

Chapter 6

Taming real world flow control experiments with MLC

"An approximate answer to the right problem is worth a good deal more than an exact answer to an approximate problem."

- John Tukey

In Chapter 2, MLC was introduced as a generic method to identify optimal control laws for arbitrary dynamics. In Chapters 4 and 5, MLC has been applied to the control of low-dimensional dynamical systems. In these examples, we have shown (1) that it is comparable to optimal linear control design for linear dynamics, (2) that it outperforms linear control methods in the case of a weak nonlinearity, and (3) that it can even identify the enabling strongly nonlinear actuation mechanism in the case that the linear dynamics are uncontrollable.

In this chapter, we describe and exemplify the application of MLC to real-world turbulence control experiments. These come with the additional challenges of high-dimensional dynamics, long time delays, high-frequency noise, low-frequency drifts and, last but not least, with the non-trivial implementation of the algorithm in the experimental hardware. In experiments, MLC is executed in the same way as for the dynamical system plants in Chapters 4 and 5:

1. MLC provides a generation of control laws to be evaluated by the plant.
2. The plant is used to evaluate and grade these individuals with respect to the given cost function.
3. MLC evolves the population.
4. The process stops when a pre-determined criterion is met.
5. After this learning phase, the best control law can be used.

The only difference between a simulation and an experiment is the need to interrogate an experimental plant. This is a technical challenge but not a conceptual point of departure from MLC. Running MLC on an experiment using an existing code is

a matter of a few days to a week of work, if the experimental hardware and software is ready for closed-loop control.

We choose three configurations: a laminar flow over a backward-facing step in a water tunnel (Sec. 6.1), separating turbulent boundary layers in wind tunnels (Sec. 6.2), and a turbulent mixing layer (Sec. 6.3). These examples encompass key phenomena encountered in most flow control problems: boundary layers, separation, mixing layers and a recirculation zone. Different kind of sensors, actuators and time scales are used and illustrate the versatility of MLC. Section 6.4 highlights the limitations of model-based linear control for the turbulent mixing layer. In Sec. 6.5, we focus on implementation issues with respect to software and hardware. Section 6.6 suggests reading on a spectrum of flow control aspects. Our interview (Sec. 6.7) addresses past and future developments in experimental closed-loop turbulence control with Professor Williams, a pioneer and leading scholar of this field.

6.1 Separation control over a backward-facing step

The first example is the control of the recirculation zone behind a backward-facing step in a water tunnel. This experiment has been performed by Nicolas Gautier and Jean-Luc Aider in the PMMH laboratory at ESPCI, Paris. This first application of MLC in separation control is described in detail in [111].

6.1.1 Flow over a backward-facing step

The flow over a backward-facing step is an important benchmark configuration in fluid mechanics. From a practical perspective, this flow represents cold mixing in a simplified combustor configuration. Increases in cold mixing, i.e. smaller recirculation zones, indicate better combustion in corresponding reacting flow. From a more academic perspective, the flow features an incoming boundary layer, which undergoes a geometrically forced separation at the edge of the step [9, 141, 22]. Excluding creeping flow, the fluid particles are unable to follow the 90° turn at the step and separate. Thus, a mixing layer is formed between the recirculation zone and the outer flow. As is typical for a mixing layer, a Kelvin-Helmholtz instability triggers vortex shedding. This advective mixing increases the width of the vorticity region downstream and determines the position of the reattachment point. This point defines the length of the recirculation zone, L_{rec} . A literature survey indicates that acting on the shear layer at the most amplified frequency contributes to the build-up of eddies, enhances mixing and, at the end, reduces the length of the recirculation zone [68, 64, 130, 112].

The described separation phenomena can be witnessed in many applications with sharp corners: flows over mountains, cars, buildings, mixing elements in microflu-

idics, just to name a few. This separation affects the lift and drag of the obstacle, the resulting noise level or the mixing.

The actuation of the experiment is set up so that there is no direct access to the shear-layer oscillation at the high-receptivity point, the corner. For a sensor signal, MLC is given the extent of the recirculation zone defined as the area with backward flow. This sensor is blind to the phase of the Kelvin-Helmholtz vortex shedding. The interest in this choice of sensor is twofold. First, MLC must now discover the shedding frequency indirectly or else MLC has to find another mechanism to control the flow. Second, from an engineering point of view, the recirculation area is much easier to infer than the state of the shear layer. For instance, optical monitoring of seeding or combustion processes can be used instead of particle image or hot-wire velocimetry. This is particularly true for hot reacting flow which rules out most velocity and pressure measurements but which may be optically assessed.

6.1.2 Experimental setup at PMMH

The PMMH water tunnel (Fig. 6.1a) operates a gravity driven flow with velocities up to 22 m/s. Its test section is $L = 0.8$ m long, $l = 15$ cm wide and $H_{\text{section}} = 20$ cm high (before the step). The step height is $h_{\text{step}} = 1.5$ cm, as depicted in Fig. 6.1b. The operating Reynolds number for the experiment results presented here is $Re = U_{\infty} \times h_{\text{step}}/\nu = 1350$, where ν is the kinematic viscosity of water. At this Reynolds number, the flow is well separated, creating a shear layer between the high-speed stream and the low-speed recirculation zone extending from the step to the reattachment point. Furthermore, at this Reynolds number, large vortices are visible in the mixing layer.

The control goal is to reduce the size of the recirculation zone. The actuation is achieved by blowing or sucking in a nominally two-dimensional slot upstream of the step, which is oriented at 45° with respect to the streamwise and wallnormal directions. By regulating the pressure in the reservoir of the slot, the exit velocity of the slot can be changed (positive for a jet, negative in case of suction). The exit velocity is taken as actuation command b .

The recirculation size is monitored by a Real-Time (RT) Particle Image Velocimetry (PIV) system which determines the flow fields at a 42 Hz frequency. This sampling frequency is over one order of magnitude larger than the characteristic Kelvin-Helmholtz frequency of around 1 Hz. However, the PIV system would be too slow for the following wind-tunnel experiments with frequencies around 10 to 100 Hz. There are many possibilities to estimate the size of a recirculation zone. We choose a simple quantity, namely the area in which instantaneous streamwise velocity is negative. The sensor signal $s(t)$ is defined as the normalized instantaneous recirculation zone,

$$s(t) = \frac{S_a(t)}{S_{a,u}}, \quad (6.1)$$

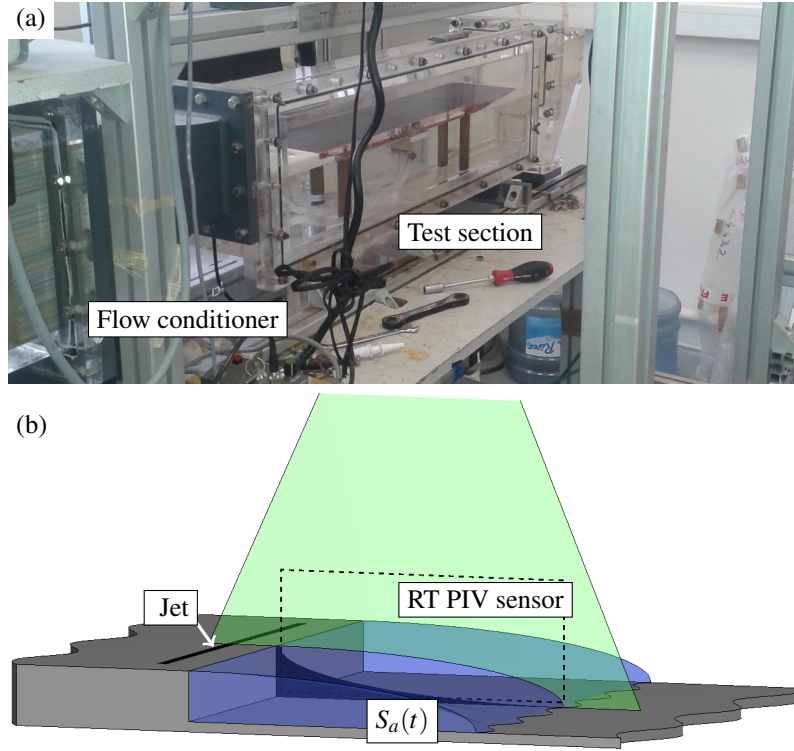


Fig. 6.1 (a) Photograph of the PMMH experiment. (b) Experimental configuration. A slotted jet is situated just upstream of the separation and performs blowing or suction in the boundary layer. A laser sheet is placed in the symmetry plane for real-time PIV and to determine the backward-flow region.

where

$$S_a(t) = \int H(-u(x,y,t)) dx dy,$$

$$S_{a,u} = \langle S_a(t) \rangle_T, \text{ without actuation.}$$

Here, u is the streamwise velocity component, H the Heaviside function and $\langle \cdot \rangle_T$ a time-averaged value of its argument over period T . Note that $S_{a,u}$ is the time-averaged recirculation area for the uncontrolled flow. The chosen sensor is not sensitive to the shear-layer vortices for the reasons mentioned above.

