm. The revised setup (figure 5.18a) consists of an unfilleted geometry on the tunnel ceiling and the large fillet on the floor. Here, the oil-flow visualisation on the floor (figure 5.18b) suggests that the smooth contour has been successful at removing flow separation. The computations in figure 5.19a also show that the flow remains attached.

In addition, the CFD results indicate that there is no longer a sidewall vortex for this configuration, as seen in figure 5.19b. This finding is supported by the numerical schlieren image in figure 5.20a, where the streamwise-coherent features visible in figures 5.11 and 5.17a have disappeared.

The experimental schlieren image in figure 5.20b shows a single streamwise-coherent feature in the upper half of the flow, which corresponds to sidewall vortices from the separation on the unfilleted

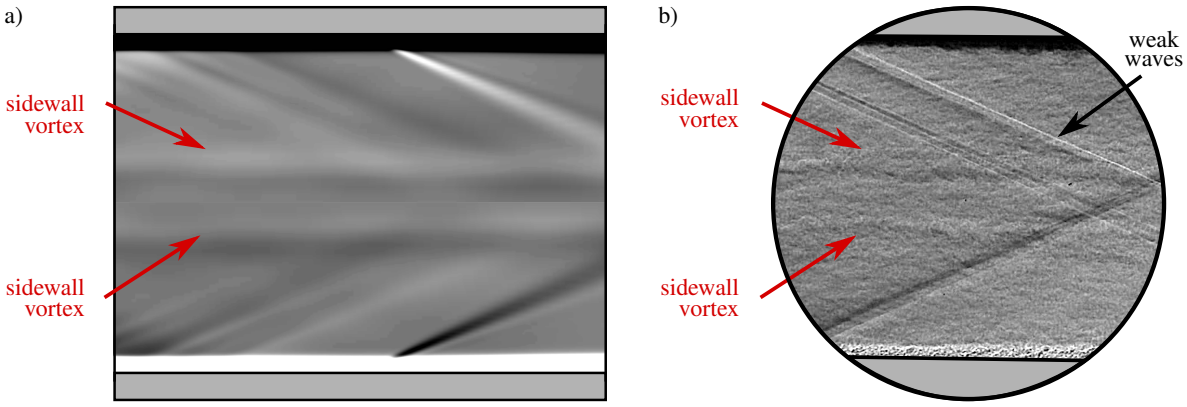

Figure 5.17: Schlieren visualisation of the flow with the small fillet geometry: a) numerical schlieren from computations (floor and ceiling: small fillet); b) experimental schlieren image with 1.1 μs exposure time (floor: small fillet, ceiling: unfilleted). Flow is from left to right.

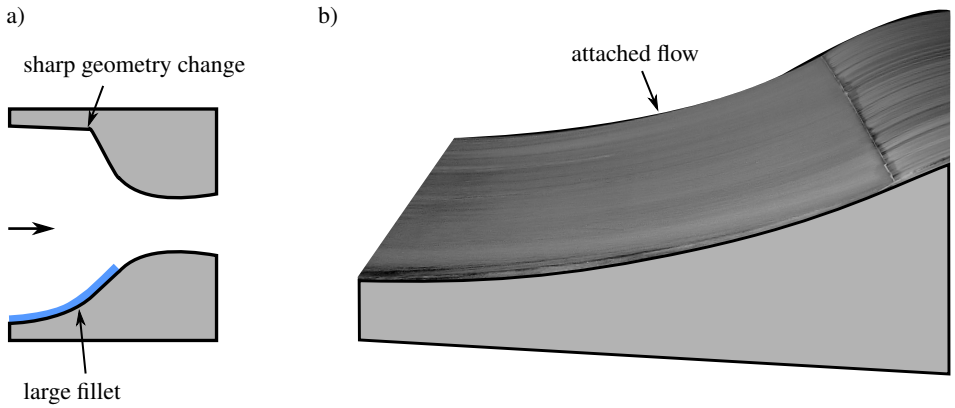

Figure 5.18: Experimental tests on the large fillet geometry on the tunnel floor, showing: a) the installed geometry, b) oil-flow visualisation on the surface marked blue in a.

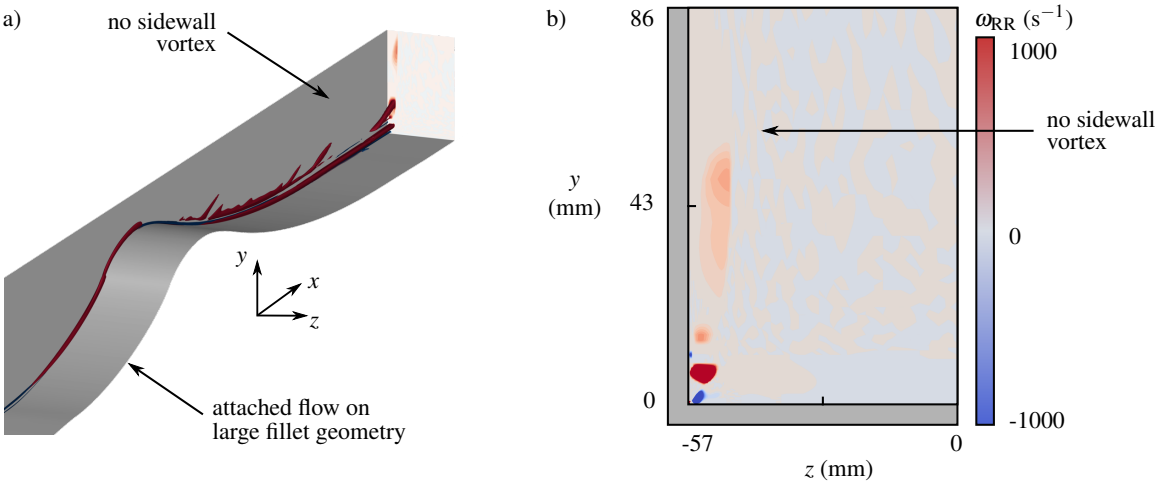

Figure 5.19: Rigid-body component of vorticity, ω_{RR}, from computations of the large fillet geometry: a) isometric view, showing isosurfaces of ω_{RR} in red and blue; b) cross-sectional distribution at $x = 120$ mm. One quarter of the wind tunnel is shown.

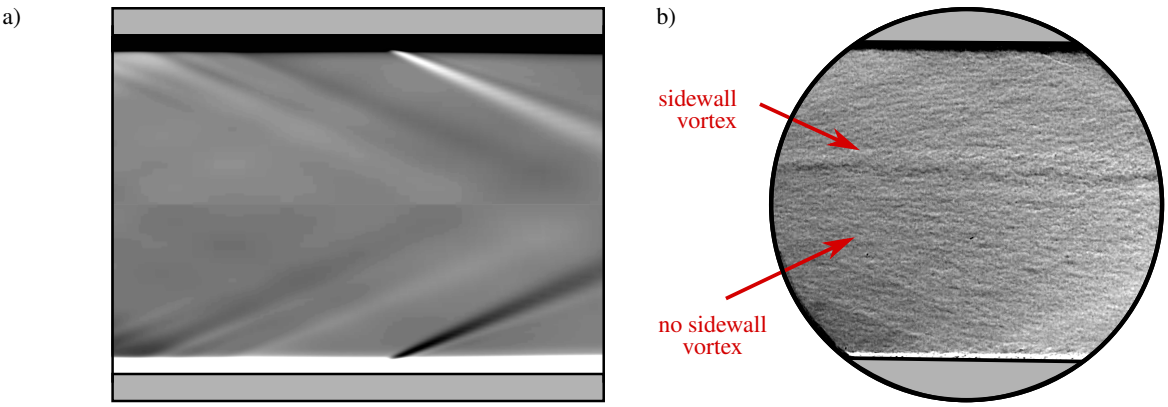

Figure 5.20: Schlieren visualisation of the flow with the large fillet geometry: a) numerical schlieren from computations (floor and ceiling: large fillet); b) experimental schlieren image with 1.1 μs exposure time (floor: large fillet, ceiling: unfilleted). Flow is from left to right.

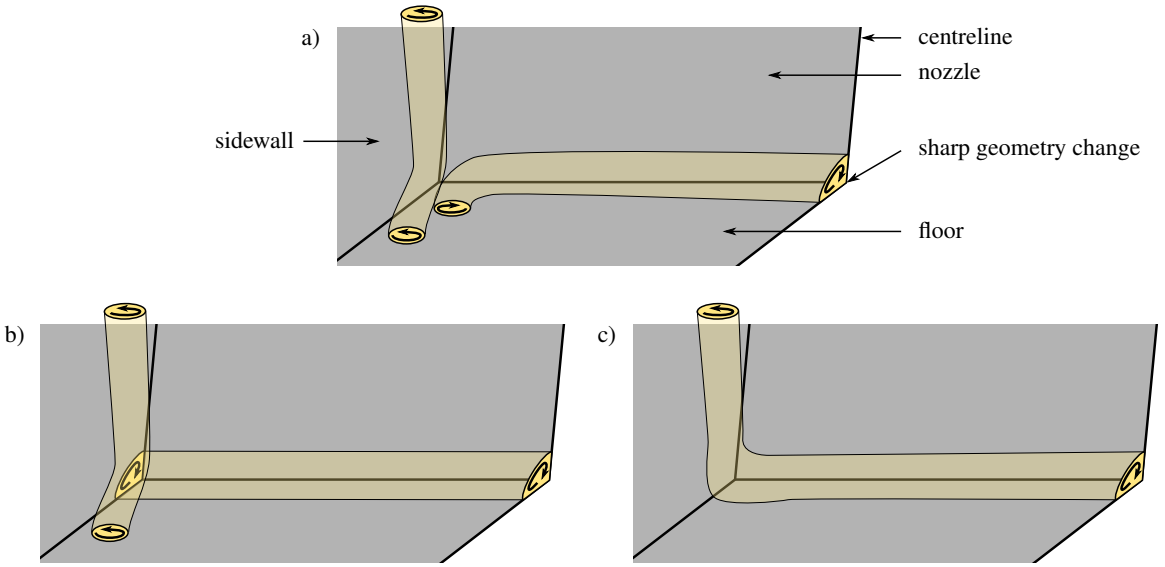

Figure 5.21: Three plausible vortex topologies associated with the separated region upstream of the nozzle.

tunnel ceiling (figure 5.18a). However, the lower half of the flow does not exhibit any streamwise features. Both computations and experiment therefore conclude that the elimination of the separation ahead of the nozzle coincides with the disappearance of the streamwise-coherent structures in the test section, providing further evidence of the close causal link between these flow features.

5.2.2 Characterisation of the sidewall vortices

The experiments and computations both provide strong evidence that the sidewall vortices are intimately coupled to the separated region ahead of the nozzle. However, a comprehensive understanding of the sidewall vortices requires more information. In particular, it is important to ascertain more precisely where the vortices come from, what impact they have on the wider flow, and whether they are likely to be found in other facilities.

Origin of the vortices

It is difficult to extract the precise production mechanism without higher-fidelity simulations or optical access to the upstream half of the test section. However, the well-established relationship between three-dimensional separation and vortex production provides some clues to the origin of the studied feature.

On the tunnel centre span, the recirculating region of the separation bubble can be treated as a spanwise-aligned vortex. In physical three-dimensional flows, it is extremely rare (or perhaps even impossible) for this vortex to start and end on the sidewalls to form a closed separation [89]. Instead, figure 5.21 shows a range of alternative flow topologies – these all result in the generation of

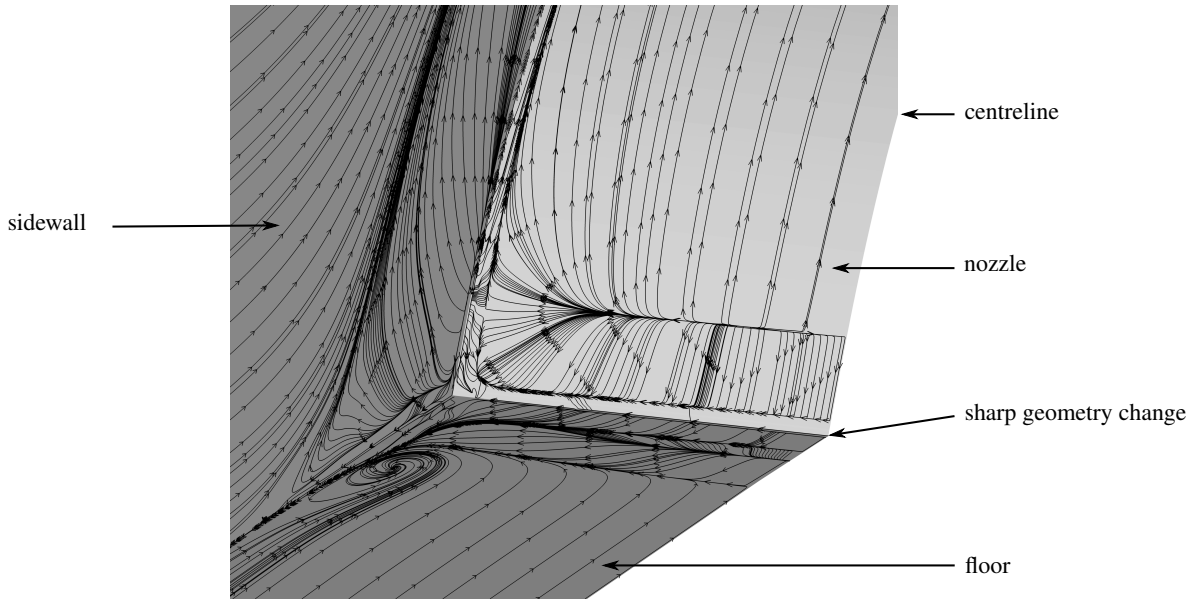

Figure 5.22: Skin-friction line topology ahead of the nozzle region, extracted from RANS computations.

a vortex along the sidewall. The sense of the generated vortex in all three cases is consistent with the computational vorticity distribution from figure 5.13.

The skin-friction line topology ahead of the nozzle is extracted from computations, as shown in figure 5.22. This structure appears to be consistent with the flow topology suggested in figure 5.21b, with critical points relating to vortices on the floor and at the sidewall–floor–nozzle corner. It is important to note, however, that the ability to accurately predict such vortex topologies using steady RANS simulations with a relatively simple turbulence model has not been verified, and so the flow structure in figure 5.22 may not be physical. Indeed, of the three possible vortex topologies, the absence of foci in the experimentally-observed oil-flow patterns on the floor (figures 5.14b and 5.15c) is consistent only with figure 5.21c. However, to confirm this flow topology requires a critical-point analysis of surface oil-flow information from the sidewalls, which is experimentally challenging and outside the scope of this study.

Whatever the precise flow topology, a streamwise-aligned vortex is produced by the separated region. This vorticity is advected by the flow through the two-dimensional nozzle, and is intensified by a vortex-stretching mechanism due to the large streamwise velocity gradient [136]. Figure 5.13 shows that, as the vortex passes through the nozzle, it also moves away from the floor towards the tunnel centre height. This trajectory is consistent with the secondary flows within the sidewall boundary layers, which were described in section 5.1.

Impact of the vortices

The influence of sidewall vortices on the overall flow can be assessed by comparing the streamwise velocity distribution between cases with and without the vortices. Unfortunately, the locations of the

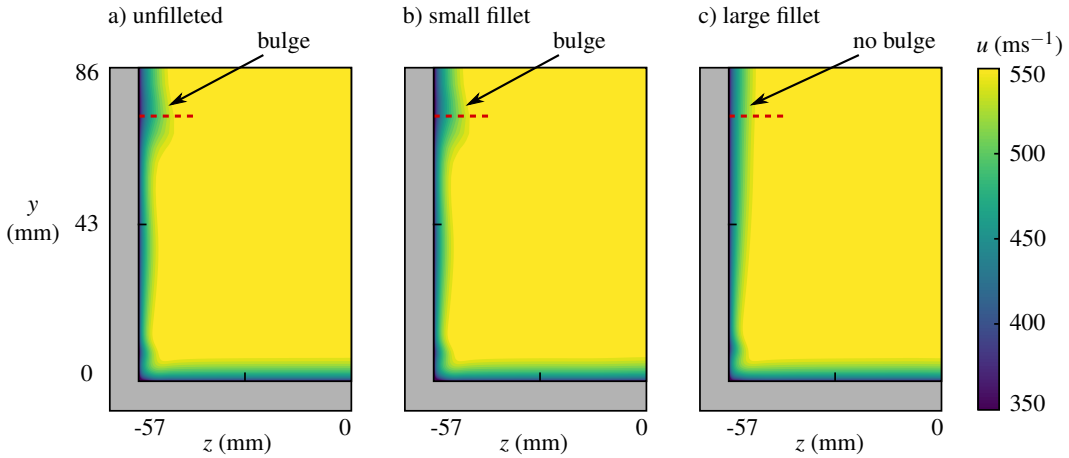

Figure 5.23: Streamwise velocity at $x = 120$ mm, from RANS simulation of wind tunnel for a) the unfilleted geometry, b) the small fillet geometry, and c) the large fillet geometry. One quarter of the wind tunnel cross-section is shown. The dashed line indicates the location of boundary-layer profiles in figure 5.24.

vortices correspond to the gaps in sidewall boundary-layer data from figure 5.3 where high-quality velocity measurements cannot be obtained. As a result, experiments cannot be used for this analysis. Instead, the likely impact of the sidewall vortices is studied using computational data provided by the Air Force Research Laboratory.

The computed streamwise velocity distribution in the two cases which do exhibit sidewall vortices (figure 5.23a for the unfilleted geometry and figure 5.23b for the small fillet) appear to be very similar to each other. These flows both exhibit a bulge in the sidewall boundary layer close to the tunnel centre height, at the location of the sidewall vortices.

Figure 5.23c displays the flow for the large fillet geometry, where there are no sidewall vortices. The core flow, floor boundary layer and corner regions have not changed significantly from figures 5.23a and 5.23b. However, the sidewall boundary layer does look slightly different. This difference is highlighted by comparing velocity profiles at $y = 70$ mm in figure 5.24, and the associated integral boundary-layer properties (table 5.1). These parameters show that the shape factor is about 1.4 for all three geometries, but there is a significant variation in the boundary-layer thickness. In fact, the sidewall boundary layer is 37% thicker when the sidewall vortices are present (the unfilleted and small fillet setups), as compared to when the large-radius fillet is used. These differences in sidewall boundary-layer profile highlight the impact of the sidewall vortices on the flow.

Therefore, the sidewall vortices are not a very strong feature, having only a local effect on the flow. This explains why these flow structures have not been noticed earlier despite the widespread use of interchangeable, two-dimensional nozzle geometries, which generally feature sharp corners. Nevertheless, these vortices are significant since they contribute to flow non-uniformity, and since they might make the local region of the sidewall boundary layer more susceptible to flow separation in adverse pressure gradients.

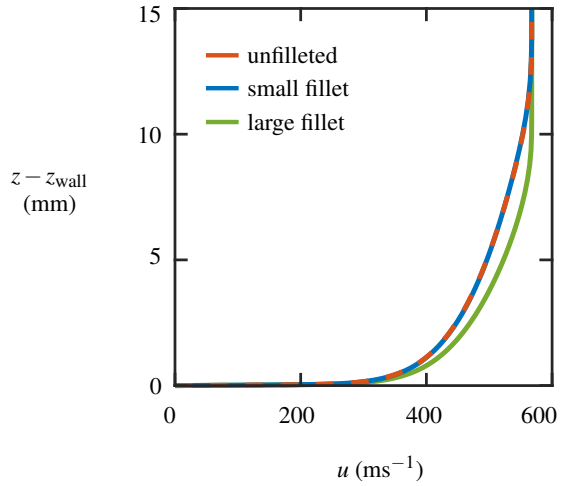

Figure 5.24: Sidewall boundary-layer profile, at $y = 70$ mm and $x = 120$ mm, from the computations of the three setups.

Table 5.1: Boundary-layer parameters for sidewall boundary-layer profiles, at $y = 70$ mm and $x = 120$ mm. These correspond to the velocity data presented in figure 5.24.

	boundary-layer thickness δ (mm)	displacement thickness δ_i^* (mm)	momentum thickness θ_i (mm)	shape factor H_i
unfilleted geometry	9.27	1.54	1.09	1.41
small fillet geometry	9.28	1.53	1.09	1.41
large fillet geometry	6.70	1.11	0.78	1.42

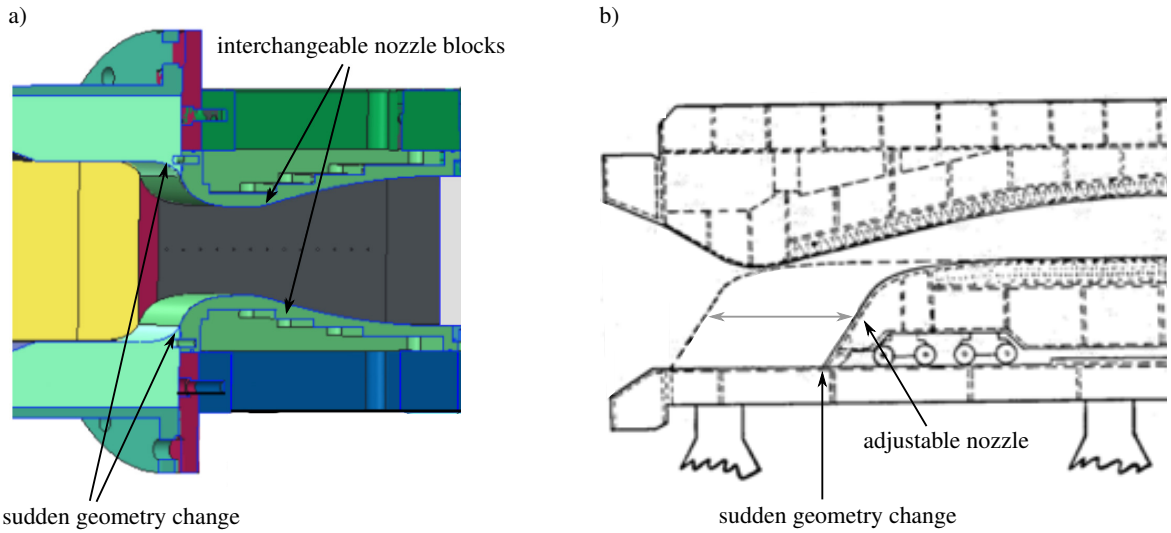

Figure 5.25: Examples of sharp geometry changes ahead of the nozzle in other facilities: a) CAD model of the supersonic wind tunnel at Imperial College [137]; b) schematic diagram of the high Mach number circuit in the Unitary Plan Wind Tunnel at NASA Langley [138].

Similar features in other facilities

The separation upstream of the nozzle, and thus the presence of sidewall vortices, are now known to be related to the sharp geometry changes at the upstream end of the two-dimensional nozzle blocks. Many supersonic wind tunnels use either interchangeable nozzle blocks, similar to the current facility, or adjustable nozzles [83]. Geometric constraints often require a sudden geometry change at the start of the nozzle in these types of facility. Examples of these sharp corners are evident in the interchangeable-nozzle supersonic wind tunnel at Imperial College London (figure 5.25a) and in the adjustable-nozzle Unitary Plan Wind Tunnel at NASA Langley (figure 5.25b).

Wind tunnels with similar sharp geometry changes are expected to be particularly susceptible to separation of the incoming boundary layer, and therefore production of sidewall vortices. Indeed, schlieren images which show similar streamwise-coherent features can be found in a number of publications originating from such facilities [132–134]. This type of sidewall vortex can be avoided by replacing sharp geometry changes in the upstream, subsonic section of the wind tunnel nozzle with a more gentle contour. When these modifications are not feasible, the interpretation of schlieren images and other data should consider the possible presence and impact of the sidewall vortices.

5.3 Structure of the corner boundary layer

The central feature of this thesis is the structure of the corner boundary layer in supersonic channel flows. A survey of the relevant literature in section 2.3 shows that subsonic corner flows typically feature a stress-induced counter-rotating pair of streamwise vortices [60]. These transfer momentum between the core flow and the boundary layers. In supersonic channel flows, the presence of

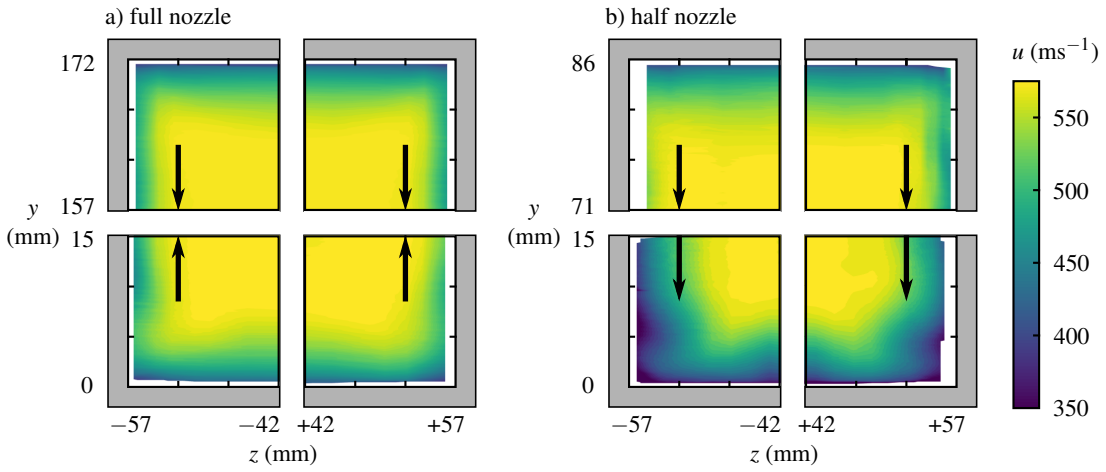

Figure 5.26: The streamwise velocity (u) in 15 mm $\times$15 mm regions around all four tunnel corners for a) the full nozzle setup, and b) the half nozzle setup. These are measured at $x = 120$ mm. Solid arrows indicate the direction of sidewall secondary flows.

streamwise corner vortices has been confirmed experimentally. However, the precise flow topology is somewhat inconsistent between different studies, as illustrated in figure 2.15 [76, 14, 13]. A key difference between the corresponding experimental setups is the geometry of the nozzles used to generate the supersonic flow (figure 2.16). It is therefore worth exploring whether the nozzle geometry may have an effect on the structure of the corner boundary layers. This can be achieved by studying and comparing the flow from both the full nozzle and the half nozzle setups.

In particular, section 5.1 has detailed mechanisms through which the pressure gradients induced by the nozzle geometry set up bulk vertical velocities within the sidewall boundary layers. It is therefore useful to examine the measured corner boundary-layer data in the context of these secondary flows.

Effect of sidewall secondary flows

The streamwise velocity component in the 15 mm $\times$15 mm regions around each corner of the channel is measured at $x = 120$ mm. This velocity distribution is shown for the full nozzle setup in figure 5.26a and for the half nozzle setup in figure 5.26b. There appear to be two distinct flow topologies. All four corners of the full nozzle setup and the top corners of the half nozzle configuration are similar in nature. Meanwhile, the bottom corners with the half nozzle have a thicker sidewall boundary layer and, therefore, lower-momentum flow in this region.

Note that the floor/ceiling boundary-layer thickness is approximately the same for all flows shown in figure 5.26. The differences in the corner boundary layer instead depend primarily on the thickness of the constituent sidewall boundary layer. This property is, in turn, directly determined by the sidewall secondary flows, which were introduced in section 5.1.

