UK Atomic Energy Authority

UKAEA-RACE-CP(19)04

Samuel Jimenez, Guy Burroughes

A conceptual framework for automated maintenance of nuclear fusion power plants

A conceptual framework for automated maintenance of nuclear fusion power plants

Samuel Jimenez, Guy Burroughes

This paper has been submitted to
CASE 2019 - International Conference on Automation Science and Engineering, University of British Columbia, Vancouver, BC, Canada, 22- 26 August 2019

A conceptual framework for automated maintenance of nuclear fusion power plants*

Samuel Jiménez, and Guy Burroughes

***Abstract*— The extreme environmental conditions inside future nuclear fusion power plants will mean that maintenance activities must be performed without human access. Current teleoperation paradigms do not enable sufficiently fast maintenance cycles, and automation is seen as a potential solution. Transitioning to an automated maintenance scheme will require acceptance from many different stakeholders with varied expertise in a very wide range of fields. Thus, communicating abstract concepts related to automation and autonomy requires clear, well-defined concepts to provide a common footing which avoids confusion and misrepresentation. Hence, this paper presents a conceptual framework which lays the foundation on which to build an automated maintenance strategy that is suitable for future fusion power plants. The aim is to provide a clear and common starting point from which to define goals, manage expectations, guide research, and assist the development of suitable standards and regulations.**

I. Introduction

Nuclear fusion is being pursued as one of the potential solutions to the world's energy challenges [1]. The most developed concept for a fusion reactor is the tokamak, a large torus-shaped vacuum vessel which houses a magnetically confined superheated plasma (~150 million °C) [2]. Here, hydrogen isotopes fuse together, producing highly energetic neutrons. These collide into surrounding structures to transfer their kinetic energy, which is converted into electricity via a conventional steam cycle. Unfortunately, these collisions also damage the materials, leading to structural weakening [3]. In future power plants this will require routine replacement of significant amounts of reactor hardware. To complicate matters, intense neutron bombardment can also cause materials to become radioactive, leading to short-lived but very high radiation levels. Hence, all maintenance must be done without human access.

The largest fusion reactor in the world is the Joint European Torus (JET). Despite being an experimental reactor, the presence of toxic chemicals and trace amounts of radiation demand that maintenance and upgrade activities inside the vacuum vessel are performed using remote handling techniques. A teleoperated, haptic master-slave servo-manipulator (MASCOT) is deployed on two 12 m long snake-like booms [3], acting as a proxy for human operators, Fig. 1 and 2. A five-person team operates the system from a control room, performing maintenance tasks such as replacing inner-wall tiles, installing new experimental devices, performing measurement surveys, removing dust by-products, etc. In total, around 50 full-time staff are directly employed in designing, operating and maintaining the JET remote handling system, which includes specialist equipment, tooling, software, processes and training.

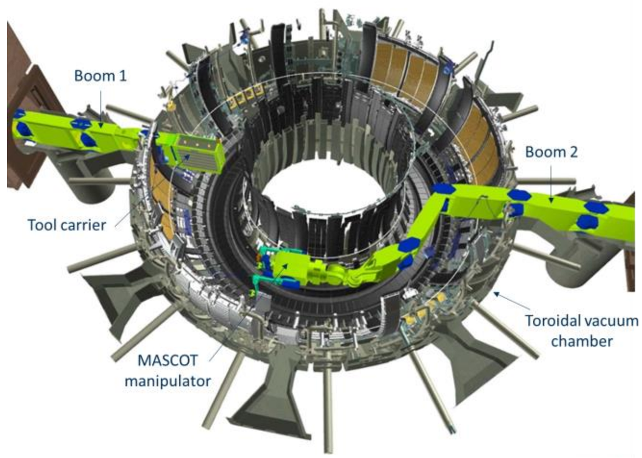

Figure 1. Cutaway of the JET reactor vessel, showing the deployment of the remote maintenance system. Image courtesy of EUROfusion.