The goal function J reads:

$$J = \langle s \rangle_T + \gamma \langle |b| \rangle_T^2, \quad (6.2)$$

where γ is a penalization coefficient for the actuation b . This parameter sets the admissible trade-off between realizing the control objective (reducing the recircu-

lation) and the cost of the actuation. A low value of γ will favor performance over economy and a high value gives preference to economy over performance.

Table 6.1 MLC parameters used for the control of the backward-facing step flow.

Parameter	Value
N_i	500
P_r	0.1
P_m	0.20
P_c	0.70
N_p	7
N_e	1
Node functions	+, -, ×, /, exp, log, tanh

Ideally, the actuation power investment should be measured against the achieved power savings. One example is aerodynamic drag reduction of a car in which the cost function comprises the saving in propulsion power minus the invested actuation power, i.e. the net energy savings. In the backward-facing step, mixing enhancement has no direct energetic benefit and we resort to an order-of-magnitude argument for choosing γ . Departure point is an optimal periodic forcing which minimizes the recirculation zone. In particular, the penalization term of this periodic forcing is set equal to the normalized size of the recirculation zone, i.e. unity. This arguably avoids the case where the actuation cost is under- or overemphasized and leads to a value of $\gamma = 3/2$. The parameters used to apply MLC are listed in the Tab. 6.1.

Advanced material 6.1 Activated options for experimental applications.

Contrary to previous implementations of OpenMLC in this book, there are the following changes for experiments:

- (a) All individuals of a new generation are evaluated, even if they have already been evaluated in previous generations. The cost function of an individual is the average value over all generations.
- (b) The 5 best individuals of the final generation are evaluated 5 times. The cost function is, now, the average value from the final generation.
- (c) Control laws that are dismissed as unpromising in the first generation (e.g. valves closed or open more than 90% of the time) are replaced until all individuals of the first generation are considered as legitimate candidates.

The options (a) and (b) are extremely important, for two reasons: (1) The evaluation time $T = 10$ s is made small in order to have the fastest evaluation time for a generation. Having more evaluation time means also better statistical values for the individuals that keep appearing. (2) This rule prevents inefficient individuals that have obtained an accidentally good evaluation to stay in the top of the ranking and contaminate the next generation. Option (c) ensures that the first generation contains a large spectrum of promising control laws.

6.1.3 Results

After 12 generations, MLC returned the best control law. Its behavior is illustrated in Fig. 6.2a. When the control starts, the sensor value s goes from an average of 1 (by definition) to an average of 0.28, which corresponds to more than 70% reduction. The other curve presented in the figure illustrates the control command providing a visual estimate of the actuation cost of the control law. Though it appears that most of the time the control command is set at $b = 0.5$, this individual has a cost of $J_{MLC} = 0.419$ while the best open-loop command, a periodic command at 1 Hz (corresponding to a Strouhal number of 0.2, displayed in Fig. 6.2b) has a cost of $J_{OL} = 0.423$ as defined by Eq. (6.2).

These comparable cost function values show that MLC has been able to find a control law which is as good as the best periodic forcing. A spectral analysis of the control command and the sensor signal under optimal MLC-based forcing [113] shows that the frequency which is most amplified by the shear layer is neither noticeable in the sensor signal nor exploited in the actuation command. The study suggests that MLC has found a novel way to act directly on the recirculation bubble with frequencies on the order of one tenth of the Kelvin-Helmoltz instability. This is compatible with the so-called flapping frequency of the bubble [253].

Furthermore, the resulting control law is closed-loop by construction. Thus, it should add intrinsic robustness to changing flow conditions since the control command is decided by flow events and not by a predetermined signal. Both MLC and optimal periodic forcing have been tested under off-design conditions, using the same cost function defined by Eq. (6.2). The results are displayed in Tab. 6.2. Unsurprisingly, the periodic forcing performs poorly on off-design conditions. In comparison, the MLC law performs better (in terms of minimizing the cost J) at lower and higher Reynolds numbers that were not included in the learning process.

Table 6.2 Cost function of MLC, J_{MLC} and open loop J_{OL} at different Reynolds numbers.

Re_h	J_{OL}	J_{MLC}
900	0.75	0.33
1350	0.42	0.42
1800	0.76	0.59

This constitutes a key highlight of the experiment: MLC has found a new unexpected actuation mechanism which is more robust against changing operating conditions than the optimal periodic forcing. Intriguingly, closed-loop control has also outperformed periodic forcing in experiments with simple phasor control. Examples are mixing enhancement after a backward-facing step [208] and drag reduction of a d-shaped cylinder [209].

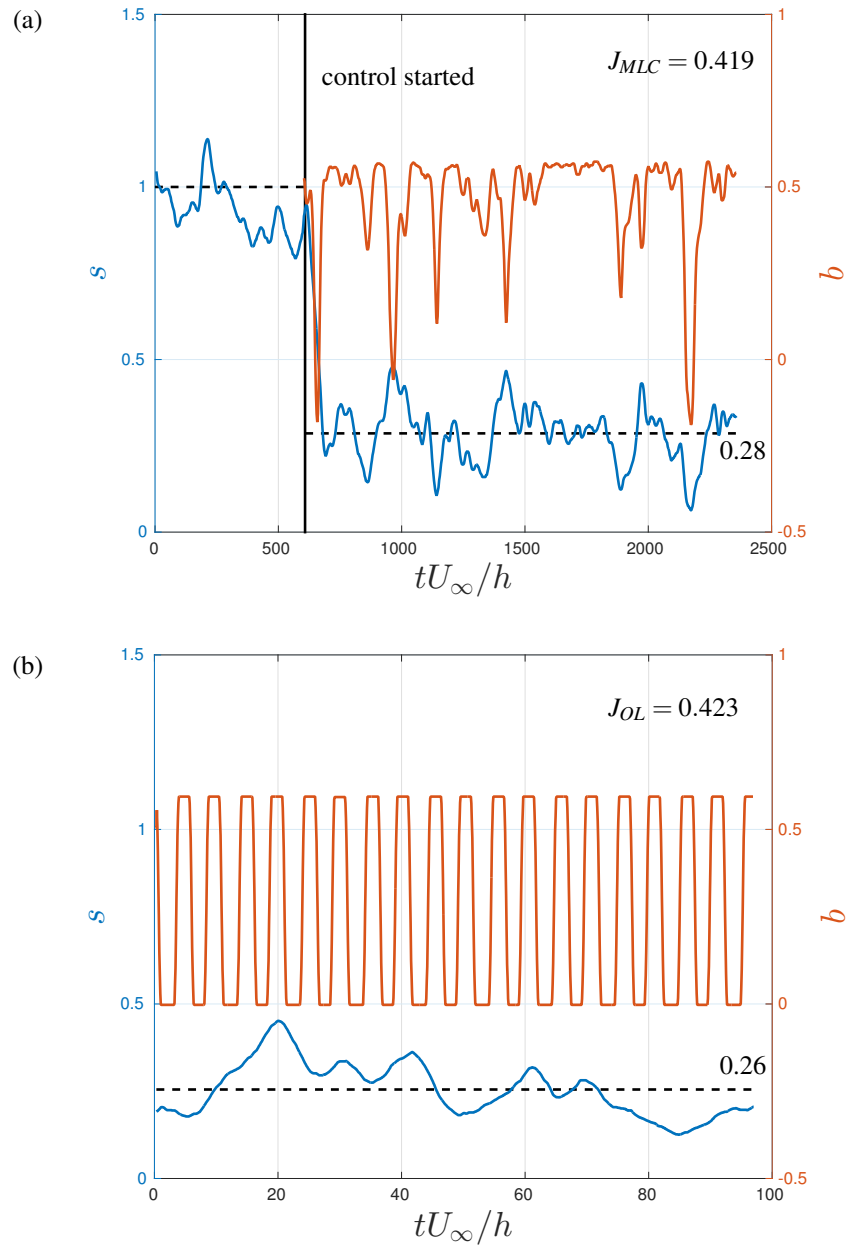


Fig. 6.2 Sensor and control command for the control of the PMMH backward-facing step flow by MLC (a) and best periodic forcing (b).

6.2 Separation control of turbulent boundary layers

In the second example, we also reduce a recirculation zone by closed-loop actuation. Yet, the geometry is a smooth ramp in which the separation point is not geometrically prescribed but physically controlled. In fact, MLC has been applied on two geometrically similar sister experiments in the Laboratoire de Mécanique de Lille, France, with Cécric Raibaud, Christophe Cuvier and Michel Stanislas and in the Laboratoire PRISME from Orleans, France, with Antoine Debien, Nicolas Mazellier and Azeddine Kourta. Both experiments feature a turbulent boundary layer which separates under the influence of an identical descending ramp. Both experiments use vortex generators (though slightly different in geometry) as actuators and hot-films as sensors. The main difference between both experiments are (1) the Reynolds number, with a factor of 10 difference between experiments, and (2) the use of different cost functions.