All corners in the full nozzle setup and the bottom corners of the half nozzle setup have the sidewall secondary flow directed vertically away from the corner (figures 5.6c and 5.7c). This bulk flow transports the low-momentum fluid in the sidewall and corner boundary layers away from the

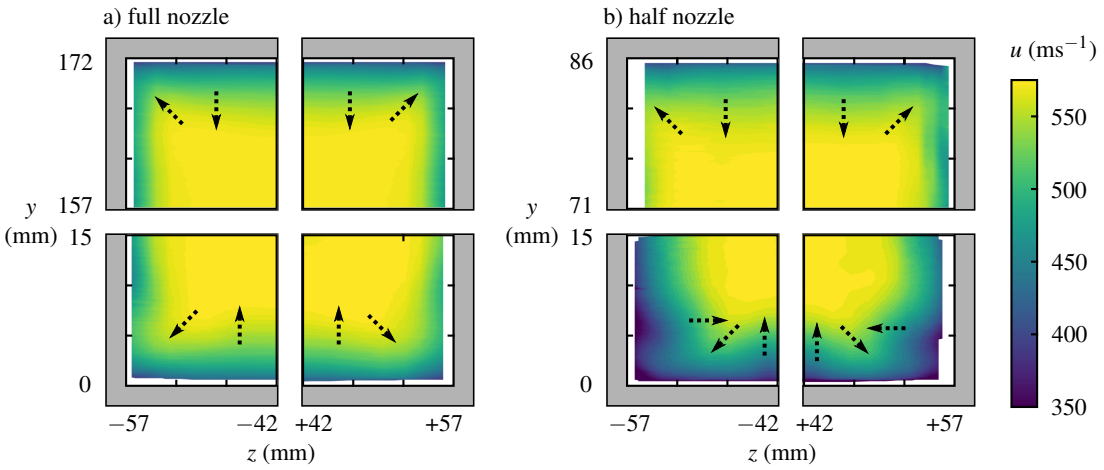

Figure 5.27: The streamwise velocity (u) in 15 mm $\times$15 mm regions around all four tunnel corners for a) the full nozzle setup, and b) the half nozzle setup. These are measured at $x = 120$ mm. Dashed arrows indicate local regions of momentum transfer.

region. As a result, in the vicinity of these corners, the sidewall boundary layers are thin. The resultant small corner boundary-layer thickness for these cases are in agreement with the velocity measurements shown in figure 5.26.

The bottom corners of the channel in the half nozzle setup are different, however. Here, the sidewall secondary flows are in the downwards direction, and so they advect low-momentum fluid into the corner region. This results in an increase in the thickness of the boundary layer, consistent with the experimental data from figure 5.26b.

The thickness of the measured corner boundary layer can therefore be understood by considering the momentum transferred into and out of these regions by the sidewall secondary flows. However, a closer examination of figure 5.26 reveals that the boundary-layer thickness is not the only defining feature of this flow. In particular, there are some regions where the high momentum core flow persists into the boundary layer and others where low momentum fluid extends locally into the core. In figure 5.27, these regions are indicated by dashed arrows which show the approximate location and direction of these momentum transfers. It is thought that these are caused by the streamwise vortices embedded in the viscous corner flow regions.

Unfortunately, the inherent limitations of a two-component LDV setup prevent direct identification of the the vortices themselves. Instead, it is necessary to infer the vortex topology by considering the regions of momentum transfer between the core flow and the boundary layers. This indirect method requires correct interpretation of the shape of the complex corner boundary layer, and so the inferred vortex topology may be prone to errors. As a result, the following discussion should be treated with some degree of caution.

The momentum transfers in the bottom corners of the half nozzle setup would be consistent with the presence of two counter-rotating vortices, while the other measured corners appear to contain only a single primary vortex residing on the floor or ceiling. These different inferred vortex

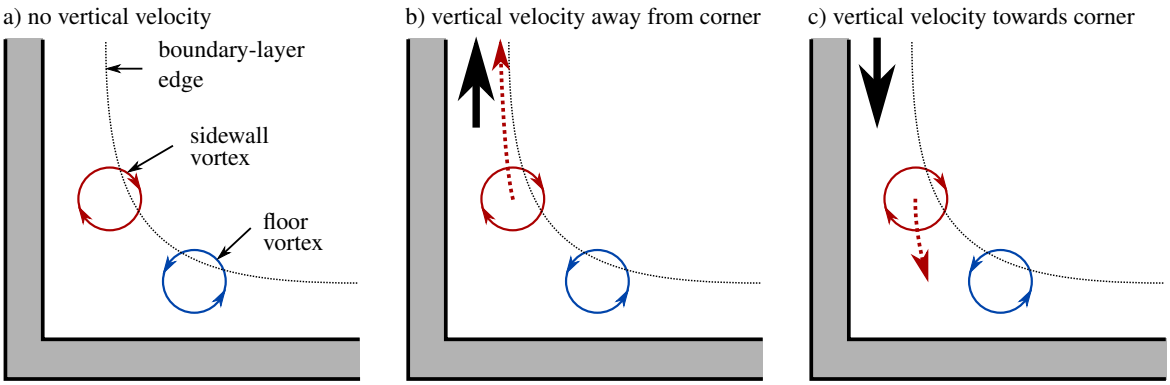

Figure 5.28: Schematic diagram indicating the effect of sidewall secondary flows on corner vortex structure. The counter-rotating vortex pair with: a) no vertical bulk flow; b) vertical velocity away from the corner, and c) vertical velocity towards the corner.

structures are another feature which can be explained by the sidewall secondary flows. In a supersonic channel with turbulent wall boundary layers but no vertical velocities, we expect the Reynolds stress anisotropy in the corner boundary layer to generate two counter-rotating streamwise-aligned vortices, symmetric about the corner bisector [60, 12]. The topology of these two vortices, referred to as the sidewall vortex and the floor vortex (figure 5.28a), are modified when the sidewall secondary flows are present. The floor vortex, residing outside the sidewall boundary layer, is largely unaffected by any secondary flow on the sidewall. However, the sidewall vortex is influenced by the local vertical velocities – the secondary flows cause advection of any pre-existing sidewall vortex, as well as the continuously-produced 'sidewall vorticity'.

The case with vertical velocity directed away from the corners is shown schematically in figure 5.28b. This corresponds to the full setup and the top corners of the half setup shown in figure 5.27. Any vorticity near the tunnel sidewall is immediately advected away from the corner by the vertical flow. This prevents the formation of a stable sidewall vortex, and so the corner region contains just one (floor) vortex.

Figure 5.28c presents the opposite case, when the sidewall secondary flow is directed into the corner. This represents the bottom corners of the half setup, shown in figure 5.27b. Now, the sidewall vortex is advected towards the corner by the bulk vertical velocity. Therefore, the region retains the counter-rotating vortex pair.

These findings may explain the apparent inconsistencies in previous experimental data. Davis *et al.*, who observed a symmetric vortex pair, conducted their study in a square duct suspended within the core flow [76]. The observed vortex topology can therefore be attributed to the absence of sidewall secondary flows, analogous to figure 5.28a. Meanwhile, the study by Peltier *et al.* was conducted in a full nozzle setup [14], where the sidewall secondary flows are expected to be away from the corners. Their observed single-vortex structure is consistent with the flow field suggested in figure 5.28b. Finally, the third study, conducted by Morajkar *et al.*, focused measurements on the bottom corners of a half nozzle setup. The secondary flows in the sidewall boundary layers are predicted to be directed

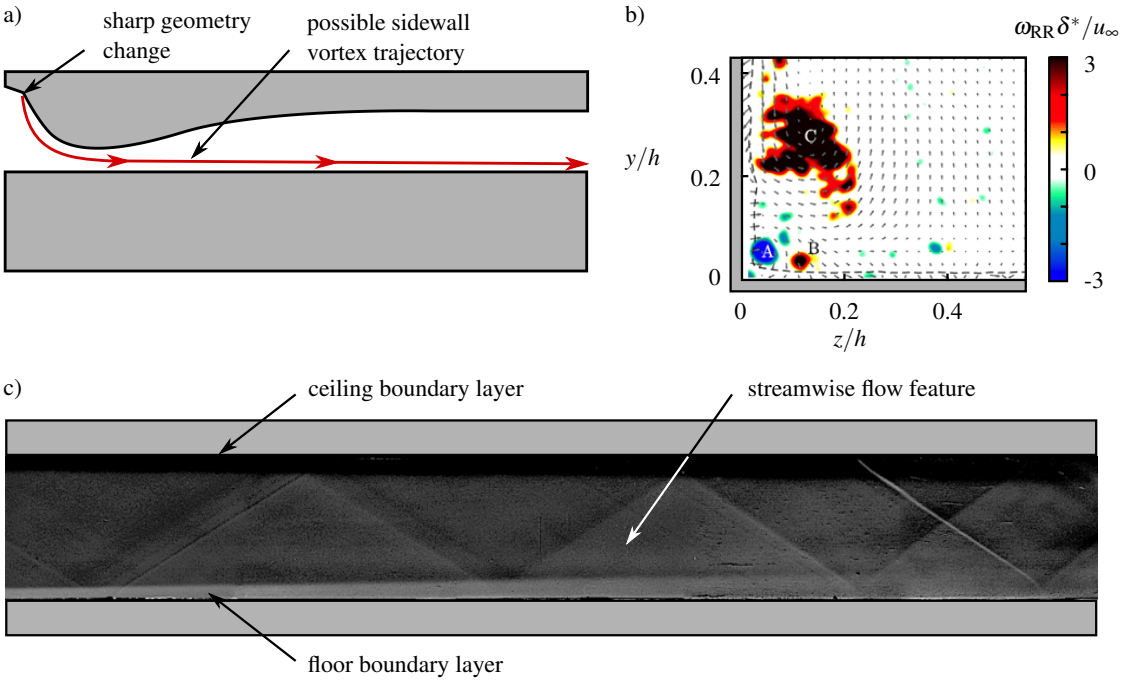

Figure 5.29: a) Schematic diagram indicating a possible trajectory for any sidewall vortex in a half nozzle geometry. b) Rigid-body component of vorticity measured by Edelman *et al.*, using particle image velocimetry in the bottom corner regions of the University of Michigan Direct Connect Isolator Facility [139]. c) Enhanced schlieren image of the flow in this facility.[2] Flow is from left to right.

towards these corner regions. The vortex pair measured in this facility is asymmetric about the corner bisector, in agreement with figure 5.28c.

Effect of sidewall vortices

Whilst the sidewall secondary flows appear to explain current and past data measured in the corner boundary layers, it is perhaps useful to consider another possible influence of nozzle geometry on the flow in these regions. This influence is related to the sidewall vortices associated with separation ahead of the nozzle, as introduced in section 5.2. For a half nozzle setup, sidewall vortices might be generated from any sharp geometry change ahead of the contoured ceiling nozzle. In the test section, the trajectory of the vortices, indicated in figure 5.29a, could lie close to the wind tunnel floor, i.e. in the immediate vicinity of the corner boundary layers. These vortices, which are known to transfer momentum between the core flow and the boundary layers, may therefore affect the momentum contained within the corner regions.

Such a phenomenon is expected to only affect those half nozzle geometries which feature a sharp change on the contoured ceiling nozzle. This is nicely illustrated using data from a recent study at the University of Michigan Direct Connect Isolator Facility, conducted by Edelman *et al.* [139]. This wind tunnel has a half nozzle geometry which produces a Mach 2 flow. Figure 5.29b shows the

[2]Private communication with Louis Edelman at the University of Michigan.

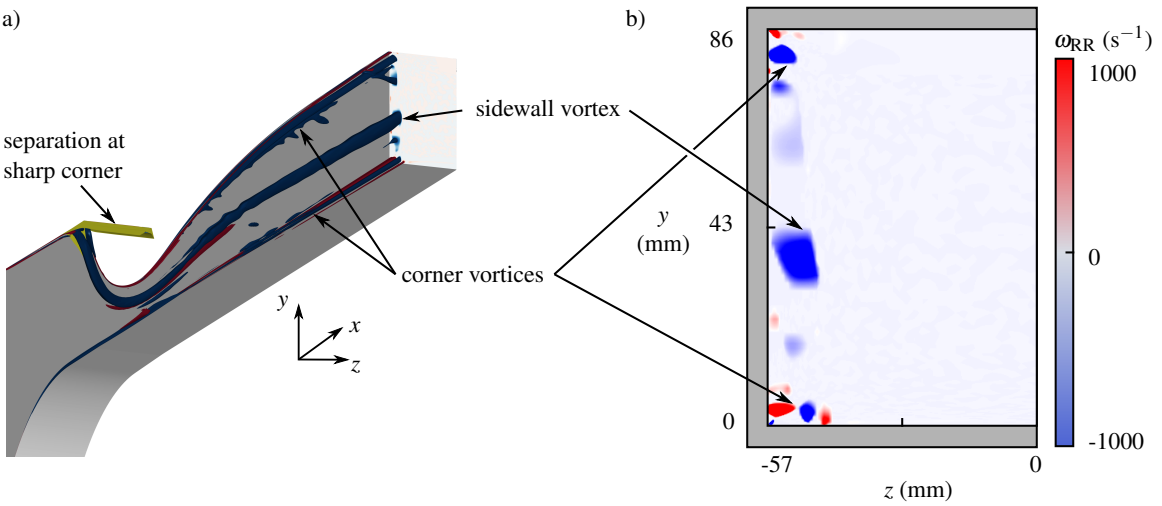

Figure 5.30: Rigid-body component of vorticity, ω_{RR}, from computations of the half nozzle setup: a) isometric view, showing isosurfaces of ω_{RR} in red and blue, with a region of separation highlighted in yellow; b) cross-sectional distribution at $x = 120$ mm. One half of the wind tunnel is shown.

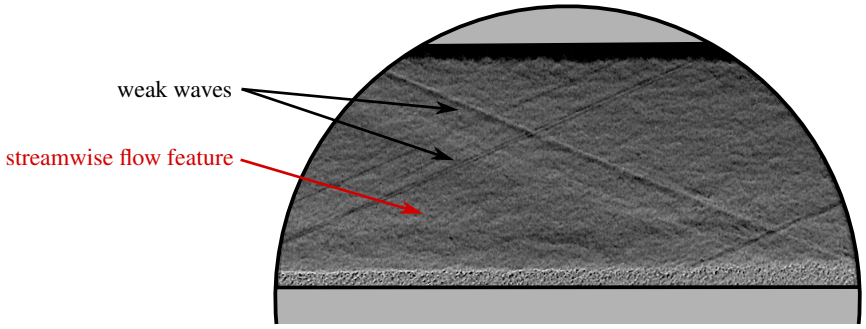

Figure 5.31: Schlieren visualisation of the flow using the half nozzle setup, with 1.1 μs exposure time. Flow is from left to right.

vorticity distribution obtained using particle image velocimetry in one of the two bottom corners of the wind tunnel. In this figure, the stress-induced corner vortices are labelled A and B. There is also a third, much larger vortex (C), whose origin is more difficult to explain. The authors suggest that this vortex was originally a corner vortex in the top corners of the wind tunnel, which has been transported downwards by the sidewall secondary flows in this direction.

However, the proposed mechanism does not explain the substantially larger size of vortex C compared to the other vortices. Instead, it is similar in appearance to the vorticity distribution for the sidewall vortex in figure 5.13b. In order to test whether vortex C might correspond to the features discussed in section 5.2, schlieren visualisation of the flow is analysed. The enhanced image, shown in figure 5.29c, appears to show a streamwise-coherent feature in the test section, analogous to the sidewall vortices. This is at the same vertical location as vortex C and its sense is consistent with the

sidewall vortices from section 5.2, providing some evidence that the origin of this vortex might in fact be a separated flow region ahead of the contoured ceiling nozzle.

Therefore, the sidewall vortices might also have an impact on the corner boundary-layer structure, but only in the somewhat restricted case of a half nozzle setup which also exhibits sharp geometry changes at the nozzle entrance. This leads to a potential fourth vortex topology in the corner boundary layers, in addition to the three cases discussed previously (figures 2.15a–c). This new corner boundary-layer structure contains three vortices: the two stress-induced nozzle vortices, as well as a sidewall vortex which originates upstream of the nozzle.

The half nozzle setup in the Cambridge facility does indeed exhibit a similar sharp geometry change ahead of the ceiling nozzle block, and so it is worth determining whether such a topology applies to the current study. The rigid-body component of vorticity is extracted from RANS computations of the half nozzle setup performed by the US Air Force Research Laboratory. This quantity is shown in figure 5.30a. The sidewall vortex can be clearly identified. Whilst the sidewall secondary flows do cause the vortex to move down towards the floor, figure 5.30b shows that it ends up embedded in the boundary layer at around $y = 40$ mm, i.e. well outside the corner boundary layer. Indeed, the 1.1 μs exposure time schlieren image of the flow in figure 5.31 also shows that the streamwise-coherent flow structure is well away from the floor boundary layer. Since the sidewall vortex is so far away from the corner, any associated perturbations to the sidewall boundary layer are not thought to affect the momentum within the corner region itself.

The corner boundary layers of supersonic channel flows have been shown in this section to depend closely on the geometry of the two-dimensional nozzle. The thickness of the corner boundary layer, which strongly affects the momentum in this region, is determined primarily by the local sidewall boundary-layer thickness. This property is influenced by secondary flows within the sidewall boundary layers that are induced by pressure gradients in the nozzle. The structure of the corner boundary layer, which can slightly modify the momentum, is governed by the vortex topology. This topology is determined by the sidewall secondary flows and, in some very specific cases, by sidewall vortices generated by sharp geometry changes ahead of the nozzle.

5.3.1 Validation of corner boundary-layer computations

The general structure of the corner boundary layer in supersonic channel flows is now somewhat better understood. It is suggested that apparently inconsistent flow topologies in the literature (figure 2.15a–c) can be attributed to nozzle geometry-induced effects. Therefore, any computations which contain knowledge of the nozzle geometry might be able to capture these differences, enabling validation to be performed on the corner boundary layers.

There are two distinct corner flow structures that can be used for these validation tests, i.e. the bottom corners of the two nozzle setups. For each case, the streamwise velocity in the 15 mm $\times$15 mm

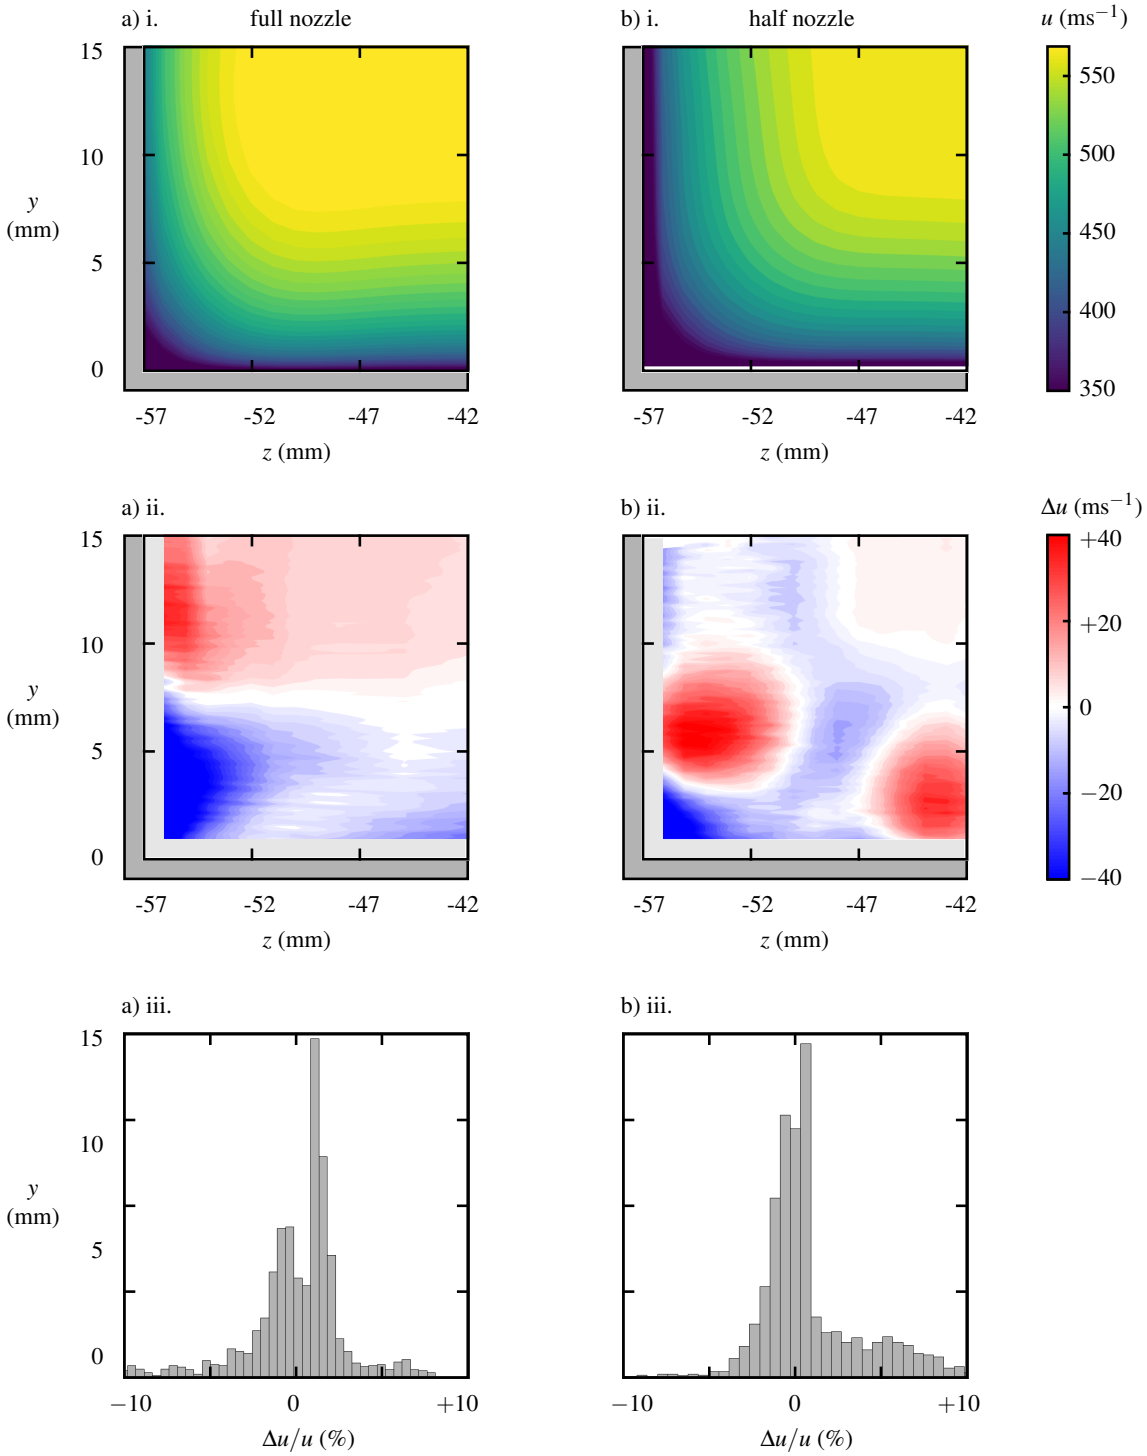

Figure 5.32: Bottom corner regions for a) the full nozzle setup, and b) the half nozzle setup: i. streamwise velocity from computations using the Spalart–Allmaras model without the quadratic constitutive relation; ii. difference between streamwise velocity from computations and equivalent experimental data; iii. histogram of percentage differences in streamwise velocity between computations and experiment.

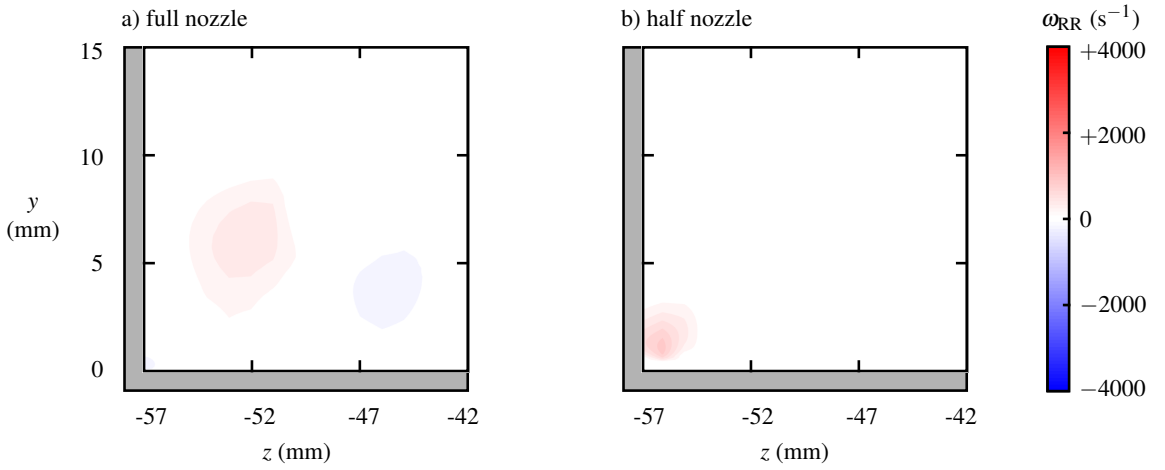

Figure 5.33: The rigid-body component of streamwise vorticity, ω_{RR}, in the bottom corner region from computations using the Spalart–Allmaras model without the quadratic constitutive relation for a) the full nozzle setup, and b) the half nozzle setup.

around the corner is extracted from the computations of the wind tunnel flow.[3] This velocity distribution is shown in figure 5.32a-i for the full nozzle setup and in figure 5.32b-i for the half nozzle arrangement.

A comparison with the equivalent experimental data (figure 5.27) shows that the constituent floor and sidewall boundary-layer thicknesses are correctly determined by these computations. The agreement of these properties can be attributed to the careful calibration of the simulations. However, for both nozzle setups, the computations show a smooth transition between the sidewall and the floor boundary layers. This contrasts with the more complex shape of the corner boundary layer from experiment, which is believed to be perturbed by embedded vortices. Indeed, figures 5.32a-ii and 5.32b-ii show that the departures from the experimental velocities are as large as 40 ms^{-1}, or 7% of the freestream velocity. From the corresponding histograms in figures 5.32a-iii and 5.32b-iii, a typical error can be calculated, defined such that 95% of data points do not exceed this deviation from the velocity measurements. This error is determined to be 6.7% with the full nozzle (figure 5.32a-iii) and 6.8% with the half nozzle (figure 5.32b-iii), much larger than the discrepancies ($\approx 3\%$) estimated in section 4.2.3 for the well-predicted regions of the flow field.

In order to further investigate whether the error can be attributed to a lack of vortices in this region, the rigid-body component of streamwise vorticity is extracted from the computations. Figures 5.33a and 5.33b for the full nozzle setup and the half nozzle setup, respectively, show that the corner region contains no vorticity greater than 500 s^{-1} $\approx 0.006\, u_\infty/\delta$, where u_∞ is the freestream velocity and δ is the nominal local boundary-layer thickness. This most likely is because the Spalart–Allmaras turbulence model used in these computations is based on a linear eddy-viscosity model. Indeed, equations 2.4 and 2.5 showed that this type of model cannot generate the streamwise vortices produced

[3]The data from these computations were provided by the Air Force Research Laboratory. Analysis was performed by the author.

