

Large-Scale 3D Printing With A Cable-Suspended Robot

Eric Barnett*, Clément Gosselin

*Département de génie mécanique, Université Laval
Québec, QC, G1V 0A6, Canada*

Email: eric.barnett.1@ulaval.ca, gosselin@gmc.ulaval.ca

Abstract

Although additive manufacturing (AM) is now a well-established industry, very few large-scale AM systems have been developed. Here, a large-scale 3D printer is introduced, which uses a six-degree-of-freedom cable-suspended robot for positioning, with polyurethane foam as the object material and shaving foam as the support material. Cable-positioning systems provide large ranges of motion and cables can be compactly wound on spools, making them less expensive, much lighter, more transportable, and more easily reconfigurable, compared to the gantry-type positioning systems traditionally used in 3D printing. The 3D foam printer performance is demonstrated through the construction of a 2.16-m-high statue of Sir Wilfrid Laurier, the Seventh Prime Minister of Canada, at an accuracy of approximately 1 cm, which requires 38 hours of printing time. The system advantages and drawbacks are then discussed, and novel features such as unique support techniques and geometric feedback are highlighted. Finally, a description of the planned system modifications is provided.

Keywords: additive manufacturing, cable-suspended parallel robot, large-scale 3D printing, polyurethane foam

1. Introduction

In recent years, additive manufacturing (AM) has transitioned from a niche industry to a truly mainstream technology: a broad range of commercial 3D printers are now available for purchase, from inexpensive do-it-yourself desktop machines to industrial printers that cost millions of dollars. Large-scale additive manufacturing, however, is a relatively new field of research, currently being investigated by only a few laboratories worldwide. The main application of large-scale AM is construction automation: the ability to construct building-sized objects rapidly, with little or no manual intervention. Additive manufacturing can also be used to produce highly customized large-scale features, with considerably less effort than that required when using traditional construction techniques.

Large-scale AM is a natural extension of 3D-printing technology. Here, the term large-scale is used to denote AM systems capable of building objects approximately

one cubic metre or larger in volume. The advantages of AM, compared to traditional manufacturing techniques, are well-established. These include the ability to construct parts of almost unlimited 3D geometry, a streamlining of the design-to-construction process, and greater automation during part construction.

These advantages are still present at the large scale, though many new drawbacks and challenges also exist. Most small- and medium-scale AM materials are expensive, costing tens to thousands of dollars per kilogram, making them cost-prohibitive at the large scale. A major obstacle that must therefore be overcome in the development of large-scale AM is the identification and/or synthesis of inexpensive and controllable materials for construction.

Earth-based materials, such as concrete and adobe, are logical choices, since they typically cost only a few cents per kilogram. Additionally, these materials are often *extrudable*, making them suitable for controlled material deposition. However, special mixes are normally needed to obtain the material properties needed for additive manufacturing. For example, regular concrete would be unsuitable for AM for several reasons. Firstly, once the concrete is mixed, there is a limited time available during which it must be deposited; mixed concrete cannot be stored in a tank in a ready-to-deposit state. Secondly, regular concrete will not stay in place when multiple layers are deposited—it will slump under its own weight.

Four earth-based large-scale AM systems have been reported in the literature. A brief description of these systems is provided here, a detailed comparison can be found

*Corresponding author

Abbreviations

AM	additive manufacturing
DOF	degree of freedom
EE	end-effector
CM	center of mass
PID	proportional-integral-derivative
SMF	surface mapping feedback

in [1]. Khosnevis et al. introduced a method called Contour Crafting, in which concrete or ceramic paste is printed and then smoothed using trowels attached to the printer head [2]. This technique has been validated through the construction of small-scale structures and housing-scale walls [3]. Buswell et al. also developed a concrete printing system [4] and conducted an extensive investigation to optimize the concrete mixture [5]. Both of these techniques are optimized for the construction industry, and most of the structures produced are walls or ruled surfaces.

Two separate techniques have been introduced which selectively bind layers of a dry bed of material, rather than depositing a wet concrete mix. Pegna described a concept for a large-scale concrete-printing system [6], whereby Portland cement is deposited over a matrix of silica material and then steam is used to selectively bind it. D-Shape was introduced by Dini [7], which selectively binds a sand bed with a chlorine-based liquid binder. These techniques are similar in concept to selective laser sintering (SLS).

A significant advantage of the D-Shape system is that it does not limit the possible 3D geometries. However, based on the information available on the project website, it seems that the printer can only produce straight-line paths, and the resulting structure is quite rough after printing has been completed and the unbound dry powder has been removed. A significant amount of manual labor is required to smooth the part following construction. It would seem that this process would destroy a lot of part detail.

A gantry-type system is used for positioning the deposition tool(s) for all of these large-scale AM systems. While several of them have been designed for relatively easy disassembly and installation on construction sites, the reliance on large linear axes for positioning limits the construction scale. Additionally, the workspace of these systems is difficult to reconfigure.

These two disadvantages are addressed in a large-scale AM concept introduced in [8], where the Contour Crafting system is modified to use a cable-driven robot for tool positioning. Extremely large workspaces are possible with cable-driven robotic systems, as is evident with the Spider-Cam and SkyCam [9] systems currently in use at nearly all major sports stadiums. Additionally, such systems are considerably more transportable and reconfigurable than traditional gantry-type systems.

Here, a large-scale 3D foam printing system is introduced, which uses a six degree-of-freedom (DOF) cable-suspended parallel mechanism for positioning. Polyurethane foam is used as a proof-of-concept material, since it is light, relatively inexpensive, and can be stored in a ready-to-deposit state for months.

This system has many novel features. Firstly, the use of a cable robot for positioning makes the system both highly transportable and reconfigurable. A *cable-suspended* robot architecture simplifies the problem of cable interference, at the cost of limiting the achievable workspace with all cables in tension.

The system can construct any 3D geometry, thanks

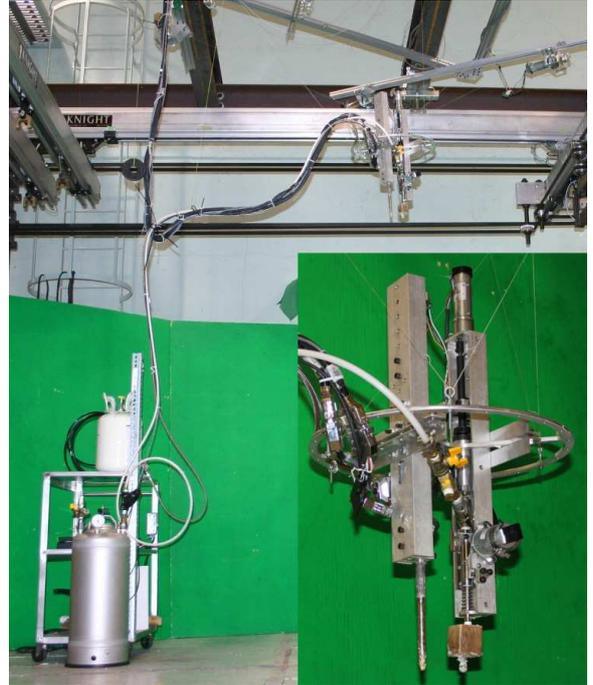


Figure 1: The cable-suspended 3D foam printer.

to a full support system and a separate support material. Additionally, geometric feedback [10] is used to detect and correct geometry errors during printing. This critical system feature greatly enhances the printing accuracy and also permits the use of printing materials that would otherwise be uncontrollable.

Of course, this system comes with a few drawbacks, compared to the traditional gantry-type system. System accuracy is obviously better with the latter system, though this factor becomes less significant as the scale is increased. As with any cable system, careful design and path planning is needed to avoid cable interference. Electrical and fluid lines attached to the end-effector (EE) complicate this problem, and also apply an external force to the EE, which varies with the position in the workspace.

In this paper, the various subsystems and capabilities of the 3D foam printer are first described. Then, a case study is presented, showing the steps involved in building a life-size polyurethane foam statue of Sir Wilfrid Laurier, the Seventh Prime Minister of Canada. Finally, the capabilities of the system and the planned improvements are discussed.

2. Experimental setup

The cable-suspended 3D foam printer, shown in Fig. 1, is composed of several subsystems, which are described in the following subsections.

2.1. Spatial Positioning and Orientation

Once a cable-driven system has been selected for positioning, several different architectures can be considered.

Obviously, such a mechanism must have a minimum of three degrees of freedom (DOF) to construct 3D objects. In addition to this, however, orientation control is needed; at a minimum the EE orientation must remain fixed in an inertial frame. For traditional 3D printers composed of linear axes, three positional DOFs are sufficient, since the EE is rigidly attached to the distal link and its orientation never changes. For a cable-driven mechanism, however, careful design is needed to ensure sufficient orientation control.

A first broad distinction exists between *cable-suspended* and *fully constrained* cable mechanisms. A cable-suspended mechanism, by definition, suspends the EE by a number of cables attached from above, using only the EE weight to maintain cable tension. Fully constrained cable mechanisms include cables attached to the EE from below as well, such that massless, or near-massless, EEs can be used. Additionally, fully constrained mechanisms can better resist off-center-of-mass forces, yield a larger workspace for the same spool footprint, and are dynamically stable at higher accelerations. Here, the term spool footprint is used to denote the area formed when projecting the spool location points downward. The usable or static workspace of a cable-suspended system is the volume within which the EE operation point can be positioned, while maintaining all cables in tension.

In [8], the fully constrained option is selected, most likely for the reasons listed above. The cable attachment points below the EE are made *movable*, to prevent interference between the lower cables and the part under construction. For these reasons, the actuator mounting configuration must be considerably more complicated, less reconfigurable, and less transportable, compared to the cable-suspended option.

Two significant advantages exist for cable-suspended mechanisms. Firstly, they are simpler, because fewer motors and cables are needed for positioning, and all actuators can be easily mounted to an overhead structure. Secondly, cable interference is a much smaller problem, and collisions between cables and objects below the EE are not an issue.

Given the advantages and drawbacks described above, the feasibility of the simpler, cable-suspended option was first considered. Three potential cable-suspended architectures were identified, as depicted in Fig. 2. The usable or static workspace boundaries shown there define the volume within which cable tensions are always positive. They are found by first determining the cable with the minimum tension, at every point in a cloud expected to contain the static workspace, according to Eq. (20), with $\dot{\mathbf{p}} = \mathbf{0}$. The interpolated surface within the point cloud that defines zero minimum cable tension also defines the static workspace.

The first option is similar to the SkyCam system, where cables are only used to control 3D positioning of the EE. As can be seen in Fig. 2a, the usable workspace is defined by the extrusion downward of the area formed by the spool

locations. In this case, the simplest configuration of three motors and three cables is used. A more useful workspace could be obtained by using a four-motor-four-cable system, with the spools defining a rectangular area.

This design works well as long as all forces on the EE pass through or very near its center of mass (CM). In practice, for an experimental system, this constraint is extremely restrictive. For example, for the 3D printing application, it is desirable to supply the print materials from a large stationary reservoir, rather than install the reservoir on the EE. It is also preferable to send electrical power and signals through cables to the EE, rather than install a power source and a signal transceiver on the EE. A large cable bundle is therefore needed, which connects to the EE and provides enough slack such that the EE can move everywhere in the workspace. The weight of the cable bundle will apply a force to the EE, which varies throughout the workspace. Some static balancing can be used to limit this force, but completely eliminating it would be very difficult. Maintaining a constant CM during the development of the various EE subsystems is also a significant design restriction: every new component added to the EE would need to be carefully balanced. Most likely, even if the cable forces pass through the CM, two-DOF yaw-pitch orientation control would be needed to keep the tools pointing down at all times.

The second option is similar to the first, except each motor drives two parallel cables. Additionally, the cables do not intersect at the EE, but are configured as shown in Fig. 2b. The advantage of this configuration is that it automatically maintains the EE orientation, while only relying on three motors for positioning. Additionally, it can resist off-CM forces to a certain degree. The main disadvantage is that the workspace is small, compared to the spool footprint, as shown in Fig. 3a.

The third option, shown in Fig. 2c, is a six-motor, six-cable system which allows for full positioning and orientation control. It is important to note that this configuration does not simply consist of that in Fig. 2b with the spools further apart and three additional motors. In addition to these changes, as shown in Fig. 3b, actuators that are *opposite* each other on the hexagonal mounting structure connect to the *same* point on the end effector. Also, this point is 60° offset from the point on the outer rim of the EE which lies in the vertical plane passing through these actuators.

