A selective recruitment strategy for exploiting muscle-like actuator impedance properties

Joshua Schultz[1], Glenn Mathijssen[2,4], Bram Vanderborght[2] and Antonio Bicchi[3,4]

***Abstract*— Two leading qualities of skeletal muscle that produce good performance in uncertain environments are damage tolerance and the ability to modulate impedance. For this reason, robotics researchers are greatly interested in discovering the key characteristics of muscles that give them these properties and replicating them in actuators for robotic devices. This paper describes a method to harness the redundancy present in muscle-like actuation systems composed of multiple motor units and shows that they have these same two qualities. By carefully choosing which motor units are recruited, the impedance viewed from the environment can be modulated while maintaining the same overall activation level. The degree to which the impedance can be controlled varies with total activation level and actuator length.**

Discretizing the actuation effort into multiple parts that work together, inspired by the way muscle fibers work in the human body, produces damage-tolerant behavior. This paper shows that this not only produces reasonably good resolutions without inordinate numbers of units, but gives the control system the ability to set the impedance along with the drive effort to the load.

I. INTRODUCTION

There is evidence that humans are successful at performing tasks in uncertain environments because they are able to modulate the physical impedance of their muscles to fit the task. To this end, several researchers have created actuation solutions that include functionality to vary the impedance of the joint in addition to imposing a force. Vanderborght et al. provides a review [1]. Most examples consist of dual servomotors that either vary the stiffness through co-contraction of antgonistic nonlinear springs [2],[3], or use one motor to impose torque, and the other to vary the impedance [4], [5]. While these examples and others can accommodate imprecision in the environment the way that humans do, they lack the resilience to failure of biological muscle.

The modular design of muscle systems has tremendous redundancy, meaning that should a few muscle cells fail, the actuator does not cease to function. This is because each cell is only responsible for a *portion* of the required force. This was realized by Du et al. [6], who studied the effect of failure

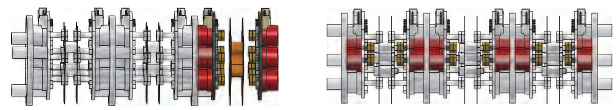

(a) Concentrating the activation in one or only a few actuation units leads to high actuator impedance

(b) Spreading the activation over all or most actuation units leads to low actuator impedance

Fig. 1: Both examples have the same activation level; active motor units are highlighted in red

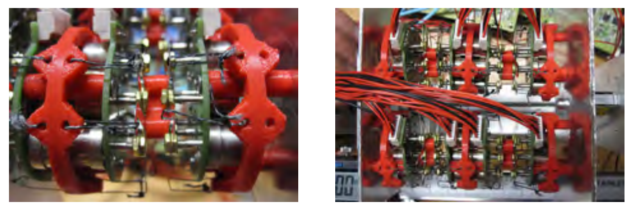

(a) A single actuation unit; mechanical stops are fiber straps

(b) A bundle of two chains of actuation units

Fig. 2: Series elastic actuation units, like those pictured here can be connected in series/parallel combinations to produce an actuator with desired performance and study recruitment phenonomena. In this implementation [8], each motor unit is a micro solenoid with a series elastic element. The series elastic elements are joined at a common hub, producing an actuation unit.

quantitatively. Secord and Asada [7] showed that in a chain of amplified piezoelectric stacks with mechanical stops, the choice of *which* "cell" to activate, while irrelevant to force, affected the resonant frequency of the device. The shift in resonant frequency is due to a change in mass distribution rather than a change in stiffness of the chain.

In the actuator of the type studied in this work, modifying the impedance is produced by exploiting the redundancy in control inputs. The variable impedance behavior is innate to the technology and requires *no additional hardware*. Fig. 1 illustrates how given a total activation level, modifying *exactly which* motor units are in the "on" state results in different impedance properties. To the authors' knowledge, this is the first example of modulating actuator impedance by varying the spatial properties of activation.

