



This is a repository copy of *Multi-objective optimisation of robotic active particle swarms for continuous repair of large scale high value structures*.

White Rose Research Online URL for this paper:
<https://eprints.whiterose.ac.uk/176757/>

Version: Accepted Version

Proceedings Paper:

Oyekan, J. (2021) Multi-objective optimisation of robotic active particle swarms for continuous repair of large scale high value structures. In: Proceedings of 2021 IEEE Congress on Evolutionary Computation (CEC). 2021 IEEE Congress on Evolutionary Computation (CEC), 28 Jun - 01 Jul 2021, Virtual Conference (Krakow, Poland). Institute of Electrical and Electronics Engineers , pp. 1312-1318. ISBN 9781728183947

<https://doi.org/10.1109/cec45853.2021.9504749>

© 2021 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other users, including reprinting/ republishing this material for advertising or promotional purposes, creating new collective works for resale or redistribution to servers or lists, or reuse of any copyrighted components of this work in other works. Reproduced in accordance with the publisher's self-archiving policy.

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>

Multi-Objective Optimisation of Robotic Active Particle Swarms for Continuous Repair of Large Scale High Value Structures

John Oyekan

Department of Automatic Control and Systems Engineering
The University of Sheffield
Sheffield, United Kingdom
j.oyekan@sheffield.ac.uk

Abstract—The manufacture and creation of large scale high value structures has been done by humans for centuries. Examples include the Egyptian pyramids, Bridges, Modern Skyscrapers to mention a few. These structures are large but also provide a high value in terms of economy, culture, display of prestige to mention a few. With advances in space technology, we are bound to see these large scale high value structures constructed in space. The vacuum of space present us with the challenges of repairing these structures. This is due to the inhospitable and dangerous environment of space. With increasing number of structures in space, there is bound to be more debris created resulting in high impact damages to these high value structures. Inspired by the biological blood clotting process and biological active particles, in this work, we propose the use of a swarm of live on artificial active particles for the purposes of continuous and timely repair of these structures. We tackle one of the challenges of artificial active particles research; the ability to navigate in crowded and obstacle filled environments. This challenge can be viewed from the perspective of a constrained multi-objective optimisation problem in which a balance between exploration of an environment and its exploitation needs to be achieved while taking into consideration the various other constraints that apply to an active particle. In this work, we show how artificial active particles could avoid obstacles in their environment through the use of an exploration mechanism and find damaged sites. Our results show that as the ability to explore increases, the active particles are able to navigate around obstacles and find a damaged site. However, there is a limit to this.

Index Terms—Active Particles, Swarm Optimisation, Autonomous Repair, Multi-Objective, Brownian motion

I. INTRODUCTION

Satellites and space assets, such as the international space station, are large scale high value assets that often suffer damages due to hyper-velocity impact collisions with micrometeoroids (natural objects such as dust and rocks) and orbital debris (human-made objects such as metallic fragments, paint chips and components from old spacecraft) at orbital velocities often exceeding 9 km/s [1] [2] (Fig. 1(c)). Such strikes lead to vital oxygen venting into space as well as render sections of the spacecraft unusable until fixed. Furthermore, damages are

difficult to repair due to the dangerous and humanly hazardous area of space vacuum. Such dangerous, repetitive, and dull nature of the repair work classifies as a job for robots.