This configuration offers a larger workspace compared to the second option, is more resistant to off-CM forces, and is therefore more *dynamically stable*. However, with twice as many motors, the control system must be more complicated. In practice, about $\pm 30^\circ$ orientation control about each axis is possible, before cable-cable or cable-EE interference occurs. If all three orientation angles are maintained at zero, however, the EE can be positioned throughout the static workspace with no interference. End-effector subcomponents above the cable attachment plane must be designed to avoid cable-EE interference. A detailed justi-

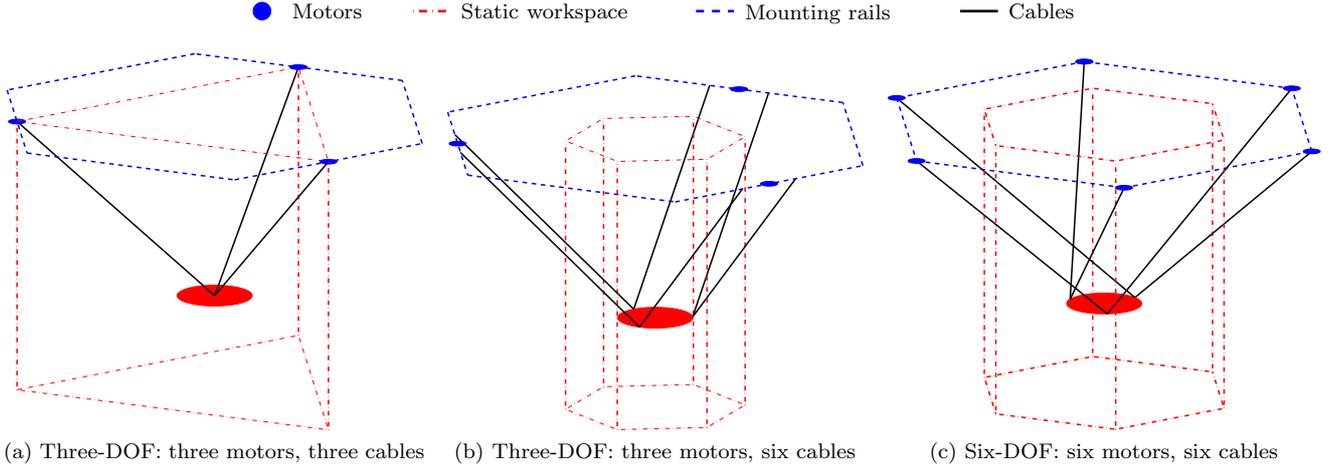


Figure 2: Options for spatial positioning.

fication and comparison of this option to alternative cable configurations is provided in [11, 12].

The six-DOF option is therefore selected as the ideal configuration, with the experimental system shown in Fig. 1. The cables are positioned by Pittman GM14904S012 motors driving 50 mm-diameter spools, as shown in Fig. 4. The motors have a maximum continuous torque of 0.88 N-m and a no-load speed of 597 RPM. These specifications translate into a maximum cable tension of 35.2 N and a maximum cable speed of 1.51 m/s. The positioning cables are made of Ultra High Molecular Weight Polyethylene (UHMWPE) kite lines¹. Kite lines are used for positioning because they tend to stretch much less than steel cables.

2.2. Kinematics

For the 3D printing application, paths are described most easily as position points in three-dimensional Euclidean space, with the orientation held constant. The inverse displacement problem (IDP) for the six-DOF mechanism at hand consists in finding the six cable lengths that will yield the desired pose (position and orientation) at P . As can be seen in Fig. 5,

$$\mathbf{b}_k = \mathbf{p} - \mathbf{a}_k - \mathbf{r}_k, \quad \rho_k = \mathbf{e}_k^T \mathbf{b}_k. \quad (1)$$

where \mathbf{b}_k is the vector from the actuator k to the attachment point of cable k on the EE, \mathbf{e}_k is the unit vector parallel to \mathbf{b}_k , \mathbf{a}_k is the vector from the inertial reference frame \mathcal{F} to actuator k , \mathbf{r}_k is the vector from the attachment point of cable k on the EE to the EE operation point P , and \mathbf{p} is the vector from \mathcal{F} to P . When the pose at P is known, the solution to the inverse displacement problem (IDP) is given by

$$\boldsymbol{\rho} = \begin{bmatrix} \rho_1 \\ \rho_2 \\ \vdots \\ \rho_6 \end{bmatrix} = \begin{bmatrix} \mathbf{e}_1^T \\ \mathbf{e}_2^T \\ \vdots \\ \mathbf{e}_6^T \end{bmatrix} \mathbf{p} - \begin{bmatrix} \mathbf{e}_1^T (\mathbf{a}_1 + \mathbf{r}_1) \\ \mathbf{e}_2^T (\mathbf{a}_2 + \mathbf{r}_2) \\ \vdots \\ \mathbf{e}_6^T (\mathbf{a}_6 + \mathbf{r}_6) \end{bmatrix}. \quad (2)$$

¹Spectra Laser Pro Gold 90#, which has a breaking force of 401 N (90 lb)

Kinematic sensitivity analysis is used to establish how errors in cable positioning affect the pose error, with the positioning error $\delta \mathbf{p}$ at P being of particular interest for AM. A small positioning error $\delta \rho_k$ is applied to Eq. (1) to obtain

$$\rho_k + \delta \rho_k = \mathbf{e}_k^T \mathbf{b}_k + \delta \mathbf{e}_k^T \mathbf{b}_k + \mathbf{e}_k^T (\delta \mathbf{p} - \delta \mathbf{a}_k - \delta \mathbf{r}_k) \quad (3)$$

which can be simplified to

$$\delta \rho_k = \mathbf{e}_k^T \delta \mathbf{p} + (\mathbf{e}_k \times \mathbf{r}_k)^T \delta \boldsymbol{\phi} \quad (4)$$

where $\delta \boldsymbol{\phi}$ is a small variation in the product of the unit vector parallel to the axis of rotation of the EE with the associated rotation angle. The relation between cable length errors and pose error can therefore be represented as

$$\delta \boldsymbol{\rho} = \mathbf{K} \delta \mathbf{x} \quad (5)$$

with

$$\delta \boldsymbol{\rho} = \begin{bmatrix} \delta \rho_1 \\ \delta \rho_2 \\ \vdots \\ \delta \rho_6 \end{bmatrix}, \quad \mathbf{K} = \begin{bmatrix} (\mathbf{e}_1 \times \mathbf{r}_1)^T & \mathbf{e}_1^T \\ (\mathbf{e}_2 \times \mathbf{r}_2)^T & \mathbf{e}_2^T \\ \vdots & \vdots \\ (\mathbf{e}_6 \times \mathbf{r}_6)^T & \mathbf{e}_6^T \end{bmatrix}, \quad (6)$$

$$\delta \mathbf{x} = \begin{bmatrix} \delta \boldsymbol{\phi} \\ \delta \mathbf{p} \end{bmatrix}.$$

If $\delta \boldsymbol{\rho}$ is regarded as the change in cable position during a short time δt , and δt approaches zero, the relation between cable velocities and twist \mathbf{t} is obtained

$$\dot{\boldsymbol{\rho}} = \mathbf{K} \mathbf{t} \quad (7)$$

where $\mathbf{t} = [\boldsymbol{\omega} \quad \dot{\mathbf{p}}]^T$ and $\boldsymbol{\omega}$ is the vector of angular velocity for the EE, given by $\boldsymbol{\omega} = \delta \boldsymbol{\phi} / \delta t$ as δt approaches zero.

The desired point-displacement sensitivity at P can be defined as the effect of cable length error $\delta \boldsymbol{\rho}$ on position error $\delta \mathbf{p}$, for a given EE pose defined by \mathbf{K} . Here,

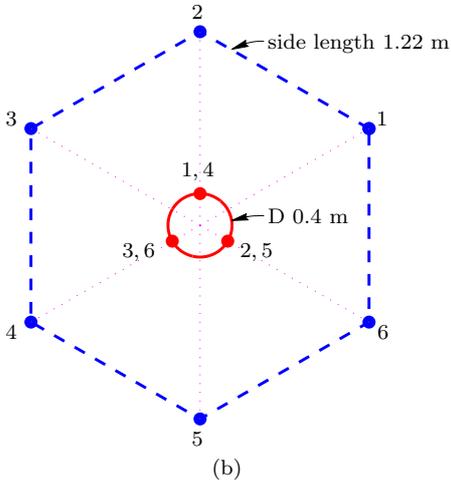
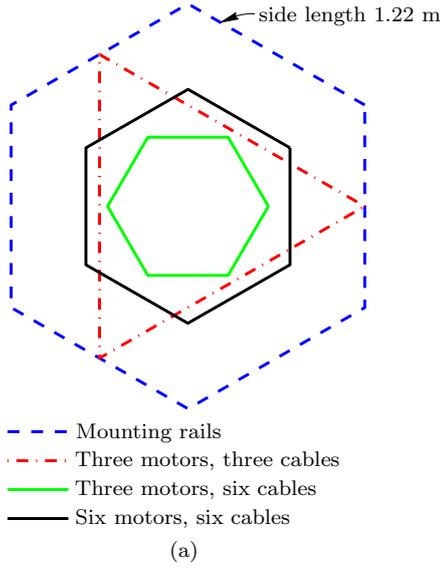


Figure 3: Cable-robot configurations: (a) static workspace comparison for the spatial positioning options; (b) cable attachment points for the six-DOF option, as seen from *below*.



Figure 4: The motor-brake assembly. The brake is the Ogura Industrial Corp. SNB 0.2G, with a torque of 1.5 lb-ft (2.04 N-m).

a point-displacement sensitivity index for dimensionally nonhomogenous \mathbf{K} is used, defined by Cardou et al. [13],

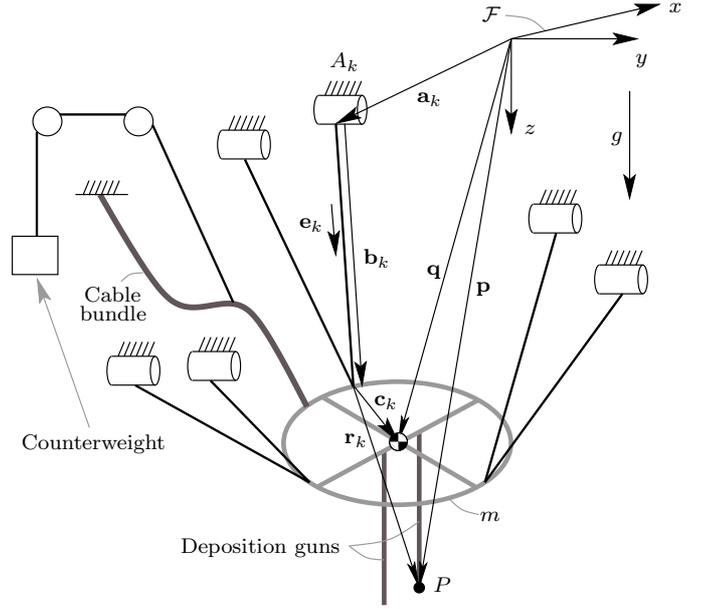


Figure 5: Spatial six-DOF cable-suspended robot.

which accounts for the coupled nature of rotations and translations and is independent of EE geometry, except for the choice of the operation point P .

As described in [13], a point-displacement sensitivity index $\sigma_{p,q}$ can be defined as

$$\sigma_{p,q} = \max_{\|\delta\rho\|_q=1} \|\delta\mathbf{p}\|_q. \quad (8)$$

where $\|\delta\mathbf{p}\|_q$ is the magnitude of point-displacement error at P , $\|\delta\rho\|_q$ is the magnitude of the vector of cable length errors, and $\|\cdot\|_q$ is the q -norm of its argument \cdot . Therefore, the *maximum* magnitude of point-displacement error for *arbitrary* cable-length errors is given by

$$\max \left(\|\delta\mathbf{p}\|_q \right) = \sigma_{p,q} \|\delta\rho\|_q \quad (9)$$

To compute $\sigma_{p,q}$, Eq. (5) is first transformed into

$$\delta\rho = \mathbf{K}_r \delta\phi + \mathbf{K}_p \delta\mathbf{p} \quad (10)$$

where $\mathbf{K} = [\mathbf{K}_r \quad \mathbf{K}_p]$ and $\delta\mathbf{x} = [\delta\phi^T \quad \delta\mathbf{p}^T]^T$, with \mathbf{K}_r and \mathbf{K}_p being the 6×3 rotational and positioning submatrices of \mathbf{K} , each of which is dimensionally homogeneous. The vector $\delta\phi$ cannot simply be set to $\mathbf{0}_3$, a vector of three zeros, since the desired index must account for the effect of rotation sensitivity on point-displacement sensitivity.