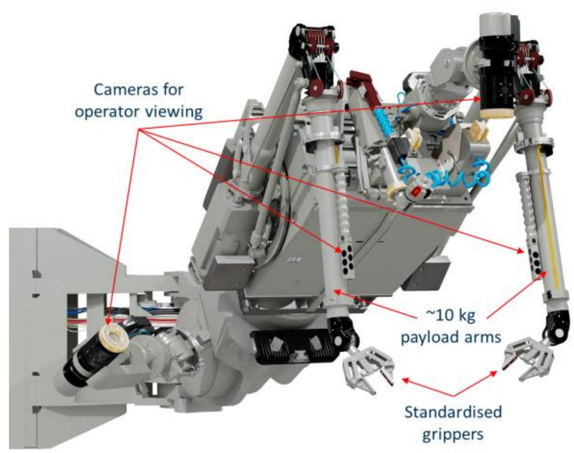

Figure 2. JET MASCOT servomanipulator. Image courtesy of EUROfusion.

As a highly versatile system, the JET model is ideal for use on experimental machines, which undergo constant modifications to meet the needs of their scientific campaigns. However, as reactor designs evolve towards future power plants maintenance durations must be reduced to a minimum

* This work has been carried out within the framework of the EUROfusion Consortium and has received funding from the Euratom research and training programme 2014-2018 and 2019-2020 under grant agreement No. 633053. The views and opinions expressed herein do not necessarily reflect those of the European Commission.

The authors are with the Remote Applications in Challenging Environments unit, United Kingdom Atomic Energy Authority, Culham Science Centre, Abingdon, OX14 3DB, UK (samuel.jimenez@ukaea.uk, guy.burroughes@ukaea.uk).

to achieve commercially viable plant availability. Hence, the maintenance systems must become faster and more efficient, with performance comparable to industrial automation.

Automated maintenance is therefore identified as a key technology for the commercial development of fusion power. However, the complexity and extreme conditions of fusion reactors make the transition from teleoperation to automation a significant technical and strategic challenge. Adopting automated maintenance strategies will have a profound impact on reactor design, so early buy-in from reactor designers is essential. Further, the deployment of automation in a high-consequence environment such as a nuclear site requires acceptance from regulators, who will expect certain guarantees of safe behavior of the automated system.

The design of maintenance strategies and systems is an inherently multidomain problem that spans the depth of the complexity hierarchy, interfacing with every aspect of the maintained system. This means that multiple different interpretations of automation concepts can coexist, leading to mixed expectations of its benefits and limitations. Thus, a first step towards the adoption of automated maintenance is to set out clearly what automation is, and what it is not, in the context of a fusion reactor. Hence, this paper presents a conceptual framework which lays the foundation on which to build an automated maintenance strategy that is suitable for future fusion power plants. The aim of this work is to provide a clear and common starting point from which to define goals, manage expectations, guide research, and assist the development of suitable standards and regulations. First, a definition of "automated system" is given, from which a model to define the automation implementation strategy is proposed. A fusion-relevant definition of "autonomy" is then provided, followed by a proposed categorization of degrees of autonomy.

II. Challenges for Automation in Fusion

The environmental conditions inside fusion power plants will be harsh. The radiation levels during operation will be such that no maintenance equipment could be permanently housed inside the machine, so that maintenance systems will have to be sufficiently mobile to deploy into the working space. This is a considerable departure from typical practice in industrial settings, for instance requiring service connection, commissioning and calibration to all be achievable without human presence. Even during maintenance, radiation levels in the order of 0.1-1 kGy/hr are predicted, sufficient to severely limit the lifetime or allowable complexity of conventional electronics [4]. Radiation levels will also impact material selection, potentially restricting standard choices such as PTFE plastic or elastomer seals. Perhaps the largest challenge to overcome is the complete lack of human access into such environments, which means that any failures must be recoverable, at the very least, through teleoperated means.

