Design Verification and Optimization of Novel Parts Feeders with the Impulse Dynamic Simulator

Dan Reznik* John Canny
dreznik@cs.berkeley.edu jfc@cs.berkeley.edu
EECS Dept., UC-Berkeley
Berkeley, CA 94720-1776 USA

***Abstract*—**
The Impulse dynamic simulator is used to generate a wide range of synthetic measurements for both a MEMs- and a longitudinally-vibrating parts feeders. Impulse allows the devices' mechanics to be modeled in detail and for results to be generated both accurately and efficiently. Obtaining such results experimentally is impractical given the number of moving parts, small dimensions, and/or complexity of part feeder interactions. Output simulation data are used to (i) debug existing problems; (ii) measure feeding performance and compare it with analytical/experimental results; (iii) optimize the design for higher feed rates; and (iv) generate simple dynamic models of their behavior.

I. Introduction

The Impulse dynamic simulator [1] is used to generate a wide range of synthetic measurements for two novel parts feeder designs. The first design is a MEMs-based actuator array called the M-Chip [2] with a surface area of a few square-cm; the second design is based on the vibratory feeder recently proposed in [3], intended for feeding at the "macroscopic" scale. Both devices are designed to convey a *part* resting on their surface along a straight line, and with a constant velocity called the *feed rate*. For the M-Chip, the part's forward motion is caused by horizontally-biased impacts against hundreds of oscillating microactuators. In the vibrations-based feeder, forward conveying is the result of frictional interactions between the part and a supporting longitudinally-vibrating plate.

The design of part feeders has remained a "black art" due to difficulties in the analysis and measurement of part-feeder interactions [7]. Experimental measurements of the device's "dynamical state" during a typical feeding run, e.g., part position/velocity, the location/magnitude of collisions, frictional forces, etc., though key in the optimization of a given design can be very difficult to obtain. This is particularly true for a device such as the M-Chip with its reduced dimensions and large number of moving parts. The goal of this work is to use dynamic simulation to generate dynamic measurements during a "virtual run" of either feeder, leading to insights on the quality and/or efficiency of their design. In particular, such simulation output could help with the following:

1. Design Verification: does the device perform as expected in both qualitative and quantitative terms?
2. Debugging: are there unpredicted glitches in the original design which decrease (or totally shut off) performance? Any clever fixes?
3. Optimization: Are there changes in the design which could lead to better performance?
4. Dynamic Characterization: can the complex interactions between part and moving (vibrating) substructure be reduced to a simple bulk dynamic model?

Dynamic simulation is ideal in addressing the above questions since it eliminates the need for experimental measurements; using it for the purposes of design is a fairly novel idea [4], mostly due to the numerical instability and inefficiency of existing simulation packages. The Impulse Tool is particularly convenient since it is the first remarkably efficient and accurate package in dealing with a large number of moving/interacting rigid bodies. Its impact-based collision model has been shown to work well on both vertex- [5] and face- (sliding) [6] types of part-surface interaction.

The remainder of this paper is organized as follows: Section II describes the principle of operation and simulation experiments performed with the M-Chip. These involve identifying design problems, trying new microactuator shapes, and dynamically characterizing the device. Section III describes the vibrations-based parts feeder, presents some performance results derived in [3], comparing these to results obtained from dynamic simulations with Impulse.

*Support provided in part by NSF Grant FD93-19412.

II. A MEMs Parts Feeder

A picture of the fabricated MEMs parts feeder – the "M-Chip" [2] – is shown in Figure 1(a). The device contains an array of approx. 10,000 micro-actuators, called *resonators*, tiled over a few square-cm of silicon substrate. All resonators are oriented towards a single direction and tiled in interleaved fashion, as shown by the electron micrograph in Figure 1(b). Each resonator is a rectangular slab of silicon supported a few μm above the substrate by a torsional rod. A metallic electrode installed underneath half of the resonator's surface is used for torsional actuation, as shown in Figure 2(a). Typically, resonators are driven at a few kHz. One end of the resonator is equipped with a ridge of several vertical poles. This asymmetry in design generates anisotropic impact forces on a small part placed over the array, causing it to feed with constant velocity towards a unique direction [2].

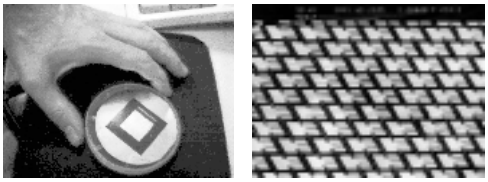

(a) (b)

Fig. 1. (a) A photo of Böhringer's hand holding his M-Chip, containing approx. 10,000 resonators. (b) Electron micrograph of a portion of the array showing the interleaved tiling of resonators.

A. Device Modeling

In modeling the M-Chip with Impulse, the focus was to preserve original nominal parameters and dimensions as closely as possible. We used the original resonator mass, geometry, and oscillation frequency as given to us by the designers. We started out by creating a 280x180x5 μm geometric model for the resonator, shown in Figure 2(a). To simplify collision detection, we modeled the set of poles installed on one extreme of the resonator collectively as a single 5 μ-high ridge. The resonator body[1] was modeled as a rectangular slab. Electrostatic actuation was replaced by (i) modeling the supporting rods as a single spring-loaded revolute joint, and by (ii) a torque control-law acting at that joint. The joint's spring and damping constants were chosen according to the torsional elasticity of silicon. The torque wave applied to the joint is a 5 kHz positive square wave, whose amplitude was chosen to produce positive oscillations of approx. 5^o (0.08 rad). Figure 2(b) shows the driving torque law superimposed on θ, the free oscillation angle, which is roughly a 5 kHz positive sinewave.

[1]Fabricated resonators posses a grating of holes, but these have no effect on collisions with the part.

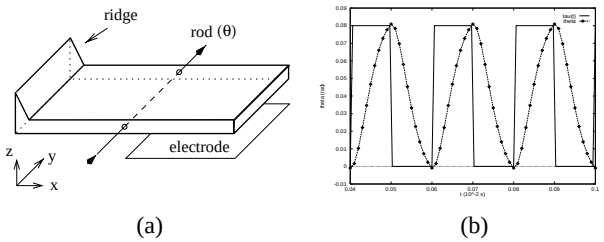

(a) (b)

Fig. 2. (a) The resonator is modeled as a rectangular slab with a ridge at one end. The supporting rods are modeled as a spring-loaded, damped revolute joint. The electrostatic actuation (provided by the electrode) is replaced by a torque control law at the joint. (b) The torque applied to the resonator is a 5 kHz positive square wave. The torque (not shown in scale) is superimposed with $\theta(t)$, the resonator angle.