When $q = \infty$, the desired index is found by solving three linear programs:

$$\sigma_{p,\infty} = \max_{i=4,5,6} \left(\max_{\mathbf{x}} \mathbf{e}_i^T \mathbf{x} \quad \text{s.t.} \quad \mathbf{L}\mathbf{x} - \mathbf{1}_{2n} \preceq \mathbf{0}_{2n} \right) \quad (11)$$

where $\mathbf{e}_i \in \mathbb{R}^6$ is formed of null components, except for the i th, which is one, $\mathbf{L} = [\mathbf{K}^T \quad -\mathbf{K}^T]^T$, \preceq denotes

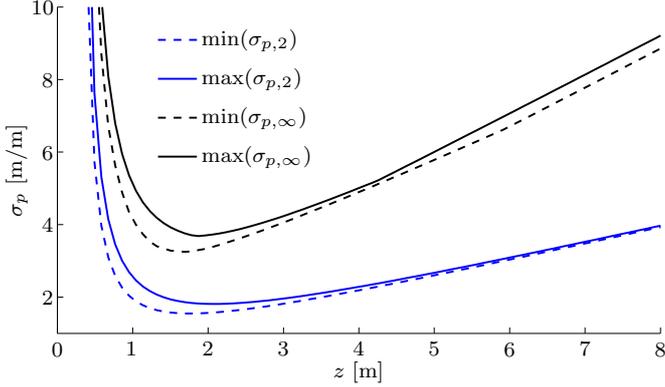


Figure 7: Kinematic sensitivity index σ_p for the cable-suspended 3D foam printer, using both the 2-norm and the ∞ -norm. Maximum and minimum curves bound σ_p , within the static workspace for the six-DOF robot, which is shown in Fig. 3a.

the component-wise inequality, $\mathbf{0}_{2n}$ is a vector of zeros of length $2n$, and $\mathbf{1}_{2n}$ is defined likewise.

When $q = 2$, a closed-form expression for the desired index is found:

$$\sigma_{p,2} = \frac{1}{\sqrt{\min(\lambda(\mathbf{K}_p^T \mathbf{P}_r \mathbf{K}_p))}} \quad (12)$$

where $\mathbf{P}_p \equiv \mathbf{1}_{6 \times 6} - \mathbf{K}_r(\mathbf{K}_r^T \mathbf{K}_r)^{-1} \mathbf{K}_r^T$, $\mathbf{1}_{6 \times 6}$ is the 6×6 identity matrix, and $\lambda(\cdot)$ are the eigenvalues of matrix \cdot .

Fig. 6 shows the kinematic sensitivity index σ_p for the cable-suspended 3D-foam printer, using both the 2-norm and the ∞ -norm. Plots are only shown for one quadrant in the horizontal plane, but are representative of the other four quadrants. A zoom of the plots is provided between z -values of 1.1 m and 2.1 m, because these are the limits of the workspace used during 3D printing. The plots show σ_p *inside* and *outside* the static workspace, defined by the hexagon of Fig. 3a, which corresponds to maximum x and y values of 0.6 m and 0.71 m, respectively. In practice, during printing, the robot is typically only used within a 0.5 m-radius cylinder. It can be seen that σ_p varies mostly with z , especially when only the static workspace is considered. Figure 7 highlights the relationship between σ_p and z .

Considering only the static workspace, between z -values of 1.1 m and 2.1 m, it can be seen that $\sigma_{p,2}$ varies between 1.6 and 2.4, and $\sigma_{p,\infty}$ varies between 3.3 and 4.7. Since the actuators are all essentially identical, a maximum positioning error magnitude can be assigned to each cable $\delta\rho_{\max}$, the worst-case scenario then being when all six cables exhibit this error. In this case,

$$\max(\|\delta p\|_2) = \sigma_{p,2} \|\delta b\|_2 = \sqrt{6} \sigma_{p,2} \delta\rho_{\max} \quad (13a)$$

$$\max(\|\delta p\|_\infty) = \sigma_{p,\infty} \|\delta b\|_\infty = \sigma_{p,\infty} \delta\rho_{\max}. \quad (13b)$$

Under these conditions, $\sigma_{p,2}$ varies between $3.9\delta\rho_{\max}$ and $5.9\delta\rho_{\max}$, and $\sigma_{p,\infty}$ varies between $3.3\delta\rho_{\max}$ and $4.7\delta\rho_{\max}$.

As mentioned in [13], the 2-norm index is regarded as an acceptable approximation of the ∞ -norm index in

most cases. This conclusion is confirmed by the results reported here, and it should also be noted that the 2-norm computations are about 25 times more computationally efficient than the ∞ -norm computations for the robot being studied. The error amplification is also acceptable, since $\delta\rho_{\max}$ is nearly always below 0.7 mm, as seen in Fig. 11, which produces a maximum point-displacement error $\|\delta p\| = 4.1$ mm, using the upper bound of 5.9 for the 2-norm. This is far less than the claimed 3D printing accuracy of 1 cm.

The relationship between cable velocities and EE linear and angular velocities is found by expanding Eq. (7) to obtain

$$\dot{\rho} = \mathbf{K}_r \boldsymbol{\omega} + \mathbf{K}_p \dot{\mathbf{p}}. \quad (14)$$

where $\boldsymbol{\omega} = \mathbf{0}_3$, because the EE orientation is kept constant for the 3D printing application, with $\mathbf{0}_n$ being a vector of n zeros. It becomes evident that the maximum speed of any one cable must always be less than $\|\dot{\mathbf{p}}\|_2$.

Both sides of Eq. (14) can be differentiated to obtain the system acceleration equation

$$\ddot{\rho} = \dot{\mathbf{K}}_r \boldsymbol{\omega} + \mathbf{K}_r \dot{\boldsymbol{\omega}} + \dot{\mathbf{K}}_p \dot{\mathbf{p}} + \mathbf{K}_p \ddot{\mathbf{p}}. \quad (15)$$

Again, with the orientation held constant, the first two terms on the right-hand side of Eq. (15) are equal to $\mathbf{0}_6$. The third term can be simplified according to

$$\dot{\mathbf{K}}_p \dot{\mathbf{p}} = \begin{bmatrix} \dot{\mathbf{e}}_1^T \\ \dot{\mathbf{e}}_1^T \\ \vdots \\ \dot{\mathbf{e}}_6^T \end{bmatrix} \dot{\mathbf{p}} = \begin{bmatrix} (\mathbf{e}_1 \times \dot{\mathbf{p}})^T (\mathbf{e}_1 \times \dot{\mathbf{p}}) / \rho_1 \\ (\mathbf{e}_2 \times \dot{\mathbf{p}})^T (\mathbf{e}_2 \times \dot{\mathbf{p}}) / \rho_2 \\ \vdots \\ (\mathbf{e}_6 \times \dot{\mathbf{p}})^T (\mathbf{e}_6 \times \dot{\mathbf{p}}) / \rho_6 \end{bmatrix}. \quad (16)$$

Each component of this term, which represents the acceleration in a cable due to its degree of *non-alignment* with $\dot{\mathbf{p}}$, is *always* positive. This may seem counter-intuitive, though if $\boldsymbol{\omega} = \mathbf{0}_3$ and $\dot{\mathbf{p}}$ is constant, $\dot{\rho}_k$ can only *increase*, and therefore $\ddot{\rho}_k$ must be positive. The fourth term represents the acceleration of each cable due to its degree of *alignment* with $\dot{\mathbf{p}}$. Applying Eq. (15) to one cable yields

$$\ddot{\rho}_k = +\mathbf{e}_k^T \ddot{\mathbf{p}} + \frac{1}{\rho_k} (\mathbf{e}_k \times \dot{\mathbf{p}})^T (\mathbf{e}_k \times \dot{\mathbf{p}}). \quad (17)$$

The preceding kinematic analysis was based on the solution to the inverse displacement problem (IDP). The forward displacement problem (FDP) consists in calculating the operation point position \mathbf{p} and orientation $\boldsymbol{\phi}$, when the vector of cable lengths $\boldsymbol{\rho}$ is known. A closed-form solution to this problem has not been formulated, thus it is advantageous to avoid real-time control solutions that rely on solving the FDP. However, the problem is similar to the forward displacement problem for a Gough-Stewart platform, for which many numerical solution techniques exist, such as polynomial methods [14, 15], iterative methods [15], and interval analysis techniques [16]. Many of these techniques apply Newton methods during the numerical solution procedure, although derivative-free techniques, such

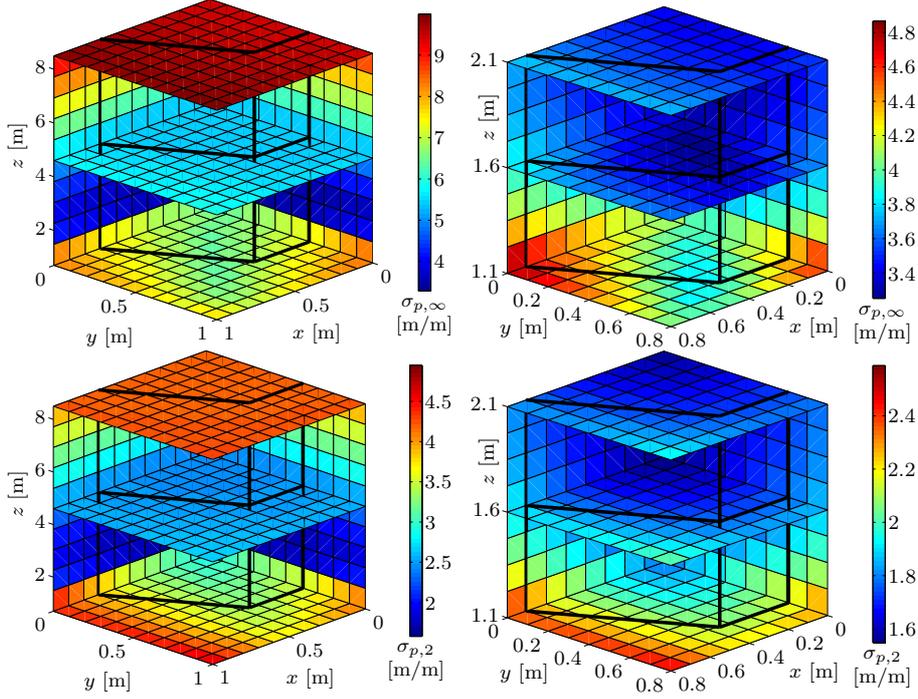


Figure 6: Kinematic sensitivity index σ_p for the cable suspended 3D-foam printer, using both the 2-norm and the ∞ -norm. Heavy black lines indicate the static workspace boundary.

as the Nelder-Mead Simplex Method [17], are used in some cases.

In our PID control technique, described in Section 2.4, cable lengths are each controlled *separately*, based on the solution to the IDP. A more complex control system could consist of controlling the pose at P , using both the IDP and the FDP. However, potential gain in positioning accuracy using this method was judged to be insufficient to justify the development of online FDP solution method.

The static compensation term h_k , defined in Section 2.3 is therefore based on the *commanded* pose at P , rather than the *actual* pose, which would require a solution to the FDP, using on the cable lengths indicated by the motor encoder positions. However, trajectory tracking error is typically less than 1 mm per cable, which produces at most 6 mm of point displacement error at P . Since the cable lengths vary between about 1 m and 2.5 m, the approximation error caused by using the commanded pose to find h_k is negligible.

2.3. Dynamics

Eq. 17 is used to characterize the relationship between cable acceleration and the Cartesian velocity and acceleration at P . It can be combined with the equations of motion for the robot to characterize the usable workspace. These equations are given by

$$\mathbf{M}\boldsymbol{\tau} = \begin{bmatrix} -\frac{1}{m}\mathbf{I}\dot{\boldsymbol{\omega}} - \frac{1}{m}\boldsymbol{\omega} \times \mathbf{I}\boldsymbol{\omega} \\ \mathbf{g} - \ddot{\mathbf{q}} \end{bmatrix}, \quad (18)$$

with

$$\mathbf{M} = \begin{bmatrix} \mathbf{e}_1 \times \mathbf{c}_1 & \mathbf{e}_2 \times \mathbf{c}_2 & \cdots & \mathbf{e}_6 \times \mathbf{c}_6 \\ \mathbf{e}_1 & \mathbf{e}_2 & \cdots & \mathbf{e}_6 \end{bmatrix} \quad (19)$$

where $\boldsymbol{\tau}$ is the vector of cable tensions, normalized by the EE mass, $\mathbf{g} = [0 \ 0 \ g]^T$ is the gravitational acceleration vector, m is the EE mass and \mathbf{I} is the EE mass moment of inertia matrix. Since the EE is held stationary, $\boldsymbol{\omega} = \mathbf{0}_3$ and $\ddot{\mathbf{q}} = \ddot{\mathbf{p}}$. For Eq. (18) to be valid, the EE must be rigid, and its center of mass (CM) is located as shown in Fig. 5. In practice, these conditions are respected relatively well, but not perfectly.