The paper is organized as follows: Section II and III introduce the basic paradigm. Section IV discusses the relationship between resolution and the number of motor units, Section V and VI explain hardening behavior (nonlinear stiffness) and the critical length (where hardening begins).

*This work was supported by NSF CNIC 1427787, ERC-291166-SOFTHANDS and ERC-337596-SPEAR

[1]Department of Mechanical Engineering, The University of Tulsa, Tulsa, OK 74104, USA joshua-schultz@utulsa.edu

[2]Department of Mechanical Engineering, Vrije Universiteit Brussel, B-1050 Brussels, Belgium glenn.mathijssen@vub.ac.be, bram.vanderborght@vub.ac.be

[3]Department of Advanced Robotics, Istituto Italiano di Tecnologia, 16163 Genova, GE, Italy

[4]Centro E. Piaggio, Università di Pisa, 56112 Pisa, PI, Italy bicchi@centropiaggio.unipi.it

Section VII-IX complete the formulation and close the paper.

II. Discretized actuation based on recruitment

The authors introduce a modular test-bed in a prior work [8], of a muscle-like basic building block that can be used to study myogenesis and recruitment concepts in a robotic context, shown in Fig. 2. It has as its motor units 6 miniature solenoids whose plungers are connected by individual series elastic spring lobes to a common hub; each solenoid acts as a discrete *displacement source* which produces a discrete force contribution to the total by virtue of its an integrated series elastic element. In that work, two terms related to actuation systems composed of discrete units are introduced:

- **actuation unit**: the smallest unit that can be added, removed or reconfigured to adjust a muscle-like actuator's characteristics (a manufactured part)
- **motor unit**: a force-producing device that can be independently activated or deactivated by a control signal

In robotic applications, it is envisioned that the input to the actuator will be a scalar feedback control signal $u(t)$, which provides only the *number* of motor units that should be recruited during each sample, not *which* ones. As these actuation systems may be large, it would be helpful to have a quick way to determine what constitutes a "favorable" vs. "unfavorable" spatial distribution of active motor units. Activation should be chosen to produce *robustness to small changes in length* and *desired impedance*.

The selective recruitment strategy described in this paper is equally applicable to any technology (such as the piezoelectric cellular actuator [7] or shape memory alloy implementations [9]) that fits the paradigm illustrated schematically in Fig. 3: A motor unit consists of a binary contractile element with a series elastic element, with all series elastic elements joined on the opposite end at a common hub.

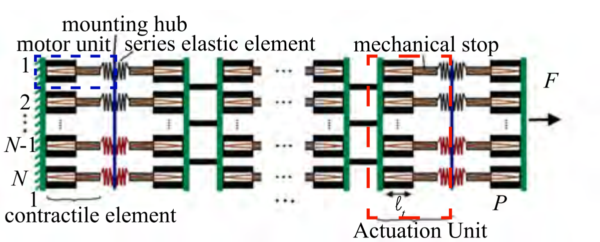

Fig. 3: Schematic of the muscle like actuator

With any series elastic actuator, changes in the load will produce changes in the length of the series elastic element, which in turn produces changes in the force. Mechanical stops (see Fig. 3) limit the contractile element's outward travel under load, preventing disengagement. Secord and Asada [7] place the stop on the mounting hub, meaning inactive units become rigid. In this paper, the mechanical stop is on the contractile element, meaning that the actuation unit conserves its spring-like behavior when the mechanical stop is engaged. By implementing the mechanical stops this way, the actuation unit's stiffness varies with activation, giving variable impedance behavior.