The next modeling step was to assemble the resonators in interleaved fashion, as in Figure 1(b). Though the actual device contains over 10,000 resonators, simulating this many moving parts is impractical with Impulse. However, by considering an 8x5 sub-array of resonators, and a rectangular array-aligned part with silicon's density, we are can remain physically consistent with the full size experiment. The final Impulse model shown as a 3D rendering is depicted in Figure 3. Notice that the part's footprint covers approx. 10 resonators, while in actuality it would cover 100s.

Fig. 3. 3D rendering of the complete model showing a 1.15x0.4x0.25 μm, 550 μg part resting at its initial position over an 8x5 resonator array.

The Impulse simulator resolves collisions between rigid bodies based on a sophisticated discrete-event impulse model. Two global parameters need to be set: the friction coefficient μ, and the restitution coefficient ϵ. These were set to the physically reasonable (and numerically-stable) values of 1.0 and 0.5, respectively. Experiments to evaluate how these parameters affect array performance have not yet been tried, though some theoretical results have been derived for vibratory feeders [7].

B. Design Optimization

The first simulated experiment performed with the array was to drive the resonators and simply drop the part on it. Figure 3 shows the part at its starting position for this experiment.

1) Part jamming: Once dropped on the array, the part is propelled forward at a constant speed of about 0.8 mm/sec, a value in close match with experimental results performed by the designers. A problem quickly discovered was that the part would jam (i.e., stop its forward motion) as soon as it encountered a new row of resonators, as shown in Figure 4(a). Jamming occurs since the part's vertical hopping is not high enough to allow it to skip over the ridges of the next row of resonators. As the part is driven up, the next row of resonators is also doing so, and the part bounces back.

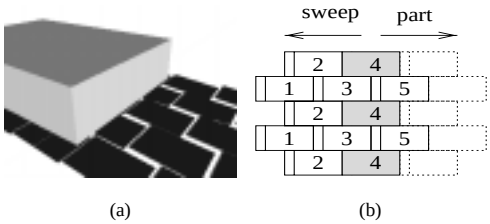

Fig. 4. The jamming problem: in (a) the part is shown unable to skip over the ridges of the encountered row of resonators. In (b) the row-strobing method is illustrated. Consecutive rows in the array are labeled from 1-8 in the direction of the part's motion (only the first 5 rows are shown). The process involves selectively turning off consecutive rows of resonators for a few oscillation cycles, in the direction opposite to the part's motion.

The first attempted solution to eliminate jamming was to drive consecutive resonator rows at different phase offsets, so that ridges in the blocking row would be going down while other resonators would be driving the part up. This idea did not work since the out-of-phase impacts occurring under the part cause lower part hops which in turn aggravate the problem of skipping a set of blocking ridges. A solution which proved successful is called *row-strobing*, illustrated in Figure 4(b). The idea is to propagate a wave of off resonators in the direction opposite to the part's motion. Label all rows from 1 to N along the array's feeding direction. The first row to be turned off is row N. That row is left in the off state during γ complete oscillation cycles, at which point it is turned back on. The process continues with row $N-1$ being off for γ cycles, and so on until row 1 is reached. At this point, the wave wraps around and the process restarts at row N. Since part position is not known, any resonator row is potentially causing jamming, thus the need for a sweeping wave (this could be alleviated with sensing). To ensure that the jamming row is located faster, the wave is propagated against the feeding direction so the relative speed between the wave and the part is higher. Choosing too small a γ may not give the part enough time to be pushed over the blocking ridges, however the average jam clearing time is proportional to γ. We found that $\gamma = 50$ oscillation cycles gave the best results. In the case of a very long array (as in the real device) the blocking row can be found faster by propagating several off wavefronts separated by a constant number λ of resonator rows. If λ is too small, too many off rows will lie under the part at any given time, reducing the feed rate considerably. In our case we set $\lambda = 4$ rows, i.e., in our 8x5 array there will be, at any given time, two off rows moving in the $-x$ direction. Since the part is about 4 resonators long, the part receives 1/4 less impacts at any given time.

As shown in Figure 5, the row-strobing method results in a motion of the part characterized by periods of constant forward feed rate (of approx. 0.8 mm/sec) interleaved with short jamming phases, where the feed rate is null.

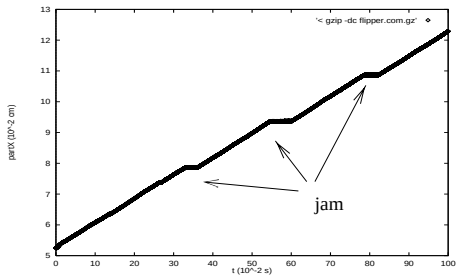

Fig. 5. Plot of the part's forward displacement as a function of time. As shown, the part's motion alternates between constant forward motion and short periods of jamming.

2) Double bouncing: The simulator treats collisions as discrete events; each collision is followed by a calculation which computes an impulse force applied at the collision point and in opposite direction to the bodies colliding. We looked at the stream of collisions between the part and the resonator located on the 2nd row and 3rd column of the array during the first non-jamming period of Figure 5, namely, for $0.05 < t < 0.2$ sec. For every such collision we recorded the associated (i) resonator angle θ, (ii) angular speed $\dot{\theta}$, and (iii) F_x, the x component (i.e., along the feeding direction) of the impulse force calculated by the simulator. Figure 6 shows a scatter plot of the $(\theta, \dot{\theta})$ pairs gathered. These pairs cluster into two separate clouds pointed to by the arrow labeled *loaded*. As expected, one of the clouds (signaled by the *free* arrow) follows quite closely the θ vs. $\dot{\theta}$ relation for the resonator's free oscillation – this is an ellipse since the free oscillation is roughly sinusoidal. However, the second cluster of

points is anomalous.

