Local Ultrastability in a Real System Based on Programmable Springs

Santosh Manicka and Ezequiel A. Di Paolo

Center for Computational Neurosciences and Robotics,
University of Sussex, United Kingdom
santosh.manicka@gmail.com, ezequiel@sussex.ac.uk

Abstract. A way to move gradually towards an objective is by making sure at every step that there is as little deviation as possible while adapting to obstacles. This has inspired us to model a local strategy to eventually attain viability (equilibrium) in a real complex dynamical system, amidst perturbations, using ultrastability to make sure that the path to viability itself is viable. We have tested this approach on a real actuator powered by a technology called "programmable springs" that allows for real-time non-linear programmable actuation. Our experiment involves a problem in adaptation similar to the pole-balancing problem. To solve it, we use ultrastability in a novel way, looking at the viability of dynamical transitions of the system in its phase space, to tweak the local properties of the actuator. Observations show that our approach is indeed effective in producing adaptive behaviour although it still requires further testing in other platforms, thus supporting the original hypothesis that ultrastability can be an effective adaptive mechanism [3] and laying a foundation for a promising new perspective in ultrastable robotics.

Keywords: Ultrastability, programmable springs, dynamical systems.

1 Introduction

The study of adaptive behaviour is central to the study and synthesis of intelligence [1]. The field has seen numerous models of adaptation, a majority of them designed from a learner-perspective. They involve a process that helps the learning system adapt to the intricacies of the task at hand based on some kind of optimisation. These models attempt to minimize the difference in the achievable performance of the system and the ideal performance for a given task [2]. In general, adaptation as a process is problem-centric.

Ashby gave a different perspective to adaptation: to learn is to act with stability in the face of environmental challenges [3]. The challenge in this case is on the viability of the system itself. In Ashby's words, there are certain 'essential variables' in a brain-like networked system that must not exceed certain 'critical values' [3]. When the environment poses a challenge to the system and some of its essential variables exceed their critical values, viability is compromised. Ashby proposed a method called 'ultrastability' to deal with such situations. As essential variables move to critical values, an adaptive mechanism kicks in and tries new connection parameters for the network that defines the system. The mechanism remains active as long as the values remain

G. Kampis, I. Karsai, and E. Szathmáry (Eds.): ECAL 2009, Part I, LNCS 5777, pp. 91–98, 2011.

critical, and stops acting once the essential variables are viable again. Ashby argued that the system in principle could find the new parameters through random search [3]. He implemented his idea in a 'Homeostat', an electromechanical device demonstrating various forms of learning [3]. It is worth noting that as a concept of adaptation, ultrastability is not performance-based, but it is based on the integrity of the system: its consequences for performance are implicit – an idea arguably closer to how adaptive mechanisms might work in natural organisms. After this breakthrough in the 1950's, only very few researchers have used ultrastability in synthetic adaptation [4]. Indeed, after the Homeostat, very few real ultrastable systems have been designed. Ultrastability can, in principle be applied to any complex dynamical system. One of our motivations is to test its generality by applying it to a different kind of electromechanical system.

A problem with practical applications of ultrastability is the efficiency of random search in large parameter spaces that characterises real complex dynamical systems, e.g., the mammalian brain. Ultrastability in such cases could take an impracticably long time to find the survival parameters. A solution to this problem could be to employ a 'controlled random' search which we describe here. The system used to test it is a real actuator powered by a novel technology called "programmable springs" developed at Sussex [5]. This technology employs a non-linear programming of the actuator's behaviour using a 'force-profile', a graph that indicates how much force should be exerted by the actuator under various conditions. In principle, they can be used to design the actuator to exhibit any complex rotary motion [5]. Programmable springs are a potentially rich technology for intelligent robot actuators as the profiles can be dynamically programmed, but one issue that needs to be addressed is that of the best method to program them adapted to specific needs. This is another motivation for our work. We describe an experiment using programmable springs for an adaptation problem where ultrastability, if used as proposed originally, might be impractical, time-wise. We implement a novel "controlled random search" flavour, wherein ultrastability is triggered *both* when the essential variables are off viability limits and their path towards viability is without a "viability structure" that we refer to as 'deviation'. When triggered, ultrastability mutates small 'pieces' of the force-profile that are deemed responsible for the deviated behaviour as observed in its real-time phase space, thus trying out new behaviour. It will also be argued that the proposed method can be generalized to other problems in adaptation as well.