6.2.1 Separating boundary layers

Every obstacle in ambient high-speed flow will generate a boundary layer on its surface. By definition, the boundary layer is a large gradient zone of the velocity between the no-slip condition at the wall and the outer potential flow. The force on the body is strongly dependent on the separation line which in turn depends on the state of the boundary layer. Hence, boundary layer separation is pivotal for the control of aerodynamic forces, like drag reduction behind a bluff body or lift increase of an airfoil. Flow separation gives rise to a shear layer, also called a mixing layer, between the slow recirculation zone and the fast outer stream.

This shear layer is prone to a Kelvin-Helmholtz instability, which generates large-scale spanwise vortices [191, 261]. Near the point of separation, the corresponding shedding frequency is characterized by a Strouhal number $St_{\Theta} = 0.012$ based on the boundary-layer momentum thickness Θ before separation [128, 280]. The vortices of the shear layer shed with a Strouhal number $St_{L_{sep}} = 0.6 - 0.8$ where L_{sep} is the separation length [62, 75, 182].

The shear layer separates the oncoming flow from the recirculation bubble extending from the separation line to the reattachment line. The reattachment point is determined by the efficiency of mixing between the high-speed oncoming flow and the low-speed recirculation bubble. The recirculation bubble is typically associated with low pressure region which increases the drag of the body. Large recirculation regions also reduce the lift force on wings.

Flow control can mainly target two mechanisms to manipulate this flow: (1) change the kinetic energy level of the boundary layer to prevent/promote separation, and (2) change the mixing properties of the separated shear layer to prevent/promote reattachment. The first mechanism is achieved by blowing/sucking the boundary layer or by using vortex generators. The second mechanism is achieved by manipulating the mixing layer. For instance, introducing perturbations at a sen-

sitive frequency may excite vortex shedding which promotes an earlier reattachment. The use of unsteady vortex generators (UVG) enables the exploitation of both mechanisms: the streamwise vortices re-distribute the kinetic energy in the separating boundary layer, while the unsteadiness promotes vortex formations in separated shear layer. These types of actuators have been used extensively for the control of boundary layers. The optimal parameters for control such as the geometry, position and frequencies are still widely discussed [173, 120, 121, 243, 123, 8].

6.2.2 *Experimental setups at LML and PRISME*

The experiments have been carried out in the Malavard closed-loop wind tunnel at the PRISME laboratory [79] and the closed-loop LML wall-turbulence wind tunnel [74]. The AVERT profile used in both experiments is the descending ramp as detailed in [74] for the LML wind tunnel. It features a sharp edge to force the position of the separation line (see Fig. 6.3). The initial slant angle of the ramp is 25° , and can be characterized by its height h_{ramp} and length ℓ . The Malavard wind tunnel has a 2 m wide and 5 m long cross section and the height and length of the ramp are $h = 100$ mm and $\ell = 470$ mm, respectively. With a free-stream velocity of $U_\infty = 20$ m/s, the Reynolds number $Re_h = U_\infty h_{\text{ramp}}/\nu$ is around 1.3×10^5 , where ν is the kinematic viscosity of air. In the LML wind tunnel, the Reynolds number is ten times smaller due to lower velocities and smaller geometry.

For control purposes, unsteady vortex generators (UVG) have been implemented one boundary layer thickness upstream of the sharp edge ramp (see Fig. 6.3). Their design, location and orientation have been chosen based on the results from [121, 248, 73]. The UVG are set up so that the vortices are co-rotating in the LML wind tunnel, and counter-rotating in the PRISME wind tunnel. The jet velocity ratio is $V_{\text{jet}}/U_\infty = 3$. The jets are made unsteady by the use of identical electro-valves which can operate in an on/off fashion up to 300 Hz.

In both experiments, the friction downstream of the separation line is monitored by hot-films (Fig. 6.3). Additionally, unsteady pressure taps are available in the PRISME experiments for the computation of the cost-matching function. In both cases, PIV is used as a post-processing tool to assess the effect of the control on the flow. Unlike the water-tunnel experiment, PIV can neither be used in real-time nor in the learning loop. The feedback sensors s_i used for the MLC control law are based on the hot-film signals:

$$s_i = \frac{h_i - h_{i,u}}{h_{i,\text{max}} - h_{i,u}} \quad \text{with } i = A, B, \quad (6.3)$$

where h_i is the raw voltage output of sensor i , $h_{i,u}$ is the average voltage for the uncontrolled case (corresponding to a separated flow and low friction) and $h_{i,\text{max}}$ is the average voltage for the most effective constant blowing leading to maximal friction.

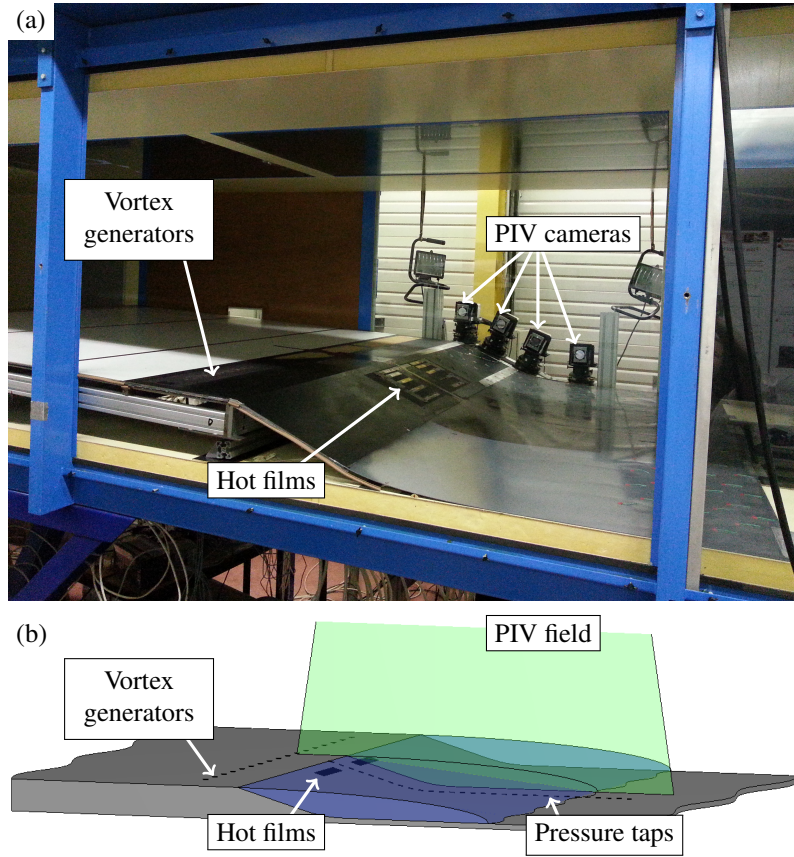


Fig. 6.3 (a) Photograph of the test section of the LML wind tunnel. (b) Experimental configuration of the separating boundary layer in the PRISME wind tunnel. The jets are placed to generate counter-rotating stream-wise vortices. Hot-film sensors are placed after the separation line and static pressure sensors are located in the symmetry plane.

The control law is subject to a cost function promoting a reduced recirculation zone and penalizing blowing:

$$J = J_{\text{HF}} + \gamma_{\text{pstat}} J_{\text{pstat}} + \gamma_{\text{act}} J_{\text{act}}, \quad (6.4)$$

with J_{HF} being an evaluation of the friction recorded from the hot-films, J_{pstat} an evaluation based on the static pressure distribution and J_{act} an evaluation of the actuation cost. $\gamma_{\text{pstat}} = 1/200$ and $\gamma_{\text{act}} = 0.6$ are chosen as penalization coefficients. The evaluation based on the friction is defined as:

$$J_{\text{HF}} = \frac{1}{N_{\text{HF}}} \sum_{i=1}^{N_{\text{HF}}} [1 - \tanh(\langle s_i \rangle_T)], \quad (6.5)$$

where $N_{\text{HF}} = 2$ is the number of hot-film sensors. The value of J_{HF} is 1 when no effect is recorded and approaches 0 as the friction increases. The evaluation based on the static pressure is defined as:

$$J_{\text{pstat}} = \frac{1}{0.1 + \sum_{i=1}^{14} \langle (p(x_i, t))_T - \langle p_u(x_i, t) \rangle_T \frac{x_{\text{max}} - x_i}{x_{\text{max}} - x_{\text{min}}} \rangle^2}, \quad (6.6)$$

with x_i being the position of the i th pressure tap after the edge, x_{min} the position of the pressure tap closest to the edge, x_{max} the furthest downstream pressure tap, $p(x_i, t)$ the static pressure recorded at position x_i and $p_u(x_i, t)$ the static pressure recorded at position x_i in the uncontrolled case. $\langle \cdot \rangle_T$ is the average over the evaluation time T . J_{pstat} is equal to 10 when controlled and uncontrolled pressure distributions coincide and approach zero when they significantly deviate from each other, with a weight which increases linearly towards the separation point. The LML experiment has no pressure taps and this term is not taken into account.

The evaluation of the actuation cost is defined by:

$$J_{\text{act}} = \begin{cases} \langle \frac{Q}{Q_u} \rangle_T & \text{at PRISME,} \\ \langle b^2 \rangle_T & \text{at LML,} \end{cases} \quad (6.7)$$

where Q is the flow-rate and Q_u the flow-rate under constant blowing. In the LML experiment the flow rate could not be integrated in the learning loop and the control command was used instead. In both cases, J_{act} is equal to 1 for constant blowing and vanishes when no actuation is recorded. The parameters of the MLC algorithm are summarized in Tab. 6.3.