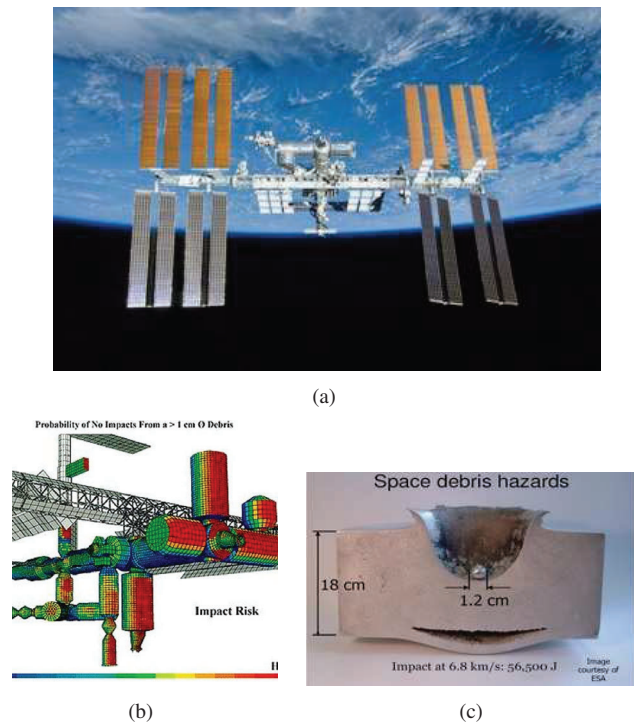


Fig. 1. Showing the international space station Fig. 1(a) with the probability of micrometeorite impact risk [2] Fig. 1(b) and a crater caused by a micrometeorite impact Fig. 1(c).

On earth, large scale high value structures/assets such as tunnels and sea-based bridges suffer continuous damage due to heavy usage, water ingress, extreme weather conditions and sea salt based corrosion. The timely and rapid repair of such structures is very important in order to prevent catastrophic loss of life. However, bridges are currently repaired by timely maintenance schedules [3] which involve humans suspended dangerously at heights while space structures often require

uncomfortable and dangerous human external vehicle space walks. In this paper, we propose that an autonomous robotic repair system comprised of a swarm of millirobots (robots measured at the millimetre range) that live on structure could repair damages as they occur. We take inspiration from the biological blood clotting system that starts healing a damage to a tissue as soon as it happens. The advantage of such approach is that it ensures that repairs are done faster thereby preventing even more damages to the structures or loss of astronaut lives in the case of long voyage space ships. With private space companies (e.g spaceX) planning to venture towards the moon, mars and deep space in a few years, being able to repair a space ship continuously, autonomously and on the spot will ensure safety to humans and longevity of space missions [4].

The Blood clotting process is a process by which blood undergoes a transformation from a liquid state to a gel based state (coagulation). The stages towards coagulation involves the following stages: Activation (chemotactic migration blood platelets to an injury site as soon as it happens); Adhesion (attachment of platelets to site); Aggregation (recruiting other platelets to site resulting in non-linear feedback of agents) deposition and maturation of fibrin by platelets [5]. Platelets, the key ingredient of blood clotting, are biological active particles. Active particles, also known as self-propelled Brownian particles are capable of taking up energy from their environment and converting it into directed motion [6]. Artificial active particles are an active research area and there have been various attempts to synthesize artificial article particles because of the promise they hold for various purposes including environmental cleaning and targeted drug delivery to mention a few.

However, [6] discusses that the current challenges of applying artificial active particles in various applications includes the precise and intelligent control of such agents, dealing with environmental constraints such as obstacles as well as navigating in the environment. These are challenges limiting their practical use. In this paper, the novelty of our work lies in deriving and applying a milli-robot controller based on the biological model of a natural active particle namely platelets for the purposes of navigating an environment to provide self-healing to structures. Our proposed milli-robot would have the capability to deal with noise in the sensor readings as well as address the challenges posed by the current state of the art in literature. Such a developed robot would ensure continuous monitoring and autonomous repair of high value structures. In this work, we focus on solving the issue of providing active particles with the ability to explore and navigate an environment filled obstacles towards finding the site of damage on a simulated structure environment.

II. METHODOLOGY

As discussed in the introductory section, the blood clotting process is the process by which blood undergoes a transformation from a liquid state to a gel based state. This involves activation, adhesion, aggregation and deposition and maturation of fibrin (See Fig. 2).

Our approach covers the first three items mostly but in future work, we plan to go onto artificial fibrin creation. In the natural blood clotting process, the first three stages are part of the primary hemostasis in which platelets immediately form a plug at the site of the injury. Secondary hemostasis occurs simultaneously where coagulation factors in the biological pathway lead to formation of fibrin strands that strengthen the platelet plug.