2 Methods

Our experimental apparatus consists of a see-saw platform. Its rotary motion about the centre is controlled by a "programmable spring" (Fig.1). It is basically a rotary electromechanical actuator with angle-sensing capability and a control software that can exert a specific force and damping in a particular direction (clockwise or anti-clockwise) when the see-saw platform is inclined at a certain angle (negative, when the left end of the platform is lower than the right end and positive, otherwise). The net exerted force determines the angular velocity, which in combination with the angle of inclination forms the axes of the system's phase space. A 'force surface' over the actuator's phase space for the profile in Fig.3 is depicted in Fig.2. The light grey areas indicate forces in the anti-clockwise direction and the dark grey areas indicate forces

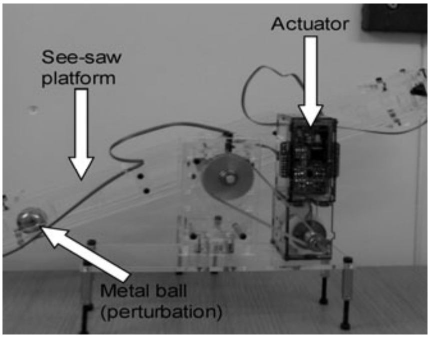

Fig. 1. The apparatus

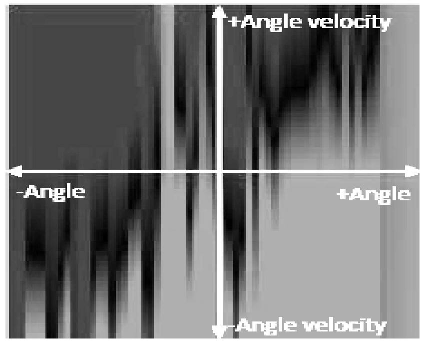

Fig. 2. Force surface for the profile in Fig.3

in the clockwise direction. The intensity of each shade indicates the intensity of the force exerted at that point. Black areas indicate zero force.

After the apparatus is switched on, the initial inclination of the platform and the shape of the force profile determine its ensuing behaviour. With a force surface symmetrical about the angle velocity axis and a uniform gradient along the angle axis, the platform reaches an 'equilibrium' point that corresponds to a horizontal inclination of the platform, regardless of where it is released from (Fig.4a). However, with the same profile, when a metal ball is introduced onto the platform, it reaches a point off the equilibrium point (Fig.4b). It can be seen that the perturbed behaviour finishes on the same side as the starting point simply because the underlying force profile is not robust enough to not give in to the weight of the ball (the light grey arrows in fig.4b show the same). In order to reach the equilibrium point in the latter case, one solution could be to restructure the force profile in such a way that the resulting behaviour looks similar to a symmetrical trail that naturally leads to stable equilibrium in the absence of perturbation (dark grey arrows in fig.4b). Thus by adopting a “natural” trail structure as a behavioural guide, a force profile that can adapt to perturbations could be found. Devising a method to accomplish this is the challenge addressed in this work.

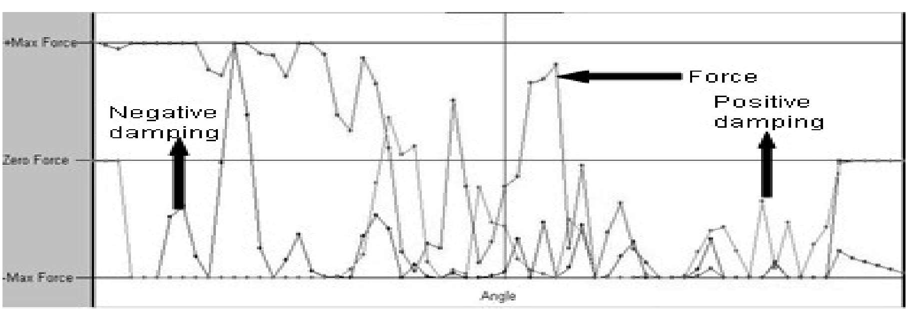

Fig. 3. A sample force profile. 'Positive damping' damps clock-wise rotation and 'negative damping' damps anti-clockwise rotation

