# Title

SuperFract™ Adaptive-Dosing Plasma Protein Fractionation System with Hierarchical Multi-Axis Micro-Dosing Control

# Technical Field

This disclosure relates to real-time, closed-loop control of plasma protein fractionation (precipitation) processes using at-line/inline analytical sensors and adaptive reagent dosing algorithms. In particular, it concerns a skid-mounted micro-dosing apparatus and integrated control system for automated bioprocessing (e.g. plasma fractionation, viral inactivation, formulation, cell culture) that delivers precise sub-microliter reagent doses under multi-loop feedback control while maintaining Good Manufacturing Practice (GMP) conditions.

# Background

Plasma protein fractionation (such as Cohn ethanol precipitation) is traditionally performed in batch mode with intermittent manual sampling. Operators withdraw 10–20 mL “grab” samples from large precipitation tanks (1,000–3,000 L) every 15–30 minutes, clarify each sample, then measure protein concentration offline (e.g. UV-280 absorbance or Kjeldahl nitrogen). Each assay incurs a 10–30 minute turnaround, during which the batch continues to stir and cool. To avoid waiting for results, operators commonly overdose the precipitating agent (ethanol or polyethylene glycol) by ~5–10% v/v and extend mixing time to “ensure” completion. This practice leads to yield loss (e.g. IgG yield reduced by 2–4 percentage points) and higher costs – excess precipitant must be removed (increasing solvent recovery energy) and filters are overloaded with unnecessary precipitate, causing premature fouling and increased filter consumption. Furthermore, offline sampling creates compliance and biosafety challenges: each dip-sample breaches the closed sterile system, requiring laminar flow hoods or isolators, disinfection, and meticulous log tracking. Manual handling and labeling of samples introduce risks of mislabeling, transcription error, and data gaps, complicating deviation investigations and regulatory audits. Critically, because sampling is infrequent, process anomalies (e.g. vortex entrainment, jacket hot-spots, or sensor drift) can go undetected for an entire batch, potentially compromising product quality or forcing costly reprocessing.

Similar issues arise in other bioprocesses where trace reagents (e.g. acids/bases for pH shift, viral inactivation agents, endotoxin binders, surfactants, or growth factors) are added in tiny increments. Historically, such additions have been done manually or with open-loop pumps, resulting in variability and suboptimal control. There is a clear need for an in-situ, self-validating dosing system that provides continuous, second-by-second feedback control, minimizes overshoot of reagents, maintains aseptic containment, and digitally records all process data without manual sampling. The present invention addresses these needs by integrating online analytics with an adaptive multi-axis dosing apparatus to tightly control precipitation and similar bioprocessing steps in real time.

# Summary

The invention provides a SuperFract™ adaptive-dosing system for plasma protein fractionation that dynamically adjusts precipitant addition based on live feedback from process analytical technology (PAT) sensors. In one aspect, a sidestream of the batch (process liquor) is withdrawn from the precipitation vessel through a sterile isolation valve and pump into an at-line flow cell containing optical and acoustic sensors. For example, a turbidity sensor and an ultrasonic transducer measure protein aggregation and solute concentration in real time. Optionally, a slip-stream can be filtered through a micro-clarifier and passed through a UV absorbance detector to measure protein content or clarity. A sensor-fusion algorithm in the system's controller continuously computes a quantitative precipitation metric (e.g. % protein precipitated) from these signals. The controller then modulates a reagent dosing pump accordingly – for instance, increasing or decreasing the ethanol/PEG addition rate – to maintain the optimal precipitation rate or to reach a target endpoint without overshoot. As precipitation proceeds, an integrated cross-flow filtration unit may be used downstream; pressure sensors on the filter measure the filtrate's trans-membrane pressure (TMP) slope, providing feedback on precipitate load and filtration health. The system uses this information to decide when to terminate or recycle the batch (e.g. if excessive TMP rise is detected). All of these operations occur in a closed loop without manual intervention, thereby minimizing reagent overdose and maximizing IgG yield.