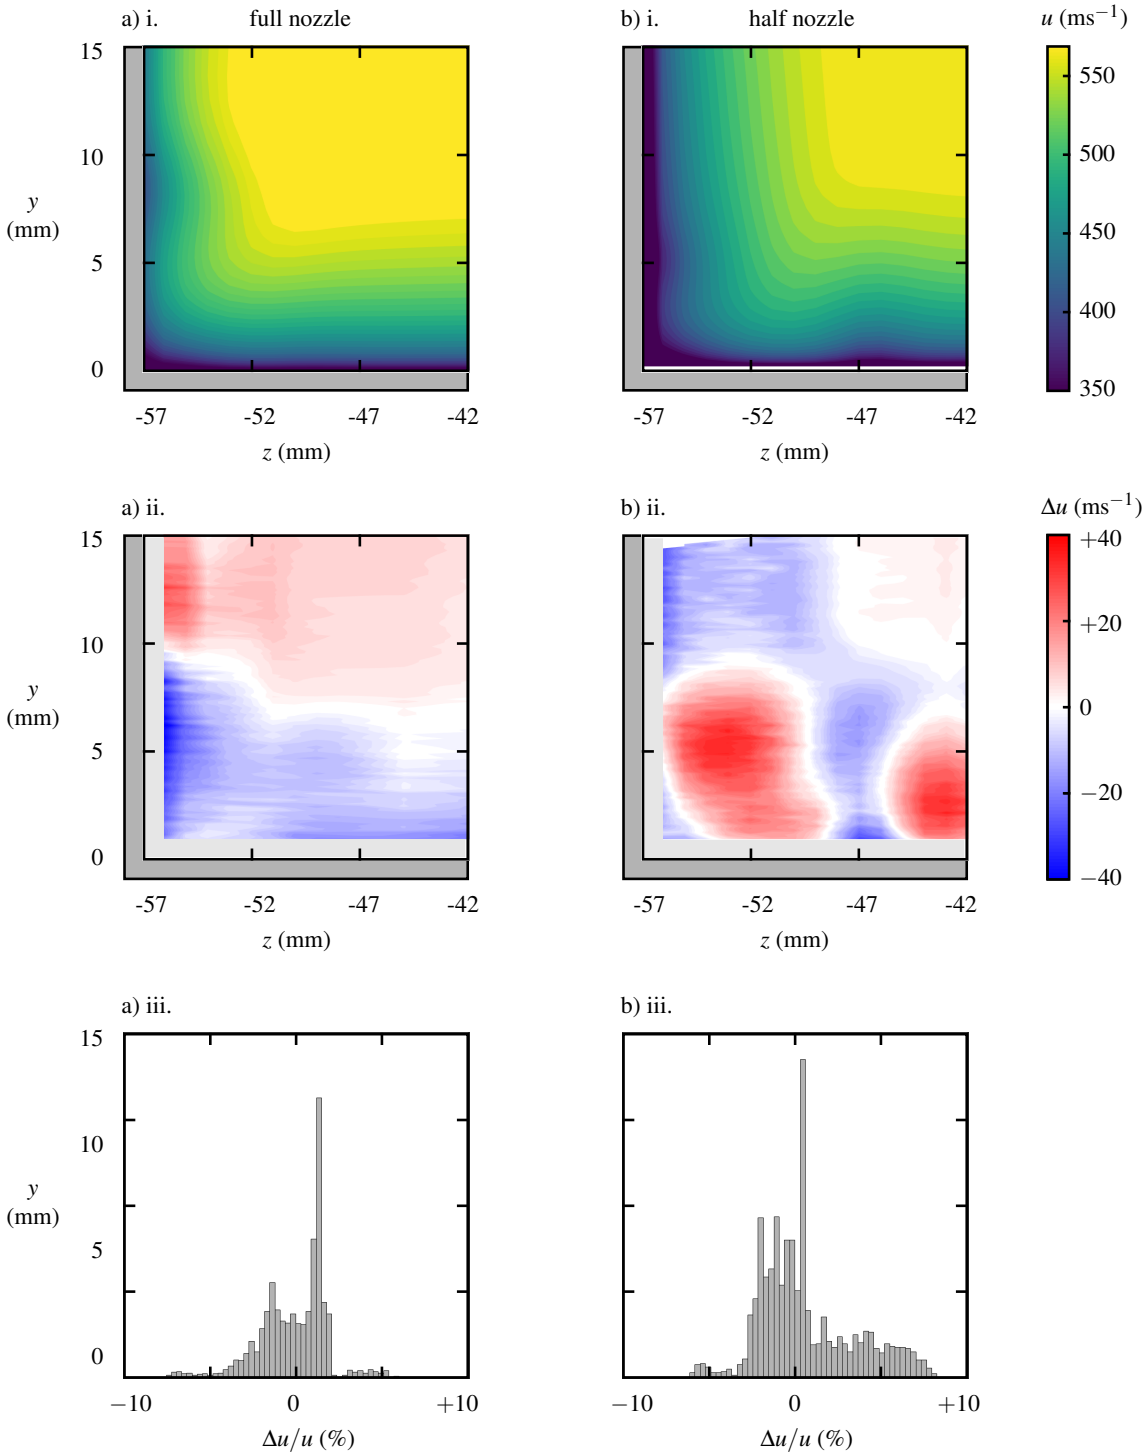

Figure 5.34: Bottom corner regions for a) the full nozzle setup, and b) the half nozzle setup: i. streamwise velocity from computations using the Spalart–Allmaras model with the quadratic constitutive relation; ii. difference between streamwise velocity from computations and equivalent experimental data; iii. histogram of percentage differences in streamwise velocity between computations and experiment.

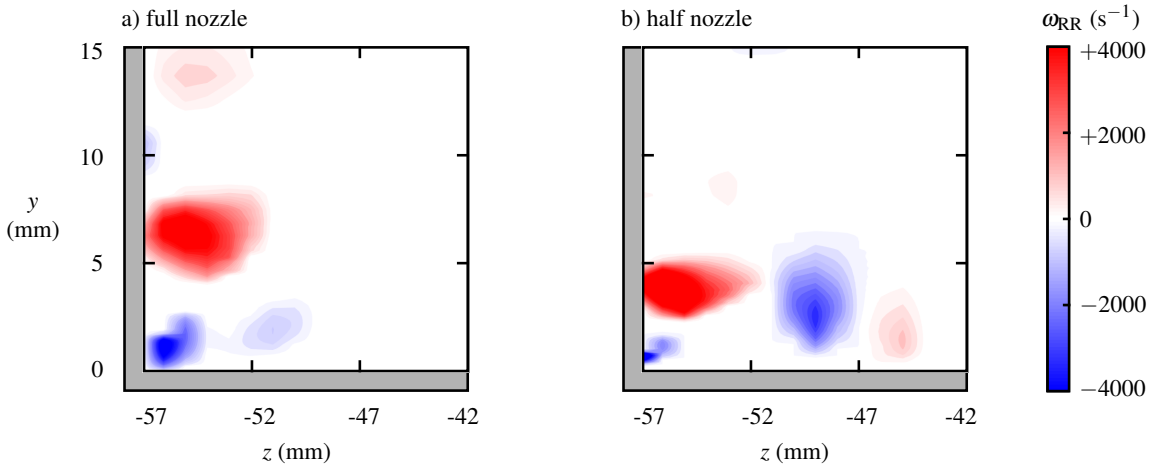

Figure 5.35: The rigid-body component of streamwise vorticity, ω_{RR}, in the bottom corner region from computations using the Spalart–Allmaras model with the quadratic constitutive relation for a) the full nozzle setup, and b) the half nozzle setup.

by turbulent stress anisotropies, and so the poor prediction of the flow in this region is somewhat unsurprising.

Equivalent computations were also performed by the Air Force Research Laboratory using the quadratic constitutive relation [71].The streamwise velocity distribution in the corner region, shown in figures 5.34a-i and 5.34b-i, suggests that the sidewall and floor boundary-layer thicknesses are still well-predicted. More importantly, the shape of the corner boundary layer now seems to better correspond to the measurements. Indeed, figures 5.34a-ii and 5.34b-ii suggest that the deviation from experimental values is reduced considerably, with velocity differences generally not exceeding 25 ms^{-1}. This discrepancy corresponds to 4.1% for the full nozzle setup (figure 5.34a-iii) and 5.6% for the half nozzle setup (figure 5.34b-iii).

The improvements in flow prediction are thought to be due to the corner vortices being simulated more accurately when the quadratic constitutive relation is used. This is confirmed by the rigid-body component of vorticity within the corner region, shown in figure 5.35. In contrast to figure 5.33, there are now obvious vortices within the corner boundary layer. Note that in figure 5.35a, the full nozzle setup shows a floor vortex close to the corner, and a sidewall vortex displaced upwards away from the corner. This is consistent with the mechanism proposed in figure 5.28b. Similarly, the configuration of vortices with the half nozzle (figure 5.35b) agrees well with the asymmetric vortex pair structure suggested in figure 5.28c.

Effect of changing the QCR constant

The computations with QCR in the preceding section use the recommended value of the constant, $c_{cr1} = 0.3$, which was calibrated in the outer part of an equilibrium turbulent boundary layer [71].

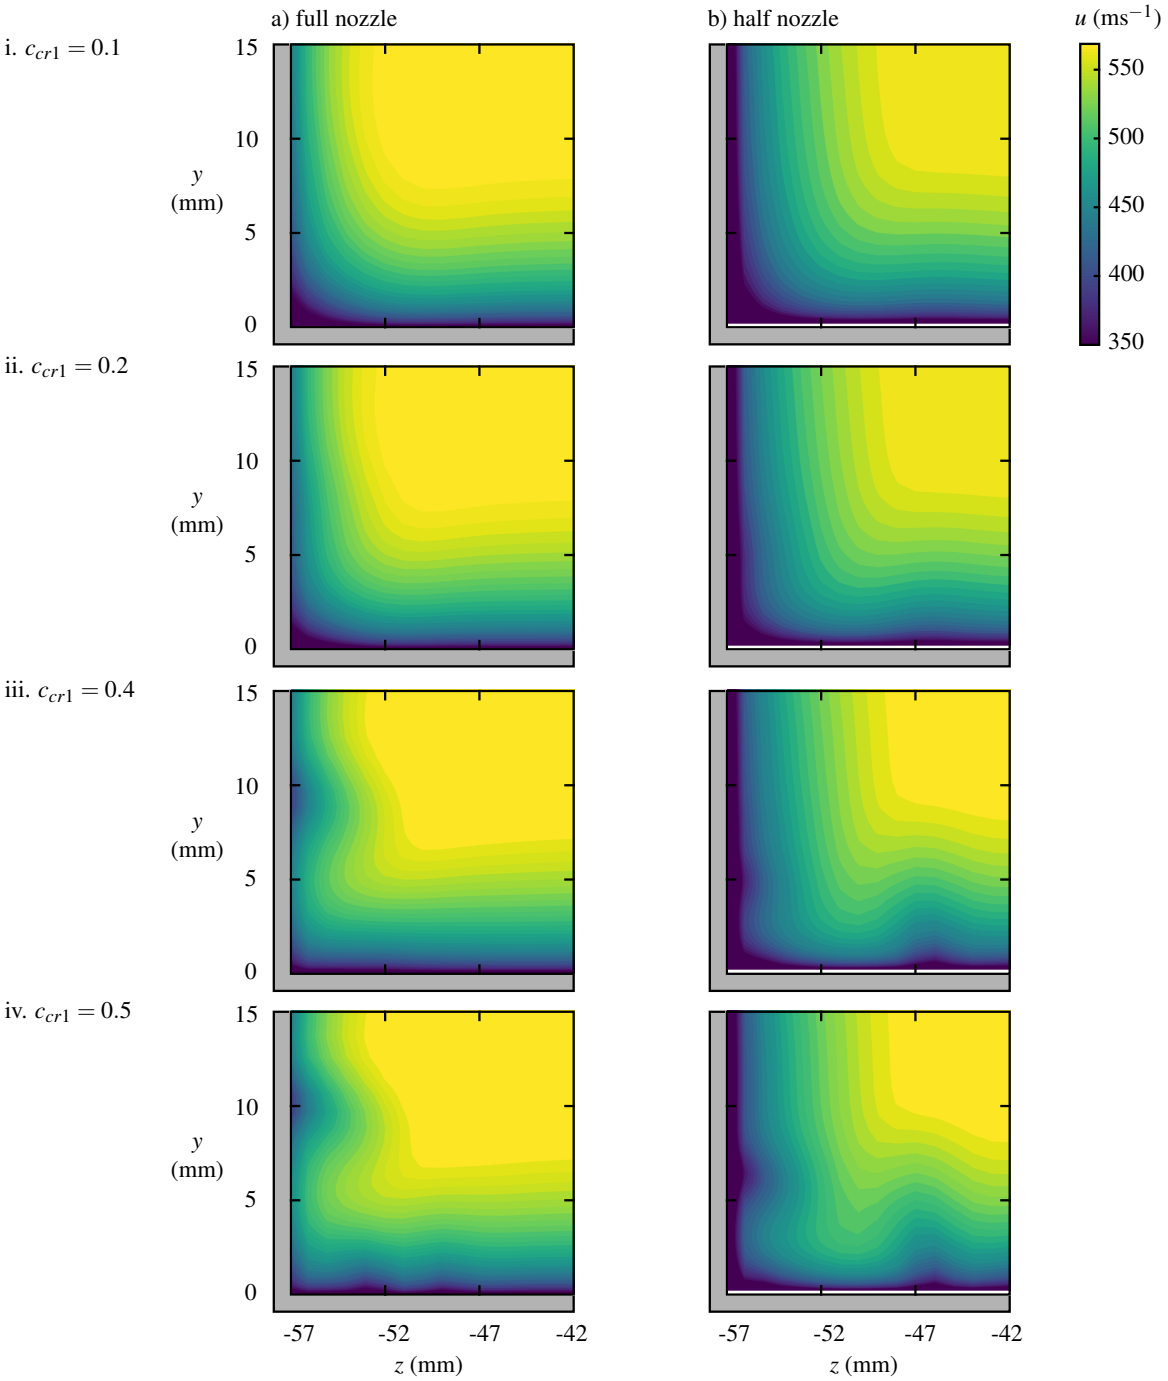

Figure 5.36: Effect of varying c_{cr1} on the streamwise velocity in the bottom corner regions for a) the full nozzle setup, and b) the half nozzle setup: i. $c_{cr1} = 0.1$; ii. $c_{cr1} = 0.2$; iii. $c_{cr1} = 0.4$; iv. $c_{cr1} = 0.5$.

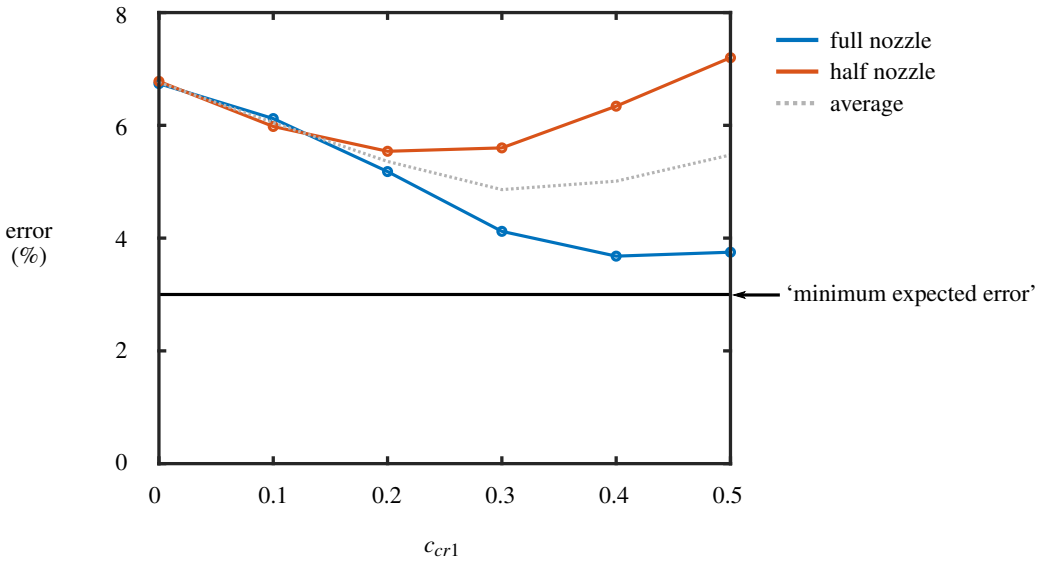

Figure 5.37: Typical error between computations and experimental data as a function of c_{cr1} for both nozzle setups. The minimum expected error is determined from velocity errors between experiment and computation in the wider flow field (section 4.2.3).

However, in practice, this 'constant' is generally used as a tuning parameter. Computations with different values of c_{cr1} were therefore investigated.[4]

The streamwise velocity distribution in the bottom corner region is presented in figure 5.32 for $c_{cr1} = 0$ and in figure 5.34 for $c_{cr1} = 0.3$. Equivalent velocity distributions with $c_{cr1} = 0.1$, 0.2, 0.4, and 0.5 are given for the full nozzle and the half nozzle setups in figures 5.36a and 5.36b, respectively. The corner boundary-layer shape shows increasing distortion as the value of c_{cr1} is increased.

In order to assess how well these computations match the experimental data (figure 5.27), a typical error is calculated for each case, defined such that 95% of data points do not exceed this deviation from the experimentally-determined velocities. Figure 5.37 plots this error as a function of c_{cr1} for both nozzle setups. In both cases, as c_{cr1} increases, the error initially gets smaller before reaching a minimum and increasing again. The minimum error occurs at $c_{cr1} = 0.4$ for the full nozzle and at $c_{cr1} = 0.2$ for the half nozzle. In both these cases, the recommended c_{cr1} value of 0.3 does not produce the best fit to experimental data but, when the error is averaged over the two nozzle setups, the minimum does indeed lie at $c_{cr1} = 0.3$. This finding therefore supports the use of the recommended value of the QCR constant in computing such supersonic corner flows.

It is important to note, however, that this finding relates specifically to the corner boundary layers of supersonic channel flows measured in the current study, and so the behaviour of coefficients is valid only for this flow field. Therefore, the results of this analysis should not be used directly to choose QCR coefficients in computations of any other types of flow. In addition, the error metric used for the current analysis (defined such that 95% of data points lie within this deviation from the measured

[4]The data from these computations were provided by the Air Force Research Laboratory. Analysis was performed by the author.

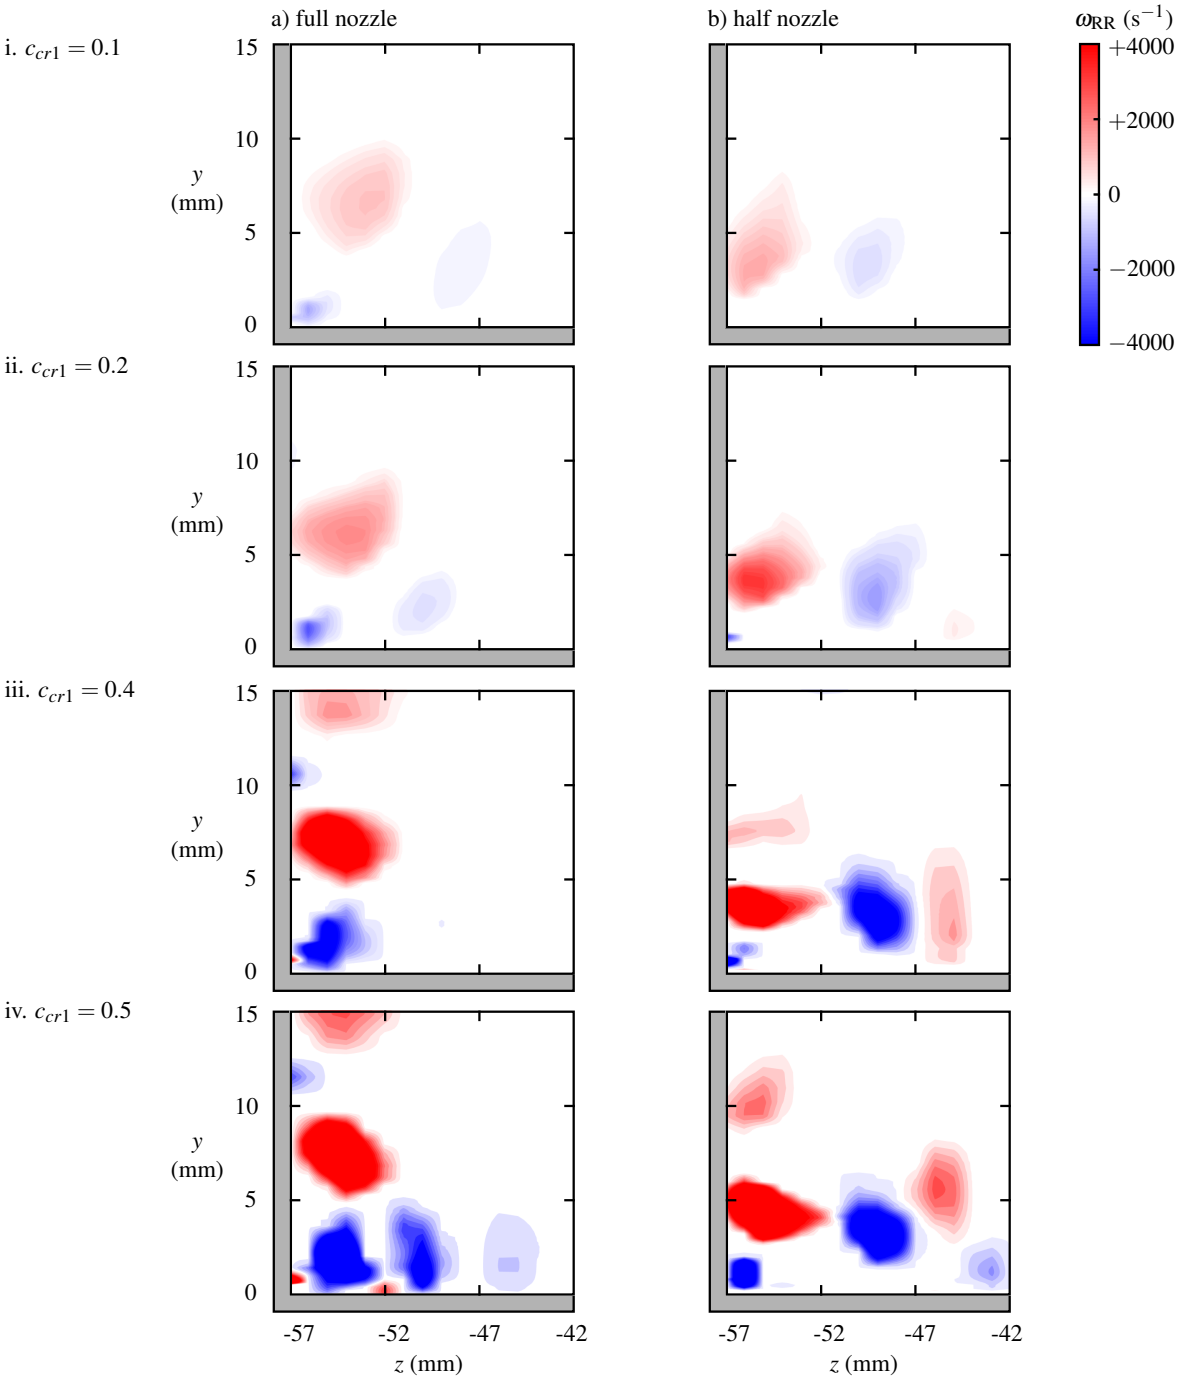

Figure 5.38: Effect of varying c_{cr1} on the rigid-body component of streamwise vorticity in the bottom corner regions for a) the full nozzle setup, and b) the half nozzle setup: i. $c_{cr1} = 0.1$; ii. $c_{cr1} = 0.2$; iii. $c_{cr1} = 0.4$; iv $c_{cr1} = 0.5$.

Table 5.2: Spatial correlation coefficient, $\tilde{C}$, calculated for the velocity data presented in figures 5.32, 5.34 and 5.36.

c_{cr1}	$\tilde{C}$	
	full nozzle	half nozzle
0 (no QCR)	0.936	0.969
0.1	0.948	0.975
0.2	0.962	0.978
0.3 (default value)	0.975	0.977
0.4	0.981	0.970
0.5	0.979	0.962

velocity) is only one possible measure of the agreement between experimental and computational data. This method of estimating error is particularly useful because it can be directly compared to the 'minimum expected error' determined in section 4.2.3. However, another insightful metric is the spatial correlation coefficient:

$$\tilde{C} = \frac{\left\langle u_{\text{exp}}\, u_{\text{comp}} \right\rangle - \left\langle u_{\text{exp}} \right\rangle \left\langle u_{\text{comp}} \right\rangle}{\left[\left(\left\langle u_{\text{exp}}^2 \right\rangle - \left\langle u_{\text{exp}} \right\rangle^2 \right) \left(\left\langle u_{\text{comp}}^2 \right\rangle - \left\langle u_{\text{comp}} \right\rangle^2 \right) \right]^{1/2}} , \tag{5.1}$$

where u_{exp} and u_{comp} represent the experimental and computational streamwise velocity data, respectively, and the angled brackets denote a spatial average over the validation region. The correlation coefficient takes values from -1 to 1 by construction, with $\tilde{C} = 0$ corresponding to no correlation and $\tilde{C} = 1$ indicative of perfect agreement between computation and experiment. The calculated values of $\tilde{C}$ for each computational case are listed in table 5.2. These coefficients highlight the subtle nature of the effect of QCR but show the same broad trends as the error metric from figure 5.37, with maximum correlation observed at $c_{cr1} = 0.4$ for the full nozzle and at $c_{cr1} = 0.2$ for the half nozzle.

Whatever the choice of error metric, it is useful to investigate the topology of streamwise corner vortices in order to better understand the shape of the error curves in figure 5.37 . For each case from figure 5.36, the distribution of the rigid-body component of streamwise vorticity is plotted in figure 5.38. Here, the strength of the counter-rotating vortex pair is seen to initially strengthen as c_{cr1} is increased. As the strength of these vortices approaches those in the physical flow, the corner boundary-layer shape is better predicted and the difference between computation and experiment is reduced, as seen in figure 5.37.

However, when c_{cr1} reaches approximately 0.5 for the full nozzle and 0.3 for the half nozzle, the appearance of additional vortices can be observed. Such vortices, thought to be non-physical, have been reported in similar computations at high values of c_{cr1} by Leger *et al.* [75]. These spurious vortices modify the shape of the corner boundary layer in a manner which does not reflect the true flow. As a result, the calculated error in figure 5.37 for both setups is seen to increase for c_{cr1} large enough that the additional vortices are produced.

5.4 Summary

This chapter describes the identification of two different flow features, both related to the geometry of the two-dimensional nozzle blocks. Firstly, vertical pressure gradients in the nozzle region induce a secondary flow within the sidewall boundary layers, which take the form of bulk vertical velocities. These flows persist into the test section and affect the thickness distribution of the sidewall boundary layers, which grow from the corners to the centre height for the full nozzle and which grow continuously downwards for the half nozzle setup. Second, a region of separated flow at a sharp corner ahead of the nozzle produces vortical structures, which are advected into the test section embedded within the sidewall boundary layers. These quasi-streamwise vortices can locally influence the thickness of the sidewall boundary layers by as much as 37%.

