

Inexpensive multi-patient respiratory monitoring system for helmet ventilation during COVID-19 pandemic

The Princeton Open Ventilation Monitor Collaboration[¶], Philippe Bourrienne¹, Stanley Chidzik², Daniel Cohen¹, Peter Elmer², Thomas Hallowell³, Todd J. Kilbaugh³, David Lange², Andrew M. Leifer^{2,4*}, Daniel R. Marlow², Peter D. Meyers², Edna Normand^{4,5,6}, Janine Nunes¹, Myungchul Oh², Lyman Page², Talmo Periera⁴, Jim Pivarski², Henry Schreiner⁷, Howard A. Stone¹, David W. Tank^{2,4}, Stephan Thiberge⁴ and Christopher Tully²

1 Department of Mechanical and Aerospace Engineering, Princeton University, Princeton, NJ, 08544, USA

2 Department of Physics, Princeton University, Princeton, NJ, 08544, USA

3 Department of Anesthesiology and Critical Care Medicine, Children's Hospital of Philadelphia, Philadelphia, PA 19104, USA

4 Princeton Neuroscience Institute, Princeton University, Princeton, NJ, 08544, USA

5 Department of Molecular Biology, Princeton University, Princeton, NJ, 08544, USA

6 Rutgers Robert Wood Johnson, New Brunswick, NJ, 08901 USA

7 Princeton Institute for Computational Science and Engineering, Princeton University, Princeton, NJ, 08544, USA

[¶]Membership list can be found in the Acknowledgements section.

* leifer@princeton.edu

Abstract

Helmet non-invasive ventilation (NIV) is a form of continuous positive applied pressure that has emerged as a useful therapy for COVID-19 patients who require respiratory support but may not require invasive ventilation. Helmet NIV has seen an increase in use during the COVID-19 pandemic because it is low-cost, readily available, and provides viral filters between the patient and clinician. Helmet NIV may also provide better patient outcomes by delaying or eliminating the need for invasive ventilation. Its widespread adoption has been limited, however, by the lack of a respiratory monitoring system that is needed to address known safety vulnerabilities and to provide clinicians with a respiratory profile of the patient. To address this safety need, we have developed an inexpensive respiratory monitoring system that is based on readily available commercial components and is suitable for rapid local manufacture. The system is designed for use in conjunction with the COVID-19 Helmet developed by Sea-Long Medical Systems [1], but is modular and can be used with other ventilation systems. The monitoring system comprises one or more flow and pressure sensors and a central remote station that can be used to remotely monitor up to 20 patients simultaneously. The system reports flow, pressure, and clinically relevant metrics including respiratory rate, tidal volume equivalent, peak inspiratory pressure (PIP), positive end-expiratory pressure (PEEP) and the ratio of inspiratory time to expiratory time (I:E). The device will sound alarms based on clinician-set thresholds. In bench tests using a commercial ventilator and artificial lung system, our device performs comparably to a commercial single-patient respiratory monitor. Results are presented from human-subject tests on a healthy volunteer undergoing helmet non-invasive ventilation. Detailed design and manufacturing documents are provided.

Key words: COVID-19, non-invasive ventilation, CPAP, helmet, respiratory profile monitor, remote monitoring, critical care, low-cost ventilator, emergency ventilator.

Contents

1	Introduction	2	1
2	Materials and methods	5	3
2.1	Design and construction of the device and software	5	4
2.1.1	Differential pressure sensor	5	5
2.1.2	Gauge pressure sensor	5	6
2.1.3	Flow-sensor assembly enclosure	7	7
2.1.4	Flow element design	7	8
2.1.5	Flow element construction	7	9
2.1.6	Interface box	8	10
2.1.7	Remote Monitoring Station	11	11
2.1.8	Bill of materials	11	12
2.1.9	Software overview	12	13
2.1.10	Interface box software	12	14
2.1.11	Calculating derived quantities	13	15
2.1.12	Software detection of Alarms	15	16
2.1.13	Software on Remote Monitoring Station	15	17
2.1.14	Device discovery and pairing	16	18
2.2	Methods for testing and evaluation of the device and software	16	19
2.2.1	Software testing	16	20
2.2.2	Flow-sensor assembly calibration tests	16	21
2.2.3	Relationship between ΔP and Q	17	22
2.2.4	Gauge pressure test	19	23
2.2.5	Methods for bench test comparison to commercial medical systems	19	24
2.2.6	Quality control and acceptance testing of final assembled devices	20	25
2.3	Methods for human subject test	20	26
3	Results	21	27
3.1	Flow, pressure and tidal volume agree with commercial systems	21	28
3.2	Human volunteer study	23	29
3.3	External sources of variability	24	30
3.4	Results of production run	24	31
4	Discussion and Conclusion	24	32
5	Acknowledgements	25	33
6	Supporting information	28	34

1 Introduction

Non-invasive ventilation systems, such as continuous positive application of pressure (CPAP) and high flow nasal canula, have emerged as important tools for treating coronavirus disease 2019 (COVID-19) patients who need respiratory support but not intubation [2,3]. Such non-invasive approaches conserve traditional ventilators and improve patient outcomes by reducing or eliminating the need for invasive ventilation [4]. Additionally, COVID-19 patients show surprisingly poor outcomes on

invasive ventilation [5], making non-invasive ventilation (where applicable) uniquely valuable during the pandemic.

Helmet non-invasive ventilation (NIV) is a form of CPAP that uses a clear, polyvinyl chloride (PVC), bubble-like helmet, attached to a soft collar that seals around a patient's neck [6]. The helmet delivers an air-oxygen mixture to the patient at slightly higher than atmospheric pressure, to keep the patient's airways open and the patient oxygenated. Helmet NIV is appealing for pandemic use because it can be run directly from a constant flow of air-oxygen, which is readily available in hospitals [7]. Helmets cost less than 300 USD, are straightforward to manufacture, and have been available during the pandemic, despite supply chain disruptions for other ventilators [1]. Importantly, Helmet NIV is an enclosed system that uses viral filters to protect clinicians and other patients from droplets or aerosolized SARS-CoV-2 viral particles shed by the patient [8]. Helmet NIV is permitted by the FDA for COVID-19 treatment [9].

Despite these advantages, lingering safety concerns have prevented wider adoption of the helmet in the United States. Steady airflow is required to clear the patient's expired CO₂ from accumulating in the helmet [10]. An unexpected drop of airflow caused by a disruption in gas supply or blockage in the circuit could lead to rebreathing of CO₂ and asphyxiation [11]. Therefore, one of the greatest concerns when using helmets is the lack of an effective monitoring and alarm system. Antisuffocation valves can help mitigate this risk, but ultimately a dedicated monitoring and alarm system is needed to employ helmet NIV safely [12]. Partly for this reason, helmets in the US have been primarily restricted to the ICU, where clinician-to-patient ratios are high enough to allow clinicians to observe patients around-the-clock.

We have addressed these safety risks and clinical needs by developing an inexpensive respiratory profile monitoring system, which we call the Princeton Open Ventilator Monitor (POVM), that is based on readily available commercial components and is suitable for rapid fabrication during a pandemic. The system is designed for use in conjunction with non-invasive helmet ventilator systems, such as the COVID-19 Helmet developed by Sea-Long Medical Systems, LLC [1], but is modular and can be used in other systems. The system comprises one or more flow and pressure sensors per helmet, and crucially includes a central station that can be used to remotely monitor up to 20 patients simultaneously. The system reports to clinicians the flow and pressure in the helmet, the patient's equivalent tidal volume (\approx TV), respiratory rate (RR), ratio of inspiratory to expiratory time (I:E ratio), peak inspiratory pressure (PIP), and peak end expiratory pressure (PEEP). The system allows helmets to be deployed more broadly on the wards, since a single clinician can now monitor many patients simultaneously. The respiratory profile provided by the system allows clinicians to track disease progression and informs treatment decisions, including decisions about when to intubate.

We have produced and tested 50 such devices in a matter of weeks at a marginal cost of under 300 USD. Commercial alternatives, like the Philips Respirionics NM3 monitor, offer some of the functionality of our device, but those systems are currently unavailable due in part to disruptions in the medical supply chain. In addition, the commercial alternatives are single-patient devices that cost 20 to 40 times more than our device, thus making them impractical for rapid scale-up and deployment in a pandemic setting. To facilitate local manufacture, we provide detailed parts lists, computer-aided design files, software, and instructions for assembly and testing.

