




Real-world applications of BIM and immersive VR in construction

Downloaded from: <https://research.chalmers.se>, 2024-04-20 10:46 UTC

Citation for the original published paper (version of record):

Johansson, M., Roupé, M. (2024). Real-world applications of BIM and immersive VR in construction. *Automation in Construction*, 158. <http://dx.doi.org/10.1016/j.autcon.2023.105233>

N.B. When citing this work, cite the original published paper.



Real-world applications of BIM and immersive VR in construction

Mikael Johansson^{*}, Mattias Roupé

Chalmers University of Technology, Department of Architecture and Civil Engineering, Division of Construction Management, SE-412 96 Gothenburg, Sweden

ARTICLE INFO

Keywords:

Building information modeling
BIM
Virtual reality
VR
openBIM
IFC
Real-time rendering

ABSTRACT

The integration of immersive Virtual Reality (VR) and Building Information Modeling (BIM) has many applications within the Architecture, Engineering, and Construction (AEC) industries. In this context VR is often highlighted for its ability to convey scale and details, especially when compared to non-immersive visualizations. However, despite being an active area of research, there is currently a lack of real-world studies exploring immersive VR in a construction-oriented context. In addition, there are still technical challenges and barriers for an efficient integration, such as rendering performance and interoperability issues. This paper addresses these issues by investigating the use of immersive, single- and multi-user VR within the openBIM ecosystem. The contribution is two-fold: (a) an in-depth presentation of algorithms and technical details of a multi-user VR application for immersive visualization of large and complex BIMs and (b) an evaluation of this VR system on several real-world construction projects. In all cases the VR visualization has been directly realized from the design teams IFC-models and the multi-user sessions has been performed both co-located as well as fully remote. The results show that multi-user VR improves communication, understanding, and collaboration, and by letting staff with knowledge and experience from construction production review the project in VR, design errors and constructability issues can be identified and resolved before reaching the actual production stage. Moreover, the use of VR is helpful regarding sequencing and planning, and to identify alternative design solutions.

1. Introduction

Immersive virtual reality (VR) and Building Information Modeling (BIM) have emerged as powerful tools for the Architecture, Engineering, and Construction (AEC) industries. When integrated, they allow for applications in a number of different areas, including design review [67], production planning [56], and construction safety [32]. The argument often put forward when compared to non-immersive, desktop visualization, is that immersive VR provides a better understanding of scale and detail and allow people to enter and inspect environments in a similar way as they would do in real life [34,82]. More recently, immersive VR has been extended to support multi-user sessions, where several participants can experience the same model at the same time [25]. For design review sessions and model inspection this has been shown to enhance communication and improve collaboration among participants [36].

However, in practice, much of the current research concerning both single- and multi-user VR within the AEC industries are still around developing various prototypes – often using game engines – that are yet to be tested and evaluated in real construction project [26,77]. Even in

the case of commercial, standalone direct-to-VR applications most studies are not performed in the context of real-world projects [12,31,42]. A few exceptions include elevator machine room planning [78], collaborative 4D-planning [76], end-user design review [69], and MEP design review [84]. Still, in the vast majority of cases, end-user evaluations or field studies are often overlooked – despite the general understanding around the importance of evaluating new technology and software systems with actual users, real-world data, and in the actual work process [17,70]. In fact, previous research has shown that certain behavior cannot even be observed or extracted in a laboratory settings, but instead has to involve field studies in the end-users real-world environment [27]. For instance, the motivational component of using and engaging with the system is largely left out when doing tests in a laboratory setting or with synthetic data [75]. In other words, it is totally different to let users speculate on how to use a system during a test or demonstration compared to observing how they actually will use it in their own real-world situation [47]. Consequently, only by validating the technology in its real setting and with real projects and stakeholders, can it become clear how they actually will use it and what contribution the technology actually gives. We therefore argue that there is limited

^{*} Corresponding author.

E-mail addresses: jomi@chalmers.se (M. Johansson), roupe@chalmers.se (M. Roupé).

research and a gap in knowledge concerning benefits and use-cases of both single- and multi-user VR in a construction-oriented context, such as for design and constructability review, sequencing, and job planning.

As an effect of the former much of the research conducted within both single- and multi-user immersive VR have not been based on the use of complete building models taken from real construction projects, and are therefore not always representative when it comes to technical challenges and barriers of using the technology. In fact, those studies that do build on real projects show that typical dataset are often very large and not easily used for VR without further processing [25,30,41,84,85]. In addition, interoperability issues, such as missing metadata, are a recurring theme when considering the BIM-to-VR transfer, often in the context of using game engines [58,85]. Regardless, complete BIMs taken from real projects pose a challenge, both in terms of interactivity and rendering performance, as well as interoperability and data management through openBIM. Although massive model visualization is an active research topic by itself, surprisingly little attention has been given to the specific case of visualizing large BIMs in real-time – especially in the context of VR and stereo rendering [83].

In this paper we address the current situation by focusing on two main aspects concerning BIM and VR research within the AEC industries: real-world cases, and real-world models. Given this context, we have adopted design science as the overall research framework [39,53]. Design science is fundamentally a problem-solving, technology-driven paradigm that revolves around the creation of an artifact with the purpose of providing utility in a practice-based environment. This research approach mainly contains three activities; Design, Build, and Evaluate. Design and Build is the process of constructing an artifact for a specific purpose, and evaluate is a process of validating how well the artifact performs in its real-world settings. Thus, the contribution of our work is two-fold. First we present the development and technical details of the artifact – *a multi-user VR system specifically designed for immersive visualization of large and complex BIMs*. In addition to algorithms for efficient real-time rendering we have developed techniques and functionality that are required in a construction-oriented context, such as efficient IFC-import, property-driven filtering, and accurate distance measuring. Secondly, we present our findings from evaluating this VR system on multiple occasions in eight (8) real-world construction project. We have explored both single- and multi-user sessions and provide identification of real-world benefits and practical uses cases of the VR technology in a construction-oriented context.

