

Gas flow loop FIC-101 — Operating principles

Contents

1. Orifice flow measurement	2
1.1 Why flow is not measured directly	2
1.2 The signal path	2
1.3 Why the square root, and what it costs	2
1.4 Why P/T/Z compensation	3
1.5 The delays in the chain, and why they are not a defect	3
1.6 Settings linked to the measurement	3
1.7 A note for the trainee — what the model simplifies	4
2. The final element: control valve FV-101	4
2.1 Why a control valve is not a tap	4
2.2 Kv, and why gas is not sized like a liquid	4
2.3 Choking: past a certain drop, pushing does not help	5
2.4 The inherent characteristic: three different valves in the same body	5
2.5 Stroke time and stiction	5
2.6 A note for the trainee — what the model simplifies	5
3. The PID controller	6
3.1 Why a controller is needed	6
3.2 Reverse or direct action	6
3.3 The three actions	6
3.4 Derivative is taken on the measurement, not on the error	7
3.5 Automatic, manual and bumpless transfer	7
3.6 Anti-windup	7
3.7 A note for the trainee — what the model simplifies	7
4. Identification and tuning	7
4.1 Why one does not tune by trial	7
4.2 The reaction curve	8
4.3 Two tunings from the same curve	8
4.4 Loop gain	8
4.5 How a tuning is judged	8
4.6 A note for the trainee — what the model simplifies	9
5. The protections: alarms, trip and relief valve	9
5.1 Three layers, in order	9
5.2 Alarms: a condition, not a message	9
5.3 UZV-101, the trip	10
5.4 PSV-101, the mechanical protection	10
5.5 All the settings	10
5.6 A note for the trainee — what the model simplifies	11
6. The gas, the pressures and the disturbances	11
6.1 Why a gas is not handled like a liquid	11
6.2 The compressibility factor Z	11

6.3 Network pressure and user back pressure	12
6.4 Temperature and the Joule-Thomson effect	12
6.5 A note for the trainee — what the model simplifies	12

A document for those using the simulator in training. It does not describe the web page (buttons, cards, sliders): it explains why the process behaves the way it does and on which control logic the model is built, so that whoever practises understands the physical phenomenon and the control strategy, not just the sequence of commands to press.

One chapter per subsystem. We start from the measurement, because a control loop can never be better than the measurement it closes on: everything that comes after — valve, controller, tuning — works on the number the transmitter delivers, right or wrong as it may be.

1. Orifice flow measurement

Instruments: FE-101, FT-101, PT-101, TT-101

1.1 Why flow is not measured directly

There is no instrument that counts the molecules going past. A differential pressure meter does something else: it deliberately creates a known obstacle — the orifice plate FE-101, a drilled disc fitted between two flanges — and measures the pressure difference the fluid generates as it goes through. That difference is linked to the flow by a known relation, and the flow is worked backwards from it.

It is an indirect method, and this explains almost everything the trainee sees on the panel: the square root, the limited range, the compensation, the sensitivity to density. None of these is a whim of the instrument: they are consequences of the fact that a pressure is measured and a flow is computed.

1.2 The signal path

The *Measurement chain* box on the mimic shows the complete path, in the same order the flow computer executes it:

ΔP at the taps → square root extraction → flow at line conditions → P/T/Z compensation → normalised flow in Sm^3/h

Every stage is computed at every scan of the simulator, not set by hand: if the network pressure changes, if the gas warms up or if the valve moves, the effect propagates along the whole chain with its own response times.

The ΔP the transmitter reads is not the permanent pressure loss of the orifice. The taps are flanged, upstream and downstream of the disc, and read the pressure drop at the point where the fluid vein is most contracted. The permanent loss — the one the plant actually pays — is a fraction of that ΔP , about 57 % with the diameter ratio used here. This is why an orifice “costs” pressure: the measurement is obtained by sacrificing part of the available head.

1.3 Why the square root, and what it costs

ΔP is not proportional to flow, but to its **square**: doubling the flow makes the ΔP four times larger. The flow computer must therefore take the square root of the signal to get back to a quantity proportional to the flow.

The practical consequence is that measurement accuracy **is not uniform across the span**. At full flow the ΔP is large and an error of a few millibar counts for little; at a quarter of the flow the ΔP is sixteen times smaller, and the same error in millibar weighs far more as a percentage of the reading. This is why the useful range of an orifice is taken as about 3:1, not 10:1 as for a mass meter.

