
Heater F01 — Operating principles

Trainee document · burner management system

DOCUMENT	Operating principles of the heater and its burner management system
REVISION	01 — first issue
PURPOSE	Training of instrument and operations technicians

How to read this document

The simulator reproduces a **ten-burner process heater** and the system that governs its lighting and its safety: two blowers shared with the adjacent heater, four inlet dampers, two stack dampers, ten pilots each with its own valve and flame detector, ten burners, and four process passes crossing the combustion chamber.

Every number the panel shows comes out of a model solved ten times a second. There are no fake animations: close a damper and the air flow really stops; lose a pilot flame and the matching burner cannot stay lit. This document explains **why** each system exists and **how** it reacts, so that what appears on the panel can be anticipated rather than endured.

Instrument tags (UZV-140-3, HV-109-7, BXD-140, PIT-126...) are not translated in any language: they are identifiers, not words. The last chapter states what the model simplifies, because knowing where the simulator ends is part of the training.

1. Heater F01: ten burners on natural draft

1.1 What it heats

Four separate process lines run through the combustion chamber, each with its own flow meter: FIT-138, FIT-139, FIT-140, FIT-141. They are four independent passes, not four branches of one circuit — a detail that matters, because it means the heat must be spread along the whole chamber rather than concentrated in one zone. That is why there are ten burners and not one large one, and why flame loss on **adjacent** burners is treated differently from flame loss on scattered ones (chapter 6).

The **firing rate** shown in the header is simply the fraction of burners lit out of the ten available, filtered with a 6 second time constant to represent the thermal inertia of the chamber.

1.2 Why the blowers serve the purge only

This is what surprises most people looking at the panel for the first time: **light the pilots and the blowers stop**. It is not a fault, it is the correct behaviour.

The heater runs on **natural draft**: combustion air enters the windbox because of the stack, not because something pushes it. The two blowers K102A and K102B exist for one purpose — **sweeping the chamber before lighting** — and once the purge is over they have no reason to run. More than that: blowing air into the chamber with the pilots open is exactly what the hardwired interlock of chapter 3 exists to prevent.

On the mimic, when a blower is commanded to run but held by the interlock, an amber **ITL** indicator appears next to its symbol. It distinguishes a machine that is *held* from one that is merely *stopped*: two different diagnoses.

1.3 Signals in this chapter

TAG	DESCRIPTION	WHERE TO READ IT
FIT-138 / 139 / 140 / 141	Flow of the four process lines	Mimic and analogue values
PIT-126	Fuel gas pressure at the header	Analogue values and mimic
K102A / K102B	Purge blowers shared with heater F02	Mimic, command column

2. The purge: why it is done and why it takes so long

2.1 The problem it solves

Before presenting a spark in the chamber you must be certain the chamber holds no flammable mixture. A gas leak accumulated with the heater down, however small, produces a violent ignition at the first attempt at a spark: not a flame, a deflagration.

The purge is the air change that removes that possibility. It is not a formality and it cannot be shortened: it is the only barrier between an undetected leak and an explosion in the chamber.

2.2 The five permissives

The purge can start only if **all** of these are true together:

1. **all pilot valves closed** — XL-PILOTS-RDY;
2. **all burner valves closed** — XL-BURNERS-RDY;
3. **both stack dampers open**, HV-105 and HV-106 — XL-RDY-DMPR: with no way out the air does not cross the chamber, it just circulates;
4. **no trip active and no shutdown command present**, from the local pushbuttons or from the ESD;
5. **at least one valid air path** (next paragraph).

When all five are satisfied **XL-RDY-PURGE** lights and the HS-PURGE pushbutton becomes pressable. Until then the panel does not explain the reason with a message: the lamps explain it, one per permissive. This is an interface convention the trainee has to learn to use — you look for the lamp that is missing, you do not look for a text.

2.3 A valid air path, and why 25 or 55 minutes

A path is valid when all three hold:

- the **F01 inlet damper** on that side is open (103A or 103B);
- the **adjacent heater F02 damper** on the same side is closed (104A or 104B), otherwise the air splits between the two heaters;
- the **blower on that side is actually running**, not merely commanded.