The remainder of the paper is structured as follows: In the next section we review related work followed by an outline of the research approach. Section 4 presents technical details of the developed VR system. In section 5 we describe the evaluation and case projects, and in section 6 the results are presented and discussed. Finally, section 7 and 8 discuss limitations and conclude the paper.

2. Related work

2.1. BIM and VR in design and construction

The use of VR in architectural design and construction goes back already to the early nineties and *The Walkthrough Project*, where users could experience a building model made with 30,000 polygons in stereo using a Head Mounted Display (HMD) with room-size tracking that was updated at 25 frames per second [15,16]. Since then, the term VR has come to include a broad spectrum of display techniques, ranging from non-immersive desktop VR or semi-immersive VR using 3D-glasses and Powerwalls, to fully immersive solutions using CAVEs or HMDs [13,73], and much of the early research around integration of Building Information Modeling (BIM) and VR was mainly done without focus on HMDs [45]. However, the introduction of the Oculus Rift HMD paved the way for a new generation of affordable, high-quality HMDs mainly directed at the consumer market. Together with increased use and adaptation of

Building Information Modeling (BIM) that occurred at this time, a new paradigm of VR within the AEC industries started [52]. With BIM, all of the 3D-data required for visualization was directly available, and with portable and affordable HMDs, like Oculus Rift and HTC Vive, the VR-technology became much more accessible in practice [44].

Typical applications for VR today include construction safety planning and training [2,8,28,32], production planning [29,56,65], virtual showrooms [63,64], as well as design review sessions [67,82,84]. During design review session performed in real-world projects VR has been shown to clarify aspects of the design that is difficult to extract from traditional design documents, such as lack of space for installations and maintenance or physical clashes between components [84]. Similarly, construction workers, MEP fitters, and people that work in service and maintenance has been found to prefer VR when compared to desktop BIM-viewers as it better mimics their work environment [47,82].

A common theme when comparing these applications with their non-immersive counterparts is the overall improved understanding of size, scale, and details. This, in turn, helps with communication and inter-party understanding for applications and tasks that typically involve multiple participants, such as design review [67]. Regarding design review specifically, Umair et al. [79] found better performance (i.e. detecting more design issues) in immersive VR when compared against both paper-based (2D) and monitor-based (3D) design review. Also, this doesn't necessarily require a high level of photorealism (e.g. advanced lighting calculations, GI, reflections). In fact, even with the use of comparably simpler visualization models, there is clear evidence that immersive VR - in comparison to non-immersive VR - enhance spatial perception and presence, features that benefit design review and increase overall productivity [59].

To get a better understanding of VR use from the contractors' perspective, Ozcan-Deniz [58] conducted a comprehensive multiple case study where data was collected from twenty-seven (27) cases from eighteen (18) construction companies in the US. Use cases and benefits were found within design review, project coordination, and planning. Improved review processes allowed changes or corrections to the design to be made in the pre-construction stage, with lower costs and improved schedules as a result. Challenges were mainly found in interoperability issues. Unfortunately, details are sparse regarding VR- and BIM-software, model complexity, and any conversion processes, which makes it difficult to pinpoint both success factors and challenges with regards to interoperability, interaction, and model quality (e.g. geometric level-of-detail, structure of information, number of disciplines).

However, despite all the documented benefits, immersive VR using HMDs has always been a tool mainly for individual immersion. While other variants of VR, like CAVEs and Powerwalls, never could compete in terms of immersion, they have always supported collaboration. As such, it makes sense that the next logical step for immersive VR using HMDs was to support multi-user sessions where several participants could experience that same model at the same time.

2.2. Multi-user VR

In order to improve communication among stakeholders, BIM-based immersive VR was extended with multi-user functionality [25,71]. Although this concept had been explored before [18], this was among the first prototypes that combined BIM and a new generation of HMDs. The prototype was developed in the Unity Game Engine, and although the process from BIM to VR was cumbersome it demonstrated the potential and paved the way for further research and also multi-user implementations in several commercial applications, such as IrisVR Prospect and Fuzor. The authors later extended their work to allow for live updates from Revit to be transferred directly to VR after initial model transfer [26]. This had actually been done before, in the commercial software Enscape, but without the support for multi-user.

Multi-user functionality has since been explored for a number of different applications, including design review [31,69,77,84],

construction planning [76], safety [9,72], and overall it has been shown to enhance communication and improve collaboration among participants. Some research suggests that this can be explained by the combination of the embodiment properties of multi-user VR (i.e. use of virtual avatars) and the immersive, full-scale properties as this contributes to both understanding and presence [1,78]. Also, in comparison to other types of remote meetings, multi-user VR effectively “shields” each user from the real environment and therefore forces them to focus and concentrate on the task without the possibility to do other activities in the background [36].

As with single-user VR there are mainly two types of studies when considering multi-user VR; those that involve the creation of a custom-made prototype, often using a game engine, or those using a dedicated commercial application. Belonging to the category of custom-made prototypes with Unity, Prabhakaran et al. [61] developed and evaluated a BIM-based, multi-user VR-environment for the furniture, fixture and equipment (FFE) design sector. Expert interviews during the evaluation found the system highly useful for FFE stakeholders and revealed that the multi-user functionality was seen as a game-changer for communication in physically spread-out teams. Similarly, Tea et al. [77] conducted an experiment with students to investigate design review performance using a custom-made prototype in Unity with multi-user functionality in VR. Using a BIM with discrepancies intentionally placed, students were able to detect more design errors in a multi-user VR environment when compared to using a more traditional approach with desktop VR navigation and online video conference tool. Still, as in the case of Prabhakaran et al. [62], this study is based on the evaluation of a prototype without tasks that are connected to a real-world project. Nevertheless, use of students for evaluation in this case is quite common. Shi et al. [72] developed a Unity-based multi-user VR system with motion tracking function to simulate hazardous scenarios in order to study how social influence affects construction workers' safety behaviors. The system was tested on 126 students and university staff, and showed that other people's unsafe behavior affected participant's behavior – a sign of presence in VR. Social influence, but in a different context, was also identified when students inspected an architectural design in both single- and multi-user VR [12]. Although the model was exactly the same in both experiments, it was perceived as having higher fidelity in the multi-user version. Also with students, but with the commercial VR-software Prospect, Haahr & Knak [31] investigated multi-user VR design review. They found that multi-user VR improved communication and understanding around design issues, mainly due to the common frame of reference, and when compared to performing design review on 2D or desktop 3D (i.e. Navisworks) immersive VR improves this process. However, geometrical (hard) clash detection was still considered to be much more efficient to do in Navisworks.