If $\mathbf{h} = \mathbf{M}^{-1} [\mathbf{0}_3 \ \mathbf{g}]^T$, Eq. (18) can be rewritten as

$$\boldsymbol{\tau} = \mathbf{h} - \ddot{\mathbf{p}}. \quad (20)$$

Combining the component version of Eq. (20) with Eq. (17) produces

$$\mathbf{e}_k^T \ddot{\mathbf{p}} = h_k - \tau_k - \frac{1}{\rho_k} (\mathbf{e}_k \times \dot{\mathbf{p}})^T (\mathbf{e}_k \times \dot{\mathbf{p}}) \quad (21)$$

which can be used to characterize the relationship between the cable tension, the EE velocity, the EE acceleration, and the location in the robot workspace. Assuming the minimum and maximum cable tension constraints τ_{\min} and τ_{\max} , the inequality relations below are established:

$$\|\ddot{\mathbf{p}}\| \leq h_k - \tau_{k,\min} - \frac{1}{\rho_k} \dot{\mathbf{p}}^T \dot{\mathbf{p}} \quad (22a)$$

$$\|\ddot{\mathbf{p}}\| \leq \tau_{k,\max} - h_k. \quad (22b)$$

Table 1: System constants and variables

Parameter	Description	Value	Units
R	motor gearbox ratio	5.9	RPM/RPM
K_{mot}	motor torque constant	0.061	N-m/A
i_{max}	maximum motor current	3.2	A
r_s	spool radius	0.025	m
m	EE mass	6	kg
K_{enc}	counts per rev.	2000	counts/rev.
K_p	proportional gain	230	A/m
K_i	integral gain	0	A/m
K_d	derivative gain	1.31	A/m
c_{enc}	encoder counts		counts
b_{home}	cable length at home pos.		m
b_{com}	cable length at start pos.		m
b_{enc}	change in cable length since start pos.		m
h_k	static cable tension		N/kg

Eq. (22) can be used to restrict the robot workspace such that a minimum acceleration magnitude at P is possible in any direction. Negative $\|\ddot{\mathbf{p}}\|_{\text{max}}$ indicates that a constraint violation has occurred, and P cannot accelerate in all directions. The maximum acceleration along cable k occurs when $\ddot{\mathbf{p}}$ and \mathbf{e}_k point in the same direction and $\dot{\mathbf{p}}$ is perpendicular to \mathbf{e}_k . The minimum acceleration along cable k occurs when $\ddot{\mathbf{p}}$ and \mathbf{e}_k point in opposite directions and $\dot{\mathbf{p}}$ and \mathbf{e}_k are parallel. It should be noted that h_k and ρ_k both depend on the pose of P , and $\tau_{\text{min}} = 0$, since the cables must remain in positive tension. The maximum tension depends on the spool actuator according to

$$\tau_{\text{max}} = \frac{RK_m i_{\text{max}}}{r_s m} \quad (23)$$

with all variables defined in Table 1. The maximum cable speed is 1.51 m/s, corresponding to a maximum possible value for $\dot{\mathbf{p}}^T \dot{\mathbf{p}}$ of 2.28 m²/s². However, during 3D printing, the maximum speed at P is 0.1 m/s, corresponding to a maximum $\dot{\mathbf{p}}^T \dot{\mathbf{p}}$ of 0.01 m²/s², which is negligible.

Figure 8 shows the variation of $\|\ddot{\mathbf{p}}\|_{\text{max}}$ within the section of the robot workspace typically used for 3D printing. It is evident that the constraint-respecting volume corresponds very closely to the static workspace shown in Fig. 3a, when the maximum cable speed is 0.1 m/s. However, the usable workspace decreases significantly as this speed is increased. Within the plot boundaries of Fig. 8, $\|\ddot{\mathbf{p}}\|_{\text{max}}$ is *always* restricted by the τ_{min} constraint Eq. (22a). However, if the maximum cable speed is sufficiently high, or if the minimum z -value is sufficiently low, Eq. (22b) will restrict $\|\ddot{\mathbf{p}}\|_{\text{max}}$ at some locations.

2.4. Control

The PID control loop for one actuator controlling one cable, shown in Fig. 9, attempts to minimize the position error ϵ , given by

$$\epsilon_k = \rho_{\text{com},k} - \rho_{\text{enc},k} - \rho_{\text{start},k} \quad (24)$$

where $\rho_{\text{com},k}$ is the desired position, or setpoint, and the actual position is given by the sum of the start position

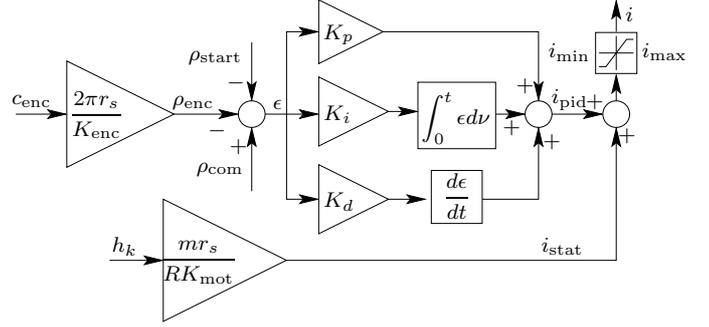
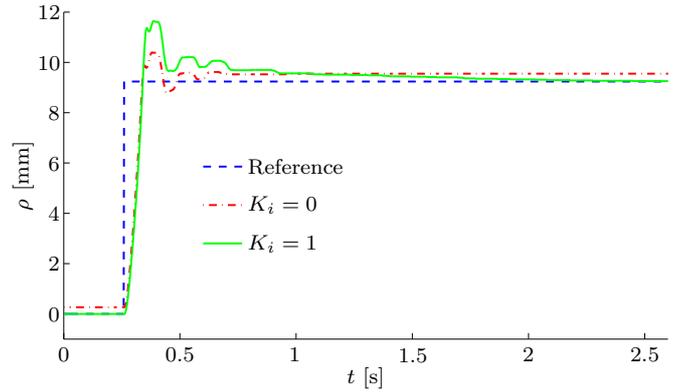


Figure 9: PID control for one motor of the six-DOF cable-suspended robot.


 Figure 10: Step response for one cable, calculated using motor encoder data, with $K_p = 230$ and $K_d = 1.31$.

$\rho_{\text{home},k}$ and that indicated by the encoder $\rho_{\text{enc},k}$. The Ziegler-Nichols method was used to tune the PID gains [18], followed by manual adjustments to minimize ϵ_k along typical robot paths, resulting in the values shown in Table 1. With these gains, the positioning error for each cable is typically less than 1 mm, at speeds up to 0.3 m/s and accelerations up to 0.3 m/s².

The step response of the PID controller for one cable is shown in Fig. 10. The cable position response data are calculated from motor encoder data that were recorded while the robot was commanded to go from $z = 1.68$ m to $z = 1.8$ m, with $x = 0$ and $y = 0$. The response with the PID gains of Table 1, which correspond to those used during 3D printing, are compared to those with $K_i = 1$. In both cases, the response is stable, with $K_i = 1$ exhibiting more overshoot but less steady-state error, as expected.

Based on these results, it might be expected that the optimal PID controller has nonzero K_i . However, for 3D printing, the steady-state error is of little importance; the more important objective is to minimize the positioning error *while tracking trajectories*. If K_i is increased sufficiently to have a noticeable effect on this, the controller output becomes unstable. Additionally, nonzero K_i can have destabilizing effects. For example, if one of the foam deposition guns is pressing against the part being constructed, integral windup can occur.

Figure 11 shows the error ϵ , calculated using Eq. (24),

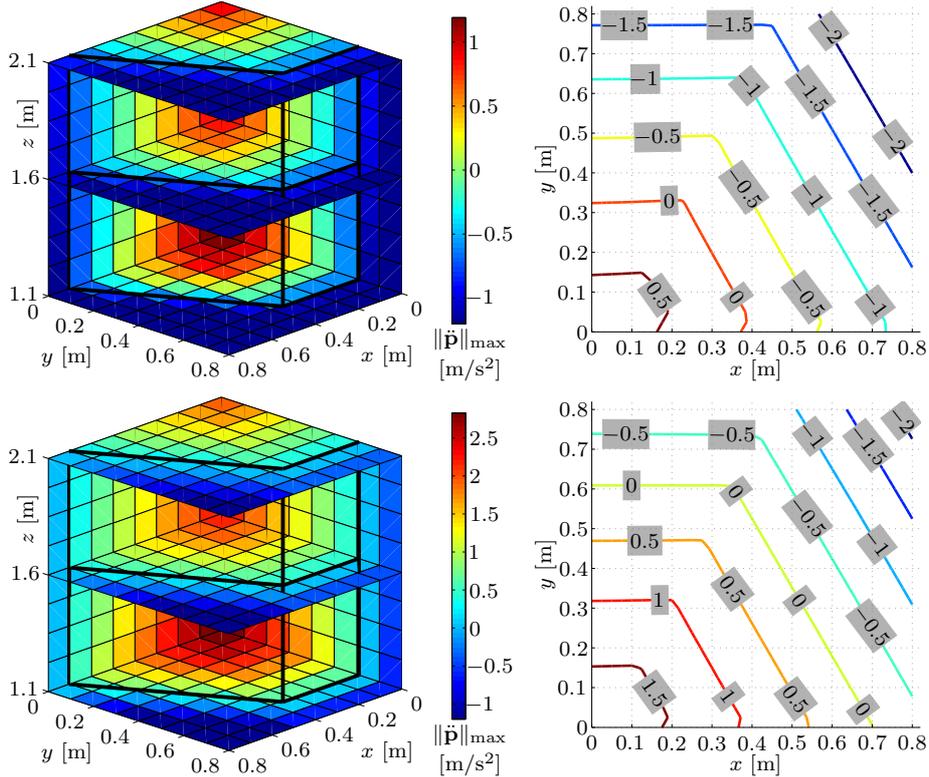


Figure 8: The maximum acceleration magnitude $\|\dot{\mathbf{p}}\|_{\max}$: with $\dot{\mathbf{p}}^T \dot{\mathbf{p}} = 2.28 \text{ m}^2/\text{s}^2$ for the upper plots; with $\dot{\mathbf{p}}^T \dot{\mathbf{p}} = 0.01 \text{ m}^2/\text{s}^2$ for the lower plots; left hand volumetric plots are shown for $x > 0$ and $y > 0$, which is representative of the other four quadrants; right-hand plots are shown for $z=1.1 \text{ m}$, which contains the most restrictive zero-contour within the 1.1 to 2.1 m z -range.

exhibited along typical robot paths. The response is shown with and without the feedforward term h_k , which predicts the static tension in each cable needed to balance the weight of the EE, for any location in the workspace. It is evident that the amplitude of the error is about the same in both cases, but with static compensation, the maximum error magnitude is approximately halved, showing that h_k is reliably compensating for the static force created by the weight of the EE. In Fig. 11, mean error μ and standard deviation σ are calculated according to

$$\mu = \frac{1}{n} \sum_{i=1}^n \epsilon_i, \quad \sigma = \sqrt{\frac{1}{n-1} \sum_{i=1}^n (\epsilon_i - \mu)^2}. \quad (25)$$

When the robot is restricted to the hexagonal-prism-shaped usable workspace of Fig. 8, the maximum static compensation current possible is $i_{\text{stat}} = 1.99 \text{ A}$. Since the maximum motor current is 3.2 A , the robot can reach any location in the workspace. However, it should be noted that the robot is not controllable *near* the hexagonal boundary, where $h_k = 0$, because any negative value of i_{PID} will make a cable go slack. In practice, a cylindrical workspace about 1 m in diameter is used to avoid this problem.

Fault detection and response is particularly important for the system introduced here, which has a relatively heavy EE, compared to the inertia of the motor-brake-

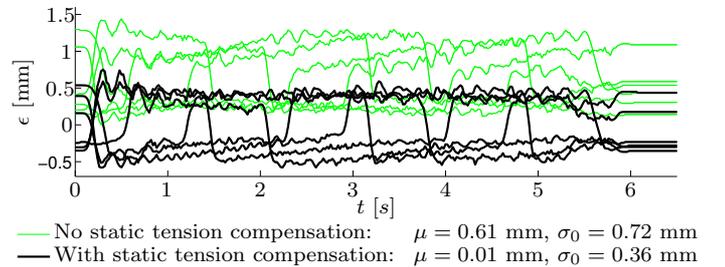


Figure 11: The position error ϵ for all six cables while the spatial six-DOF cable-suspended robot follows a 0.5-m-diameter horizontal circular path, with $K_p = 230$, $K_i = 1$ and $K_d = 1.3$.

spool assemblies. In fact, the EE weight is sufficient to overcome the inertia of all six motor assemblies and make the EE fall to the ground at a high acceleration, potentially damaging it, or breaking the positioning cables.

A simple, robust fault response technique for this type of system is to use failsafe brakes, as shown in Fig. 4. These brakes must be powered to be disengaged, and will engage automatically on loss of power, locking the EE in place. The brakes can be engaged using soft faults, which consist of error states within the control program, or hard faults, such as a total loss of power, or a motor fault.

2.5. Deposition materials and mechanisms

As discussed in the introduction, the feasible printing materials for large-scale additive manufacturing are very limited because their volumetric cost must be low. For this reason, all examples of large-scale AM to date have used concrete or other earth-based materials for printing. Unfortunately, most of these materials have a limited shelf-life in the ready-to-print state. For example, once concrete has been mixed, it is typically extrudable only for a few hours. Additionally, concrete flow characteristics gradually change over time: an ever-increasing amount of force is required for extrusion.