With **both** paths active the purge lasts **25 minutes**; with **one only**, **55 minutes**. The ratio is not arbitrary: for the same chamber volume to be exchanged, half the flow needs more than twice the time, because the exchange is not a piston but a dilution.

2.4 Purge memory: 12 hours, three attempts, one total shutdown

When the purge completes the system records a **valid purge memory** and the countdown starts: the panel shows the time left. The memory decays in three ways.

- **After 12 hours.** A chamber swept yesterday is not a swept chamber: a valve may have passed in the meantime.
- **After three failed pilot ignitions.** Each failed attempt has admitted gas into the chamber without lighting it. Three attempts are the limit beyond which that gas must be swept out before trying again.
- **On a total shutdown.** The total shutdown closes pilots and burners and invalidates the memory: resetting and relighting is not enough, the purge must be repeated.

While the memory is valid the pilots can be lit; once it decays, they cannot.

2.5 Purge failure

If even one permissive is lost during the cycle — someone closes a stack damper, the interlock stops the blower, a shutdown command arrives — the purge does **not pause: it fails**. Alarm XA-PURGE-F01 appears and the cycle must be run again from the beginning.

This is strict on purpose. A purge interrupted halfway and then resumed gives no guarantee at all about the number of air changes actually achieved, which is the only thing that counts.

2.6 Times in this chapter

CONDITION	DURATION	NOTE
Two air paths active	25 min	Both blowers running, F02 dampers closed
One path only	55 min	More than twice: the exchange is a dilution
Memory validity	12 h	Countdown visible on the panel
Failed pilot ignitions tolerated	3	At the third the memory decays

3. The hardwired interlock of the shared blowers

3.1 Two blowers for two heaters

K102A and K102B do not belong to this heater: they are **shared** with the adjacent heater F02. Each blower serves one side — A or B — and on each side there are two inlet dampers: one towards F01 (103A, 103B) and one towards F02 (104A, 104B).

That creates a dangerous possibility: **sending air into a heater that is firing**. If blower A is running and damper 103A is open while the F01 pilots are lit, air is being blown onto pilot flames that have minimal power. The result is extinction, and an extinction with the gas valves still open means unburnt gas in the chamber.

3.2 The rule

A blower is **enabled to run** only if, for its side, both hold:

- the damper towards F01 is **closed**, **or** all F01 pilot valves are closed;
- the damper towards F02 is **closed**, **or** all F02 pilot valves are closed.

Put positively: *a blower may run only towards a heater that has no open pilots.*

3.3 Why it is hardwired, and why it is a continuous permissive

The interlock **does not go through the PLC**: it is hardwired logic. That is an architectural choice, not a convenience. A protection that prevents a dangerous air admission must not depend on a program's scan cycle, nor on the program being the right one: it must hold even with the PLC failed.

It is also a **continuous permissive**, not a simple start consent: it is evaluated at every instant, including with the blower already running. That is why, the moment the pilot valves open, the blower **stops**. It is not a sequence someone wrote: it is the interlock withdrawing its consent. And that is correct, because by then the purge is over and the heater carries on under natural draft.

4. Lighting: pilots first, then burners

4.1 Why a pilot

A main burner does not light from a spark. The gas flow is too high and the efflux velocity too fast: the spark cannot stabilise a flame front, and meanwhile gas keeps entering. The **pilot** is a small burner, at reduced flow, that does light from a spark and stays lit: it gives the main burner an ignition source that is **continuous and already proven**.

Hence the basic rule: **no burner may open its valve unless the pilot of the same position has flame confirmed**. The burner pushbutton stays disabled until that condition is true.

4.2 The pilot sequence

STEP	DURATION	WHAT HAPPENS
Opening	5 s	Valve UZV-140-n is commanded open
ZSC check	instantaneous	The closed limit switch must no longer be active
Spark	10 s	Ignition transformer IGN-140 energised
Flame check	instantaneous	Detector BXD-140 must confirm flame

If the ZSC check fails, or if at the end of the spark the detector sees no flame, the position goes to **fault**: the valve closes again, the alarm for that position is raised and the failed ignition counter advances.