Both the thickness and the complex flow structure of the corner boundary layers are influenced by the sidewall secondary flows and, for some specific geometries, by the separation-induced streamwise vortices. The computations are able to capture both these flow features, and so are used to validate numerical methods. The comparison between computational data and validation measurements shows that the use of the quadratic constitutive relation significantly improves the prediction of corner flows, mainly as a result of generating vortices in this region.

However, even with the use of this model, the corner flows are not predicted perfectly. There is still a residual error of approximately $4-6\%$ which can, perhaps, be attributed to inaccuracies in the position and the strength of the corner vortices. Furthermore, computations which use too large a value of the constant, c_{cr1} produce additional, spurious vortices which distort the corner boundary layer in a non-physical manner. Chapter 6 aims to better understand the capabilities and limitations of this model by revisiting the formulation of the quadratic constitutive relation.

Chapter 6

Revisiting the quadratic constitutive relation

A key conclusion from Chapter 5 is that, whilst using the quadratic constitutive relation can help computations to significantly improve the simulation of corner flows, this model does not produce a perfect prediction of the corner boundary-layer shape. In order to better understand the strengths and limitations of the quadratic constitutive relation (QCR), this chapter looks in detail at the formulation of the model. This allows an assessment of how the quadratic constitutive relation influences the way computations use mean flow data to estimate the turbulent stresses, which govern features like streamwise vorticity production in corner geometries.

6.1 The conventional quadratic constitutive relation

The quadratic constitutive relation, introduced in some detail in section 2.3.2, is designed to improve the accuracy of RANS computations. These methods cannot obtain the turbulent stresses directly, and so estimate them using a prescribed turbulence model to obtain the eddy viscosity, μ_t, from the mean flow data. From this, an 'eddy-viscosity model' is used to determine the turbulent stresses from μ_t. Many RANS computations use a linear eddy-viscosity model (LEVM), which is defined for one-equation turbulence models in equation 2.3: $\tilde{\tau}_{ij} = 2\mu_t S^*_{ij}$. In this equation, $\tilde{\tau}_{ij}$ is the estimated Reynolds stress tensor and $S^*_{ij} = S_{ij} - \frac{1}{3}\frac{\partial u_k}{\partial x_k}\delta_{ij}$, where $S_{ij} = \frac{1}{2}\left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i}\right)$.

However, the resulting distribution of normal stresses is too simple to accurately represent the true values in physical flows. The quadratic constitutive relation therefore aims to better estimate these stresses with only a mild increase in complexity by considering terms that are quadratic in the mean vorticity and strain tensors. The formulation of the relation for one-equation models, QCR-2013 [71],

This chapter uses velocity statistics obtained from publicly-available DNS data sets, listed in table 6.1. Analysis of the data was performed by the author.

is defined in equation 2.7:

$$\tilde{\tau}_{ij} = 2\mu_t S^*_{ij} - \frac{4c_{cr1}\mu_t}{\sqrt{S_{kl}S_{kl}+\Omega_{kl}\Omega_{kl}}}\left[\Omega_{ik}S^*_{kj} - S^*_{ik}\Omega_{kj}\right] - c_{cr2}\mu_t\left[\sqrt{2S^*_{kl}S^*_{kl}}\right]\delta_{ij}\,,$$

where the strain tensor $\Omega_{ij} = \frac{1}{2}\left(\frac{\partial u_i}{\partial x_j} - \frac{\partial u_j}{\partial x_i}\right)$. The constants c_{cr1} and c_{cr2} were proposed as 0.3 and 2.5, respectively, after calibration in the outer part of an equilibrium turbulent boundary layer [71].

The quadratic constitution relation has shown some success in the prediction of flow in geometries featuring streamwise corners, as demonstrated in section 5.3.1. This reflects the findings from other researchers, who have shown that QCR is able to successfully predict the presence of a vortex pair in the corner flow [71, 72]. The improvement in flow prediction has been shown to result in a better estimate of corner separation [73], including shock-induced separation in transonic and supersonic flows [74]. However, the comparisons with experiment in section 5.3.1 showed that even these predictions are limited in their accuracy. Furthermore, for sufficiently large values of constant, c_{cr1}, additional, non-physical vortices are produced, consistent with similar findings from Leger *et al.* [75].

With the exception of the successful prediction of large streamwise vortices in Couette flow [140], the improvement in the computation of other complex flow fields has been somewhat modest. Of course, a perfect prediction of turbulent stresses for all types of flow is not attainable using an approach with a small number of calibration constants. However, the limited success in mean flow prediction suggests that the inclusion of QCR terms does not significantly improve estimates of Reynolds stresses. In order to better understand the predictive abilities of this eddy-viscosity model, the next section starts by revisiting the formulation of the quadratic constitutive relation.

6.2 A more general version of the quadratic constitutive relation

The Reynolds stress modification for QCR-2013 is not, in fact, the most general quadratic form that uses only mean-flow properties. There are three terms, quadratic in S and Ω, which satisfy the fundamental contraction ($\overline{u'_i u'_i} = k$) and symmetry ($\overline{u'_i u'_j} = \overline{u'_j u'_i}$) properties [57, 141]. Assuming that these terms all scale like the quadratic term in QCR-2013, we obtain:

$$\begin{aligned}\tilde{\tau}_{ij} = 2\mu_t S^*_{ij} &- \frac{4c_{cr1}\mu_t}{\sqrt{S_{kl}S_{kl}+\Omega_{kl}\Omega_{kl}}}\left[\Omega_{ik}S^*_{kj} - S^*_{ik}\Omega_{kj}\right] - c_{cr2}\mu_t\left[\sqrt{2S^*_{kl}S^*_{kl}}\right]\delta_{ij} \\ &\underbrace{-\frac{4c_{cr3}\mu_t}{\sqrt{S_{kl}S_{kl}+\Omega_{kl}\Omega_{kl}}}\left[S^*_{ik}S^*_{kj} - \frac{1}{3}S^*_{kl}S^*_{kl}\delta_{ij}\right]}_{\text{A}} \underbrace{+\frac{4c_{cr4}\mu_t}{\sqrt{S_{kl}S_{kl}+\Omega_{kl}\Omega_{kl}}}\left[\Omega_{ik}\Omega_{kj} + \frac{1}{3}\Omega_{kl}\Omega_{kl}\delta_{ij}\right]}_{\text{B}}\,.\end{aligned} \tag{6.1}$$

The two new terms (A and B) are identical for a parallel mean shear flow, $u(y)$, where $v = w = 0$. More generally, these terms also tend to the same expression for any fully-developed flow where $u \gg v, w$. Whilst a long-term aim is to discriminate between these two terms, for now the problem is

Table 6.1: Direct numerical simulations of incompressible flows analysed in this paper. Re_τ is the skin-friction Reynolds number and Re_θ is the Reynolds number based on boundary-layer momentum thickness.

geometry	authors	Reynolds number
channel	Bernardini, Pirozzoli, Orlandi [142]	Re_τ = 2000, 4100
channel	Hoyas, Jiménez [143]	Re_τ = 2000
boundary layer	Sillero, Jiménez, Moser [144]	Re_θ = 6650
boundary layer	Schlatter, Örlü *et al.* [145]	Re_θ = 2500
boundary layer	Schlatter, Li *et al.* [146][1]	Re_θ = 4300
pipe	El Khoury, Schlatter *et al.* [147]	Re_τ = 1000
separated boundary layer	Coleman, Rumsey, Spalart [150]: case A	Re_θ = 1500
square duct	Pirozzoli, Modesti *et al.* [51]	Re_τ = 1100

therefore simplified by considering instead:

$$\begin{aligned}\tilde{\tau}_{ij} = 2\mu_t S^*_{ij} &- \frac{4c_{cr1}\mu_t}{\sqrt{S_{kl}S_{kl}+\Omega_{kl}\Omega_{kl}}}\left[\Omega_{ik}S^*_{kj} - S^*_{ik}\Omega_{kj}\right] - c_{cr2}\mu_t\left[\sqrt{2S^*_{kl}S^*_{kl}}\right]\delta_{ij} \\ &- \frac{4\tilde{c}_{cr3}\mu_t}{\sqrt{S_{kl}S_{kl}+\Omega_{kl}\Omega_{kl}}}\left[S^*_{ik}S^*_{kj} - \frac{1}{3}S^*_{kl}S^*_{kl}\delta_{ij}\right],\end{aligned} \tag{6.2}$$

with $\tilde{c}_{cr3} = c_{cr3} - c_{cr4}$. Further simplification is achieved by assuming that c_{cr1}, c_{cr2} and $\tilde{c}_{cr3}$ are constant in space, rather than a function of the production-dissipation ratio, P/ε, or indeed some other relevant non-dimensional quantity.

In order to calibrate the constants (c_{cr1}, c_{cr2}, $\tilde{c}_{cr3}$), ensemble-averaged velocity statistics from DNS data sets of incompressible canonical flows at moderate Reynolds numbers ($\mathrm{Re}_\tau > 1000$) are used. The cases, detailed in table 6.1, include channel flows [142, 143], boundary layers [144–146] and pipe flows [147]. The profiles from these different flow families exhibit quite distinct turbulent stress distributions, especially away from the wall [148].

For these parallel shear flows, the only turbulent stress component which affects the mean velocity, $u(y)$, is the shear stress, τ_{xy}. RANS methods use a prescribed turbulence model to estimate the eddy viscosity, μ_t. They then use an appropriate eddy-viscosity model with the calculated distribution of μ_t to directly relate $u(y)$ with τ_{xy}.

However, for the current analysis the DNS data provides both u and τ_{xy}. It is therefore possible to determine the distribution of μ_t that would correctly relate the shear stress to the mean flow in a RANS simulation. This approach is independent of turbulence model, as it uses the true mean flow and turbulence statistics directly from DNS. The analysis invokes the method of Spalart *et al.* [149]:

$$\mu_t = \frac{\tau_{ij}S_{ij}}{2S_{kl}S_{kl}} = \frac{\tau_{xy}}{\frac{\partial u}{\partial y}}\,. \tag{6.3}$$

[1] This study was performed using well-resolved large-eddy simulations.

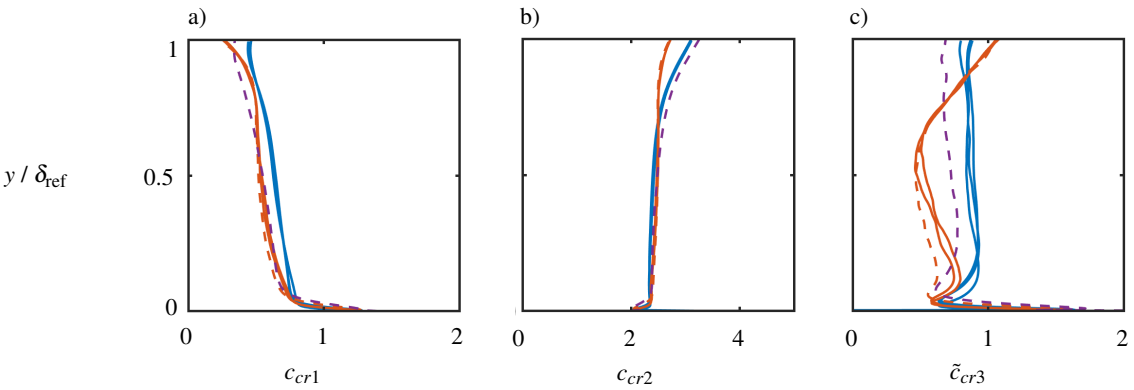

Figure 6.1: Calibration of a) c_{cr1}, b) c_{cr2} and c) $\tilde{c}_{cr3}$ along the wall-normal direction (scaled by the flow thickness, δ_{ref}), to produce the true Reynolds stresses. Flow fields: channel flow (—) [142, 143]; boundary layers (—) [144–146]; pipe flow (—) [147].

The extracted distribution of μ_t can then be used to determine the unknown calibration constants, c_{cr1}, c_{cr2} and $\tilde{c}_{cr3}$. By evaluating the diagonal terms of equation 6.2 for a parallel mean shear flow, we obtain:

$$\begin{bmatrix} \tau_{xx} \\ \tau_{yy} \\ \tau_{zz} \end{bmatrix} = -\mu_t \frac{\partial u}{\partial y} \begin{bmatrix} 2 & 1 & 1/3 \\ -2 & 1 & 1/3 \\ 0 & 1 & -2/3 \end{bmatrix} \begin{bmatrix} c_{cr1} \\ c_{cr2} \\ \tilde{c}_{cr3} \end{bmatrix}, \tag{6.4}$$

where τ_{xx}, τ_{yy} and τ_{zz} correspond to the true normal stresses, which are extracted from DNS. Equation 6.4 is then inverted to find the values of c_{cr1}, c_{cr2} and $\tilde{c}_{cr3}$ which give the correct normal stresses at each point:

$$\begin{bmatrix} c_{cr1} \\ c_{cr2} \\ \tilde{c}_{cr3} \end{bmatrix} = -\frac{1}{\mu_t \frac{\partial u}{\partial y}} \begin{bmatrix} 2 & 1 & 1/3 \\ -2 & 1 & 1/3 \\ 0 & 1 & -2/3 \end{bmatrix}^{-1} \begin{bmatrix} \tau_{xx} \\ \tau_{yy} \\ \tau_{zz} \end{bmatrix}. \tag{6.5}$$

This matrix equation directly determines the constants from the true mean flow and turbulent stresses, which are known from DNS. Equation 6.5 is then solved at each point for the canonical parallel shear flows described above.

The values of the constants as a function of the wall-normal coordinate, y, for the three different flow cases are shown in figure 6.1. These 'constants' appear to be fairly uniform across much of the flow, and are similar between the different flows. Whilst the constants should, in theory, vary with non-dimensionalised wall-normal coordinate, the problem is simplified by assigning a single value. A reasonable compromise between the various canonical flow fields gives $c_{cr1} = 0.7$, $c_{cr2} = 2.5$ and $\tilde{c}_{cr3} = 0.8$. This modified version of the conventional formulation is termed the 'general quadratic constitutive relation', or general QCR.

6.2.1 Application of the general relation

In order to test the effect of the proposed modifications, it is necessary to conduct RANS simulations with these different turbulence models on a wide range of flows, and to compare the mean flow with corresponding validation data. Such an involved process is outside the scope of this thesis. Instead, we focus on the turbulent stresses, which can have a significant impact on the mean flow.

The data from the direct numerical simulations listed in table 6.1 are used to provide the 'correct' mean flow as an input for the eddy-viscosity models (equations 2.3, 2.7, and 6.2) which RANS simulations would use to estimate the Reynolds stress distribution. These estimated turbulent stresses are then compared with the 'true' values obtained directly from DNS. Note that since RANS may not exactly compute the correct mean flow, this procedure does not quite equate to an assessment of how well an actual RANS calculation would predict the real turbulent stresses. However, the validity of this process improves with increased accuracy of RANS simulations and the fundamental aim of RANS is to produce the correct mean flow, so this approach can still be considered instructive.

The normal turbulent stress components (τ_{xx}, τ_{yy} and τ_{zz}) are estimated from the mean flow of a DNS solution using equation 2.3 (linear eddy-viscosity model), equation 2.7 (QCR-2013) and equation 6.2 (general form of QCR). For this process, the eddy viscosity, μ_t, is calculated using equation 6.3.

Figure 6.2 shows how the estimated normal Reynolds stresses compare with the 'true' values, extracted from the direct numerical simulations of parallel shear flows. This shows that QCR-2013 improves on the turbulent stresses predicted with linear eddy-viscosity models. In turn, the more general form of QCR proposed here predicts the turbulent stresses better than QCR-2013. Note that these improvements are achieved with values for the constants that are neither spatially dependent nor flow specific.

This analysis can also be extended to more complex cases in table 6.1, such as the separated boundary layer [150] and the flow in a square duct [51]. In contrast to the parallel shear flows from figure 6.2, the mean velocity in such flow fields (which have non-zero cross-flow, and which vary in x or z) are no longer determined solely by τ_{xy}. Therefore, it is not possible to use the eddy-viscosity distribution from equation 6.3 to correctly relate shear stress with the mean flow, and so equation 6.3 is not physically meaningful for these flows. Instead, a different method is used to evaluate a relevant eddy-viscosity distribution, μ_t, which provides the best fit for a weighted balance of all six Reynolds stress components:

$$\mu_t = \frac{\tau_{ij}\tilde{\sigma}_{ij}}{\tilde{\sigma}_{kl}\tilde{\sigma}_{kl}} \,, \tag{6.6}$$

where τ_{ij} is the true Reynolds stress tensor from DNS and $\tilde{\sigma}_{ij} = \tilde{\tau}_{ij} / \mu_t$. Here, the definition of $\tilde{\tau}_{ij}$ is taken from equation 2.3, 2.7 or 6.2 for the LEVM, QCR-2013 and general QCR, respectively. Note that $\tilde{\tau}_{ij}$ is proportional to μ_t for all three definitions, and so $\tilde{\sigma}_{ij}$ is independent of μ_t.

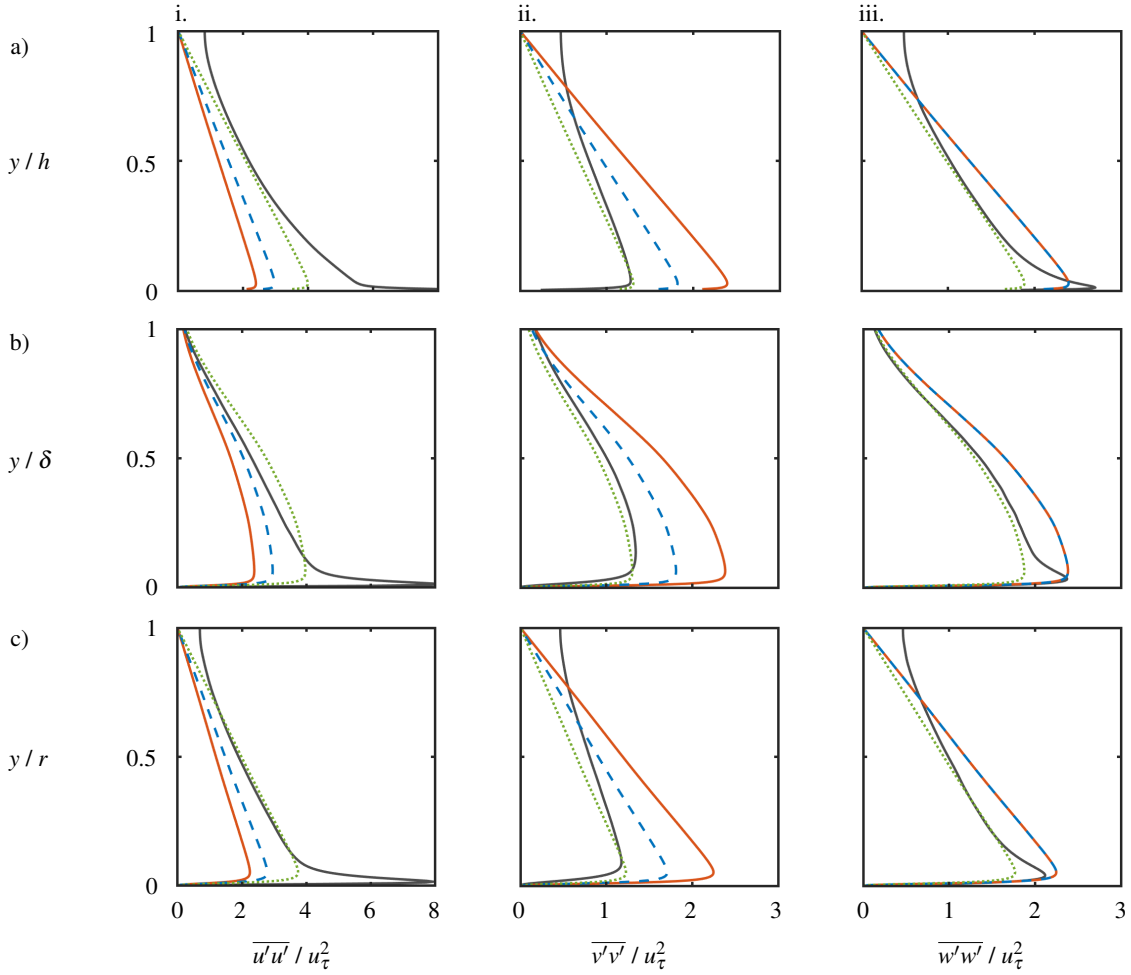

Figure 6.2: Effect of different Reynolds stress approximations on estimating the normal turbulent stresses (i. $\overline{u'u'}$, ii. $\overline{v'v'}$, iii. $\overline{w'w'}$), scaled by the friction velocity, u_τ. The true stress from DNS (—) and estimated stress from an LEVM (—), QCR-2013 (- - -) and general QCR (· · ·). Flow fields: a) channel flow [142]; b) boundary layer [144]; c) pipe flow [147].

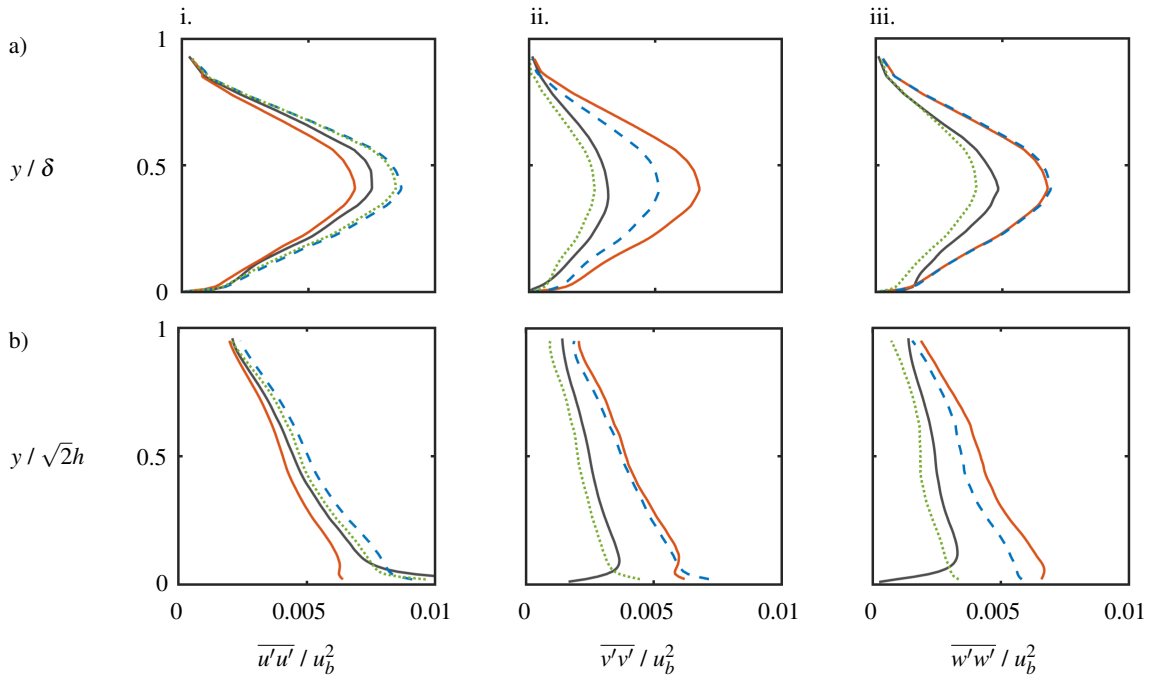

Figure 6.3: Effect of different Reynolds stress approximations on estimating the normal turbulent stresses (i. $\overline{u'u'}$, ii. $\overline{v'v'}$, iii. $\overline{w'w'}$), scaled by the bulk velocity, u_b. The true stress from DNS (—) and estimated stress from an LEVM (—), QCR-2013 (- - -) and general QCR (· · ·). Flow fields: a) separated boundary layer [150]; b) corner bisector of square duct flow [51].

Figure 6.3 shows the estimated normal turbulent stresses for these flows. Whilst there are no real differences between QCR-2013 and the LEVM in turbulent stress prediction for these flows, the new general QCR does appear to provide substantial improvements. In the near-wall region, however, these simple models continue to struggle in predicting the stresses, and capturing this wall-normal dependence remains a challenge for the future.

The general form of QCR appears to lead to promising improvements in the estimate of Reynolds stress distribution. Note, however, that the accuracy of mean flows computed using corresponding RANS simulations has not yet been evaluated. Such an assessment would be the ultimate test of this proposed modification to a linear eddy-viscosity model, and would permit its numerical stability to be tested.

6.3 The quadratic constitutive relation in corner flows

The analysis of direct numerical simulations has shown that conventional QCR produces a somewhat inaccurate estimate of turbulent stresses. Despite this, conventional QCR appears to more accurately calculate the flow along corner geometries than the corresponding linear eddy-viscosity models [71, 72, 74]. In order to better understand this apparent contradiction, the turbulent stress estimates in corner flows are analysed more closely.

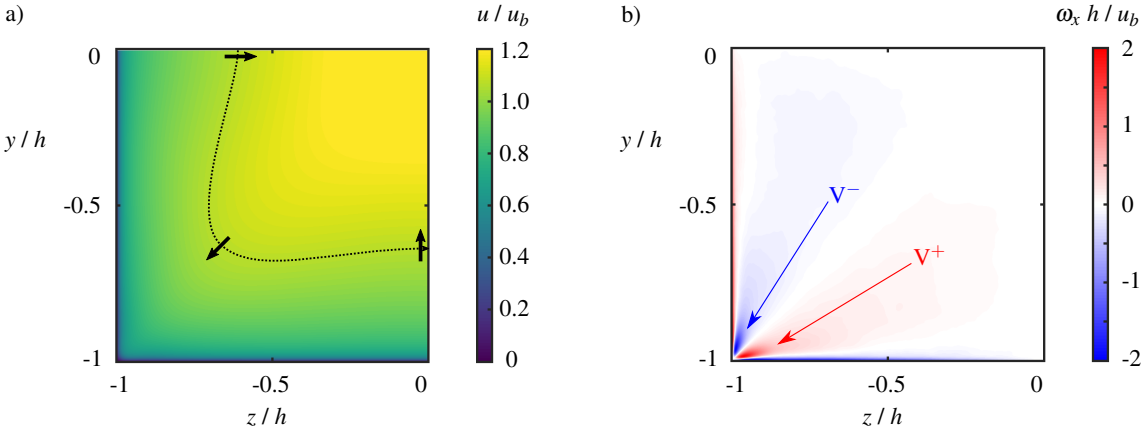

Figure 6.4: Direct numerical simulation of the flow in a quarter of a square duct [51] with half-height, h, and bulk velocity, u_b: a) the mean streamwise velocity component, u; b) the mean streamwise vorticity, ω_x. The dotted line is a contour of velocity $u/u_b = 1$, representative of the boundary-layer edge shape.

In section 2.3.1, the generation of streamwise vorticity in corner geometries was attributed to the production of Prandtl vortices of the second kind, related to anisotropies in the turbulent stresses. In particular, the two terms labelled C and D in the mean streamwise vorticity equation (equation 2.1) were seen to be dominant [45]. These terms from equation 2.1 are given again below for convenience:

$$\text{term C}: \quad \frac{1}{\rho}\left(\frac{\partial^2}{\partial y^2} - \frac{\partial^2}{\partial z^2}\right)\tau_{yz}\,,$$

$$\text{term D}: \quad \frac{1}{\rho}\frac{\partial^2}{\partial y \partial z}\left(\tau_{zz} - \tau_{yy}\right)\,.$$

For a linear eddy-viscosity model, these vorticity production terms equate exactly to zero. This explains why RANS simulations with LEVMs are unable to generate the quasi-streamwise vortices which exist in corners, resulting in a poor prediction of these flows.