So far, it is clear that most research around multi-user VR has been done in experimental settings, often with students. However, there are also a few examples regarding actual use of the technology in a real-world setting. During MEP design review in a real project VR was found to clarify aspects difficult to comprehend from conventional media and multi-user allowed participants to actually follow each other and review the systems collaboratively [84]. Similar benefits were identified by Sateei et al. [69], but for architectural design review in a public school project with building end-user representatives – several previously overlooked critical design issues could be identified and resolved in collaboration. Truong et al. [78] explored multi-user VR for elevator machine room planning in a fully remote, real-world setting, and found that it improved planning accuracy, collaboration, and user-satisfaction and brought significant economic benefits to the business. However, the study also identified interoperability issues and technical limitations between BIM and VR.

To some extent we can summarize that much of the current BIM-and-VR research within the AEC industries are still around developing various prototypes – often using game engines – that are yet to be tested and evaluated in real construction project. In fact, even when

considering standalone direct-to-VR application most studies are not performed in the context of a real-world project. As such, we argue that there still exists a research gap when considering real-world use of immersive VR in a construction-oriented context.

2.3. Technical challenges and barriers

Moving beyond the lack of real-world studies, there are several additional challenges and barriers – including technical – for an efficient integration of BIM and VR. Prabhakaran et al. [62] did a comprehensive review and analysis of challenges facing use of immersive technologies within the AEC industries and identified multiple categories. Similarly, Chen et al. [21], identified several VR-AEC adoption barriers and their corresponding mechanisms. From these studies we found recommendations, challenges, and barriers within three recurring technical areas that we explore in more detail in the following subsections; *Interaction* (e.g. user interface, tools), *Interoperability* (e.g. BIM-to-VR conversion, data transfer, metadata), and *Interactivity* (e.g. frame rate, rendering performance).

2.3.1. Interaction

Although we are not yet at an industry-standard for VR in terms of interaction and navigation techniques, some common features among different software applications have started to appear, such as teleport navigation, miniature models for overview and navigation, and the concept of a tools palette. For game engine-based applications these similarities are perhaps also due to both Unreal and Unity providing their respective VR template and interaction toolkit with ready-made interaction functionality [55]. Still, even in these areas there are questions regarding actual implementations, such as strategies for avatar and model scaling in multi-user sessions, and choices of design for UI-widgets like sliders and buttons [80]. With multi-user capabilities there is also the question on how to represent avatars and the importance of representation for communication and collaboration [36,37]. Heidicker et al. [35] compared full-body avatars with only head and hands avatars in immersive social VR and came to the conclusion that full-body avatars was not needed for co-presence in that setting. In fact, only head and hands were actually found to be better than a full-body avatar with idle animations. This has then been used as argument for simpler avatars in other studies and might also, at least to some degree, point to the Uncanny Valley effect [74]. However, there is far from any consensus around avatar representations and when adding full-body motion-capture, users has instead been found to prefer a photorealistic full-body avatar [60].

Several applications now have functionality for taking measurements or making markups in 3D. However, compared to similar functionality in conventional BIM-viewers, like Solibri and BIMCollab Zoom, measurement in VR is still very basic. In fact, much of the functionality that AEC professionals have learned to expect from desktop or mobile BIM-viewer, are not yet common features in VR system targeting design and constructability review [24,66]. Therefore it actually make sense to look at what functionality and features that are considered required in non-VR BIM viewers and applications, such as exact measurements with snapping, advanced color-coding and filtering color-coding, metadata and non-geometrical BIM-object access, and sectioning tools [68]. From our previous work we can confirm the request from AEC professionals regarding color-coding as well as more advanced measurement capabilities [47]. To satisfy the needs in a construction-oriented context a measurement tool in VR should ideally offer snapping functionality and c/c dimensioning. Systems like Arkio and Gravity Sketch do support this type of interaction in VR already today, but as they should be seen as more of full-fledged 3D modeling tools than a BIM-viewer application for VR, we still consider it an unsolved issue in practice.

2.3.2. Interactivity

A fundamental concept of any real-time rendering system is to

provide interactivity in the form of a sufficiently high frame rate. For modern HMDs this means between 72 and 120 frames per seconds depending on manufacturer and model, and failure to reach this target will eventually lead to dizziness, fatigue, and/or motion sickness for the user [81]. Because of the immersion provided by HMDs, the effects of a fluctuating or to low frame rate are typically worse than for non-immersive visualizations, such as navigating a 3D-model on a regular screen. Whether or not this interactivity demand can be fulfilled depends on the hardware (e.g. CPU and GPU performance), software (e.g. computing efficiency and acceleration techniques), and model complexity (e.g. number of objects and number of triangles). As such, the same BIM on the same hardware might behave totally different depending on VR application – from perfect to unusable [45].