Hence also the **low flow cut-off**: below a minimum fraction of full scale the computer forces the reading to zero instead of amplifying noise. Without that cut-off, with the valve nearly shut the panel would show a dancing flow generated by transmitter noise alone.

1.4 Why P/T/Z compensation

The relation between ΔP and flow contains the **density** of the gas. The flow computer cannot ignore it, because the density of a gas is not a constant: it depends on pressure, temperature and the compressibility factor Z .

With compensation **in service**, the computer reads line pressure and temperature from transmitters PT-101 and TT-101, computes the real density and obtains the true flow. With compensation **bypassed** it uses a fixed density, the design one, and is wrong every time the real conditions move away from it:

Real condition	Real density	Reading without compensation
Pressure higher than design	higher	low (reads less than the truth)
Pressure lower than design	lower	high
Gas warmer than design	lower	high (reads more than the truth)
Gas colder than design	higher	low

The error is not small: taking the network from 8 to 10.5 bar g with compensation bypassed makes the reading drop by about 12 % against the true flow, and the panel writes how far off it is under the compensation control.

It is the most treacherous fault of all, because **the loop keeps working beautifully**: the controller brings the measurement onto set-point and holds it there, steady. It is simply controlling a wrong number, and the user is getting a different flow. The alarm list marks the bypass separately for exactly this reason: a protection or a correction that has been switched out is invisible to anyone looking only at the values.

1.5 The delays in the chain, and why they are not a defect

Between the flow actually going down the pipe and the number shown on the panel there are three delays, all real and all present in the model:

- the **transport delay**, the time the pressure signal takes to reach the transmitter through the impulse lines;
- the **transmitter damping**, a deliberate filter the instrument technician sets to make the reading legible rather than jumpy;
- the **line capacity**, that is the volume of gas that has to empty or fill before the flow really changes downstream.

They are the reason a loop can be destabilised. A process without delay would tolerate any gain: K could be raised as far as one liked and control would only get faster. It is the delay that sets a limit — the correction arrives when the situation has already changed — and that is why tuning is a craft and not a trivial calculation.

On this simulator those three parameters are **fixed and declared**, not adjustable: noise 0.6 % of span, damping 2 s, transport delay 0.8 s. On a real plant the trainee suffers those numbers, does not move them.

1.6 Settings linked to the measurement

Tag	Meaning	Setting	Effect
FAL-101	Flow low	400 Sm ³ /h	alarm
FAH-101	Flow high	4600 Sm ³ /h	alarm
FDEV-101	Persistent deviation between set-point and measurement	8 % for 25 s	alarm, in automatic only
PDAH-101	Orifice ΔP above the transmitter range	full scale	alarm
FY-101	P/T/Z compensation bypassed	—	marked as a bypass

The deviation alarm is inhibited with the controller in manual: in manual it is the operator who decides the output, and a departure from set-point is not an anomaly but a choice. That is control room practice, not a shortcut of the simulator.

1.7 A note for the trainee — what the model simplifies

- The pressure and temperature measurements used by the compensation are taken as instantaneous and noise-free: in the model only the ΔP transmitter has noise, delay and damping. On a plant PT-101 and TT-101 have their own uncertainties too, which add to that of the ΔP .
- The compressibility factor Z is computed with a virial equation truncated at the second coefficient, valid below 20 bar for methane. A custody transfer computer would use fuller correlations.
- The “plugged impulse line” fault makes the reading fall progressively, as a partial blockage getting worse would. A dry, immediate blockage of one single tap, which would give an instant jump, is not represented.

2. The final element: control valve FV-101

2.1 Why a control valve is not a tap

An isolating valve has two useful positions, open and shut; a control valve works **continuously** in every intermediate position, and its merit is not how much goes through when fully open, but how **regularly** the flow follows the travel of the plug.

This difference explains why the line carries two valves with roles that are not interchangeable: UZV-101 is a shutdown valve and sits fully open or fully shut, FV-101 is a control valve and moves continuously. A shutdown valve used for control would work badly and wear out quickly; a control valve used for isolation does not guarantee tightness.

2.2 Kv, and why gas is not sized like a liquid

The capacity of a valve is expressed by its **Kv**: the water flow in m³/h that passes with one bar of pressure drop. For a liquid the formula is simple and the flow grows with the square root of ΔP .

