Proceedings of the 33rd European Safety and Reliability Conference (ESREL 2023)
Edited by Mário P. Brito, Terje Aven, Piero Baraldi, Marko Čepin and Enrico Zio
©2023 ESREL2023 Organizers. *Published by* Research Publishing, Singapore.
doi: 10.3850/978-981-18-8071-1_P042-cd

Designing Future Control Environments: A Feasibility Study Approach Developed at IFE FutureLab

Lars Hurlen

Digital Systems, IFE, Norway. E-mail: lars.hurlen@ife.no

Gyrd Skraaning

Digital Systems, IFE, Norway. E-mail: gyrd.skraning@ife.no

In this paper we present lessons learned from 10 years of research into future control environments performed at the IFE FutureLab. The paper focuses on a research-based feasibility study approach that we have developed, based on a perceived need to more effectively perform future-oriented research in this area. The motivation and rationale behind it is discussed, inspired in particular by ideas on “wicked problems” and “design thinking” approaches. Practical insights from utilizing this type of study in two concrete project cases are presented. The ambition of the paper is to contribute to the ongoing discussion on future approaches to safety.

Keywords: Future Control Environments, Control Room design, Human System Interface design, Proof-of-concept stydy, Feasibility study, Human Automation Interaction.

1. Background

Since the mid-eighties, Institute for Energy Technology (IFE) in Norway has worked with complex, safety-critical industries to help improve the way control environments are shaped and used. In response to complex accidents caused by a combination of human, technological and organizational (HTO) factors, a key research focus has been to identify dimensions that influence system performance, and explore effective design improvements.

A key asset in IFEs research capability has been the HAMMLAB – a full scope powerplant control room simulator environment with state-of-the-art experimental facilities, where detailed human performance data are captured and analyzed. This lab has been used to study a wide variety of human performance issues, particularly within the international Halden Project hosted by IFE (Øwre, 2012).

As the pace of digitalization sped up at the turn of the century, there was a growing awareness both of new industry needs and new technological opportunities, motivating novel approaches to how one might shape control environments and organize work. At the same time, these industries were under increasing pressure to operate more efficiently while maintaining safety.

In response, concepts of operations were changing rapidly: Higher degrees of automation, passive safety systems, centralization, multi-unit operation, as well as distributed and remote operations were some of the common approaches. Also, control room technology evolved rapidly – graphical user interfaces were replacing analogue control panels, enabling more compact control environments as well as new ways of visualizing information and interacting with industrial systems.

New questions were asked, and IFE got involved in a wide range of conceptually oriented projects, particularly within transportation, energy production and distribution as well as the health domain. These were exploratory and multidisciplinary R&D projects, focusing on proof-of-concept arguments rather than human performance assessment to validate well-defined solutions. For this kind of early concept development, the research questions are open-ended, while simulators of imagined future systems are typically unavailable.

We realized that the classical experimental methods and techniques that we were accustomed to from HAMMLAB studies simply didn't fit these research questions. The methods were either too rigid, best suited for constrained problems that could be studied in detail – providing quantitative but largely irrelevant results in our context, or they were too speculative – offering relevant insight and ideas but without the design rationale and justification that was needed within the domains that we were operating.

Throughout this paper we will use two concrete case examples to illustrate these challenges as well as the research approaches that we have followed. The first example is a research project on multi-unit operation within the energy production domain, named “MultiKon” (Hurlen, Eitrheim, Rindahl and Hepsø, 2022). In this case Equinor, the leading oil & gas operator on the Norwegian continental shelf, had identified a desirable business case related to operating multiple offshore petroleum facilities from a single, onshore operation center. This was a novel idea with many unanswered questions related to team structure, operator competence, control center design, HMI design, as well as broader organizational factors. The company needed to come up with an effective operational concept that would be accepted internally, and to adequately demonstrate its feasibility to regulatory authorities. Equinor made it clear that they needed applicable concepts that could be realized in practice, not merely a list of human factors considerations related to multi-unit operation.