In addition, hardware geometries will be constrained by plasma physics requirements, and space near the reactor core will be at a premium. This leads to confined, complex spaces in which handling systems will have to execute intricate kinematic paths. Some robotic systems are likely to become contaminated by chemical and radiological hazards so will have to be decontaminated after use [5]. Some of these "maintenance of maintenance system" operations will also require remote or automated capabilities. Industrial automation will typically rely on a wealth of experience and in-depth understanding of the process being automated. However, this knowledge does not yet exist for fusion reactors, introducing a further engineering design challenge.

A further challenge is that systems deployed in fusion power plants, being nuclear regulated sites requiring high capital investment, will need to demonstrate safe behavior to designers, regulators and investors. This is a significant challenge for automation involving novel Artificial Intelligence (AI) techniques.

III. Definition of Automated System

The concept of "automated system" can take on many different meanings and has resisted attempts at formalizing a definition [6]. When discussing automation of complex systems of systems, several interpretations tend to coalesce onto the same design, creating misunderstandings about the capabilities and scope boundaries. To overcome this we define "automated systems" for the fusion maintenance case in terms of uncertainty management. This is a broad concept, widely applicable to any engineered system, which provides a common starting point:

For a given task, a system is considered automated if during the task execution it can manage uncertainty without human input.

Examples of areas where uncertainty is often encountered in a fusion maintenance context include:

- The position, geometry, or mass of a handled component.
- The duration or sequencing of a series of tasks.
- Health condition of components and equipment.
- Environmental conditions such as temperature, radiation levels, etc.

With this definition, the scope of automation is tied to the scope of the task under question. A robotic system tasked with "grip this object" might be considered automatic if no human was required to, for instance, position the gripper, actuate it and confirm the gripping was successful. To execute the task automatically, the robot would primarily be managing the positional uncertainty of the gripper relative to the object. The task of "sort these items into these boxes, by color" is more complex and has a broader scope. It could feasibly include a series of subtasks, among them "grip this object". For the subtask of assigning colors to each object, the system might have to manage uncertainty such as their surface roughness, lighting conditions, occlusions, etc. If no human intervention were needed during the complete execution, the system would be considered automated for the task in question.

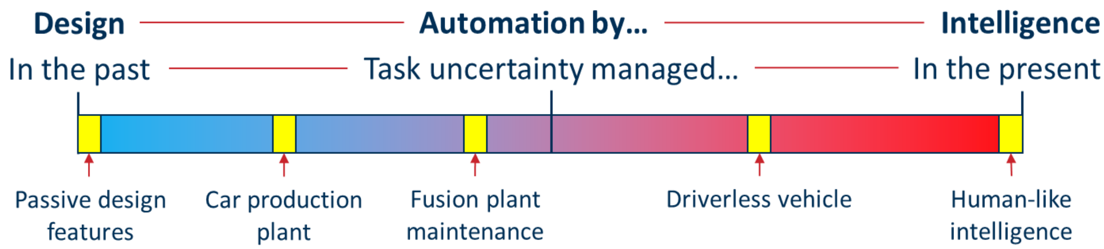

Figure 3. Automation implementation model, shown as a continuous spectrum between Automation by Design and Automation by Intelligence.

It follows that the categorization of a system as automatic is not absolute. A system may be considered automatic when executing one task but require human input for others. This implicit bounding of automation capabilities is important: it avoids overly enthusiastic extrapolation to out-of-scope tasks, or, conversely, an unfair judgement of performance based thereon

IV. Automation Implementation Model

The way in which automation of a task is implemented can have profound impacts on its acceptability for stakeholders. Thus, we introduce a model to describe automation from the perspective of its implementation strategy, stemming from the notion of uncertainty management as a fundamental aspect of automation.

Two broad approaches to automation are considered: Automation by Design (AbD) and Automation by Intelligence (AbI). These are viewed as forming opposite ends of a continuous spectrum. This is illustrated in Fig. 3, which considers qualitatively the point in time when most of the uncertainty associated with a task is managed: whether in the past or in real-time. In general, any task that is automated will have some mix of both approaches.