Another aspect of the invention is a hierarchical closed-loop micro-dosing skid that serves as the reagent delivery module for the above system. In one embodiment, the skid comprises a manifold of multiple micro-annular gear pumps (e.g. N = 2–21 pumps, each <6 mL/min capacity) arranged in parallel. These positive-displacement pumps have chemically resistant 316L stainless steel/PEEK wet-ends and are magnetically coupled to servo drives on a high-speed fieldbus, enabling precise metering of microliter-scale doses. The outlet of each pump connects to a sanitary injection ring that mounts on the main process line. The injection ring has multiple radial nozzles (quills) merging into a common aseptic feed, such that micro-volume reagent shots are rapidly dispersed into the main stream with minimal coefficient of variation (e.g. <1% CV for a 5 µL injection). The pump manifold and associated fluid lines are designed for full Clean-in-Place/Sterilization-in-Place (CIP/SIP) sanitization – for example, the pump heads and ring can withstand 121–130 °C steam sterilization for ≥20 minutes to achieve sterility between batches. In some implementations, the pump ports can also accept single-use disposable cartridges (e.g. pre-filled Luer-lock reagent vials) to combine sterile disposability with the accuracy of positive displacement pumping.

The dosing skid may further integrate a compact analytical sensor spool downstream of the injection point. This spool can house one or more inline sensors such as: a quartz crystal microbalance with dissipation (QCM-D) to detect microscopic precipitation or biofilm formation on a sensor crystal, a microfluidic endotoxin detector (e.g. a limulus amebocyte lysate sensor) for real-time impurity monitoring, and/or a redox (ORP) probe to monitor chemical conditions. These inline signals feed into the control logic (e.g. a TwinCAT™ or PLC function block) to automatically adjust pump dosing and ensure each reagent addition achieves its intended effect. The controller records every dose event with its corresponding reagent ID, volume, lot, time stamp and any PAT sensor readings, storing this data as a unique batch lineage record in a historian database. This electronic batch record enables full traceability and is compliant with data integrity regulations (e.g. EU Annex 11 and 21 CFR §11) regarding audit trails and genealogy.

In operation, the adaptive control system orchestrates multiple feedback loops in a hierarchical manner. A primary control loop (Yield Optimizer) adjusts the precipitant addition rate to maintain the process on target (e.g. aiming for a specific turbidity or precipitation fraction) using the fused sensor input. A secondary loop (Quality Controller) monitors critical quality parameters – for example, if an endotoxin sensor or other critical process parameter exceeds a safety threshold, a Vital Override layer immediately pauses all pumps to prevent contamination or product loss. Meanwhile, a background self-diagnostic routine performs periodic auto-calibrations: for instance, the system can momentarily divert flow through a self-test cartridge with known reference solutions or empty chamber to validate that sensors read zero or baseline correctly. If sensors disagree beyond a tolerance (e.g. >2σ deviation between redundant sensors) or if expected sensor updates are missed, the controller triggers a fail-safe mode that freezes dosing and alerts operators. The control algorithms themselves can range from model-predictive control (utilizing a state-space model of precipitation) to rule-based or machine-learning approaches. In some embodiments, a machine-learning anomaly detection model continuously checks the sensor data for abnormal patterns and can invoke fail-safe actions or adaptive tuning if a deviation is detected (e.g. unexpected turbidity fluctuations). The system is also cloud-enabled: process data can stream to a cloud platform where a remote machine-learning model (digital twin) analyzes performance and sends back updated control parameters or model weights to the local controller. This cloud supervision allows the control algorithm to “learn” from each batch and improve over time, while still executing in real time on an edge PLC at the process site.

By combining multi-modal sensing, intelligent control software, and a flexible micro-dosing hardware platform, the invention ensures precipitation steps are executed with optimal precision and reproducibility. It eliminates the need for manual sampling and overdosage buffers, thus improving yield and reducing process time. The closed-loop system maintains sterile boundaries at all times and automatically logs comprehensive batch data for compliance. In summary, the disclosed adaptive-dosing platform transforms traditionally open-loop, labor-intensive precipitation operations into a fully automated, self-correcting process that is safer, more efficient, and more consistent.

# Brief Description of the Drawings

- FIG. 801: Schematic of a side-loop dual-sensor control circuit (Embodiment E-1). This figure shows the precipitation tank and a recirculation loop with an isolation valve, loop pump, and flow cell housing an optical turbidity sensor and an ultrasonic sensor. A main dosing pump introduces precipitant into the tank, and a return valve directs the side-stream back to the vessel.

- FIG. 802: Schematic of a tri-sensor arrangement with inline UV clarifier (Embodiment E-2). It extends FIG. 801 by adding a splitter valve that diverts a portion of the side-stream through a 0.2 µm clarifier filter and a UV-280 absorbance detector. This allows concurrent measurement of optical density and clarified protein concentration.