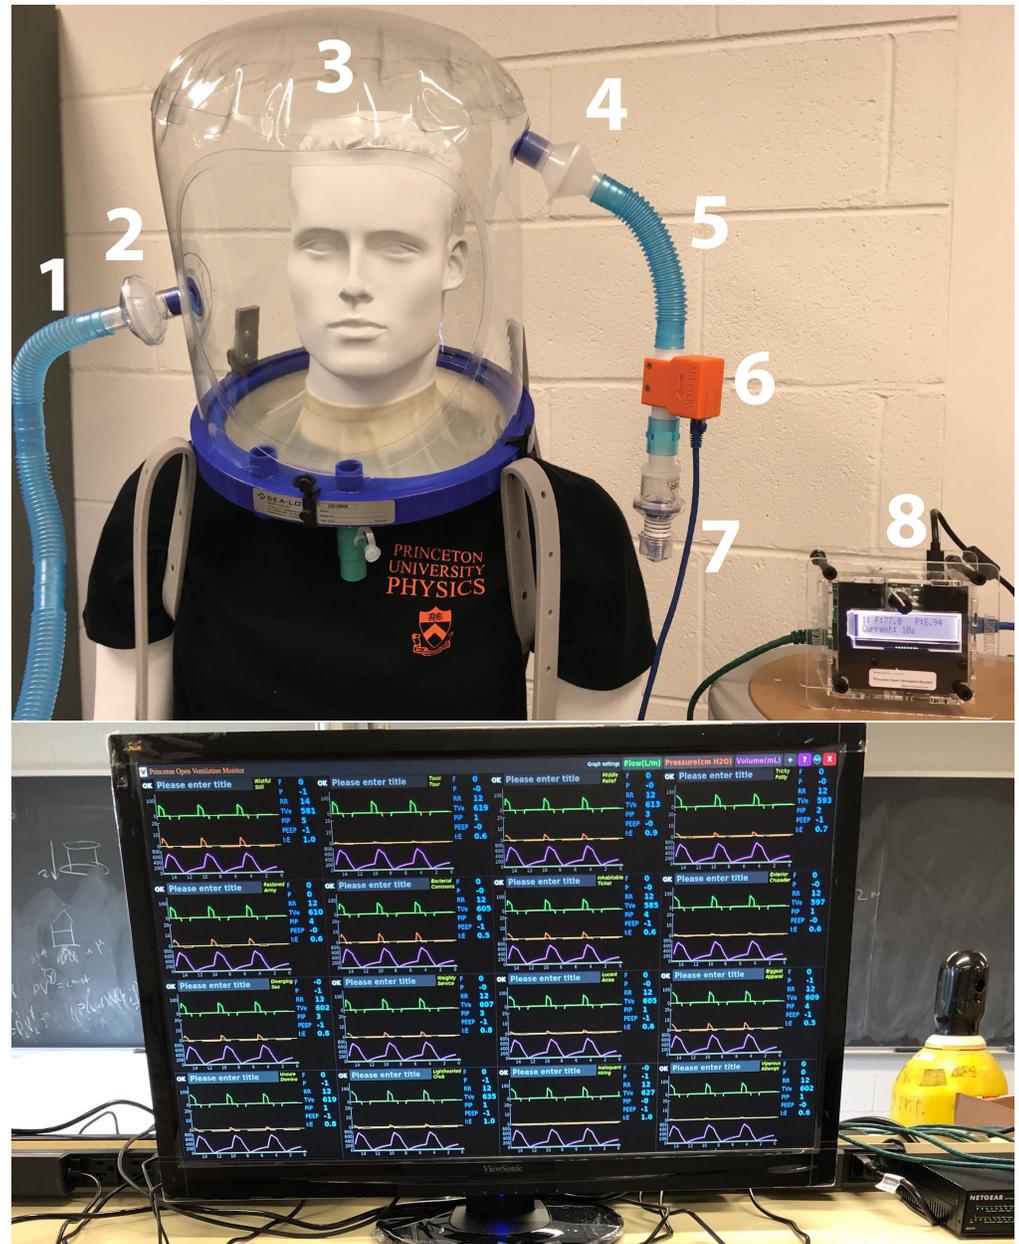


Fig 1. Device monitors respiratory features by measuring pressure and flow of air leaving a helmet. Top: Device is shown installed in a typical circuit. (1) Inspiratory path. (2) Filter. (3) Patient in helmet. (4) Filter. (5) Expiratory path. (6) Flow-sensor assembly. (7) PEEP valve. (8) Interface box. The interface box reports basic respiratory information and announces audible and visual alarms if flow, pressure or respiratory rate cross clinician defined thresholds. Bottom: The remote monitoring station displays flow, pressure, and volume waveforms and clinically relevant quantities including equivalent tidal volume, respiratory rate, I:E ratio, PIP and PEEP from up to 20 devices. Here 16 devices in a test-circuit are being monitored simultaneously.

2 Materials and methods

2.1 Design and construction of the device and software

The Princeton Open Ventilation Monitor device consists of a flow-sensor assembly and an interface box located at the patient's bedside, as well as a multi-patient remote monitoring station, as shown in Fig 1. When used with a helmet for monitoring patients undergoing non-invasive ventilation, the flow-sensor assembly is inserted into the expiratory path. A schematic of the system is shown in Fig 2.

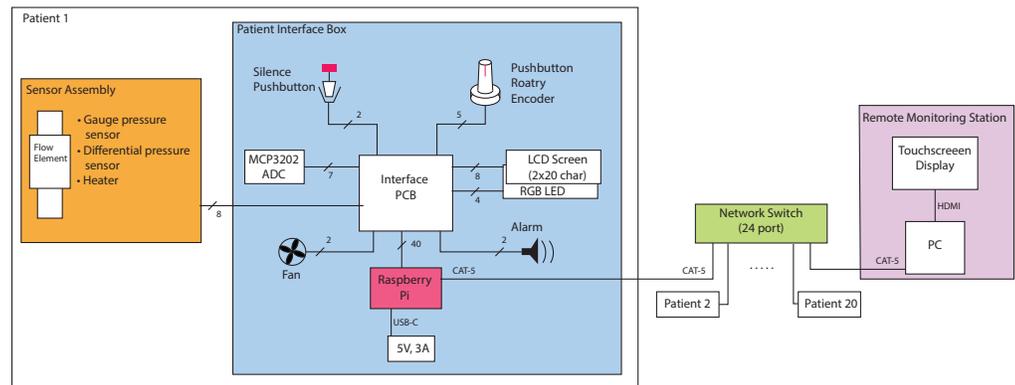


Fig 2. Schematic of system design.

2.1.1 Differential pressure sensor

The flow meter consists of a differential pressure sensor that measures the pressure drop across a flow channel. This pressure drop is then used to calculate flow through the channel. The differential pressure sensor was chosen to be capable of measuring relatively small pressure drops (± 500 Pa) with good dynamic range. Stable operation in a humid environment was also a key consideration, as were cost and availability. The device selected was the Sensirion SDP31 differential pressure sensor. This part works by diverting a small portion of the flow and measuring the change in temperature as the gas passes over a small internal heating element. The temperature change in the gas as it flows through the sensor is related to the mass flow rate of the diverted flow stream, which in turn is related to the pressure difference across the sense ports. Readings from the SDP31 are routed to the controller over an I2C bus at a rate of 120 readings per second.

The SDP31 flow sensor, which comes in a surface mount package, is mounted to a small printed circuit board (PCB, see Supplementary Design File 1) that is fixed to the flow block as shown in Fig. 3. To prevent condensation from forming in the narrow flow-sensor channel, a heater has been added. Heat is generated in a TIP32 transistor housed in a TO-220 package. The TIP32 is mounted on the opposite side of the printed circuit board from the SDP31. Operating the transistor at a power of less than 1 W suffices to raise the temperature in the SDP31 sensor to 45°C , which is above the dew point of the air in the system. The operating temperature is regulated using a pulse-width modulation signal from the readout controller.

2.1.2 Gauge pressure sensor

A second sensor, the MP3V5004G from Freescale Semiconductor, is included to measure the gauge pressure in the flow block. The MP3V5004G measures pressures in the range $0 < p < 3920$ Pa (approximately 0 to 40 cm H_2O) and provides an analog output in the



Fig 3. Flow-sensor assembly. A flow and pressure sensor are exposed to the flow element. The electronic sensor elements are mounted on circuit boards that connect to an external RaspberryPi controller via a standard RJ-45 jack.

range $0.6 < V < 3.0$ V. The analog output is digitized using a Microchip MCP3202 12-bit analog to digital converter. PCB layout and schematics are included in Supplementary Design File 1.

2.1.3 Flow-sensor assembly enclosure

A 3D printed enclosure houses an RJ-45 jack and surrounds the PCBs and flow element. The enclosure was printed from PLA by a commercial fabrication house (3DGence), see Supplementary Design File 1.

2.1.4 Flow element design

The flow element, shown in Fig. 4, is a 19-mm-diameter tube, partially obstructed by a 10 mm thick wall perforated by nineteen 3.175-mm-diameter channels arranged in a hexagonal pattern. One of the channels is partially obstructed and is used to guide a small fraction of the air flow through the flow sensor. The obstruction is only partial, as a portion of the air flow can directly cross the central wall through a small aperture. We tuned the diameter of this aperture in order to match the dynamic range of the sensor to the typical flow rates to be measured. We found that a 1.59-mm-diameter aperture is well adapted to a 500 Pa dynamic range differential pressure sensor.

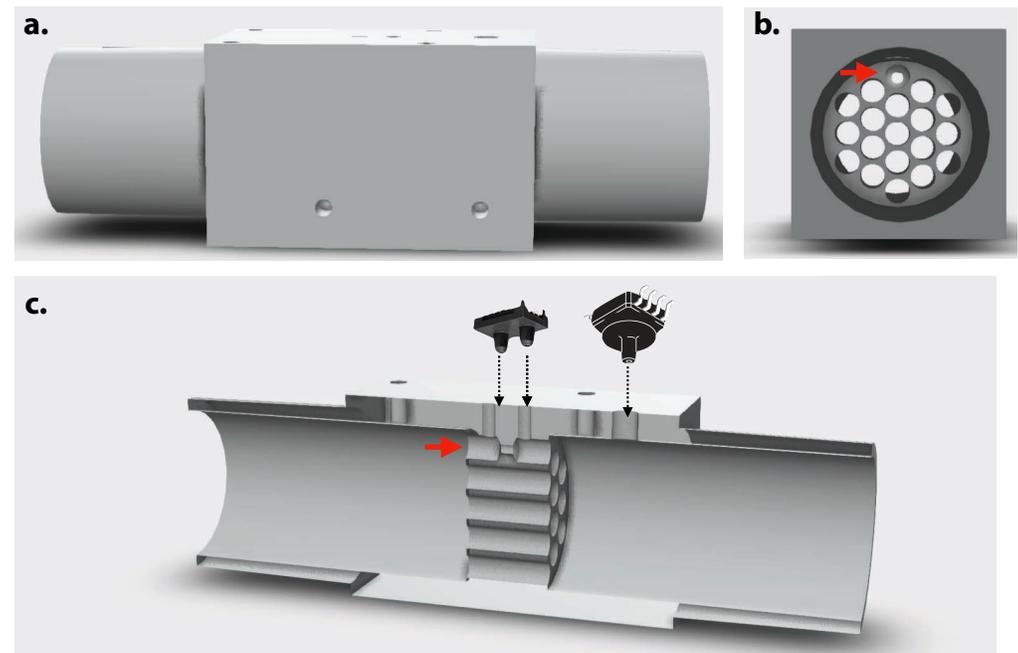


Fig 4. Flow element. The flow element is a 19 mm diameter tube of length 82 mm with an obstruction to create a pressure drop. **a.** Exterior view. **b.** View along the direction of airflow. A honeycomb pattern of holes decreases turbulence. **c.** Interior view. Flow is detected as a pressure drop across a small obstruction in one of the channels (red arrow), as measured by a differential pressure sensor (two pronged device, pictured). An additional sensor measures gauge pressure.