A. Automation by Design

In AbD the bulk of task uncertainty is managed by humans during the design phase. Critical design parameters are identified with the aid of engineering analysis tools, and design features are then included to reduce or accommodate the associated uncertainty.

An example of extreme AbD is the addition of lead-in features to aid a peg-in-hole insertion task, Fig. 4. Without them, a handling system might not have sufficient positional accuracy to reliably complete the task, so would need the ability to monitor the insertion and make adjustments, nominally provided by a human. By adding the lead-in the designers have not eliminated the positional uncertainty of the peg, but have equipped the system with the ability to manage it without human intervention at the time of task execution, thus enabling automated behavior.

AbD forms the backbone of conventional industrial automation as components, robots and plants are all designed to interact with each other in a way that keeps uncertainty within levels acceptable to the automated systems. Deviations from nominal operation will typically fall beyond the scope of automation, so require human intervention.

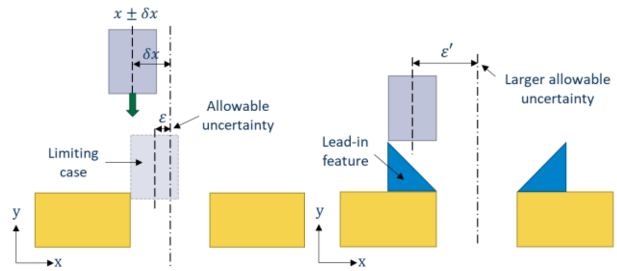

Figure 4. Schematic showing an example of Automation by Design. Left) Problem description showing a peg with horizontal coordinate x±dx, whose insertion will fail if dx exceeds the ε allowable uncertainty. Right) The addition of a lead-in increases the allowable uncertainty that the system can manage without human intervention.

B. Automation by Intelligence

In AbI most of the task uncertainty is managed by the system at the time of task execution. To achieve this, operational information is obtained in real-time using sensors. By interpreting their data, the uncertainty is reduced, bounded by sensor resolution and data processing capability. Then, an appropriate response action is executed following a control law. “Intelligence” here refers to any general control process, not necessarily AI techniques. Hence, an AbI approach to automating the insertion of the peg could be to infer its position with cameras to reduce the positional uncertainty, and use an integral controller to minimize the squared distance error to the hole centerline.

At the extreme end of the spectrum, AbI would denote a system with human-level intelligence and beyond, for instance capable of designing and manufacturing new tooling to accomplish a task. The uncertainty managed at the design phase by humans in this case would be minimum, as the system could independently provide suitable solutions to completely unpredicted scenarios.

AbI is very common in industrial automation, although typically only taking the form of simple sensing and control that supports a wider AbD strategy. For instance, position switches may be used to determine when a component has been correctly loaded into a jig in a production line, but most of the uncertainty of the manufacturing task will have been reduced by humans at an earlier stage though the careful design of the cell layout, the component geometries, the preprogrammed robot motion, etc.

With the increasing adoption of more advanced sensing and control techniques higher degrees of AbI are becoming possible. This is evident in the development of driverless vehicles which, although reliant on existing traffic rules and road infrastructure, are able to manage the uncertainty of unexpected scenarios such as pedestrians crossing the road unpredictably [7].

C. The right balance for fusion maintenance

An essential requirement for applying AbD is a good understanding of the process to be automated, such that all critical uncertainties can be identified and addressed early on. This is usually aided by simplifying tasks as far as possible, making the problem more tractable. In the fusion case, this process is made difficult as existing experimental reactors are not representative of power plant conditions. Experience can be extrapolated from analogous systems, such as industrial manufacturing plants, but the complexity of fusion reactors and the challenges outlined in section II suggest that automated maintenance systems may need to manage higher levels of uncertainty than commonly found in industrial automation. Thus, these systems will require higher levels of intelligence to be able to perform tasks under conditions not fully predicted by the designers.