Following the notation used in Akenine-Möller et al. [5], there are mainly three different techniques to improve rendering performance; *simplification* (e.g. mesh decimation, LOD, reduce the area of interest), *culling* (e.g. view-frustum culling that skips objects outside of the view, occlusion culling that skips objects that are hidden by other objects, Fig. 1), or *pipeline optimizations* (e.g. batching objects together to reduce draw calls). MEP-models are often characterized by high geometric complexity and therefore candidates for mesh simplification [41,85]. However, it can also be difficult to find a suitable level of decimation, and not introduce errors when done fully automatically [20]. For typical building models – which naturally exhibit a lot of occlusion – occlusion culling has been found to provide a suitable choice [45]. Within this category of acceleration technique several different algorithms and implementations exist, mainly separated by if they require pre-computation or not. For instance the built-in occlusion culling system in Unity Game Engine requires a pre-computation step. *Online* occlusion culling, on the other hand, require no preprocessing and also supports dynamic objects, but has historically required a rather advanced implementation to function efficiently [3,54]. However, with recent improvements and extensions to graphics APIs like OpenGL, the possibilities to implement efficient occlusion culling fully on the GPU has emerged using so-called GPU-driven rendering [14].

On the one hand, game engines typically have lower performance out-of-the-box (i.e. when simply importing a 3D/BIM-model), but on the other hand, are flexible enough to allow additional acceleration techniques, such as geometry simplification, LOD, or occlusion culling, to be implemented and used. For instance, Du et al. [25] found that activating occlusion culling in their test project increased frames per second with 50%. Still, any such technique will bring additional processing time, which typically also has to be done again for new versions of the building model.

On the contrary, dedicated direct-to-VR-solution, like IrisVR, Fuzor, or Twinmotion, are often optimized to provide a better performance out-of-the-box. However, if those systems fail to deliver enough performance, there is little opportunity other than restricting the area of interest (e.g. by using the section box in Revit) or to have certain modeling strategies that keep object and triangle count low (e.g. rectangular railings instead of cylindrical, less detailed components) [30,84]. For

instance, when trying to view a BIM of the Library of the Canadian Parliament with ~45,000,000 triangles in VR, neither IrisVR nor Twinmotion could handle the complexity and produced extremely low frame rate. Instead, the model had to go through a number of optimizations steps in Rhino and 3ds Max – that reduced the dataset to ~2,000,000 triangles – before it could be transferred to Unity Game Engine and viewed in VR [30]. Similarly, Zaker & Coloma [84] had to use the section box in Revit on a 40,000,000 triangle model in order to guarantee sufficient rendering performance in Fuzor. For some of the direct-to-VR-solutions that build upon a game engine, like Unity Reflect and Twinmotion, it is possible transfer the scene to the respective game engine for further processing and optimization if performance is not sufficient. Regardless, large and complex BIMs are not without problems when considering immersive VR.

One might ask if these examples of large BIMs are representative of typical projects or just super extreme cases. However, 10 million triangles for a single discipline are far from an extreme case today [41]. In fact, when looking at federated BIMs from real-world projects like residential buildings, high-rise office buildings, or hospitals anything from 10 to 100 million triangles is possible, as we will later see in this paper.

2.3.3. Interoperability

Interoperability issues encountered during the BIM-to-VR process typically manifests itself in two different ways; a) either some of the data cannot be transferred [85], or b) there is a cumbersome or additional step to actually transfer all of the data [49]. Sometimes there can also be a combination; the process is complex and inefficient, but still, not all of the data is transferred. In general, missing geometry tend not to be a problem so when referring to (missing) data in this case, the usual meaning is BIM-object metadata, such as properties, or material definitions and textures [85].

However, the content typically produced in BIM authoring software, such as Revit and Tekla, also includes non-physical objects, like spaces, rooms, and levels, and also relations and constraints among objects. As of today, the only standardized way to exchange and transfer digital building models is through the Industry Foundation Classes (IFC) file format. The major difference between IFC and general 3D file formats, such as FBX, is that it fully supports all BIM content, including spatial components, storeys, systems, resources, costs and scheduling data, relations, properties, openings, advanced solid geometry, complex hierarchies, assemblies, classification systems, global positioning, and much more. Together with other open standards, such as BCF for issue management, and mvdXML for model validation, IFC represents a fundamental component in the openBIM ecosystem [43]. Still, only a few VR projects have adopted IFC as the primary exchange protocol, one being the immersive safety training CPVR framework which utilizes IFC-files directly and therefore allow user to quickly upload new scenarios without further processing [9]. In addition, both Hilfert and König [40] and Ojala et al. [57] have presented prototype VR viewers made with Unreal for primary use of IFCs, but they have never been tested in real construction projects.




Fig. 1. Culling explained.

In practice, the actual BIM-to-VR transfer is either file-based (i.e. importing a previously exported file), or addin-based, where the VR session is initiated directly from the BIM authoring software using an application-specific plugin [6]. Addins – which has the *potential* to extract virtually all BIM-data – are more common with direct-to-VR software, like IrisVR or Fuzor, whereas applications and prototypes that are created using game engines often rely on a file-based model transfer using the FBX file format. Although the FBX file format support object properties per se, any BIM metadata is typically not added to the file when exporting from Revit or Navisworks, which is an often mentioned issue in the literature regarding BIM-to-VR interoperability [49,85]. Import speed is another factor, as simply importing a large building model (~90,000,000 triangles) from Revit as an FBX-file into the Unity Game Engine can take more than 4 h [85]. In addition, the game engine route often contains additional step, like mesh processing and simplification, materials and texture processing, or other procedures [48,85]. These issues are perhaps the reason why interoperability and the lack of a smooth and lossless from-BIM-to-VR process often is stated as a problem when using a game engine, even if the exact details are sometimes omitted [6,78]. Although the use of Unity Reflect and Unreal Datasmith have been shown to make this process easier, additional optimization steps such as occlusion baking or mesh simplification may still be needed with a game engine-based approach [48,62,85]. Also, even if object metadata (i.e. BIM-data) is preserved, other types of BIM-objects, such as levels and rooms/spaces are not transferred into the game engine environment. As we will show later, these types of BIM-objects have an important meaning in a construction-oriented context and can therefore be used to add other qualities to VR.