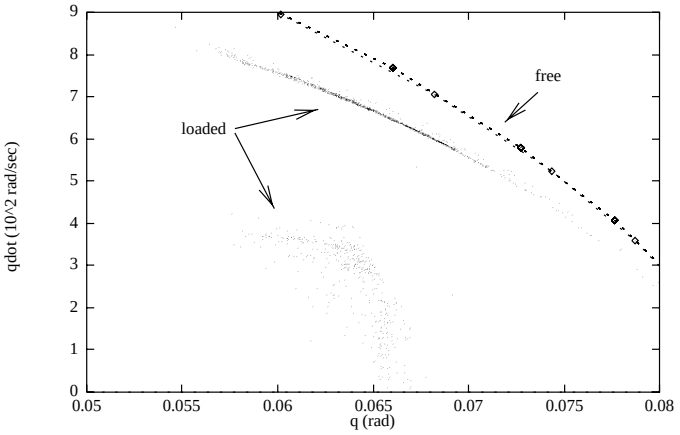

Fig. 6. Scatter plot correlating resonator angle θ (x-axis) w/ the angular speed $\dot{\theta}$ (y axis). The *free* arrow shows this relationship for free oscillation – roughly an ellipsoid. The *loaded* arrow points to $(\theta, \dot{\theta})$ pairs generated by individual collisions when the part is over a chosen resonator. Some of the loaded points follow the free oscillation curve, while others lie in an anomalous region of low values of $\dot{\theta}$.

What phenomenon could be generating the lower cloud of points? That cluster indicates that many collisions are occurring at an angle θ with a much slower $\dot{\theta}$ than that of free oscillations. We found that these collisions were being caused by secondary bounces of the resonator on the part occurring shortly after a normal 5 kHz bounce. To show that, we correlated θ, $\dot{\theta}$, and F_x with the collision's *inter-arrival time*, which measures the time elapsed between the current collision and the one last occurring in the simulation. This is shown in Figure 7(a,b,c). Though one expects collisions with a single resonator to be spaced by 1/5 kHz = 0.2 ms (i.e., the part receives one impact per resonator cycle), the graphs show that collisions cluster over .02, .18, and .2 ms interarrivals. The collisions occurring at .02 ms after the normal .2 ms ones are *double bounces*; after the first collision, the ridge is still being driven upward by the torque control law, causing the resonator to ricochet one or more times against the part. After the double bounce, the resonator will tend to resynchronize with the driving square wave, so that the next collision occurs within .18 sec: $0.18 + 0.02 = 0.2$ ms. The (a) plot shows that the .02 collisions occur at lower values of θ than the normal .2 ms ones. This indicates that a downward motion of the part favors the double bounces. The (b) plot shows that double bounces occur at much slower $\dot{\theta}$ than normal bounces, implying that the former transmit less impact energy to the part than the latter. The (c) plot shows that most double bounces are associated with an impulse pointing in the negative feeding direction, i.e., they act as brakes!

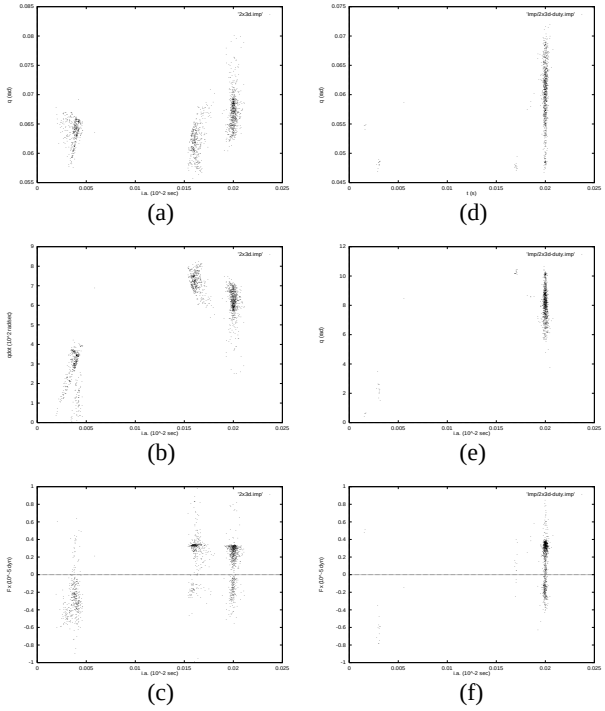

Fig. 7. Graphs (a,b,c) show θ, $\dot{\theta}$, and F_x (plotted along y) clustering over three specific collision inter-arrivals (plotted along x): .02,.18, and .2 ms. Graphs (d,e,f) show the re-clustering of the same quantities, this time solely over 0.2 ms inter-arrivals, when a shorter duty cycle is used.

Double bounces were eliminated by reducing the *duty cycle* of the square wave torque driving the resonators from 50% to 36%. This shuts off the driving torque law just before a double bounce is likely to occur (i.e., .02 ms after the average time normal collisions occur). The results of this change are shown in Figure 7(d,e,f), which side by side with the previous plots illustrate how the cloud of .02 and .18 ms collisions coalesce into a single cluster over 0.2 ms inter-arrivals. These graphs also show that the new impulses occur in average at a higher value of $\dot{\theta}$, i.e., they transfer more momentum at every collision. This simple reduction in duty cycle increases the feed rate from 0.8 to 1.0 mm/sec, i.e., a 30% improvement. As an interesting note, the points in Figure 7(f) are split evenly in the positive and negative F_x range, indicating that the total force applied to the part over the period considered is zero – the part is feeding forward at a constant speed.

3) Optimizing the resonator's shape: The feed rate ν was measured against three parameters: (i) the part's mass M, (ii) the ridge's distance L from the resonator's midpoint, and (iii) the ridge height H. The last two parameters are illustrated in Figure 8.

The plots in Figure 9 show the results of these experiments. Plots (a,b,c) show, respectively, results from the

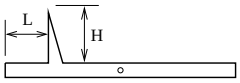

Fig. 8. The shape of the resonator is parametrized according to ridge distance L, and ridge height H.

mass, ridge distance, and ridge height experiments. The x-axis labels the parameter being varied; the feed rate is plotted along y. The collision angle θ, angular velocity $\dot{\theta}$, and the height of the part's center of mass PartZ (averaged over an entire feeding task) are superimposed over the feed rate; their numeric values have been omitted for the sake of clarity.