Table 6.3 MLC parameters used for the control of the PRISME and LML separating boundary layer.

Parameter	Value (PRISME)	Value (LML)
γ_{pstat}	1/200	0
γ_{act}	0.6	2
N_i	100	500
P_r	0.1	0.1
P_m	0.2	0.25
P_c	0.7	0.65
N_p	7	7
N_e	1	1
Node functions	+, -, ×, /, exp, log, tanh	

6.2.3 Results

Both experiments led to successful separation mitigation of the boundary layer such as displayed in Fig. 6.4. In the case of the LML experiment, the best open-loop control reference is a constant blowing of the UVG at $V_{jet}/U_\infty = 3$ leading to $J_{OL} = 3$. MLC achieves a control with similar performance but at a reduced actuation cost leading to $J_{MLC} = 2.1$ (Fig. 6.5). In this case, γ_{act} directly selects the operating point of the system as the dominating mechanism seems to be strongly linked to the kinetic energy injection in the boundary layer. The mechanism can be summarized as follows: the more one blows, the more the boundary layer attaches. Vanishing penalization leads to constant blowing while high penalization leads to the unforced state.




Fig. 6.4 Cartographies of back-flow coefficient χ in the LML experiment. This coefficient is defined in analogy of the intermittency factor as the average percentage of the occurrence of a negative streamwise velocity at the considered point [249]. (a) Uncontrolled flow. (b) MLC modified flow. The recirculation zone has been drastically reduced. For both cases the iso-line at $\chi = 50\%$ has been traced.

This dominance of the kinetic energy injection has also been demonstrated in the LML experiment [80]. In that case, the pressure distribution obtained in controlled




Fig. 6.5 Time series of sensor A (a,b,c) and low-pass filtered control signal (e,f,g) for uncontrolled (a,d), open-loop constant blowing control (b,e) and MLC control (c,f) in the LML experiment. The level of friction obtained appears to be proportional to the amount of kinetic energy injection.

cases is compared to one of the baseline flow in Fig. 6.6 (a). Both constant blowing and MLC schemes lead to a reduction of the mean recirculation region since the recovery region associated with the pressure plateau is shifted upstream. Noticeable is the acceleration of the flow induced by the UVGs upstream of the sharp edge location ($x/h_{\text{ramp}} = 0$) as emphasized by the strong decrease in pressure. However, pressure distributions computed for the best open-loop actuation and MLC almost coincide. This implies that the efficiency of both control approaches are approximately equivalent. This is confirmed by the measurement of the separation length L_{sep} from the PIV dataset which is reduced by about 40% when control is applied (see Tab. 6.4). Nevertheless, the actuation cost to achieve the same recirculation zone reduction is significantly lower ($\approx 20\%$) for MLC as evidenced by the momentum coefficient

$$c_{\mu} = \frac{S_j d_c V_j^2}{1/2 S_{\text{ref}} U_{\infty}^2},$$

where S_j represents the cross section of the jets, d_c the duty cycle and S_{ref} the surface of the ramp (defined by the flow exposed surface of the descending part of the ramp) reported in Tab. 6.4.

Intriguingly, the reference open-loop control and MLC are distinctly different in terms of the frequency distribution of the actuation command. The frequency distribution is computed from a zero-crossing algorithm applied to the mass-flow

Table 6.4 Cost function values, separation and actuation properties for the natural, best open-loop and MLC cases in the PRISME separating boundary layer experiment.

Case	natural	open-loop	MLC
J	50.4	0.291	0.32
L_{sep}/h_{ramp}	5.4	3.14	3.16
$c_{\mu} (\times 10^{-4})$	-	16.51	13.66

controller signal (see Fig. 6.6 (b)). Unlike open-loop periodic forcing for which the blowing frequency is fixed at 30 Hz, the frequency distribution of the MLC is broadband. More surprisingly, the frequencies of MLC are significantly larger than for the best periodic forcing. Thus, MLC and periodic forcing yield similar separation mitigation by significantly different underlying actuation mechanisms. A more detailed study of the flow physics is reported in [80].

6.3 Control of mixing layer growth

We finally present a feedback control of a turbulent mixing layer built and operated in the ANR Chair of Excellence ‘TURbulence CONTROL using Reduced-Order Models (TUCOROM)’ (ANR-10-CHEX-0015) at Institute PPRIME, France. This experiment comprises all feedback control challenges described in Chapter 1: The complex vortex dynamics between actuation and sensing leads to high-dimensional dynamics, strong nonlinearity and large time delays. The goal of this experiment is to optimize the mixing between both sides of the shear layer. To date, this is arguably the most challenging implementation of MLC in an experiment.

6.3.1 Mixing layer flows

Mixing layers or, equivalently, shear layers, arise when two streams of flow at different velocities interact. Such mixing layers can be observed in almost every flow. The recirculation zone of bluff-body wakes is bounded by shear layers. The near-field dynamics of a jet is determined by surrounding mixing layers. Any separation phenomenon, e.g. in a diffuser or an airfoil, leads to mixing layers. The convectively unstable dynamics of these flows is particularly rich, as the streamwise evolution is a noise amplifier of upstream perturbations with large frequency bandwidth. The Kelvin-Helmholtz instability leads to large-scale vortex formation in laminar and turbulent mixing layers [134]. These vortices merge in the streamwise direction to form larger and larger structures which decay in a turbulence cascade to increasingly small vortices. Thus, the mixing layer leads inevitably to high-dimensional,

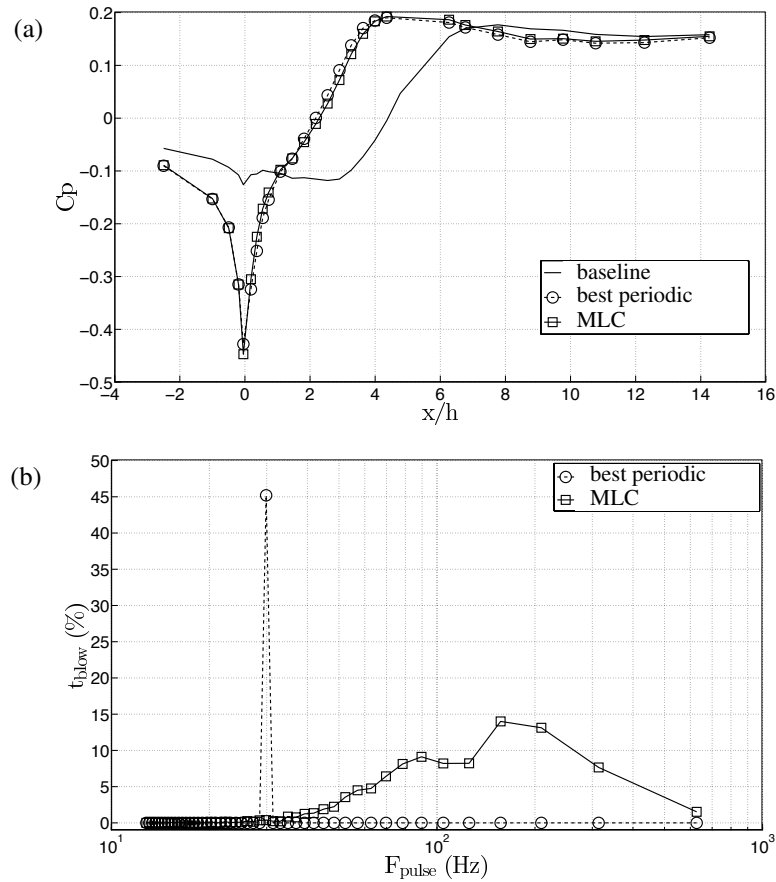


Fig. 6.6 PRISME experiment: (a) Pressure distribution along the ramp. (b) Frequency distribution of the blowing.

multi-scale vortex dynamics. Already by phenomenology it is clear that the possibility of low-dimensional or linear models are very limited as will be corroborated in Sec. 6.4

Nonetheless, simple feedback flow control strategies have been successfully applied to mixing layers [209]. From early periodic forcing studies [134], it is observed that actuating at the frequency of the main structures reinforces said structures and invigorates the mixing, while actuation at much higher frequencies interferes with the production of large structures and mitigates the width of the shear layer.