We focus on primary hemostasis in this work. Our proposed approach results in a rapid recruitment of other active particles to the site of damage for the next stage in self-healing. The area of the injury results in an adaptive recruitment of active particles leading to a nonlinear feedback cycle. This results in active particles going from an exploration state to a liquid state then gel state towards deposition for repair. We describe the stages in the diagram shown in Figure 3. We now explain each of the stages below:

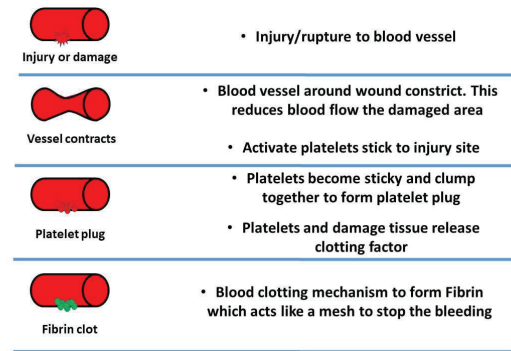


Fig. 2. The blood clot process.

- **Release state:** This is the stage in which the smart active particles are released when a damage is detected or to replenish particles that have been used in the repair process. This leads to exploration of the large structure towards finding damaged sites. The larger the distance between the release site and the damaged site, the longer it will take for the active particles to get to the damaged site.
- **Exploration state:** In this stage, the active particles explore the structure. The parameters of each individual particle can be assigned and tuned before release towards controlling exploration behaviour. As will be seen in the experimental section, the higher the rate of exploration the increased likelihood that particles will find the damaged site. Taking an analogy of fluids in gas states, the higher mobility of particles enable them to achieve a gas like behaviour that enables them to "diffuse" and explore the environment.
- **Stuck state:** During exploration, it is highly likely that particles would encounter obstacles on structures. This could cause the agents to be stuck. The more obstacles in the environment, the more difficult it will be for the particles to get to the damage site. As a result, a strategy

must be discovered to ensure that they escape from stuck scenarios.

- **Liquid state:** The liquid state happens when the particles get attached to a damaged site. In this state, the particles start a count down time that transitions them into a gel state after a threshold value. In this liquid state, the particles also move more slowly in order to navigate to the areas of intense damage. This ensures that in the case of limited number of particles, priority based repair take place with the most damaged site repaired more intensively.
- **Recruit state:** In this state, particles recruit other nearby active particles in order to create a non-linear feedback that ensures repairs start taking place faster. The rate of recruitment is controlled by the field of influence generated by individual active particles.
- **Gel state:** The gel state is not implemented in this work. However, the aim for this state is that particles either start secreting repair agents such as polymers to start the repair process [7].




Fig. 3. The blood clot process.

A. Assumptions

We consider the following assumptions in this current work. In the future, these assumptions will be relaxed as progress is made.

- Since active particles take energy from the environment to do some work, we assume that the robotic active particles can take induction energy from the structure surface (e.g space craft hull) or solar energy from the sun to power their work.
- In the case of the international space station, since the spacecraft goes round the earth every 90 minutes, in the worse case scenario, this is how long it should take for repairs to take place when using direct solar energy.
- We assume in this work that the active particles have a way of attaching themselves easily to the hull. This could be via suction, magnetic, electrostatic and other forms of attachment.
- At this stage of the research, we assume the robots have a differential or omnidirectional drive.
- Robots have minimal sensors and in this case to measure the condition of the structure as well as to detect their

neighbours. They find their way around by stochastic motion similar to brownian motion or diffusion behaviour.

- It is possible to use externally generated electromagnetic forces to inform agents of a damage as used in micro-robotics research [8] [9]. This can generate a potential field for the agents to flow and navigate towards the problem. However, this could present challenges when obstacles are encountered. In this work, we assume that there is no global field that pulls agents towards the site of damage. Instead, there are only localised damages that generate local fields. As a result, agents have to explore an environment while avoiding obstacles to discover damaged sites.