For a gas things change, because the gas **expands** through the valve: the density at the outlet is lower than at the inlet and the formula must account for it with an expansion factor. The simulator uses the formulation of IEC 60534-2-1, the same one as the sizing sheets:

$$W = 31.6 \cdot K_v \cdot Y \cdot \sqrt{x \cdot p_1 \cdot \rho_1}$$

where W is the mass flow, x the ratio of pressure drop to upstream pressure, Y the expansion factor and ρ_1 the upstream density.

2.3 Choking: past a certain drop, pushing does not help

When the ratio x exceeds a critical value, the gas velocity in the restricted section reaches the speed of sound and **the flow stops increasing**, however much the downstream pressure is lowered. The valve is *choked*.

It is a behaviour that surprises anyone coming from liquids, and on the simulator it can be felt directly: lowering the user back pressure below a certain point, the flow stops responding. The controller, which still sees an error, opens the valve further and gets nothing — and that is one of the situations in which the output saturates and the integral starts to accumulate.

2.4 The inherent characteristic: three different valves in the same body

The **inherent characteristic** is the relation between plug travel and Kv. The simulator lets you change it, and that is no cosmetic detail: it changes the gain the controller sees.

Characteristic	Travel → Kv relation	Process gain from 35 % to 85 % opening
Equal percentage	exponential	0.58 → 3.67 %/% — grows sixfold
Linear	proportional	1.64 → 1.37 %/% — essentially constant
Quick opening	square root	1.25 → 0.71 %/% — falls with load

The consequence is the most important lesson in the whole simulator: **a tuning that is good at one operating point is not good everywhere**. With the equal percentage characteristic, a controller tuned at low flow becomes too aggressive at high flow, and the same value of K that held the loop steady makes it oscillate.

It has to be said that on this plant the pressure drop across the valve varies little as the flow changes: that is why here it is the **linear** characteristic that gives constant gain. On plants where the valve ΔP collapses as the flow grows — a pump against a long pipeline, for instance — it is the equal percentage that compensates that collapse and linearises the loop. The choice of characteristic depends on the plant, not on a universal rule.

2.5 Stroke time and stiction

The valve does not move instantly: full travel takes 6 s, and the controller has to live with that delay like the others.

Stiction is a separate fault and is worth understanding well, because it is the most frequent defect on real loops and the most misunderstood. The plug stays stuck until the command exceeds a stick band; once past it, it breaks away with a jump and **overshoots** the target. The controller sees the error reverse, corrects the other way, the valve sticks again, and so on: a **limit cycle** appears, a permanent oscillation of almost constant amplitude.

The teaching point is that this cycle **cannot be removed by tuning the PID**. It can be slowed down by lengthening the integral time, but it stays: the cause is mechanical, and the cure is the valve, not the controller. Recognising a stiction cycle and stopping blaming the tuning is one of the skills this simulator sets out to train.

2.6 A note for the trainee — what the model simplifies

- The stiction model has a fixed stick band and breakaway jump; on a real valve they vary with position, with differential pressure and with the state of the packing.

- The positioner is not represented as a loop of its own: the valve follows the command with a stroke time and a rate limit, without its internal position loop.
- The “valve jammed” fault freezes the position where it was: it does not represent a valve failing safe by opening or closing on loss of air.

3. The PID controller

FIC-101, ISA form

3.1 Why a controller is needed

Without a controller the operator sets a valve opening and the flow is whatever follows from it. As long as nothing changes it works; but the network pressure swings, the gas temperature changes, the user varies its own back pressure — and the flow drifts away, with nobody bringing it back.

A controller closes the loop: it measures, compares with the set-point, and moves the valve in the direction that reduces the gap. It is no longer the opening that is imposed: it is the flow.

3.2 Reverse or direct action

It is the first thing to set and the easiest to get wrong. On this loop, raising the output opens the valve and the flow **goes up**. If the measurement is above set-point, the output must therefore be **lowered**: **reverse** action.

Getting the action wrong does not leave a permanent error, which could be corrected: it makes the loop run away. The controller corrects in the wrong direction, the error grows, the correction grows with it, and the output ends up at one end of the scale. It is a defect recognisable in seconds, and it is worth trying on the simulator so as to recognise it in the field.