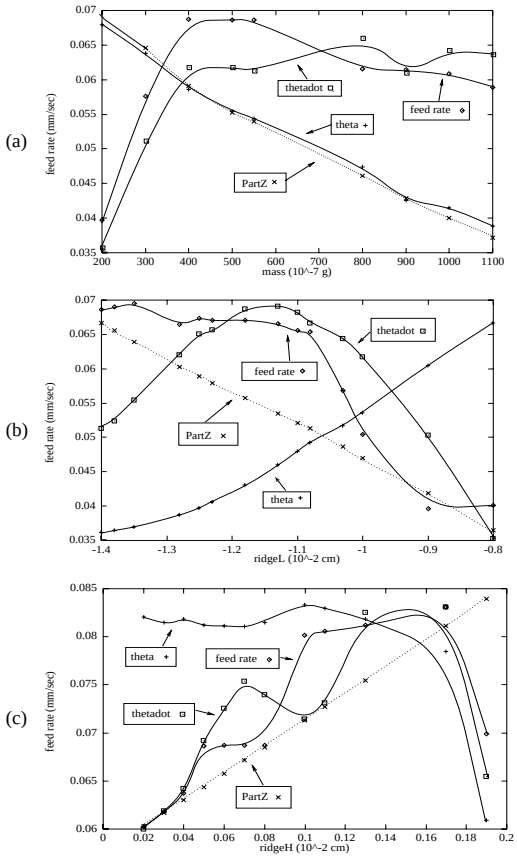

Fig. 9. Results of the performance studies. Each graph's y-axis presents simultaneously ν, θ, $\dot{\theta}$, and PartZ (averaged over a feeding experiment) plotted against the parameter being varied. For (a), (b), and (c), the parameters are, respectively, M, L, and H. The y-axis is labeled in feed rate units, numeric values for the other quantities have been omitted for simplicity.

Figure 9(a)'s optimal ν at $M = 55$ μg is reached since (i) θ is monotonically decreasing (as it gets heavier, the part "sinks" into the array, tracking PartZ) and (ii) $\dot{\theta}$ tends to level off (both these effects contribute to less momentum transferred on the $+x$ direction).

For Figure 9(b), $\dot{\theta}$ increases with L by a simple lever-like effect (ridge gets closer to revolute joint). This effect explains a monotonically decreasing PartZ, since as the ridge gets closer to the center, it transfers less momentum at every impact. These two opposing trends cause the $\dot{\theta}$ curve to go through a maximum at approx. $L = -125$ μm, which in turn causes ν to reach its highest value shortly after $\dot{\theta}$'s maximum.

The first obvious fact in Figure 9(c) is that PartZ moves linearly up with an increase in the ridge height. For the lower region of H values, θ remains constant while $\dot{\theta}$ increases monotonically, and so does ν. At $H = 16$ μm both θ and $\dot{\theta}$ decrease sharply, also decreasing ν. At present we haven't been able to justify what are the geometric/dynamic reasons for this effect.

C. Dynamic Characterization

1) Array ballistics: The plot in Figure 9(a) shows a linear relationship between M and the PartZ, i.e., the array acts as a linear spring over which the part bounces. This linear relationship yields a spring constant $k = 294$ kdyn/cm. For $M = 55$ μg, this mass-spring system resonates at 370 Hz. We computed the Discrete Fourier Transform (DFT) of PartZ for $0.05 < t < 0.2$ sec, and found two major frequency components: one at 360 and one at 5 kHz, corresponding, respectively, to the spring-like oscillations, and the normal resonator impacts. The DFT also revealed that the 360 Hz component was 5 times larger than the 5 kHz one, suggesting that too much energy is being spent in the upward vibrational mode (more on this below).

2) Friction modeling: To understand the type of friction experienced by the part as it "slides" on the array, we performed the following experiment. A 55 μg part is dropped on the array as in Figure 3. Enough time is waited so that the part achieves its equilibrium feed rate of 1.1 mm/sec. At that point ($t = 0.05$ sec), an external force F_{ext} in the $-x$ direction is applied to the part's center of mass. The experiment consists in observing the resulting feed rate ν' for different values of F_{ext}. Figure 10 shows the part's x position vs. time, before and after F_{ext} is applied – this graph reveals a linear relationship between the new feed rate and F_{ext}, i.e., the array acts as a *viscous* medium, akin to a fluid! This effect motivated us to look at this problem in depth, but for lack of space we direct the reader to another publication [3]. In a nutshell, it is caused by the velocity-independent coulomb frictional force combined with the temporally asymmetric stream of part-resonator impacts.

The viscous model prescribes a frictional force proportional to (i) the part's weight, and (ii) the differ-

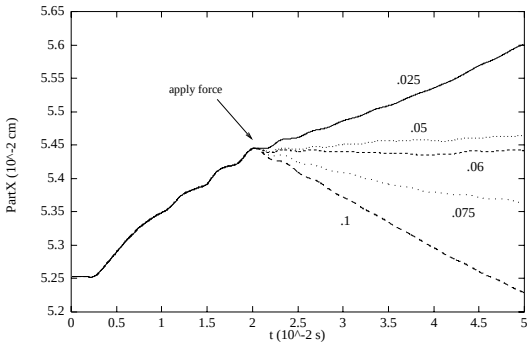

Fig. 10. The change in ν for different values of F_{ext}. After the force is applied, the part reverts to a new constant velocity ν' linearly related to F_{ext}, as in a viscous fluid.

ence between part velocity and normal feed rate, i.e., $F_{fric} = \mu M g(\nu' - \nu)$, where μ is the the coefficient of static friction. This relation was used to compute values for μ for various externally applied forces, noting that at equilibrium $F_{fric} = F_{ext}$. The results are tabulated in Table I. As shown, the model assumption explain the data quite well, with μ nearly independent of F_{ext} and approx. equal to Impulse's global coefficient of friction.

F_{ext} (dyn)	ν' (mm/sec)	μ
0	1.1	N/A
.025	.6	1.16
.05	.06	.99
.06	0	0.93
.075	-.36	1.02
.1	-.88	.99

DOF	E_{avg} (10^{-5} dyn cm)	Stdev
LinX	1.78	.71
LinY	.1	.15
LinZ	1.38	1.86
AngX	.65	.86
AngY	.63	1.28
AngZ	.0081	.045
PotZ	50.1	14.3

TABLE I

LEFT: FRICTION COEFFICIENT (μ) COMPUTED USING A VISCOUS FRICTION MODEL, FOR VARIOUS EXTERNALLY APPLIED FORCES. RIGHT: ENERGIES AND STANDARD DEVIATIONS PRESENT IN EACH OF THE PART'S INDEPENDENT DOF'S, OVER $0.05 < t < 0.2$ SEC.