Another disadvantage of earth-based materials is their weight: moving large volumes of these materials requires a mechanism with heavy-lifting capacity. Such a cable-driven mechanism could be designed and constructed, but it would not be ideal for a prototype, proof-of-concept system. Based on these and other considerations, a large-scale AM material for a cable-driven 3D printer should have the following characteristics:

- inexpensive
- lightweight
- can be stored in a ready-to-deposit state
- can be deposited in a controlled manner

One class of materials that satisfies these characteristics is *foams*. For the object material, only foams that produce a solid, permanent structure after deposition need to be considered. With these restrictions, one-component polyurethane foam emerges as a clear candidate. It is relatively inexpensive, at \$600/m³, and it weighs about 25 kg/m³. It is available in 10-lb and 16-lb cylinders that can be stored for months in a ready-to-deposit state. Shaving foam was selected as the support material, because it does not permanently bond to polyurethane foam and is easily removed by spraying the part with water following construction.

With the deposition materials selected, a method is needed for precise, machine-controlled deposition. Manual dispensing guns, specifically designed for polyurethane foam, are readily available for purchase, such as the one shown in Fig. 12. The foam flowrate is controlled by pulling a variable-position trigger to retract a teflon-coated needle within the gun barrel. The needle and barrel have cup-and-cone mating surfaces, allowing for relatively precise flow control. When the trigger is released, a return spring shuts the gun, pushing the needle against the front of the barrel, stopping the flow immediately at the gun tip. Since this foam gun is specifically designed for precise flow control of polyurethane foam, a reasonable approach is to simply retrofit it for automated deposition.

Such an automated foam deposition system was developed, as shown in Fig. 13. The gun is rigidly clamped in place, with the trigger and spring return removed, to provide easy access to the back end of the needle. The



Figure 12: The Great Stuff Pro 14 foam dispensing gun, manufactured by the Dow Chemical Company.

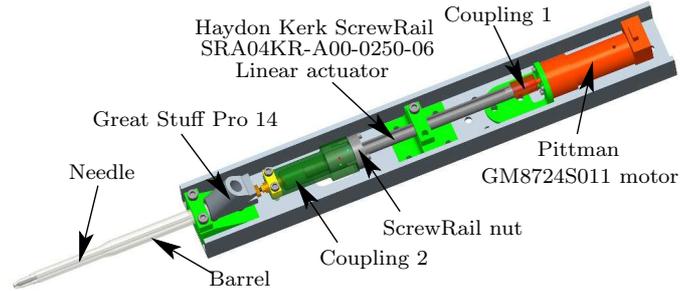


Figure 13: Foam gun assembly.

needle is rigidly coupled to a SRA04KR-A00-0250-06 Haydon Kerk ScrewRail linear actuator, which has a lead of 5.35 mm/rev (0.25 in/rev) and a maximum axial load of 111 N (25 lb). A Pittman GM8724S011 DC servo gearmotor is used to drive the linear actuator, with a maximum speed of 720 RPM, and a maximum continuous torque of 0.102 N-m, with both specifications applying *after* the 6.3:1 gearbox. The maximum possible speed of the linear actuator, and thus the maximum axial speed of the foam gun needle, is 64 mm/s; the maximum continuous axial force is 19 N. This system provides relatively robust needle position control: with a PID position controller, a dynamic positioning error of 0.05 mm is readily achieved. This is the average error observed while the needle follows a fully specified trajectory, with needle positions specified at the controller time step of 0.002 s.

The relationship between needle position and foam flowrate is also needed. To establish this, the flowrate was found at different positions by measuring the mass of foam deposited during a certain period of time. The flowrate vs. needle position curves for both polyurethane foam and shaving foam are shown in Fig. 14. Ideally, the position-flowrate relationship would be linear, though in reality, it is highly non-linear, with a relatively large deadband. However, the positioning range within which most foam flow occurs is sufficiently large to produce acceptable flow control. The deadband can be eliminated by changing the zero position of the needle, though this tends to cause some foam leakage. The shaving foam is much less controllable than the polyurethane foam, due to the rapid initial flowrate increase, starting at a needle position of about 0.7 mm. The plots only show the flowrates up to a needle position of

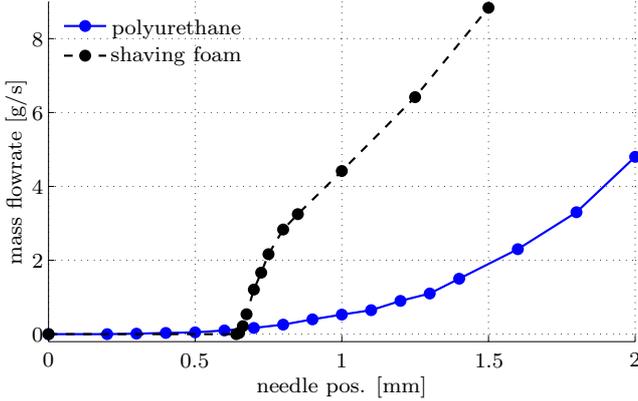


Figure 14: Variable flowrate through the foam deposition guns.

2 mm, as higher flowrates are never used during foam 3D printing. Material flowrate is continuously varied in synchronization with the continuously varying EE speed. A relatively constant bead size is achieved using this method, as can be seen in the real-time trajectories shown in the video of Fig. 18, for the online version of the article.

The mechanical design shown in Fig. 13 works relatively well, though it has several disadvantages. Firstly, because the barrel is held rigidly in place, significant shock can occur when the needle contacts it as the gun is shut. Secondly, the system no longer has a failsafe shut mechanism: if power is lost during deposition, the gun will stay open and foam will continue to flow. Thirdly, for a system with multiple deposition guns, inactive guns can touch and destroy the material deposited with the active gun.

These flaws are addressed in a modified failsafe foam gun actuation mechanism, shown in Fig. 15. Here, the barrel is made *movable* along the needle axis, by mounting it to a linear rail. A locking mechanism, actuated with a solenoid, is added to provide two gun modes: (a) the locked position, shown in 15a, where the solenoid is powered, the lever is engaged, and the motor controls the axial position of the needle within the barrel, and thus the foam flow; (b) the unlocked position, shown in Fig. 15b, where the solenoid is unpowered, the barrel spring holds the barrel against the needle, and the motor positions the needle and the barrel simultaneously.

With this design, the barrel is not rigidly fixed, and the barrel spring dampens the shock of closing the gun. Secondly, foam flow immediately stops in the case of a fault, because the solenoid must be powered to keep the lever engaged and the gun open. In case of a fault and/or loss of power while depositing foam, the solenoid will stop exerting force on the lever, the lever return spring will cause the lever to disengage, and the barrel return spring will push the barrel axially into the needle, shutting off the foam flow.

Under normal conditions, at most one foam gun will be armed for deposition at a time, with all other guns retracted. To arm a gun for deposition, the motor is used to first push the needle and barrel past the position where

the lever can be engaged. The lever is then engaged, and then the needle and barrel are retracted to the zero position. Any further retraction will move the needle only, causing foam to flow, since the barrel is held in place by the lever.

The failsafe foam gun mechanism is also more suitable for soft faults, when motor and solenoid control are still available. For example, under certain fault conditions, it may be desirable to immediately turn off foam flow. This is achieved almost instantaneously by shutting off the solenoid, whereas, with the rigid mechanism, the motor would need to be actuated to drive the needle against the barrel.

Both the rigid and failsafe foam gun mechanisms are designed to achieve precise, *variable* foam flowrates. Constant flowrates would significantly limit the performance of the 3D printer: either constant-speed deposition paths would need to be used, or accuracy would need to be sacrificed when the EE is not traveling at the nominal speed.

A variable-flow-control deposition system is also capable of correcting *local* geometry errors, meaning it can be used to vary the flow along a deposition path, to correct part-height error. This is to be distinguished from *global* flowrate adjustment, whereby the average flowrate would be adjusted based on the average part-height error. A local geometry-error correction system, called surface mapping feedback, has been implemented on the 3D foam printer, and is described in detail in the next subsection.

2.6. Geometric feedback

A common problem for many 3D printing systems is the accumulation of error that is proportional to the part height. This problem can lead to the rejection of many materials that would otherwise be suitable for 3D printing. A geometric feedback algorithm, called surface mapping feedback (SMF), was developed specifically for these types of materials [10]. With SMF, the geometry error is measured periodically, typically between layers, and deposition parameters are adjusted for subsequent layers to correct the errors detected. Figure 16 shows the SMF concept.

SMF is designed to correct *local* geometry errors. Therefore, a system capable of specific, location-targeted geometry correction is needed, such as that introduced in the previous subsection. Additionally, a device capable of accurately measuring the geometry errors is necessary.

A linear one-dimensional laser displacement sensor was selected for measuring geometry error, the Keyence IL-300, with the IL-1000 sensor amplifier. The measurement range is 160–450 mm from the emission point, the linearity is 0.25% of full scale, and the sampling rate is up to 3000 Hz. The IL-1000 outputs distance measurements as analog voltage or current.

The advantage of using a one-dimensional displacement sensor is that it is reasonably inexpensive and provides an accurate mapping of the surface. The disadvantage is that the laser must be physically displaced to map the entire

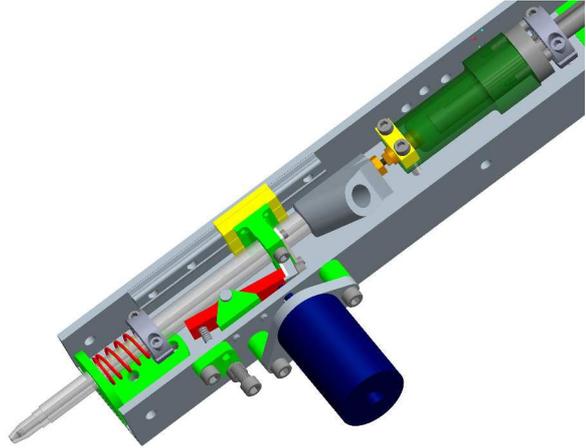
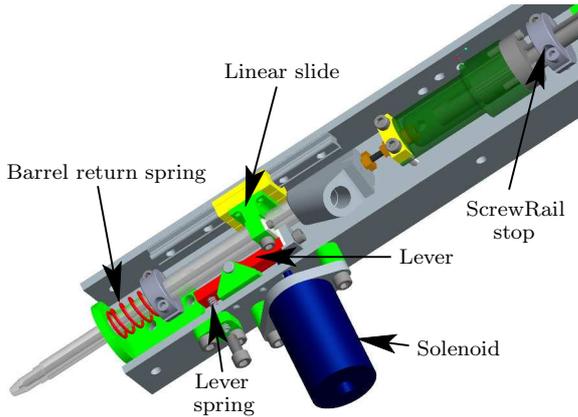


Figure 15: Foam gun failsafe mechanism: (left) solenoid on, foam flow enabled; (right) solenoid off, foam flow disabled.

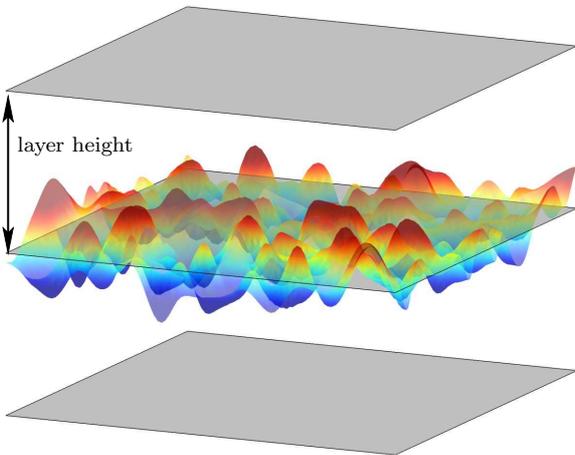


Figure 16: Surface Mapping Feedback (SMF).

top surface of the part, a process that takes about the same time as depositing the material for one layer. More sophisticated sensors could map swaths or areas of the part all at once, though these are typically much more expensive. Additionally, these sensors will not be able to accurately map the top surface if it is undulating, since line of sight to the entire surface region to be mapped might not be possible from the laser emission location.

SMF attempts to minimize the measured error at all locations of the top surface of the part, by implementing PID control loop

$$\mathbf{U}_m = K_p \mathbf{E}_m + K_i \sum_{i=1}^m \mathbf{E}_i + K_d (\mathbf{E}_m - \mathbf{E}_{m-1}) \quad (26)$$

where K_p , K_i , and K_d are the proportional, integral, and derivative gains, respectively, and \mathbf{E}_m is the *error surface* for layer m , represented by a matrix of height error data, several hundred units long on each side. The output surface \mathbf{U}_m is treated as a surface *function*, which can be evaluated at any point in the horizontal plane according to $U(x, y)$. Therefore, for a particular point (x_i, y_i) along

a path in a future layer, the flowrate Q is given by

$$Q = C \dot{s}_i [1 - U(x_i, y_i)] \quad (27)$$

where C is an experimentally determined material-dependent constant of proportionality and \dot{s}_i is the speed at point i . More information on SMF can be found in [10] and [19].