The second case example comes from the Halden Project (restructured as the OECD NEA Halden Human-Technology-Organization Project in 2021). This is a joint, ongoing research project hosted by IFE since 1958, performing safety-oriented research for the nuclear industry to be used as technical basis for modernization projects and newbuilds in this industry. Members come both from the regulatory, utility and vendor side. A key research area for many years has been human-automation collaboration. Increasing levels of automation present new opportunities but also new concerns in terms of the potential for errors, misunderstandings, and human out-of-the-loop unfamiliarity (see e.g. Wickens, 1992).

Following a number of serious accidents particularly in the aviation domain, significant efforts have been directed into studying human performance in settings where operators are required to collaborate with various types of automatic systems.

A concrete example from this work is a large HAMMLAB experiment that we performed in 2009 to investigate the effect of automation transparency on human performance in next-gen, highly automated power plants. Previous studies had shown a positive effect of making automation goals, activities, and effects more apparent to operators, and we expected this experiment to provide similar results.

However, the futuristic nature of the plant-wide automation that we conceptualized – one that we had first invented, then developed and visualized in a full-scope nuclear process simulator – produced few such effects (see Skraaning & Jamieson, 2019). Quite contrary, operators often struggled to keep up with novel automatic features, and sometimes executed control actions in direct opposition to the activities of automation.

We were unhappy about these results. There seemed to be fundamental weaknesses in the way that the human-automation team was designed. And not only did we feel the need to further address these fundamental weaknesses, but we also needed a more effective research approach to support conceptual development of future control environments. A full-scope simulation environment is conceptually constraining in a lot of ways, and making fundamental changes to it is time-consuming. Thus, we started to search for alternatives.

2. Towards a new approach

Safety-driven research in complex industrial domains commonly separates technical and non-technical (i.e., human and organizational) efforts. In this situation there is an obvious risk of compartmentalization where different disciplines focus exclusively on their part of the problem. This might be an effective approach if the different pieces are later integrated into a well-functioning solution. Unfortunately this is not always the case. At its worse, the parts that seem to work in isolation may be combined into something that has hidden weaknesses. We were motivated to exploit what we saw as an underutilized opportunity: Performing research that addresses conceptual design questions in a holistic manner, taking an integrated HTO perspective in research targeted more towards *forming* hypothesis than testing them.

One reason why our previous methods had disappointing limitations is illustrated in an article by Norman and Verganti (2014), where it is argued that established methods on human-centered design are well suited to incrementally improve upon a concept or idea (climb the hill you´re on), but ill-suited to conceptualize a new and better idea (find a higher hill to climb). In line with what these authors claim, we couldn´t expect the intended users of future control environments to simply tell us what the best solutions would be. In future-oriented projects users´ previous experience and knowledge has limited applicability because of the radically new setting. Also, in industry-driven projects like our MultiKon case example, fulfilling user needs may only be one of many desired characteristics of the operational concept to be developed. As important were other business objectives related to efficiency and safety, such as cost, implementability, maintainability and scalability.

We let ourselves further inspire by the ideas of Horst Rittel on wicked problems. He claims that there are two kinds of problems: One type is well-structured, easily framed, and although it may be characterized as complex from a certain standpoint, it lends itself to established scientific methods of inquiry. Rittel identified such problems as "tame". The other kind is highly complex, unpredictable, changing, is subject to conflicting interests, and do not lend itself to simple procedures, or even easy characterizations. He called them “wicked”. In the words of Stoltermann & Nelson (2014, pp. 16-17).

> *Very few everyday situations of any importance can be described as tame problems. For instance, there is never only one best solution to such problems. There are only solutions that are good or bad. There is no one correct approach or methodology for solving these problems, and it is not possible to formulate one comprehensive and accurate description of a problematic situation from the beginning…By treating a wicked problem as a tame problem, energy and resources are misdirected, resulting in solutions that not only are ineffective, but also can create more difficulty because the approach used is an intervention that is, by necessity, inappropriately conceptualized…If taken seriously, the wicked nature of these types of problems leads to paralysis. This paralysis is most often skirted by the assumption that wicked problems can be simplified and recast as tame problems. This, of course, exacerbates the original wicked problem situation and creates an even greater mess.*