B. Individual Active Particle Controller Definition

In this section, we describe a foraging chemotactic controller capable of using stochastic behaviour to explore the environment as well as use chemotactic behaviour to navigate up a noise gradient once a damage site is found. This is inspired by the biological active particle, bacterium, whose motion can be described according to Equations 1 to 3.

$$\tau = \begin{cases} \tau_o \exp(\alpha \frac{dP_b}{dt}) & \text{if } C(x, t) > 0 \\ \tau_o & \text{else } C(x, t) = 0 \end{cases} \quad (1)$$

$$\frac{dP_b}{dt} = \int_{-\infty}^t \frac{dP_b}{dt'} \exp(\frac{t-t'}{\tau_m}) dt', \quad (2)$$

$$\frac{dP_b}{dt} = \frac{k_d}{(k_d + C(x, t))^2} \frac{dC}{dt} \quad (3)$$

Assuming that we represent the damaged sites by a spatial function $C(X)$ in the environment $S(X)$, $C(x, t)$ is a spatial measurement obtained by the biological active particle at x and time t with $x \in X$. τ is the adaptive mean run length value, τ_o is the mean run length in the absence of concentration gradients and α is a amplification constant of the bacterium chemotactic pathway. P_b is the fraction of the receptor bound of the agent when measuring concentration $C(x, t)$, k_d is the dissociation constant of the bacterium chemoreceptor and controls the chemical sensitivity of the bacterium. $\frac{dP_b}{dt}$ is the rate of change of P_b and $\frac{dP_b}{dt}$ is the weighted rate of change of P_b .

Equation 2 was implemented in the discrete form using a memory length of 4 according to a bacterium's chemosensory pathway [10]. The above Equations determine the time between tumbles and hence the length of runs between tumbles. In this work, during the tumble phase, the agent can randomly choose an angle in the uniform distribution set $\sigma \in \{0, \dots, 360\}$.

We define a control law given by Equation 4 below.

$$motion = \begin{cases} tumble() & \text{if } counter > \tau \\ run() & \text{else } counter < \tau \end{cases} \quad (4)$$

where *counter* is a variable that is incremented every time step and gets reset when it is greater than τ .

According to Equation 1, in the absence of spatial readings, that is $C(x, t) = 0$, $\tau = \tau_o$. As a result, the agents will still move about in the environment in search of a spatial function.

C. Relationship Between Individual Motion and a Spatiotemporal Distribution

In this section, we shall attempt to show that the proposed method is capable of providing a priority based coverage to both spatial and spatiotemporal damages in the environment. In both [11] and [12], the authors discussed that the motion of active cells can be described as a Langevin Equation. Active particles are representations of living organisms that are capable of moving about in the environment for various purposes including foraging but are also affected by noise from the environment. We write the Langevin Equation for our proposed bacterium algorithm as in Equation 5.

$$m \frac{\delta v_i}{\delta t} = -\lambda - \tau + \sqrt{2v^2\tau_o\epsilon(t)} \quad (5)$$

where m is mass, the term λ is a frictional term that depends on the reading $C(x, t)$ obtained from the environment. We shall define the λ term as Equation 6

$$\lambda = \frac{v\beta}{C(x, t)} \quad (6)$$

where β is a tuning value and v is the velocity of the agent. τ is the chemotactic term in the spatial function field $C(X)$ with the third right hand term simulating the effect of random changes in the agent's direction where ϵ represents white noise. We introduce $\omega = v^2\tau_o$ as a diffusion coefficient in this work.

This leads to the Fokker Planck relationship of Equation 7 [13] and the stationary distribution of Equation 8. The stationary distribution shows that as k_d increases, the coverage (exploitation) of the spatial function $C(X)$ provided by the agent increases while as τ_o increases, the exploration attribute of the agent increases [14]. From an optimisation perspective, this points to a constrained multi-objective optimisation problem in which one has to obtain a balance of exploration and exploitation while ensuring the agent can optimally and feasibly complete the task within its physical constrains [15] [16].

$$\frac{\partial P(X, t)}{\partial t} = -[\lambda + \tau] \frac{\partial P(X, t)}{\partial t} \nabla C(X) + v^2\tau_o \frac{\partial^2 P(X, t)}{\partial t^2} \quad (7)$$

$$P(X, \infty) = \exp(-[\lambda + \tau]) \frac{C(X)}{v^2\tau_o} \quad (8)$$

This shows that if the function $C(X)$ changes with time (spatiotemporal function), then the spatial distribution $P(X, t)$ would also change with time. In order to do this however, it must be ensured that the speed of our agent is faster than the rate at which the spatiotemporal function changes or damages in the environment propagates. This is because it is not possible for an agent to track a function that is faster than itself. In the case of cracks that develop slowly

overtime, it means that this approach can be used to heal the cracks and stop them. The Equation 8 also shows that the healing behaviour of the agent would be dependent on the distribution of the spatial function $C(X)$ in the environment thereby suggesting that our agent is adaptable.

