Roke

INNOVATION THOUGHT LEADERSHIP

ROBOTS AT SCALE

Abstract:
What does the future of autonomy look like? We see a tension between different models of uncrewed vehicles: high-value, high-capability singletons vs commoditised systems at massive scale. This paper makes the argument for robots at scale and digs into some of the surprising reasons why that's hard to achieve. We consider not just scale in terms of numbers, but increasing levels of autonomy, co-operation and system complexity. We're sharing our vision, the benefits and the challenges, to encourage collaboration and realise the potential that we can see. The future of autonomy is unlimited.

CONTENTS

1

INTRODUCTION

In this paper we explore what it means to operate robots at scale. First, we'll look at this from the perspective of the operator, to understand what's difficult. Then we'll see how we can provide solutions to these challenges.

Imagine you're the operator. Choose a mission. It could be surveying a large area, searching an urban disaster zone, or a military surveillance mission. Something spread over a wide area and probably hostile to people. Give yourself a deadline - you're looking for a survivor or another mission depends on this - time matters. You're going to use uncrewed autonomous vehicles - robots - to meet the mission goal.

Let's start by thinking how we might tackle this without robots at scale. The task is too big for any kind of consumer technology. We're assuming that you'd need at least a few high capability (and high cost) robots. Let's assume that this is important enough that you can get them assigned. It's likely there will need to be a risk assessment: these are expensive assets, quite possibly in high demand. Now you're going to task each of them, possibly through a proprietary control interface, specific to the robot. If the job is taking too long, well, let's hope that you can get another robot assigned.

Is there a better way? We think so.

Now let's assume that you've got access to a large pool of low to mid capability robots (tens or even hundreds strong). They're a mix of ground and airborne; most with a core set of sensors, but some with one or more specialised payloads. How would you hope that this would play out?

If we were deploying such robots at scale today, what might that look like? Hard work is the most likely answer. Yes, there are solutions for deploying large numbers of drones. Those tend to be for displays, though. Each drone works from a meticulously pre-planned set of instructions. That's no good for this sort of mission. Your starting point is to try to work out what's available for your use. What state are all your robots in? Check if they've got batteries; if they have, check the charge level. Check the equipment fit. You probably need to connect to each of them manually to configure them for the mission. Different groups of UxVs may have different control and supervision interfaces. The robots may make assumptions about available network connectivity. The robots may not actually be that smart. It's not obvious why this is better - but that's because we're looking at today's technology.

1

Let’s jump to the near future and see how robots at scale should work.

There’s a pool of sixty robots of varying types. In a single view, you can see the status of them all. Most are at a high state of readiness and are good to deploy now. A few are on scheduled maintenance or have flagged issues through self-test, so are unavailable. A handful are not fully charged but are on a self-managed charging schedule; they’ll be ready within the hour.

You specify the mission aims to a planner. The planner suggests a few combinations of robots. You reject a couple of options as impractical and adjust one or two constraints. After two iterations of the planner, you’ve got three good looking plans. One plan takes advantage of some additional sensor packs, suggesting they be paired with some of the selected vehicles. You task the human maintainer to attach the sensors: it’s all completely plug-and-play.

Once you commit to the plan, the robots are configured for the mission: the system downloads their local squad objectives, the communications plan, and any other data they need. The planner shows the robots as assigned. Once they’re ready, they self-load for deployment.

Once deployed, the focus shifts to the supervision system. One pane of glass shows the entire mission status. One robot is running low on power: it’s been using a radar sensor with a relatively high power drain. It may be that it’s done everything it needs to, in which case it can return to base. Otherwise, it’ll fly to a ‘tanker’ ground vehicle and recharge.

Another robot gets tangled in something and must be considered effectively lost. That’s no problem, because there are lots of robots, so the others simply cover.

If we now look back at the operator view, we can unpick the technical challenges. The challenges we can derive centre around:

- Asset management;
- AI-assisted planning;
- Co-operating, smart robots;
- Modular robots;
- Built-in self-maintenance such as charging and testing;
- Common supervision protocols; and
- Dynamic communications.