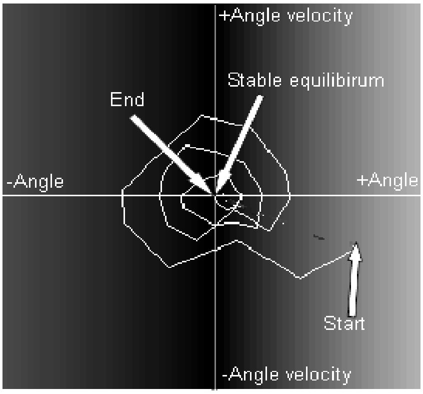

Fig. 4a. Actuator's behaviour without perturbation

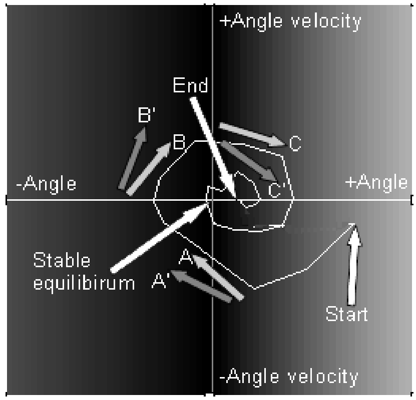

Fig. 4b. Actuator's behaviour with a permanent perturbation

In all the explanations that follow, the metal ball shall be considered a part of the apparatus. The proposed adaptation method consists of 3 steps: (1) actuation, (2) behaviour analysis and (3) profile mutation. Starting with a random profile, the actuator is activated and its trails in the phase space are captured for a predetermined length of time. The actuator is then paused and its behaviour during that period is analyzed by breaking the trail into individual 'transitions' like the light and dark grey arrows in Fig.5 below. In the analysis, we consider a random subset of these transitions with a predetermined number of samples (Parameter *V*). Any transition similar to the dark grey arrows (henceforth referred to as 'ideal transitions') are considered deviating with respect to a 'guide trail' as described in the previous paragraph, following which, the responsible tiny part of the profile is slightly mutated, that is randomly restructured in the spirit of ultrastability. The whole process is then repeated until a profile is found such that its trail spends a considerable length of time near the stable equilibrium point called as the 'Global Viability Zone' or 'GVZ' (the central square in Fig.5).

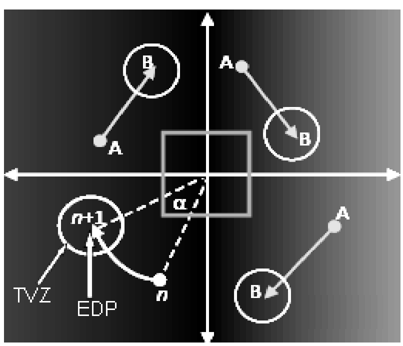

Fig. 5. A-B arrows are ideal transitions. The square is the GVZ and the circles are the TVZ's.

Each actual transition considered for analysis in a trail is compared with a corresponding ideal transition and if they are considerably different then the actual transition is considered deviating. An ideal transition at any point in the phase space is constructed as follows: if *n* is the starting point of a transition, draw an arc clock-wise from *n* with the stable equilibrium point (SEP) as the centre through a fixed angle α, the parameter *W* (Fig.5). The point *n+1* where the arc ends is referred to as the 'Expected Destination Point' (EDP). Then, a 'Transition Viability Zone' (TVZ) is constructed around the EDP whose radius is proportional (by a fixed 'TVZ radius factor', the parameter *X*) to the distance between *n* and EDP. If the actual destination point (ADP) in the actual transition is not within the TVZ, then the transition is deemed deviating. The EDP computation method captures the essence of approaching of the trail towards the SEP. The TVZ then acts as a "cushion" for the transitions so they don't have to be perfectly circular (in fact they can't be as they have to head in towards the SEP) and gives way to the formation of spiral trails. Again, the cushion lets spiral-like trails to emerge, thus giving room for complex behaviour. When a profile mutation is required because of a deviating transition, the responsible point on the force profile is looked-up and mutated as follows: if the absolute value of the angle velocity of the ADP is greater than that of EDP, then the absolute value of the force and damping values at that point are decreased and increased otherwise. The magnitude of mutation to be performed is computed as follows:

$$Magnitude = A \times B / C,$$

where, A = Random number in (0, Initial mutation range (Parameter *Z*)),

B = ADP-to-EDP distance,
C = Mutation range factor (Parameter *Y*).

The mutation range factor is manually set (an appropriate value was estimated by trial and error in our experiments). Mutations of proportionally lesser magnitude are also performed in the same direction on a few points near the chosen point in the profile in order to get a 'smooth effect', for the sake of gracefulness. After a predetermined initial mutation range is set, it is periodically tuned according to how well the adaptation is happening. After every few iterations of the actuation-analysis-mutation cycle, the current profile is scored and if it is better than the last evaluated profile, then the mutation range is narrowed a bit so the adaptations accumulated until then are more likely to be conserved, otherwise it is widened a bit so new adaptations could be explored. The method for scoring a profile is as follows:

$$Profile\ score = \frac{D - (2 \times E)}{F}$$

where, D = time spent by the trail in the GVZ,

E = standard deviation of the time spent in each quadrant of the phase space,
F = total time spent in the phase space.