4.3 The burner sequence

STEP	DURATION	WHAT HAPPENS
Opening	5 s	Valve HV-109-n is commanded open
ZSO check	instantaneous	The open limit switch is verified
Flame check	10 s	UV scanner XA-111 must confirm flame

The proving time is the part not to be misread: it is not a courtesy wait. It is the maximum time within which, if the flame has not established, the valve must close again — because for its whole duration gas is entering.

4.4 Valve position discrepancy

A closed limit switch ZSC that becomes active again **while the valve is commanded open** is a **discrepancy**: either the valve did not move, or the switch is lying. In both cases the valve position is unknown, and not knowing where a gas shut-off valve is cannot be accepted.

The simulator treats it as a **permanent condition**, not as a check made once: it acts with the pilot running, during ignition and during proving. A field fault does not disappear because the check has already gone by.

The practical consequence is that the position **cannot be recovered from the panel**: the fault is cleared in the field. What can be done from the panel is to exclude it with the MOS bypass and keep the heater in service on the other nine (chapter 5.4).

4.5 The pilot lighting permissive

A pilot may open if **one** of these holds:

- the **purge memory is valid**, or
- **at least one pilot is already lit**.

The second looks like a shortcut. It is not. If a stable flame is already present in the chamber, the mixture build-up the purge exists to prevent cannot occur: lighting a second pilot is a different operation from lighting the first. It is also the reason why, if **all** pilots are shut down, lighting one again requires a new purge.

5. The heater ON state

5.1 Why it has to be defined

Two of the automatic protections are armed **only with the heater in service**. The reason is obvious seen from the right side: "flame loss on three adjacent burners" is an abnormal condition only if those burners were supposed to be lit. On a heater being started up, with six positions still cold, the same measurement is the normal state.

So a precise definition of "heater in service" is needed, and it has to be one the system can evaluate by itself.

5.2 The definition adopted

The heater is in the **ON state** when, for **every** position 1 to 10, at least one of these holds:

- pilot **and** burner have flame confirmed, **or**
- the position is **excluded by the MOS bypass**;

and in addition **at least one position is actually lit**.

The last condition is the guard against the absurd case: without it, excluding all ten positions with the MOS bypasses would give a heater formally "in service" while completely cold, with its protections armed on nothing.

The state is shown by lamp **XS-F01-ON**.

5.3 Why it is a memory and not an instantaneous calculation

The condition is **set** when it becomes true and **cleared** when all pilots are off. It is not recalculated every scan because, if it were, a race would appear: the instant three burners lose flame, the condition "all positions lit" becomes false **in the same instant** the protection should act — and the protection would disarm itself exactly when it is needed. The memory breaks that race.

5.4 The MOS bypass

Each position has a **MOS** — a maintenance bypass that excludes the pilot and the burner of that position together. It exists to keep the heater in service when one position is faulty and cannot be repaired straight away.

It has a stated limit, though: **ten MOS active together are a partial shutdown cause**. A heater with every position excluded is not a protected heater, it is a heater with its protections defeated, and the system refuses to stay in that state. The bypass is a tool for one position, not for the heater.

6. Partial and total shutdown

6.1 Two levels, two logics

	PARTIAL SHUTDOWN	TOTAL SHUTDOWN
What it closes	Burners only	Burners and pilots
Purge memory	Stays valid	Invalidated
What it takes to restart	Reset and relight the burners	Reset, new purge , full relight

The difference is not one of severity, it is one of nature. The partial shutdown removes the heat while leaving the proven ignition sources in the chamber: service can resume in minutes. The total shutdown removes those too, and with them any certainty about the state of the chamber: from there you restart only by going through the purge again.

6.2 The causes

CAUSE	LEVEL
Partial trip pushbutton on the local panel	Partial
Partial shutdown command from the ESD	Partial
Ten MOS active simultaneously	Partial
Loss of all burner flames with the heater in the ON state	Partial
Emergency pushbutton on the local panel	Total
Total shutdown command from the ESD	Total
Flame loss on three adjacent burners and their pilots , heater in the ON state	Total

6.3 Why "adjacent", and why the pilots too

Three burners out **scattered** along the chamber are a loss of duty: a nuisance, not a hazard. Three burners out **next to each other** are an extended cold zone, and a cold zone in a chamber fed with gas is where unburnt gas collects.