- FIG. 803: Inline acoustic resonance probe setup (Embodiment E-3). A piezoelectric probe is mounted directly in the wall of the precipitation vessel. The probe's resonance circuit (driver and sensor) connects to the controller to measure solution density or sound velocity changes indicative of protein precipitation.

- FIG. 804: Inline dielectric spectroscopy cartridge (Embodiment E-4). A single-use disposable cartridge with twin electrodes is inserted into the process flow. An impedance chip measures the capacitance or permittivity of the mixture between the electrodes, which correlates with protein content and precipitation progress. Sterile Luer-lock ports allow installing or removing this cartridge as needed.

- FIG. 805: Laser-diffraction particle analyzer (Embodiment E-5). A laser scattering cell and photodiode array are connected to the side-stream. This setup measures the particle size distribution of precipitate solids in real time, feeding size data to the controller (via a data processor) to infer precipitation completeness.

- FIG. 806: Model predictive control (MPC) architecture (Embodiment E-6). This block-diagram figure depicts a state-space model module receiving a vector of sensor inputs, an optimizer computing the future dosing trajectory, and an actuator output block driving the pump. A constraint module ensures operational and safety limits (e.g. max dose rate) are not violated.

- FIG. 807: Rule-based adaptive recipe control (Embodiment E-7). Shown here is a PLC ladder logic diagram with a recipe look-up table (LUT) and set-point trim block. The controller adjusts dosing according to predefined recipe stages and fine-tunes the pump set-point based on feedback, while an alarm block handles any rule violations.

- FIG. 808: Machine-learning anomaly guard (Embodiment E-8). This figure illustrates a feature extractor processing sensor data, a Gradient Boosted Trees (GBT) or similar classifier evaluating the data for anomalies, and a failover state machine. If an anomaly is detected, the system can initiate remedial actions or revert to safe mode.

- FIG. 809: Cross-flow filtration ΔP feedback loop (Embodiment E-9). FIG. 809 shows a post-precipitation filtration unit with inlet and outlet pressure sensors. A calculator computes the TMP (ΔP) across the filter in real time. A three-way valve on the filtrate line is controlled to either send flow forward or recycle it back to the tank depending on the TMP slope (high ΔP triggers recycle).

- FIG. 810: PWM dosing control scheme (Embodiment E-10). A pulse-width modulation controller drives the reagent pump in pulses. A flow transmitter provides feedback on actual flow rate, and an anti-chatter logic block smooths the output. This allows very fine control of effective flow by adjusting duty cycle, even if the pump has a minimum stable flow rate.

- FIG. 811: Dual-stage pH and PEG cascade process (Embodiment E-11). This figure shows two dosing pumps: one (PD-210) adding an acid or base to shift pH (monitored by a pH sensor), and another (PD-211) adding polyethylene glycol (PEG). A fusion control output coordinates the two additions – the first pump brings the solution to a target pH, then the second pump introduces PEG to precipitate specific proteins once pH is adjusted.

- FIG. 812: Temperature-ramped ethanol fractionation (Embodiment E-12). Illustrated here is a jacketed precipitation vessel with a temperature control unit (TIC-102) and the main ethanol pump. The controller ramps the jacket temperature according to a predefined profile while simultaneously adjusting ethanol dosing. A ramp generator module ensures the temperature and addition rate follow the desired time trajectory to optimize fractionation.

- FIG. 813: Cloud-supervised digital twin architecture (Embodiment E-13). This diagram depicts an edge programmable logic controller (PLC) on site, an MQTT communication broker, and a cloud machine-learning server. The edge PLC sends process data to the cloud, where an ML model analyses it and an update agent returns optimized parameters or control policy adjustments to the PLC. This closed-loop cloud feedback improves control performance across batches.

- FIG. 814: Self-testing sensor cartridge (Embodiment E-14). FIG. 814 shows a specialized sensor module with an auto-zero chamber and micro-valves. The cartridge can temporarily isolate a small volume of fluid or a reference solution in front of the sensors (optical test path, acoustic test path) via built-in valves. A self-test circuit triggers these valves and compares sensor readings in the reference condition to expected values, thereby verifying calibration and performance of the sensors over time.