As it can be seen, a number of control parameters are manually set before starting each experiment. A number of experiments were conducted by varying them. The following section will present the results from the most successful experiment.

3 Experiments

This experiment was run for about 10 hours at the end of which 842 profiles were generated. The following values were chosen for the parameters: V=250, W=30°, X=2, Y=80 and Z=100. Fig.6 below shows how the factors D, E (described above) and the profile score change with time. As it can be seen, there is a gradual improvement in factor A and the profile score over time. Since B seems to be almost constant, it can be assumed that A has more or less solely contributed to the improvement in profile score. The improvement in A can be attributed to the gradual approach of the average point in each quadrant visited by the evolving trails, towards the SEP (note the circles marked in Fig.7 where this approach has been captured). Fig.8 below shows the behaviour of the 600^{th} profile that had the highest profile score of all. The trail indicates that it has developed some kind of 'one-arm strategy', with one active arm at a time (the two ellipses). Though it could give the indication that during these one-sided times, the ball is stuck on that side whereas in fact, it was not. This is supported by direct observation. Also, though not quite distinctly highlighted in fig.8, sudden drops in the angle velocity are also present. They characterise something we would like to call the 'braking behaviour' which actually let the ball roll smoothly on the platform thus performing a pause-and-pass act. Such an act in principle can enable the ball to gradually settle at the centre of the platform. Though it never fully successfully happened, the repeated braking behaviour indicates a trend to bring the ball to the centre.

Overall, the results show that the system develops a tendency towards reaching a dynamic equilibrium where the ball actuator behaves in such a way that the ball steadily sways about the centre of the platform and thus not simply succumbing to its weight. Experiments with other control parameter values resulted in different kinds of behaviours that could be in part explained by the choice of those values. Also, a slightly different approach to the mutation method was tried in it the direction of the mutation was also randomly chosen in a whole-hearted spirit of ultrastability, thus encouraging rapid exploration. The results (not presented here) showed emergence of interesting behaviours in a very short period of time. However, they quickly disappeared as the learned adaptations got easily eroded by the ensuing explorations.

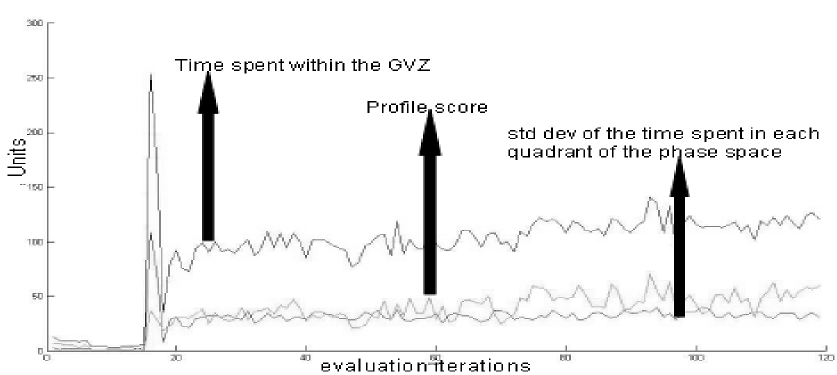

Fig. 6. Trends in profile score, time spent in the GVZ (A) and standard deviation of the time spent in each quadrant of the phase space (B)

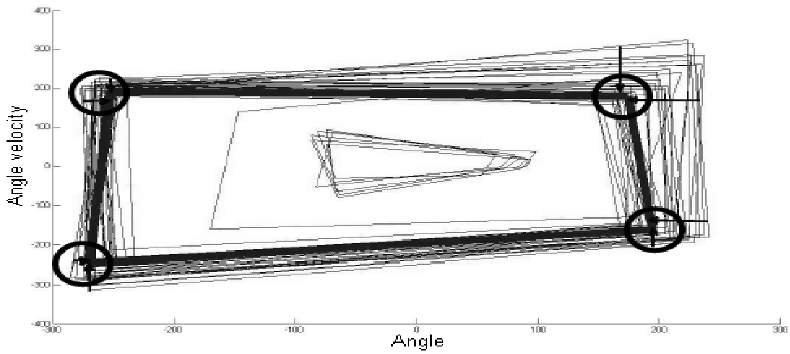

Fig. 7. A frontal view of the time evolution of the quadrant-wise average of the trails. The bold box is the latest average.