These key challenges form the underlying drivers for our work. Let’s look at them in more detail.

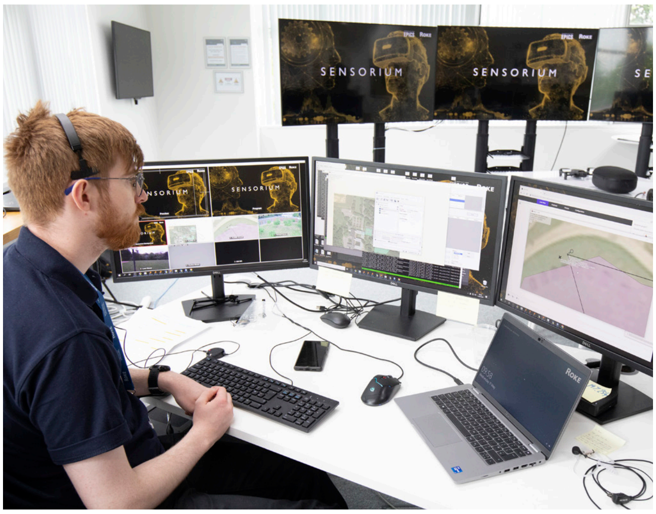

2

THE FUTURE OF ROBOT DEPLOYMENTS

Outside of drone displays, uncrewed vehicle deployments at scale are rare. The main choice today is the complexity or capability of the robot that we use. Typically, the best way to achieve more is to get a more capable robot – and that will often come with a significant increase in cost. We could improve our ability to search a complex disaster site for survivors by using a drone that can fly faster, for longer and with a better camera (or other sensors). Such a drone – if it existed – would be expensive.

Our assertion is that we should shift from this model of individual deployments requiring high levels of human control or expense, to large-scale deployment supported by automation. If we enable robots at scale, then we can run many units in parallel to cover our search area in a given time. If thermal imaging helps certain parts of the search, we need only include some suitably equipped robots.

This model gives us more flexibility in matching robot capabilities to missions, as well as trading capabilities across different missions. If there are more robots available that aren't needed elsewhere, let's do the search quicker. If we later need to redeploy some of those robots to another site, we can calculate the effect on the remaining search time. A whole suite of statistical optimisation tools is now available to us in a way that single, high-capability robots cannot benefit from.

The core idea is that smart robots don't have to be complex. We can trade the sophistication of individual robots for large groups of cheaper, commoditised units. It's not an approach that will work everywhere, but we believe that it is widely applicable and will bring a step-change in overall capability.

2

This approach drives a shift to modularity and commoditisation. Roke is adopting modular robot designs, exploiting affordable, off-the-shelf electronics. This reduces cost of manufacture and repair and enables simple reconfiguration for different roles. Open innovation and collaboration also reduces costs and expedites progress; Roke believes in an open and non-proprietary approach. This is particularly important for payloads. Core functions are likely already integrated in hardware and software. Changing any of the navigation sensors, even moving its location on the robot body, needs integration and testing. Adding payloads to scan or interact with the physical world should be trivial. We expose standard interfaces (serial, USB, etc.) allowing everything from near-infrared cameras, magnetometers, or chemical sensors to be easily added (and removed). This lets us tailor a squad of robots for specific tasks. With commoditised robots at scale, we can keep a variety of units ready to go and intelligently deploy the ones best suited to the task at hand.

By making use of modular, commodity hardware and understandable, explainable algorithms for autonomy, we can achieve much more than singleton robots with rigid purposes. These systems should make use of open architectures to maximise options for interoperability, scalability and extensibility. We can then reuse capability across a range of tasks, free up human operators for critical decisions, and achieve co-operation through interoperability.

MANAGEMENT OF ROBOTS IN CHALLENGING ENVIRONMENTS

3

Existing swarm management systems show that robots can be coordinated for specific tasks, with sufficient pre-planning. Agricultural drones already optimise farming practices through crop monitoring and precise pesticide spraying. Underwater sensor networks employ swarms of aquatic robots for seabed mapping, marine life monitoring, and data collection in harsh underwater environments. Today's operations tend towards predefined tasks in well-defined, uncontested spaces needing limited interaction. Robots operate based on simple rules and local information. In more complex, dynamic, and contested environments things get harder. We can split the types of robot management we're considering into system and mission management.