The situation is different when quadratic modifications are introduced, however. To understand the differences, we will consider the Reynolds stress modification defined in equation 6.2, noting that the proposed values of $c_{cr1} = 0.7$, $c_{cr2} = 2.5$ and $\tilde{c}_{cr3} = 0.8$ give general QCR while the combination $c_{cr1} = 0.3$, $c_{cr2} = 2.5$ and $\tilde{c}_{cr3} = 0$ corresponds to QCR-2013. Evaluating the relevant turbulent stress combinations from terms C and D in equation 2.1 for a corner flow, $u(y,z)$, we find that:

$$\tau_{zz} - \tau_{yy} = \frac{(2c_{cr1} - \tilde{c}_{cr3})\,\mu_t}{\sqrt{\left(\frac{\partial u}{\partial y}\right)^2 + \left(\frac{\partial u}{\partial z}\right)^2}}\left[\left(\frac{\partial u}{\partial z}\right)^2 - \left(\frac{\partial u}{\partial y}\right)^2\right]\,, \tag{6.7}$$

$$\tau_{yz} = \frac{(2c_{cr1} - \tilde{c}_{cr3})\,\mu_t}{\sqrt{\left(\frac{\partial u}{\partial y}\right)^2 + \left(\frac{\partial u}{\partial z}\right)^2}}\left[\left(\frac{\partial u}{\partial y}\right)\left(\frac{\partial u}{\partial z}\right)\right]\,. \tag{6.8}$$

These turbulent stress combinations only depend on the constant combination $2c_{cr1} - \tilde{c}_{cr3}$. This suggests that the generation of streamwise vorticity depends only on the value of $2c_{cr1} - \tilde{c}_{cr3}$ (i.e. some measure of the strength of the quadratic modification) rather than the individual contributions from the different terms. Therefore, while corner flows do show the importance of quadratic terms in turbulence modelling, they are poor at discriminating between them. This suggests that evaluating the mean flow along a streamwise corner is not a suitable test for comparing the capabilities of different quadratic modifications to eddy-viscosity models.

Furthermore, it is interesting to note that $2c_{cr1} - \tilde{c}_{cr3} = 0.6$ both for QCR-2013 and for general QCR. This suggests that vorticity production in corner flows is essentially equivalent between the two forms and provides an explanation for the success of conventional QCR in such flows despite somewhat poor Reynolds stress prediction. An additional consequence of this finding is that extending QCR to its more general form is unlikely to detrimentally impact the key success of conventional QCR, i.e. the prediction of the mean flow along streamwise corners.

6.3.1 Analysis of direct numerical simulations

The analysis of corner flows discussed above can be illustrated using ensemble-averaged velocity statistics from direct numerical simulations. A fully-developed, incompressible square duct flow was computed by Pirozzoli *et al.* for a skin-friction Reynolds number, $\mathrm{Re}_\tau = 1100$ [51]. Figures 6.4a and 6.4b show the streamwise velocity and vorticity distributions, respectively, for a quarter of the duct. As well as capturing the shear in the boundary layers, the vorticity distribution displays the pair of counter-rotating vortices (labelled V^+ and V^-) either side of the corner bisector. The distorted boundary-layer shape, denoted by the dotted line in figure 6.4a, demonstrates how the vortices affect the streamwise velocity profile through momentum transfer between the core flow and the boundary layers.

Figure 6.5a presents the Reynolds stresses (τ_{yz}, τ_{yy} and τ_{zz}) which appear in the streamwise vorticity equation (equation 2.1), extracted directly from DNS. These are compared with equivalent estimates using equations 2.3, 2.7 and 6.2, which correspond to linear eddy-viscosity model, QCR-2013 and general QCR, respectively.

A comparison of figure 6.5a with figure 6.5b shows that linear eddy-viscosity models cannot accurately predict the distribution of relevant turbulent stresses. When the QCR-2013 modification is used (figure 6.5c), there is an improvement in τ_{yz}. However, the estimates of τ_{yy} and τ_{zz} are still inaccurate in both magnitude and topology. Figure 6.5d shows that general QCR appears to better predict all three stresses, though the true stress distributions from figure 6.5a are still not perfectly replicated.

Figure 6.6 shows the equivalent comparison for the stress combination $\tau_{zz} - \tau_{yy}$, which appears in term D of equation 2.1. Linear eddy-viscosity models which obey equation 2.3 assume that $\tau_{yy} = \tau_{zz}$, and so the difference between the two stress components evaluates identically to zero (figure 6.6b). More interestingly, figures 6.6c and 6.6d appear to be equivalent – QCR-2013 exactly emulates the

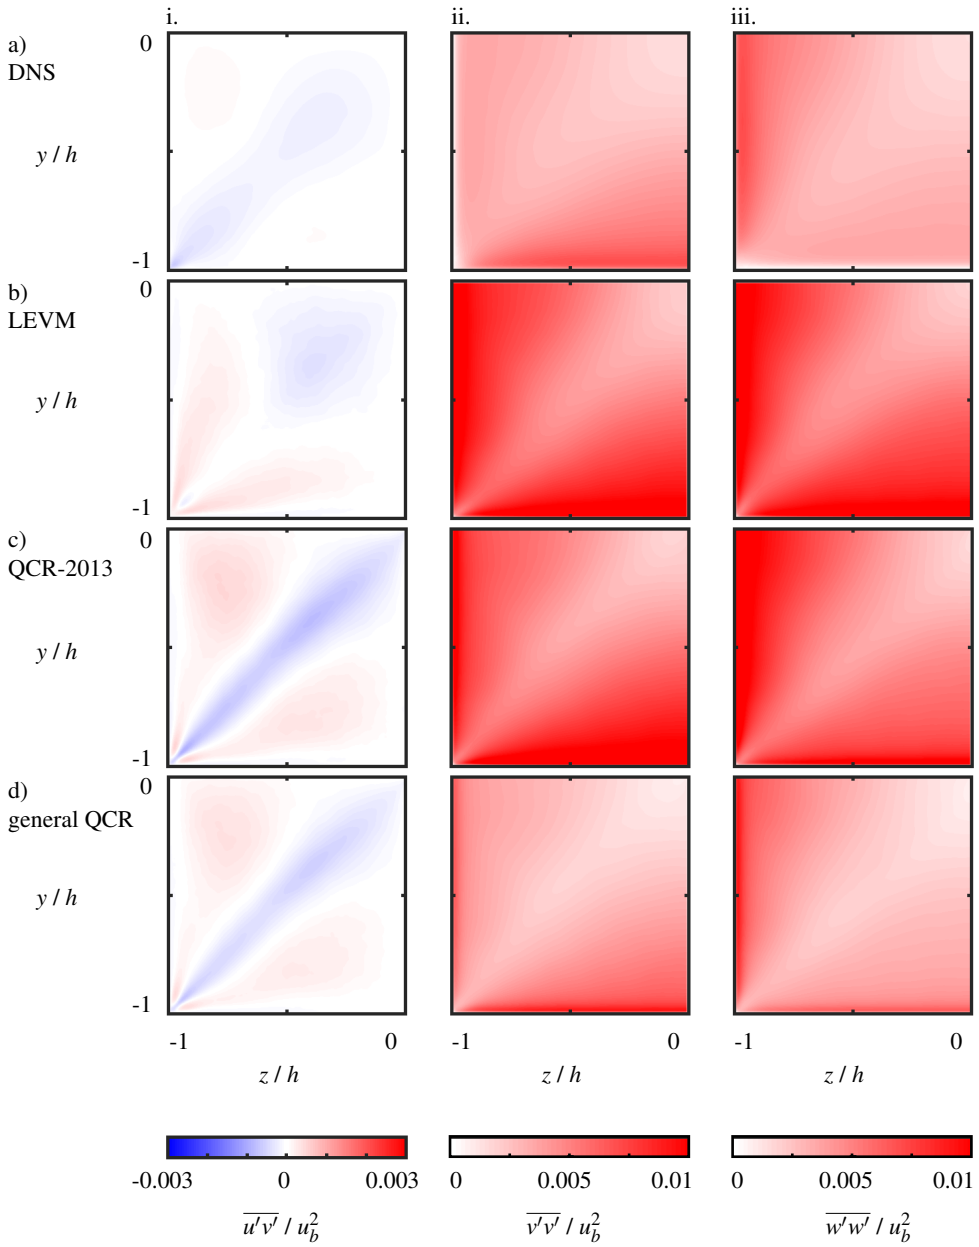

Figure 6.5: Turbulent stresses relevant for streamwise vorticity production, in the square duct [51]: i. $\overline{u'v'}$, ii. $\overline{v'v'}$, and iii. $\overline{w'w'}$. These stresses are a) directly extracted from the DNS turbulence statistics, or are estimated using b) a linear eddy-viscosity model, c) QCR-2013, and d) general QCR.

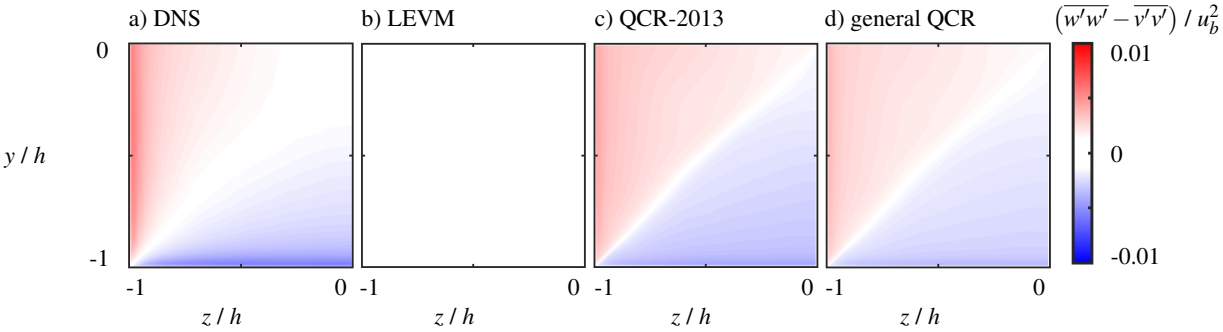

Figure 6.6: Turbulent stress combination $\overline{w'w'} - \overline{v'v'}$, in the square duct [51]. These stresses are a) directly extracted from the DNS turbulence statistics, or are estimated using b) a linear eddy-viscosity model, c) QCR-2013, and d) general QCR.

estimates from general QCR. This is a direct consequence of the finding that, for these quadratic modifications, $\tau_{zz} - \tau_{yy}$ depends only on $2c_{cr1} - \tilde{c}_{cr3}$. This combination of constants is 0.6 for both QCR-2013 and general QCR. For the same reason, the distribution of τ_{yz} in figure 6.5c (QCR-2013) and in figure 6.5d (general QCR) appear to be equivalent.

It is now possible to evaluate terms C and D in equation 2.1, individually and in combination. These terms evaluated for the stress distributions from the direct numerical simulation are displayed in figure 6.7a. This shows one lobe each side of the corner bisector (P^+ and P^-), of opposite sign. Both terms C and D affect these lobes, with term C contributing more towards the corner bisector and term D influencing the behaviour close to the walls. As expected from the analysis of the streamwise vorticity equation, linear eddy-viscosity models (figure 6.7b) cannot produce vorticity via this mechanism.

The production terms for QCR-2013 (figure 6.7c) and general QCR (figure 6.7d) produce very similar distributions to one another. This is a consequence of these two quadratic modifications displaying equivalent spatial distributions for the relevant turbulent stress combinations τ_{yz} (figure 6.5i) and $\tau_{zz} - \tau_{yy}$ (figure 6.6). Therefore, the following discussion does not distinguish between QCR-2013 and general QCR.

A closer analysis of figure 6.7 shows that the estimated vorticity production for QCR exhibits primary lobes of the correct sign, labelled P_1^+ and P_1^-. These correspond to P^+ and P^- from figure 6.7a. This rough agreement in the structure, sign and magnitude of the production terms explains the improvement in simulation of corner flows due to QCR.

However, the topology of the production terms for the quadratic models do exhibit deviations from the 'true' distribution taken directly from DNS. A comparison of the first column in figures 6.7c and 6.7d to figure 6.7a shows that these quadratic models underestimate the contribution from term C. Instead, the primary lobes in the quadratic models (P_1) are determined mainly from term D.

Therefore, the distribution of overall vorticity production is poorly predicted by quadratic models, evident from an examination of the third column in figure 6.7. QCR predicts the P_1 lobes to be elongated close to the corner bisector, whilst the true production distribution (P^+ and P^-) is concentrated towards the walls. This suggests that, whilst QCR can predict the presence of corner

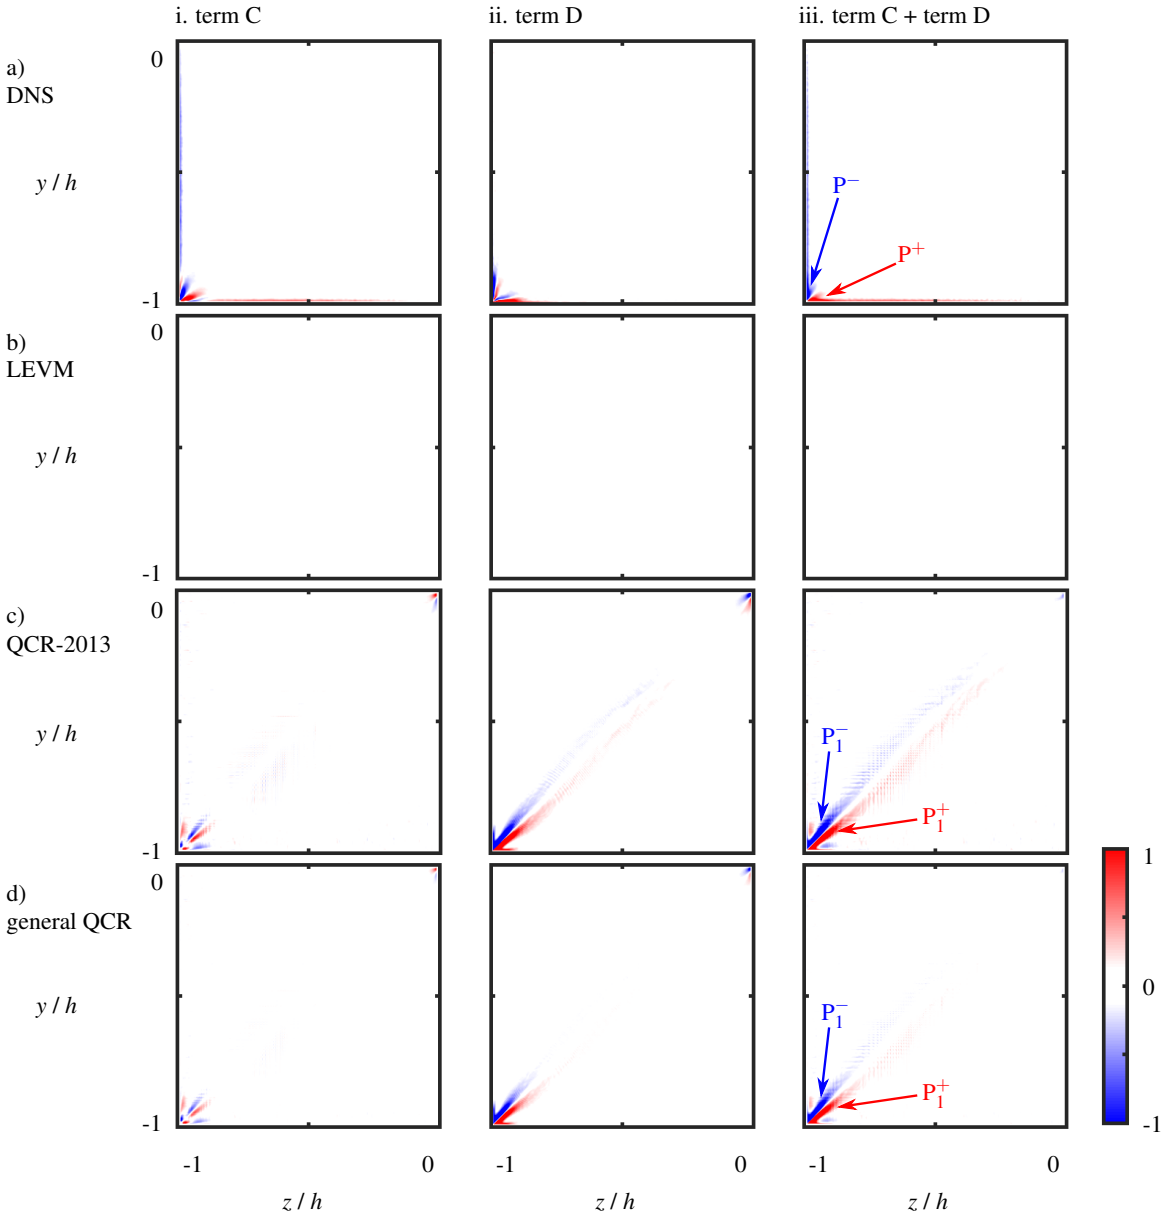

Figure 6.7: The streamwise vorticity production terms from equation 2.1, in the square duct [51]: i. term C $\left(\frac{h^2}{u_b^2}\left(\frac{\partial^2}{\partial y^2}-\frac{\partial^2}{\partial z^2}\right)\left(-\overline{v'w'}\right)\right)$, ii. term D $\left(\frac{h^2}{u_b^2}\frac{\partial^2}{\partial y \partial z}\left(\overline{v'^2}-\overline{w'^2}\right)\right)$, and iii. term C + term D. These terms are a) directly extracted from the DNS turbulence statistics, or are estimated using b) a linear eddy-viscosity model, c) QCR-2013, and d) general QCR.

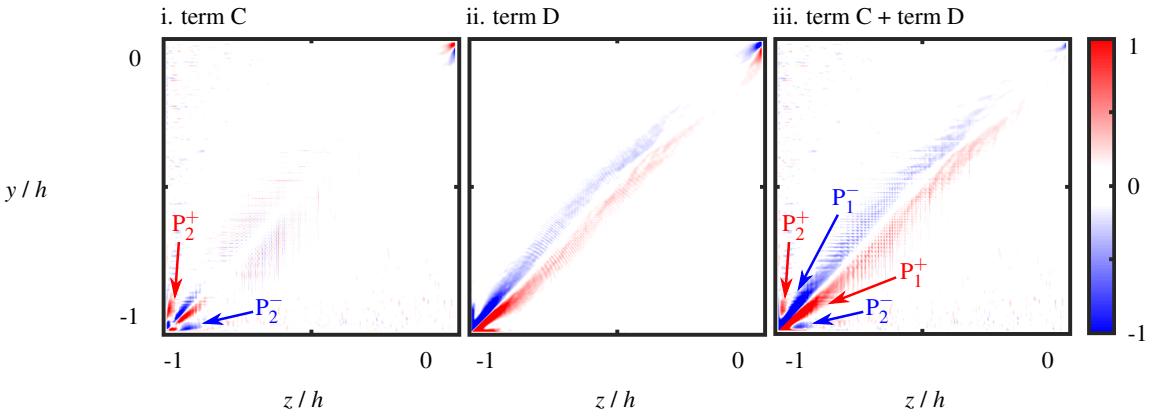

Figure 6.8: The streamwise vorticity production terms from equation 2.1, in the square duct [51]: i. term C $\left(\frac{h^2}{u_b^2}\left(\frac{\partial^2}{\partial y^2}-\frac{\partial^2}{\partial z^2}\right)\left(-\overline{v'w'}\right)\right)$, ii. term D $\left(\frac{h^2}{u_b^2}\frac{\partial^2}{\partial y \partial z}\left(\overline{v'^2}-\overline{w'^2}\right)\right)$, and iii. term C + term D. These terms are obtained with $c_{cr1}=0.5$, $c_{cr2}=2.5$ and $\tilde{c}_{cr3}=0$ in equation 6.2.

vortices and the resultant boundary-layer shape, the properties of the vortices (such as their position) may not quite correspond to those in the physical flow.

Another observed limitation of QCR in corner flows is the generation of additional non-physical vortices when simulations are run with c_{cr1} greater than the recommended value of 0.3 in order to, for example, better match the mean flow [75]. To investigate this further, the production terms are estimated for $c_{cr1}=0.5$ and $\tilde{c}_{cr3}=0$, i.e. $2c_{cr1}-\tilde{c}_{cr3}=1.0$ rather than 0.6. The primary effect, shown in figure 6.8iii, is to strengthen all the turbulent stresses, and thus vorticity production. This figure also reveals the appearance of additional lobes away from the corner bisector, labelled P_2^+ and P_2^-. These lobes are of opposite sign to the primary P_1 lobes. The P_2 lobes also appear in figure. 6.8i, and so seem to be associated with term C (rather than term D) in equation 2.1. Note that weak analogues to these P_2 lobes can be identified in figures 6.7c and 6.7d (where $2c_{cr1}-\tilde{c}_{cr3}=0.6$). The lobes in figure 6.8, however, are strong enough to be prominent in the distribution of overall vorticity production. The presence of the P_2 lobes may correspond to the non-physical vortices observed for high values of c_{cr1} in section 5.3.1.

In this way, the appearance of the spurious vortices can be linked to term C in equation 2.1 and, in turn, the inaccuracy of estimating the Reynolds shear stress component τ_{yz}. However, the dependence of these production terms on the parameter $2c_{cr1}-\tilde{c}_{cr3}$ suggests that any simple quadratic modification would face similar issues with corner flow prediction. To overcome this, it is necessary to add complexity either by capturing the wall-normal dependence of the tuning constants or by estimating the Reynolds stresses using a different, more involved technique.

6.4 Summary

The analysis of velocity statistics from DNS data sets in this chapter shows that the conventional form of the quadratic constitutive relation does not, in general, accurately compute the turbulent stresses. Therefore, it is unsurprising that this quadratic modification does not significantly improve mean flow predictions for most flow fields. Indeed, the observed improvements in predicting the flow along streamwise corners cannot be attributed to a much better prediction of turbulent stresses.

Instead, this success is due to the curious feature that corner flows are sensitive to the presence of quadratic modifications in general, rather than to the specific quadratic modification in use. Moreover, the analysis of direct numerical simulations of the square duct flow shows that, whilst QCR does generate streamwise vorticity in the corner regions, the position and strength of the produced vortices may not be exactly correct and there is a possibility of creating artificial, non-physical vortices near the walls.

This chapter also proposes a more general version of QCR with one additional quadratic term. The corresponding coefficients are calibrated using DNS data for boundary layers, pipe flow and channel flow. The resulting eddy-viscosity model is seen to significantly improve turbulent stress prediction in many flow fields, including the separated boundary layer and the square duct flow. This finding suggests that general QCR may improve mean flow predictions in flows influenced by turbulent stress anisotropies, though RANS computations using the new model are required to demonstrate this hypothesis conclusively.

Chapter 7

Impact of the corner boundary layers on shock-induced separation

Chapter 5 has shown that the geometry of two-dimensional supersonic wind-tunnel nozzles can have a strong influence on the structure of the corner flows in the test section. However, the influence of the corner boundary layer on separation in the local region (and, therefore, in the wider flow field) when it encounters an adverse pressure gradient remains yet to be established. This chapter aims to address this gap in knowledge by investigating the response of the tunnel's floor boundary layer to an oblique shock reflection for two setups with very different corner flow structures.

7.1 Experimental tests of the oblique shock reflection

In order to test the effects of differing corner boundary layers, it is necessary to produce two configurations which have identical flow fields except for the corner regions. Two such setups are designed by exploiting knowledge about the effects of nozzle geometry on the corner boundary layer, discussed in section 5.3. A half nozzle geometry is used for both setups. For setup A, shown in figure 3.3a, the half nozzle is in the conventional orientation, with the contoured surface on the ceiling. Meanwhile, setup B has an inverted half nozzle setup so that the contoured surface is on the floor, as depicted in figure 3.3b. Each of these configurations will then be exposed to an oblique shock reflection on the floor of the test section so that the only substantial difference between them is the direction of the secondary flow along the sidewalls, and the resulting effect on the bottom corner boundary layers.

The flow in the empty wind tunnel, with no shock-generating wedge and with a straight floor and ceiling, is evaluated first. Schlieren images are shown in figures 7.1a and 7.1b, for setups A and B, respectively. At first glance, the flows in these two setups appear to be very similar in nature. A prominent oblique wave, originating from a junction between constituent liner blocks, is visible in both setups. However, the static pressure measurements shown in figure 4.9 suggest that this wave is weak. Velocity profile measurements at several streamwise locations (figure 4.12) show further that the wave does not appear to perturb the boundary layers.

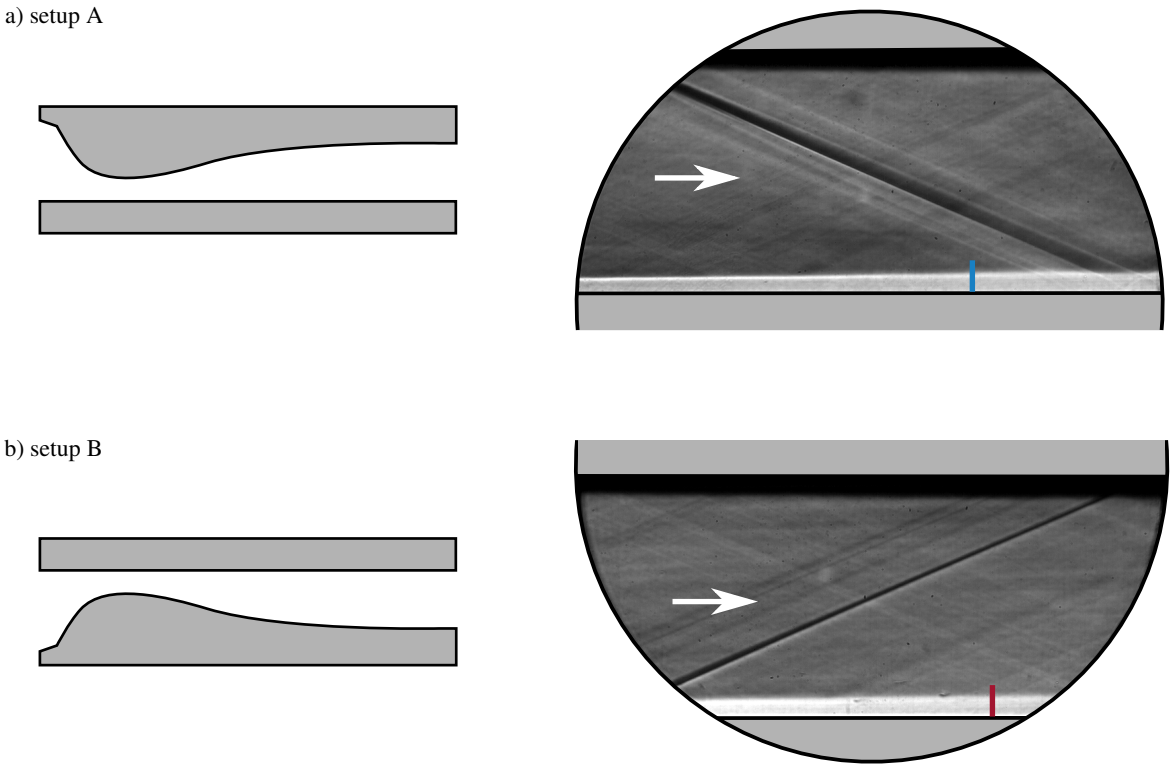

Figure 7.1: Schlieren images with the empty tunnel for a) setup A, and b) setup B. The solid lines indicate the location of laser Doppler velocimetry measurements.