III. Model of force-length behavior

Schultz and Hawrylak [10] developed a model for series-elastic units with binary operation configured in a bundle of chains. If the units are identical (same spring constant, same contraction difference) and the length of the bundle L is known, a closed-form expression for the force F produced as a function of activation is as follows:

$$F = \frac{k}{n}\left[nL - \sum_{j=1}^{m}\sum_{i=1}^{n} \ell_{i,j}^{s}\right], \tag{1}$$

k is the spring constant of each series elastic element, n is the number of units in each chain, m is the number of chains in the bundle. In addition to being summation variables, i and j represent the row and column indices of a given unit in the bundle. $\ell_{i,j}^{s}$ is a binary constant that depends on the activation state: it can be either ℓ^0, the inactive resting length, or $\ell^a < \ell^0$, the active resting length.

Looking at Equation (1), one can see that if all units in the bundle are of identical characteristic, it is only the total *number* active that matters; *exactly which* units are recruited (made active) will not affect the force. A similar closed-form expression exists which allows k, ℓ^0, and ℓ^a to vary for each node. Equation (1) does not handle the following phenomena present in the physical implementation:

1) each actuation unit contains multiple motor units
2) the actuation units behave differently in extension vs. contraction
3) length-based switching condition due to mechanical stops

These phenomena introduce the notions of hardening and critical length, discussed in sections Section IV-Section VIII. Each actuation unit will belong to one of the three classes below. The following assumptions are made:

1) if active, the external load will not overpower the motor unit (internal contraction)
2) the mechanical stops never break
3) there is always some preload, however small, placing the actuator in tension

A. Class I

An actuation unit in Class I has all of its motor units inactive. For a single chain, this means that the mechanical stop is engaged. The resting length is then ℓ^0, and the stiffness is that of all of its springs in parallel. A small modification to treat bundles of chains will be addressed in Section VII.

B. Class II

An actuation unit in Class II has at least one motor unit active; the inactive ones "float" – they slide or change length freely and exert no force. The length of the actuation unit in question is between ℓ^a and ℓ^0. Conceptually, the actuation unit becomes a single series elastic element with a stiffness equal to all the *active* motor units' springs in parallel, with a resting length of ℓ^a. If an actuation unit has all motor units active, it is always in Class II, without exception.

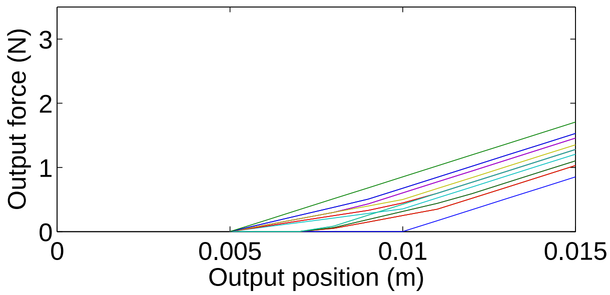

(a) Muscle-like actuator with N=2 and P=2.

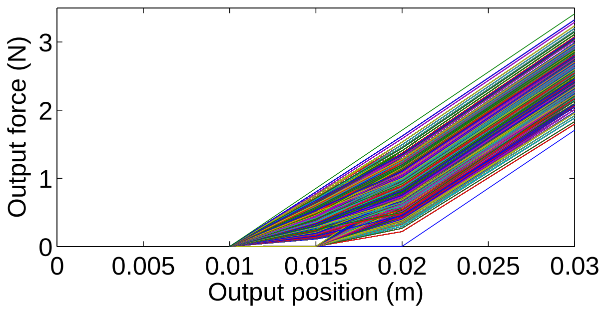

(b) Muscle-like actuator with N=4 and P=4.

Fig. 4: Possible output forces as a function of the output position for two muscle-like actuators with different number of motor units per actuation unit (N) and various numbers of actuation units in series (P). Each line represents a different activation pattern.

C. Class III

An actuation unit has at least one motor unit active, and has been elongated substantially by the external load. Inactive motor units do not float, but encounter the mechanical stop, meaning they also impose a spring force. In Class III the length of the actuation unit is greater than ℓ^0.