System management is about making sure that we have our complement of robots ready to go, whenever they're needed. Like any response team, military or emergency response, robots can be at different readiness levels. Like any team of crewed vehicles (ambulances, armoured vehicles), there will be a maintenance schedule to keep everything mechanically sound. Our robots are also computers so, like an IT estate, we need to keep track of patch state and installed software versions. As we dig into this, it is apparent and perhaps surprising that robots at scale come with a substantial asset management challenge. Anything that eases the burden of fault finding and diagnostics will also help. There is a fundamental need to keep track of all our robots, their capabilities and their status; and to use this information to keep them mission ready.

This is where we can leverage autonomy. Rather than expecting a human to keep batteries charged, robots become self-charging. Robots can self-test and flag potential issues to a human maintainer. This self-management performed by individual robots enables scale through reducing management required by a human operator.

3

Mission management is the other core component. Mission management matches robots to tasking, deconflicts requests and configures the selected units. A planner works with the given objective and information about the capabilities and readiness of the available robots. Hierarchical planning may help tame the complexity of the planning problem. Our planning system needs to present options, allow the user to adjust constraints and iteratively refine plans. Automating the generation of potential mission assignments and execution tactics supports and improves the human operator's decision making.

Once we've identified the set of robots that will be used to meet the objective, the mission system needs to configure and task the robots. Attention then shifts to the supervision system that monitors and provides input to the squad. Planning and management systems (centralised or distributed) divide large-scale operations into more manageable squads; breaking down a mission into more understandable and executable chunks. This approach is described in our earlier paper on squads vs swarms.

To achieve scale that enables us to address complex missions, we require autonomous management of individual robots, a simple interface for resource management, and an intelligent planning system. These capabilities allow us to go beyond simple, pre-defined tasks to proactive and reactive systems that may not have the luxury of significant time to plan well in advance of deployment.

EMERGENT BEHAVIOURS OF INTELLIGENT, COLLABORATIVE, LARGE-SCALE ROBOTICS

4

To build systems that will react without complete human direction requires intelligence at both the individual and group level. This enables us to create capability that can receive a high-level order and execute the mission, without needing every stage rigidly defined. Much like how we might set out a requirement for a team of humans, we want to state the end goal and for the team to self organise and carry out the tasks required.

Nature continues to be a source of inspiration for complex, large-scale coordination. We can learn from research into, for example, cluster-based underwater wireless sensor networks inspired by murmurations. This and similar research shows how we can achieve a co-ordination without massive computational power. For more complex computational problems, there is substantial research into heuristics and approximation algorithms.

Some intelligence will come from machine-learned algorithms. In this we face two related issues of scale: training and co-operation. To train a machine learning algorithm, a scalable simulation environment is essential. For approaches like reinforcement learning, the overhead and risk of standing up a large swarm of robots for every training cycle will limit training time or increase costs, or both. Iterating simulation and real-world testing help us understand the level of simulation fidelity that we need to achieve the functionality we need. We have experimented with everything from simple ‘grid worlds’ through to sophisticated simulations in the ROS2 ecosystem. These have wildly different processing loads and can be used effectively for training at different levels of abstraction.

Combining approaches lets us efficiently layer machine intelligence.

There are several other properties that we want robots at scale to exhibit. Our challenge is to design the individual units such that these properties emerge.