2.1.5 Flow element construction

Flow elements are manufactured on a CNC mill from 2.54×2.54 cm² white Delrin Acetal resin rod (McMaster-Carr part number 8739K92). The material is commonly

used in FDA approved medical devices, meets ASTM D6100 standards, and is readily available. The CAD machining file is available in Supplementary Design File 1. One of the goals of the machining was to make the blocks repeatable enough that the flow versus pressure drop of all blocks could be characterized with a just a few parameters, eliminating the need for individual calibration. It was particularly important to maintain the uniformity and spacing of the eighteen 3.175 mm-diameter holes (for example, see Fig. 10). Drill wander was minimized by using stiff bits and starting the hole by “pecking.” Deburring is also important, as there is a tendency for residual Delrin “hair” to form at the entry and exit of the hole, see 5. This was reduced by touching up the holes with a slightly larger bit. With proper jiggling and after some practice, one block could be machined in under 20 minutes. The yield is about 80%.

Cleaning the blocks to “medical grade” required some care. The primary goal was to remove all residual Delrin hair. This was most successfully done by a commercial company [13], using a frozen CO₂ deburring technique, see Fig. 5. The hair could also be melted away by application of a heat gun. In addition, we had success by mounting a 3.175-mm-diameter “low scratch” tube brush dipped in 80 grit water-based slurry in a hand drill and then polishing each hole. These last two methods require some care with the spinning brush being preferable. After the mechanical processes, blocks were cleaned with soap and water and then sonicated in isopropanol for 10 minutes.

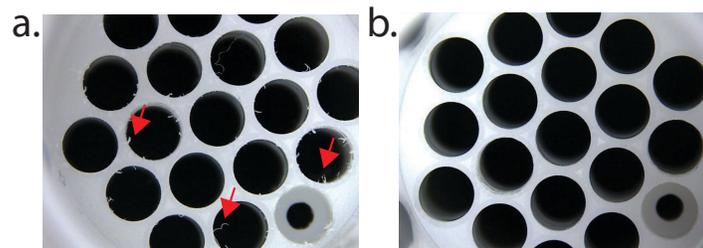


Fig 5. Delrin hair removal. Flow block is show a) before and b) after frozen CO₂ deburring. Note the small delrin hair filaments (red arrows) are gone.

2.1.6 Interface box

The interface box reads out the flow and pressure sensors, provides an alarm system with a human interface, and transmits data to a central remote monitoring station for remote, multi-patient monitoring. The core element of the interface box is the microcontroller, which was chosen to be the Raspberry Pi (RPi) Model 4B, running the Raspbian operating system. The RPi is low in cost, widely available, and enjoys the support of a broad user community. It also provides flexibility and allows for additional sensors and interfaces. Wired and wireless network ethernet interfaces are available on the RPi. The human interface was chosen to be simple and minimal. The setting and management of alarm functions is accommodated by a two-line 20-character-per-line LCD screen with individual Red/Green/Blue LED backlights, a pushbutton rotary encoder, and a red-colored “Silence” pushbutton.

On the sensor readout side, the digital bus interfaces were restricted to comply with 3.3V I2C and SPI and the control elements to 3.3V PWM and GPIO. The data transmission distance of the I2C bus imposes a proximity requirement on the interface box to the flow element of approximately 2 m. The time-base of the readout is derived from the RPi system clock using timer interrupts. The human interface is based on asynchronous signal interrupts.

The SDP31 differential pressure sensor is placed into continuous, average until read, mode. The RPi reads the SDP31 over I2C at a rate of 50 Hz, where each reading is the



Fig 6. Interface box. Raspberry Pi based interface box sits at patient bedside. It provides an alarm system with a human interface and it relays sensor information to the central remote monitoring station.

arithmetic average of 40 individual readings taken at 0.5 ms intervals. The MP3V5004G pressure sensor has an analog output that is digitized by an MCP3202 12-bit ADC from Microchip. The MCP3202 is read out at a rate of 200 Hz over SPI. A 4-sample arithmetic average is formed and recorded at a rate of 50 Hz on the same time-base as the SDP31. The temperature of the SDP31 sensor is read out at 1 Hz.

The preferred operating temperature of the flow block sensors is 45°C to avoid condensation from hot, humid vapor. A heating circuit located on the flow block is adjusted using a convergent “take-back-half” servo algorithm operating on a PWM control signal. The heating circuit has a maximum heating power of 1 W and reaches a maximum temperature of 64°C. A convergent algorithm enables the circuit to operate at a fixed heating value approximately 200 seconds after power-up, with infrequent adjustment. Small shifts of the analog signal from the pressure sensor are observed when the heating power is adjusted.

The addition of a single, customized two-layer PCB was found to be sufficient to make the connections between the interface box hardware and the 40-pin GPIO through a short ribbon cable, see Supplementary Design File 1. This construction made it simple to design an enclosure to host the RPi and interface box components:

- 2x20 character LCD screen with RGB backlight
- pushbutton rotary encoder
- red-colored silence pushbutton
- RJ45 connector for the flow element cable
- alarm buzzer
- cooling fan

Laser cut design files for the enclosure and an assembly video are in Supplementary Design File 1 and Supplementary Video 1.

The initial system software for the RPi is installed on a 32GB microSD memory card with a separate data partition. The disk image is then replicated using multiple USB3 hubs running in parallel to provide identical systems for all interface boxes. The custom enclosure allows the microSD card to be removed and replaced, but to minimize handling of the microSD card, the preferred method of software update is through a local repository on the remote monitoring station.

The alarm is based on a 10 s running average of the flow rate and pressure, as well as the average respiratory rate, calculated using a decaying weight on the last several breaths. Maximum and minimum values are entered directly via the pushbutton rotary knob. Rotation of the knob cycles through alarm setting menus, and a “push and turn” operation is used to change the alarm threshold settings.

The sounding of the alarm generates a PWM volume-controlled buzzer and introduces a red back light to the LCD display. A red silence pushbutton on the interface box may be used to temporarily silence the audible alarm, turning the back light to orange. Silence can be invoked preemptively or post-alarm. The silence duration defaults to 120 s, but can also be manually adjusted upon invoking the silence button. The alarm status and all monitoring data is continuously streamed to the remote monitoring station through a wired ethernet interface. The unique digital identification of the flow element sensors, RPi ethernet MAC address, and a unique interface box naming system allows the remote monitoring station to unambiguously assign alarms to patients, as described in Section 2.1.14.

2.1.7 Remote Monitoring Station

The remote monitoring station is the primary means of providing clinicians with detailed patient information. The remote monitoring station comprises a networked computer and a 27 inch monitor. Both an Intel NUC miniPC and a Raspberry Pi were tested, although any modern PC running Python on any operating system will suffice.

The remote monitoring station provides a simultaneous overview of all patient interface box data and alarm states. For each patient it shows flow, pressure and volume waveforms and a set of derived quantities, including respiratory rate and equivalent tidal volume (see Fig. 1). In a “drilldown” view it provides full-screen plots of the waveforms for a single patient as well as a pressure-volume plot and derived quantities, while still displaying alarm status information on all patients, see Fig. 15.

The remote monitoring station receives data from interface boxes over a wired or wireless network. In a hospital setting we expect an air-gapped wired network to be used. Our system was tested on an ethernet network with a 24 port hub. The monitoring station acts as an always-on appliance and handles the automatic discovery, pairing and logging of patient interface boxes, described in Section 2.1.14. Because the interface boxes lack battery-backed real-time clocks, the remote monitoring station maintains a master time record for recorded data streams.

2.1.8 Bill of materials

A summary level bill of materials is shown in Table 1 and detailed bill of materials is in Supplementary Design File 2.

Subsystem	Item	Cost
Flow-Sensor Assembly	Flow element	\$ 34.10
	SDP31 Sensor	\$ 25.00
	MPV35 Pressure Sensor	\$ 9.29
	Circuit Boards	\$ 27.32
	Enclosure	\$ 12.85
	Flow-Sensor Assembly Total	\$ 108.56
Interface Box	Microcontroller, LVPS, and disk	\$ 76.00
	Display and Encoder Knob	\$ 11.00
	Circuit Board	\$ 30.00
	Enclosure	\$ 8.98
	Ethernet cable	\$ 8.00
	Interface Box Total	\$ 133.98
Remote Monitoring Station	Intel NUC Computer	\$ 359.00
	Touch Screen Monitor	\$ 560.99
	Ethernet Cables	\$ 147.80
	Network Switch	\$ 80.00
	Remote Station Total	\$ 1147.79

Table 1. Summary level bill of materials. Full details are available in Supplementary Design File 2. Circuit board costs include assembly and minor components. Each remote monitoring station accommodates up to 20 patients.

2.1.9 Software overview 247

Software is released as open source under the MIT License and is available at <https://github.com/Princeton-Penn-Vents/princeton-penn-flowmeter>. The software integrates together the data collection and handling in the interface boxes and remote monitoring stations, provides an analysis suite for monitoring and alarm determination, drives the rotary-dial interface of the interface box and implements the graphical user interface of the remote monitoring station. It has been designed to work with minimal computational resources, to be easily portable, to have few external dependencies, and to use standard, well supported, tools when external dependencies are used. 248-256

The software is written entirely in Python and uses only open-source components. On top of a base Python 3.7 installation, additional requirements are currently NumPy, SciPy, PyQt5, pyqtgraph, pyzmq, confuse/pyyaml, zeroconf and the Python Raspberry Pi tools. Qt was selected for both its ability to display real-time graphics with minimal computation resources, and for its portability across computing platforms including for embedded devices. The sections below discuss the software design and implementation of the interface boxes, the remote monitoring station and the communication between them. 257-264

2.1.10 Interface box software 265

The interface box simultaneously runs two python applications called device loop and patient loop. 266-267

The device loop collects raw data from the sensors and converts it to JSON, which it both transmits through a ZeroMQ broadcast to port 5556 and logs to a simple file. A basic JSON “packet” is produced at 50 Hz containing the monotonic clock, pressure, and flow data. The flow data is averaged over four readings taken at 200 Hz. Every second, an extended packet is produced, which includes the device temperature measurement, heater duty cycle, sensor identification number, and the path of the current file used for logging. This application does not read the main configuration file, and is completely stand-alone from the rest of the system. It is started when the device starts up and runs until the device is powered down. 268-276

The second application is the patient loop This application is responsible for reading the ZeroMQ stream from the device loop, powering the on-device LCD display, taking input from the rotary and silence button, applying calibrations, computing the alarm and respiratory analysis, and producing a new 50Hz JSON ZeroMQ broadcast stream for the remote monitoring station to ingest. The new 50 Hz stream is similar to the original stream, except calibrations have been applied. A different type of JSON structure, with rotary information (including the clinician-set alarm thresholds for respiratory rate, flow and pressure), unique per-device MAC address, current on-device timestamp, silence timer, and last interaction time is sent at 1Hz or whenever the rotary is changed. A configuration file in YAML format allows most settings to be changed or overridden without changing the Python code. 277-287

The patient loop runs several threads. There is a collector thread, which reads the input stream and broadcasts to the output stream (the actual reads and writes to the stream are buffered inside the internal ZeroMQ thread). There is also an analysis thread, which runs a lightweight analysis several times per second and a full breath analysis every few seconds (both rates are configurable); this thread can update the display if a displayed value has changed. The rotary and silence button are interrupt driven, though the silencer does launch timers which run in a thread. Locks are used in several places to protect communication of rotary settings and analysis products. 288-295

All communication from the device loop to the patient loop is one-way; and all 296

communication from the patient loop to the remote monitoring station is also one-way. This provides a simpler architecture that was easy to make robust, and ensures that one box can be monitored by multiple stations if needed. One side effect of this is that the realtime clocks on the patient stations are not synchronized, so synchronization is made possible by the timestamps on the remote monitoring station.