## Detailed Description

Overall System Configuration (FIGS. 801–802): The adaptive fractionation control system will first be described with reference to FIG. 801 (Embodiment E-1). A main precipitation vessel T-101 (801-1) holds the plasma or process liquid in which proteins are being precipitated (e.g. via ethanol addition). The vessel is equipped with conventional agitators and a cooling jacket (not shown) to maintain the required temperature. A side-loop circuit withdraws a small portion of the tank contents for analysis: an isolation valve XV-301 (801-2) on the tank outlet opens to let liquid flow through a side-loop pump P-201 (801-3). Pump P-201 is a sanitary centrifugal or positive-displacement pump sized to circulate a slip-stream (for example, ~0.5–5% of total volume) through the loop. Downstream of the pump, a flow cell FC-301 (801-4) houses at least two analytical sensors. In this embodiment, an optical turbidity sensor NTU-301 (801-5) is installed to measure the suspension's turbidity (optical density due to precipitate) in real time, and an ultrasonic sensor US-302 (801-6) measures acoustic properties (e.g. sound velocity or attenuation) of the liquid. The two sensors provide complementary data: turbidity correlates with bulk protein aggregate formation, while ultrasonic readings can indicate changes in solvent composition or protein concentration. The flow cell may also include an optical reference element PRX-303 (801-7) for calibration (e.g. a reference photodiode or known optical target for turbidity baseline).

After passing the sensors, the side-loop stream returns to the main tank via return valve XV-302 (801-8), unless an alternate path is engaged. Embodiment E-2 (FIG. 802) introduces an optional alternate path: a clarifier branch for sample conditioning. In FIG. 802, a splitter valve XV-311 (802-2) can divert a portion of the side-loop flow through a clarifier filter cartridge CF-401 (802-3), such as a 0.2 µm sterile filter, to remove solids. The clarified liquid then passes through a UV absorbance detector UV-401 (802-4) that measures protein concentration at 280 nm (or other wavelength). This configuration (Embodiment E-2) thus adds a third sensor modality – UV spectrophotometry – which directly quantifies dissolved protein in the supernatant. The clarifier branch has its own return/waste valve XV-312 (802-5) to either send the analyzed, clarified sample back to the process or discard it as needed. In normal operation, the system can continuously toggle between the main flow cell (for turbidity and ultrasonic readings) and the clarifier/UV loop to gather a wide range of real-time data. All side-loop components are arranged in a compact skid close to the tank, with sanitary connectors to ensure the sample loop is representative of tank conditions and quickly returns samples after analysis.

Reagent Dosing Module (FIGS. 801 & 810): The controlled addition of precipitant (such as ethanol) is achieved through a dedicated dosing pump PD-202 (801-9). Pump PD-202 draws the reagent from a source (e.g. ethanol storage) and injects it into the main vessel or an upstream mixing point. In FIG. 801, PD-202 is the main dosing pump feeding directly into T-101. The operation of PD-202 is governed by the control system to achieve adaptive dosing. In one simple implementation, PD-202 may be a variable-speed pump (e.g. gear pump or peristaltic) that the controller throttles up or down continuously. In Embodiment E-10 (FIG. 810), a more advanced control of dosing is depicted: the pump is driven in a pulse-width modulation (PWM) manner. In FIG. 810, the reagent pump (810-1) is rapidly switched on and off by a duty-cycle controller (810-2) to effectively meter small amounts of reagent with high precision. A flow transmitter FT-204 (810-3) provides feedback on the actual delivered flow rate, and an anti-chatter logic block (810-4) smooths out the pulses to avoid any oscillation in the tank. This PWM approach allows the system to achieve effectively very low dosing rates (by using short pulses at controlled frequency) even if the pump's minimum stable flow rate is relatively high.

The invention’s dosing hardware is not limited to a single pump. Referring now to FIG. 805 and FIG. 806 in combination with FIG. 801, the system may incorporate the aforementioned multi-axis micro-dosing skid for more complex control. In one embodiment, the single pump PD-202 is replaced or augmented by a manifold of multiple micro-pumps (FIG. 805 can be considered conceptually here, though it focuses on particle analysis). For example, a rack of micro-annular gear pumps driven by servomotors can be used to deliver precise reagent doses. Each pump in the rack can be assigned to the same reagent for redundancy/high-flow, or different pumps can handle different reagents. The pumps are preferably sized for fine dosing resolution (e.g. each pump having <6 mL/min max flow with accuracy down to a few microliters). They are constructed of biocompatible materials (e.g. 316L stainless steel pump heads and rotors, and PEEK static seals) to withstand corrosive reagents and meet pharmaceutical cleanliness standards. The pump drives are electronically controlled (for instance, via an EtherCAT® or similar high-speed bus) allowing the central controller to command each pump’s speed or stroke individually in real time.