If it is assumed that the spatial function $C(X)$ covers the entire environment $S(X)$, then Equation 8 provides guaranteed complete coverage as time $t \rightarrow \infty$. Future work would investigate this possibility further and the expected time of convergence.

D. Macroscopic Equations of Swarm Population Dynamics

Since we plan to deploy a swarm of artificial active particles onto a structure, we now present a series of Equations that enable us to understand the dynamics of the swarm at a macro-level.

There are two ways by which artificial active particles could be added onto the structure: Continuously or at once. In this work, we consider the case of all the particles added at once. In this case, the Equation 9 shows that the total number of particles N_T is made up of:

- Number of active particles stuck, N_S
- Number of active particles in liquid state, N_L
- Number of active particles in exploration state, N_E
- Number of active particles in gel state, N_G
- Number of active particles in recruitment state, N_R

$$N_T = N_S + N_L + N_G + N_R + N_E \quad (9)$$

We assume that a certain number of robots are released at the beginning and as such is fairly constant. This is represented by the Equation 10. In this work, ω is a diffusive term which we define as $\omega = v^2\tau_o$; where v is the velocity of the active particle, $\frac{\partial N_E}{\partial t}$, $\frac{\partial N_S}{\partial t}$, $\frac{\partial N_L}{\partial t}$ and $\frac{\partial N_G}{\partial t}$ are the rate of exploration, rate of the number of active particles stuck, rate of the number of active particles in liquid state, and rate of the number of active particles in gel state respectively.

$$\frac{\partial N_E}{\partial t} = N_0 - \left[\frac{\partial N_S}{\partial t} + \frac{\partial N_L}{\partial t} + \frac{\partial N_G}{\partial t} \right] \omega \quad (10)$$

The second term of Equation 10 is given by Equation 11 where ϕ captures the geometry of obstacles as well as the number of obstacles in the environment as a function $f(N_{obs}, G_{obs})$. The third term of Equation 10 is given by Equation 12.

$$\frac{\partial N_S}{\partial t} = N_0 - \phi \left[\frac{\partial N_E}{\partial t} + \frac{\partial N_L}{\partial t} + \frac{\partial N_G}{\partial t} \right] \quad (11)$$

$$\frac{\partial N_L}{\partial t} = \gamma \frac{N_0}{d} N_L f(G_{DS}, M_{DS}) - \left[\frac{\partial N_E}{\partial t} + \frac{\partial N_G}{\partial t} + \frac{\partial N_S}{\partial t} \right] \quad (12)$$

Equation 12 captures the possibility of the damage site geometry (G_{DS}) and magnitude (M_{DS}) or severity of damage affecting the number of robots recruited to the damaged site. d is the distance of the damaged site from the site where the

self-healing robots were released. When the agents find the damaged site, they recruit others. The recruitment rate to the damaged site is given by Equation 13, where we define σ as Equation 14.

$$\frac{\partial N_R}{\partial t} = \sigma N_L \quad (13)$$

$$\sigma = (G_R^i - G_A^i) \exp\left(\frac{\| -r \|}{k}\right) \quad (14)$$

The gain G_R^i repels agents when they are too close to each other while G_A^i attracts agents thereby recruiting them to the damaged site; r is the distance between each agent and k is a constant that can be used for tuning purposes.

The more the agents in the liquid state, the more local influence they generate and this attracts more agents to the damage site. However, the repulsion term G_R^i ensures that the damaged site just gets enough agents to repair the damage leading to some agents breaking away to explore and find other damaged sites. Once in the liquid state, we define η (in Equation 15) as a timeout parameter that moves an agent's motion from liquid state to gel state. In the gel state, it irreversibly starts bounding with the damaged site at a position due to its self-sacrificing repair secretions.

$$\frac{\partial N_G}{\partial t} = \eta N_L \quad (15)$$

III. EXPERIMENTS AND RESULTS

We use a simulation environment to test our methodology with green tiles as shown in Figure 6 indicating damaged sites while white tiles were undamaged. Damaged sites were created using pixels so that if a tile was damaged, a boolean value was set to 1 and 0 otherwise. An agent was able to read the states of 10 by 10 tiles at its location. The value at a damaged site was obtained by adding up the boolean states of these tiles thereby resulting in a damaged value range of $\{0 \dots 100\}$.