Configuration is handled through several YAML files. There is a `processor/config_defaults.yaml` file that stores the defaults for most settings. There is a top-level file `povm.yaml` that stores copies of settings that might be commonly configured in a deployment. And then, when running, the patient software stores a `povm-live.yaml` file in the data directory that keeps track of runtime changes to the alarm limits, so that the alarm settings get restored when the system is restarted. This file is not intended to be edited directly, and is only searched for the rotary settings. If this file were to be corrupted, this is not a failure - the settings are just not restored. A dedicated thread stores the rotary settings every 10 seconds if they have changed.

2.1.11 Calculating derived quantities

Patient breaths and respiratory features are analyzed both locally on the interface box and on the remote monitoring station.

The interface box sensors produce two direct observables as a function of time: absolute pressure (cm-H₂O) and air flow (mL/min). These two quantities are presented with a digital readout on the bedside interface box. The remote monitoring station displays pressure and air flow as continuous time-series waveforms, along with air volume (mL), integrated from the flow. Since the helmet is designed to have a large constant flow through the circuit, the raw integral of flow steadily increases with time. To obtain a timeseries that shows volume fluctuations associated with breathing, we apply a weak high-pass filter (Butterworth critical frequency of 0.004 Hz) that removes the slow drift while preserving the shape on timescales relevant to breathing. See Figure 7-left.

We call our estimate of tidal volume an “equivalent Tidal Volume” because it represents our best estimate of the tidal volume as measured on top of the constant flow through the helmet. To estimate equivalent tidal volume ($\approx TV$), the volume per-breath is needed. Here we use a slightly different scheme to estimate the volume. We subtract off the mean of the flow before integrating and then use a weaker high-pass filter (Butterworth critical frequency of 0.0004 Hz). This improves accuracy in our per-breath estimate of tidal volume when used in an invasive ventilator context, and it offers a good compromise between controlling drift and distorting the shape of each waveform when computing derived quantities.

Many derived quantities of interest, such as equivalent tidal volume ($\approx TV_i$ and $\approx TV_e$), respiratory rate (RR), inspiratory-to-expiratory time ratio (I:E), peak inspiratory pressure (PIP), and peak expiratory endpoint pressure (PEEP), are functions of a breath cycle. Breaths can vary in duration, so we need to identify landmarks on the time series that define points along the breath cycle. We do this by first smoothing the absolute pressure and air flow time-series using locally estimated scatterplot smoothing (LOESS) with a Gaussian kernel that has a 0.2 second standard deviation, then finding the roots of the smoothed curve and its derivative as the four turning points of the breath cycle. These four turning points correspond to the time points during a breath cycle when the lungs are least inflated, mid-inhale, most inflated, and mid-exhale, as shown in Figure 7. Once these turning points have been identified, each quantity of interest is derived by analyzing an extremum or a difference of pressure, volume, or time between these points. Breath-based derived quantities are listed in Table 2.

Breath records are logged and accumulate as a growing list. Cumulative

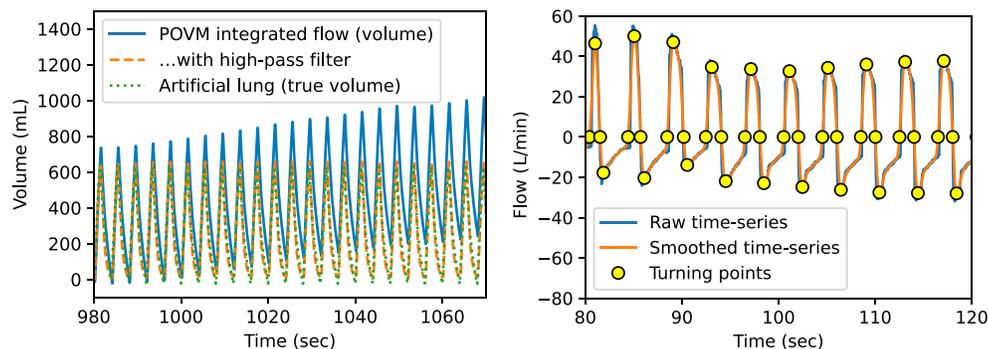


Fig 7. Signal processing to calculate volume and identify breaths. Left: effect of a high-pass filter in removing drift in volume (integral of flow). Recording here is with a test lung. In a helmet circuit the integral drifts because the ventilation system is designed to continuously vent expired air. Right: identifying breath cycles in the time-series data via turning points: roots of smoothed flow and its derivative.

When flow is minimal (inhaling) or maximal (exhaling):

Time (sec)

Flow, as measured (L/min)

Flow, smoothed by 200 ms (L/min)

Change in pressure per unit time, smoothed by 200 ms (cm H₂O)

Compliance: change in smoothed flow per unit smoothed pressure (L/min/cm H₂O)

When volume is minimal (empty lungs) or maximal (full lungs):

Time (sec)

Exhale time: difference between empty and previous full

Pressure, as measured (cm H₂O)

PIP: maximum pressure achieved

PEEP: pressure when lungs are empty

Volume, derived with 0.004 Hz Butterworth frequency (mL)

Inspiratory tidal volume (TV_i): full volume minus previous empty

Expiratory tidal volume (TV_e): empty volume minus previous full

Time between breaths, measured at successive empty lungs (sec)

Respiratory rate: inverse of time between breaths (1/min)

Table 2. Quantities derived for each breath from the original timeseries (flow, pressure, and volume versus time).

measurements are derived from these breath-based measurements as exponentially weighted moving averages (EWMA) of breath records with $\alpha = 0.3$, meaning that it takes approximately $1/0.3 = 3.3$ consecutive breaths for a new trend to emerge. In particular, the respiratory rate, inspiratory and expiratory-endpoint pressures (PIP and PEEP), inspiratory and expiratory tidal volume equivalents ($\approx TV_i$ and $\approx TV_e$), and compliance are all reported as moving averages.

2.1.12 Software detection of Alarms

Alarms are raised when a monitored quantity exceeds a predefined threshold (either too high or too low) and are immediately removed once the quantity returns to a suitable value. Thresholds are set via the rotary dial on each interface box and are passed to the analysis software together with other data. Some alarms are raised on time-window averaged time series, specifically, flow and pressure, while others are raised on EWMA averages of derived quantities, currently only respiratory rate. In addition, “stale data” alarms are raised if any derived quantity is not received after a time window, which could indicate a bad connection.

The analysis functionality is implemented using the rolling buffer functionality described above together with numerical functions from the NumPy and SciPy Python packages. It produces Python dictionaries for breath measurements, cumulative measurements, and alarm states. The analysis is run on both the interface box and the remote monitoring station.

2.1.13 Software on Remote Monitoring Station

The back-end software is shared between the interface box and the remote monitoring station. A Generator is the core class, which handles the analysis thread and most of the data structure. The subclass that does the on-device collection and rebroadcast thread as described in Section 2.1.10 is the Collector. There is also a RemoteGenerator subclass of Generator, and this instead only reads from an interface box broadcast in a thread. This is used by the main Graphical User Interface (GUI) to monitor each connected interface box.

The remote monitoring station GUI is implemented in PyQt [14]. The GUI main window displays a grid of patient stations, each of which shows waveforms for flow, pressure, and volume and derived quantities. The grid can be rearranged via drag and drop. Clinicians can adjust the waveform axes. Clicking the status button shows advanced options, such as removing a disconnected sensor. Alarm status is displayed using both colors and symbols, as described in Table 3.

Clicking any of the grid tiles on the main screen will bring up a “drilldown” screen of information for that interface box (see Fig. 15). The drilldown screen shows a larger set of waveforms and additional derived quantities. It also includes a Volume vs. Pressure plot and additional information about alarm limits and averaging times, the duration elapsed since the last seen breath, the time elapsed since the last interaction at the patient station, and the countdown until the end of any silence mode. An overview status bar shows status and alarm information from all of the other interface boxes.

The device logs both software state and sensor data. Software-related logging is performed by the main thread and at the level of each Generator. Python errors, changes to the rotary knob, alarms and silencing events are all logged. Sensor data is stored in a CSV file, and breath data is stored in JSON files.

Alarm Status	Color	Color Applied To	Status Window Text
No alarm	Black	Status window Title background Figure background	OK
No alarm; Silence mode	Black	Status window Title background Figure background	S
Alarm	Red	Status window Title background Figure background	!
Alarm; Silence Mode	Yellow	Title background	S!
Device Disconnected	Blue	Status window Title background Figure background	D

Table 3. Alarm status GUI scheme for remote monitoring station.

2.1.14 Device discovery and pairing

When patient interface boxes are connected to the network, they are automatically discovered by the remote monitoring station using the python-zeroconf framework. Clinicians then need only to associate the interface box with a given patient by entering the patients name, room number or other identifier.

To reduce the chance of error when associating interface boxes with patients, each interface box has an effectively-unique human-friendly name. The interface box and remote monitoring station share a deterministic naming algorithm that assigns adjective-noun combinations to the patient interface box's ethernet address. Two word names are chosen from 4096 curated adjectives and nouns each (for a total of 16.7 million possible names), using a pseudorandom algorithm seeded by the bits of the ethernet address. Unlike assigned names, which could be lost due to memory or disk errors, and which require careful records to avoid collisions, the device names are fully determined from the hardware ethernet address. Names are displayed on each interface box, and also on the remote monitoring station.