The PID gains were selected using the Ziegler-Nichols method, followed by manual adjustment, with the primary objective being to minimize $\|\mathbf{E}\|$. In regions where the part is too low, the height error will be negative as indicated by the blue regions in Fig. 16. For the subsequent layer, SMF will modify the foam gun actuator positions at path points in these regions, such that the foam flowrate will be higher than the nominal rate obtained according to linear interpolation of the position-flowrate relationships shown in Fig. 14. In regions where the error is positive, shown by the red in Fig. 16, less material will be deposited in the next layer. If the measured height is consistently low in the same region, the integral term will compensate for this by causing the system to deposit even more material to this region in the next layer. If the error in a particular region decreases significantly from one layer to the next, the derivative term will cause a further increase in deposition, for this region, in the next layer. The derivative term thus helps the system respond to disturbances and dampen error oscillation.

2.7. Control hardware

In the previous subsections, the mechanical hardware for the 3D foam printer was introduced. Here, the control systems are described, including the motor drives and computers. Adapting an off-the-shelf solution for controlling the six-DOF cable-suspended robot would be difficult, since much of the low-level control code cannot typically be modified by the user for such systems. Therefore, a customized control approach is much more reasonable, since a mathematical model for the system can be built from scratch.

RT-Lab was selected as the development environment for the 3D foam printer because it offers a great degree of design flexibility without the need for extensive low-level programming. RT-Lab is a real-time simulation software environment, fully integrated with MATLAB/Simulink. A typical RT-Lab Simulink model consists of Console and Master subsystems. The Console is a GUI that allows the user to interact with the system while it is running. The Master consists of mathematical models of the hardware to be controlled, along with the implementation of all other real-time operations. After the model is compiled, the Console runs on a PC, in MATLAB/Simulink, and the Master runs on the target node, a computer with a real-time operating system, such as QNX, installed. The real-time computer is connected to all of the physical hardware, and communicates with the PC over an ethernet network. A diagram of the 3D foam printer control systems is shown in Fig. 17. In this case, a Slave target node is also present, to increase the number of axes that can be controlled.

The motor drive boxes receive control commands from the target nodes and send current to the motors using pulse-width modulated (PWM) current control. Each drive box contains four (4) BE40A20 PWM servo drives, made by Advance Motion Controls. Two drive boxes are used, providing control of eight motors: six for EE positioning axes and two for deposition axes. Motor encoder signals, which indicate angular position, are routed through the motor drive box to the target nodes, where they are used in the motor PID position control loops, described in Section 2.4.

The target nodes communicate with the motor drive boxes using digital and analog control signals. They have the following input-output cards installed: Acromag APC8620 carrier board, SBS/Greenspring IP-Quadrature for reading motor encoder signals, Acromag IP220-8 analog out for sending the current control signal to the motor drives, Acromag IP320-8 analog in for reading the laser position signal, and Acromag IP470 digital in/out for control signals between the target node and the motor drive box.

It is important to note that only the Master and the Slave shown in Fig. 17 are truly real-time subsystems. The Console communicates with the Master and Slave over an ethernet network, but it is only suitable for operations such as on/off toggling or modification of certain parameters. RT-Lab control is designed to be tolerant of lags or interruptions in communication between the Console and the real-time subsystems.

While the Console can be thought of as loosely disconnected from the Master and Slave, it is still a Simulink model with a fixed-time-step solver, and communication with these subsystems is rarely delayed by more than a few milliseconds. On the other hand, the Matlab Supervisory Code is quite atypical for an RT-Lab implementation: it is completely asynchronous with the other three subsystems.

The Matlab Supervisory code serves two main func-

tions. Firstly, it periodically supplies trajectory data to the Master, to avoid transferring the data for an entire part over the network at once. Secondly, and most importantly, it implements surface mapping feedback (SMF), described in the previous subsection. Many of the Matlab functions and toolboxes applied for SMF are incompatible with real-time code.

As shown in Fig. 17, the Matlab Supervisory Code communicates with the other subsystems by reading and writing to the same files. Although, some of this communication could be replaced by more direct communication such as TCP/IP, restricting the communication to file IO helps to isolate system crashes. For example, if the Matlab Supervisory Code crashes, the Master will continue to run but will get stuck while waiting for a file to appear in a specific location. Upon observing the crash, a user can then switch to manual mode in the Console and control the robot, in many cases also successfully resuming the part construction. More direct communication between the Supervisory Code and the other subsystems is more likely to produce a cascading system crash.

3. Part construction process

Here, the part-construction process is described, starting with an STL or PLY input file, and ending with the constructed polyurethane foam object. The process is demonstrated through the construction of a foam statue of Sir Wilfrid Laurier, the seventh Prime Minister of Canada. Figure 18 shows the original plaster statuette, the scanned PLY file, and the constructed polyurethane statue. In the online version, the video shown in this figure summarizes the construction process, while also showing timelapse photography of the foam deposition.

Construction begins with the creation and verification of the STL or PLY file for the object to be built. This file could be created from scratch using CAD software, downloaded online², or produced via 3D scanning. This file should have no holes in the mesh and no non-manifold edges or vertices.

The 0.725-m-high Laurier statuette, created by Louis-Philippe Hébert in 1898, was loaned to the authors by the Museum Collections of Laval University. A 1-mm-resolution 3D scan of the statuette was produced at the Computer Vision and Systems Laboratory of Laval University, using the Go!Scan 3D by Creaform. MeshLab [20] was used to manually remove noise from the raw scanned data, and then a screened Poisson surface reconstruction [21] was used to close holes. MeshLab was also used to eliminate non-manifold edges and other geometry errors, and to produce the final watertight mesh composed of 1 052 706 triangular faces, stored in the PLY binary data format.

The original statuette is scaled up by a factor of three to produce the constructed foam statue, 2.16 m high. Since

²e.g. <http://graphics.stanford.edu/data/3Dscanrep/>

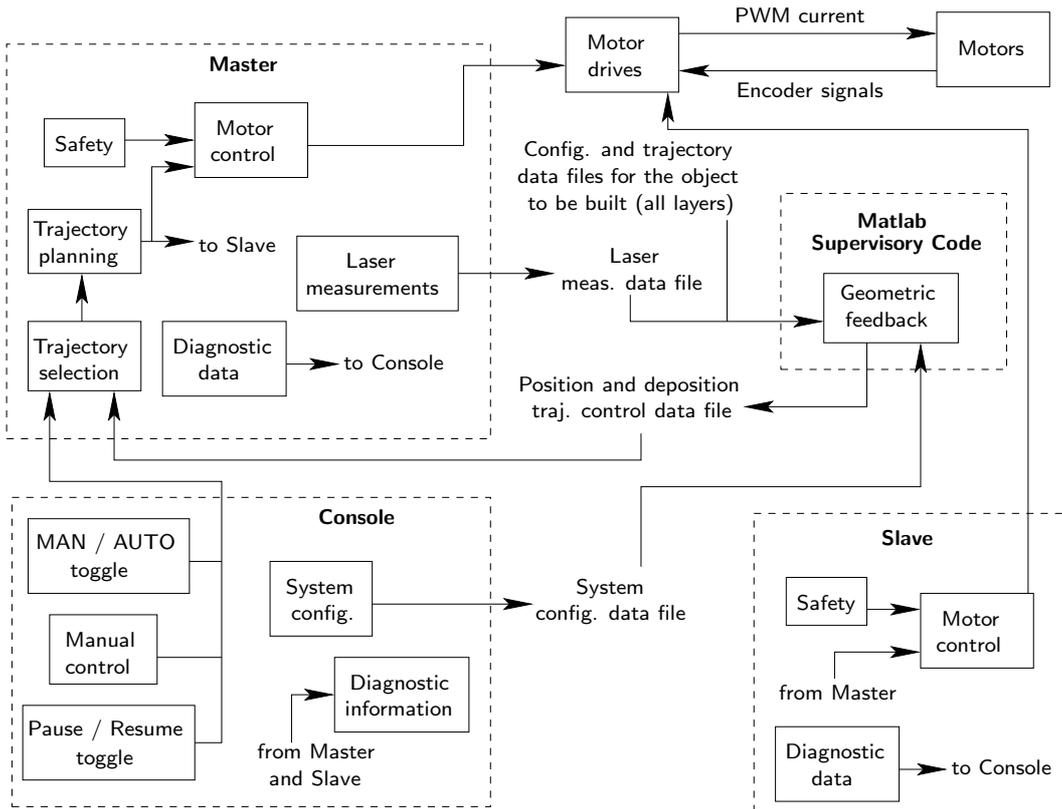


Figure 17: Control diagram for the 3D foam printer.

Figure 18: Video summarizing the construction of the Laurier statue. The still image shows the original plaster statue, the PLY file rendering, and the constructed foam statue.

the usable workspace for the 3D foam printer is shaped like a circular cylinder, about 1 m high and 1 m in diameter, the statue is built in three sections, which are then glued together.

3.1. Tool-path generation

An input text file is used to control tool-path generation parameters such as scaling, rotation, path height, and path width. Tool paths are generated, based on the CAD file, in the following order: object boundaries, support boundaries, and fill paths. More information about the tool-path generation process can be found in [22].

The path height and path width used for constructing the Laurier statue are 10 mm and 12 mm, respectively, for both the polyurethane foam and the shaving foam. These dimensions were chosen based mainly on the objective to build as large an object as possible, in a reasonable amount of time. However, other factors include a desire to maximize construction resolution and optimize foam gun performance. With these dimensions, the Laurier statue is built in 216 layers, or 72 layers per section. Figure 19 shows the object and support boundary paths for section 2 of the Laurier statue, along with the boundary and fill paths for one layer.

A 3D printer must have a complete, robust support solution to be able to construct *any* 3D structure. As mentioned in the Introduction, most of the large-scale 3D printers developed to date can only produce relatively simple shapes such as extrusions and parts that do not require support.

The ideal support solution is material-dependent. One potential large-scale full-support solution is the selective layer binding used by Dini [7], also described in the Introduction. The advantages of this technique are fast deposition and the recycling of the support material. However, a large volume of support material is needed and the material containment structure is necessarily larger than the largest object to be printed.

One of the main advantages of using a cable-suspended robot for 3D printing is its transportability and relatively easy workspace configuration conversion. With a selective layer binding support technique, these advantages are greatly diminished, when compared to a traditional, gantry-type 3D printer. For this reason, a deposition-based support technique was chosen for the 3D foam printer.

As described in Section 2.5, shaving foam was selected as the support material for the 3D foam printer. Since this material does not form a rigid structure and will slump under its own weight, two novel support techniques were developed, shown two-dimensionally in Fig. 20.

Figure 20a shows the first method, the polyurethane retaining shell, where shaving foam occupies all of the volume encountered when projecting the part downward, except the part volume itself. A polyurethane shell is formed to contain the shaving foam and prevent it from collapsing. To prevent the shell from contacting the object, the shell

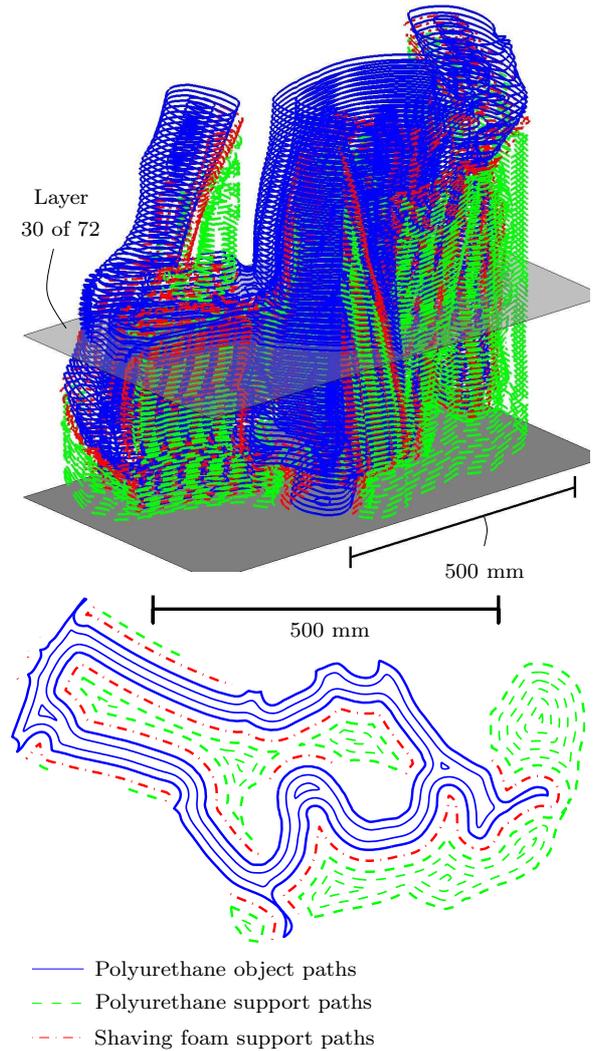


Figure 19: Deposition paths for the Laurier statue, section 2 of 3: (above) boundary contours for all 72 layers; (below) boundary and fill contours for layer 30 of 72.