When considering all of the interoperability issues we find in relation to the BIM to VR process, it is somewhat surprising that the IFC file format hasn't received more attention, especially for use in real-world projects. In fact, when design and construction projects require a vendor-neutral BIM exchange format, IFC is the proven, go-to solution. In Scandinavia, IFC is the format used by-contract in BIM-based building projects, and we now even have examples of so-called Total BIM projects where the BIM - in the form of a federated IFC model – is used as the legally binding construction document [24].

However, the reason for why IFC is not used more for VR can perhaps be explained by the observation that interoperability issues and rendering challenges are often seen as two separate barriers [62], when they are in fact much more related and connected. For instance, the need to process and simplify large BIM datasets often comes from the challenge of providing sufficient rendering performance (i.e. frame rate) for complex 3D-datasets [30,85]. As this process typically involves additional software, such as 3ds Max, Blender, and also eventually a game engine, additional export/import “gates” has to be crossed, which can then introduce interoperability issues (i.e. the more processing steps, the more risk for interoperability issues and data loss). Hence, by providing an efficient rendering system, additional optimization may not be needed, which would potentially lead to less interoperability issues. To some extent the relation between interoperability challenges and rendering challenges could also be seen as reversed. As less sophisticated 3D-model formats are used (which itself is a source for missing data and interoperability issues), such as FBX or OBJ, any rendering or graphics system used cannot take advantage of semantic data or logical relation between objects to improve rendering performance, but instead has to treat the input 3D dataset as a generic set of triangular meshes with no further logic than a transform hierarchy. Thus, providing a solution for either one of these issues could potentially help solve the other one.

Our technical contribution to address this problem is conceptually simple – by (A) using the IFC file format as the only data source we inherently minimize all interoperability issues that a “BIM-to-VR” transfer process could encounter, and by (B) utilizing an efficient real-time rendering system we can directly visualize complex BIMs without the need for geometry simplification or other pre-processing steps. In the following sections we will describe the process and the technical details

to achieve this, starting with the overall research approach.

3. Research approach

For the research presented in this paper, design science has been used as the overall research framework. Design Science Research (DSR) constitutes a problem-solving paradigm that revolves around the creation of an artifact that addresses a problem and provides utility in an environment or setting where it is instantiated [38,39]. The result is thus the artifact itself, but also the knowledge around how and why it enhance or improves on the application context in which it is utilized. In practice, this research approach is commonly defined as consisting of three interrelated activities: Design, Build and Evaluate. Together, these activities form the design cycle, which represents an iterative process where design alternatives are generated and evaluated. The design cycle is further supported by the relevance cycle and the rigor cycle, which helps connect the design activities to the environment (i.e. relevance) as well as the scientific knowledge base (i.e. rigor). As already stated, the artifact in this particular case is a VR software system for use in a construction-oriented context.

However, although the designed artifact and the knowledge produced around it is the results of an iterative and chronologically stretched out process, both are presented in a more compressed and final state in this paper, making it appear more like a single design-build-evaluate iteration. In practice, however, this process is represented by three (3) main iterations. To provide a better context and deeper understanding around the full process, the three main design cycle iterations are illustrated in Fig. 2. Within each iteration a number of steps are performed, including *artifact development*, *internal evaluation*, and *external evaluation* (Fig. 2, Top). The result from one iteration is then fed into the next iteration. The artifact has been tested and evaluated during field studies conducted at eight case projects (A-G), which will be described in detail later in this paper (Section 5) together with technical features of the developed VR system (Section 4). As is common in design science projects, design criteria and artifact requirements are often updated based on discoveries from a previous design-evaluation loop together with any changes in the environment and knowledge base. For instance, addition of BCF-functionality and mini-model for navigation is a direct result from the evaluations during the first iteration. However, the most notable addition in this context is perhaps the multi-user functionality which was added due to a major change in the environment – the COVID-19 pandemic. With the pandemic leading to restrictions around physical site visits (remote) multi-user functionality became a fundamental feature in order for the research project to progress.

4. The VR system

As the technical platform in this study we have used and further developed BIMXplorer [11,46]. The final version of this artifact is a high-performance VR-viewer that supports the IFC file format (IFC2x3, IFC4) and creation of federated building models (i.e. typically using a single IFC-file per discipline). IFC-files are imported directly and no other preparation or optimization is needed before entering the project in immersive VR. The developed user interface mainly consists of a tool palette with several different tools as seen in Fig. 3. The following is some examples of tools that exist; Measurement and dimensioning with snapping (also c/c), filtering and color-coding, 3D-markups, object information (BIM-properties) and 3D-labels, section planes, miniature model, multi-user functionality (e.g. gather, goto), and BCF snapshots (presented more in subsections 4.3–4.6).

BIMXplorer is developed in C++ and uses OpenGL as rendering API. For Oculus HMDs (e.g. Quest and Rift) the Oculus API is used, all other HMDs (e.g. HTC Vive) are connected through OpenVR. In the following subsections we describe the development of some of the most important technical components in more detail, in particular the rendering engine




Fig. 2. Design science research process. *Top*: All the steps performed and (re)routes followed during a single design iteration. *Bottom*: Step-by-step, start-to-end illustration of the three main design cycles performed during this research project, including technical features and case projects (A-H) for each iteration.




Fig. 3. VR-screenshots from BIMXplorer showing different tools.

and the algorithm for efficient, GPU-driven occlusion culling.

4.1. IFC-support and data management

IFC-files are imported using the xBIM framework [51]. xBIM extracts all elements, relations, properties, and quantities, and the actual geometry as indexed triangular meshes. The API is designed so that an in-memory representation of the full IFC-file is constructed. That is, users of the API can programmatically navigate and query the contents of the file exactly according to the IFC-specification. The import speed of xBIM is

typically around 75–85% compared to Solibri depending on discipline. For instance, importing only the architectural model (524 MB IFC-file) in project A (see Fig. 11) takes ~120 s.