The names of the first 20 devices we produced are given in Table 4.

2.2 Methods for testing and evaluation of the device and software

2.2.1 Software testing

Software was designed with the following testing strategies in place: Static type analysis is performed to ensure basic validity of the code on every commit. Some portions of the code, such as the custom rolling window, have unit tests. A simulation and a data recording playback feature provide tools to perform integration tests for the hardware, analysis, and the clinician GUI. An alternative patient loop that runs using Qt instead of requiring the hardware rotary, screen, and silence button enables testing with the playback even without a physical device present.

2.2.2 Flow-sensor assembly calibration tests

The flow-sensor assembly was calibrated by using a mass flow controller (Alicat Scientific MCR-100SLPM-D). As sketched in Fig. 8c, the inlet port of the mass flow controller (MFC) was connected to compressed air and the outlet port was connected to the flow element by a long segment (1.5 m) of medical tubing. To minimize the effect of

Ethernet address	Device name
dc:a6:32:8f:67:58	cloudy pat
dc:a6:32:8f:67:46	time tooth
dc:a6:32:82:f6:a2	valuable canyon
dc:a6:32:83:73:79	experimental chunk
dc:a6:32:8f:55:64	regal tactile
dc:a6:32:83:73:4f	eastern wardrobe
dc:a6:32:8f:67:8b	untenable text
dc:a6:32:87:49:e7	unfamiliar outset
dc:a6:32:8f:67:0e	abundant creek
dc:a6:32:83:73:b5	ceaseless doubter
dc:a6:32:83:71:84	dressy clavicle
dc:a6:32:80:73:99	glum seeder
dc:a6:32:83:6f:fb	warlike bull
dc:a6:32:87:4a:11	timorous puzzle
dc:a6:32:83:73:3d	such feather
dc:a6:32:87:4c:a2	liable ability
dc:a6:32:84:e4:2e	conformable lily
dc:a6:32:83:71:45	cohesive macro
dc:a6:32:71:76:71	international clock
dc:a6:32:82:99:ba	unhurried camp

Table 4. Hardware ethernet addresses and associated human-readable names for the first 20 patient interface boxes produced.

tubing curvature in the flow calibration (see tests below), the tubing upstream of the flow element was maintained horizontal and straight, and downstream the flow element is directly connected to atmospheric pressure. Successive steps of constant flow Q ranging from 10 L/min to 100 L/min were imposed while the differential pressure ΔP at the flow sensor was measured using our device (see Fig. 8). The details of the sweep were chosen to probe both the steady-state relation between Q and ΔP and temporal dynamics during changes of applied flow. The 0 to 100 L/min flow range accessible to our MFC is clinically relevant to the helmet NIV application and corresponds to differential pressure ΔP readings from 0 to 150 Pa (Fig. 8b). The relation between ΔP (Fig. 8b, red) and Q (Fig. 8a, blue), averaged over many flow elements, gives the calibration in Fig. 9b. Flow elements with obvious manufacturing defects had Q - ΔP relations that deviated (see Fig. 10), and were removed from the global calibration.

Similar calibrations were carried out in high relative humidity environment ($RH = 100\%$) for several hours (Fig. 8d) to verify the resilience of the sensor to humid air and possible condensation. The effect of the presence of a PEEP valve was also tested. The effect of the curvature of the tubing was systematically estimated by running calibration tests while forcing a 90° bend at a distance L of the input of the flow element (Fig. 8e).

2.2.3 Relationship between ΔP and Q

Our device calculates the reported flow Q from a differential pressure measure ΔP across a portion of the flow element. A single global lookup table is used for all devices to convert from ΔP to Q (see Fig. 9). The relationship between Q and ΔP matches existing knowledge from fluid dynamics.

The geometry of the flow element and the Reynolds number determine this relation. The Darcy–Weisbach equation relates flow (here described as a flow velocity V) to

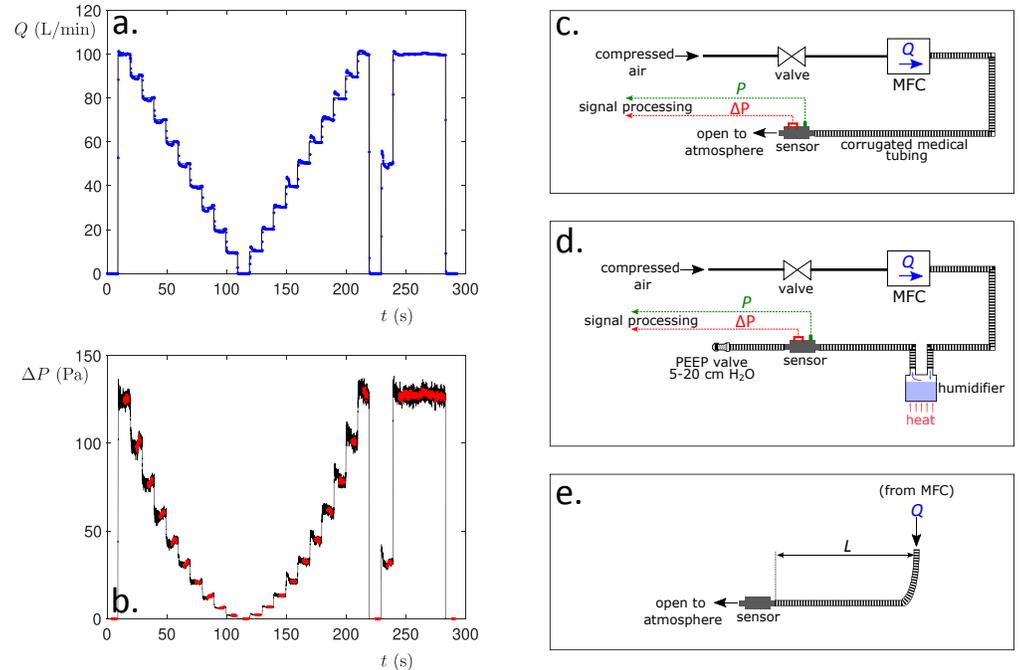


Fig 8. Calibration of the flow sensor. a) Standard flow sweep used to calibrate the flow sensor. Successive steps of flow Q (L/min) were imposed – the blue points are the flow measured in the mass flow controller (MFC). b) The differential pressure ΔP (Pa) as measured by our flow sensor. Black points are raw data measured by the sensor at 50 Hz. Red points are derived from the black points by eliminating the beginning and end of each step and averaging to match the 10-Hz of data from the MFC. c) Sketch of the standard configuration of our calibration setup. The MFC delivers a controlled flow Q to our flow element. d) More tests were performed by adding a humidifier and a PEEP valve. e) The influence of the tubing curvature was systematically tested by forcing a 90° bend at a distance L of the flow block.

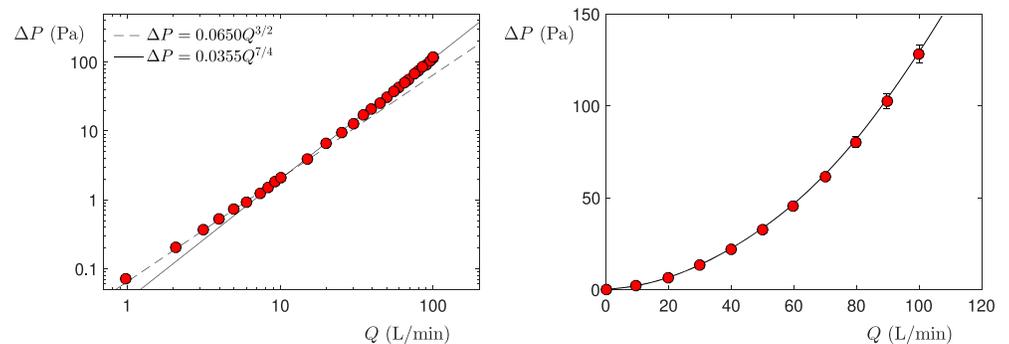


Fig 9. Calibration curve for measuring flow from differential pressure. Left: Measurements of ΔP as function of the imposed flow Q for a single flow element. For low flow (below 15 L/min), $\Delta P \propto Q^{3/2}$ whereas $\Delta P \propto Q^{7/4}$ for larger flows. Right: Averaged measurements of ΔP as a function of the flow Q on 42 flow blocks machined by commercial firm Xometry. The error bars represent standard deviations. The solid line shows the lookup table that is installed in all of our devices.

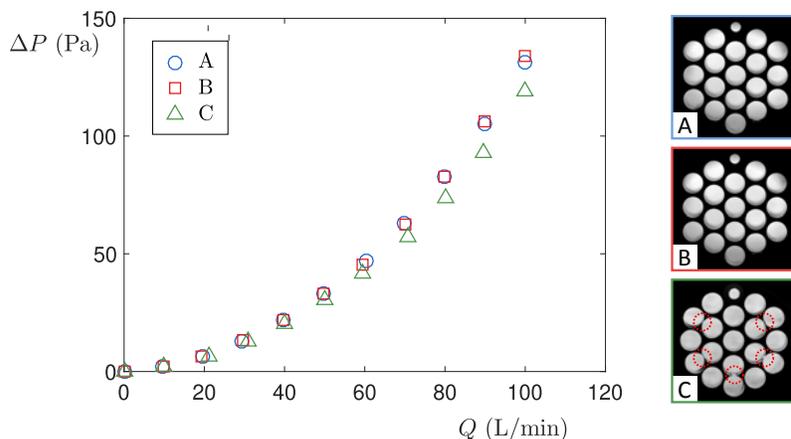


Fig 10. Influence of large machining defects on flow calibration.

Measurements on flow element A and B overlap nicely, but measurements on C are affected by the obvious defects in the honeycomb pattern (red dashed circles in the bottom picture).

differential pressure ΔP ,

$$\frac{\Delta P}{\frac{1}{2}\rho V^2 L/D} = f_D, \quad (1)$$

given the length L and diameter D of a pipe, and a Darcy friction factor f_D that depends on the Reynolds number. At 100 L/min our device operates at an intermediate Reynolds number of $Re \approx 1800$, which is neither in the high ($Re \gg 2000$) nor low ($Re \ll 2000$) turbulence limits. ($Re = VD/\nu$, where the kinematic viscosity of air $\nu \simeq 2 \times 10^{-5} \text{ m}^2/\text{s}$. At $Q = 100 \text{ L/min}$, air through one of the 19 honeycomb channels of diameter $D \simeq 3 \text{ mm}$ travels at approximately $V_1 \simeq 12 \text{ m/s}$.)