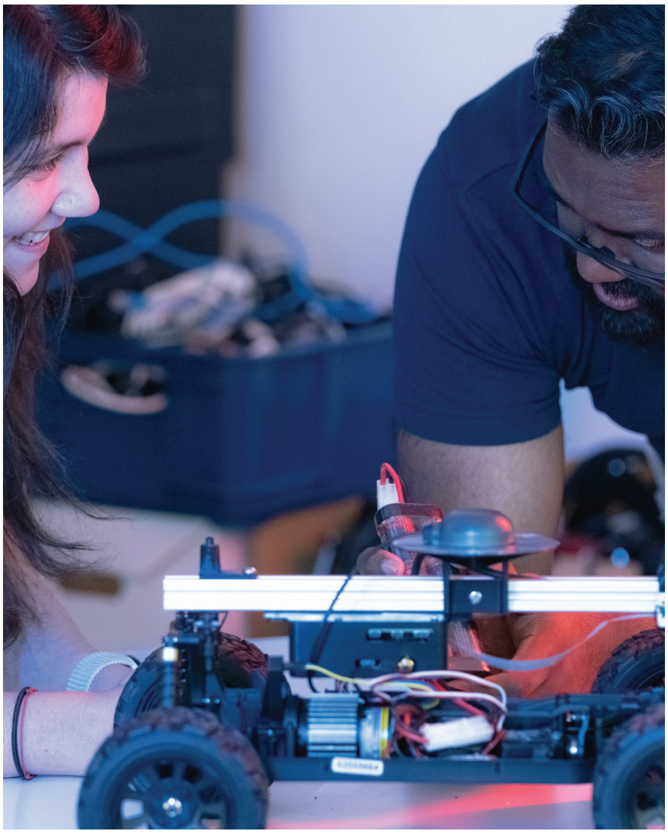

4

We shift away from both highly centralised command-and-control and any requirement for fully connected communications between robots. Otherwise, our challenge of scale includes bandwidth and information flow. Devices will communicate mostly with peers and then distribute less time-critical messages by more efficient, multi-hop diffusion. Similarly, command and control messages can be collated or flooded to avoid every unit trying to maintain a direct connection to base. Communications will self organise, switching modes as needed, with some units providing backhaul or acting as local hubs. Visual cues between devices might be useful for co-operation, as well as enabling better human-robot collaboration. A multi-faceted approach will bring us closer to seamless autonomous operation.

We then need to add layers of behaviours and strategies to meet a range of mission goals. For some large-scale deployments, co-existence might be enough; in some search scenarios, the robots can operate near independently in a defined region. A more complex search scenario may require a different approach. Consider a search space that consists of both indoor and outdoor locations. There may be terrain that is difficult to navigate, stairs, doors, or even chemical contamination that make areas impassable to some members of our squad. A human will avoid the contaminated zone, whereas a small, commodity UGV cannot navigate stairs. In this scenario, we require co-operative and collaborative behaviours, including human-robot collaboration. We need to be sure that our approaches to collaboration can scale. Whether we're deploying five or fifty robots, they must be able to work together. For some behaviours that may mean dynamically switching the algorithm or model. Switching may even happen during a mission if the squad composition changes dramatically due to a loss or addition of members.

Exploiting research to solve real-world problems is part of our approach. One example of research into the complexity of scaling squads is Heterogeneous Decentralised Receding Horizon Control (HD-RHC), a novel control technique for UAV swarms that optimises mission planning and movement by considering individual capabilities and limitations of different UAV types within the swarm . Japan has experience of using several types of remote-controlled robots in work environments that humans cannot enter directly. Their solution achieves flexible remote control using Robot Service Network Protocol (RSNP) to improve cooperation. Roke have similarly used the SAPIENT standard to task autonomous surveillance systems. Interoperability using common protocols and open standards will be critical for future deployments of robots at scale, and at speed.

SUSTAINABLE OPERATION

5

To manage and operate any infrastructure at scale, key challenges will emerge due to the sheer amount of hardware in operation. Taking lessons from large-scale storage and compute facilities, we must eliminate single points of failure and build in high availability, resilient, and intelligent resource management. That includes designing for fault detection and repair in both hardware and software, through self-test and patching.

To operate robots at scale, we need to design for device failure – as the number of robots increases, the probability of device failure goes up. Any algorithms and architectures that we use must be able to cope with such a failure. For example, we use distributed decision making or, at least, self-organisation that is resilient to device failure; if a single decision-making robot fails, it cannot compromise the rest of the team.