All of the vertex data is stored in a single, large array. Individual geometries then reference this array with *StartIndex* and *Count*. Normals are compressed and represented using the format *GL_INT_2_10_10_REV* which uses a single 32 bit integer per normal vector. In order to speed-up intersection tests, we maintain a *dynamic AABB tree* similar to the implementation in the Bullet Physics Engine that has very fast insertion, removal, and update-of-nodes operations [23]. During

IFC-import, objects are added to this tree one-by-one. We have debug-functionality to move objects in the scene, and even for the largest models in our tests (A and B, Fig. 11), we can move around objects in real-time. Update speed in this case is an effect of having a GPU-driven renderer (see below) as practically all of the CPU is available for other processes. Typically, the imported scene is saved to an internal binary file format, which is compressed on-the-fly using Zstandard compression [86], which is both faster and has more efficient compression than the zlib and gzip alternatives.

4.2. Rendering engine

High performance even for massive models is provided by GPU-driven indirect rendering and occlusion culling that takes advantage of modern OpenGL features like indirect draw calls and Shader Storage Buffer Object (SSBO) and atomic counters. The occlusion culling algorithm and implementation is based on the work of [14] but extended to also support visibility- and selection state, materials, and semi-transparent geometry. With occlusion culling only potentially visible objects are processed by the GPU. The algorithm is explained in the following subsections.

4.2.1. Indirect rendering

We first describe *indirect rendering*, and then how this is extended with *occlusion culling*. With indirect rendering we do not issue individual draw calls for each geometry object but instead upload a large DRAW INDIRECT BUFFER with all the draw calls and then issue a single MultiDrawIndirect-call which renders everything. For this to work we use a single vertex array (i.e. containing all of the positions for all of the geometry in the model), and each indirect draw call contains *first* and *count* that reference the single vertex array (i.e. first and count being similar to a non-indirect draw call). As illustrated in Fig. 4, we keep all transformations and materials in two separate arrays. To “connect” the indirect draw call with the transformation and the material we use two additional arrays with indices. During the indirect draw call the *base-Instance* value will fetch those indices which we use in the vertex shader to lookup the correct transformation and material. These indices are *instanced attributes* and therefore per-draw (i.e. per-batch), and not per-vertex. Finally we use a third array with a per-draw selection state, and if this value is one (1), we modulate the material with a selection color. This approach of rendering selection inherently forces us to update the selection state array as soon as something in the scene is selected or deselected by the user. Still, this is actually very fast, even for very large scenes. Also, in practice we render indexed triangular meshes, but in the illustrations we use draw arrays to make them easier to read. A typical BIM-object, like a door, may contain two different materials (e.g. one for the frame, and one for the blade), and as such are rendered in two different batches. For a scene with n batches, the sizes of the corresponding indirect draw buffer, transformation and material indices buffers and the selection state buffer is n. The size of the buffer holding the transformations is usually smaller, like (3/4)*n, as all batches that belongs to a single object share the same transformation. The materials buffer is also much smaller than n, as each scene only contains a limited

amount of unique materials. Size in this context refers to the number of individual items in an array, not the data size. For example, a single item in the transformations array is made up of 16 floats.

4.2.2. GPU-driven occlusion culling

With a GPU-driven rendering pipeline, adding occlusion culling essentially boils down to making sure that the DRAW INDIRECT BUFFER only contains visible objects. To achieve this we have to introduce some additional rendering steps and some additional arrays, most importantly one integer array which keeps a record of all the objects and if anyone is visible (1) or not (0). As illustrated in Fig. 5, we start (Step 1) each frame by filling the DRAW INDIRECT BUFFER with objects that was visible last frame. Essentially, we process the visible-objects-array and add an indirect draw call if the record is “1”, i.e. visible. An atomic counter keeps track of the number of draw calls we add to the indirect buffer. Step 2 renders the previously visible objects with a single MultiDrawIndirect-call. Next (Step 3), we render a bounding-box for each object in the scene. These boxes are rendered as points in a single draw call and expanded to boxes in the Geometry Shader (GS). Use of a Geometry Shader is often not advocated due to performance reasons, however, we also evaluated rendering unit-boxes with instancing, but it was not faster. Nevertheless, if a bounding-box passes the depth test it means it’s visible and the Fragment Shader (FS) will add “1” to a (secondary) visible-objects-array at the correct index (all positions in the visible-objects-array is cleared to zero “0” before Step 3, and as this array is bound as an SSBO it is possible to perform scattered writes in the Fragment Shader). For Step 4 a similar process as in Step 1 is used, but with the main difference that only visible objects that has not already been rendered in Step 2 are added to the DRAW INDIRECT BUFFER. Finally, Step 5 renders any newly visible objects with a single MultiDrawIndirect-call. The secondary visible-objects-array (these are swapped from frame-to-frame) now contains the full record of visible objects and is used again in Step 1 the next frame. This marks the end of the opaque, non-transparent pass. Any semi-transparent objects are rendered in a secondary, similar pass, but as objects are not sorted by distance, we use the *weighted average transparency* rendering technique, which is order-independent [10].

4.2.3. GPU-driven frustum culling

In comparison to the more advanced occlusion culling algorithm, frustum culling is a simpler process as it does not depend on inter-object relations. The basic idea is simply to make sure that the DRAW INDIRECT BUFFER is only populated with objects that intersect the view frustum. As illustrated in Fig. 6, this can be performed in three steps. First (Step 1)all the bounding-boxes are rendered as points, but in comparison to the occlusion culling algorithm they are never rasterized, only processed by the vertex shader to see if they intersect the view frustum. If so, the corresponding position in the visible-objects-array (VB1) will be marked with “1”. After that, Step 2 builds the DRAW INDIRECT BUFFER, and then Step 3 renders all the opaque objects (i.e. that intersects the view frustum) in a single draw call. Any semi-transparent objects are rendered in a secondary, similar pass, but using a order-independent rendering technique.




Fig. 4. Indirect rendering.