3.3 The three actions

Action	Parameter	What it does	What happens if overdone
Proportional	K	Corrects in proportion to the present error	Growing oscillation
Integral	T _i (s)	Accumulates the error over time and removes the permanent offset	Slow hunting, overshoot
Derivative	T _d (s)	Anticipates from how fast the measurement is moving	Amplifies noise, makes the valve dance

Proportional alone never closes the error: to hold the valve in a position other than its rest position, a proportional controller needs an error to command it there. After a disturbance the flow settles at a value near the set-point but not equal to it, and that offset stays until the integral action steps in. That is exactly scenario 4 of the training module.

The **integral time** is read as follows: it is the time the integral action takes to repeat on its own what the proportional action has already done. A short T_i means aggressive integral action; a long one, slow. Taken to maximum it amounts to switching it off.

3.4 Derivative is taken on the measurement, not on the error

In the simulator the derivative action acts on the **filtered measurement**, not on the error. The difference shows when the operator moves the set-point: the error changes as a step, and a derivative computed on it would give an instant kick to the output, with a jerk on the valve that serves no purpose. Taking the derivative on the measurement, which cannot change as a step, that kick disappears and the derivative action carries on doing its job on real transients.

That said, on a **flow** loop the derivative is rarely used: the measurement is inherently noisy and fast, and the derivative turns that noise into continuous plug movement, wearing it out. On flow and level one almost always works in PI.

3.5 Automatic, manual and bumpless transfer

In **manual** it is the operator who imposes the output; the controller, however, is not idle: it **tracks**, that is it continuously recomputes its own integral term so that its arithmetic adds up exactly to the output the operator has set. When switching to automatic, the controller starts from there and the valve does not move a millimetre: that is **bumpless transfer**.

Without tracking, switching to automatic would jump the output to whatever value the controller had in memory — a blow on the valve and a pointless transient. It is a function nobody notices on a plant until it is missing.

3.6 Anti-windup

When the output is already at full scale and the error persists — unreachable set-point, valve fully open, insufficient network pressure — the integral action keeps accumulating to no purpose: the valve cannot open any further. This accumulation is called **windup**.

The trouble comes on recovery. If the integral has grown to hundreds of percent, when the set-point becomes reachable again the output stays pinned at 100 % until that accumulation has unwound, and the flow overshoots the new set-point by a wide margin before coming back.

Anti-windup prevents the accumulation: as soon as the output saturates, the integral term is computed backwards so as to stay consistent with the real output. On the simulator the difference is measured: with anti-windup, recovery after saturation takes about 10 s, without it more than 100.

3.7 A note for the trainee — what the model simplifies

- The controller has neither dead band nor output rate limit: two functions common on industrial controllers, left out here so as not to add parameters.
 - The derivative filter is fixed at 1 s and is not exposed as a control.
 - Output saturation is fixed at 0-100 %: operator-settable output limits are not represented.
-

4. Identification and tuning

4.1 Why one does not tune by trial

A loop can be tuned by turning the parameters and watching what happens, and sometimes it works. But without knowing **what the process is like** one does not even know which way to go, and a sluggish loop and an unstable one need opposite corrections that by eye are easily confused.

The alternative is to measure the process before tuning. The measurement is done with a **step test**, as old as control itself and still the most used in the field.

4.2 The reaction curve

The controller is put in **manual**, everything is allowed to settle, a known step is applied to the output and the response of the measurement is recorded. Three numbers are read off the curve:

Symbol	Name	What it tells
K_p	Process gain	How much the measurement changes, in % of span, per % of output
τ	Time constant	How “slow” the process is to respond once it has started
θ	Dead time	How long passes before the measurement moves at all

The ratio θ/τ tells how difficult a loop is: the more dead time there is compared with the time constant, the less gain can be used. On this simulator the test returns values around $\tau \approx 2.0$ s and $\theta \approx 1.4$ s — a fast flow loop with a dead time that is not negligible.

4.3 Two tunings from the same curve

From the same three numbers different tunings can be derived, depending on what is wanted. The simulator proposes two, side by side:

Method	Gain K	Integral time T_i	Expected behaviour
Ziegler-Nichols, reaction curve	$0.9 \tau / (K_p \theta)$	3.33θ	Fast response, about 20 % overshoot
λ (IMC)	$\tau / (K_p (\lambda + \theta))$	τ	Slower, without overshoot

Ziegler-Nichols was born to answer a disturbance quickly and accepts overshooting the target; λ tuning was born never to overshoot and accepts taking longer. It is not that one is right and the other wrong: **it depends on what costs more** on the plant, oscillation or slowness.