6.3.2 Experimental setup of the TUCOROM wind tunnel

The TUCOROM experiment has a double turbine wind tunnel aimed at creating a turbulent mixing layer. The wind tunnel generates two flows that can be set at different velocities and which are separated by a plate. The optically accessible test

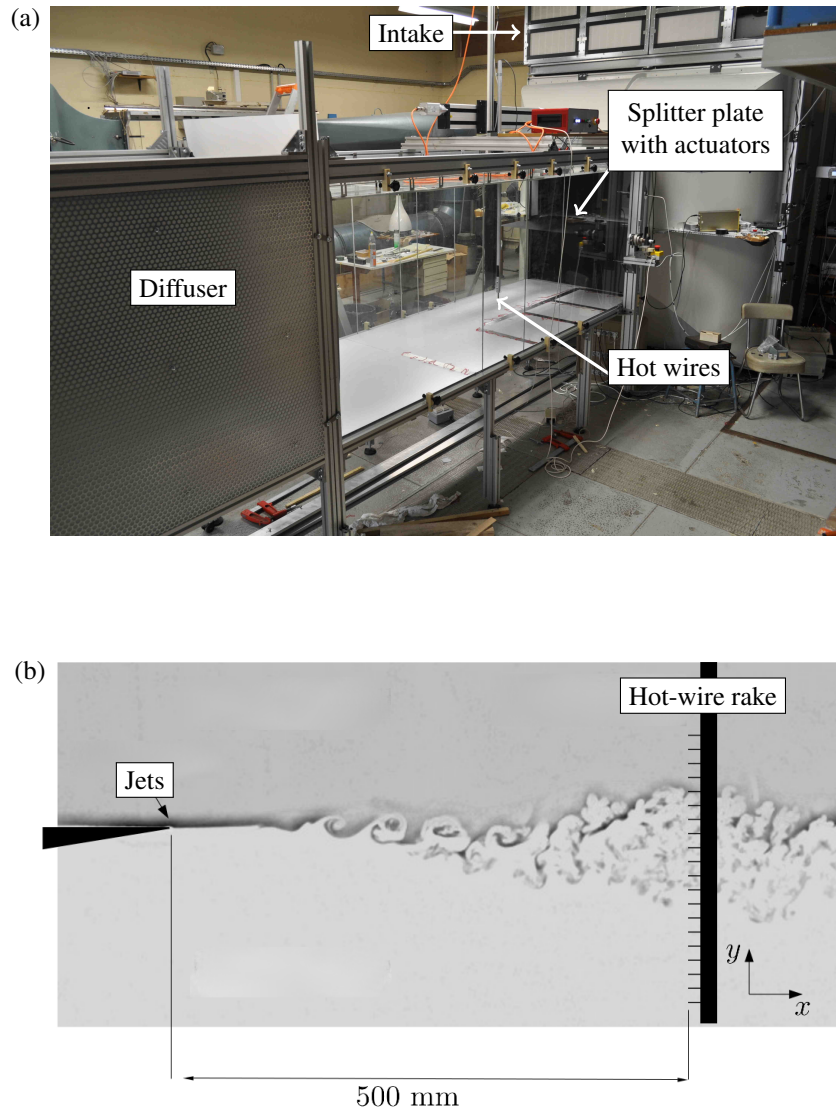


Fig. 6.7 TUCOROM mixing layer demonstrator: (a) Photo of the test section. (b) Experimental setup of the mixing layer. The hot-wire rake is placed at 500mm downstream of separating plate to capture the structures in the shear layer. The spacing of the hot-wire probe is $\Delta y = 8\text{ mm}$.

section has a cross section of $1 \times 1 \text{ m}^2$ and a length of 3 m. The separation plate ends with a triangular profile with a tip 3 mm thick which includes 96 holes for actuation jets blowing in the streamwise direction. The jets can be operated independently. A vertical rake of 24 hot wires, operated at 20 kHz, serves as a sensor array and can be placed at any position downstream in the test-section. The hot wires, vertically separated by a 4 mm offset, map the mixing layer profile. The closed-loop control is implemented on a Concurrent[®] RT system which combines the use of Simulink[®] loop design and Fortran home-made adaptor functions to encode sensor acquisition, compute control decisions and command actuation signals up to a maximum rate of 10 kHz.

The unforced and forced mixing layers are described in detail in [206]. For the results presented here, the wind tunnel is operated with a velocity ratio of $U_1/U_2 = 3.6$, and the Reynolds numbers based on the initial mixing layer thickness is 500 (laminar) for the learning process and can be set to 2000 (turbulent) to test off-design conditions. The hot-wire rake is placed at $x = 500 \text{ mm}$ downstream from the separation plate. The sensors are based on the fluctuations of the raw hot-wires sensors:

$$s_i(t) = h_i(t) - \langle h_i(t) \rangle_{T_{rms}}, \quad (6.8)$$

where $h_i(t)$ is the raw velocity measured at hot-wire number i and $T_{rms} = 1 \text{ s}$ is the time interval used in order to compute the moving average of the velocity signal.

The goal of the control is to increase mixing. In this Chapter, we chose the width of the shear layer profile as a quantity to maximize. This width is approximated by the ratio of the integral of the average fluctuation energy profile over the maximum of this profile. Thus, the cost function reads:

$$J = 1/W \quad \text{with } W = \frac{\sum_{i=1}^{24} \langle s_i^2(t) \rangle_T}{\max_{i \in [1,24]} (\langle s_i^2(t) \rangle_T)}, \quad (6.9)$$

where $T = 10 \text{ s}$ is the evaluation time for a control law. This time corresponds to approximately 950 Kelvin-Helmoltz periods of the initial shear layer.

The 96 actuators are operated simultaneously as a single input, although they could have been actuated independently. Studies with periodic forcing of different actuator configurations have indicated that this synchronous operation is best for enhancing the mixing layer width. The control law should only return a binary value (0 or 1) as the valves commanding the jets can only be operated in an on/off mode. The maximum jet velocity is kept constant by the pressure tank. Details on the implementation of the control law by MLC are given in the next section.

The parameters used for the operation of MLC in this experiment are detailed in Tab. 6.5.

The output of the expression trees can be an arbitrary function of the sensors $s_i(t)$, which could take any value in \mathbb{R} . In our case, the valves commanding the actuation jets can only be opened or closed. Hence, the output of the trees is compared to zero, and thus $b(t) = 0$ if the output of the tree is negative, and $b(t) = 1$ otherwise. This operation is achieved inside the Fortran part of the RT loop, after interpretation of the

Table 6.5 MLC parameters used for the control of the TUCOROM mixing layer.

Parameter	Value
N_i	1000 (first generation) 100 (other generations)
P_r	0.1
P_m	0.25
P_c	0.65
N_p	7
N_e	1
Node functions	+, −, ×, /, sin, cos, exp, log, tanh

expression trees, and before sending back the actuation command to the Labview[®] framework (see Sec. 6.5.4).

Evaluating an individual costs 3 s transient time and $T = 10$ s evaluation time. Thus the first generation with 1000 individuals requires around 3.5 h and subsequent generations require around 20 min each. This represents around 8 – 10 hours for a typical run with 10-15 generations. The reader is strongly advised to implement health monitoring of the experiment and simple automatic reset procedures.

6.3.3 Results

Though the experiment was run until the 15th generation, the best control law was already obtained after the 8th generation. The effect of the MLC-controlled flow is shown in Fig. 6.8 — together with the unactuated flow and the best periodically forced flow. The periodic-forcing benchmark has been obtained through an extensive parametric study of harmonic forcing with many frequencies ranging from 1 to 500 Hz and various duty cycles. A periodic forcing at frequency of 9 Hz and 50% duty cycle was found to yield the largest increase in mixing-layer width W of a 54%. In terms of performance, the MLC controlled flows exhibits a 70% increase of the shear-layer width, which outperforms the periodic-forcing benchmark by 16% with respect to this width or 30% with respect to increase of the width by periodic forcing. Looking back at Fig. 6.8, the similarity between the periodic forcing and the

Advanced material 6.2 Varying the population size to minimize the time of the runs.

The MLC of this experiment follows the experimental changes of advanced material 6.1. In addition, the first generation contains 1000 individuals while the following ones are restricted to 100 individuals. The large first generation (with removal of uninteresting individuals) ensures that the search space contains effective control laws. The significantly reduced size of further generations reduces the time investment in an inefficient exploration.




Fig. 6.8 Pseudo-visualizations of the TUCOROM experimental mixing layer demonstrator [206] for three cases: natural baseline (top, width $W = 100\%$), the best periodic-forcing benchmark (middle, width $W = 154\%$), and MLC closed-loop control (bottom, width $W = 170\%$). The velocity fluctuations recorded by 24 hot-wires probes (see Fig. 6.7) are shown as contour-plot over the time t (abscissa) and the sensor position y (ordinate). The black stripes above the controlled cases indicate when the actuator is active (taking into account the convective time). The average actuation frequency achieved by the MLC control is comparable to the periodic-forcing benchmark.

effect of the MLC law is evident. A spectral analysis of the MLC control command shows that MLC also operates around 9 Hz. Both actuations amplify the large-scale eddies at the downstream measurement position. The difference between both optimal periodic forcing and MLC is fourfold:

1. MLC is a sensor-based closed-loop control and does not explicitly depend on time. In other words, MLC schedules actuation based on flow events and not on a predefined ‘clockwork’.
2. The shear-layer width associated with MLC actuation is significantly larger.
3. The effective duty cycle exhibited by MLC is much smaller, leading to a 48% decrease in the actuation cost. Intriguingly, this benefit arises although actuation was not penalized in the cost function (6.9).
4. The actuation effect of MLC dramatically outperforms periodic forcing for changing operating conditions, such as changing the maximum stream velocity. The situation is comparable to the PMMH experiment in Sec. 6.1 or drag reduction of the D-shaped body with phasor control [209]

Feeding back sensor signals is particularly important to optimize the control law. The periodic forcing acts like a clockwork on the noise amplifier and performs well as long as no perturbations are modifying the underlying base flow. On the other hand, MLC triggers the actuation based on what is sensed. MLC adapts with a time delay of the order of $2 \times x / (U_1 + U_2)$ which is comparable to the period of maximum excitation. Thus, MLC realizes a phasor-type control as in [209] which is more ‘in phase’ with the flow physics and more robust to changing operating conditions.