3) Energetics of part motion: We measured the average energy present in each of the part's degrees of freedom: three translational kinetic energies along X, Y, and Z, denoted LinX, LinY, and LinZ; three angular kinetic energies about the X, Y, and Z axes, denoted AngX, AngY, and AngZ; the part's potential energy PotZ (with $g = 981$ cm/sec^2) with respect to its resting height. The values for these energies averaged over an $0.05 < t < 0.2$ sec (and the corresponding standard deviations) are shown in Table I.

The above data shows that this type of array consumes a disproportionate amount of energy to keep the part at an average potential energy, rather than for forward motion. This suggests that an array with longitudinal rather than vertical actuators would be more energy efficient (e.g., see [8] for a novel type of design based on thermally-actuated cilia). Notice also that a small portion of the energy is *equipartitioned* between AngX and AngY, implying that part motion along these DOF's is pretty much chaotic. The low value of AngZ shows that the array (as expected) is unable to accelerate the part about the Z axis.

4) An abstract model: From the above studies, the following simplified model of part/array dynamics can be derived: the array acts as a springy conveyor belt over which the part hops. The average part height (how much it sinks into the array), denoted d, is a function of the part's mass M and the supporting spring's stiffness k. Friction with the conveyor belt is viscous with coefficient μ. The belt feeds at a rate ν, also a function of d. This model is depicted in Figure 11.

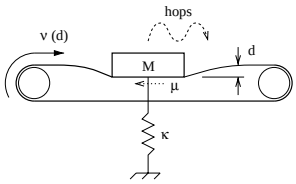

Fig. 11. The array acts as a spring-loaded conveyor belt over which the part hops. The friction is viscous with coeff. μ. The average part height d is a function of the part's mass M and the supporting spring's stiffness k. The belt's feed rate ν is a function of d.

III. A Longitudinally-Vibrating Parts Feeder

The parts feeder design originally proposed in [3] and illustrated in Figure 12 is considered next. It consists of a flat *plate* actuated by a linear *motor*. The motor is assumed to convert an input voltage into an acceleration of the plate. A small *part* is placed on the plate's surface. Coulomb friction is assumed to be the only force acting on the part. The plate's motion (see below) is assumed to be such that the part is always *sliding* over it (no sticking). The problem of contact point indeterminacy [9] between the plate and the part's footprint is avoided by treating the latter as a point mass.

The feeders' principle of operation is based on one interesting feature of the Coulomb frictional force: its magnitude is independent of the part's velocity relative to the plate. Namely, the plate's oscillatory motion is made to spends more time moving forward (positive velocity) than backward (negative velocity). Imagine a part initially at rest on the plate's surface. During one such time-asymmetric vibration, the part will experience forward force of magnitude μmg for *a longer time* than it

will experience the same force on the opposite direction, resulting in net a forward momentum added to the part after the cycle.

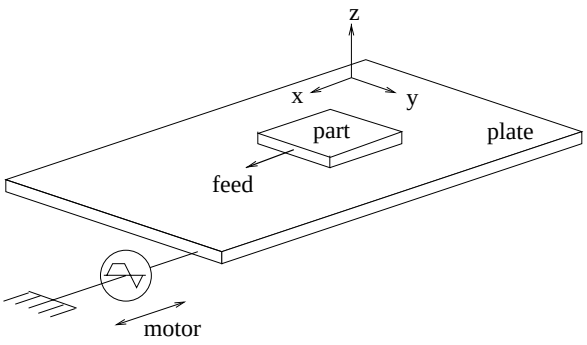

Fig. 12. Model of the longitudinally vibrating parts feeder. A linear motor commands in-plane, time-asymmetric oscillations of the plate, which cause a part resting on its surface to move forward at a constant velocity called the *feed rate*.

A. Plate Motion

One cycle of the motor's input voltage is shown in Figure 13(a). It causes the plate to accelerate in sequence at $\{a_{max}, 0, -a_{max}\}$ m/s^2. Shown in Figure 13(b) is the plate's *velocity profile*, consisting of a positive (forward displacement) semicycle of length F followed by a negative (backward displacement) semicycle of period B. Notice that $F > B$, as required for forward feeding. The integral of this curve is the plate's displacement, shown in Figure 13(c).

The positive semicycle is an isosceles trapezoid with parallel bottom and top sides of lengths F, and P, respectively. The trapezoid's height is the plate's maximum forward velocity, ν_{max}. The negative semicycle is an isosceles triangle of base B and height $|\nu_{min}|$, the plate's maximum backward velocity. The period T is equal to $F + B$. We call $d = F/T$ the *duty cycle* of the velocity profile. Constraining the acceleration waveform to have a null average value derive [3]:

$$P = T\sqrt{2d - 1} \tag{1}$$

B. Part Dynamics

The part is idealized to a point of mass m. The plate's vibratory motion gives rise to an instantaneous sliding frictional force f_{fric} acting on the part. The Coulomb friction model states that f_{fric} is (i) in the direction of the plate's motion relative to the part, and (ii) of magnitude μmg, where $g = 981\ cm/s^2$ is the acceleration of gravity and μ is the coefficient of friction of the plate/part interface.