Multiple pumps enable simultaneous multi-reagent control. For instance, in Embodiment E-11 (FIG. 811), two pumps are used: one (811-1) doses an acid or base for pH adjustment, while another (811-3) doses PEG for protein precipitation. The system coordinates these additions: first bringing the solution to a target pH (monitored by a pH sensor PT-103, 811-2), then initiating PEG addition once pH is in range. The ability to run two dosing pumps in parallel shortens process time and improves control over sequential precipitation steps (for example, Fraction II and Fraction III in plasma fractionation can be executed back-to-back without manual intervention by using one pump for acid and one for ethanol/PEG). Another scenario is Embodiment E-12 (FIG. 812), where precipitation is influenced by temperature ramping. In this case, one subsystem (812-2/TIC-102) controls the jacket temperature of the vessel, while pump PD-202 (812-3) adds ethanol according to a time-based profile. The controller (812-4, a ramp generator module) synchronizes the cooling rate and dosing rate to achieve an optimal precipitation trajectory. This hierarchical control ensures the temperature set-point and reagent addition are adjusted in tandem – effectively a multi-input-multi-output (MIMO) control problem handled by the system.

Micro-Dosing Manifold Details (FIG. 805, FIG. 813): The micro-dosing skid mentioned can have significant hardware innovations to facilitate flexible, GMP-compliant operations. In one embodiment, the skid includes a sanitary injection ring as described in the Summary. This ring (not explicitly shown in earlier figures, but part of the dosing module) has multiple radial injection quills that enter the main flow at evenly spaced intervals. By delivering reagent into the flow at multiple points, it ensures rapid mixing and reduces local concentration spikes. The ring is typically installed in-line on the feed to the precipitation tank or in the recirculation loop. All wetted components of the manifold and ring are designed for easy sterilization: for example, the entire fluid path can be subjected to steam sterilization at 121–130 °C for at least 20–30 minutes (SIP protocol) without disassembly. This allows re-use in multi-batch manufacturing while maintaining sterility. In some implementations, single-use disposable options are employed: the micro-pumps can accept disposable reagent cartridges (e.g. small sterile bags or vials that plug into the pump inlet via Luer locks). After a batch, the cartridges (and possibly the pump heads if they are part of the cartridge) are simply discarded, eliminating cleaning requirements. This hybrid approach offers the accuracy of fixed pumps and the convenience of disposable fluid paths.

The pump manifold is controlled by a dynamic assignment logic. Each pump can be directed to draw from any of k available reagents, and the controller can reassign pump roles on the fly. For example, if one reagent is no longer needed mid-batch, that pump can be flushed and switched to start dosing a different reagent. A small-volume flush (e.g. <50 μL of water-for-injection) is pumped through to purge any residue, and the inline QCM-D sensor can verify that no significant material remains (e.g. a frequency shift Δf <2 Hz indicates negligible carryover). Only after this verification does the controller permit the pump to dispense the new reagent, thereby preventing cross-contamination. All pump usage data (reagent identity, volumes dispensed, flush volumes, etc.) are logged with timestamp and lot information, building a genealogy record for each pump's contribution. This allows full traceability of which reagent went through which pump at what time, in compliance with electronic batch record regulations.

To accommodate different process scales, the manifold is scalable and reconfigurable. Additional pump modules can be added to the rack to increase dosing throughput or handle more reagents in parallel. Likewise, the injection ring diameter and the number of nozzles can be increased for larger flow rates, while maintaining the same dosing precision by adjusting nozzle diameters proportionally. In one variant, a dual-ring architecture is employed: one ring (inner) is fed by a subset of pumps (e.g. those carrying aqueous reagents) and a second ring (outer) fed by others (e.g. those carrying organic solvents). This separation, illustrated conceptually in an alternate design, permits simultaneous dosing of two reagent types without mixing them until they enter the main stream. For example, ethanol could be injected via the outer ring and an aqueous buffer via the inner ring at the same time, which might be useful in certain staged precipitations. The rings are concentric and engineered so that their flows quickly merge in the main pipe.