IV. Precision of Actuator Force Produced

Although the muscle-like actuator (i.e. a configuration of actuation units in parallel and series) can only be fired discretely it still offers numerous possible force-displacement operating points. The number of unique discrete output forces that can be produced depends on the number of motor units per actuation unit (N) and the number of actuation units in the configuration (P). This is shown in Fig. 4 where each point of each line represents a position and force that can be achieved by the actuator for some activation pattern M on each actuation unit. This is simulated for an actuator with N=P=2 and for N=P=4. Even with only a modest increase in N and P, the density of operating points produced increases greatly.

A. Increasing the length of the chain

The Force-displacement-activation relationship for an individual actuation unit was given in [8] and is repeated here:

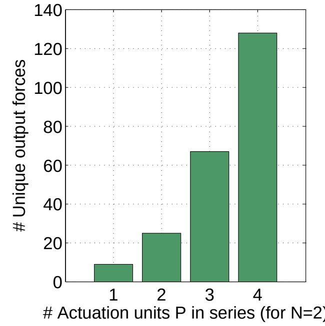

(a) Unique discrete forces with increasing chain length

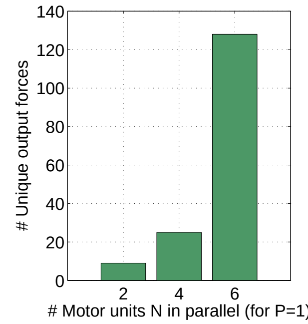

(b) Unique discrete forces with increasing motor units

Fig. 5: The number of unique output forces that can be produced by the actuator increases strongly with increased chain length P, as well as the number of motor units per actuation unit M. The output force calculated by the solver was rounded to 4 digits after the decimal point.

$$F = \begin{cases} (\ell_t - x) \sum_{i=1}^{M} k_i, & x \geq 0 \\ (\ell_t - x) \sum_{i=1}^{M} k_i \quad - \quad x \sum_{j=1}^{N-M} k_j, & x < 0 \end{cases} \tag{2}$$

where k_i, k_j are the spring constants from each motor unit to the mounting feature, x is the position of the mounting feature and ℓ_t is the contraction distance of a motor unit. M is the number of motor units active within an actuation unit.

Fig. 5a shows the force produced where the number of motor units N per actuation unit is 2, and the number of actuation units P in the chain increases from 1 to 4.

B. Increasing number of motor units per actuation unit N

The number of unique discrete output forces that can be produced by one actuation unit depends on the number of motor units within an actuation unit. This is shown in Fig. 5a where P is 2, and the number of solenoids in parallel N increases from 2 to 6.

C. Occupation rate of actuation range

The *occupation rate* is a metric that characterizes the variation in force production with activation patterns. We define it as the number of unique output forces divided by the maximum number of output forces in a certain range to within a certain tolerance. An example can be found in Fig. 6. In this example the output range is indicated by the red stars and equal to 1.02 N. The accuracy desired is 0.001 N. The number of unique output forces is 340. As such, the occupation rate is equal to $\frac{340}{1020}$=33 %. If a motor unit or entire actuation unit fails, the new occupation rate can be calculated ommitting the damaged portion, making it a measure of damage tolerance.

V. Variable Impedance: Deliberately Placing Actuation Units in Class I, II, or III

The analysis of this section permits us to quickly and easily categorize actuation units into class I, II, or III,

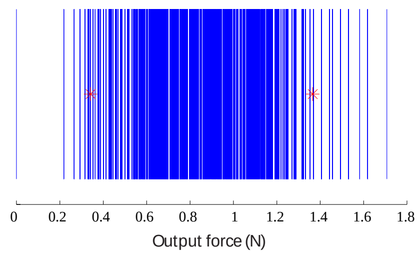

Fig. 6: The occupation rate is studied for N=4 and M=4 and output position 0.02 m. Between the two red stars the occupation rate is 33% for an accuracy of 0.1% of the maximum output range.

based on only the activation state (easily stored), and the overall actuator length (easily measured). Using this method, calculation of the individual lengths of each actuation unit is not necessary, which is useful when considering large configurations.

