Change Is the Ultimate Self-Adaptive Challenge

(Owen) Shang-Wen Cheng
owen.cheng@gmail.com

***Abstract*—This paper argues that change alone is sufficiently difficult; simply, handling change, i.e., anticipated changes, remains the ultimate challenge of self-adaptation.**

I. Introduction and Position

A system able to handle unanticipated changes, whether from itself, its dependent systems, its users, or its environment, is arguably a holy grail of self-adaptation. But, are unanticipated changes the ultimate challenge of self-adaptation? This paper argues that change alone is sufficiently difficult; simply, handling change—more to the point, anticipated changes—remains the ultimate challenge of self-adaptation.

II. Journey of Systems

Let us embark on a short, imaginative journey and consider a few self-adaptive systems across different domains, donning self-adaptation lenses. Imagine that self-adaptive systems have become prevalent and power the civilized world. As a quick reality check, this is not too far-fetched, since any sufficiently smart software-intensive system is conceptually nearly indistinguishable from a self-adaptive system: one composed of some *base functionality* with *self-adaptive capabilities*. Where relevant with each system, we briefly consider its self-adaptive features, then home in on how each might handle anticipated and unanticipated changes, to assess what makes these hard.

A. Smart Home: Eco-Intelligent Thermostat

Installed as a component of a smart home, it monitors the temperature of all major spaces in the home, remembers the occupant-desired comfort settings, checks local weather for current conditions and temperature trends, and queries the power grid for cost trends. It interacts with the occupant via a phone app to obtain comfort settings and just-in-time feedback. It initially assesses home efficiency data to determine how quickly heat dissipates, or enters, the house. It controls both the central heating and air conditioning units. Using these data and control, it aims to maximize comfort while minimizing cost and ecosystem impact, by analyzing seasonal temperature and utility cost trends, and planning optimal heating or cooling commands to achieve user-desired comfort levels.

In designing its self-adaptive capabilities, the primary challenges are the sophistication of model in, and data available to, the well-interconnected thermostat. Sudden fluctuations in temperature or cost trend should be an *anticipated* change, but would it be adequately handled by the thermostat's existing models? Can the thermostat's seasonal models work in all climate zones and all countries, which are arguably anticipated changes? If the owner installs higher-efficiency windows, the designer might have partially anticipated this by allowing the thermostat to be user-triggered to re-run a home efficiency analysis. However, incorporating new, fusion energy cost analysis would be *completely unanticipated*, possibly requiring a redesign. While the last example seems very challenging, proactively designing for it hardly seems worthwhile, whereas coming up with a thermostat design resilient to handling all the anticipated changes already seems difficult enough.

B. Smart Appliance: Self-Directed Vacuum Cleaner

Installed in one's home with a power home base and multiple fast-charge way stations, once initiated, the vacuum cleaner scans each reachable room, maps out the target area, plans a coverage path, vacuums the floor, and empties its bin or recharges at a way station as needed. Cleaning can be initiated by the occupant on demand or per schedule, or it can be triggered by the indoor dust level exceeding a configurable threshold. Limited by its small wheels, it avoids any obstacles or precipitous drops. It switches its cleaning brush and varies vacuum strength to adapt to the types of surface it is vacuuming. It monitors its own battery level to ensure sufficient juice to return to base. Using its knowledge of room sizes, optimal path, floor types, and thus the cumulative energy needed, it ensures floor cleanliness while minimizing vacuum-run duration to maximize battery use and optimize overall energy utilization.

At first glance, this automated vacuum cleaner seems like a typical path-planning problem, with goals and obstacles as typical considerations. Factors complicating the path plan include battery use, schedule constraint, handling variation in floor types, and planning over more than one run to tackle multiple spaces of the house. Variations in dust and debris ought to be *anticipated* change, to a degree. If a large vase suddenly fell and shattered, the cleaner would have to make multiple runs just to clean the affected area, resulting in suboptimal energy use.

While floor types do not change often, adding or relocating a single rug should be *anticipated*, and would require the knowledge state of the cleaner to update on the next vacuum run. Encountering a floor type change can severely affect battery use and invalidate the optimal plan created at the start of a run, making achieving global optimum infeasible. If the cleaner encounters an *unanticipated* floor surface, what default brush and strength setting makes sense? As example considerations, choosing the maximum setting may destroy a delicate wool rug, but treating it as an obstacle and avoiding it risks missing the mark on cleanliness. A house with reconfigurable spaces arguably imposes *unanticipated* changes on the cleaner, but if it were designed to handle arbitrary room shapes and sizes, and not rely heavily on past path plans, then this becomes a non-issue. Whereas, designing the cleaner to vacuum arbitrary room shapes and sizes, which belongs in the *anticipated* category, poses a significant challenge to get right: How many configurations should be tested? Are more than one way station needed? How many runs are needed? Can it be guaranteed to make all scheduling constraints?