Sensor Modalities and Fusion (FIGS. 803–805): As noted, the system can utilize multiple sensor types to monitor the state of the precipitation. Embodiments E-3, E-4, and E-5 illustrate three additional modalities beyond turbidity and UV. In E-3 (FIG. 803), an inline acoustic resonance probe is used. This probe 803-1 is a piezoelectric device mounted flush with the vessel wall so that one face is in contact with the liquid. The probe's driver and sensing circuitry 803-2 excite it over a frequency range and measure the resonance characteristics. The resonance frequency and damping can shift based on fluid density, viscosity, and composition (for instance, as proteins precipitate out, the density of the remaining solution changes). These measurements are sent via an interface 803-3 to the fusion controller FUS-900, which combines them with other sensor inputs. In E-4 (FIG. 804), a dielectric spectroscopy cartridge is shown. This cartridge 804-1 is a small flow-through cell with disposable construction. It contains twin electrodes 804-2 that contact the fluid and connect to an impedance analyzer chip 804-4. By measuring impedance or capacitance across a range of frequencies, the system can detect changes in the polarization properties of the solution, which relate to protein concentration and aggregation state. The cartridge may be installed via sterile connectors (804-5) and can be single-use to avoid cleaning. Its readings, too, feed into the main controller. In E-5 (FIG. 805), a laser diffraction particle size analyzer is integrated into the side loop. The arrangement includes a laser source 805-1 directing a beam through the sample flow (in scattering cell 805-2), and a photodiode array 805-3 to detect scattering at various angles. A data processor 805-4 computes the size distribution of particles (precipitate flocs) from the scattering pattern. This information (particle size and potentially count) is invaluable for understanding the progress of aggregation; for instance, an increase in average particle size may indicate that flocs are growing and approaching a filter-clogging size. The controller uses a feedback link 805-5 to incorporate this into dosing decisions (e.g. slowing addition if particles are getting too large too fast).

The sensor fusion algorithm (controller FUS-900) combines all available sensor inputs to derive a robust estimate of the state of the precipitation. Some sensors provide direct measurements of the target variable (e.g. UV gives protein concentration in solution, turbidity correlates with precipitate mass), while others provide indirect or predictive information (ultrasonic, dielectric, particle size, etc.). The fusion controller can be implemented as a software library or function block that runs continuously. It may use methods such as a Kalman filter or other state estimator that takes into account sensor noise and dynamics. By fusing signals, the system mitigates the limitations of any single sensor (for example, turbidity may saturate at high solids content, but acoustic or pressure measurements might still indicate ongoing precipitation). The fused output could be a calculated value such as “estimated % precipitation complete” or a rate of change of that value, which is then used by the dosing controller.

Control Algorithms (FIGS. 806–808): The dosing controller applies advanced algorithms to decide how much reagent to add and when, based on the fused sensor feedback. In Embodiment E-6 (FIG. 806), a Model Predictive Control (MPC) approach is depicted. The MPC takes in the sensor vector 806-1 (which could be the fused state estimate or individual sensor readings) and uses a state-space model 806-2 of the precipitation process. This model is essentially a set of equations (or a trained machine learning model) that predicts how the precipitation will respond to reagent addition over time (capturing dynamics like mixing delay, nucleation kinetics, etc.). The optimizer 806-3 then solves for the optimal future addition trajectory (over a prediction horizon) that will achieve the desired endpoint (for instance, 90% precipitation) without overshooting. It outputs an actuator control 806-4 which is the immediate pump rate command, and repeats this optimization at each time step, continuously adjusting the plan. A constraint module 806-5 ensures that certain limits are respected (e.g. not exceeding a maximum pump rate, avoiding temperature dropping below a threshold, etc.). MPC is well-suited to multivariate control, so it can handle multi-input situations like E-11/E-12 (coordinating pH and PEG or temperature and ethanol) within one unified framework.

Embodiment E-7 (FIG. 807) illustrates a rule-based adaptive recipe control strategy, which might be simpler to implement in some settings. In FIG. 807, a PLC ladder logic diagram 807-1 is shown as the control backbone. The system has a pre-programmed recipe LUT (Look-Up Table) 807-2 that defines target set-points for different phases of the process (for example: add ethanol at 100 mL/min until turbidity reaches X; then hold for Y minutes; then add more until Z; etc.). The set-point trim block 807-3 adjusts these targets in real time based on the feedback – for instance, if the turbidity is rising faster than expected, the set-point might be trimmed to a lower value to avoid overshoot. Essentially, the controller still follows a recipe, but can make minor adjustments on the fly (“adaptive tuning”). An alarm block 807-4 monitors for any rule violations or unexpected behavior (like if turbidity doesn't rise when it should, or any sensor reading goes out of acceptable range) and triggers interlocks or alerts. Rule-based control is easier to validate and understand (from a regulatory perspective), though it may be less optimal than MPC or ML-based control.