6.4 Alternative model-based control approaches

To motivate the use of machine learning control using genetic programming, we consider the inability of linear system identification to capture the strongly nonlinear dynamics in the TUCOROM mixing layer experiment from Sec. 6.3. In particular, we attempt to identify a linear input–output model from actuation to hot-wire sensors using the ERA/OKID algorithms described in Sec. 3.5.1.

As mentioned earlier, it is often difficult to obtain clean impulse response experiments, especially in turbulent fluid systems with stochastic fluctuations. Instead, a pseudo-random actuation sequence is constructed by repeatedly turning the jets on or off for a random interval of time. The hold time is sampled from a discrete Poisson distribution with parameter $\lambda = 4$. This distribution is then scaled so that the mean hold time is 0.05 seconds. Using pseudo-random excitation sequences for system identification of fluid systems is inspired by [158]. The actuation sequence is shown in Fig. 6.9 (top).




Fig. 6.9 (a) Pseudo-random actuation sequence for system identification and (b) velocity measurements from hot-wire rake. The actuation sequence consists of Poisson blowing at the leading edge of the splitter plate with average hold duration of 0.05 s. Nineteen hot-wire probes are used at a downstream location of $x = 500$ mm.

A particular pseudo-random blowing sequence is used for 300 identical experimental runs, and the sensor measurements from a rake of nineteen hot wires at a downstream location of $x = 500$ mm are collected and phase averaged. The resulting phase-averaged velocity measurements are shown in Fig. 6.9 (bottom). Immediately, coherent structures may be observed as correlations in the hot-wire signals; however, phase-averaging removes important stochastic fluctuations that account for a significant amount of the kinetic energy. Therefore, even perfect model reconstruction of the phase-averaged velocity measurements would only account for a limited aspect of the full nonlinear flow.




Fig. 6.10 The ERA/OKID models for the input–output response of each hot wire may be combined to produce a frequency response plot. The color code indicates the corresponding hot wire and is adopted from Fig. 6.9. The phase information (b) strongly suggests a time delay, while the magnitude plot (a) indicates weak flow resonances.

The phase-averaged measurements are then used in the ERA/OKID algorithm to obtain input–output models for the various hot-wire signals. A frequency response of the various models for the nineteen hot-wire signals is shown in Fig. 6.10. The magnitude plot shows moderate flow resonance at certain frequencies, and the phase plot captures the time delay between actuation and measurements. However, when analyzing the model reconstruction of a particular hot wire in Fig. 6.11, it is clear that the model error is on the same order of the signal strength. This indicates that the phase-averaged measurements still contain strong, reproducible nonlinear flow effects that cannot be captured by the linear models.

To summarize, linear model identification fails to capture two features of the mixing layer flow field: stochastic fluctuations and coherent nonlinear phenomena. Taken together, these effects account for a significant portion of the kinetic energy, and are likely important for flow control efforts. However, it may still be interesting to test the performance of robust model-based control, considering that many nonlinearities may be considered as model uncertainty or disturbances.




Fig. 6.11 The true hot-wire signal u is compared with the model-reconstructed signal u_r for the 9th hot wire near the middle of the rake. Note that the error Δu is nearly as large as the original signal.

6.5 Implementation of MLC in experiments

This section details the software and hardware implementation of the experiments described in the previous sections. Following Chapter 2, two closed loops need to be set up: a fast inner real-time (RT) control loop (Sec. 6.5.1) and a ‘slow’ outer MLC learning loop. Section 6.5.1 describes the real-time control loop which extracts information from the experiment through sensors, computes the actuation command and passes it to the actuator. Details about the hardware and software implementation of MLC for the PMMH, PRISME and TUCOROM experiments are provided in Secs. 6.5.2, 6.5.3 and 6.5.4, respectively.

6.5.1 Real-time control loop - from sensors to actuators

Any experimental RT control loop consists of four elements: the plant, sensors, actuators and a computational unit. The sensors get information about the plant. The computational unit processes these signals and sends a control command to the actuators which, in turn, will act on the plant. All of these elements must be able to operate in a frequency range of the actuation mechanism at work. Stabilizing an unstable oscillator, for instance, typically requires a sampling rate which is significantly larger than the eigenfrequency. In fact, the Nyquist theorem would impose the need for a sampling frequency which is at least twice that frequency and a rule of thumb is to use a sampling frequency of at least ten times the eigenfrequency.

The RT loop constitutes a classical feedback experiment for a given controller. The slow outer MLC loop provides new controllers for this experiment after each cost function evaluation. This implies that the computational unit needs to be programmable, either by allowing it to read and interpret the individuals provided by the learning loop (at every time step), or by allowing the learning loop to change its content. Two approaches have been used to achieve this goal:

- The outer loop is able to change (and compile) the controller code on the RT-loop computation unit at least once per generation (see Sec. 6.5.3 as an example).
- The use of an on-board control law parser: the controller implemented in the computational unit is re-loaded every sampling cycle and could thus change even during one evaluation period (see Sec. 6.5.4).

As the control law provided by the MLC learning loop is arbitrary, the RT loop has to be protected against control commands that provide values outside the safe operation range of the experiment. This means that the controller — implemented as a parser or as an externally modified part — should be encapsulated in a framework that can protect the actuators from operating out of safety margins.

The last aspect to be taken into account is that the experiment may evaluate hundreds or thousands of control laws and may hence take more time than a standard control experiment where only a single control law is tested or optimized. The time scales for MLC experiments may range from hours in wind tunnels to days in water tunnels. Hence, all components of the experiment should be able to operate in an unsupervised manner at least to evaluate a few generations. This might imply the addition of other regulation loops for temperature, calibration updates, tank refill, to name only few examples. It is also advisable to design some telltales in order to detect any unusual behavior of the experiment. Feedback loops have a reputation of pointing at weak spots of experiments and this is amplified by the exploratory nature of the MLC algorithm.

6.5.2 MLC implementation in the PMMH flow over a backward-facing step

In the water tunnel, the characteristic frequency is about 1 Hz. With such a low characteristic frequency, the controller can be implemented on any modern computer (Fig. 6.12). The real-time loop, including the RT PIV part, is implemented in Labview[®]. Thus, the Labview[®] project contains the real-time PIV acquisition, the command signal to the actuator and the controller. The outer MLC loop employs OpenMLC introduced in Sec. 2.3.

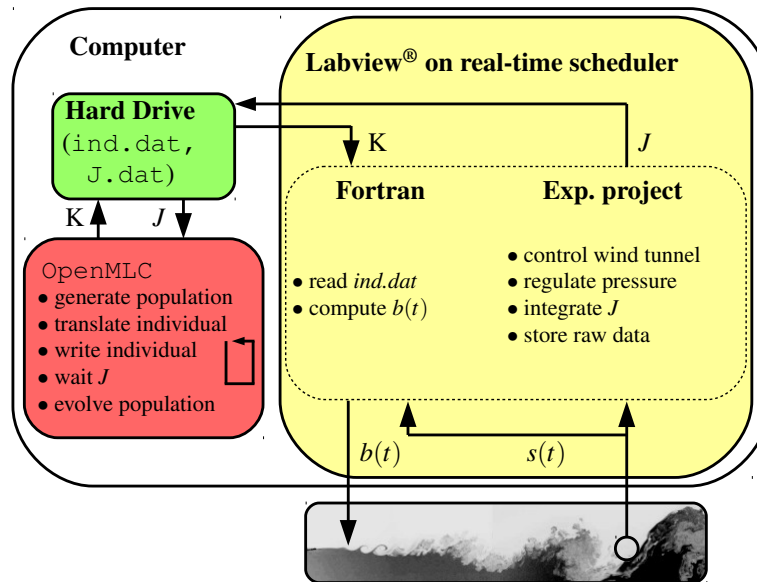


Fig. 6.12 MLC implementation architecture in PMMH backward-facing step flow.

The control law is first translated from LISP to a classical symbolic expression by OpenMLC then transmitted through file exchange on the computer hard-drive to the Labview[®] project. Each time OpenMLC evaluates an individual, it writes a file (e.g. *ind.dat*) containing the symbolic expression describing the control law. The appearance of the file is detected in one of the RT-loop iterations. The file is read, the string variable containing the expression is then updated and the file is deleted. A parser function is then used to transform the sensor value to a control command inside the controller block.