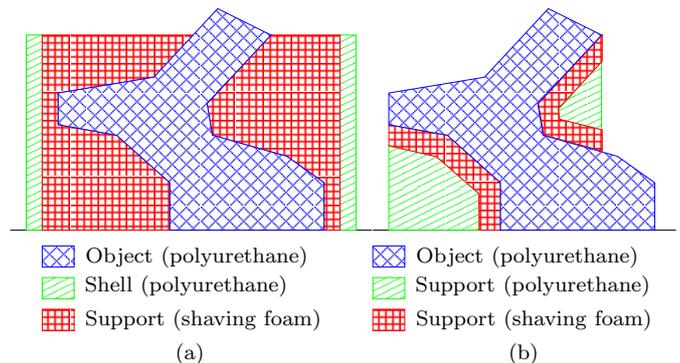


Figure 20: Support techniques: (a) polyurethane retaining shell; (b) shaving foam buffer.

is offset outward, with the additional space also filled by shaving foam.

Figure 20b shows the second method, the shaving foam

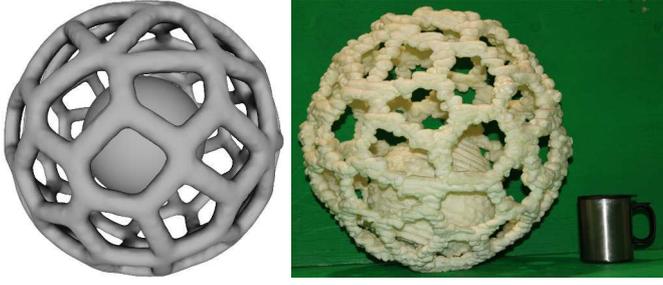


Figure 21: Voronoi sphere, 40 cm in diameter, with trapped solid sphere: (left) rendering of PLY file; (right) foam model constructed with the 3D printer.

buffer, where, nominally, polyurethane is used as both the object and support materials. A buffer of shaving foam is created in between the object and support sections, to permit support removal. Since there are no large concentrations of shaving foam, surface tension forces are sufficient to hold it in place.

The polyurethane retaining shell method can be used to construct any 3D part, but it is very material-intensive. For this reason, the shaving foam buffer technique was also developed, which is much more efficient, though might seem to have limited applicability due to the difficulty of support removal. However, since polyurethane foam is *elastically deformable* to a certain extent, this method can be used for most part geometries. Indeed, support removal for the Laurier statue was accomplished quite easily, as can be seen in the video of Fig. 18, for the online version of the article. The shaving foam buffer technique can even be used for parts such as the Voronoi sphere shown in Fig. 21. In this case, the polyurethane support was *intertwined* through the holes in the sphere, but it was successfully removed in about one hour by cutting it into small pieces.

3.2. Trajectory generation

Paths, defined in *space*, are converted to trajectories by also defining them in *time*. Solving this problem with an online controller is unnecessary, since all paths are known prior to construction. An offline solution is therefore used for trajectory generation: an optimization problem is defined for each path, whereby the objective is to minimize the time of travel, subject to the kinematic constraints

$$\begin{aligned} \dot{s}_{\max} &= 100 \text{ mm/s}, & a_{\max} &= 300 \text{ mm/s}^2, \\ \ddot{s}_{\max} &= 1200 \text{ mm/s}^3 \end{aligned} \quad (28)$$

where \dot{s}_{\max} is the maximum speed, a_{\max} is the maximum total acceleration, given by

$$a = \sqrt{\ddot{x}^2 + \ddot{y}^2}, \quad a \neq \ddot{s}, \quad (29)$$

and \ddot{s}_{\max} is the maximum jerk.

These constraints are quite conservative, and the foam guns typically deposit at only 0–10% of their maximum

flowrates along deposition paths. Conservative parameters were chosen at this stage of the foam printer development in order to focus on the accuracy and robustness of the system. There is considerable room for optimizing these parameters to increase the printing speed. However, if the maximum acceleration is significantly increased, cable tension constraints become necessary, according to the equations developed in Subsec. 2.3.

In [23], a trajectory planning technique for a 3-DOF positioning cable-suspended system, with tension constraints, is introduced. A similar technique is used for the 3D foam printer, except that the trajectories are two-dimensional, in the horizontal plane, and there are no tension constraints. Since the 3D foam printer uses 6-DOF cable positioning, the analytical cable tension constraint equations introduced in [23] cannot be used. However, numerical versions of these equations can be developed to extend their applicability to the 6-DOF system.

Figure 22a shows the input data for a single path, decimated by a factor of five to make individual points more discernible. Figure 22b shows output path data, also decimated by a factor of five, equally spaced in time by 0.1 s. The path speed, computed using this data, is shown in Fig. 22c.

For every path point, a foam gun position is computed according to Eq. (27) and Fig. 14, such that all trajectory data points have coordinates (x, y, a) . A trajectory file consists of a list of such points, equally spaced by 0.02 s, preceded by a header that indicates the number of trajectories, trajectory end points, and trajectory materials. During deposition, online cubic spline interpolation is used to reduce the point spacing to 0.002 s, the fixed time step of the controller. This technique is used to minimize the data needed to describe the trajectories, while at the same time ensuring smooth velocity, acceleration, and jerk while the trajectories are followed.

3.3. Foam deposition

Once trajectory planning is complete, all trajectory data for the entire part has been generated. Positioning data are not changed during part construction, but deposition data are adjusted after every layer, according to the geometric feedback algorithm, SMF.

If the foam printer is offline, the hardware systems shown in Fig. 17 are first turned on, and then the Simulink model, which consists of the Console, Master, and Slave, is loaded and executed. Then, while in MAN mode in the Console, the foam deposition systems are initialized: tank levels are verified, valves are opened, and flow through the guns is ensured. Part construction is started by executing the Matlab Supervisory Code, and then switching the Console to AUTO. While printing, the following operations are executed in a loop until the part is complete:

1. Measure the surface formed by the deposited material from the previous layer.
2. Adjust the deposition control data for the next layer.

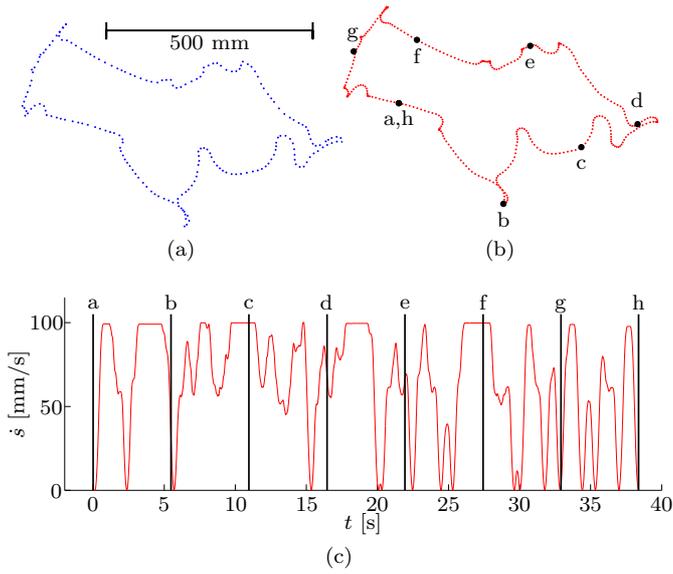


Figure 22: Trajectory planning for one 2.5 m-long path: (a) input points, decimated by factor of five; (b) output points, 0.1 s spacing; (c) trajectory speed.

3. Deposit the polyurethane foam trajectories for the next layer.
4. Deposit the shaving foam trajectories for the next layer.
5. Wait for the prescribed time.

A minimum wait time per layer is needed to ensure that the polyurethane foam cures sufficiently before the next layer is deposited. Without this time, if the layers are sufficiently short, the uncured foam concentrate will coalesce, and large gas bubbles will form. The minimum wait time needed depends on the maximum polyurethane path concentration, i.e., the maximum number of adjacent polyurethane paths for the layer. For a polyurethane shell, consisting of isolated paths that do not touch when deposited, the minimum layer wait time is less than a minute. For layers with many adjacent paths, such as those used to build the pedestal of the Laurier statue, the minimum layer wait time is about twelve minutes. Figure 23 shows the Laurier foam statue, at various stages of construction.

4. Results and discussion

The 3D printing time for building all three sections of the Laurier polyurethane foam statue was 38 hours. The 2.16-m-tall statue has a volume of 182 L; the support consisted of 93 L of polyurethane foam and 23 L of shaving foam. The printer performance can be divided into three broad categories: accuracy, robustness, and speed.

Thus far, maximizing the accuracy has been the main focus of the printer development. One indication of accuracy is the measurements made with the laser for the geometric feedback algorithm. However, these will only show

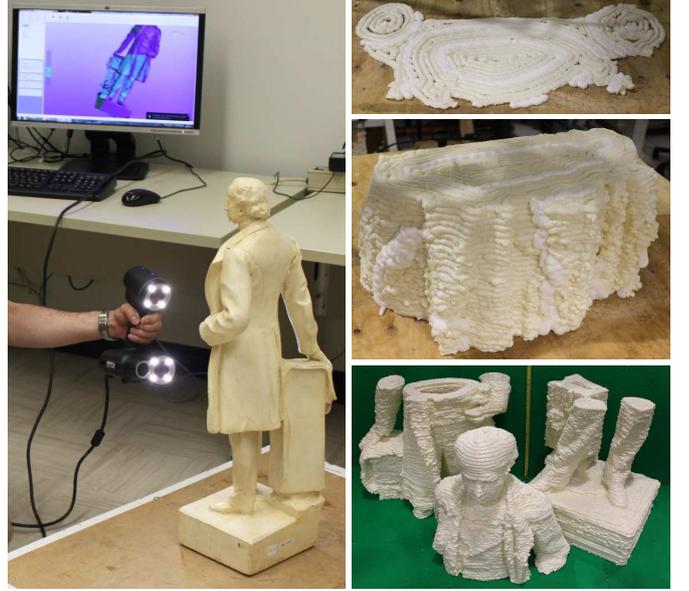


Figure 23: Laurier statue construction: (left) 3D scanning; (right) pictures taken during the construction process.

vertical error and will not indicate error due to EE positioning, because the laser is mounted on the EE. Even with these drawbacks, the laser measurements provide a good indication of the printing accuracy. Additionally, they can be used to analyze the performance of the geometric feedback system.

Surface mapping is accomplished by following all deposition paths for the previously deposited layer, in the laser reference frame, while reading in laser height measurements at the model fixed time step, 0.002 s. Measurements are recorded in memory for the entire layer, and then saved at a resolution of 0.02 s, in the same format as the trajectory files, except laser measurements replace foam actuator positions.

Figure 24 shows the interpolated error surface for six layers of section 2 of the Laurier statue. The Matlab function TRISCATTEREDINTERP is used to produce surface data, based on the laser height measurements. In this figure, layer 0 represents the measurement of the substrate before deposition starts; this is used to establish the initial height of the part. It can be seen that the error rarely exceeds 10 mm above or below the nominal part height. This behavior is representative of the performance seen when constructing the other two sections.

The height error can also be quantified by finding the weighted mean μ^w and standard deviation σ^w of the error for each layer m , according to

$$\mu^w = \frac{\sum_{i=1}^{n^*} \dot{s}_i e_i}{\sum_{i=1}^{n^*} \dot{s}_i}, \quad (30a)$$

$$\sigma^w = \left[\frac{\sum_{i=1}^{n^*} \dot{s}_i (e_i - \mu)^2}{\sum_{i=1}^{n^*} \dot{s}_i} \right]^{1/2} \quad (30b)$$

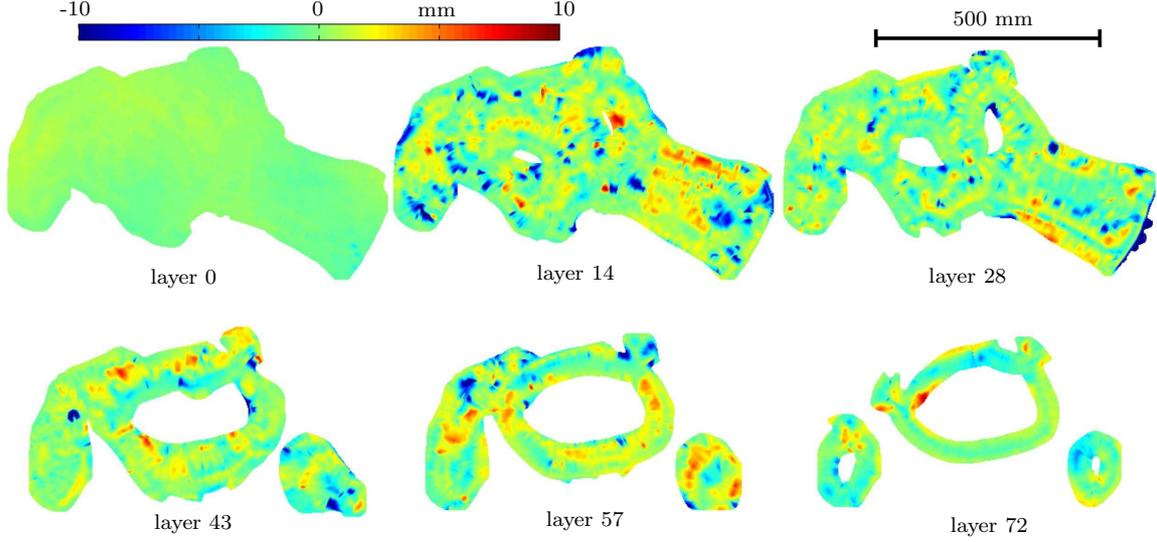


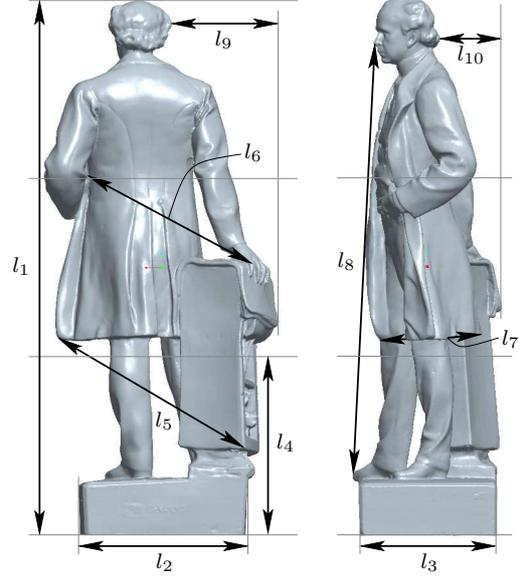
Figure 24: Vertical error measured with the laser displacement sensor for the Laurier statue, section 2 of 3.

where n^* is the number of measured points for layer m , e_i is the error at point i , and \dot{s} is the speed at point i . The point errors are weighted by the path speed to obtain a mean and standard deviation in space, rather than in time.