Embodiment E-8 (FIG. 808) focuses on the machine-learning anomaly detection guard that can overlay any control approach. In FIG. 808, a feature extractor 808-1 processes the raw sensor data to derive features (e.g. slopes, deviations, correlations between sensors). These features feed a trained ML model, such as a Gradient Boosted Tree (GBT) classifier 808-2, which outputs a confidence metric or anomaly score 808-3 indicating how likely the current process behavior is normal versus aberrant. For instance, if sensors start diverging or show a pattern that was associated with past failures (like a sudden drop in ultrasonic signal while turbidity is flat), the model may flag it. The system then can engage a failover state machine 808-4 that implements a predefined response to anomalies. The response could be as simple as pausing reagent addition and notifying an operator, or switching to a more conservative control mode. This ML guardian helps catch subtle issues that rule-based alarms might miss, increasing the robustness of the process.

Integration of Filtration Feedback (FIG. 809): After or during precipitation, the mixture typically goes through filtration to separate solids (precipitates) from the liquid. Embodiment E-9 (FIG. 809) shows a cross-flow (tangential flow) filter unit F-501 with pressure sensors PT-501A (809-1) on the feed/inlet and PT-501B (809-2) on the filtrate/outlet. The difference between these pressures gives the trans-membrane pressure (TMP) across the filter, computed at 809-3. The rate of TMP increase is a key indicator of how quickly the filter is clogging with precipitate. The system uses this feedback in a couple of ways. First, a steady increase in TMP is expected as filtration proceeds; however, if the slope exceeds a threshold (meaning the filter is clogging too fast), the controller can adjust the process: for example, it might recycle the flow back to the tank via recycle valve XV-403 (809-4) instead of forcing it through the filter. In practical terms, the system could temporarily stop filtration and keep mixing the batch (perhaps adding a small dose of solvent to re-dissolve a bit of precipitate) to relieve the filter, then resume filtration. Another use is to define the end-point of precipitation: if TMP rise slows down and stabilizes, it might indicate that no more new precipitate is being formed (i.e. the precipitation reaction is complete). At that point, the controller can conclude that the batch has reached the target precipitation (often ~90% of the target protein precipitated) and can stop further reagent addition. By monitoring ΔP, the system prevents severe over-precipitation that would unnecessarily burden the filters. The filter unit's integration is part of the multi-loop control: while the inner loop handles dosing to reach turbidity targets, an outer loop monitors filtration performance to make higher-level decisions (continue/stop/recycle). For example, the invention may include a rule like: "If TMP > 200 kPa within first 5 min of filtration then halt dosing and recycle batch flow, add 0.5% extra ethanol to dissolve aggregates, then retry filtering" – an automated contingency that improves yield and avoids filter failure.

Cloud Connectivity (FIG. 813): Modern bioprocessing benefits from data-driven optimization. In Embodiment E-13 (FIG. 813), the system is connected to cloud-based analytics. The on-site controller (an edge PLC 813-1) publishes process data (sensor readings, pump statuses, etc.)

to a remote server via a secure MQTT broker 813-2 or similar messaging protocol. In the cloud, a machine-learning server 813-3 (which may host a sophisticated process model or AI that has learned from many batches, possibly across multiple facilities) analyzes the data in near-real-time. The cloud can compare current batch performance to historical trends, detect subtle shifts in raw material behavior, or continuously refine the precipitation model. It then sends back updates via an update agent 813-4 to the local controller. These updates could be new model parameters for the MPC, revised set-points, or even direct control adjustments. The local PLC applies these updates to improve control accuracy. Because connectivity is two-way and continuous, the system effectively implements a “digital twin” of the precipitation process: the physical process is mirrored by a virtual model in the cloud that guides it. Importantly, the system is designed to fail safe – if connectivity is lost or cloud instructions are outside expected bounds, the local controller can ignore them or fall back to its last best parameters, ensuring reliable operation even without cloud support.

Data from the process is also stored for batch reporting and regulatory compliance. Each completed batch yields a rich dataset (time series of sensor readings, pump rates, events, etc.) which is archived and can be reviewed or used for further learning. The cloud link allows multi-batch and multi-plant learning – for example, improving precipitation yield by analyzing dozens of batches. It also offloads heavy computational tasks (like retraining a machine learning model) to the cloud where more processing power is available, while the critical real-time control loop on the edge PLC remains fast and deterministic (only applying the results of those computations).