Since there are many different activation patterns that can produce the same force, activations can be distributed over actuation units in Class II in different ways so as to vary the impedance. Increases in the load will cause the elastic elements to extend. As the length increases, individual units will move from Class II into Class III, and the actuator will exhibit hardening behavior. Hardening of a series elastic actuator has been found to be useful in robotic locomotion [11].

In most practical implementations, the physical characteristics of the actuation unit, its activation state, and the length of the actuator L will be readily available. Determination of which actuation units are in Class I is trivial. Dividing units into classes II and III is more difficult.

There is a strong motive to avoid Class III: if all the units are Class I or Class II, the force can readily be calculated in closed form [8] with knowledge of L and the activation state. This is done by treating each *actuation unit* as a node with the equivalent resting length and stiffness as described in Section III.

It is also prudent to consider what happens with small changes in length of the actuator. If a large number of units are near the Class II/Class III boundary, there may be "chatter" as units rapidly move between the two classes, meaning large changes in force and impedance come from only small changes in displacement.

VI. Distinguishing between Class II and Class III

The number of units in Class II and in Class III can be calculated if the activation state of each and every unit is known. However, the critical length, L_{crit}, the length of the bundle L at which the first actuation unit transitions from Class II to Class III, can be estimated with less information. We will begin by analyzing a single chain and extend the analysis to a bundle of chains.

A. *Concept of active and inactive actuation units*

Consider a chain under load of n actuation units with the activation state of all motor units known. p denotes units of that chain that have *at least one* motor unit active. The length of each individual actuation unit does not depend on the order in which they are placed, so the chain is equivalent to two chains in series; one p units long of "active" units and the other $n-p$ units long of "inactive" units. Any unit in the same chain with the same stiffness will be in the same class. If all units are identical, this means that all units in the chain with the same activation state will be in the same class. This is why calculating each individual actuation unit's length is unnecessary.

The goal of this analysis is to quickly estimate L_{crit}. In order to contact the mechanical stops and enter Class III, at least one active actuation unit will need to elongate from a length of ℓ^a to a length of ℓ^0. The stiffness of active actuation units, while below that of inactive units, is not drastically so, meaning it is reasonable to assume that the inactive chain undergoes *some* deformation when the chain is stretched to length L.

The stiffness of an active actuation unit equals the stiffness of all springs of its *active* motor units in parallel; If it is in Class II, those from inactive motor units do not contribute. Therefore, for $L < L_{crit}$, the actuation unit with the lowest activation state has the minimum stiffness. It is tempting to estimate L_{crit} by setting the stiffness of the active chain to be the minimum possible stiffness, say, assuming that every unit in the inactive chain has only one unit active. After all, this means that the inactive chain would deform by the least amount, forcing the active chain to its longest possible length. This has a tendency to overestimate L_{crit}, particularly for higher activation levels. The reason is thus: if all the active units are at minimum stiffness, the deformation is borne equally by each unit in the active chain, decreasing the likelihood that any of them will go into Class III. It turns out that heterogeneous activation patterns, where some units are stiffer than others, will force the more compliant actuation units into Class III at shorter overall actuator lengths, because these "weak links" undergo most of the total deformation. This will be demonstrated below.