That control law is used during the evaluation period and at the same time the cost function is integrated. The Labview[®] program then writes the cost function value of the individual in a file (e.g. *J.dat*) in the exchange directory of the computer. OpenMLC detects the presence of the file, reads and deletes it, which tells Labview[®]

that the transmission is complete. The next individual can then be evaluated in the same way.

The exchange by files may sound like an unnecessary and time-consuming way to proceed. Yet it is relatively simple, as only the read/write protocol needs to be implemented. The file exchange requires that a parser can be implemented on the RT loop.

6.5.3 MLC implementation in the LML and PRISME experiments

In the LML and PRISME experiments, no dedicated, high-performance and highly expensive computer with real-time capacities was available. The RT loop has instead been implemented in an Arduino micro-controller (Fig. 6.13). In contrast to the water-tunnel experiment in Sec. 6.5.2, a parser on the Arduino would have led to a significant performance loss. The water-tunnel experiment at PMMH had a large characteristic time scale of 1 seconds and the communication between RT and MLC loop happened over file exchange. The time scales of wind-tunnel experiments are one to two orders of magnitude faster and we chose a more efficient communication strategy. `OpenMLC` has been set up so that (1) the evaluation function generates code for the Arduino with all control laws for one generation, (2) compiles it, and (3) burns it on the micro-controller. By abandoning the possibility of changing the individual on the fly, the computational load achieved on the board has been reduced to a bare minimum. This speed-up enables a control update frequency of 1 kHz without even having to implement low-level programming on the Arduino.

The burned code contains all individuals of one generation and sends back the cost function values. The outer loop is controlled by by Matlab[®] code which generates the codes for the individuals and retrieves the cost function values in addition to reading any other information which may be fed back by the Arduino. This code can interrupt the experiment, realize a calibration, and re-burn and re-start the evaluation if needed.

6.5.4 MLC implementation in the TUCOROM experiment

The controller is implemented on a high performance PC with 1 TeraFlop computing power (Fig. 6.14). The real-time loop is designed in Labview[®] and compiled to run at the desired frequency enabled by a real-time scheduler. The reading of the sensors, the signal processing, and the calculation of the control command is achieved by a Fortran function whose output is then transferred inside the Labview[®] designed project. The Labview[®] project encompasses all other commands needed to operate the wind tunnel, from the turbine operation, pressure regulation for the jets, calibration, temperature monitoring, and so on.

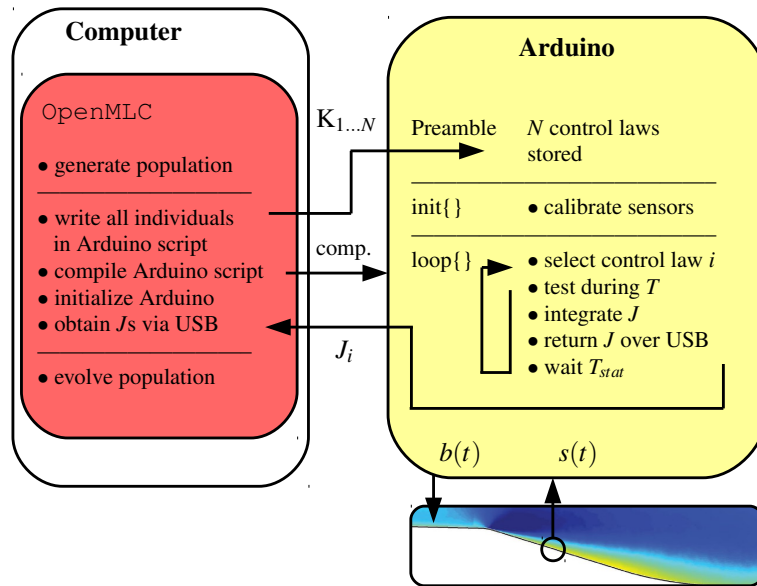


Fig. 6.13 MLC implementation architecture in LML and PRISME boundary-layer experiments.

The interaction between `OpenMLC` and the RT loop is achieved through file exchange as in the PMMH water-tunnel experiment (Sec. 6.5.2). This has the advantage of changing the control law on the fly. Here, the loop described by the Labview® project is in charge of monitoring the experiment and act accordingly in case of an unexpected event, such as failing actuators or sensors.

Advanced material 6.3 LISP interpreter pseudo-code.

All LISP parsers operate in the same way and are rather easy to build, through a recursive function: function def:

```
translated_string=read_my_LISP(LISP_string):
```

```
(1) Detect outer parenthesis
```

```
if no parenthesis: translated_string = LISP_string
```

```
(2) Detect spaces
```

```
operator=LISP_string(par1:space1)
```

```
argument#n=LISP_string(space#n:space#n+1)
```

```
(3) Translate arguments
```

```
tr_args#n=read_my_LISP(argument#n)
```

```
(4) Replace strings
```

```
if op is a function
```

```
translated_string=operator(argument#1,...)
```

```
if op is an operation
```

```
translated_string=argument#1 operator argument#2
```

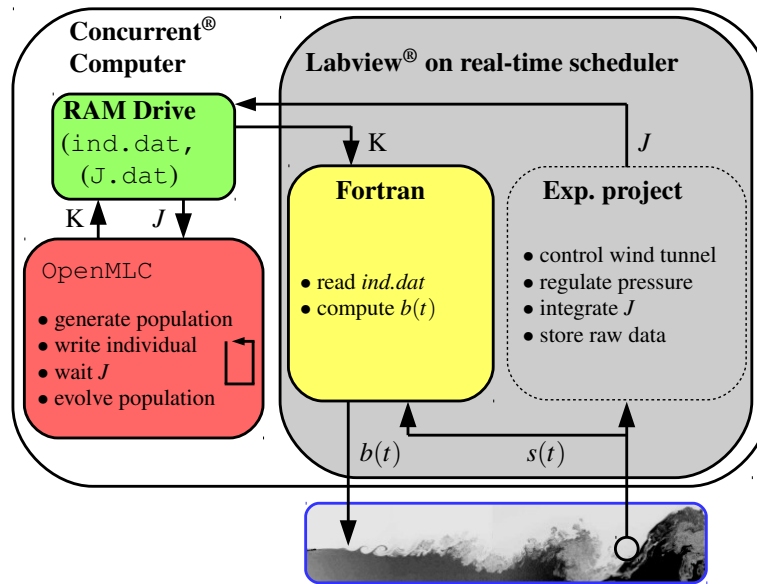


Fig. 6.14 MLC implementation architecture in TUCOROM mixing layer.

6.6 Suggested reading

- (1) **Modern developments in flow control**, by M. Gad-el-Hak, *Applied Mechanics Reviews*, 1996 [108].

This review provides an early perspective on flow control with a number of future directions that have since been developed.

- (2) **Control of turbulence**, by J. L. Lumley and P. N. Blossey, *Annual Review of Fluid Mechanics*, 1998 [179].

This review is among the earliest overviews summarizing efforts on modeling for turbulence control.

- (3) **Feedback control of combustion oscillations**, by A. P. Dowling and A. S. Morgans, *Annual Review of Fluid Mechanics*, 2005 [85].

This review considers combustion control using linear control theory.

- (4) **Dynamics and control of high-Reynolds number flows over open cavities**, by C. W. Rowley and D. R. Williams, *Annual Review of Fluid Mechanics*, 2006 [231].

This review describes the control of flow over open cavities, which provides an illuminating success story of flow control design.

- (5) **A linear systems approach to flow control**, by J. Kim and T. R. Bewley, *Annual Review of Fluid Mechanics*, 2007 [161].

This review provides a clear summary of modern techniques applying linear control theory to fluid flow control.

- (6) **Control of flow over a bluff body**, by H. Choi, W.-P. Jeon, and J. Kim, *Annual Review of Fluid Mechanics*, 2008 [65].
This review considers the problem of controlling bluff body flows from physical and mathematical perspectives.
- (7) **Optimal and robust control of fluid flows: Some theoretical and computational aspects**, by T. T. Medjo, R. Temam, and M. Ziane, *Applied Mechanics Reviews*, 2008 [186].
This review provides an overview of mathematical and numerical considerations in modern flow control.
- (8) **Input-output analysis and control design applied to a linear model of spatially developing flows**, by S. Bagheri, J. Hoepffner, P. J. Schmid, and D. S. Henningson, *Applied Mechanics Reviews*, 2009 [13].
This review explores linear model-based control of fluid systems with the Ginzburg-Landau equation as an illuminating example.
- (9) **Actuators for active flow control**, by L. Cattafesta, *Annual Review of Fluid Mechanics*, 2011 [54].
This review considers one of the most important factors in experimental flow control: the actuators.
- (10) **Adaptive and model-based control theory applied to convectively unstable flows**, by N. Fabbiane, O. Semeraro, S. Bagheri, and D. S. Henningson, *Applied Mechanics Reviews*, 2014 [101].
This review provides an overview of research on active control for transition delay with example code for the linearized Kuramoto–Sivashinsky equation.
- (11) **Analysis of fluid systems: stability, receptivity, sensitivity**, by P. J. Schmid and L. Brandt, *Applied Mechanics Reviews*, 2014 [238].
This review investigates flow analysis in terms of cost optimization.
- (12) **Closed-loop turbulence control: Progress and challenges**, by S. L. Brunton and B. R. Noack, *Applied Mechanics Reviews*, 2015 [43].
This review explores the state-of-the-art of closed-loop turbulence control and provides an emphasis on modern methods from machine learning.