We tested two scenarios for active particle deployment: (i) dealing with multiple damaged sites as a single agent and (ii) dealing with a single damaged site with a population of active agents. This is discussed in detail below. For the approach in which we used a population of agents, 20000 contaminated tiles were deployed randomly at $(x, y) = (800, 550)$ with standard deviations of $(\sigma_x, \sigma_y) = (30, 30)$ (See Figure 6). If a tile has been declared damaged by the random process previously, it was redeclared as damaged. As a result, two or more damaged tiles at a pixel would read as one damaged tile. Placing the damaged tiles randomly in this way introduced some noise into the readings obtained by the agent.

Agents were placed in a bounded environment with dimensions $(x, y) = (1000, 800)$. Whenever an agent encounters a boundary, it uses the tumble phase to change its direction. This is equivalent to using a proximity sensor to detect obstacles and subsequently changing direction.

During individual experiments, we estimated how many iterations to afford our artificial particles in ensuring complete coverage of the environment. This was estimated by using a

simple back and forth motions in the bounded environment [17]. The Equation 16 was used to calculate how many iterations it would take to provide optimal coverage of our simulated environment using $\Delta S = 5$ and $v = 20$. Δ is the distance between each parallel trajectory and v is the velocity of the agent.

$$t = \frac{x(1+y)}{\Delta S v} \quad (16)$$

The value obtained was 8100 iterations for our simulated environment. If the value of ΔS were made any larger, the resolution of the cleaning would suffer resulting in areas that are not cleaned or covered. We estimate that in order to be fair, our algorithm should take ≤ 8100 iterations for an environment of this size. We decided to set the iteration limit to a value of 5000. However, during our population based experiments, the agents were allowed to run for a maximum of 4500 iterations [14].




Fig. 4. Showing an environment with obstacles and five damaged locations.




Fig. 5. Showing the results of the agent's distribution in the presence of obstacles.

A. Dealing with Multiple Damaged Sites as a Single Active Agent

As discussed so far, obstacles and boundaries will occur on a large structure in which the agent is likely to be deployed as well as potentially multiple damaged sites. As a result, we introduce obstacles as well as multiple damaged sites as shown in Figure 4. When the agent encounters a boundary, the agent stays at the boundary until the tumble behaviour of the bacterium gives it a randomly chosen direction to head away from the obstacle. In this way, the sensor requirements were kept to the minimum. According to Equation 4, it should be possible for our active particle to find distant separated concentrations of contaminated tiles if given enough time. In order to present a challenge to our proposed approach and proof that our active particle was capable of finding damage sites in the presence of obstacles, a spatial distribution with obstacles was created as shown in Figure 5. The spatial distribution was obtained by using means of $(x, y) = (150, 150)$, $(x, y) = (850, 150)$, $(x, y) = (450, 400)$, $(x, y) = (150, 550)$ and $(x, y) = (800, 550)$ with standard deviations of $(\sigma_x, \sigma_y) = (30, 30)$ for 20000 randomly distributed contaminated tiles. In results shown in Figure 5, the agent was able to find the damaged site with the distribution at $(x, y) = (150, 150)$ receiving more coverage. This is because the active particle was closest to this distribution at the beginning of the experiment. However, this also shows that as distance between the release site and damaged site increases, the less coverage damaged sites will receive. This would call for optimally placing active particle release sites strategically in order to ensure adequate coverage of the large scale structure.

Furthermore, Figure 5 shows that despite the maze, the agent is capable of exploring the environment and finding the various damaged sites in the environment.




Fig. 6. Environment showing 3 obstacles and swarm of active particles.