B. *The 3-spring model*

With a modest increase in complexity, the critical length can be calculated accurately in closed form without calculating the individual lengths of each unit in the chain. This model assumes several things:

1) actuation units with the same activation level in a given chain will have the same stiffness
2) actuation units with the same activation level in a given chain will consequently deform by the same amount
3) actuation units with the lowest activation level in a given chain will have the lowest stiffness, and therefore deform by the largest amount

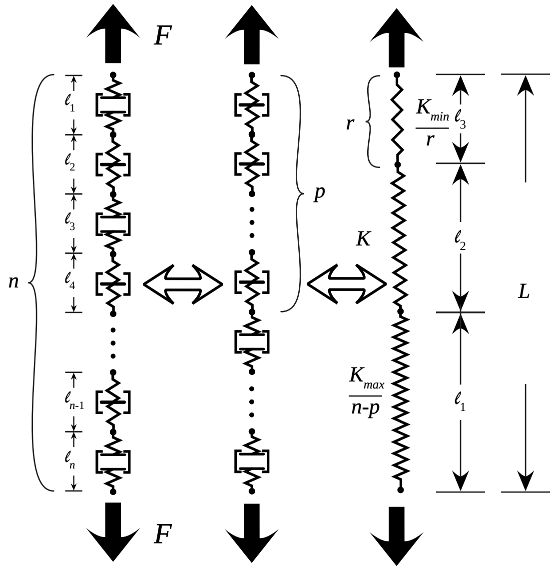

Fig. 7: Since the geometric distribution of activation in a chain of actuation units has no bearing on what class they will fall into they can be reordered as shown. A given chain in an actuator can be divided into 3 sub-chains. Stiffnesses, number of units in the chain, and the lengths of each are shown.

4) as such, these will be the first units to reach Class III with increases in length, and all are equally likely to enter Class III for a given length

To determine whether a chain is in danger of entering Class III, it is sufficient to consider only those actuation units that are at the *minimum* activation state for the chain. Suppose that the chain contains r of these units. This divides the active chain into two "subchains", or the 3-spring configuration shown in Fig. 7. Again we can assume that all r units are on the end of the chain. Each of these r units will have a stiffness K_{min}. The stiffness of the subchain is then K_{min}/r. There will be $p-r$ actuation units that have a greater number of motor units active. Collectively, *this* subchain will have a stiffness K, resulting from the remaining actuation units with non-uniform activation levels.

Applying equilibrium at the two nodes where the subchains join and noting that the total length must sum to L results in the following system of equations:

$$\begin{bmatrix} -\frac{K_{\max}}{n-p} & K & 0 \\ 0 & -K & \frac{K_{\min}}{r} \\ 1 & 1 & 1 \end{bmatrix} \begin{bmatrix} \ell_1 \\ \ell_2 \\ \ell_3 \end{bmatrix} = \begin{bmatrix} \ell_a K_{\min} - \ell_0 K_{\max} \\ \ell_a \left(K_{\min} + K(r-p)\right) \\ L \end{bmatrix} \quad (3)$$

If ℓ_3 reaches a length of $r\ell_0$, all r of these units will be in Class III. The expression for ℓ_3 can be found using Cramer's rule. Setting this equal to $r\ell_0$ and solving it for L gives the critical length.

$$L_{crit} = (\ell_0 - \ell_a)\left(\frac{K_{min}(n-p)}{K_{max}} + \frac{K_{min}}{K} + r\right) + p\ell_a + \ell_0(n-p) \quad (4)$$

K can be calculated directly from the activation state of all actuation units in the chain, and r from a simple search. It is worth noting that the last two terms in Equation (4) are equal to the resting length of the actuator for that activation pattern.

If one is unwilling to query the entire activation state of the chain, substituting $K = K_{max}$ and the lowest possible value r_{lb} that r *could* be gives a lower bound for L_{crit}. This only requires knowledge of p and K_{min}.

Fig. 8 shows the critical lengths by evaluating Equation (4) for a chain of 36 actuation units of 12 motor units each with a series of 10000 random activation states, plotted against the total activation level for the chain. It is clear that *even with the same total number of motor units active* very different values of L_{crit} can be produced. Thus if one wants to remain in Class II at all costs, it is advantageous to choose the activation pattern that results in the highest possible L_{crit}. Conversely, if hardening behavior associated with reaching the safety stops is desirable, a pattern could be chosen with a lower L_{crit}.