As we scale up, the challenge of resource management such as power also increases because of the number of individual units. We cannot escape the need to charge (and know the charge level) of a lot of devices. Intelligent task allocation, energy harvesting techniques, or improved battery technologies can help, but to reduce the burden on a human operator, we need robots that self-charge. Some domestic robots, like vacuum cleaners, already do this. However, robots at scale will need to co-ordinate their charging schedules and manage their power needs as a group for each deployment. In the squad model, we can imagine 'tanker' robots that exist to provide roving recharge points, especially to constrained devices, like an airborne vehicle.

Security is another concern with an increasing number of concurrently operating robots. Trust is critical. If we assume that each device needs to be able to securely identify itself, then every robot needs a unique cryptographic identity. Secure key management makes assumptions about available communication methods: when and how can devices be keyed. Keys may be set during manufacture or programmed (and re-programmed) later. The details of key management must fit the deployment lifecycle. When we recover a collection of robots after an operation, especially from a hostile environment, we will assess how much we trust the state of the device. We might not have maintained communication throughout the operation – we won't know everything that's happened to the robot. Secure attestation of the robot's state becomes a useful way to prove that the device integrity has been maintained.

5

Sustainability and ethical considerations are important, and transparency of this will be crucial to growing public trust and acceptance. At Roke, we design robots for durability, repairability, and responsible material usage throughout their lifecycle to support sustainability. Ethically, we need to integrate safeguards within the management systems to ensure adherence to legal and ethical standards while operating autonomously. Operations need to fit with regulations, such as those designed to manage spectrum use and prevent interference. Certain jurisdictions may limit operations due to concerns over noise or visual pollution, particularly in the case of drones. For aerial robotics, airspace use regulation is currently a significant constraint.

We must also address affordability to realise the potential of large-scale robotic systems. Commercial Off-The-Shelf (COTS) components, as exemplified by our Silverfish platform, are the main path to affordability. The Silverfish, an off-the-shelf remote-control vehicle enhanced with off-the-shelf computing and sensor capabilities, provides a model for cost-effective, large-scale robotic deployment. Using readily available, modular components enables adaptation to the unique requirements of each mission; an economical solution without sacrificing functionality.

Modular design principles significantly contribute to affordability: interchangeable parts streamline the manufacturing process, simplifying production and maintenance. This approach allows for the easy replacement of faulty components without servicing or replacing the entire robot, saving both time and expense. These have been areas of research for some time, and we are now exploring this at Roke to move research into the field.

Cost reduction is also achieved through economies of scale and commoditisation. Large-scale production leads to better prices for components, optimisation of manufacturing processes, and distribution of fixed costs over a greater number of units. Design simplicity, concentrating on essential capabilities while avoiding unnecessary complexity, further controls costs while maintaining performance.

Considerations of affordability extend beyond initial production and deployment phases. Strategies for maintenance and lifecycle management influence the total cost of ownership. Designing for disposability or ease of repair can ensure the provision of long-term value from robotic systems. Where we need custom hardware and software, we design for re-use and use open interfaces to extract the most value from anything bespoke. Intelligent algorithms allow lower-cost robots to deliver impressive performance, reducing dependencies on expensive hardware.

A pragmatic, cost-aware approach paves the way for broader adoption of large-scale robotic systems. We believe this offers the ability to do more and do it faster than could be achieved by specialised, high-cost platforms.

ROKE'S RESPONSE TO ROBOTS AT SCALE

6

Roke has a multifaceted approach to tackle the challenges of large-scale robotic operations. Our strategy coalesces around three central pillars:

6.1 DEVELOPING VERSATILE PLATFORMS

Central to our approach is the development of adaptable platforms that allow exploration of robots at scale. Our in-house robotic platform, Silverfish, is one such capability. Silverfish is a cost-effective robotic platform, repurposed from a commercially available remote-control car. This allows us to conduct large-scale experiments at low cost. It accelerates our understanding of operating robots at scale, distributed intelligence, and collective decision-making processes. In addition, by developing an interoperable software platform for the Silverfish using the SAPIENT protocol, we can redeploy intelligence to other UxVs under a common control system. This allows supervision of a broad range of autonomous robotic capability through a single command and control interface, as well as enabling rapid integration.