Fig. 25 shows μ^w and σ^w for each material at every layer of the Laurier statue. Both parameters are rarely more than 5 mm from the nominal level and both are well-controlled throughout the construction. The shaving foam performance is inferior because it is less controllable, as is evident in Fig. 14.

Based on Figs 24 and 25, it can be concluded that SMF is successful in controlling the vertical error for the 3D foam printing system. Before SMF was implemented, many iterations of construction were needed to produce even the simplest of parts. Even then, parts were limited to a height of approximately 20 layers before the error became too large and there was significant interference between the foam gun and the part, or the deposition clearance was so great that the foam was deposited far from the correct location.

A better indication of the printer accuracy can be obtained from the comparison of physical measurements on the printed foam statue to the corresponding virtual measurements on the CAD model. Ten representative measurements were selected, shown graphically and listed in Fig. 26. The difference between the physical and virtual measurements is nearly always less than the claimed printing accuracy of 10 mm. Measurements l_1 and l_8 are expected to be larger, because they each cross two section boundaries, so additional error is expected due to section alignment and gluing error. Measurements l_9 and l_{10} were taken to characterize whether the statue has the proper vertical orientation. The large error for l_9 is observable in the picture of the foam statue shown in Fig. 18. This error is likely due to an EE positioning bias that warrants



Measurement	CAD model [±3 mm]	Foam statue [±5 mm]	Error [mm]
l_1	2156	2137	19
l_2	618	619	1
l_3	545	540	5
l_4	720	718	2
l_5	861	864	3
l_6	711	704	7
l_7	404	413	9
l_8	1749	1727	22
l_9	435	485	50
l_{10}	242	245	3

Figure 26: Distance measurements for the Laurier statue.

further investigation.

A second measure of the printer performance is robustness. During the construction of the Laurier statue, major

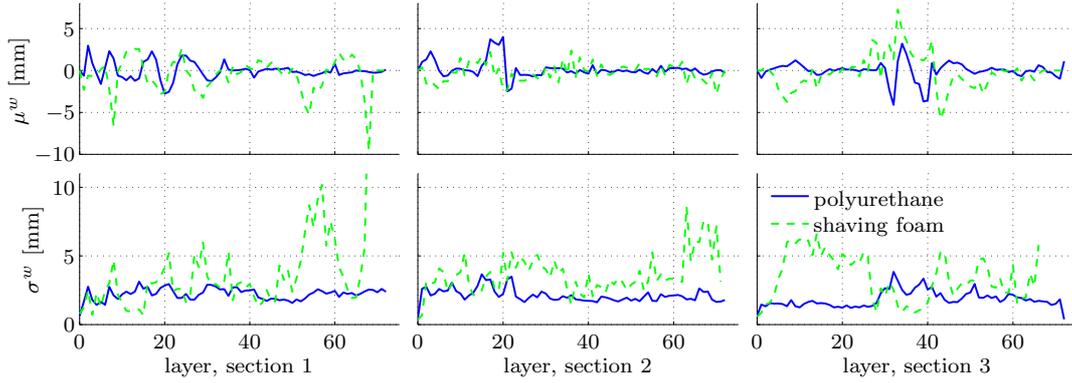


Figure 25: Mean error and standard deviation of laser surface height measurements, for the Laurier statue, for sections: (a) 1 of 3; (b) 2 of 3; (c) 3 of 3.

manual intervention was required about five times, for reasons such as removing a blockage in a foam supply hose and debugging a motor drive fault. Minor manual adjustments were also made to enhance the print quality, such as removing accumulated foam residue on the foam gun tips. Otherwise, the printing process proceeded automatically. Obviously, a fully automated system is desirable, and this is an area for future work.

The print speed was relatively conservative, because of the initial focus on accuracy. It is likely that the speed could be doubled without making any hardware changes; with a few minor improvements the speed could likely be increased much more.

It can also be observed that sharp edges and narrow features, such as the coat lapel of the Laurier statue, are reproduced poorly. This is caused by the error associated with approximating the 2D areas with 1D deposition paths, which can likely be greatly reduced by refining the tool-path planning.

The bulk properties of the completed parts are also an important measure of printer performance. The bonding strength of the polyurethane foam between layers is about as strong as that within layers; individual layers cannot be peeled apart after the foam has cured. As polyurethane foam is an insulating material, it intended to fill enclosed, dark spaces. However, when used as a construction material for additive manufacturing, the finished product might be expected to degrade, since it is stored in the open at room temperature, exposed to air and light. However, the authors have observed that long-term exposure effects are minor, since parts stored for over a year in these conditions have changed minimally. The most observable effect is the transition from a light yellow or off-white color, initially, to a darker yellow.

Proper handling of a constructed part for the first few days following construction is critical. If the shaving foam not entirely removed, the leftover residue degrades the polyurethane surface over time, creating a viscous, orange coating. This problem is easily eliminated by thoroughly spraying the part with water following construction. After deposition and scaffolding removal, overhanging features

of the part can sag under their own weight, with deformations of up to 3 cm having been observed. This problem is addressed by supporting these features from below for a few days, until the foam is fully cured.

Build orientation affects several part properties, and thus cannot simply be chosen to optimize material use and construction speed. Firstly, part features that are in contact with the shaving foam have a rougher finish. Secondly, in the constructed orientation, the geometry seen from the side is generally more visually pleasing than that seen from above or below. For example, a head will look better when built upright than on its side, even if the printing accuracy is the same in both cases.

5. Conclusion and future work

In this paper, a large-scale 3D foam printer was introduced. The system includes several novel features, most notably, the use of a cable-suspended mechanism for 3D printing. The support techniques, which permit the construction of any 3D geometry, are also unique. Finally, the geometric feedback system proved to be a critical component of the 3D printer, greatly enhancing the system robustness and accuracy.

One area for future work is to enlarge the robot workspace. This will enable the construction of parts like the Laurier statue in a single piece. However, it should be noted that about 100 L more support would be required to build the statue in this way, using the shaving-foam-buffer technique. Also, the proper suspension of the cable bundle, which supplies power, signals, and deposition materials to the EE, becomes significantly more difficult when enlarging the workspace. This bundle must be carefully configured, such that it does not interfere with the positioning cables, never exerts a large force on the EE, and never collides with the part being built.

The accuracy, robustness, and speed of the printer will also be improved, as described at the end of Section 4. Finally, different printing materials will be investigated.

Acknowledgements

The authors gratefully acknowledge the support of The Social Sciences and Humanities Research Council of Canada (SSHRC), under Strategic Research Grant 890-2011-0039, led by Prof. Aaron Sprecher, McGill University. This work was also supported by Le Fonds de Recherche du Québec – Nature et Technologies (FRQ-NT), in the form of a post-doctoral fellowship for E. Barnett. The authors thank the Museum Collections of Laval University for providing the plaster statuette of Sir Wilfrid Laurier, and the Computer Vision and Systems Laboratory of Laval University for 3D scanning the statuette.

References

- [1] S. Lim, R. A. Buswell, T. T. Le, S. A. Austin, A. G. F. Gibb, T. Thorpe, Developments in construction-scale additive manufacturing processes, *Autom. Constr.* 21 (2012) 262–268.
- [2] B. Khoshnevis, Automated construction by Contour Crafting - related robotics and information technologies, *Autom. Constr.* 13 (2004) 5–19.
- [3] J. Zhang, B. Khoshnevis, Optimal machine operation planning for construction by Contour Crafting, *Autom. Constr.* 29 (2013) 50–67.
- [4] R. A. Buswell, R. Soar, A. Gibb, A. Thorpe, Freeform construction: mega-scale rapid manufacturing for construction, *Autom. Constr.* 16 (2007) 224–231.
- [5] T. T. Le, S. A. Austin, S. Lim, R. A. Buswell, A. G. F. Gibb, T. Thorpe, Mix design and fresh properties for high-performance printing concrete, *Mater. Struct.* 45 (2012) 1221–1232.
- [6] J. Pegna, Exploratory investigation of solid freeform construction, *Autom. Constr.* 5 (1997) 427–437.
- [7] E. Dini, Method for automatically producing a conglomerate structure and apparatus therefor, US Patent 8,337,736 (2012).
- [8] P. Bosscher, R. L. Williams II, L. S. Bryson, D. Castro-Lacouture, Cable-suspended robotic Contour Crafting system, *Autom. Constr.* 17 (2007) 45–55.
- [9] L. L. Cone, Skycam, an aerial robotic camera system, *Byte* 10 (10) (1985) 122–132.
- [10] E. Barnett, J. Angeles, D. Pasini, P. Sijkkes, Surface mapping feedback for robot-assisted rapid prototyping, in: *Proc. IEEE Int. Conf. Robot. Autom.*, Shanghai, China, 2011, pp. 3739–3744.
- [11] J.-D. Deschênes, P. Lambert, S. Perreault, N. Martel-Brisson, N. Zoso, A. Zaccarin, P. Hébert, S. Bouchard, C. Gosselin, A cable-driven parallel mechanism for capturing object appearance from multiple viewpoints, in: *Proceedings of the 6th International Conference on 3-D Digital Imaging and Modeling*, Montréal, Canada, 2007, pp. 367–374.
- [12] C. Gosselin, S. Bouchard, A gravity-powered mechanism for extending the workspace of a cable-driven parallel mechanism: application to the appearance modelling of objects, *Int. J. Autom. Tech.* 4 (4) (2010) 372–379.
- [13] P. Cardou, S. Bouchard, C. Gosselin, Kinematic-sensitivity indices for dimensionally nonhomogeneous Jacobian matrices, *IEEE Trans. Robot.* 26 (1) (2010) 166–173.
- [14] P. Nanua, K. Waldron, V. Murthy, Direct kinematic solution of a Stewart platform, *IEEE Trans. Robot. Autom.* 6 (4) (1990) 438–444.
- [15] J.-P. Merlet, Direct kinematics of parallel manipulators, *IEEE Trans. Robot. Autom.* 9 (6) (1993) 842–846.
- [16] J.-P. Merlet, Solving the forward kinematics of a Gough-type parallel manipulator with interval analysis, *Int. J. Robot. Res.* 23 (3) (2004) 221–235.
- [17] J. C. Lagarias, J. A. Reeds, M. H. Wright, P. E. Wright, Convergence properties of the nelder–mead simplex method in low dimensions, *SIAM J. Optim.* 9 (1) (1998) 112–147.
- [18] J. G. Ziegler, N. B. Nichols, Optimum settings for automatic controllers, *Trans. ASME* 64 (1942) 759–768.
- [19] E. Barnett, The Design of an Integrated System For Rapid Prototyping With Ice, PhD Thesis, McGill, Montreal, QC, Canada (2012).
- [20] P. Cignoni, MeshLab, Visual Computing Lab ISTI - CNR, <http://meshlab.sourceforge.net/>, accessed: 2014-12-16.
- [21] M. Kazhdan, H. Hoppe, Screened Poisson surface reconstruction, *ACM Trans. Graphics* 32 (3), article 29.
- [22] E. Barnett, J. Angeles, D. Pasini, P. Sijkkes, A heuristic algorithm for slicing in the rapid freeze prototyping of sculptured bodies, in: J. Angeles, B. Boulet, J. J. Clark, J. Kovacs, K. Siddiqi (Eds.), *Brain, Body, and Machine*, Springer-Verlag, Berlin, 2010, pp. 149–162.
- [23] E. Barnett, C. Gosselin, Time-optimal trajectory planning of cable-driven parallel mechanisms for fully specified paths with G1-discontinuities, *ASME J. Dyn. Sys.* 137 (7), 12 pages.