However, with increasing levels of intelligence also comes the increased ability for the automated system to behave in unanticipated, and potentially negative, ways. This is one of the primary underlying concerns with the deployment of AbI, as validation and verification to guarantee safe behavior remains an active field of research [8]. An excellent analysis of Critical Barriers to Assurance and Regulation (C-BARs) is provided by the Assuring Autonomy International Programme [9]. In contrast, automation of the AbD type is used in safety-critical systems in nuclear fission plants, such as emergency control rod insertion, so that regulatory precedents exist [10]. This suggests that AbD-heavy systems would, today, be more readily acceptable to regulators.

Thus, automated maintenance systems for fusion power plants will need to rely on AbD as far as possible to aid regulatory acceptance, while having sufficient intelligence to tackle the complexity of their environment with only reduced human support. This is illustrated qualitatively by the "Fusion plant maintenance" label in Fig. 3.

To successfully maintain a fusion reactor many different tasks will need to be executed in parallel. However, the number of different robotic systems must be kept within practical bounds to limit costs and manage complexity. It is therefore necessary to rely as far as possible on a reduced number of more capable robotic systems which can each perform multiple different tasks. When referring to systems capable of performing complex series of tasks without human intervention, the term "autonomy" is often used.

V. Definition of Autonomy

In the context of fusion reactor maintenance, we propose the following definition:

An autonomous system can execute a range of different tasks automatically, and adapt the sequencing of those tasks in response to changing circumstances.

With this definition, there is no restriction on the approach used to automate each of the tasks that the autonomous system can execute, whether AbD or AbI. However, being able to change the sequence of which tasks are to be performed, and in what order, does require an awareness of the world state and the ability to reactively manage uncertainty. This implies that autonomous systems have a higher reliance on intelligence. For example, a robotic system may be able to automatically "move to a location X", a task which could be fully preprogrammed or rely on intelligent navigation. But, in order to "avoid collision" with an object unexpectedly lying across its path, the system must rely on intelligence to detect the obstacle and sequence the collision avoidance task.

The choice of the term "automated" rather than "autonomous" maintenance follows this logic. Fusion reactors are likely to require robotic systems which are autonomous to some degree. However, the decisions associated with what maintenance activities are performed, how and when, are very complex and could feasibly impact the plant safety case. Thus, it is not proposed that they are delegated to machines, at the very least for first generation fusion power plants.

VI. Level of Autonomy

The degree to which a complex system of systems is considered autonomous has wide ranging implications for stakeholders, including regulators, designers, investors, site managers, etc. However, quantifying this abstract notion is a known challenge. One of the more pressing critical barriers to the deployment autonomy is creating a common language to allow communities of multidisciplinary experts to discuss issues around control, verification and validation of autonomy.

A. Definitions in other industries

Acknowledging the importance of this, other industries, such as space and connected autonomous vehicles have laid out a common groundwork by producing definitions for levels of autonomy. In these conceptual frameworks levels of autonomy increase with complexity and capabilities for their respective field, but also demand more testing and regulation as the level increases.

The European ECSS Space Segment Operability Standard [11] defines on-board autonomy as: "On-board autonomous functions that provide the space segment with the capability to continue mission operations and to survive critical situations without relying on ground segment intervention". **Error! Reference source not found.** Table 1 outlines the levels of autonomy for space, focusing on the autonomy required to be able to complete certain mission goals. It has been used to define roadmaps for development and help communicate to experiment holders the level of autonomy required by the platform developers.

TABLE I. ECSS LEVELS OF AUTONOMY

Level	Description	Functions	Naming
E1	Mission execution from ground control; limited onboard capability for safety issues.	Real-time control from ground for nominal operations.	Tele-Command
E2	Execution of pre-planned, ground-defined, mission operations on-board.	Capability to store time-based commands in an on-board scheduler.	Pre-planned
E3	Execution of adaptive mission operations on-board	Event-based autonomous operations.	Semi-autonomous
E4	Execution of goal-oriented mission operations on-board	Goal-oriented mission (re-planning).	Goal-Oriented Operation

These notions can be extended to produce different classifications for individual subsystems such as: information interpretation; autonomous guidance, navigation and control; autonomous mission decision making; payload autonomy; and autonomous fault detection, isolation and recovery. This increased granularity allows more nuanced discussion where necessary [11].