6.2 EXPERIMENTING WITH DIVERSE OPERATIONS

We are exploring the dynamics of heterogeneous robots operating together in real-world conditions. Bringing Silverfish together with other commercial robot platforms gives us a way of exploring squads as well as scale. Moving away from lab tests to real-world experimentation gives us a fundamentally different understanding of what matters. Through this, we better understand the challenges and potential of operating robot types and capabilities in unison. This will lead to insights into multi-robot cooperation, interoperability of diverse systems, and the development of adaptive control algorithms that can work at scale.

6.3 ACCELERATING INNOVATION THROUGH COMPETITION

To foster innovation, we are creating competitive environments (our 'Squad Game'). Pitting teams against each other creates an environment to hot-house strategies for effective robotic squads. Through this initiative, and by building partnerships, we want to push the development of robotics control strategies, cooperation algorithms, and mission planning. We are also using simulators alongside real robots to enhance our experimental process. Simulation is important for safe and cost-effective development, enabling training, testing and prototyping. It also allows us to use an iterative process to benchmark simulated results against real-world operations, ensuring their reliability and validity.

Roke's approach to large-scale robotics reflects our commitment to innovation, collaboration, and solutions. We are excited about unlocking the potential of large-scale robotic intelligence.

7

CALL TO ACTION

The advancement of robotics demands collective effort. Collaborative development and open-source implementations contribute to shared innovation, reduced development costs, increased knowledge sharing and accelerated progress. Roke invites researchers, technologists, and policy-makers to join us in this next stage of robotic operations and development.