B. Applying a Swarm of Active Particles to an Obstacle Filled Environment

A swarm of 25 active particles were then deployed in an environment with multiple obstacles and a single damaged site as shown in Figure 6. Different values of the exploration parameter v ($v = 7$; $v = 14$; $v = 21$) were tested. The graph in Figure 7 shows that when v increased in value from $v = 7$ to

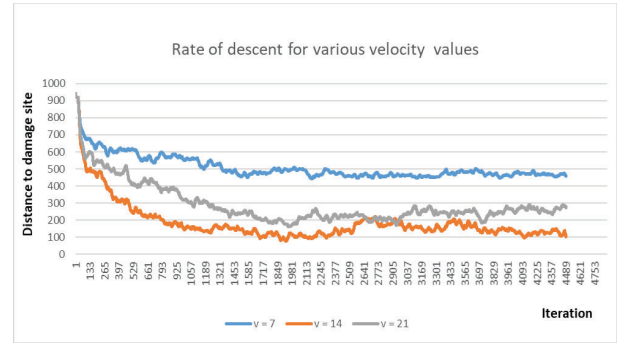


Fig. 7. The average rate of active particles finding the damaged site with increases in v values.

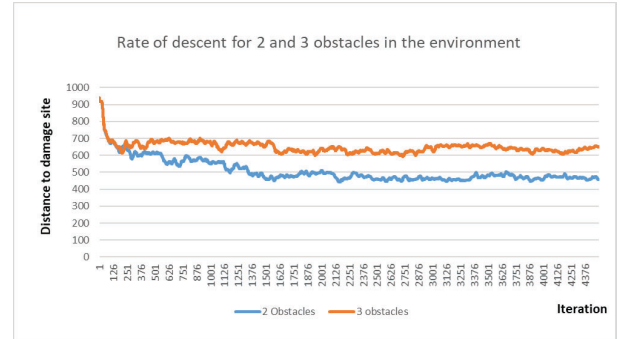


Fig. 8. The average rate of active particles finding the damaged site with an increase in the number of obstacles.

$v = 14$, the rate at which active particles increased. However, at $v = 21$, the rate reduced compared to $v = 14$. This results show that there is an optimal exploration term that needs to be applied to the swarm of active particles to find damage sites quickly.

Also, we tested the consequences of adding more obstacles to the environment. Figure 8 shows that as the number of obstacles in the environment increases, the number of active particles finding the damaged site reduces. This shows the impact of obstacles on our approach. However, as discussed previously, if given enough time, a single active particle as the potential to find the damaged sites in an environment. Once an active particle finds a site, it recruits passing particles to it using the Equation 14. Once recruited, the distance between individual active particles was regulated by the Equation 14.

IV. CONCLUSIONS

In this work, we have presented the application of active particles to the healing of large scale high value structures. We have also addressed one of the challenges raised by the active particle research community, that of dealing with obstacles in the environment. According to [6], this is one of the challenges that need to be addressed in the use of active particles for various applications. Other challenges include how to engineer useful emergent behaviours, identify compact propulsion mechanisms and energy supplies capable of lasting

for the whole particle life cycle and how to scale down active particle dimensions towards the nanoscale. The above can be viewed from the perspective of a constrained multi-objective optimisation problem in which a balance between exploration of an environment and its exploitation needs to be achieved while taking into consideration the technical feasibility (sensing and actuating at nanoscales) and optimality (in terms of energy usage) of an active particle.

In this work, we have focused on understanding how active particles behaviour could be controlled in crowded and complex environments filled with obstacles. We presented a minimalist active particle controller that uses a random exploration behaviour and a sensor to detect the condition of the environment. We have shown in our results that given enough time, a minimalist active particle has the potential to navigate an environment filled with obstacles and find multiple damaged sites. As seen in the results, as the number of obstacles increase, the number of agents finding a damaged site reduces. This supports the challenge highlighted in [6]. Nevertheless, this can be solved by increasing the exploration capability of active particles. An increased exploration capability gives them a more diffusive behaviour and an almost "gas" like behaviour enabling them to navigate around obstacles. Nevertheless, we also discovered that there is an optimal value for this exploration term beyond which performance drops.

In future work, we will investigate how to deploy this concept on physical agents. It will be interesting to also see how we can deploy healing materials to a damaged site as well as how agents respond to a generated field from the site of damage. Currently, we have investigated a scenario where the damaged site has a local influence. With a generated field influencing the recruitment ability of the agents to a site, it is possible that the rate at which agents find the site increases. However, this might lead to difficulties in navigating the obstacles on a structure due to the need to resolve the conflict between the pull of the generated field and escaping from the obstacle. Furthermore, the applications of constrained multi-objective optimisation might offer some insights into the research and development of nanoscale active particles.