We clearly recognized the challenges highlighted by Stoltermann & Nelson (2014) above. We had experienced first-hand how analytically oriented approaches, or even user-centered design processes, sometimes created as much confusion as clarity in complex projects, often resulting in requests for ever more analysis in a prolonged and largely unfruitful search for exact answers to ill-defined problems. We had also seen how effective concepts and solutions often came about *despite* the methods utilized rather than because of them. As their ineffectiveness proved obvious, we would simply improvise a different way of approaching the issue. But rather than acknowledging this dilemma we would instead rationalize our initial choice, arguing in hindsight that by adjusting these methods based on our now improved perspective they would provide the desired results next time around, thus there was apparently nothing wrong with the methodological approach.

For the development of future operational concepts, it is our experience that the distinction between tame and wicked design problems is seldom introduced, and much less appropriate approaches for addressing these distinctively different problem types. Rather, there seems to be a common notion that answers to the challenges one is facing can be *uncovered* rather than *created.* Or at the very least, that once the problem is properly understood through some sort of analysis, good solutions will be obvious and can be materialized using a set of predefined techniques. This might have something to do with the background of many practitioners in this area. For example, it is our impression that many “human factors engineers” are biased towards *observation* away from imagination. There is the tendency to consider new concepts or ideas as “discoveries” rather than an act of creation. There are of course many relevant insights that can be gained by studying complex human-made systems, from a distance so to speak. But one cannot observe one's way towards the making of novel operational concepts for future technologies. It is like cooking - by following a recipe one can make a delicious dish, but there is no recipe for making recipes.

Also, there are mercantile issues. Buyers of research projects will require a plan, and unless it is organized in a sufficiently convincing manner it is difficult to get potential customers to commit. We had seen how we ourselves sometimes leaned towards familiar and “proven” methods for control room design in the planning phases of future-oriented projects, like the ISO-11064 – even though we knew we would need to largely modify our approach later – simply because we didn´t have a sufficiently mature and presentable alternative.

Thus, although our existing methods didn't fit the conceptual design goals, the absence of a method was not an option either.

3. The IFE FutureLab

We needed a less rigid and proceduralized research approach to support conceptual design for future control environments, that would still be sufficiently solid to inspire confidence and provide reasonable scientific rigor, predictability, and direction, as well as producing useful and applicable results. It was our experience that good solutions to what we now identified as wicked problems were *created* in a deliberate but often seemingly unorderly process consisting of iterative cycles of analysis, creative exploration, and design evaluation. Through this cyclic process the problem itself got increasingly better understood and became ever more approachable.

We took inspiration from the domains of design, architecture, and engineering – disciplines that are used to combining analytic and creative efforts to reach novel responses to problems or new opportunities (innovation).

A common technique is to dedicate physical spaces to host creative work and to exhibit evolving prototypes and mockups. A key idea is that making explicit representations of the problems and solutions we are exploring, facilitate shared understanding and inspire creativity in multidisciplinary teams (Kelley, 2001, pp. 119-147). At IFE, we already had good experiences with labs for dedicated work within a specific area, so our first idea was to establish a "FutureLab" – a development and test environment that would serve as a base for this research. Thus, in 2009 we set such a lab up to be highly flexible and configurable, featuring a rich set of Human Machine Interface technologies, suited for creating early prototypes and mock-ups.

The second initiative was to work towards a methodological framework or approach that would suit the types of projects that we were engaged in. This ambition was fulfilled over time, and it took us quite a few iterations to provide and formulate solutions that felt reusable and sufficiently solid. We have chosen to call this approach a "feasibility study". Feasibility is a term that has a lot in common with the notion of a proof-of-concept, which belongs to the early stages of the much-used technological readiness scale (Mankins, 1995). There are different types of feasibility, such as technical, legal, economical, and operational feasibility. The projects that we have been involved with typically have a focus on operational feasibility, although MultiKon is an example where other types were certainly in play.