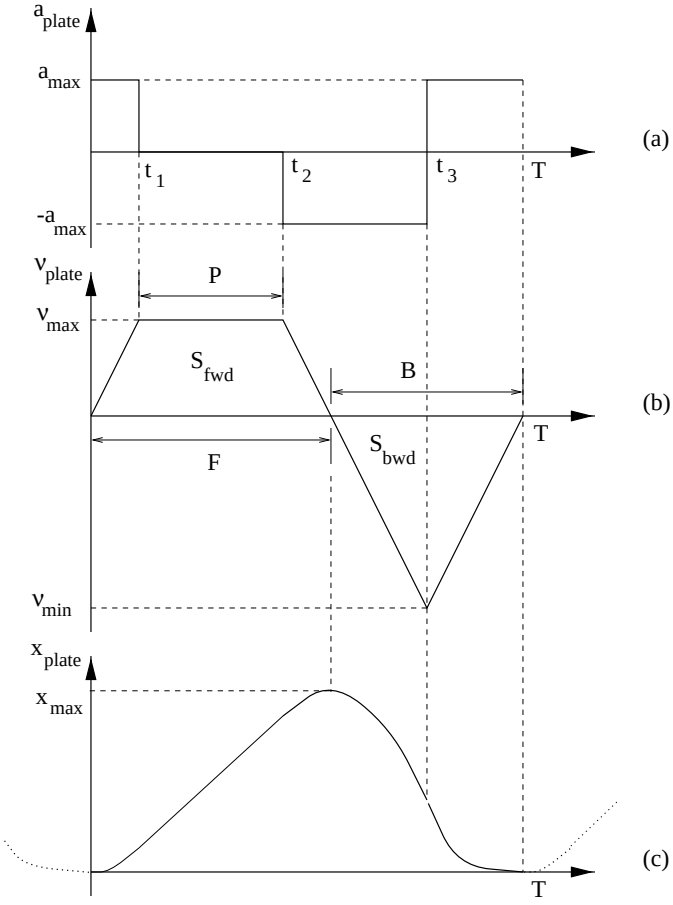

Fig. 13. Details of plate motion. (a), (b), and (c) show respectively one cycle of the plate's acceleration, velocity, and position.

1) Quasi-static analysis: The following *quasi-static* assumption is made: the part's maximum change in velocity within one vibration cycle is assumed negligible compared to the plate's maximum velocity, i.e.:

$$\mu g T << \nu_{max} \tag{2}$$

A given part velocity ν_{part} splits the velocity profile $\nu_{plate}(t)$ into two portions: (i) $\nu_{plate} > \nu_{part}$, (ii) $\nu_{plate} < \nu_{part}$. This can be visualized by superimposing ν_{part} (a horizontal line according to quasi-statics) over the ν_{plate} graph, as illustrated in Figure 14. Denote phase (i)'s duration by t_{fwd}; phase (ii)'s duration is then simply $T - t_{fwd}$. During (i), $f_{fric} = \mu mg$, and during (ii) $f_{fric} = -\mu mg$. The part will experience a net positive (resp. negative) change in momentum over the entire cycle if $t_{fwd} > T/2$ (resp. $t_{fwd} < T/2$), i.e., the ν_{part} horizontal line slides up (resp. down) over the velocity profile, Figure 14, resulting in a narrower (resp. wider) cross-section of the positive

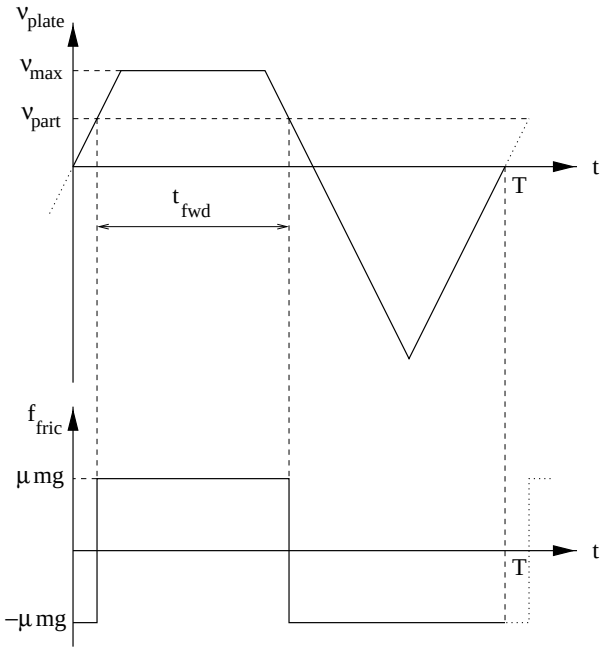

Fig. 14. Frictional forces acting on the part. The top diagram shows the plate's velocity profile superimposed on the part's velocity (assumed constant during the cycle). The bottom graph shows that while $\nu_{plate} > \nu_{part}$ (which lasts for t_{fwd}), the frictional force $f_{fric} = \mu mg$; during the rest of the cycle $\nu_{plate} < \nu_{part}$ and therefore $f_{fric} = -\mu mg$. In the cycle shown, the part will de-accelerate since $t_{fwd} < T/2$.

trapezoid. This will cause the part to stop accelerating at an equilibrium velocity ν_{eq} such that $t_{fwd} = T/2$.

If $P > T/2$, the part will accelerate to ν_{max} and reach a *saturation* point after which it can no longer accelerate, yielding an expression for the maximum duty cycle d_{sat} for which ν_{eq} is smaller than ν_{max} [3]:

$$d_{sat} = \frac{5}{8} = 0.625$$

For $d < d_{sat}$, the following expression yields a value for ν_{eq}:

$$\nu_{eq} = \frac{a_{max}T}{4}(2d - 1) \tag{3}$$

Notice that for $d = 0.5$, $\nu_{eq} = 0$, i.e., as expected the part doesn't feed since the velocity profile is perfectly symmetric, with $P = 0$, and $t_{fwd} = T/2$. An "optimal" value d_{opt} for the duty cycle is derived [3] such that it places ν_{eq} midrange between 0 (no feeding velocity) and ν_{max} (the saturation point):

$$d_{opt} = 2 - \sqrt{2} \cong 0.586$$

So the final expression for the equilibrium velocity over the entire d range (0.5 to 1.0) becomes:

$$\nu_{eq} = \begin{cases} \frac{a_{max}T}{4}(2d-1) & 0 < d < 5/8 \\ \frac{a_{max}T}{2}(d - \sqrt{2d-1}) & 5/8 < d < 1 \end{cases}$$

Notice that for $d = 5/8$, ν_{eq} is maximal and equals $a_{max}T/16$. The graph for this expression in terms of the duty cycle is shown in Figure 15(a), for $a_{max} = T = 1$.