The relation between a Darcy friction factor f_D and Reynolds number can be looked up in standard Moody diagrams [15]. For turbulent flow in sufficiently long smooth tubes, $f \propto Re^{-1/4}$, which implies $\Delta P \propto Q^{7/4}$. When our flow is lower, the Reynolds number drops. At a flow of 10 L/min ($Re = O(200)$), the flow in a shorter tube will have a growing boundary layer. In such laminar flow cases we expect $\Delta P \propto Q^{3/2}$. We note that our calibration lookup table follows the $\Delta P \propto Q^{7/4}$ for high flows and the $\Delta P \propto Q^{3/2}$ for lower flows in agreement with expectations (Fig. 9).

2.2.4 Gauge pressure test

The gauge pressure sensor was calibrated by imposing a constant pressure with a pressure controller (Fluigent MFCS-EZ). The sensor saturates at a value around 3000 Pa and exhibits a linear response, see Fig. 11.

2.2.5 Methods for bench test comparison to commercial medical systems

To compare our device against commercial systems, the April 22nd version of our device was tested on a commercial test lung (IngMar ASL5000) driven by a commercial ventilator (GE Avance CS2 Anaesthesia System) in series with a commercial respiratory monitoring system (Philips Respironics NM3). Our device and the commercial monitoring system were both situated in-line on the inspiratory path of the ventilator circuit. Simultaneous recordings of pressure and flow were made from our device, the commercial monitor, and the test lung system. Ground-truth volume information was

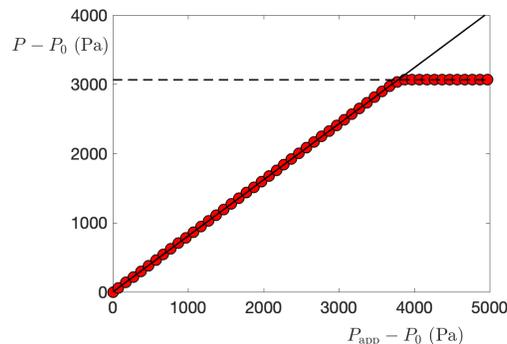


Fig 11. Pressure sensor calibration. The pressure P measured varies linearly (solid line of slope 0.81) with the pressure P_{app} applied with the Fluigent pressure controller and saturates above 3000 Pa.

also recorded from the test lung and precisely time-aligned. Pressure information from our device was also compared to the commercial monitor (NM3). Respiratory rate, PIP and PEEP were also recorded from our device and occasionally compared to the display of the commercial ventilator. Our device was tested under a variety of ventilation modes including Pressure Control Ventilation, Volume Control Ventilation, and Synchronized Intermittent-Mandatory Ventilation with Pressure Control (SIMV PC). The test lung's compliance was varied from 80 mL/cm H₂O, typical for a healthy adult, to 20 mL/cm H₂O which is more typical of a diseased lung. Respiratory rate was set on either the ventilator or the lung to be approximately 15 breaths per minute.

2.2.6 Quality control and acceptance testing of final assembled devices

We developed a quality control protocol that could be used to test finished assembled devices prior to delivering to patients. The test demonstrates that each flow-sensor assembly's combination of specific flow element and sensors gives measured results for flow and pressure within a set range, for example $\pm 10\%$ of actual values. This was done by providing a known flow (using an Aalborg GFC47 mass flow controller) or known pressure (measured at a low flow using a water manometer sampling the flow just upstream of the flow element), recording sensor data with an interface box, analyzing the data using the fixed average calibrations discussed in Sec. 2.2.3 and 2.2.4, and comparing the recorded data to the known values.

2.3 Methods for human subject test

Human subject tests were conducted in accordance with IRB protocol # 12857 approved by Princeton University's Institutional Review Board. Written consent was obtained from a healthy adult volunteer. The subject was placed in a helmet (Sea Long Model PN5404) fed by 80 L/min of medical air (AirGas) configured as shown in Fig 1 with the PEEP valve set to its lowest limit, nominally 5 cmH₂O. A one way valve (Teleflex Model 1644) was added to one of the helmet ports to serve as an anti-suffocation valve. The subject was in a standing position. The subject's spontaneous breathing was recorded on our respiratory monitor device over two 20 min trials as the subject was instructed to breath normally. A water manometer in our air delivery system appeared to contribute to dynamic changes to the volume in the helmet circuit coincident with the subject's breathing, so the manometer was inactivated during portions of the recording.

3 Results

3.1 Flow, pressure and tidal volume agree with commercial systems

We placed a prototype of our device in a circuit with a commercial single-patient respiratory monitoring system (Philips NM3) on a test lung (ASL-5000) driven by a commercial ventilator. We compared recordings made by our device to those made by the test lung and the commercial monitoring system. Recordings from our device show strong agreement with the test lung for flow and volume. See Figs. 12, 13 and 14.

Our device's flow measurement showed closer agreement to the test lung than the commercial monitoring system that was measuring flow simultaneously. There was a scaling factor of about 5% discrepancy between our device and the test lung's measured flow, while the commercial monitor reported flows that differed from our device by a scaling factor of about 15% and an offset of 2 L/min. The commercial respiratory monitor also would occasionally show a flat line for flow or pressure during a breath, while our system captured dynamics of all breaths (see artifact at approximately 158 s in Fig. 12 and as a horizontal spread around zero in Fig. 14).

In measures of pressure, our system had close agreement to the commercial respiratory monitor. We note that our system seemed to avoid quantization artifacts found in the commercial monitoring system at low pressures (see steps in pressure around 80 s in NM3 recording in Fig. 12).

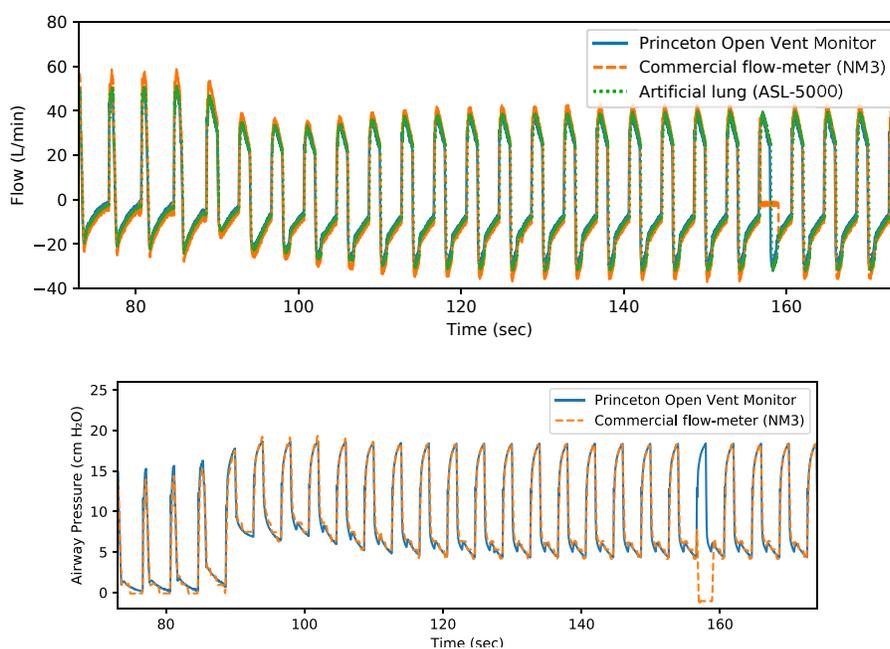


Fig 12. Comparison to commercial monitor. Overlay of flow and pressure measured simultaneously from the Princeton Open Vent Monitor, a commercial respiratory monitor (NM3), and an artificial lung (ASL-5000) driven by a commercial ventilator.

We also compared derived quantities—the equivalent tidal volume, PIP, PEEP, and breathing rate—with those reported by the ventilator at the same time. Our system largely agrees with the ventilator's own settings (see Table 5).

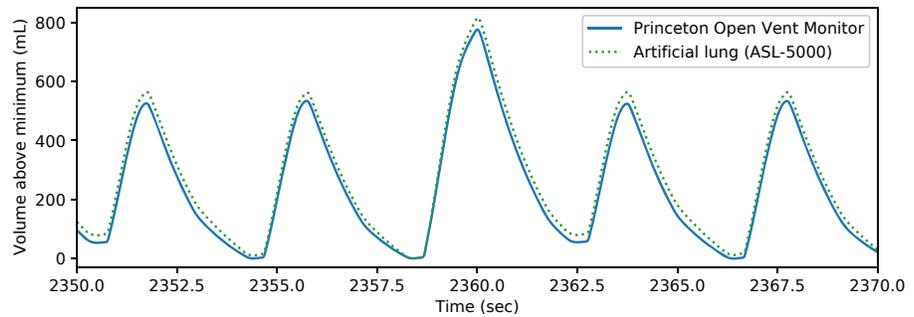


Fig 13. Volume comparison to test lung. Overlay of derived volume from the Princeton Open Vent Monitor (using the high-pass filter with the weaker 0.0004 Hz Butterworth critical frequency) and measured volume from an artificial lung (ASL-5000), driven by a commercial ventilator.

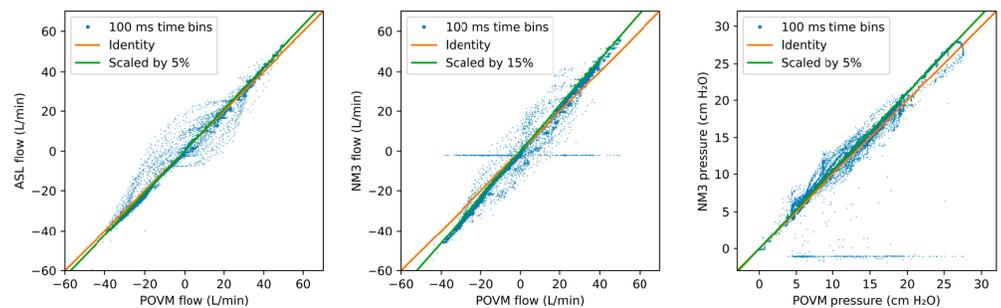


Fig 14. Fine comparison of Princeton Open Ventilation Monitor (POVM), the artificial lung simulator (ASL), and a commercial respiratory monitor (NM3) in matched time-bins (100 ms wide). The commercial monitor occasionally drops to zero and has discretization effects at low pressure. POVM and ASL agree in flow up to a 5% scale factor; POVM and NM3 agree in flow up to a 15% scale factor, and POVM and NM3 agree in pressure up to a 5% scale factor.