The objective of a FutureLab feasibility study is to develop and test concepts, and consists of four main phases: Framing, concept development, a feasibility test, and the generation of results (see Fig 1). For the first two phases we build on ideas from the design domain, particularly on "design thinking". This is a term that is increasingly used within the Human Computer Interaction (HCI) community to describe principles of integrative and user-centered enquiries that include iterative cycles of prototyping and evaluation/reflection (see e.g. Cross, 2023).

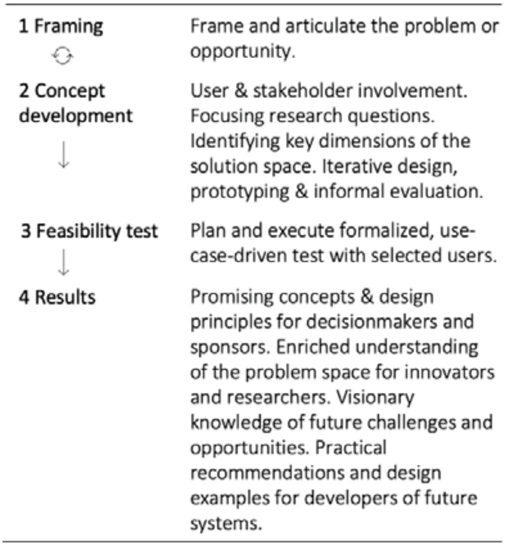

1 Framing ↻	Frame and articulate the problem or opportunity.
2 Concept development ↓	User & stakeholder involvement. Focusing research questions. Identifying key dimensions of the solution space. Iterative design, prototyping & informal evaluation.
3 Feasibility test ↓	Plan and execute formalized, use-case-driven test with selected users.
4 Results	Promising concepts & design principles for decisionmakers and sponsors. Enriched understanding of the problem space for innovators and researchers. Visionary knowledge of future challenges and opportunities. Practical recommendations and design examples for developers of future systems.

Fig. 1: Phases of the IFE FutureLab feasibility study

When broadly applied, design thinking may sometimes rely on overly proceduralized approaches, but an article by Kees Dorst struck a particular chord with us (Dorst, 2011). Here he introduces the idea of a "frame" – a perspective or standpoint from where to tackle complex problems. Many different frames are possible, but they will potentially lead to very different results.

In the MultiKon case for example, following a set of analytically oriented exercises that involved a study of multi-unit operation in other domains and a series of interviews with Equinor team leaders, as well as creative workshops in the FutureLab, our chosen frame centered on the operational similarity between units to be operated from a single control center, as well as their level of HMI integration.

In the concept development phase, we emphasize prototype development; mock-ups, VR-models or other control environment representations that will help the multidisciplinary team to gain a shared understanding of both challenges and opportunities. In both these phases we utilize the physical lab space to gather users and other project stakeholders in a series of workshops.

For example, in the MultiKon project participants came from the owner side, the various technical or organizational disciplines that were involved, the unions, as well as representatives of potential end users. We would use the lab to identify, stage and play out expected critical and typical situations, and use cases that the future control environment needed to support, discussing ideas and issues that this brought up (ref Fig. 2). In our work on human-automation collaboration within the Halden Project, we often took advantage of visiting members and guest researchers, involving them in ongoing feasibility studies as we saw fit. Since the lab was ready to go and easy to set up, relatively little work was required to make the necessary preparations each time.

The objective of the test phase (see Fig 1) is to evaluate the feasibility of the ideas that have been developed and provide a rationale and justification for the recommendations provided by the study. In this phase, prototypes are developed to a point where users can make a reasonable assessment of the operational concepts under exploration, but not necessarily further. From a practical standpoint, the test environments in both of the example projects described in this paper were made by linking several semi-interactive PowerPoint presentations combined with cardboard and paper images to emulate the imagined control environment.

Test participants are recruited based on assumed suitability. From an evaluation perspective this is somewhat controversial because of potential bias issues. Still, we realized that we needed participants that were both willing and able to “play the future game” – to accept the *setting*, if not the ideas. We therefore recruit somewhat selectively among potential end users (that had not participated in the concept development work) and downplay test results from participants that did not inform the feasibility objectives of the study. Some participants might, for example, be fixated on the operation of existing control environments and therefore less capable of grasping the novelty of a proposed concept.