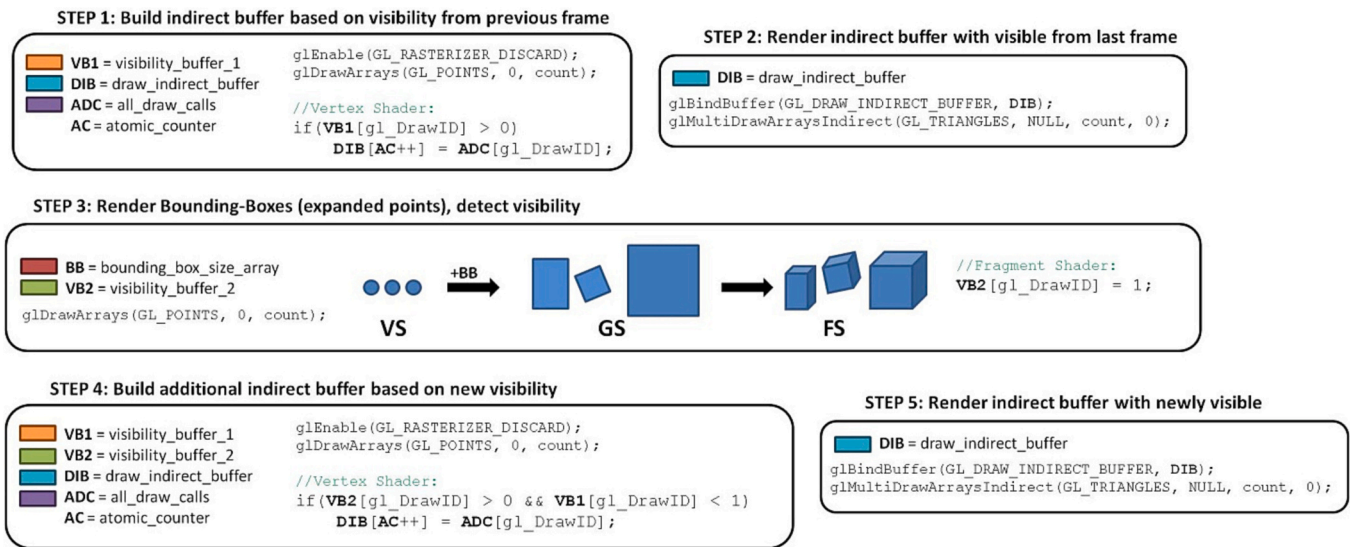


Fig. 5. Indirect rendering with occlusion culling.

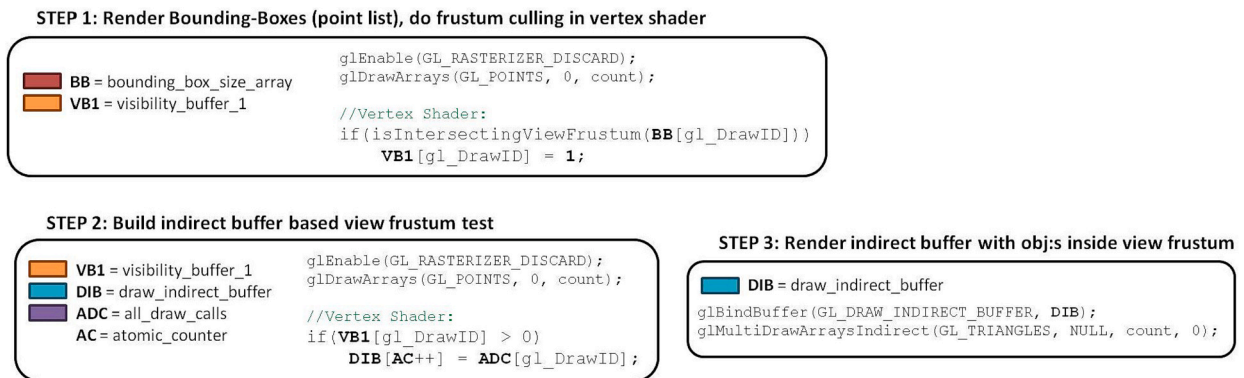


Fig. 6. Indirect rendering with frustum culling.

4.3. Section plane outlines

With a GPU-driven rendering system in place, it makes sense to explore other rendering techniques than those traditionally used. One such area is the rendering of section outlines, a rendering feature that almost has become a must-have attribute in desktop BIM-viewers. Although details are unknown as to how the feature is typically implemented in desktop viewers, it is assumed that the edges found by intersection-testing on the CPU are stored in an array and then rendered as lines. We instead developed a fully GPU-driven approach. As illustrated in Fig. 7 it consists of two passes; A) rendering the scene with the clip plane active, and B) adding the GPU-created outlines. The first pass (A) is integrated with the occlusion culling algorithm and bounding-boxes rendered in Step 3 of that algorithm are discarded if fully outside the clip-plane, which means that the corresponding objects are not added to the list of indirect draw calls. The second pass (B) activates polygon mode “LINE” (i.e. `glPolygonMode(GL_LINE)`) and then places two (2) clip planes near each other, with opposite direction (and slightly offset from the one in the first pass), in order to make a “corridor”. Due to the rasterizing rules with “LINE”, the polygons will not only be clipped, but also produce an edge – the *outline edge*. Also, in practice, the second pass (B) is preceded by a culling pass which collects objects whose bounding boxes intersect the clip plane. As such, only a subset of the objects is actually rendered. However, to what degree this technique is actually better than “CPU-intersection-and-line-rendering” is up for question. We have not yet done a performance test between the two

approaches, and although our technique is fully GPU-driven it is likely that rendering lines is actually faster on typical BIMs today. The benefit of our approach is that no memory has to be allocated for the line geometry, and also it only spends processing power on outlines that are actually visible. On fast GPUs and with very large BIMs this might turn out to be preferred characteristics.