	POVM derivation	Ventilator reading	NM3 reading
TV	440 mL	500 mL (setting)	460 mL
PIP	28.0 cm H ₂ O	29 cm H ₂ O	
PEEP	4.5 cm H ₂ O	5 cm H ₂ O	
RR	15.0 bpm	15 bpm	
I:E	0.66		

Table 5. Comparison of derived quantities. Princeton Open Vent Monitor (POVM) derived tidal volume equivalent (\approx TV), peak inspiratory (PIP) and end-expiratory (PEEP) pressure, and breathing rate (RR) compared with contemporaneous readings/settings from the ventilator and readings from the commercial monitor (NM3).

3.2 Human volunteer study

A healthy consenting adult volunteer was placed in a helmet and received medical air at 80 L/min while undergoing monitoring with our device. The device captured reasonable waveforms for flow, pressure and equivalent tidal volume, as shown in Fig. 15. The measured respiratory rate matched that reported by the subject. Measured flow waveforms appear qualitatively similar to previous reports of human subjects in a helmet [10]. Recorded device data stream is available in Supplementary Data Files 1.

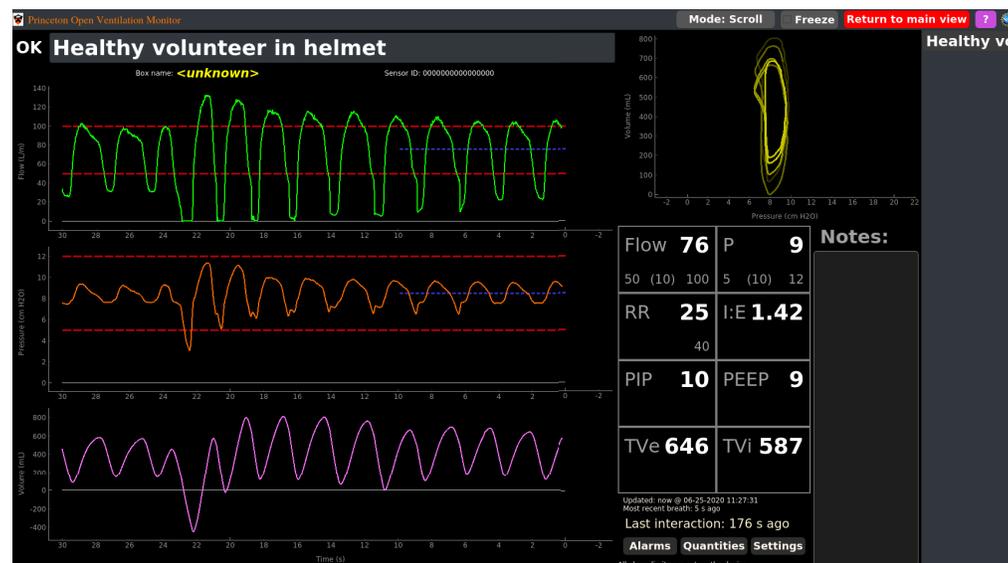
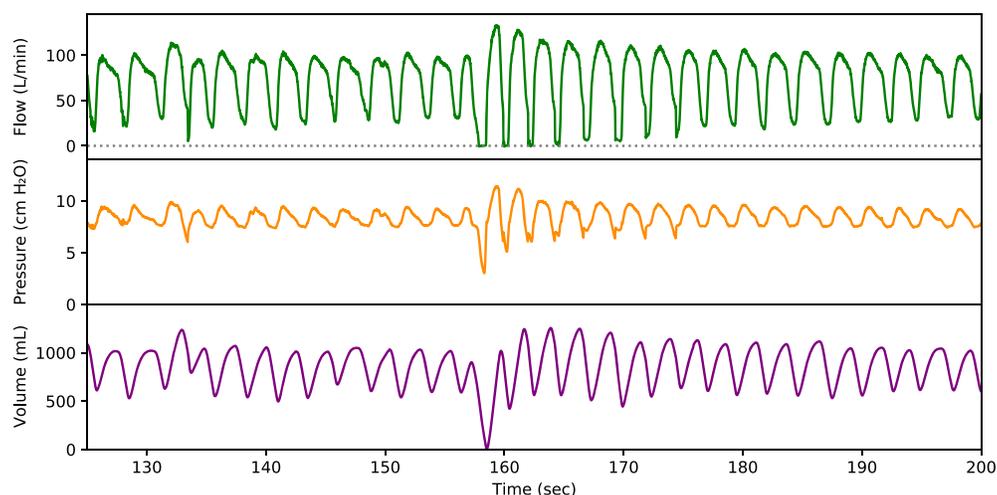


Fig 15. Respiratory profile of healthy human subject in helmet. Top: Pressure, flow and integrated flow for gas on the expiratory path of a helmet worn by a healthy human subject is shown. Positive flow is defined to be gas leaving the helmet through the expiratory path. Bottom: Derived quantities, including tidal volume, respiratory rate, etc are visible on the remote monitoring station (shown here in “drilldown” screen).

3.3 External sources of variability

We tested the influence of relative humidity on our sensors, Fig. 8(d). We challenged our device by measuring 20 hours of continuous flow of saturated humid air ($RH = 100\%$). Both sensors performed even in the presence of significant condensation in the flow element, which suggests that our heater circuit functions effectively. In tests with variable humidity, we observed that the flow Q increased slightly by less than 3% while the humidity rose from 5 to 100%.

Finally, the presence of curvature in tubing has been reported to affect measures of downstream flow [16,17]. This is relevant because our device may dangle off the helmet at an angle, as can be seen in Fig. 1. We introduced a 90° bend into the tubing a variable distance L upstream of our device and observed the change in our device's reported flow Q under a constant imposed flow of 100 L/min, Fig. 8(e) and Fig. 16. We observed a systematic increase in the flow Q reported by our device for bends less than 30 cm upstream of our device. The largest deviations from the imposed flow occurred when the bend was closest to the device (5% for $L = 5$ cm).

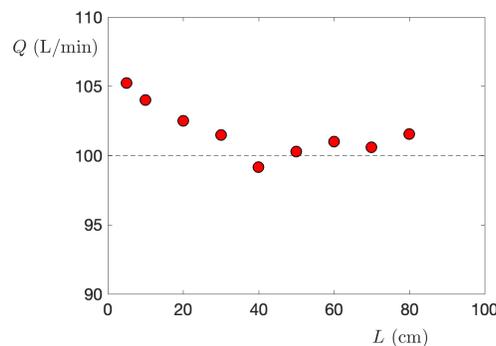


Fig 16. Effect of bent tubing. Flow Q measured by our flow sensor in the presence of a 90° bend positioned a distance L upstream of the flow block. A 100 L/min flow was imposed by a Mass Flow Controller.

3.4 Results of production run

We completed a production run of 50 devices. Producing devices at scale posed additional challenges regarding calibration and quality control. In particular manufacturing defects in the flow element altered the relationship between the measured pressure difference ΔP and the flow Q across the device, as shown in Fig. 10. Initial determination of ΔP vs. Q for each flow element with a fixed sensor array allowed us to eliminate outliers from the determination of the average calibration. After each flow element is integrated into its final flow-sensor assembly, it will undergo the acceptance testing described in Sec. 2.2.6. An example test result is shown in Fig. 17.

4 Discussion and Conclusion

We designed and built an inexpensive respiratory monitoring device from conception to production in a six week period coincident with the initial peak of the COVID-19 pandemic in our region. Our device shows similar, and in some aspects better, agreement with a test lung than a commercial monitoring system that is currently unavailable. Measurements of a healthy volunteer in a helmet taken by our device are reasonable and match prior reports. Our device has the potential to improve patient

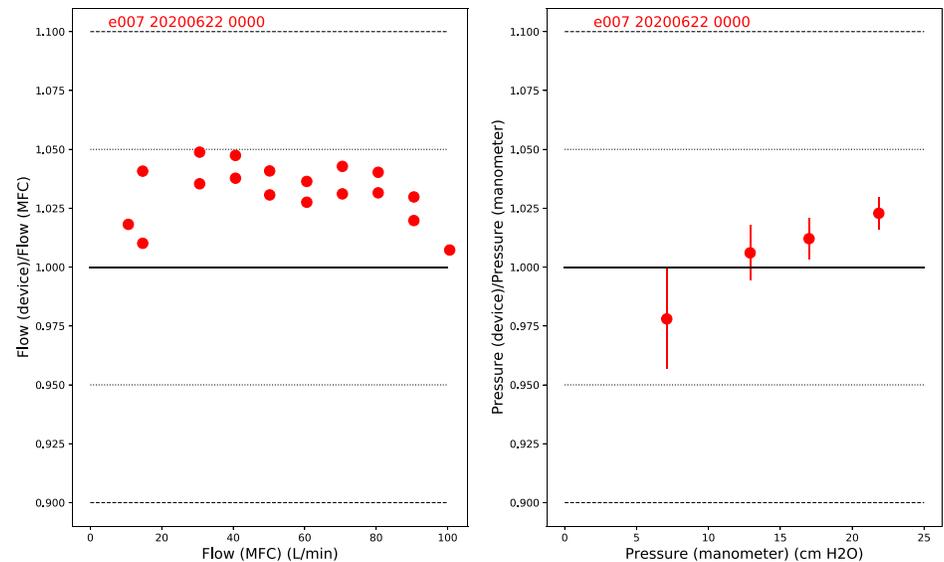


Fig 17. Results of acceptance testing of a final flow-sensor assembly. Left: comparison of calibrated flow reported by the device under a test to flow set by a Mass Flow Controller. Right: comparison of calibrated pressure reported by the device under test to pressure measured with a water manometer. Red text uniquely identifies the flow-sensor assembly.

outcomes and to mitigate well-documented safety concerns regarding CO₂ re-breathing that have prevented the helmet's widespread adoption. Our device also introduces the ability to monitor many patients simultaneously. Importantly, our device uses off-the-shelf components that have shown themselves to be robust to supply chain disruptions. The device is well suited for local manufacture during a pandemic. The device is designed for helmet systems but we have also successfully tested it in traditional invasive ventilator settings. We currently are seeking FDA emergency use authorization.