REFERENCES

- [1] T. Ghidini, "Materials for space exploration and settlement," *Nature materials*, vol. 17, no. 10, pp. 846–850, 2018.
- [2] R. Putzar, M. Gulde, D. Sabath, H. Fiedler, G. Drolshagen, A. Braukhane, A. Horstmann, C. Wiedemann, M. Schimmerohn, and F. Schäfer, "Measuring impact craters on the iss columbus module," *69th International Astronautical Congress*, 2018.
- [3] S. Rizkalla, M. Dawood, and D. Schnerch, "Development of a carbon fiber reinforced polymer system for strengthening steel structures," *Composites Part A: Applied Science and Manufacturing*, vol. 39, no. 2, pp. 388–397, 2008.
- [4] T. Wareham, "Designing robot teams for distributed construction, repair, and maintenance," *ACM Transactions on Autonomous and Adaptive Systems (TAAS)*, vol. 14, no. 1, pp. 1–29, 2019.
- [5] E. Petito, S. Momi, and P. Gresele, "The migration of platelets and their interaction with other migrating cells," in *Platelets in thrombotic and non-thrombotic disorders*. Springer, 2017, pp. 337–351.
- [6] C. Bechinger, R. Di Leonardo, H. Löwen, C. Reichhardt, G. Volpe, and G. Volpe, "Active particles in complex and crowded environments," *Reviews of Modern Physics*, vol. 88, no. 4, p. 045006, 2016.

- [7] S. Gupta, H. W. Kua, and S. D. Pang, "Autonomous repair in cementitious material by combination of superabsorbent polymers and polypropylene fibres: A step towards sustainable infrastructure," *arXiv preprint arXiv:1706.02680*, 2017.
- [8] Z. Ren, T. Wang, W. Hu, and M. Sitti, "A magnetically-actuated untethered jellyfish-inspired soft milliswimmer," in *Robotics: Science and Systems*, 2019.
- [9] H. Zhao, J. Leclerc, M. Feucht, O. Bailey, and A. T. Becker, "3d path-following using mrac on a millimeter-scale spiral-type magnetic robot," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 1564–1571, 2020.
- [10] K. M. Passino, "Biomimicry of bacterial foraging for distributed optimization and control," *IEEE control systems magazine*, vol. 22, no. 3, pp. 52–67, 2002.
- [11] W. Ebeling and F. Schweitzer, "Self-organization, active brownian dynamics, and biological applications," *arXiv preprint cond-mat/0211606*, 2002.
- [12] F. Schweitzer, "Brownian Agent Models for Swarm and Chemotactic Interaction Brownian Agents," in *Fifth German Workshop on Artificial Life*, 2002, pp. 181–190.
- [13] W. Ebeling and F. Schweitzer, "Self-Organization, Active Brownian Dynamics, and Biological Applications," *Nova Acta Leopoldina*, vol. 88, no. 332, pp. 169–188, 2003.
- [14] J. Oyekan, D. Gu, and H. Hu, "Stochastic ant agent for priority-based coverage," in *2013 IEEE/ASME International Conference on Advanced Intelligent Mechatronics*. IEEE, 2013, pp. 1140–1145.
- [15] W. Ning, B. Guo, Y. Yan, X. Wu, J. Wu, and D. Zhao, "Constrained multi-objective optimization using constrained non-dominated sorting combined with an improved hybrid multi-objective evolutionary algorithm," *Engineering Optimization*, vol. 49, no. 10, pp. 1645–1664, 2017.
- [16] Z. Fan, Y. Fang, W. Li, J. Lu, X. Cai, and C. Wei, "A comparative study of constrained multi-objective evolutionary algorithms on constrained multi-objective optimization problems," in *IEEE Congress on Evolutionary Computation (CEC)*, 2017, pp. 209–216.
- [17] I. Rekleitis, A. P. New, E. R. Samuel, and H. Choset, "MAC Versus PC: Determinism and Randomness as Complementary Approaches to Robotic Exploration of Continuous Unknown Domains," *The International Journal of Robotics Research*, vol. 19, no. 1, pp. 12–31, 2009.