Similarly, SAE International have published a taxonomy describing the full range of levels of automation in on-road motor vehicles, as defined in **Error! Reference source not found.**[12]. This has enabled developers, investors, regulators, and governments to discuss the relative advancement of system as it affects deployment, products, and regulations.

TABLE II. SAE AUTOMATION TAXONOMY

Level	Description	Naming
SAE1	The driving mode-specific execution by a driver assistance system of either steering or acceleration/deceleration with the expectation that the human driver performs all remaining aspects of the dynamic driving task.	Driver Assistance
SAE2	The driving mode-specific execution by one or more driver assistance systems of both steering and acceleration/deceleration with the expectation that the human driver performs all remaining aspects of the dynamic driving task.	Partial Automation
SAE3	The driving mode-specific performance by an automated driving system with the expectation that the human driver will respond appropriately to a request to intervene.	Conditional Automation
SAE4	The driving mode-specific performance by an automated driving system, even if a human driver does not respond appropriately to a request.	High-Automation
SAE5	The full-time performance by an automated driving system of all aspects of the dynamic driving task under all roadway and environmental conditions that can be managed by a human driver.	Full Automation

B. *Levels of autonomy for nuclear fusion*

The ECSS definitions for space applications form a solid basis for extrapolation to fusion. However, they are lacking on detail for teleoperation missions, which are not as viable in a space environment due to communication latency for systems located far apart. Teleoperation is central to fusion, and even in maintenance systems with high autonomy it would play a key role, for example in rescue and recovery scenarios.

The SAE definitions for driverless vehicles have a higher reliance on intelligence-based methods. This is partly because the feasibility of handing back control to the human driver in a fault case relies on limited human reaction times. Thus, driverless cars must be able to deal with a wide range of fault scenarios without human assistance if they are to be practically deployed. In the fusion maintenance case reaction times can be significantly longer, so it is expected that handing back control to human operators is more feasible. The SAE definitions, do, however, offer useful definitions for assistive autonomy, which could be key to a progressive phasing in of autonomy, or a staged regulatory acceptance programme.

With these considerations, we propose a set of levels of autonomy focused on fusion maintenance, Table 3.

TABLE III. AUTONOMY LEVELS FOR FUSION REACTOR MAINTENANCE

Level	Description	Operators per robot	Naming
N1	Real-time control enabling human operators to execute tasks through robotic proxy.	5:1	Teleoperation
N2	Real-time control through robotic proxy, with intelligence-based assistance to reduce cognitive loading of human operators.	2:1	Assisted Teleoperation
N3	Automatic execution of pre-planned, validated tasks directly monitored by humans operators.	1:1	Pre-planned automation
N4	Automatic execution of pre-planned, validated tasks without direct monitoring by humans operators.	1:5	Unsupervised automation
N5	Automatic execution of goal-oriented task sequences.	1:50	Full autonomy

The discretization of a continuous concept such as degree of automation is a compromise between descriptive accuracy and ease of communication. Nonetheless, the amount of intelligence required by the autonomous systems can be indicated approximately by the progressive reduction of number of operators required to control each robot.

Although in teleoperation modes higher level intelligence is provided by a human operator, servo-manipulators such as the JET MASCOT rely on a significant amount of sensing and control to respond to human commands and provide adequate feedback. For this reason, they are the first, lowest form of autonomy. Purely mechanical through-wall manipulators commonly found in the nuclear industry would not be considered to have automated features, so would not class as N1.

It should be noted that the proposed scheme does not necessarily reflect the complexity of the control algorithms required at each level, so that some instances of “Assisted teleoperation” may be more technically challenging than some cases of “Unsupervised automation”.