Fig. 8 shows that the lower bound is reasonably tight, making it an acceptable conservative estimate of the critical length. Note that the critical lengths decrease with increased activation. This stems from the decreased resting length of the actuator with increased activation. The critical lengths are separated into a few parallel bands (naturally with some scatter), the upper bands likely being cases where the activation happens to be distributed more evenly. The majority of activation patterns' critical lengths end up in the lower bands. These upper bands tend to disappear at higher activation levels, suggesting that there are fewer options which yield the longer critical length. For higher activation levels, it is more likely that there will be stiffer actuation units in the active chain, forcing the more compliant ones to absorb the majority of the deformation and they hit Class III at shorter lengths.

VII. Bundles of Chains

A single chain will seldom produce enough force for a given application, so multiple chains will be collected into parallel bundles in order to move a load. The force produced by each chain is additive. In addition, the stiffness of each chain is additive.

Since the actuator is a bundle of chains joined at the ends, each chain has the same length, L. To determine whether the actuator has any units in Class III, one simply checks each chain one by one. The critical length is therefore:

$$L_{crit}^{total} = \min_j L_{crit}^j \quad (5)$$

where L_{crit}^j is the critical length of chain j.

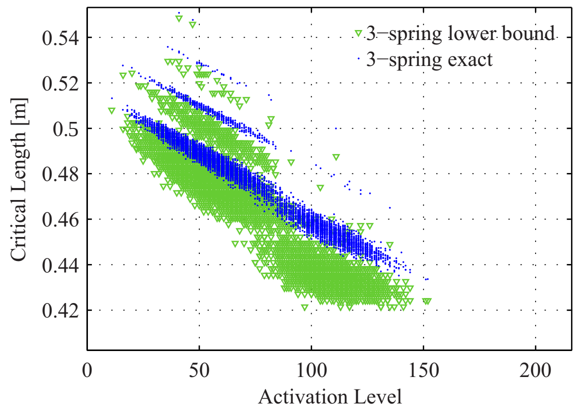

Fig. 8: Critical lengths vs. number of motor units active for a randomly-generated set of activation patterns

Unlike the single chain case, there is also a critical length on the *lower* side. Imagine that chain j has a particularly high activation state, and chain k has a particularly low activation state ($p << n$). In this case, chain j could pull L below $p_k\ell^a + (n_k - p_k)\ell^0$, where the subscript k refers to the values for chain k. In this case the length of chain k could be below its resting length. This chain will behave like a *slider*. Further activation of this chain will not contribute to the force until chain k "catches up" with the more active chain or the load increases, which will cause L to go up. The user may want to avoid this situation by distributing the activation more equally across the chains in the bundle. Conversely, this behavior may rather be exploited to control the stiffness during operation.

VIII. Controlling the Impedance and Computing Actuator Force

In order to determine the force-length behavior of a muscle-like discretized actuator, the first step is to determine whether there are units in Class III or not. This can be estimated reasonably accurately using only the total number of actuation units in each chain p that have at least one motor unit active, and the minimum activation state in the chain. If every actuation unit in the entire actuator is in Class II, the Schultz-Hawrylak formulation can be used to determine the actuator force. Once L_{crit} is exceeded, the actuator will continue to harden at discrete intervals with increasing length as more units enter Class III. The calculation of force is less elegant, involving a matrix multiplication. It is the nonlinearity of the mechanical stops and movement from one class to another that gives variable impedance behavior. Activation patterns can be modified while still preserving the same overall activation level for the actuator to vary the actuator's stiffness by trading off between Class I and Class II Class II or forcing actuation units into or out of Class III to exploit the hardening behavior.