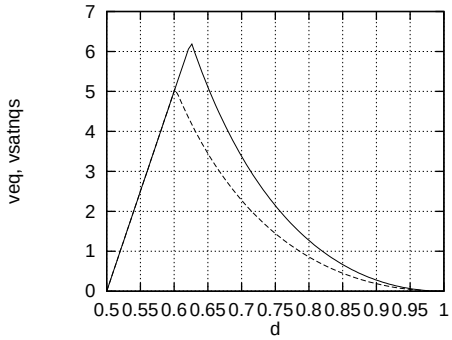

Fig. 15. The part's feed rate ν_{eq} plotted against the duty cycle d, assuming $a_{max} = T = 1$. (a) Quasi-static assumption (solid line): ν_{eq} is linear up to $d = 0.625$ after which it tracks ν_{max}. (b) Non-quasi-static feed (dotted line): Non-quasi-static saturation occurs before the quasi-static one; from that point the non-quasi-static feed rate is strictly less than the quasi-static one.

At the quasi-static equilibrium feeding velocity, the total force applied to the part per cycle is zero. If an external force f_{ext} is applied to the part in the X direction, the latter will seek a new equilibrium velocity ν' such that $f_{fric} = f_{ext}$. In particular [3]:

$$\begin{aligned} f_{ext} &= -K(\nu' - \nu_{eq}) \text{ with:} \\ K &= \frac{4\mu mg}{a_{max}T} \end{aligned} \tag{4}$$

The above equation states that around a neighborhood of ν_{eq} the plate's frictional force (averaged over sufficiently many cycles) is proportional to the part's velocity deviation from the feed rate, i.e., the part behaves as if it were immersed in a viscous fluid moving forward at the feed rate (where K is the constant of viscosity).

2) Non-quasi-static analysis: When Equation 2 does not hold, the velocity of the part within each vibration cycle can be expected to vary appreciably with respect to the plate's. As stated above, the part will accelerate (or de-accelerate) at a constant rate equal to f_{fric}/m, i.e., μg.

Surprisingly, the expression for the feed rate in the non-quasi-static case prior to saturation is identical is indentical to the result derived the quasi-static assumption (Equation 3), as shown in [3].

In the non-quasi-static case, the onset of saturation occurs when the part's acceleration within the positive semicycle takes ν_{part} to the upper right corner of that semicycle, as shown in Figure 16. The particular duty cycle $d_{sat,nqs}$ above which this occurs is given by:

$$d_{sat,nqs} = \frac{5a_{max}^2 + \mu g(\mu g - 2a_{max})}{8a_{max}^2} \tag{5}$$

Notice that if μg is small compared to a_{max} (another phrasing of the quasi-static assumption, Equation 2), then $d_{sat,nqs}$ reduces to 5/8, which matches the saturation duty cycle for the quasi-static case.

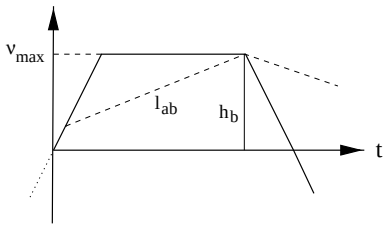

Fig. 16. The onset of saturation in the non-quasi-static case. This occurs when part acceleration in the positive semicycle takes part velocity to the upper right corner of that semicycle.

Equation 5 can be solved for the ratio $a_{max}/\mu g$ to give the following alternative characterization of the onset of saturation:

$$\frac{a_{max}}{\mu g} > \frac{1 + 2\sqrt{2d-1}}{5 - 8d} \tag{6}$$

If the inequality in Equation 6 is strictly satisfied, then the situation becomes as depicted in Figure 17, namely, part acceleration within the positive semicycle hits the ν_{max} hard limit before the end of the plateau phase, and the part follows the plate kinematically for a period of L seconds, as shown in Figure 17. An expression has been derived for the feed rate above saturation for the non-quasi-static case [3]:

$$\nu_{sat} = \frac{a_{max}T}{2}\left[\frac{a_{max} - \mu g}{a_{max} + \mu g}\right]^2 (d - \sqrt{2d-1}) \tag{7}$$
$$d > d_{sat,nqs}$$

A plot of the equilibrium velocity predicted by both models (quasi- and non-quasi-static) is presented in Figure 15(b). This Figure shows that beyond the saturation point (where the curves intersect) ν_{sat} is strictly less than ν_{eq}, meaning that non-quasi-static behavior reduces the feed rate beyond the saturation point $d_{sat,nqs}$. Furthermore, the saturation point itself is strictly less than the quasi-static one, so the peak feed rate predicted by the quasi-static model is never reached.

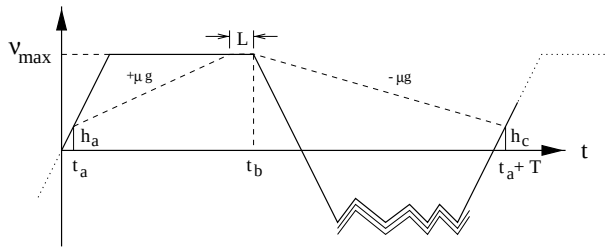

Fig. 17. The part's velocity profile beyond the non-quasi-static saturation point: the part accelerates rapidly enough to hit the ν_{max} barrier before the end of the plateau phase.

C. Dynamic Simulation

As a tool for validating the analytical results derived in the previous sections, we ran various rigid-body dynamic simulations using the Impulse tool [1]. A rectangular part 1x1x0.1 cm and mass $m = 0.1g$ was used though the analytical results were derived assuming the part was a point mass. The vibrating plate was defined as a one link robot with a single degree of freedom along the X axis. The acceleration wave described earlier was implemented by associating an Impulse *control law* to the plate. The simulation studies conducted involved measuring the equilibrium feed rates with respect to (i) the duty cycle, and (ii) externally applied forces.

1) Choosing design parameters: A frequency of 50 Hz ($T = .02$ s)was chosen for the plate, corresponding to the typical operating frequency of existing vibratory bowl feeders [7]. The plate's oscillation amplitude x_{max} was set at 1/4 cm. At $d = 0.586$, (for which $\nu_{eq} = \nu_{max}/2$), Equation ?? prescribes $a_{max} \cong 15,000 cm/s^2$, i.e., approx. $15g$. At this level, and using $\mu = 1.0$, Equation 5 yields $d_{sat,nqs} \cong 0.61$, so at $d = 0.586$ we are well in the non-saturating region.

2) ν_{eq} vs. d: Figure 18 shows a plot of the measured equilibrium velocity with the above parameters for a_{max} set to 10, 15, and 20 $10^3 cm/s^2$. The data points are superimposed with the analytical curves derived in previous sections, showing good predictive behavior.