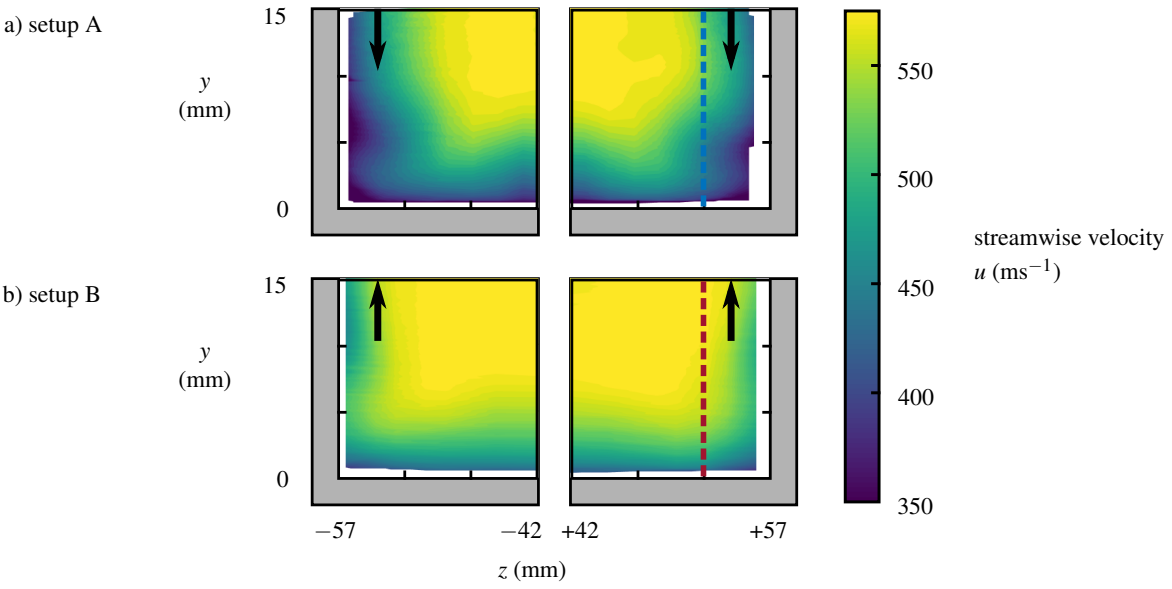

Figure 7.2: The streamwise velocity (u) in 15 mm$\times$15 mm regions around the bottom tunnel corners for a) setup A and b) setup B. These are measured using LDV at $x = 120$ mm. The directions of the sidewall secondary flows are identified by solid arrows. The dashed lines indicate the location of the profile measured in figure 7.3b.

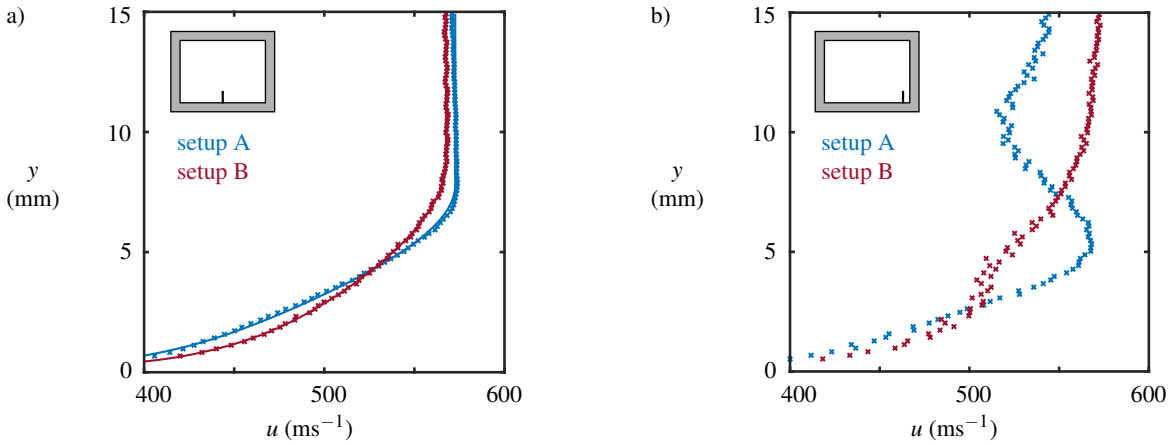

Figure 7.3: Comparison of boundary-layer profiles, measured using LDV at $x = 120$ mm: a) on the tunnel centre span, and b) at $z = 52$ mm.

Despite their similar appearance, the discussion from section 5.1 suggests that two setups exhibit quite different sidewall boundary layers. In a half nozzle arrangement, the sidewall secondary flows, induced by the pressure gradients in the nozzle, are directed from the contoured surface towards the flat surface. For setup A, this corresponds to the downwards direction whilst for setup B, the sidewall flows are in the upwards direction. These secondary flows serve to transport the low-momentum fluid within the boundary layers, and so produce distinctly different flow structures in the bottom tunnel corners for the two configurations. This difference is illustrated by the streamwise velocity distributions presented in figure 7.2. In particular, setup A has much thicker corner boundary layers than setup B, even though the floor boundary-layer thickness is the same for both.

The floor boundary layers are measured on the tunnel centre span at $x = 120$ mm using laser Doppler velocimetry. Figure 7.3a suggests that the velocity profiles for the two setups are similar but not identical. This is consistent with the apparent floor boundary-layer thickness in the schlieren images (figure 7.1). The boundary-layer parameters corresponding to the velocity profiles are listed in table 7.1. These show that the boundary-layer displacement thickness and shape factor differ by 13% and 5%, respectively. Corresponding velocity measurements at $z = 52$ mm, within the corner boundary layer, are shown in figure 7.3b. The non-equilibrium nature of the boundary layer here means that it is not possible to calculate integral parameters through profile fitting. It is, however, clear that the differences in velocity profile between setups A and B are significantly larger in the corner regions than on the tunnel centre line.

Setup A and setup B therefore correspond to two otherwise-identical setups with distinct corner boundary layers. In order to test the impact of these corner regions on the separated flow field, an oblique shock is introduced. Figure 3.3 shows that a full-span wedge is placed on the tunnel ceiling. This produces a flow deflection angle of 8°, and thus generates an oblique shock.

Note that in figures 3.3a and 3.3b, the wedge is placed below the upper surface of the tunnel to ensure that the ceiling boundary layer disappears into the gap and that the shock is generated in

Setup	δ / mm	δ_i^* / mm	θ_i / mm	H_i
A	7.7	1.00	0.75	1.33
B	8.7	0.87	0.69	1.27

Table 7.1: Incompressible floor boundary-layer parameters, measured at $x = 120$ mm and $z = 0$ mm, for experimental profiles presented in figure 7.3a.

clean flow. The gap for setup B (8 mm) is larger than for setup A (5 mm). This is made necessary by differences in the upper corner flows between the two setups. The 5 mm gap in setup A is sufficient to capture the ceiling boundary layer across the tunnel span. However, the flat ceiling nozzle for setup B is associated with thicker sidewall and corner boundary layers. Therefore, a wedge placed 5 mm below the ceiling would still be in boundary layers across much of the tunnel span. Instead, by dropping the wedge a further 3 mm, the shock is generated in clean core flow.

The different wedge position has one main effect – the shock in setup B impinges on the floor boundary layer approximately 6 mm (about one local boundary-layer thickness) upstream compared to setup A. Whilst the slight difference in boundary-layer displacement thickness between these two locations is not thought to be significant, any direct comparison between the two setups does require an alternative coordinate system. We therefore define $\tilde{x}$ as the streamwise coordinate relative to the inviscid shock-reflection location. In other words, $\tilde{x} = x - 141$ mm for setup A and $\tilde{x} = x - 135$ mm for setup B.

7.1.1 Measurements of the separated flow field

Figure 7.4 shows a schlieren image of the flow both for setup A and for setup B. In both cases, the wedge is visible at the top of the image. The flow deflection caused by the leading edge and by the bottom corner of the wedge generate an oblique shock and an expansion fan, respectively. The latter waves, denoted the 'wedge expansion', complicates the flow field, so we restrict our attention to the features upstream of the first expansion wave.

The wave pattern in both setups is representative of a typical two-dimensional separated oblique shock–boundary-layer interaction, described in some detail in section 2.1.2. The incident shock causes the incoming floor boundary layer to separate. The flow deflection at the upstream edge of the separation produces a separation shock, the curvature of the separation bubble causes a series of expansion waves, and the turning of the flow as it reattaches results in a reattachment shock.

The incident shock, separation shock, expansion fan and reattachment shock are all at approximately the same angles and relative position in figures 7.4a and 7.4b. The only obvious difference between the two schlieren images is the absolute location of the interaction. This is due to the 3 mm difference in vertical position of the wedge between the two setups.

A streamwise LDV traverse 15 mm from the tunnel floor on the centre span is shown in figure 7.5. For both setups, the streamwise velocity, u, shows a profile which agrees with the waves captured

a) setup A

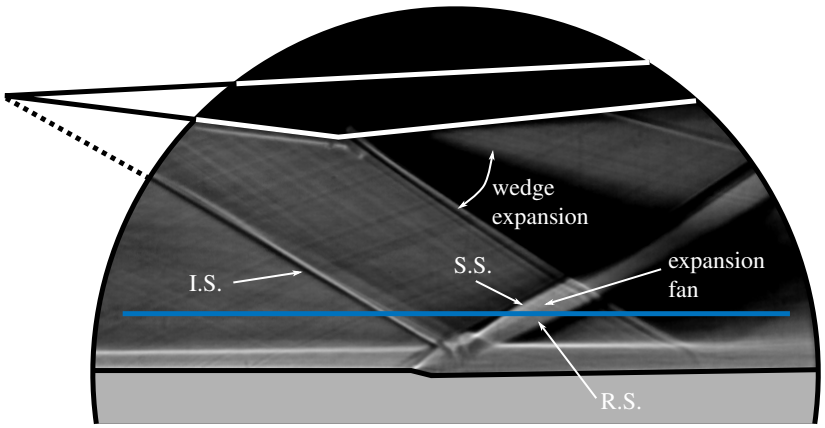

b) setup B

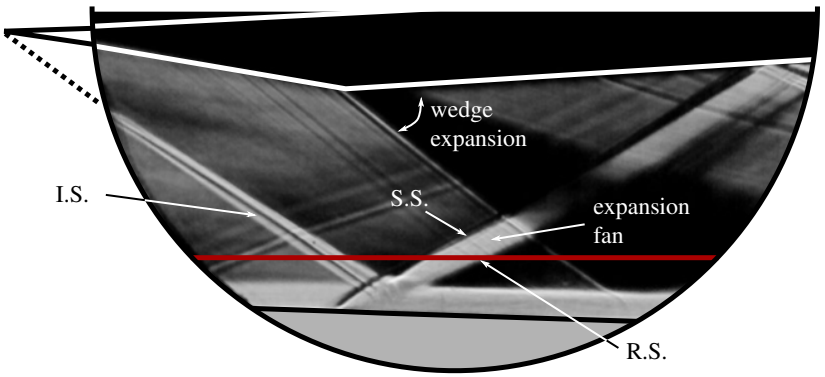

I.S. incident shock
S.S. separation shock
R.S. reattachment shock

Figure 7.4: Schlieren image of the oblique shock–boundary-layer interaction for a) setup A, and b) setup B. The horizontal line denotes the LDV traverse location for figure 7.5.

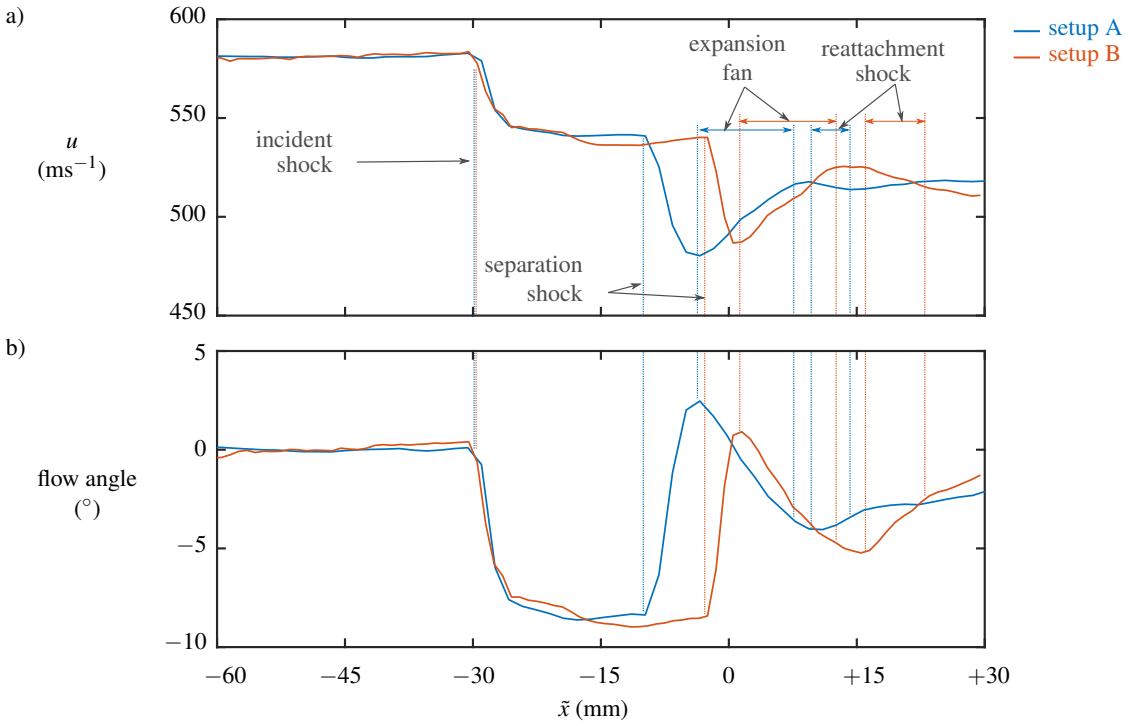

Figure 7.5: Streamwise LDV traverse at $y = 15$ mm and $z = 0$ mm, for both setup A and setup B. The figure shows a) streamwise velocity (u), and b) the local flow angle, measured upwards from horizontal.

in the figure 7.4. The velocity drops as it passes through the incident shock and the flow deflects by $8.2 \pm 1.1^\circ$ in setup A, and by $7.6 \pm 1.1^\circ$ in setup B. In both cases, these error bounds encompass the flow deflection angle of 8° set by the wedge. The following two deceleration regions, which bound an area of re-acceleration, correspond to the separation shock and the reattachment shock respectively. The reattachment shock appears to be weak in strength and smeared, even outside the boundary layer. Note that even prior to the wedge expansion, the flow has not returned to horizontal, but still has a downwards deflection angle of 2°. There is, however, a gradual turning of the flow towards horizontal, perhaps associated with weak waves generated due to the displacement effect as the boundary layer recovers.

Whilst the velocity profiles for both setups in figure 7.5 are similar, there are some subtle differences between them. The separation shock, expansion fan, and reattachment shock in setup A all occur further upstream compared to the corresponding waves from setup B. The apparently obvious conclusion of this difference might be that setup B has a larger separation. However, this is not necessarily the case because the exact velocity profile measured at $y = 15$ mm is determined by a number of factors. In particular, the streamwise position for a given wave is governed by: the streamwise position of the wave origin; the wave angle, determined by how much the flow deflects; the vertical position of the wave origin, i.e. the distance of the sonic line from the floor.

The similarities in incoming boundary layer (figure 7.3a) indicate that the sonic line for the origin of the separation shock is at a similar height from the wall for the two setups. Therefore, measurement of a more upstream separation shock in setup A is consistent with either a more upstream separation point and/or a steeper wave angle, indicating a more severe flow deflection at separation. Meanwhile, the expansion fan measurements also being more upstream in setup A provides some evidence that the separation bubble in this flow is larger in height, since the incident shock has reflected from the sonic line further from the wall. It is not possible, however, to make any similar inferences from the measured reattachment shock position, since the state of the boundary layer (and thus the position of the sonic line) downstream of the interaction could be substantially different for the two setups.

Figure 7.6 shows the surface oil-flow visualisation on the tunnel floor for both setups. Several separated regions, identified from a close study of the skin-friction line topology, are highlighted in the figure. This topology appears to be largely similar between setup A and setup B. In both cases, there is significant separation on the centre span, which is relatively two-dimensional. Separated by a narrow channel of attached flow are regions of corner separation on each side. The spatial extents of the separated regions are listed in table 7.2 for both setups.

For setup A, figure 7.6a shows that the central separation extends 25.2 mm in the streamwise direction and covers 68% of the tunnel span. Meanwhile, the regions of corner separation begin about 24 mm further upstream than the central separation and each cover about 17% of the tunnel span. These measures of the separated flow field agree extremely well with those obtained from previous similar experiments in the same facility, conducted by Xiang [6]. On the other hand, the oil-flow visualisation for setup B (figure 7.6b) has a central separation which extends 16.6 mm in the streamwise direction and which covers 69% of the tunnel span. The regions of corner separation

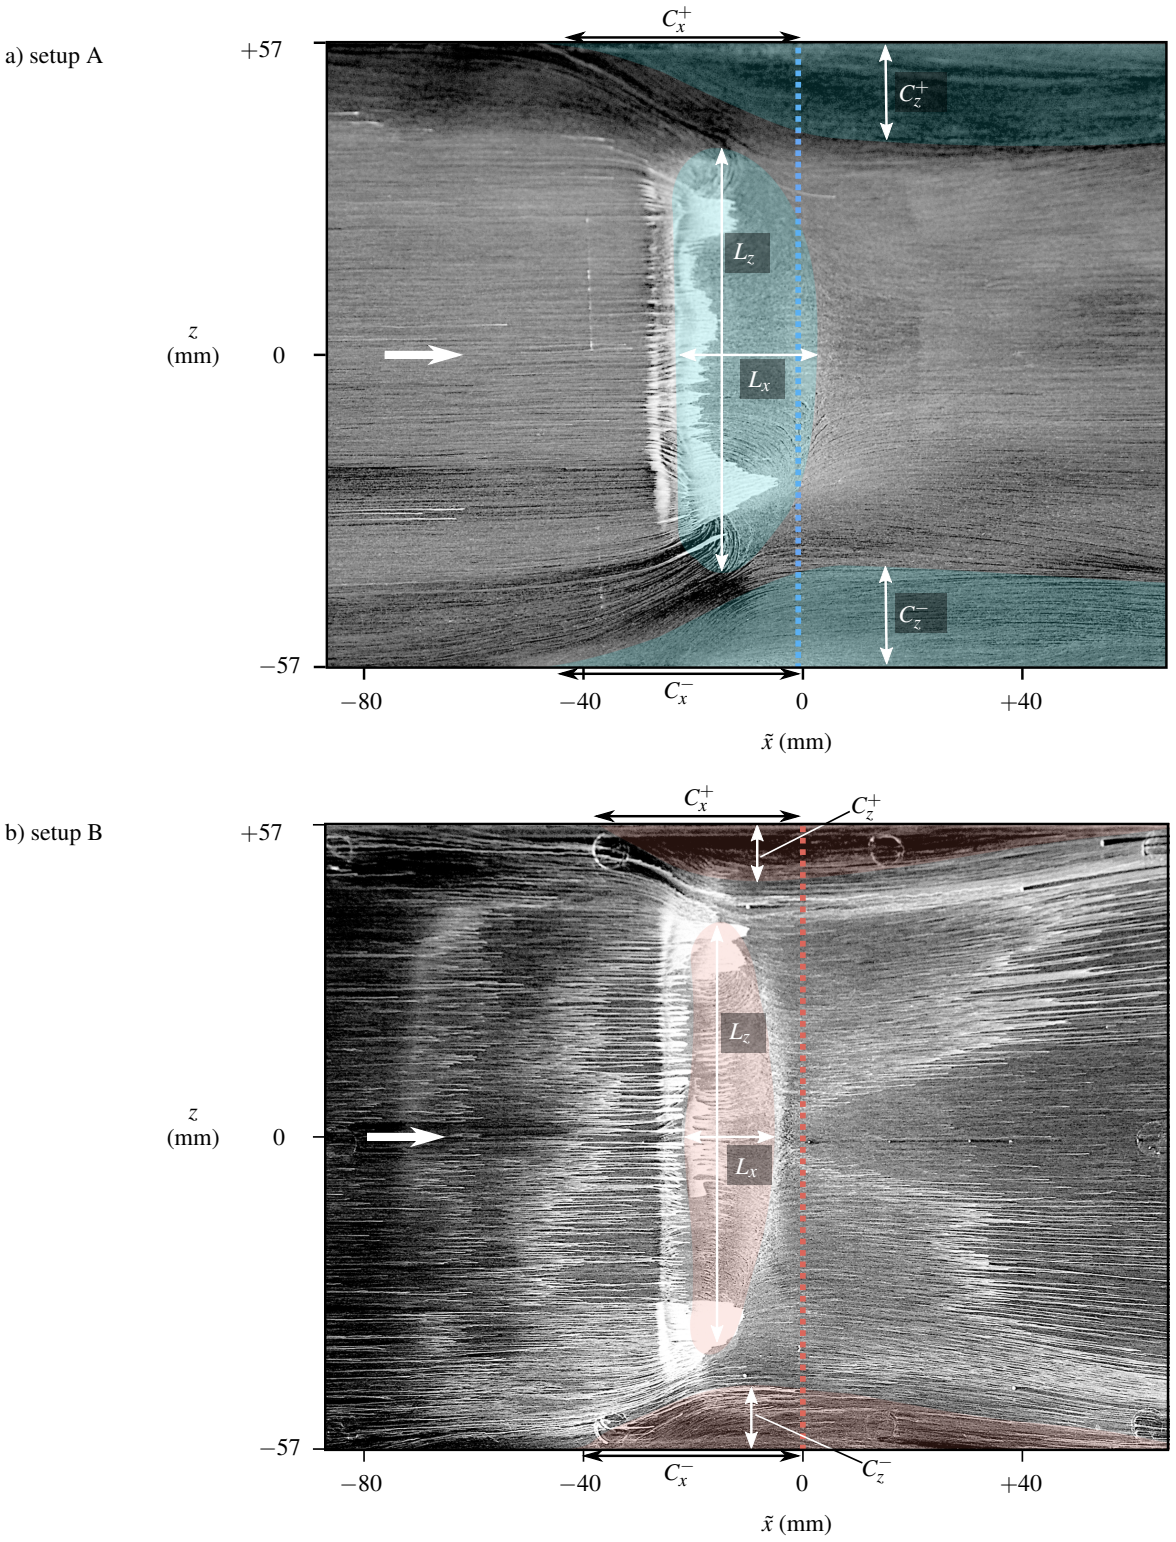

Figure 7.6: Oil-flow image of the separated flow field on the tunnel floor with a) setup A, and b) setup B. The regions of separation are highlighted and the inviscid shock reflection line is marked by a dotted line. The labelled separation parameters are detailed in table 7.2.

Setup	L_x / mm	L_z / mm	C_x^+ / mm	C_z^+ / mm	C_x^- / mm	C_z^- / mm
A	26.4	77.2	46.6	18.7	47.7	19.7
B	16.6	78.2	36.3	10.9	38.2	11.9

Table 7.2: Separation parameters determined from oil-flow visualisation presented in figures 7.6a and 7.6b for setup A and for setup B, respectively. The parameters are defined in figure 7.6.

begin roughly 19 mm further upstream than the central separation and each cover about 10% of the tunnel span.

A comparison between the two oil-flow images in figure 7.6 highlights some important differences between the two setups. Setup B has a corner separation which starts about 5 mm further downstream than in setup A. The corner separation is also roughly 40% less wide for setup B than for setup A. In addition, the width of the corner separation in figure 7.6a for setup A remains close to its maximum extent up to $\tilde{x} = 60$ mm, the end of the measurement region. On the other hand, the corner separation in setup B has a footprint on the tunnel floor which becomes much thinner towards the downstream region of figure 7.6b. A final key difference between the two setups is the streamwise length of the central separation, which is about 35% shorter in setup B than in setup A.

The steady-state surface pressure distribution, measured using pressure-sensitive paint, is presented in figures 7.7a and 7.7b for setup A and for setup B, respectively. In both cases, there is a uniform pressure region upstream of the interaction. Across the SBLI, there is a rapid pressure rise. This is two-dimensional across much of the tunnel span. However, the pressure rise does extend further upstream near the corners, corresponding to a less severe pressure gradient. For both setups, the location and size of the pressure rise correlates well with the separated regions identified from oil-flow visualisation (figure 7.6).

The pressure rise on the tunnel centre span for the two setups are compared in figure 7.7c. The sharp pressure rise corresponding to the separation shock is very similar in both cases. Downstream of the separation shock, instead of a plateau underneath the separation bubble followed by a second pressure rise, the measurements show instead a smeared, more gradual increase in pressure. This pressure rise occurs over a smaller streamwise distance in setup B than in setup A, corresponding to a shorter interaction length.