Self-Testing and Maintenance (FIG. 814): Maintaining sensor accuracy over time is crucial for an autonomous system. Embodiment E-14 (FIG. 814) addresses this with a self-testing probe cartridge. The cartridge integrates small auto-zero chambers 814-1 that can temporarily isolate a miniature volume of fluid adjacent to the sensors. Using micro-valves 814-2, the system can trap either a sample of current process fluid or introduce a known reference fluid into these test chambers. For instance, one chamber might contain a clean solvent (baseline) and another might contain a standard solution of known properties. The sensors (optical, acoustic) have dedicated test paths (814-3 for optical, 814-4 for acoustic) within the cartridge where they can measure these reference fluids. A self-test circuit 814-5 controls the timing of these tests and records the sensor responses. If a turbidity sensor reads anything other than the expected zero when exposed to pure solvent, the system knows that sensor has drifted (e.g. fouled optics) and can recalibrate or mark it for maintenance. Similarly, an acoustic sensor could be tested in air or known liquid to check its baseline. These self-checks can be performed at the start of each batch or periodically (e.g. every few hours) during a long batch, pausing dosing briefly for the test (or using a parallel sensor so one monitors while the other tests). The controller can automatically adjust calibration factors based on the test (for example, subtracting any offset detected in the turbidity reading). In addition to sensor self-tests, the pumps and valves have monitoring for predictive maintenance. Each pump’s run hours, total volume dispensed, and number of start/stop cycles are logged in the system’s historian. A predictive model (which could be a simple threshold or a machine learning model such as a LightGBM regressor) uses this data to predict when a pump’s performance will drift out of spec (e.g. due to seal wear). For instance, it might estimate that after X million revolutions, the pump’s flow accuracy will deviate by >0.2% CV, and thus schedule a maintenance or replacement for that pump at the next downtime. A watchdog safety function is also implemented: if the controller ever fails to receive an expected sensor update twice in a row (indicating a possible sensor fault or communication issue), or if the actual dosing deviates from commanded by more than a set tolerance, the system will automatically disable all pumps and halt the process. This fail-safe ensures that no uncontrolled reagent addition occurs in the event of a hardware or software fault.

GMP and Implementation Considerations: All wetted materials in the system are suitable for biopharmaceutical use (e.g., 316L stainless steel surfaces with <0.5 µm Ra finish, USP Class VI-certified polymers). For embodiments with reusable flow paths, standard CIP/SIP procedures apply (e.g., flushing with water, chemical cleaning, followed by 121 °C steam for ≥20 min). Disposable components (like sensor cartridges or pump cartridges) are delivered sterile (γ-irradiated at ≥25 kGy) and are simply installed prior to use. The system monitors Critical Process Parameters (CPPs) such as temperature, pH, ethanol concentration, etc., to keep them within validated ranges. For example, temperature might be maintained at -5 °C ±2 °C (alarm if >+3 °C deviation) and ethanol at 8% ±0.5% v/v. If a CPP drifts out of permissible range, the controller triggers alarms and can transition to a safe state (e.g., hold addition until corrected). The entire skid and control software are designed to be compliant with regulatory guidelines for equipment and data integrity. An OPC-UA server interface is available, exposing key process data and allowing integration with plant Distributed Control Systems (DCS) or Manufacturing Execution Systems (MES). This means batch reports and real-time status can be directly communicated for electronic batch records and supervisory control.

In summary, the disclosed system provides a comprehensive solution for adaptive precipitation control in bioprocessing. By combining a flexible hardware platform (capable of precise, multi-channel micro-dosing) with smart analytics (multi-sensor fusion, advanced control algorithms, self-diagnostics, and cloud connectivity), it ensures that each precipitation step is performed optimally and consistently. The result is higher yield of target proteins, reduced process times, improved safety (no open handling of biofluids or solvents), and rich data capture for every batch. The various embodiments (E-1 through E-14) illustrate how different features can be implemented or combined: from different sensor setups (optical, acoustic, dielectric, etc.) to different control schemes (MPC, rule-based, etc.) to additional functionalities like self-testing and cloud supervision. All these fall within the scope of the invention, which can be tailored to specific process needs. The claims below seek to protect these novel features and their combinations, ensuring exclusive rights to the key mechanisms that enable real-time adaptive dosing across multiple sensor and control architectures.