Fig. 2. The IFE FutureLab in use during the MultiKon project: Concept development (top), and feasibility testing (bottom).

In the results phase (see Fig 1), we draw on lessons learned from all previous phases. We summarize findings from the test and discuss promising conceptual directions and success factors. Common outputs of a feasibility study include identified objectives and goals (as linked to overall system performance and/or business objectives), important challenges or risks, as well as conceptual recommendations in an HTO perspective. Depending on the needs of the study, the results can provide an enriched understanding of the problem and solution space, visionary knowledge of future challenges and opportunities, as well as describe key design principles and illustrate practical examples from the concept development work useful for later detailing phases.

For example, within the Halden Project research output includes tested concepts for novel ways of visualizing process information (Hurlen et al, 2015), performing automated procedures and supporting outage teamwork (Hurlen & Le Darz, 2015). Lessons learned from this research led to a feasibility study where we redesigned the whole human-automation system from the ground up in a further effort to shape what we had found to be promising concepts into design principles that would be useful as a starting point for novel concepts of operations for future plants. Strategies for employing agent-oriented automation structure might be the most important safety-related insight gained from this research. The study suggests that his could enable improved diagnostic support and enable a more function-oriented interaction style, both with a potential for improving operator and system performance (Hurlen & Skraaning, 2022).

In the Multikon study results include practical strategies for maintaining cross-unit situation awareness, alarm-handling and avoiding unit confusion. Tested concepts provide a starting point for further detailed design, and key design principles and design dilemmas provide technical basis for company Human System Interface guidelines and technical requirements for multi-unit operations (Hurlen, Eitrheim, Rindahl and Hepsø, 2022).

4. Concluding remarks

One possible critique of the described feasibility study approach is that it does not contain the experimental rigor that research in safety-critical domains often require. The exploratory approach described in the previous section could for some imply a lack of focus, suited only for providing "fluffy" results. We agree that this method will not provide results with the same certainty as e.g. more experimentally oriented research. But it is important to note that neither is it intended to: The objective is to explore the problem or opportunity, gain a deeper understanding of it, and make conceptual recommendations for directing further work – work that might well include more rigorous experimental methods suited for more detailed development phases, but is now better motivated and directed.

Another question that could be raised, is whether this approach might simply recast the wicked problem into a tame one, as critiqued in Section 2. As mentioned, we feel that a structured approach is required for the type of projects that IFE is involved in, the challenge is probably to apply the right amount of rigor. It is our experience that the four phases we outline represent an appropriate level of detail, providing the structural support needed for planning and coordinating work, while at the same time being sufficiently flexible, allowing it to be tailored to the specific needs of each project and its participants.

Another issue concerns prototype and test environment fidelity. Our ambition in the feasibility approach has typically been to provide participants with a rich context, first to support exploration of concepts and later to test their feasibility. We have found it a success factor to develop fairly detailed mock-ups in order to sufficiently illustrate the issues under exploration, and to ground concepts in the complex industrial reality where they need to survive. This way of "staging" ideas with end users during a conceptual enquiry is a common practice within the design domain (see e.g. Kurniawan, 2004, p. 306). It seems however somewhat underutilized in the control environment domain, probably because of the extra labor and practical complexities involved.

We see this approach as more ambitious than comparable techniques that are more typically used, such as tabletop discussions based on sketches and diagrams, but still far away from the labor-intensiveness of providing a more realistic simulated R&D environment. We believe that the studies we have performed in FutureLab shows how the feasibility study approach can lead to novel and useful results in a reasonably quick and cost-effective way, and participants frequently commented that being outside their normal working environment in our FutureLab setting was conducive to them adopting an open mindset. Feasibility studies, as described here, are however considerably more demanding than the simpler methods mentioned. It is fair to say that we sometimes struggled to find the time and resources to make this extra effort, and that the appropriate level of test environment fidelity was an ever-resurfacing debate.