The two buttons on the simulator are there to make that comparison, not to give the answer: the tuning found by hand should be set beside these two and judged on the same indices.

4.4 Loop gain

What matters for stability is not K alone, but the **loop gain**: the product of the controller gain and the process gain. If the process triples its own gain — and that is what an equal percentage valve does going from 35 % to 85 % opening, ch. 2.4 — the loop gain triples with it, and a loop that was stable starts hunting without anybody having touched the controller.

That is why a tuning has to be verified **at the operating point where the plant will actually sit**, and why on loops with large load swings valve characteristics are chosen specifically to keep that product constant.

4.5 How a tuning is judged

By eye it is judged badly. The simulator measures four indices, three continuous and one per transient:

Index	What it measures	What it tells
IAE over 60 s	Integral of the absolute error	How much error has accumulated: lower is better
Swing over 20 s	Difference between maximum and minimum of the measurement	Whether residual oscillation is present
Time within $\pm 2\%$	Fraction of time inside the band	How much the loop stays where it should
Overshoot	How far the measurement passes the set-point, in % of the step	Whether the tuning is too aggressive

Alongside these, for every transient the panel reports rise time 10-90 % and settling time within $\pm 2\%$. A tuning is judged by comparing these numbers between two runs, not by remembering what the curve looked like.

4.6 A note for the trainee — what the model simplifies

- The simulator step test uses a single step in one direction only. In the field it is normally done out and back, to reveal any valve hysteresis.
- Identification assumes a first order model with dead time. It is the model used in practice, but it is an approximation: the real loop has several time constants in series.
- The performance indices are computed on the compensated measurement, that is on what the controller sees: if compensation is bypassed, the indices say the loop is working well even when the true flow is wrong.

5. The protections: alarms, trip and relief valve

5.1 Three layers, in order

The protections of this line are arranged on three levels, acting one after the other as severity grows:

1. **The alarm** warns the operator and leaves time to act. It does nothing by itself.
2. **The trip** acts when the operator has not acted or could not: PAHH-102 closes the isolating valve UZV-101 and the line stops.
3. **The relief valve** PSV-101 is the last defence, purely mechanical: it discharges to the vent when everything else has not been enough.

The order is not accidental and is one of the things this simulator wants to drive home: if the relief valve lifts during an exercise, it means the two previous layers did not work or were not used. **That is why a PSV lift takes away the pass from any scenario**, even if all the other criteria are satisfied.

5.2 Alarms: a condition, not a message

The list follows ISA-18.2. Every alarm is a **condition evaluated at every scan**, not a message written once: hence the three states.

State	Meaning
ACT	The condition is present and the alarm has not been acknowledged

State	Meaning
ACK	The condition is still present but the operator has taken it on
RTN	The condition has cleared but the alarm has not yet been acknowledged

Two additional marks matter as much as the states:

- **First-out** marks the *first* cause that triggered the trip. When one trip drags others along in a cascade, the list fills with alarms and without this mark it is no longer possible to see where it started.
- **Bypass** marks the functions that have been switched out — P/T/Z compensation bypassed, anti-windup bypassed. An invisible bypass is a protection that is not there: anyone looking only at process values has no way of noticing.

5.3 UZV-101, the trip

UZV-101 is an isolating valve: fully open or fully shut, never halfway. It closes on operator command (emergency pushbutton) or automatically when the downstream pressure exceeds the PAHH-102 setting.

Reset **does not reopen** while the cause persists: pressing it with the pressure still high does nothing. That is the correct behaviour of a shutdown system and not a discourtesy of the interface — a reset that reopens on a cause still present is the best way to earn a second immediate trip.

5.4 PSV-101, the mechanical protection

PSV-101 protects the meter run from network pressure excursions, discharging to the safety vent. It does not go through the control system: it is a spring set on the bench and sealed.

Quantity	Value	Meaning
Set pressure	11.0 bar g	Pressure at which the valve starts to open
Blowdown	7 %	How far the pressure must fall before it reseats
Overpressure	within 10 %	How far the pressure may rise above the setting while relieving

Blowdown deserves an extra line: if the valve reseated exactly at the setting, the pressure would rise again as soon as it closed and it would reopen, in a chattering that destroys the seat within hours. Blowdown is the margin that prevents that chattering.

The PSV-101 setting appears in the simulator's settings table **but cannot be changed**, unlike all the others. It is not a limitation of the panel: it is a faithful representation of the fact that this is the only protection layer the operator cannot move. Seeing it fixed next to the movable ones is half the lesson.