C. Deep Space Asteroid Miner

Launched toward a lithium-rich asteroid in the Asteroid Belt for autonomous operation, the miner captures the asteroid and brings it into earth-proximal solar orbit while locating one or more optimal mining zones on the asteroid to initiate the mining machinery. It then monitors mining progress and drill wear to coordinate the optimal launch of ore offload and mining supply missions from Earth.

In space, one can anticipate and avoid most obstacles known to astronomers ahead of time, but also monitor and plan trajectories around unanticipated objects. Cosmic radiation of high energy particles is another major threat, whether from solar flares, other cosmic objects, or unanticipated cosmic events, but one can cope using known error-correction techniques. What if the asteroid does not have a minable zone because its hardness exceeded that of the drills? Or, if the density of surrounding asteroids suddenly grew and blocked out solar energy for the spacecraft? Having considered these possibilities, one can design the miner with a number of backup plans, including picking a nearby asteroid that has lithium and is minable, or carrying a mini radioisotope thermoelectric generator to carry on for short durations without solar energy.

D. Self-Driving City Bus

Operating from city bus depots, each bus follows a designated route of the day, safely avoids other road users, watches for road detours, monitors its own fuel efficiency, and tracks ridership trends, incorporating daily, weekly, and seasonal cycles to tune trip frequency and optimize revenue, ridership, and energy use. Each day, a transport authority decides the initial route and frequency of each bus, but the city bus has autonomy in increasing or reducing frequency depending on demand. Each bus possesses historical data of daily, weekly, and seasonal peaks and troughs in ridership, and it knows the minimum riders required per trip to minimally afford its fuel use. *Anticipated* changes might include:

- A child emerging suddenly from between parked cars
- Sinkhole in the road, or a bag hovering above the road
- Fluctuation in ridership, from empty to full bus
- Route disruption and detour, perhaps due to transient event closure, or longer term construction closure
- Mechanical breakdown
- Fluctuations in fuel cost, whether gradual or sudden

Although anticipated, these changes are nonetheless challenging to design for. Detecting and avoiding collisions with small humans or balls or unidentified floating objects are advanced capabilities of self-driving technology and difficult to achieve with zero miss. Handling sudden fluctuations in any of the design parameters requires careful modeling, and intentional trade-off between objectives. Or, no detour may be available or known, forcing the bus to abandon route and riders and return to base.

Some *unanticipated* changes might include:

- Ridership surge due to a football game or convention
- A collapsed bridge, or an emergency aircraft landing
- An occupant seeking emergency medical attention, forcing the bus to either (a) stop and wait for an ambulance, or (b) route immediately to a nearby hospital
- Vandalism leaving the bus unable to continue operation
- An electromagnetic pulse shutting down the bus...and city
- A meteor impact eliminating all routes to the end point

While some of these examples seem admittedly extreme, they are arguably foreseeable outcomes. Three questions are worth pondering, although the author does not attempt to address them in this paper: First, when does it become unreasonable to consider and design for the more far-fetched possibilities? Second, once a possibility has been considered, is it still unanticipated? Third, might reasonable default strategies exist to address the unanticipated changes, thereby rendering them just part of the pursuit to tackling all *anticipated changes*? On the third point, except for the more destructive cases, many of the examples above could be addressed by a reroute, or simply an emergency stop. In other words, without knowing the exact unanticipated states, experienced engineers could conceivably design the bus with default responses to a range of reasonably unforeseeable outcomes.

III. Discussion and Concluding Remarks

This paper surveyed four self-adaptive systems—a simple thermostat, a robotic vacuum cleaner, an asteroid mining spacecraft, and an autonomous city bus—and considered a number of both anticipated and unanticipated changes, as well as the challenge of handling some of them. The key point conveyed through these examples is that engineers already have their work cut out for them simply handling the *anticipated changes*, whereas unanticipated changes could well be handled with well-designed default strategies. Deeply engaging conversations ought to occur around whether all unknown-unknowns could be so handled and, crucially, *should* be handled! To illustrate the point, if we engineered the self-driving city bus to handle unanticipated bombardment of meteors, but it cannot avoid collision with all road users, then we have failed at our primary objective of designing a fit-for-purpose city bus. On the other hand, solving for the anticipated changes would likely simultaneously get us much closer to tackling the handling of unanticipated changes.

Acknowledgement

The author would like to acknowledge and thank Bradley Schmerl for providing his guidance, insight, and feedback to help improve this paper.