7.2 Discussion of flow physics

In both setups, the separated flow field exhibits departures from a two-dimensional separation due to corner effects, which are consistent with the mechanisms described in section 2.2.2. Figure 7.6 shows that the corner separation begins ahead of the central separated region. Pressure waves produced by the displacement effect of this corner separation are evident in the surface pressure distribution presented in figure 7.7. These waves are thought to modify the pressure gradient imposed by the incident shock. For example, the narrow channels of attached flow between the central and corner

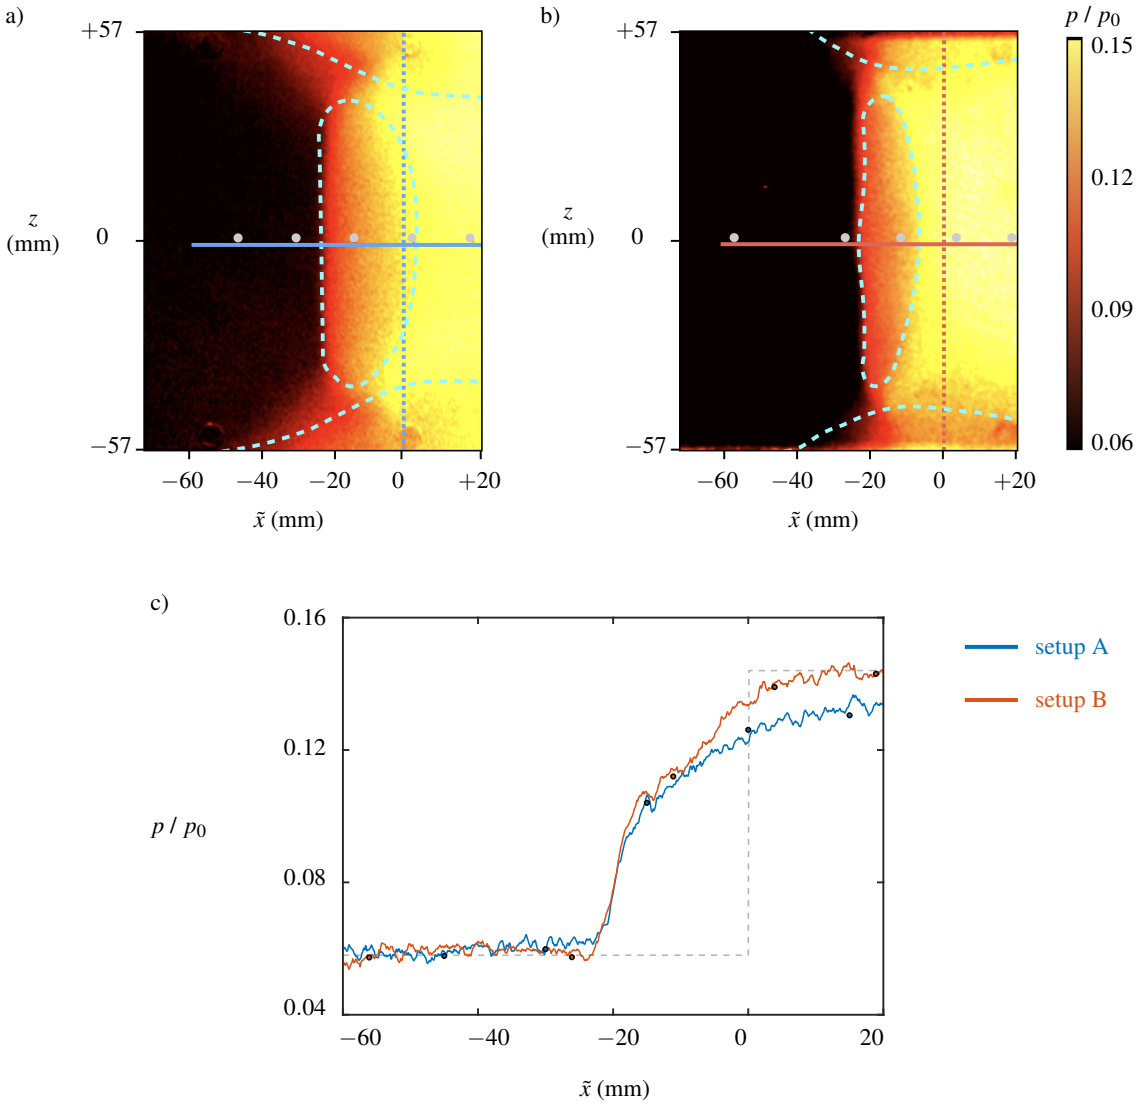

Figure 7.7: Static pressure distribution on the tunnel floor for a) setup A, and b) setup B: the dotted line shows the inviscid interaction location, the dashed line shows separation regions extracted from oil flow and dots mark the locations of pressure taps. c) Pressure measurements at $z = -1$ mm, including a comparison with pressure taps (dashed line shows the inviscid pressure distribution).

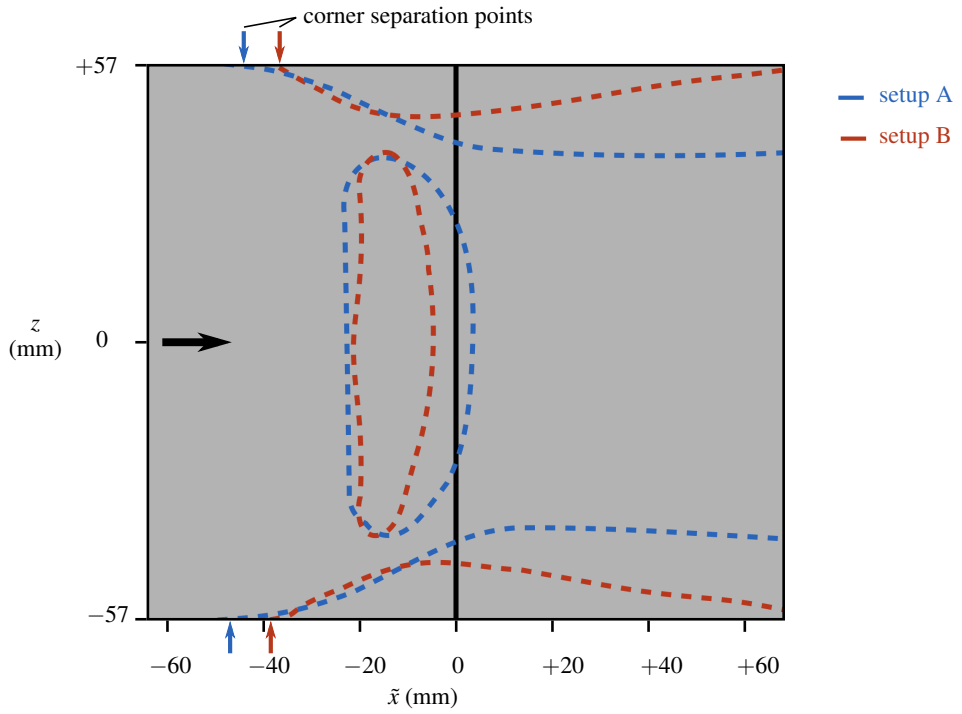

Figure 7.8: Comparison of floor separation topology between setup A (blue) and setup B (red). The topology is determined from the oil-flow visualisation presented in figure 7.6a (setup A) and in figure 7.6b (setup B). The onset of corner separation is marked by arrows, and the solid line indicates the inviscid interaction location.

separations in figure 7.6 are due to the smearing of the overall pressure rise when these waves are ahead of the interaction.

The main differences between the two setups can be summarised by figure 7.8, which shows a comparison of separated regions, determined by the oil-flow visualisation in figure 7.6. Setup B, which has a higher corner momentum, exhibits a more downstream onset of corner separation compared to setup A. Moreover, the corner separation in setup B covers a smaller spanwise extent than setup A. This behaviour is in agreement with an intuitive understanding that a corner boundary layer with higher momentum would be more resistant to separation, and so would exhibit a less severe corner separation.

The differences in the corner flow behaviour affect the rest of the flow field through the waves generated due to the displacement effect of the corner separation. In general, these waves take the form of compression waves followed by expansion waves, as shown for the two setups in figure 7.9. However, the expansion waves in both these setups occur too far downstream to have an effect on the main interaction itself. Therefore, we need only consider the compression waves for now.

The streamwise extent of the central separation, L_x, is about 35% shorter in setup B than in setup A. This effect can be explained by the differences in corner separation. Figure 7.9a shows that in setup A, the more upstream corner separation point causes the compression waves to begin earlier.

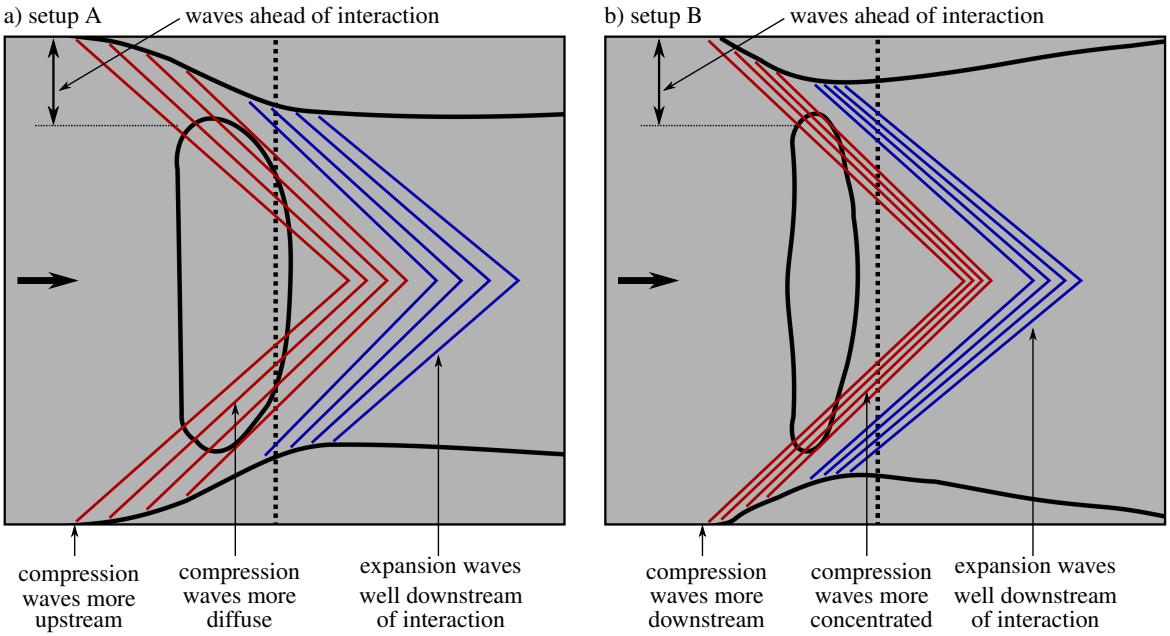

Figure 7.9: Schematic illustration of the footprint of hypothesised corner waves on the tunnel floor for a) setup A, and b) setup B. Red lines correspond to compression waves and blue lines denote expansion waves. Separated regions are marked with a solid line and the dashed line is the inviscid shock reflection location.

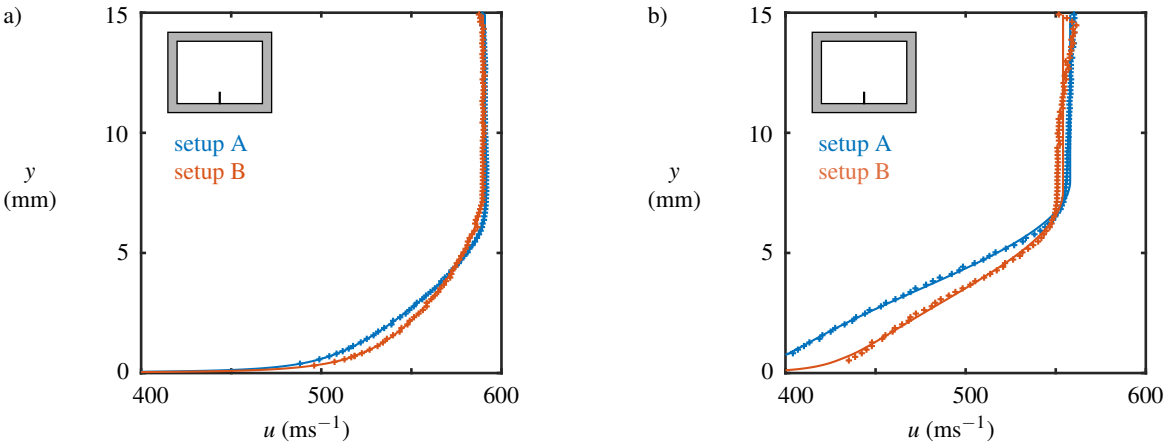

Figure 7.10: Comparison of the floor boundary-layer profile for setup A (blue) and setup B (red). Measurements performed on the centre span at: a) $\tilde{x} = -40$ mm (upstream of the interaction), and b) $\tilde{x} = 20$ mm (downstream of the interaction).

Table 7.3: Incompressible floor boundary-layer parameters, along the tunnel centre line, both upstream and downstream of the interaction for setups A and B. These correspond to the profiles presented in figure 7.10.

	$\tilde{x} = -40$ mm (upstream)				$\tilde{x} = 20$ mm (downstream)			
setup	δ (mm)	δ_i^* (mm)	θ_i (mm)	H_i	δ (mm)	δ_i^* (mm)	θ_i (mm)	H_i
A	6.94	0.92	0.68	1.36	8.00	2.33	1.24	1.88
B	7.74	0.80	0.61	1.29	7.58	1.65	1.05	1.57

These waves impose an additional pressure rise on the central separation, particularly around the expected reattachment location. Whilst the compression waves in figure 7.9a may appear to intersect downstream of the reattachment point, they are still expected to influence the centre-span pressure profile for two reasons. Firstly, moving downstream, the corner waves will enter a regime where the local Mach number is reduced due to the impinging shock. Thus, the waves will be refracted to shallower angles and will intersect further upstream than indicated in figure 7.9a. Secondly, any pressure rise is thought to be transmitted quite effectively in the spanwise direction within the separation bubble, as evidenced by the spanwise-uniform pressure within the central separation in figures 7.7a and 7.7b. Therefore, even on the tunnel centre span the overall pressure rise is increased by the corner compression waves, and so the central separation is extended in the streamwise direction. On the other hand, figure 7.9b shows that setup B has a more downstream corner separation. The trajectories of these waves intersect the very edge of the central separation in setup B, and so they have only a minor impact. There is therefore a smaller additional pressure rise on the tunnel centre span, resulting in a shorter streamwise extent of the central separation.

In contrast, the spanwise extent of the central separation is similar for the two setups with a difference in width, L_z, of only 1.5%. The insensitivity of this parameter is consistent with two separate effects, both related to the corner wave mechanism presented in figure 7.9. These effects relate to the region where the corner waves are ahead of the interaction, in which the overall pressure rise is smeared. The pressure-smearing in setup A takes place over a larger streamwise distance and so is highly effective, helping the incoming boundary layer to stay attached. Whilst the compression waves are more concentrated in setup B, resulting in less efficient pressure-smearing, the expansion waves from the corners propagate through this region much sooner, and reduce the adverse pressure gradient experienced by the flow. These two mechanisms appear to have approximately the same effect, and so the spanwise extent of the central separation is relatively insensitive to the size of corner separation.

The impact of any differences in central separation between the two setups are assessed by comparing the centre-span floor boundary-layer profile at two separate locations. Upstream of the inviscid interaction location, at $\tilde{x} = -40$ mm, figure 7.10a shows that the two setups exhibit very similar profiles. However, there are large differences between the two setups when equivalent

measurements are taken about three boundary-layer thicknesses downstream of the interaction, at $\tilde{x} = 20$ mm (figure 7.10b).

A more quantitative comparison involves analysis of the boundary-layer parameters, listed in table 7.3. The differences in incoming boundary-layer displacement thickness and shape factor are only 13% and 5%, respectively. However, downstream of the interaction, the boundary layer in setup A is 41% thicker, and has a shape factor 20% larger, than setup B. These parameters indicate that the boundary layer in setup A is significantly less healthy. It is worth noting, however, that the recovery of a boundary layer downstream of an SBLI is fastest immediately after the interaction [15]. Whilst the measurements are taken at the same position relative to the inviscid interaction, this measurement location is further downstream of the reattachment line in setup B than in setup A. Therefore, the boundary layer in setup B has been recovering over a greater distance, which exacerbates any differences between the two setups.

Nevertheless, the differences between the boundary-layer profiles in figure 7.10b show that central separation of setup A, with thicker corner boundary layers, has a more severe impact on the incoming flow. Therefore, while the differences between the two setups in the footprint of the interaction on the tunnel floor may appear to be small, the effect on the flow is in fact quite noticeable when considering the state of the boundary layer itself.

7.3 Summary

This chapter investigates the separation behaviour of the floor boundary layer in response to an oblique shock reflection for two otherwise-identical setups with distinctly different corner flows. Even small differences in the momentum within the corner boundary layers are shown to have a substantial impact on the flow field developing in response to an adverse pressure gradient. The corner momentum does not just influence local separation but can also affect the wider flow field.

For instance, differences in the corner flow can significantly influence the streamwise extent of the central separation and the state of the floor boundary layer downstream of the SBLI, even on the centre span of the wind tunnel. This experiment also provides a non-intrusive method to test the mechanisms proposed by Xiang [6], which describe the effects of corner separation in terms of the generation of pressure waves. The measured results are consistent with these mechanisms, and so provide further evidence that the behaviour of the flow in channels is affected by these corner waves.

Chapter 8

Concluding remarks

An extensive review of the existing literature showed the importance of the corner regions of transonic or supersonic channel flows experiencing shock–boundary-layer interactions. The low-momentum corner boundary layer tends to separate earlier than the rest of the floor boundary layer. The compression and expansion waves produced by the displacement effect of the corner separation propagate away from the sidewalls and influence the pressure gradient experienced by the rest of the flow. This results in significant departures from two-dimensionality of the separated floor boundary layer, which affect even the flow on the tunnel centre span. For example, the streamwise separation length on the centre line can vary by a factor of five in response to differences in corner separation.

To predict such flow fields accurately therefore requires a reliable estimate for the corner separation, which depends on the momentum within the corner boundary layer. Unfortunately, there are few comprehensive studies on the supersonic corner boundary layers themselves. Furthermore, the experimental studies that have been performed on the corner boundary layers have reported apparently conflicting flow structures. The causes for these inconsistencies have not been isolated so there is no single corner flow topology that can be used to validate commonly-used numerical methods, which typically struggle to compute these flows reliably. The resulting difficulties in understanding the key physics of these corner flows inhibits supersonic aircraft inlet design, affects the interpretation of data from wind tunnel experiments, and hinders the development of relevant numerical techniques. This thesis addresses this gap in knowledge by performing validation-quality experiments targeted at the corner regions of the Mach 2.5 flow in the rectangular test section of a supersonic wind tunnel.

8.1 Summary of findings

A validation study requires detailed flow characterisation data in order to set up simulations, so that they are representative of the physical flow. Therefore, the wind tunnel experiments were designed to carefully characterise the wind tunnel flow. An emphasis was placed on capturing any inherent ‘imperfections’ so that computations can accurately reproduce these in the simulated flow. The characterisation process involved careful quantitative measurements of the as-installed tunnel

geometry and flow field, whilst critically testing assumptions about the experimental methods and setup.

This more careful procedure, necessitated by the demands of a validation study, confirmed that the measurement techniques are reliable within stated experimental error. For example, a floor boundary-layer profile from laser Doppler velocimetry agrees with Pitot probe measurements to within 1%, and traverses conducted in opposite directions exhibit differences in boundary-layer integral parameters of less than 5%. The detailed measurements performed during the characterisation process also revealed some interesting findings about the flow within the tunnel. For example, the velocities in the test section reflect the symmetry set by the tunnel geometry to within 2%. However, any large asymmetry in the left–right direction is a good indicator of potential leakage from the surrounding air into the wind tunnel. This is a particularly useful feature, since leaks can otherwise be difficult to identify.

The measured flow characterisation data take the form of sidewall static pressures, velocities across the tunnel cross-section, and floor boundary-layer profiles at several streamwise locations. These measurements were used to set up RANS simulations of the wind tunnel flow, which were performed by researchers at the US Air Force Research Laboratory. The close communication maintained throughout this process is a crucial aspect of the design, execution, and analysis of a validation study and led to several synergetic benefits.

The experimental-computational collaboration helped to configure a system to inject seeding particles into the corner regions, which was positioned using the informed guidance of computations. This system enables velocities to be measured in the corner regions to validate relevant numerical methods, as well as in the sidewall boundary layers to set up the computations. The latter measurements revealed the presence of secondary flows, which take the form of vertical velocities within the sidewall boundary layers. These secondary flows are generated by vertical pressure gradients in the nozzle region, and are therefore dependent on the geometry of two-dimensional wind-tunnel nozzles. The vertical velocities transport low-momentum sidewall boundary-layer fluid and so affect its thickness distribution. With a symmetric, 'full' nozzle setup the sidewall boundary layer is thickest at the centre height of the tunnel but an asymmetric, 'half' nozzle arrangement results in the sidewall boundary layer growing continuously from top to bottom.

Another finding from the careful flow characterisation process was the observation of streamwise-coherent flow features in short exposure schlieren images. Engineers at the Air Force Research Laboratory identified these as streamwise vortices, embedded within the tunnel's sidewall boundary layers. Analysis of the computational data performed by the author showed that these vortices perturb the flow in their vicinity and can change the local sidewall boundary-layer thickness by as much as 37%. The vortices are caused by a region of separated flow, due to a sharp geometry change ahead of the two-dimensional nozzle block used in the facility. When this sharp corner was replaced with a more gentle contour, the flow no longer separated and, in turn, the vortices along the tunnel sidewalls disappeared.

There are therefore two ways in which the two-dimensional nozzles can influence the sidewall boundary layers, through both the sidewall secondary flows and through the geometry-induced

vortices. In fact, both these mechanisms can influence the flow structure in the corner boundary layers. The bulk vertical velocities in the sidewall boundary layers strongly influence the thickness of the corner boundary layer. The momentum within this region can be further modified by the topology of streamwise vortices in this region, influenced by both the sidewall secondary flows and the sidewall vortices (though the latter only affects some half nozzle setups which have sharp corners ahead of the contoured nozzle surface). These mechanisms help to explain the apparent inconsistencies in corner flow structure reported in the literature, where experiments were performed in facilities which use different nozzle geometries.

This strong influence of nozzle geometry was exploited to test the impact of corner boundary-layer structure on separation induced by an oblique shock reflection in a channel. Different nozzle geometries were used to produce two setups which have quite distinct corner regions, but are otherwise the same. In response to an incident oblique shock with 8° flow deflection, the setup with thinner corner boundary layers experienced a corner separation which started further downstream and covered a smaller spanwise extent. This behaviour can be directly attributed to a higher-momentum corner flow being more resistant to the adverse pressure gradient of the oblique shock.

The differences in corner separation between the two setups also influence the central separated region of the floor boundary layer. The more benign corner separation in the thin corner boundary-layer setup generates fewer compression waves from the displacement effect of the corner separation. The floor boundary layer at the centre span therefore experiences a smaller additional adverse pressure gradient on top of that imposed by the oblique shock. As a result, the streamwise extent of the central separation is reduced by 35%. This consequently results in the centre-span floor boundary layer downstream of the interaction being 41% thinner and exhibiting a fuller profile.

RANS computations of the wind tunnel conducted by the Air Force Research Laboratory were set up with the help of the detailed flow characterisation data. The author used data from these computations to validate numerical models relevant in the corner regions of supersonic channel flow. Simple turbulence models, based on a linear eddy-viscosity model, do not accurately reproduce the measured corner flow exhibiting velocity differences up to 7%. This can be attributed to their inability to capture vortices within the corner boundary layer, produced by anisotropies in the turbulent stresses, which modify the local momentum. When a quadratic modification to the linear eddy-viscosity model (termed the quadratic constitutive relation) is introduced, the vortices and, in turn, flow structure in the corner regions are better predicted. The discrepancy from experimental velocity data is reduced to about $4-5\%$, but the measured boundary-layer structure is not exactly captured by these simulations. Therefore, whilst vortices are generated, their strength and position appear to differ from the physical corner vortices.

A closer examination of the quadratic constitutive relation reveals that there are two other quadratic terms which were neglected in the original formulation. All three quadratic terms have an identical effect on the turbulent stress combinations which are associated with streamwise vorticity production in corner geometries. Therefore, corner flows demonstrate the importance of these quadratic terms but cannot discriminate between them. A more general version of the quadratic constitutive relation is

proposed which includes two out of the three possible quadratic terms. The relevant constants are tuned using direct numerical solutions of important canonical flows, leading to a better prediction of the anisotropies between turbulent stresses.

Therefore, the key conclusion of this thesis is that the corner boundary layers of supersonic wind tunnels with rectangular cross-section are strongly influenced by the geometry of the two-dimensional nozzles used to produce the supersonic flow. This affects both the thickness and the structure of the corner boundary layers. The governing mechanisms are associated with secondary flows in the sidewall boundary layers induced by vertical pressure gradients in the nozzle and, in some cases, by vortices originating from a region of separated flow ahead of the nozzle. The resulting differences in corner flow, though subtle, can have a substantial impact on the response of the floor boundary layer to an oblique shock reflection.

RANS computations of these flow fields capture both the pressure-driven secondary flows and the geometry-induced vortices. Therefore, difficulties faced by these simulations in computing corner flows appears to be based on their inability to reliably capture the vortices produced by anisotropies in the local turbulent stresses. The quadratic constitutive relation is able to somewhat improve mean flow prediction by generating vorticity in the corner regions, albeit with strength and position slightly different to the physical vortices.

8.2 Implications of findings

The work presented in this thesis highlights the significant impact nozzle geometry has on the corner boundary layers of supersonic wind tunnels. Whilst the full and half nozzle setups are generally considered to be equivalent, this is not quite correct. More specifically, it is important to account for nozzle geometry when interpreting data from any experiments which are focused on the corner regions of the test section or on the response of the tunnel boundary layers to an adverse pressure gradient. This is especially true when comparing data across facilities, where different nozzle geometries can lead to noticeable differences in flow field. Furthermore, any computations which are quantitatively compared with experimental data require a knowledge of the nozzle geometry applied.

A better understanding of the vortex structure within corner boundary layer for different nozzle geometries also enables a more informed approach to flow-control experiments which target corner separation. For example, the position and direction of vortex generators can be configured to produce vortices which enhance or weaken existing vortical structures, as required to increase the local momentum and thus delay separation.

This thesis has also identified a region of separated flow ahead of the nozzle block, which produces vortices in the test section. These vortices can perturb the sidewall boundary layers in which they reside, affecting their local thickness by as much as 37%. The separation associated with the vortices is related to the two-dimensional geometry of the nozzle, typical of supersonic wind tunnels with rectangular cross-section. Facilities with interchangeable nozzle blocks often feature sharp corners in

the subsonic section of the flow, making them potentially susceptible to this problem. Any sidewall vortices can be avoided by geometry modifications of the upstream, subsonic section of the wind tunnel nozzle. Where these modifications are not feasible, experimentalists should consider the presence of sidewall vortices when interpreting schlieren images and other data, particularly relevant in the context of the increasing sensitivity of high-speed cameras.

The findings in this thesis are not entirely restricted to wind tunnels, however. The inlets of supersonic aircraft feature oblique shock waves, which generate vertical pressure gradients that could produce vertical secondary flows in the sidewall boundary layers. These are expected to affect the corner boundary layers and, in turn, the wider flow field. Alongside the different pressure profiles experienced by the floor and ceiling boundary layers, this mechanism could also contribute to additional distortion of the flow entering the engine.

The quadratic constitutive relation has been shown to improve the prediction of corner flows, providing evidence to support its use where such flows are likely to be important. However, the results from these computations should be treated with caution since the position and strength of the generated vortices do not exactly match those in physical flows. This limitation is not specific to the quadratic constitutive relation formulated by Spalart but is a feature of any eddy-viscosity model limited to quadratic terms, as required for simplicity and low computational cost.

8.3 Recommendations for future work

This thesis bridges some gaps in our understanding of supersonic corner flows and their effects, but some open questions remain. Many of these questions are focused on how the structure of the corner boundary layer relates to the separation in this region when the flow encounters a shock.

The experimental results in Chapter 7 show that the momentum of a corner boundary layer affects its local separation by directly comparing two setups. However, this does not allow the quantitative relationship between local momentum and the size of corner separation to be established. Such information requires the experiment to be performed for a range of corner boundary-layer states, from which a predictive model can be extracted. High-quality velocity data obtained around (and inside) these separated flow regions would further contribute to an understanding of the key physical mechanisms.

The results presented in Chapter 5 show that RANS computations can predict corner flows reasonably well. In order to test whether this success extends to accurate calculation of separated flow fields, it is necessary to perform similar validation-type simulations on the flow with an oblique shock reflection. A direct comparison with wind tunnel experiments would aim to establish the ability of various turbulence models and eddy-viscosity models to capture the governing flow physics.