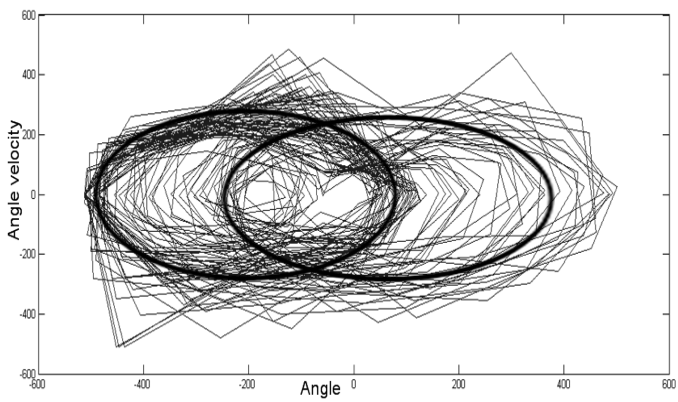

Fig. 8. Behaviour of a top scorer profile in the phase space. The bold ellipses mark approximately the average trajectories, each one slightly biased to one of the two sides of the platform.

4 Discussion

The objective of this work was to test ultrastability as an adaptive mechanism in a real system and to tailor it to suit the structure of the powering mechanism based on "programmable spring". Due to the continuous nature of the force profile and the need to adaptively shape the profile in a controlled fashion, the notion of viability zone in ultrastability was extended to the properties of the *dynamical transitions* in the phase space rather than defining it *only* for the state of the whole system. This is a novel conception of how ultrastability may be used. Our method was tested with a balancing problem. The ideal outcome of the experiments is a force profile adapted to the 'structure' of the perturbations (the dynamics of the moving ball) wherein the ball would be brought to a stand still in the middle of the platform, regardless of its initial position. With respect to this, the actual outcome could be treated as only partially successful.

One reason for this is the highly sensitive physical aspects of the system – even a slight movement of the ball off the centre can throw the system off balance and trigger behavior that keeps the ball in continuous motion. Repeated use also creates mechanical noise that cannot be easily removed and that the system has to adapt to further. Moreover, interesting behaviours might emerge if the EDP (Fig.5) were to be computed from a 'spiral-in' towards the SEP rather than from a circle as proposed in this work as the inward bias might encourage a stronger control around the SEP. Future research in this direction could take into consideration these design factors in order to achieve improved performance. With regard to the time to adaptation, our current approach takes quite a long time to find a suitable profile because the actuation, analysis and mutation steps happen one after the other. If they could be performed in parallel then the performance should in principle improve.

To summarize, it was shown that ultrastability can be used as an effective adaptive mechanism in a programmable spring based real electromechanical system. Our extended version of ultrastability can shape the dynamical phase space transitions of the system, moving beyond the conventional approach of monitoring only the state of the system. We believe that this approach can be used as an adaptive mechanism for other problems too with a proper, choice of the various parameters used in the method. For instance, if the actuator were to be used in a wheel that has to rotate continuously, then the SEP would be a line horizontal to the angle-axis depending on the desired velocity and the direction of rotation and the shape of a TVZ may be a rectangle inclined at various angles to the SEP line. We thus believe that the combination of a generic programmable spring technology and a generic adaptation mechanism based on ultrastability is a promising and workable tool for designing adaptable robots.

References

1. Beer, R.D.: Intelligence As Adaptive Behavior: An experiment in Computational Neuroethology. Academic Press Professional, Inc., San Diego (1990)
2. Landau, D.Y.: Adaptive Control: The Model Reference Approach. The "Control and System Theory" Series, vol. 8. Marcel Dekker, New York (1979)
3. Ashby, R.W.: Design for a brain: The origin of Adaptive Behaviour (2e). Chapman and Hall, London (1960)
4. Di Paolo, E.A.: Homeostatic adaptation to the inversion of the visual field and other sensorimotor disruptions. In: Meyer, J.A., Berthoz, A., Floreano, D., Roitblat, H., Wilson, S.W. (eds.) SAB 2000. From Animals to Animals, Proc. of the Sixth International Conference on the Simulation of Adaptive Behavior, pp. 440–449. MIT Press, Cambridge (2000)
5. Bigge, B., Harvey, I.: Programmable Springs: Developing actuators with programmable compliance for autonomous robots. Robotics and Autonomous Systems 55, 728–734 (2007)