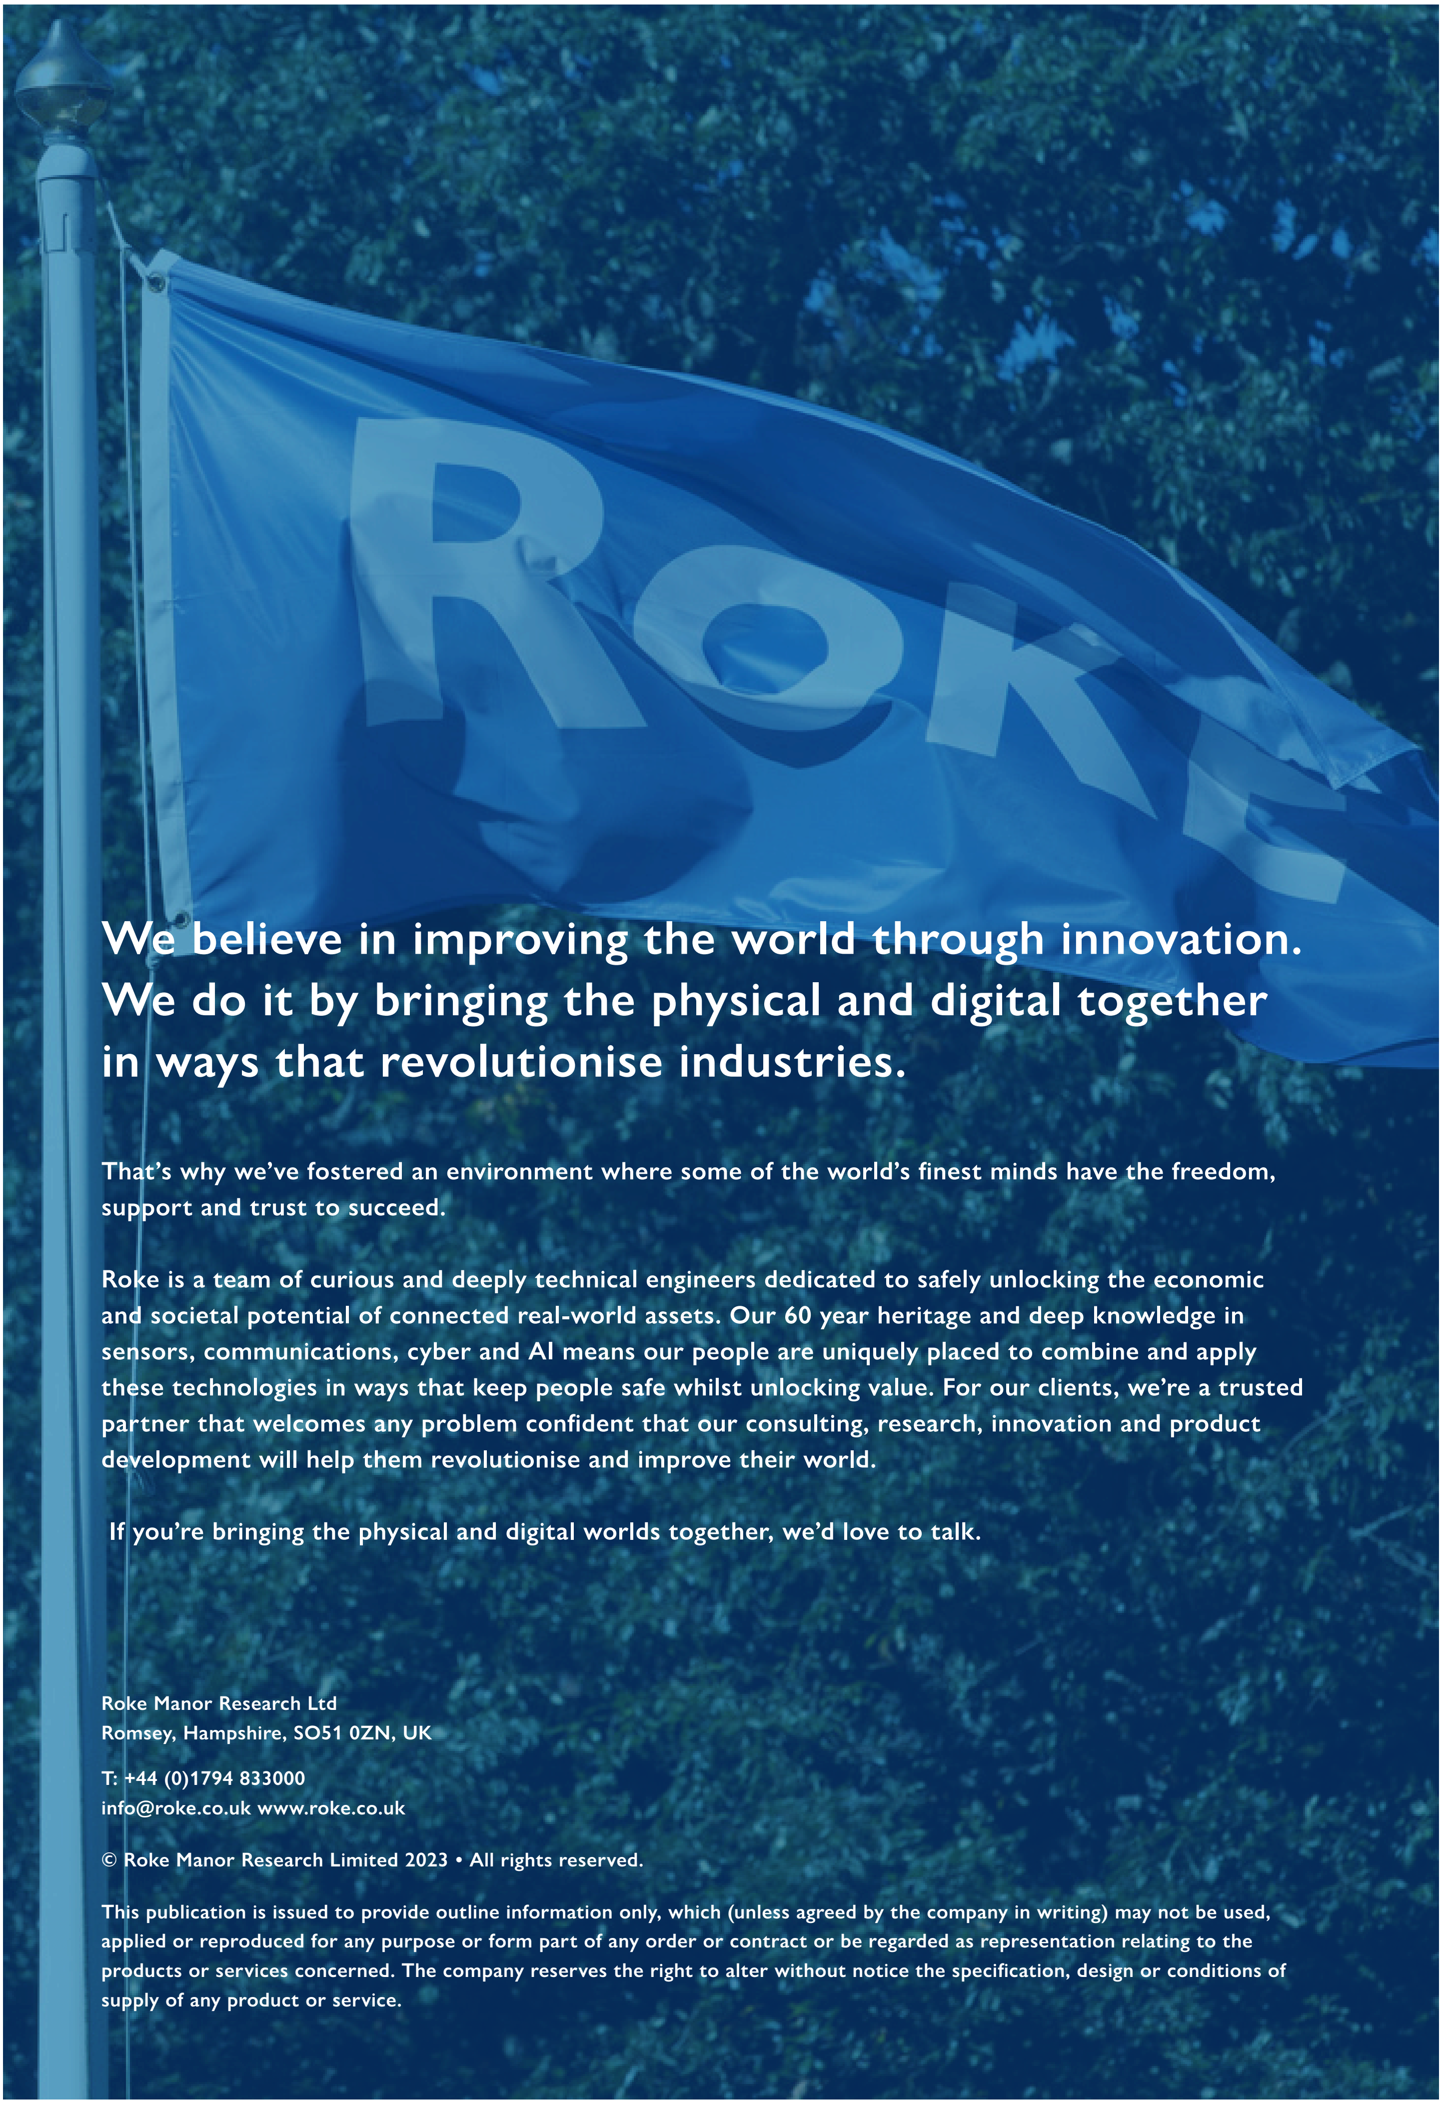

We believe in improving the world through innovation. We do it by bringing the physical and digital together in ways that revolutionise industries.

That’s why we’ve fostered an environment where some of the world’s finest minds have the freedom, support and trust to succeed.

Roke is a team of curious and deeply technical engineers dedicated to safely unlocking the economic and societal potential of connected real-world assets. Our 60 year heritage and deep knowledge in sensors, communications, cyber and AI means our people are uniquely placed to combine and apply these technologies in ways that keep people safe whilst unlocking value. For our clients, we’re a trusted partner that welcomes any problem confident that our consulting, research, innovation and product development will help them revolutionise and improve their world.

If you’re bringing the physical and digital worlds together, we’d love to talk.

Roke Manor Research Ltd
Romsey, Hampshire, SO51 0ZN, UK

T: +44 (0)1794 833000
info@roke.co.uk www.roke.co.uk