That the condition also requires the loss of the **matching pilots** is just as meaningful: if the pilots of those three positions are still lit, a continuous ignition source still exists in that zone, and any gas arriving there burns instead of accumulating. It is the simultaneous loss of both levels that constitutes the hazard.

6.4 Loss of all flames, and how it differs from a shutdown

"Loss of all burner flames" must not trip when it is the operator shutting the burners down from the pushbutton. The simulator tells the two apart by comparing the state with the **previous scan** — there were flames and they were commanded, now there are none — instead of relying on a latch. Stopping the burners from the panel raises no trip, and that is exactly what the test suite checks.

6.5 First-out

When a heater trips, many alarms arrive within seconds: flames drop, valves close, derived signals act. Knowing **which cause came first** is the difference between understanding the failure and guessing at it.

The alarm list marks the first critical cause with a **First-out** label and states it in plain words next to the alarm — not a code, the sentence. The marking stays until a reset is performed.

7. ISA-18.2 alarms

7.1 The three states

STATE	MEANING
ACT	Active: the condition is present and the alarm has not been acknowledged
ACK	Acknowledged: the operator has taken it on, the condition is still present
RTN	Returned: the condition has cleared, but the alarm has not been acknowledged yet

The distinction between ACK and RTN is the one most often forgotten. A **returned** alarm stays in the list precisely so the operator can notice that something happened while they were looking elsewhere: it disappears only on acknowledgement.

7.2 Blinking

In the pilot and burner matrix, a position in fault **blinks red until the alarm is acknowledged**, then stays steady red. It is the same convention as the banner and the common alarm lamp: *blinking means not yet taken on*, it does not mean "more severe". A steady red fault is a fault somebody has seen.

7.3 Lamp test

The lamp test pushbutton lights **every** panel lamp, the fifty matrix LEDs, the beacon and the horn, for as long as it is held down. It is deliberately a held-button function: a lamp test that starts and runs on its own leaves the panel lit with nobody watching it, and for that duration no genuine alarm would be distinguishable.

8. The analogue values

PIT-126 — fuel gas pressure. The header below the burners starts around 4.2 bar with the heater down and **falls by about 0.11 bar for each burner lit**: at ten burners it settles around 3.1 bar. That is the behaviour of a supply network: more consumers open, less pressure at the header. It is the signal that reveals a gas supply anomaly before it becomes a flame loss.

FIT-138 ... FIT-141 — process flows. The four lines swing slowly around 62 m³/h with a ± 14 m³/h excursion and different phases, representing a process load that varies. They are not controlled by the burner system: they are the **disturbance** the heater works against.

Firing rate. Fraction of burners lit, filtered at 6 seconds.

9. The time scale

The simulation runs on a **fixed 100 ms scan** of process time. The header offers a **×1 / ×10 / ×60** accelerator that multiplies process time, not clock time: it exists so nobody waits 25 or 55 real minutes for a purge.

In the training module the distinction is explicit: the accelerator shortens the wait on processes, but the **exercise clock and the dwell time the criteria require are measured in real time**. Speeding up does not pass a test any sooner.

10. What the model simplifies

Knowing where the simulator ends is part of the training. The stated simplifications:

- **The adjacent heater F02 is not simulated.** It appears only through its two inlet dampers, which contend for the blowers. That is enough to make the interlock real; it does not reproduce that heater's sequences.
- **Combustion is not modelled chemically.** There is no excess air, no flue oxygen, no chamber temperature and no efficiency: the flame is present or absent.
- **Draft is implicit.** The stack dampers are a purge permissive and a way out for the flow, but no chamber draft is computed and none is controlled.
- **The ten positions are identical.** No difference in setting, wear or location relative to the process passes.
- **Process flows are an imposed disturbance,** not the result of a heat balance: they do not react to the duty developed by the burners.
- **Gas pressure is an algebraic relation** between the number of burners lit and header pressure, not a network model with pressure drops.
- **Flame detectors are ideal.** There are no marginal readings, no fouled optics and no false signals: the flame is either seen or not seen.

None of these simplifications touches what the panel exists to train: understanding what a purge is for and why it is not shortened, telling a hardwired interlock from a fault, carrying out a staged light-off, knowing when the protections are armed and when they are not, telling a partial shutdown from a total one, and reconstructing a cause from the first-out.