6.7 Interview with Professor David Williams

David Williams is Professor of Mechanical & Aerospace Engineering at the Illinois Institute of Technology, IL, USA. Since 1999 he has been the Director of the Fluid Dynamics Research Center at IIT. In addition, Professor Williams is an active member of the American Institute of Aeronautics and Astronautics, and is currently serving his third term on the fluid dynamics technical committee.



Professor Williams is a leading expert on experimental closed-loop flow control, with numerous experimental success stories on challenging flow configurations, such as the reduction of acoustic tones in aircraft cavities or separation control over an airfoil in flight maneuvers. He is a highly collaborative researcher and draws researchers from around the globe to work with him in his advanced wind tunnel facility.

Professor Williams has received numerous national and international honors for his research accomplishments and excellence in teaching. He is Fellow of the American Physical Society and was awarded with the prestigious Alexander von Humboldt Fellowship and the Honeywell Advanced Technology Achievement Award. He is also a licensed commercial pilot with more than 1000 hours of flying time.

Authors: You are a leader in experimental flow control, with many successful demonstrations of closed-loop control for aerodynamic applications. It would be an understatement to say that there are issues that arise in experimental flow control that are not usually present in simulations and theory. Can you discuss key enablers in experimental flow control that have been developed in the last two decades? Which challenges do you see as guiding foci of relevant future research?

Prof. Williams: The “key enablers” for experimental active flow control (AFC) have come primarily from the adaptation of modern methods of closed-loop feedback control algorithms into experiments, and the rapid development of sensors and digital signal processing (DSP) hardware, e.g., dSPACE and Arduino DSP’s. This combination enables AFC to become much more than an exercise in actuator development and open-loop control. It is now possible to adapt AFC to changing flight/flow conditions, and to interact directly with internal flow instabilities, which leads to far more robust and effective controllers than can be achieved with open loop control or by changing a base flow state. With closed-loop active controllers the actuator power requirements can be minimized, which is not possible with open-loop methods.

The focus in experimental fluid dynamics is shifting from open-loop demonstrations of actuator performance at steady flow conditions to closed-loop control

during unsteady flow conditions. Open-loop actuator demonstrations compete with passive flow control devices. For example, separation control on wings will likely continue to be done with passive vortex generators, because the active devices are more expensive and less reliable.

Future research in experimental AFC will continue to emphasize the development of low-dimensional models that accurately represent the flow's response to external disturbances and the flow's response to actuator input, i.e., disturbance models and plant models. These can be used to design practical and effective feed-forward and feedback controllers that run in real-time on the digital-signal processing hardware. Examples of active flow control systems that have successfully employed these techniques include the reduction of unsteady loads in gusting flows on wings and road vehicles, thermo-acoustic combustion instability suppression, control over the reattachment length in separated flows over steps, reduced drag on bluff bodies, reduce acoustic tone amplitudes in cavity tones, control the rotational motion of slender bodies at high angles of attack, and reduce lift hysteresis on pitching wings.

Another important enabler in experimental flow control has been in sensor development. MEMs-based sensor technology is reducing the cost and size, and increasing the performance of pressure sensors, motion sensors with IMU's, gyro's, and accelerometers. This enables distributed sensing capability, such as multiple pressure sensors in a diffuser or over the surface of a wing, for better flow state estimation in practical applications. Shear stress sensors continue to be developed and applied with varying degrees of success in the laboratory, but have not developed to practical application levels. Full-field sensing techniques, such as real-time PIV, enable us to explore new methods of integrating experiments with numerical simulations. Data assimilation methods combine experiments and simulations as a way of correcting simulations.

Actuator technology is sufficiently mature for solving a large number of flow control problems, but one finds that our understanding of how the flow will respond to an arbitrary actuator input is somewhat limited. System Identification methods have proven to be useful for developing black-box models that are effective in modeling flows, but they are not as reliable as models developed from first principles.

Another challenge for all flow control investigators is to recognize that fluid dynamic time scales introduce time delays into the system that we are trying to control. These delays limit the achievable bandwidth of the overall system control, so actuator bandwidth is not a major issue for most flow control applications.

Authors: In your experience, what are some of the biggest gains in experimental flow control that came about from simulations or models? In contrast, what do you see as the key limitations of models in capturing real-world fluid phenomena?

Prof. Williams: Most of the significant advances in flow control were the result of collaborations between experiments, theory and simulations, so it is difficult

to attribute advances in experimental flow control as being the direct result of simulations or models. For example, in what I consider to be the first modern flow control experiment, Prandtl used suction to suppress separation from one side of a cylinder, which provided evidence for the role of the boundary layer in flow separation and became the principal motivator for the development of boundary layer flow control. In this case theory preceded the experiment, but experimental observations led to the theory.

In my experience the key limitations of simulations in capturing real-world are matching Reynolds number, computing an accurate actuator input to the flow, and matching the dimensionality of the experiment. Two-dimensional simulations of the flow's response to actuation often produce inaccurate results, because the real-world actuators typically introduce three-dimensional disturbances into a flow. It is difficult for actuators used in experiments to produce two-dimensional disturbances.

In the case of closed-loop active flow control, models of the flow field response to external disturbances and models for the flow field response to actuator input are very useful for achieving effective control.

Authors: Can you comment on the role and benefits of experimental flow control in the future as computer simulations become more powerful?

Prof. Williams: The collaboration between experiments and simulations will become even stronger as simulations become more powerful. Fortunately the engineering community has outgrown the notion from the 1990's that simulations will make experiments obsolete. The emerging area of research known as "data assimilation" seeks to capitalize on integration of large-scale, real-time, experimental data sets with numerical simulations. Although the techniques being developed are not specifically focused on flow control applications, it seems likely that the approach will be useful for the entire flow control community. For example, imagine real-time experimental data being used to "correct" a full-scale DNS simulation of a particular flow. The DNS can provide the full-state feedback information to a controller that would not be available from the experimental data. Simulations done correctly provide more detailed information about a flow than can be obtained by experiment. Experiments, on the other hand, can explore a wider parameter space than simulations. Simulations and models often identify instabilities that are susceptible to control. In principle, simulations (adjoint methods) can identify spatial locations where actuator placement can be most effective.

Authors: In the coming decades, where do you see the big future gains in our everyday lives and in industry due to flow control?

Prof. Williams: The Holy Grail of flow control would be the ability to control all scales of turbulence in shear flows, such as, in turbulent boundary layers, jets, wakes and mixing layers. Imagine being able to 'flip a switch' and laminarize a turbulent boundary layer. Fuel savings for the aerospace industry would be enor-

mous. However, I don't envision that capability in the near-term future, primarily because we don't have sufficient understanding of the turbulence production mechanisms, or ways in which actuators can be used to interact with those mechanisms. However, pursuing the goal of turbulence control drives the development of new technology in numerous areas. New approaches to flow modeling (e.g., resolvent modes), advances in sensor and actuator technology, and novel control algorithms will lead to improved understanding of the fundamental phenomena governing turbulence.

Linear thinking about the current trends in flow control suggests that AFC will result in more efficient energy recovery via improved aerodynamics on aircraft and wind turbines. Reduced drag and improved gust tolerance on commercial aircraft can be expected. I believe that it is possible for conventional flight control surfaces to be replaced by AFC actuators, and these aircraft will fly without vertical stabilizers and will fly with substantially larger aerodynamic efficiency, i.e., L/D ratio. AFC will enable tailored combustion processes that will reduce pollutants from combustion processes.

The rapid development of low cost and ever more powerful DSP's will enable more intelligent control of flows and systems. The application of distributed sensing will improve state estimation.

Will there be a break through in our understanding of turbulent flow mechanisms that enables laminarization of turbulent shear flows? Linear extrapolation says 'no', but fortunately, nature does not follow linear extrapolations.

Authors: Do you envision these being facilitated by better hardware, increasing model fidelity, access to more complete data, or some combination thereof?

Prof. Williams: Improvements in a combination of all of the above may lead to the better understanding that is required to control turbulent shear flows.

There is a difference between what we can do with flow control and what is of interest to industry. For example, we know how to delay airfoil stall with various types of leading edge actuation, such as, pulsed-pulsed blowing jets, but industry continues to use mechanical vortex generators for stall control, even though there is a continuous drag penalty. From industry's perspective the complexity and reliability issues of an AFC stall control system outweigh the benefits of reduced drag. It is obvious that AFC techniques must buy their way onto a system by providing greater benefits than the cost of added complexity. Sometimes those benefits are unexpected, such as, the use of pneumatic AFC for flight control of an aircraft. The primary function of the pneumatic actuator is to provide roll and yaw control, which takes bleed air from the engine. At the same time a secondary effect occurs, where the AFC increases the efficiency of the flight vehicle, so that the range and endurance are better with AFC than without. In this case, the overall system benefits from the application of AFC.

Authors: We thank you for sharing your insights in this interview!