The experience of the JET reactor [13] has shown that a N1 teleoperated system is feasible, and can safely manage the complexity and uncertainty associated with maintenance of a fusion reactor. Any further increase in autonomy will require a programme of testing that can demonstrate safe behavior, if it is to be accepted in a nuclear setting. In the case of smart devices introduced to nuclear environments this acceptability is obtained using tools like statistical testing [14]. However, when it comes to autonomy of complex systems the control space is too large and dynamic for traditional statistical testing. Verification and validation of autonomous systems therefore emerges as one of the key challenges to overcome before automated maintenance of fusion reactors can become a reality.

VII. Conclusions and Further Work

This paper has proposed a conceptual framework for automated maintenance of nuclear fusion power plants. The aim is to provide a common grounding to aid communication of abstract concepts of automation and autonomy with plant stakeholders with disparate professional backgrounds.

Future work will study fusion maintenance use cases which can validate the definitions and models proposed in this work. From here, plant designers can be engaged to determine what levels of autonomy may be required for different classes of maintenance task, and hence identify test cases which can act as first demonstrators of the technical feasibility of automated maintenance. Key questions that emerge from this work for further exploration are: "What types of demonstration and proofs of safe behavior will regulators demand for different classes of automation?" and "What verification and validation schemes can prove autonomy of complex systems of systems?".

References

[1] S. Cowley, "The quest for fusion power", *Nature Physics*, vol. 12, no. 5, pp. 384-386, 2016.

[2] J. Ongena, R. Koch, R. Wolf, and H. Zolm, "Magnetic-confinement fusion", *Nature Physics*, vol. 12, no. 5, pp. 398-410, 2016.

[3] R. Buckingham, Rob, and A. Loving, "Remote-handling challenges in fusion research and beyond." *Nature Physics*, vol. 12, no. 5, pp. 391-393, 2016.

[4] A. Loving, O. Crofts, N. Sykes, D. Iglesias, M. Coleman, J. Thomas, J. Harman, U. Fischer, J. Sanz, M. Siuko, M. Mittwollen, "Pre-conceptual design assessment of DEMO remote maintenance", *Fusion Engineering and Design*, vol. 89, no. 9-10, pp 2246-2250, 2014

[5] Oya, Yasuhisa, et al. "Decontamination of candidate materials for ITER remote handling equipment exposed to tritiated moisture", *Fusion technology*, vol. 39, no.2, pp. 1023-1027, 2001.

[6] S. Nof (Ed.), Handbook of Industrial Robotics (Vol. 1), John Wiley & Sons, 2009

[7] Bandyopadhyay, C. Z. Jie, D. Hsu, M. H. Ang Jr, D. Rus, and E. Frazzoli, "Intention-aware pedestrian avoidance," in: *Experimental Robotics*, J. Desai, G. Dudek, O. Khatib, V. Kumar, Eds., Heidelberg: Springer, 2013, pp. 963–977

[8] M. Fisher, L.A. Dennis, and M. P. Webster, "Verifying autonomous systems", *Communications of the ACM*, vol. 56, no. 9, pp 84-93, 2013.

[9] University of York, *Critical Barriers to Assurance and Regulation (C-BARs)*, https://www.york.ac.uk/assuring-autonomy/research/barriers-to-assurance/ [Accessed 15 Mar. 2019].

[10] International Atomic Energy Agency, "The role of automation and humans in nuclear power plants", No. IAEA-TECDOC-668, 1992.

[11] Gao, Yang, ed., *Contemporary Planetary Robotics: An Approach Toward Autonomous Systems*, John Wiley & Sons, 2016.

[12] SAE International, "Taxonomy and definitions for terms related to on-road motor vehicle automated driving systems", No. J3016, 2013

[13] A. Rolfe, "A perspective on fusion relevant remote handling techniques", *Fusion Engineering and Design*, vol. 82, no. 15-24, pp. 1917-1923, 2007.

[14] T.L, Chu, A. Varuttamaseni, and J.S. Baek. "A Statistical Testing Approach for Quantifying Software Reliability; Application to an Example System", No. BNL-112743-2016, Brookhaven National Laboratory, 2016.