Chapter 6 of this thesis proposes a more general quadratic constitutive relation, which appears to better predict the Reynolds stress anisotropies for a wide range of flows. For simple geometries, these anisotropies are not expected to influence the mean flow, which is the primary concern of RANS simulations. However, in more complex cases, we often do expect the turbulent stresses to influence

the mean flow, perhaps through production of secondary flows. In these cases, the better prediction of Reynolds stress anisotropy by a more general quadratic constitutive relation might result in a more accurate mean flow in computations. This provides compelling reasons to conduct RANS simulations with the general form of the relation. These computations would permit evaluation of the numerical stability of this modification, and would enable the effect of the proposed modification on mean flow calculations to be assessed.

References

[1] S.A. Carioscia, J.W. Locke, I.D. Boyd, M.J. Lewis, and R.P. Hallion. Commercial development of civilian supersonic aircraft. Technical report, Institute for Defense Analyses, 2019.

[2] R.Y. Tang, B. Elias, L. Luther, and D. Morgan. Supersonic passenger flights. Technical report, Congressional Research Service, 2018.

[3] A. Kharina, T. MacDonald, and D. Rutherford. Environmental performance of emerging supersonic transport aircraft. Technical report, The International Council on Clean Transportation, 2018.

[4] I.H. Rettie and W.G.E. Lewis. Design and development of an air intake for a supersonic transport aircraft. *Journal of Aircraft*, 5(6):513–521, 1968.

[5] J. Benek. Corner flows: Motivation and objectives. Comments made at the *10th Annual Shock Wave/Boundary Layer Interaction (SWBLI) Technical Interchange Meeting*, 2017.

[6] X. Xiang and H Babinsky. Corner effects for oblique shock wave/turbulent boundary layer interactions in rectangular channels. *Journal of Fluid Mechanics*, 862:1060–1083, 2019.

[7] N. Titchener and H. Babinsky. Shock wave/boundary-layer interaction control using a combination of vortex generators and bleed. *AIAA Journal*, 51(5):1221–1233, 2013.

[8] H. Babinsky, J. Oorebeek, and T. Cottingham. Corner effects in reflecting oblique shock-wave/boundary-layer interactions. In *51st AIAA Aerospace Sciences Meeting*, 2013-0859.

[9] B.F. Carroll and J.C. Dutton. Characteristics of multiple shock wave/turbulent boundary-layer interactions in rectangular ducts. *Journal of Propulsion and Power*, 6(2):186–193, 1990.

[10] G. Kurth and C. Bauer. Air intake development for supersonic missiles. In *44th AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit*, 2008-5263.

[11] K.M. Arun, S. Tiwari, and A. Mani. Three-dimensional numerical investigations on rectangular cross-section ejector. *International Journal of Thermal Sciences*, 122:257–265, 2017.

[12] F.B. Gessner, S.D. Ferguson, and C.H. Lo. Experiments on supersonic turbulent flow development in a square duct. *AIAA Journal*, 25(5):690–697, 1987.

[13] R.R. Morajkar, J.F. Driscoll, and M. Gamba. Experimental study of supersonic turbulent corner flow evolution in a low aspect ratio rectangular channel. In *53rd AIAA Aerospace Sciences Meeting*, 2015-0542.

[14] S.J. Peltier, B.E. Rice, N.J. Bisek, C.K. McKenna, and J.W. Hofferth. Structure of secondary motion in a Mach 2 boundary layer. In *2018 AIAA Aerospace Sciences Meeting*, 2018-0583.

[15] H. Babinsky and J.K. Harvey. *Shock wave–boundary-layer interactions*, volume 32 of *Cambridge Aerospace Series*. Cambridge University Press, 2011.

[16] J. Ackeret, F. Feldmann, and N. Rott. Investigations of compression shocks and boundary layers in gases moving at high speed. Technical report, National Aeronautics and Space Administration, 1947.

[17] F.M. White. *Fluid mechanics*. McGraw-Hill Boston, 1999.

[18] E.L. Houghton and P.W. Carpenter. *Aerodynamics for engineering students*. Butterworth-Heinemann, 2003.

[19] S.M. Hirt and B.H. Anderson. Experimental investigation of the application of microramp flow control to an oblique shock interaction. Technical report, National Aeronautics and Space Administration, 2009.

[20] J.W. Kooi. Experiment on transonic shock-wave boundary layer interaction. Technical report, Nationaal Lucht-en Ruimtevaartlaboratorium, 1975.

[21] J.B. Abbis, L.F. East, et al. A study of the interaction of a normal shock wave and a turbulent boundary layer using a laser anemometer. *RAE TR*, 75141, 1976.

[22] J. Délery, J.G. Marvin, and E. Reshotko. Shock-wave boundary layer interactions. Technical report, Advisory Group for Aersopace Research and Development (France), 1986.

[23] D.R. Smith and A.J. Smits. The effects of successive distortions on a turbulent boundary layer in a supersonic flow. *Journal of Fluid Mechanics*, 351:253–288, 1997.

[24] P.J.K. Bruce, D.M.F. Burton, N.A. Titchener, and H. Babinsky. Corner effect and separation in transonic channel flows. *Journal of Fluid Mechanics*, 679:247–262, 2011.

[25] I.E. Alber, J.W. Bacon, B.S. Masson, and D.J. Collins. An experimental investigation of turbulent transonic viscous-inviscid interactions. *AIAA Journal*, 11(5):620–627, 1973.

[26] M. Sajben, M.J. Morris, T.J. Bogar, and J.C. Kroutil. Confined normal-shock/turbulent-boundary-layer interaction followed by an adverse pressure gradient. *AIAA Journal*, 29(12):2115–2123, 1991.

[27] R.M. Chriss, W.R. Hingst, A.J. Strazisar, and T.G. Keith. An LDA investigation of the normal shock wave boundary layer interaction. *La Recherche Aérospatiale (English Edition)*, no. 2:1–15, 1990.

[28] P.J. Bruce, H. Babinsky, B. Tartinville, and C. Hirsch. Corner effect and asymmetry in transonic channel flows. *AIAA Journal*, 49(11):2382–2392, 2011.

[29] J. Benek, C. Suchyta, and H. Babinsky. The effect of tunnel size on incident shock boundary layer interaction experiments. In *51st AIAA Aerospace Sciences Meeting*, 2013-0862.

[30] X. Xiang and H. Babinsky. Corner effects in oblique shock wave/boundary layer interactions in rectangular channels. In *55th AIAA Aerospace Sciences Meeting*, 2017-0984.

[31] H. Babinsky. Corner effects in SBLIs. Comments made at the *10th Annual Shock Wave/Boundary Layer Interaction (SWBLI) Technical Interchange Meeting*, 2017.

[32] A. Weber, H.-A. Schreiber, R. Fuchs, and W. Steinert. 3D transonic flow in a compressor cascade with shock-induced corner stall. In *ASME Turbo Expo 2001: Power for Land, Sea, and Air*. American Society of Mechanical Engineers Digital Collection, 2001.

[33] T. Handa, M. Masuda, and K. Matsuo. Three-dimensional normal shock-wave/boundary-layer interaction in a rectangular duct. *AIAA Journal*, 43(10):2182–2187, 2005.

[34] P. Doerffer and R. Szwaba. UFAST experiments data bank: Unsteady effects of shock wave induced separation. *Gdańsk*, pages 131–177, 2009.

[35] H. Babinsky, Y. Li, and C.W. Pitt Ford. Microramp control of supersonic oblique shock-wave/boundary-layer interactions. *AIAA Journal*, 47(3):668, 2009.

[36] Y. Zhang, H. Tan, F. Tian, and Y. Zhuang. Control of incident shock/boundary-layer interaction by a two-dimensional bump. *AIAA Journal*, 2014.

[37] P. Bookey, C. Wyckham, and A. Smits. Experimental investigations of Mach 3 shock-wave turbulent boundary layer interactions. In *35th AIAA Fluid Dynamics Conference*, 2005-4899.

[38] A. Sathianarayanan and S. Verma. Experimental investigation of a Mach 4 shock-wave turbulent boundary layer interaction near an expansion corner. In *53rd AIAA Aerospace Sciences Meeting*, 2015-0112.

[39] W.E. Eagle and J.F. Driscoll. Shock wave–boundary layer interactions in rectangular inlets: three-dimensional separation topology and critical points. *Journal of Fluid Mechanics*, 756:328–353, 2014.

[40] X. Xiang. *Corner effects for oblique shock wave/turbulent boundary layer interactions in rectangular channels*. PhD thesis, University of Cambridge, 2018.

[41] J.A. Benek, C.J. Suchyta, and H. Babinsky. Simulations of incident shock boundary layer interactions. In *54th AIAA Aerospace Sciences Meeting*, 2016-0352.

[42] M.P. Pizzella, S. Warning, M.W. McQuilling, A. Purkey, M. Scharnhorst, R.and Mani, J.A. Benek, C.J. Suchyta, and H. Babinsky. On the effect of test section aspect ratio for shock wave-boundary layer interactions. In *55th AIAA Aerospace Sciences Meeting*, 2017-0535.

[43] B. Wang, N.D. Sandham, Z. Hu, and W. Liu. Numerical study of oblique shock-wave/boundary-layer interaction considering sidewall effects. *Journal of Fluid Mechanics*, 767:526–561, 2015.

[44] P.J.K. Bruce, H. Babinsky, B. Tartinville, and C. Hirsch. Corner effect and asymmetry in transonic channel flows. *AIAA Journal*, 49(11):2382–2392, 2011.

[45] F.B. Gessner and J.B. Jones. On some aspects of fully-developed turbulent flow in rectangular channels. *Journal of Fluid Mechanics*, 23(4):689–713, 1965.

[46] S.G. Rubin. Incompressible flow along a corner. *Journal of Fluid Mechanics*, 26(1):97–110, 1966.

[47] N. Tokuda. Viscous flow near a corner in three dimensions. *Journal of Fluid Mechanics*, 53(1):129–148, 1972.

[48] A.J. Baker and J.A. Orzechowski. An interaction algorithm for three-dimensional turbulent subsonic aerodynamic juncture region flow. *AIAA Journal*, 21(4):524–533, 1983.

[49] K.N. Ghia. Incompressible streamwise flow along an unbounded corner. *AIAA Journal*, 13(7):902–907, 1975.

[50] S.G. Rubin, P.K. Khosla, and S. Saari. Laminar flow in rectangular channels. Part II: Numerical solution for a square channel. *Computers & Fluids*, 5(3):151–173, 1977.

[51] S. Pirozzoli, D. Modesti, P. Orlandi, and F. Grasso. Turbulence and secondary motions in square duct flow. *Journal of Fluid Mechanics*, 840:631–655, 2018.

[52] G.M. Bragg. The turbulent boundary layer in a corner. *Journal of Fluid Mechanics*, 36(3):485–503, 1969.

[53] M. Phillips, S. Deutsch, A. Fontaine, and S. Yavuzkurt. Three dimensional measurements of instantaneous flow field in corners with smooth surfaces under a zero pressure gradient. In *ASME Turbo Expo 2004: Power for Land, Sea, and Air*, pages 1019–1027, 2004.

[54] M. Zamir and A.D. Young. Experimental investigation of the boundary layer in a streamwise corner. *The Aeronautical Quarterly*, 21(4):313–339, 1970.

[55] P. Bradshaw. Turbulent secondary flows. *Annual Review of Fluid Mechanics*, 19(1):53–74, 1987.

[56] E. Brundrett and W.D. Baines. The production and diffusion of vorticity in duct flow. *Journal of Fluid Mechanics*, 19(3):375–394, 1964.

[57] S. B. Pope. Turbulent flows, 2001.

[58] L. Prandtl. *Essentials of fluid dynamics: With applications to hydraulics, aeronautics, meteorology and other subjects*. Blackie & Son, 1953.

[59] R. Pecnik and G. Iaccarino. Predictions of turbulent secondary flows using the v^2-f model. *Center for Turbulence Research, Annual Research Briefs, Stanford University*, pages 333–343, 2007.

[60] F.B. Gessner. The origin of secondary flow in turbulent flow along a corner. *Journal of Fluid Mechanics*, 58(1):1–25, 1973.

[61] H.J. Perkins. The formation of streamwise vorticity in turbulent flow. *Journal of Fluid Mechanics*, 44(4):721–740, 1970.

[62] T. von Kármán and H.S. Tsien. Boundary layer in compressible fluids. *Journal of the Aeronautical Sciences*, 5(6):227–232, 1938.

[63] D. Coles. The turbulent boundary layer in a compressible fluid. *The Physics of Fluids*, 7(9):1403–1423, 1964.

[64] D. Modesti, S. Pirozzoli, and F. Grasso. Direct numerical simulation of developed compressible flow in square ducts. *International Journal of Heat and Fluid Flow*, 76:130–140, 2019.

[65] J. Boussinesq. *Essai sur la théorie des eaux courantes*. Imprimerie nationale, 1877.

[66] W. Rodi. A new algebraic relation for calculating the Reynolds stresses. In *GaMuM Workshop, Paris, France*, volume 56, 1976.

[67] S.S. Girimaji. A galilean invariant explicit algebraic reynolds stress model for turbulent curved flows. *Physics of Fluids*, 9(4):1067–1077, 1997.

[68] R. Cécora, A. Eisfeld, B.and Probst, S. Crippa, and R. Radespiel. Differential reynolds stress modeling for aeronautics. In *50th AIAA Aerospace Sciences Meeting*, volume 465, 2012.

[69] T.B. Gatski and C.G. Speziale. On explicit algebraic stress models for complex turbulent flows. *Journal of Fluid Mechanics*, 254:59–78, 1993.

[70] P.R. Spalart. Strategies for turbulence modelling and simulations. *International Journal of Heat and Fluid Flow*, 21(3):252–263, 2000.

[71] M. Mani, D. Babcock, C. Winkler, and P.R. Spalart. Predictions of a supersonic turbulent flow in a square duct. In *51st AIAA Aerospace Sciences Meeting*, 2013-0860.

[72] N.J. Bisek. High-fidelity simulations of the UTSI Mach 2 test section. In *2018 AIAA Aerospace Sciences Meeting*, 2018-1097.

[73] C.L. Rumsey, J. Carlson, and N. Ahmad. FUN3D Juncture Flow computations compared with experimental data. In *AIAA Scitech 2019 Forum*, 2019-0079.

[74] J. Dandois. Improvement of corner flow prediction using the Quadratic Constitutive Relation. *AIAA Journal*, 52(12):2795–2806, 2014.

[75] T.J. Leger, N.J. Bisek, and J. Poggie. Supersonic corner flow predictions using the Quadratic Constitutive Relation. In *22nd AIAA Computational Fluid Dynamics Conference*, 2015-3432.

[76] D.O. Davis and F.B. Gessner. Further experiments on supersonic turbulent flow development in a square duct. *AIAA Journal*, 27(8):1023–1030, 1989.

[77] I. Bermejo-Moreno, L. Campo, J. Larsson, J. Bodart, D. Helmer, and J.K. Eaton. Confinement effects in shock wave/turbulent boundary layer interactions through wall-modelled large-eddy simulations. *Journal of Fluid Mechanics*, 758:5–62, 2014.

[78] N.J. Bisek. High-order implicit large-eddy simulations of supersonic corner flow. In *52nd Aerospace Sciences Meeting*, 2014-0588.

[79] D.B. Helmer, L.M. Campo, and J.K. Eaton. Three-dimensional features of a Mach 2.1 shock/boundary layer interaction. *Experiments in Fluids*, 53(5):1347–1368, 2012.

[80] B.E. Rice, N.J. Bisek, S.J. Peltier, and J.W. Hofferth. Investigation of secondary motion in high speed flow. In *55th AIAA Aerospace Sciences Meeting*, 2017-0526.

[81] D.O. Davis, F.B. Gessner, and G.D. Kerlick. Experimental and numerical investigation of supersonic turbulent flow through a square duct. *AIAA Journal*, 24(9), 1986.

[82] T. Davidson. *Effect of incoming boundary layer state on flow development downstream of normal shock wave–boundary layer interactions*. PhD thesis, University of Cambridge, 2016.

[83] A. Pope and K.L. Goin. *High-speed wind tunnel testing*. John Wiley & Sons, 1965.

[84] S.P. Colliss. *Vortical structures on three-dimensional shock control bumps*. PhD thesis, University of Cambridge, 2014.

[85] G.S. Settles. *Schlieren and shadowgraph techniques: Visualizing phenomena in transparent media*. Springer Science & Business Media, 2001.

[86] Photron. *Datasheet: FASTCAM NOVA S6/S9/S12*, 2020.

[87] J.L. Brown and J.W. Naughton. The thin oil film equation. Technical report, National Aeronautics and Space Administration, 1999.

[88] A. Stanbrook. The surface oil flow technique for use in high speed wind tunnels. AGARD 70:39–49. Technical report, NATO AGARD, 1962.

[89] J. Délery. *Three-dimensional separated flow topology: Critical points, separation lines and vortical structures*. John Wiley & Sons, 2013.

[90] L.C. Squire. The motion of a thin oil sheet under the steady boundary layer on a body. *Journal of Fluid Mechanics*, 11(2):161–179, 1961.

[91] A. Thom and C.J. Apelt. The pressure in a two-dimensional static hole at low Reynolds numbers. Technical report, Aeronautical Research Council, 1957.

[92] P.J.K. Bruce. *Transonic shock/boundary layer interactions subject to downstream pressure perturbations*. PhD thesis, University of Cambridge, 2008.

[93] R. Shaw. The influence of hole dimensions on static pressure measurements. *Journal of Fluid Mechanics*, 7(4):550–564, 1960.

[94] R.E. Rayle. An investigation of the influence of orifice geometry on static pressure measurements. Master's thesis, Massachusetts Institute of Technology, 1949.

[95] N. Titchener. *An experimental investigation of flow control for supersonic inlets*. PhD thesis, University of Cambridge, 2013.

[96] T. Liu, M. Guille, and J.P. Sullivan. Accuracy of pressure-sensitive paint. *AIAA Journal*, 39(1):103–112, 2001.

[97] J.W. Gregory, K. Asai, M. Kameda, T. Liu, and J.P. Sullivan. A review of pressure-sensitive paint for high-speed and unsteady aerodynamics. Proceedings of the Institution of Mechanical Engineers. *Part G: Journal of Aerospace Engineering*, 222(2):249–290, 2008.

[98] B.G. McLachlan and J.H. Bell. Pressure-sensitive paint in aerodynamic testing. *Experimental Thermal and Fluid Science*, 10(4):470–485, 1995.

[99] M. Sajben. Uncertainty estimates for pressure sensitive paint measurements. *AIAA Journal*, 31(11):2105–2110, 1993.

[100] S. Koike and H. Babinsky. PSP measurements using a commercial digital camera applied to shock wave-boundary layer interactions with a streamwise vortex. Presented at the *6th German-Japanese Joint Seminar*, 2017.

[101] K. Sami. *Physics of three-dimensional normal shock wave/turbulent boundary layer interactions in rectangular channels*. PhD thesis, University of Cambridge, 2012.

[102] F.A. MacMillan. Experiments on Pitot-tubes in shear flow. Technical report, Aeronautical Research Council, 1957.

[103] B. McKeon et al. *Velocity, vorticity and Mach number*, pages 215–471. Handbook of Experimental Fluid Mechanics. Springer, 2007.

[104] F. Durst, A. Melling, and J.H. Whitelaw. *Principles and practice of laser Doppler anemometry*. Academic Press, London, 1976.

[105] W. Lai, J. Shakal, and D. Troolin. Accuracy, resolution, and repeatability of Powersight PDPA and LDV systems. Technical report, P/N 5001520, TSI Incorporated, 2013.

[106] H.C. van de Hulst. *Light scattering by small particles*. John Wiley & Sons, 1957.

[107] R.A. Humble, G.E. Elsinga, F. Scarano, and B.W. Van Oudheusden. Three-dimensional instantaneous structure of a shock wave/turbulent boundary layer interaction. *Journal of Fluid Mechanics*, 622:33–62, 2009.

[108] P. Buchhave and W.K. George. Bias corrections in turbulence measurements by the laser Doppler anemometer. In *Laser velocimetry and particle sizing: Proceedings of the Third International Workshop on laser Doppler anemometry*, 1978.

[109] M.V. Morkovin. Effects of compressibility on turbulent flows. *Mécanique de la Turbulence*, pages 367–380, 1962.

[110] N. Titchener, S. Colliss, and H. Babinsky. On the calculation of boundary-layer parameters from discrete data. *Experiments in Fluids*, 56(8):159, 2015.

[111] J.M. Oorebeek. *Comparison of distributed suction and vortex generator flow control for a transonic diffuser*. PhD thesis, University of Cambridge, 2014.

[112] C.C. Sun and M.E. Childs. A modified wall wake velocity profile for turbulent compressible boundary layers. *Journal of Aircraft*, 10(6):381–383, 1973.

[113] E.R. Van Driest. Turbulent boundary layer in compressible fluids. *Journal of the Aeronautical Sciences*, 18(3):145–160, 1951.

[114] A.J. Musker. Explicit expression for the smooth wall velocity distribution in a turbulent boundary layer. *AIAA Journal*, 17(6):655–657, 1979.

[115] P.G. Buning et al. Overflow user's manual. Technical report, NASA Langley Research Center, 2002.

[116] E.F. Toro, M. Spruce, and W. Speares. Restoration of the contact surface in the HLL-Riemann solver. *Shock Waves*, 4(1):25–34, 1994.

[117] B. Koren. Upwind schemes, multigrid and defect correction for the steady Navier-Stokes equations. In *11th International Conference on Numerical Methods in Fluid Dynamics*, pages 344–348. Springer, 1989.

[118] R. Tramel and R. Nichols. A highly efficient numerical method for overset-mesh moving-body problems. In *13th Computational Fluid Dynamics Conference*, 1997-2040.

[119] P. Spalart and S. Allmaras. A one-equation turbulence model for aerodynamic flows. In *30th Aerospace Sciences Meeting and Exhibit*, 1992-0439.

[120] S.R. Allmaras and F.T. Johnson. Modifications and clarifications for the implementation of the Spalart-Allmaras turbulence model. In *Seventh International Conference on Computational Fluid Dynamics*, ICCFD7-1902, 2012.

[121] J. Benek, J. Steger, and F.C. Dougherty. A flexible grid embedding technique with application to the Euler equations. In *6th Computational Fluid Dynamics Conference*, 1983-1944.

[122] Pointwise Version. 18.0R4. *Pointwise Incorporated*, 2017.

[123] F. Menter and C. Rumsey. Assessment of two-equation turbulence models for transonic flows. In *Fluid Dynamics Conference*, 1994-2343.

[124] F.R. Menter, M. Kuntz, and R. Langtry. Ten years of industrial experience with the SST turbulence model. *Turbulence, Heat and Mass Transfer*, 4(1):625–632, 2003.

[125] D.C. Wilcox. Reassessment of the scale-determining equation for advanced turbulence models. *AIAA Journal*, 26(11):1299–1310, 1988.

[126] D.C. Wilcox. *Turbulence Modeling for CFD*. DCW Industries Incorporated, 3rd edition, 2006.

[127] W.L. Oberkampf and T.G. Trucano. Verification and validation in computational fluid dynamics. *Progress in Aerospace Sciences*, 38(3):209–272, 2002.

[128] D.P. Aeschliman and W.L. Oberkampf. Experimental methodology for computational fluid dynamics code validation. *AIAA Journal*, 36(5):733–741, 1998.

[129] W.L. Oberkampf and C.J. Roy. *Verification and validation in scientific computing*. Cambridge University Press, 2010.

[130] P.J. Roache. *Fundamentals of verification and validation*. Hermosa Publishers, 2009.

[131] C.J. Roy and W.L. Oberkampf. A comprehensive framework for verification, validation, and uncertainty quantification in scientific computing. *Computer Methods in Applied Mechanics and Engineering*, 200(25-28):2131–2144, 2011.

[132] B.W. Skews. Aspect ratio effects in wind tunnel studies of shock wave reflection transition. *Shock Waves*, 7(6):373–383, 1997.

[133] S. Ristić. Flow visualization techniques in wind tunnels: Optical methods (Part II). *Scientific Technical Review*, 57(2):38–49, 2007.

[134] F. Lu, S. Ali, E. Braun, and L. Maddalena. A modern compressible flow laboratory experience for undergraduates. In *49th AIAA Aerospace Sciences Meeting*, 2011-0274.

[135] V. Kolář. Vortex identification: New requirements and limitations. *International Journal of Heat and Fluid Flow*, 28(4):638–652, 2007.

[136] H. Tennekes and J.L. Lumley. *A first course in turbulence*. MIT Press, 1972.

[137] M. Gramola, P.J.K. Bruce, and M. Santer. Experimental FSI study of adaptive shock control bumps. *Journal of Fluids and Structures*, 81:361–377, 2018.

[138] C.M. Jackson, W.A. Corlett, and W.J. Monta. Description and calibration of the Langley Unitary Plan Wind Tunnel. Technical report, National Aeronautics and Space Administration, 1981.

[139] L.M. Edelman and M. Gamba. Structure of a Mach 2 shock train from experimental measurements. In *AIAA Scitech 2020 Forum*, 2020-1609.

[140] P.R. Spalart, A. Garbaruk, and M. Strelets. RANS solutions in Couette flow with streamwise vortices. *International Journal of Heat and Fluid Flow*, 49:128–134, 2014.

[141] S.B. Pope. A more general effective-viscosity hypothesis. *Journal of Fluid Mechanics*, 72(2):331–340, 1975.

[142] M. Bernardini, S. Pirozzoli, and P. Orlandi. Velocity statistics in turbulent channel flow up to $\mathrm{Re}_\tau = 4000$. *Journal of Fluid Mechanics*, 742:171–191, 2014.

[143] S. Hoyas and J. Jiménez. Reynolds number effects on the Reynolds-stress budgets in turbulent channels. *Physics of Fluids*, 20(10):101511, 2008.

[144] J.A. Sillero, J. Jiménez, and R.D. Moser. One-point statistics for turbulent wall-bounded flows at Reynolds numbers up to $\delta^+ \approx 2000$. *Physics of Fluids*, 25(10):105102, 2013.

[145] P. Schlatter, R. Örlü, Q. Li, G. Brethouwer, J.H.M. Fransson, A.V. Johansson, P.H. Alfredsson, and D.S. Henningson. Turbulent boundary layers up to $\mathrm{Re}_\theta = 2500$ studied through simulation and experiment. *Physics of Fluids*, 21(5):051702, 2009.

[146] P. Schlatter, Q. Li, G. Brethouwer, A.V. Johansson, and D.S. Henningson. Simulations of spatially evolving turbulent boundary layers up to $\mathrm{Re}_\theta = 4300$. *International Journal of Heat and Fluid Flow*, 31(3):251–261, 2010.

[147] G.K. El Khoury, P. Schlatter, A. Noorani, P.F. Fischer, G. Brethouwer, and A.V. Johansson. Direct numerical simulation of turbulent pipe flow at moderately high Reynolds numbers. *Flow, Turbulence and Combustion*, 91(3):475–495, 2013.

[148] J.P. Monty, N. Hutchins, H.C.H. Ng, I. Marusic, and M.S. Chong. A comparison of turbulent pipe, channel and boundary layer flows. *Journal of Fluid Mechanics*, 632:431–442, 2009.

[149] P.R. Spalart, M.L. Shur, M.K. Strelets, and A.K. Travin. Direct simulation and RANS modelling of a vortex generator flow. *Flow, Turbulence and Combustion*, 95(2-3):335–350, 2015.

[150] G.N. Coleman, C.L. Rumsey, and P.R. Spalart. Numerical study of turbulent separation bubbles with varying pressure gradient and Reynolds number. *Journal of Fluid Mechanics*, 847:28–70, 2018.