The pre-saturation values tend to be lower than the analytical results, and this was due to an anomaly in the numerical integration of Impulse, which is still being investigated. In particular, the simulator will occasionally apply impulses to the part with a negative x component, much higher than a_{max}, showing that the package has some numerical instability problems. When these anoma-

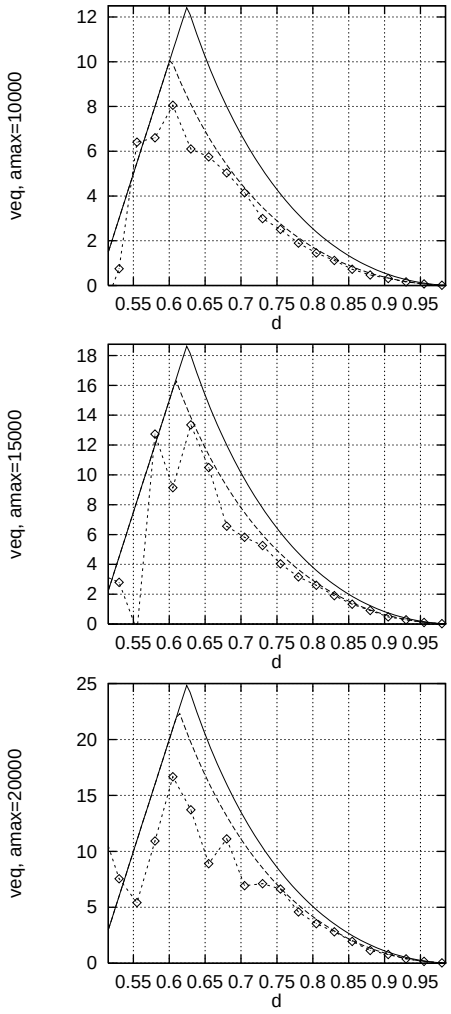

Fig. 18. Simulation results for the part's equilibrium feed rate for $a_{max} = \{10, 15, 20\}10^3 cm/s^2$ (shown from top to bottom).

lies are averaged out with the otherwise stable computations, they result in lower feed rates. The post-saturation values follow closely the non-quasi-static results, tailing off to very small feed rates as d approaches unity.

3) ν_{eq} vs. f_{ext}: The design parameters chosen above resulted in unstable computations for the experiment in which an external force is applied to the velocity-equilibrated part. Through trial and error it was determined that $T = 0.01s$, $a_{max} = 10000 cm/s^2$, and $\mu = 0.5$ produced very stable results. Plugging these parameters on Equation 4 yields a predicted value of 1.96 for K.

The simulations performed to estimate K consisted in (i) dropping the part on the feeder for 1 second and then applying a force in the -15,15 dyn range, and averaging the resulting velocity for the two following seconds. A total of 30 samples in that range were tried (spaced by 1 dyn each), and a linear relationship between f_{ext} and ν' was measured, as shown in Figure 19. A least-squares fit of the data was performed leading to the equation $y = -4.72 + 1.95x$; the measured K is the angular coefficient in this equation, 1.95 which matches within 1% the analytical value computed above.

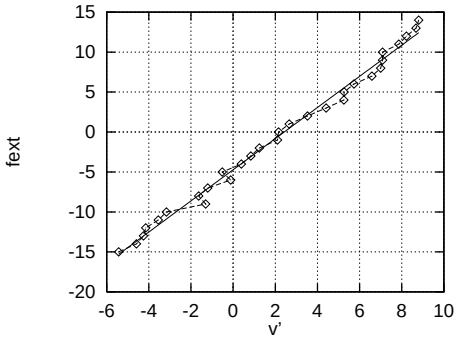

Fig. 19. Plot of f_{ext} vs. ν', for values of $fext$ ranging from -15 to 15 dyn, and spaced by 1 dyn. the graph is shown superimposed with the linear function $-4.72 + 1.95x$ obtained by a least squares fit of the data points.

IV. Conclusion

In this paper we show how dynamic simulation can be used as an effective tool in the characterization an further design optimization of two novel types of parts feeders. Design fixes and improvements have been suggested for the MEMs feeder which greatly improve its performance. Simulation results found for the vibrations-based feeder are in very good agreement with analytical results. We have found in both cases that Impulse is an invaluable tool in verification, debugging, optimization, and dynamic characterization of both feeders.

References

[1] B. Mirtich and J. Canny. Impulse-based simulation of rigid bodies. In *Symposium on Interactive 3D Graphics*, New York, NY, 1995. ACM Press.

[2] K. Böhringer, B. Donald, and N. MacDonald. Single-crystal silicon actuator arrays for micro manipulation tasks. In *IEEE Workshop on Micro Electro Mechanical Systems (MEMS)*, San Diego, CA, February 1996.

[3] D. Reznik, J. Canny, and K. Goldberg. Analysis of part motion on a longitudinally vibrating plate. In *International Workshop On Intelligent Robots and Systems*, Grenoble, France, September 1997.

[4] J. Krishnasamy, M. Jakiela, and D. Whitney. Mechanics of vibration – assisted entrapment with application to design. In *IEEE International Conference on Robotics and Automation*, Minneapolis, MN, May 1995.

[5] B. Mirtich, Y. Zhuang, K. Goldberg, J. Craig, R. Zanutta, B. Carlisle, and J. Canny. Estimating pose statistics for robotic part feeders. In *IEEE International Conference on Robotics and Automation*, Minneapolis, MN, April 1996.

[6] D. Berkowitz and J. Canny. Designing parts feeders using dynamic simulation. In *IEEE International Conference on Robotics and Automation*, Minneapolis, MN, April 1996.

[7] G. Boothroyd. *Assembly automation and product design.* Marcel Dekker, Inc., New York, NY, 1991.

[8] J. Suh, S. Glander, R. Darling, C. Storment, and G. Kovacs. Combined organic thermal and electrostatic omnidirectional ciliary microactuator array for object positioning and inspection. In *Proc. Solid State Sensor and Actuator Workshop*, Hilton Head, NC, June 1996.

[9] M. Peshkin and A. Sanderson. The motion of a pushed, sliding workpiece. *IEEE Journal of Robotics and Automation*, 4(6), December 1988.