IX. Conclusion

Discretized muscle-like actuation has obvious benefits with regard to resiliency, but also posesses innate length-dependent variable impedance behavior. By modulating exactly which units are activated for a scalar control input, the impedance of the actuator can be specified while conserving the same overall *number* of units active. This gives the controller the ability to invoke or revoke hardening behavior (nonlinear stiffness) along with specifying a general control input. This paper has presented a closed-form expression based on system constants and activation state to accurately predict the critical length where hardening begins. Even when only a few actuation units are present, the force-displacement operating points quickly become dense despite the binary nature of individual motor units.

Future work will conduct force-length-activation experiments on the hardware shown in Fig. 2 to see how accurately the models perform in the presence of manufacturing uncertainty and friction. Models of more complex configurations will be evaluated and models produced as well. Methods will be sought for more efficient calculation of the actuator force beyond the critical length.

References

[1] B. Vanderborght, A. Albu-Schaeffer, A. Bicchi, E. Burdet, D. Caldwell, R. Carloni, M. Catalano, O. Eiberger, W. Friedl, G. Ganesh *et al.*, "Variable impedance actuators: A review," *Robotics and Autonomous Systems*, vol. 61, no. 12, pp. 1601–1614, 2013.

[2] J. Hurst, J. Chestnutt, and A. Rizzi, "An actuator with physically variable stiffness for highly dynamic legged locomotion," in *Robotics and Automation, 2004. Proceedings. ICRA '04. 2004 IEEE International Conference on*, vol. 5, April 2004, pp. 4662–4667 Vol.5.

[3] R. Schiavi, G. Grioli, S. Sen, and a. Bicchi, "VSA-II: a novel prototype of variable stiffness actuator for safe and performing robots interacting with humans," *2008 IEEE International Conference on Robotics and Automation*, pp. 2171–2176, May 2008.

[4] A. Jafari, N. G. Tsagarakis, and D. G. Caldwell, "Adjustable Stiffness (AwAS)," *IEEE/ASME Transactions on Mechatronics*, vol. 18, no. 1, pp. 355–365, 2013.

[5] R. Van Ham, B. Vanderborght, M. Van Damme, B. Verrelst, and D. Lefeber, "MACCEPA, the mechanically adjustable compliance and controllable equilibrium position actuator: Design and implementation in a biped robot," *Robotics and Autonomous Systems*, vol. 55, no. 10, pp. 761–768, Oct. 2007.

[6] X. Du, R. Dixon, R. Goodall, and A. Zolotas, "Modelling and control of a high redundancy actuator," *Mechatronics*, vol. 20, no. 1, pp. 102–112, Feb. 2010.

[7] T. W. Secord and H. H. Asada, "A variable stiffness actuator having tunable resonant frequencies," *IEEE Transactions on Robotics*, vol. 26, no. 6, pp. 993–1005, 2010.

[8] J. Schultz, G. Mathijssen, B. Vanderborght, and A. Bicchi, "Toward Motor-Unit-Recruitment Actuators for Soft Robotics," in *Proceedings of the IEEE International Conference on Biomedical Robotics and Biomechatronics*, São Paulo, 2014, pp. 887–892.

[9] K. J. Cho and H. Asada, "Multi-axis SMA actuator array for driving anthropomorphic robot hand," *Proceedings - IEEE International Conference on Robotics and Automation*, vol. 2005, no. April, pp. 1356–1361, 2005.

[10] J. A. Schultz and P. J. Hawrylak, "Modular Actuation Systems: A Scalable Solution for Delivering Robotic Performance," in *Proceeding of the Robot Makers Workshop: The Future of Digital Rapid Design and Fabrication of Robots (RoMa), Robotics: Science and Systems Conference*, Berkeley, CA, 2014, pp. 1–2.

[11] M. Khoramshahi, H. Jalaly Bidgoly, S. Shafiee, A. Asaei, A. J. Ijspeert, and M. Nili Ahmadabadi, "Piecewise linear spine for speedenergy efficiency trade-off in quadruped robots," *Robotics and Autonomous Systems*, vol. 61, no. 12, pp. 1350–1359, Dec. 2013.