Another reflection that could be made is that a feasibility study of the kind described in this paper might not be for everyone. The open-endedness of the problem space, the number of disciplines that needs to be involved, the test setting to be created, as well as the seemingly disorganized nature of such an inquiry will typically challenge practitioners in unpredictable ways, requiring considerable agility. It is our experience that not everyone enjoys this kind of problem solving. It follows from our discussions in Section 2 that if a team do not share a basic level of shared understanding of the wicked nature of the problem and an appreciation for the types of results that can be achieved by the feasibility study approach, they will struggle with it. Their methodological disorientation may cause them either to give up or try recasting the problem as tame.

As an example, in the Halden Project several teams typically work on very similar research topics but using different approaches, probably for a good reason. After participating in a number of these projects it is our experience that disagreements about the subject matter – about the insights one has made, interpreting results, identifying promising solutions, etc. – are usually easy to deal with. Differing viewpoints can lead to healthy, productive discussions and deeper insights. But when teams are too divided on methodological issues, adhering to different cultures of inquiry, a productive collaboration is far more difficult to accomplish.

Thus far our usage of the feasibility study approach has been limited to the process control domain and has largely focused on design issues related to future control environments. Future work could inquire into its suitability for other types of complex sociotechnical systems. We find it likely that the described way of "staging" future safety-critical situations through prototypes in a lab setting could be effective also in other settings, either by itself or complimentary to other methods of inquiry. But there might also be characteristics of the particular wicked problem to be addressed that limits its usefulness, such as the number and characteristics of involved stakeholders.

Acknowledgements

We would like to acknowledge the OECD Halden Project for supporting the establishment of the IFE FutureLab, and Equinor for sharing the MultiKon reference case used in this paper.

References

Cross, N. (2023). Design thinking: Understanding how designers think and work. Bloomsbury Publishing.

Dorst K. (2011). The core of design thinking and its application, Design Studies Vol 32 No. 6.

Hurlen L., Le Darz P., (2015). "Improving Outage Control Centre team performance through design approach", Proceedings of the 2015 ESREL conference

Hurlen, L., Skraaning, G., Myers, B., Jamieson, G. A., & Carlsson, H. (2015). The plant panel: Feasibility study of an interactive large screen concept for process monitoring and operation. NPIC & HMIT.

Hurlen L., Skraaning G. (2022). HTOR-009: Innovative Human System Interface concepts for highly automated plants. OECD Halden HTO Project. Request can be submitted to: https://ife.no/en/project/the-halden-reactor-project/

Hurlen, L., Eitrheim, M., Rindahl, G., & Hepsø, V. (2022). Concepts For Operating Multiple Petroleum Facilities From a Single Control Centre.

ISO (2000) Ergonomic Design of Control Centres — Part 1: Principles for the Design of Control Centres (ISO-11064-1).

Kelley, T. (2001). The art of innovation: Lessons in creativity from IDEO, America's leading design firm (Vol. 10). Currency.

Kurniawan, S. (2004). Interaction design: Beyond human–computer interaction by Preece, Sharp and Rogers (2001), ISBN 0471492787. Universal Access in the Information Society, 3, 289-289.

Mankins, J. C. (1995). Technology readiness levels. White Paper, April, 6(1995), 1995.

Nelson, H. G., & Stolterman, E. (2014). The design way: Intentional change in an unpredictable world. MIT press.

Norman, D. A., & Verganti, R. (2014). Incremental and radical innovation: Design research vs. technology and meaning change. Design issues, 30(1), 78-96.

Rittel, H. W., & Webber, M. M. (1973). Dilemmas in a general theory of planning. Policy sciences, 4(2), 155-169.

Skraaning, G. & Jamieson, G. A. (2019). Human Performance Benefits of The Automation Transparency Design Principle: Validation and Variation. Human Factors, 63(3), 379-401, 0018720819887252.

Wickens, C. D. (1992). Engineering psychology and human performance (2nd ed.). New York, NY: Harper Collins Publishers.

Øwre, F. (2010). The history of HAMMLAB: 25 years of simulator based studies. In Simulator-based Human Factors Studies Across 25 Years: The History of the Halden Man-Machine Laboratory (pp. 13-42). London: Springer London.