For future versions of the device we are exploring adding CO₂- and O₂-sensors, and possibly adding hardware, such as an EEPROM, with each sensor board to enable per-sensor calibration curves.

5 Acknowledgements

This project is supported by Princeton University, including generous support by the Provost's Office, and by National Science Foundation grants OAC-1836650, PHY-2031509 and IOS-1845137. Any opinions, findings, conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation or Princeton University.

We acknowledge contributions from all members of the Princeton Open Ventilation Monitor Collaboration, including Lauren Callahan, Matthew Creamer, Sophie Dvali, Alexander E. Gaillard, Alex Glaser, Bert G. Harrop, Darryl Johnson, Julianne M. LaChance, Theodore H. Lewis III, Martina Macakova, Mala Murthy, Jonathan Prevost, P. Dylan Rich, William R. Sands and Lisa Scalice.

We would like to thank Dr. Ronah Harris for shooting and editing the supplementary video. We are grateful for productive conversations with Dr. Jacob Brenner, Dr. James Weimer, Dr. Maurizio Cereda, Dr. Barry Fuchs, Mike Frazer and Victoria Berehnolz all

of Penn Health Tech. We thank Sebastian Seung and his volunteers for feedback. We thank Theodor Brasoveanu and his students at the Princeton Day School for supplies. We thank Meghan Testerman of Princeton University Libraries for assistance.

592
593
594

References

1. Sea-Long Medical Systems website;. <https://www.sea-long.com/>.
2. Grasselli G, Pesenti A, Cecconi M. Critical Care Utilization for the COVID-19 Outbreak in Lombardy, Italy: Early Experience and Forecast During an Emergency Response. *JAMA*. 2020;doi:10.1001/jama.2020.4031.
3. Phua J, Weng L, Ling L, Egi M, Lim CM, Divatia JV, et al. Intensive Care Management of Coronavirus Disease 2019 (COVID-19): Challenges and Recommendations. *The Lancet Respiratory Medicine*. 2020;8(5):506–517. doi:10.1016/S2213-2600(20)30161-2.
4. Patel BK, Wolfe KS, Pohlman AS, Hall JB, Kress JP. Effect of Noninvasive Ventilation Delivered by Helmet vs Face Mask on the Rate of Endotracheal Intubation in Patients With Acute Respiratory Distress Syndrome: A Randomized Clinical Trial. *JAMA*. 2016;315(22):2435–2441. doi:10.1001/jama.2016.6338.
5. Richardson S, Hirsch JS, Narasimhan M, Crawford JM, McGinn T, Davidson KW, et al. Presenting Characteristics, Comorbidities, and Outcomes Among 5700 Patients Hospitalized With COVID-19 in the New York City Area. *JAMA*. 2020;doi:10.1001/jama.2020.6775.
6. Antonelli M, Conti G, Pelosi P, Gregoretti C, Pennisi MA, Costa R, et al. New Treatment of Acute Hypoxemic Respiratory Failure: Noninvasive Pressure Support Ventilation Delivered by Helmet—a Pilot Controlled Trial. *Critical Care Medicine*. 2002;30(3):602–608. doi:10.1097/00003246-200203000-00019.
7. Savickaite A. Video: Non-Invasive CPAP by Helmet Setup COVID-19, Maurizio Franco Cereda, MD; 2020. <https://www.helmetbasedventilation.com/post/video-non-invasive-cpap-by-helmet-setup-covid-19-maurizio-franco-cereda-md>.
8. Cabrini L, Landoni G, Zangrillo A. Minimise Nosocomial Spread of 2019-nCoV When Treating Acute Respiratory Failure. *The Lancet*. 2020;395(10225):685. doi:10.1016/S0140-6736(20)30359-7.
9. U S Food and Drug Administration. Enforcement Policy for Ventilators and Accessories and Other Respiratory Devices During the Coronavirus Disease 2019 (COVID-19) Public Health Emergency; Sun, 03/22/2020 - 12:57. <https://www.fda.gov/regulatory-information/search-fda-guidance-documents/enforcement-policy-ventilators-and-accessories-and-other-respiratory-devices-during-coronavirus>.
10. Taccone P, Hess D, Caironi P, Bigatello LM. Continuous Positive Airway Pressure Delivered with a "Helmet": Effects on Carbon Dioxide Rebreathing. *Critical Care Medicine*. 2004;32(10):2090–2096. doi:10.1097/01.ccm.0000142577.63316.c0.
11. Patroniti N, Saini M, Zanella A, Isgro S, Pesenti A. Danger of Helmet Continuous Positive Airway Pressure during Failure of Fresh Gas Source Supply. *Intensive Care Medicine*. 2007;33(1):153–157. doi:10.1007/s00134-006-0446-5.

12. Milan M, Zanella A, Isgrò S, Deab SAEAES, Magni F, Pesenti A, et al. Performance of Different Continuous Positive Airway Pressure Helmets Equipped with Safety Valves during Failure of Fresh Gas Supply. *Intensive Care Medicine*. 2011;37(6):1031–1035. doi:10.1007/s00134-011-2207-3.
13. Cryogenic Institute of New England, Inc.;
14. Riverbank Computing. PyQt5; 2018. Available from: <https://riverbankcomputing.com/software/pyqt/intro>.
15. Moody F. Friction factors for pipe flow;.
16. Dean WR. LXXII. The stream-line motion of fluid in a curved pipe (Second paper). *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*. 1928;5(30):673–695. doi:10.1080/14786440408564513.
17. Hellström LHO, Zlatinov MB, Cao G, Smits AJ. Turbulent pipe flow downstream of a 90° bend. *Journal of Fluid Mechanics*. 2013;735:R7. doi:10.1017/jfm.2013.534.

6 Supporting information

595

Supplementary Video 1. Instructional video showing interface box assembly. Video available online at <https://vimeo.com/433023739>.

596

597

Supplementary Design File 1. Compressed design files. A zip file, `SuppDesignFiles.zip` (1.8 MB), containing computer aided design files for laser cutting, 3D printing and CNC machining the interface box enclosure, flow-sensor assembly cap and flow element. Contents are described in Table S1.

598

599

600

601

Supplementary Design File 2. Detailed Bill of Materials. Spreadsheet in excel format, `BillOfMaterials.xlsx` (17 kB), containing detailed pricing and parts numbers for all components.

602

603

604

Supplementary Data Files 1. Human subject recording. `hum_subj.zip` (3.6 MB) Recording of a healthy human subject volunteer breathing normally in a helmet for approximately two 20 min sessions. A portion of this recording is shown in Fig. 15. The compressed zip file containing two files within it: the first, `device_output.json`, is a data stream recorded from the Princeton Open Ventilator Monitor in JSON format that includes raw sensor readings. The second, `hum_subj_waveforms.csv`, is a corresponding CSV file that contains recorded pressure, flow and volume over time in real-world units. The CSV file contains two volumes corresponding to the more- or less-aggressively high-passed filtered volumes discussed in the text. Note the volumes have undefined offsets.

605

606

607

608

609

610

611

612

613

614

Filename	FileType	Target	Component	Note
InterfaceBoxLaserCutDesign/	Folder		Interface Box	
Vent_boxLaserCut.txt	Text			Laser cutter model information
ReadyToLaserCutFile_3X_ElectronicEnclosures_1.cdr	CorelDraw	Laser Cutter	Interface Box Enclosure	3 enclosures per cut
ReadyToLaserCutFile_3X_ElectronicEnclosures.pdf	PDF			1 enclosure per cut
ReadyToLaserCutFile_1X_ElectronicEnclosures.cdr	CorelDraw			
ReadyToLaserCutFile_1X_ElectronicEnclosures.pdf	PDF			
Case_InnerLayer.pdf	PDF			
Case_InnerLayer.cdr	CorelDraw			
FlowSensorAssemblyCap3DPrintDesign/	Folder		Flow-sensor Assembly Cap	
Cover_FlowBlock8.mm.stl	STL	3D Printer		
Cap_Version8.PNG	PNG			Render for visualization
FlowElementDesign/	Folder		Flow element	
IceHoney500_v2.stp	STEP	CNC Machine		
IceHoney500_v2.pdf	PDF	Human		
IceHoney500_v2.ipt	Autodesk Inventor			
PressureBoard/	Folder		Gauge Pressure Sensor PCB	
MP3V_V3-2M-GERBER.zip	ZIP	Commercial Fab House		GERBER files
MP3_V3-Loading_Top.pdf	PDF	Human		Layout for Assembly (Top View)
MP3_V3-Loading_Bot.pdf	PDF	Human		Layout for Assembly (Bottom View)
MP3_V3-2SCH.pdf	PDF	Human		Schematic
MP3_V3-2PCB.pdf	PDF	Human		Layout
MP3_V3-2M.sch	SCH	Layout Software		Schematic
MP3_V3-2M.pcb	PCB	Layout Software		Layout
MP3_V3-2-BOM.txt	Text	Human		Bill of Materials
InterfaceBoard/	Folder		Interface Board PCB	
LCD_Interface4.sch	SCH	Layout Software		Schematic
LCD_Interface4-SCH.pdf	PDF	Human		Schematic
LCD_Interface-BOM.txt	Text	Human		Bill of Materials
LCD_Interface-4.pcb	PCB	Layout Software		Layout
LCD_Interface-2_PCB.pdf	PDF	Human		Layout
LCD_Interface.Loading-Top.pdf	PDF	Human		Layout for Assembly (Top View)
LCD_Interface.Loading-Bot.pdf	PDF	Human		Layout for Assembly (Bottom View)
Interface.2-GERBER.zip	ZIP	Commercial Fab House		GERBER Files
FlowBoard/	Folder		Flow Sensor PBC	
SDP3x-Sensor_BOM.txt	Text	Human		Bill of Materials
SDP3x_6aM-SensorSCH.pdf	PDF	Human		Schematic
SDP3x_6aM-SensorPCB.pdf	PDF	Human		Layout
SDP3x_6aM-Sensor.sch	SCH	Layout Software		Schematic
SDP3x_6aM-Sensor.pcb	PCB	Layout Software		Layout
SDP3x_6aM-LoadingTop.pdf	PDF	Human		Layout for Assembly (Top View)
SDP3x_6aM-LoadingBot.pdf	PDF	Human		Layout for Assembly (Bottom View)
SDP3x_6aM-GERBER.zip	ZIP	Commercial Fab House		GERBER Files

Table S1. Contents of Supplementary Design File 1: SuppDesignFiles.zip.