5.5 All the settings

Tag	Variable	Alarm	Trip	Action
FAL / FAH-101	Compensated flow	400 / 4600 Sm ³ /h	—	alarm
FDEV-101	Deviation SP – PV	8 % for 25 s	—	alarm, in automatic only

Tag	Variable	Alarm	Trip	Action
PDAH-101	Orifice ΔP	full scale	—	alarm
PAL-100	Network pressure low	5.0 bar g	—	alarm
PAH / PAHH-102	Downstream pressure	6.0 bar g	7.0 bar g	UZV-101 closes
PSV-101	Meter run pressure	—	11.0 bar g	relief to vent, not adjustable
YS-101	Output saturated	over 30 s	—	alarm

All the settings in the table except the PSV one can be edited on line from the panel, and the alarm/trip pair realigns itself: moving the alarm past the trip, the trip moves to keep the margin.

5.6 A note for the trainee — what the model simplifies

- The trip system is represented as a single channel. A real safety system on this service would use redundant transmitters with voting, and could take a single channel failure without losing the protection.
- The closing time of UZV-101 is not modelled: the valve is open or shut within the same scan.
- The PSV is modelled with a clean opening and a relief flow computed with the same standard as the control valve. The progressive opening of a pilot operated valve is not represented, nor the back pressure on the vent header.

6. The gas, the pressures and the disturbances

6.1 Why a gas is not handled like a liquid

A liquid is practically incompressible: a cubic metre stays a cubic metre all along the line. A gas does not. The same mass of methane occupies very different volumes at different pressures, and this has two consequences visible all over the panel.

The first: **saying “flow” without saying at what conditions means nothing**. That is why gas flow is expressed in **normal** cubic metres — Sm^3/h , the volume that mass would occupy at 0 °C and 1013.25 mbar — which is an indirect way of talking about mass using a unit of volume. The measurement chain box shows both numbers: the flow at line conditions and the normalised one.

The second: density enters every formula, from the orifice measurement to the valve sizing. Changing pressure or temperature is not a detail: it changes the process.

6.2 The compressibility factor Z

The perfect gas is an approximation. At the pressures of this line methane departs from it by about 2 %, and the factor Z measures exactly that departure: $Z = 1$ would be a perfect gas, here it is about 0.98. The panel shows it openly in the gas box because it enters the density calculation and therefore the measurement.

Quantity	Value for methane
Molar mass	16.04 g/mol
Isentropic exponent k	1.31
Density at normal conditions	0.7175 kg/m ³

Quantity	Value for methane
Joule-Thomson coefficient	0.45 K/bar

6.3 Network pressure and user back pressure

They are the two main disturbances, and they act in different ways:

- The **network pressure** (PT-100) is what the line has available. If it falls, less gas passes at the same opening: the controller has to open further to hold the set-point, and if the network falls too far the set-point becomes unreachable and the output saturates.
- The **user back pressure** (PT-102) is what the line has to overcome. If it rises, the drop available across the valve shrinks and the same thing happens. If it falls a long way, the valve chokes (ch. 2.3) and the flow stops responding.

Both are **load** disturbances: they do not touch the set-point, they touch the relation between opening and flow. A well tuned loop absorbs them without the operator noticing; that is exactly what scenarios 4 and 5 practise.

6.4 Temperature and the Joule-Thomson effect

Gas temperature enters the density, therefore the measurement, and is one of the three inputs of the compensation.

There is however also an effect the gas produces by itself: going through a restriction and losing pressure, methane **cools down**. It is the Joule-Thomson effect, about 0.45 degrees per bar of drop. On the panel it can be seen by comparing TT-101, upstream of the valve, with TT-102 downstream: the difference is not a measurement error, it is physics. On plants with large pressure drops it is the reason gas is preheated before reducing it.

6.5 A note for the trainee — what the model simplifies

- The gas is pure methane. A real network gas is a mixture, with density and calorific value changing over time: a custody transfer computer accounts for this with a chromatographic analysis or a continuous density measurement.
- The network pressure is treated as an imposed source: the model does not represent the upstream network pressure drop as the offtake grows.
- The Joule-Thomson effect is applied with a constant coefficient. In reality it depends on pressure and temperature, and for some gases it even changes sign.
- Upstream and downstream temperatures relax towards ambient with fixed time constants: heat exchange with the pipe computed on the real geometry is not represented.