4.4. Filtering and color-coding

It is possible to control visibility and color of objects based on their IFC-properties in a similar way that can be done in Solibri and BIM-Collab Zoom. By defining a set of filters, such as “Color by discipline” or “Color by classification code”, etc., the BIM can then be color-coded according to request from the user. In VR, the user can then select from a list of predefined filters and color-code the model directly in the VR interface. Creating a filter can be seen in Fig. 8. Each line is a rule and when the filter is activated, rules will be processed from top to bottom and each rule will check every object in the scene (i.e. the last rule that affects an object decides its color). Two things help us do this fast: First, during IFC-import we use a perfect hashing algorithm [22] to give each unique string (e.g. PropertySet name, Property name, value) a unique integer value. Most of the time a rule will check if a certain string “is equal to”, and with hashing we only need to compare integers, which is fast. Second, as each SceneObject is a unique, independent object in memory, we can use a parallel loop using OpenMP when processing the filters [4]. This makes activation of a filter very fast. Fig. 9 shows an

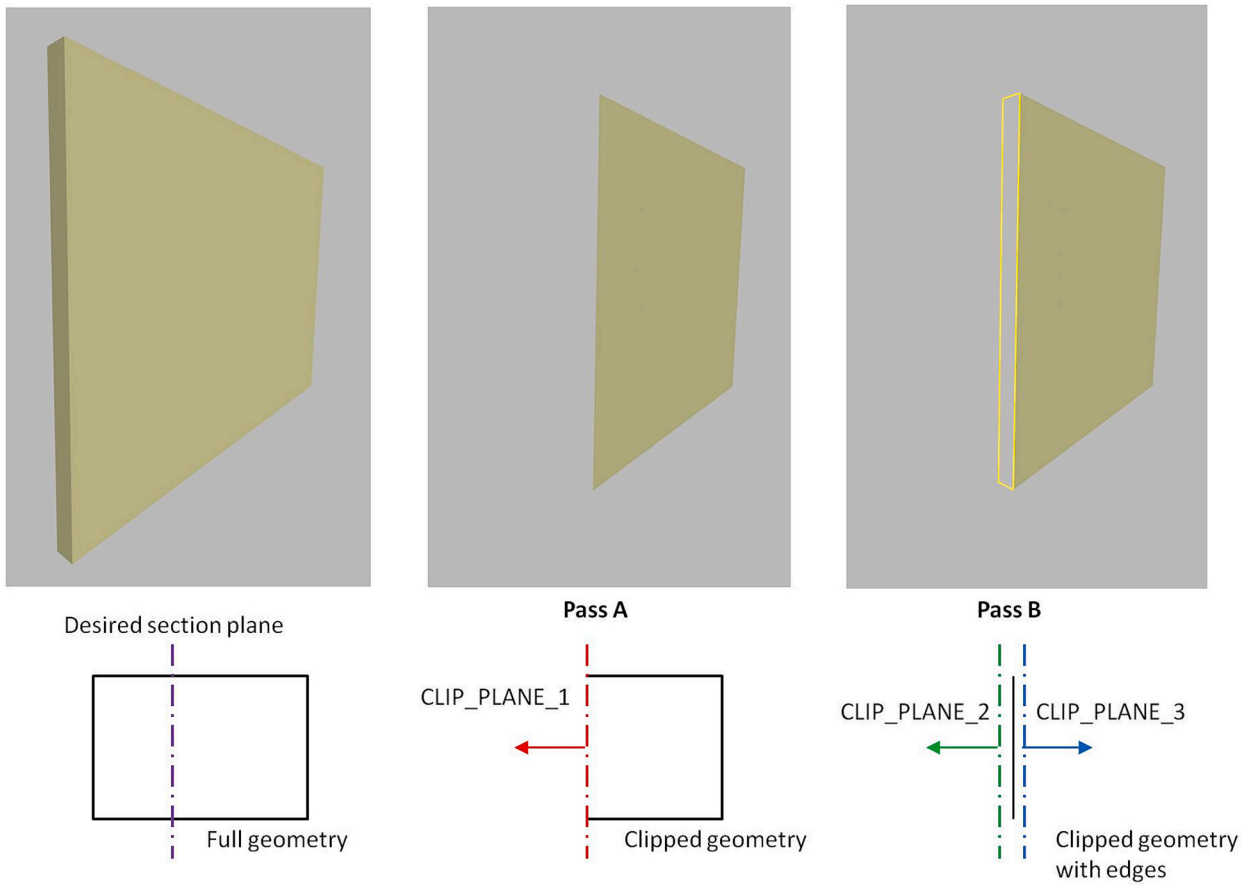


Fig. 7. Rendering section outlines (yellow). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)




Fig. 8. Creation of color-coding and filtering rules.




Fig. 9. Default IFC-colors (left), color-by-system (right). Model from project F.

example of filtering that colors MEP objects by system; the first rule makes all objects transparent, additional rules color by system name. Fig. 8 shows the desktop interface, the VR interface can be seen in Fig. 3 (top row, second from left). For the model in Fig. 9, activating the filter takes less than a second.

4.5. Snapping

The measurement tool supports vertex, plane, edge, and center snapping. Depending on which pair of measurement points that has been selected, a different measurement will be produced. For instance, when measuring between two edges, the system tries to measure the shortest parallel distance. The logic is copied from Solibri, as this is a tool many people on the construction site are used to. The following subsections explain the snapping algorithms:

4.5.1. Edge and vertex snapping

During IFC-import, the geometry is processed in order to find all the free edges per surface and then cache that information. Each object is initially seen as a “triangle soup” and surfaces are formed by comparing vertices and normals among triangles. Sometimes, surface information is already available for the geometry, and the process then becomes considerably faster as only free edges need to be found. As an example, when using Threading Building Blocks (TBB) with a parallel loop implementation [4] this only takes around 2 s for all the geometry in the architectural model in project A, see Fig. 11. During measuring, when the user points at a surface, all the (pre-calculated) free edges for that surface can be obtained from the cache and then see if any edge or vertex is inside the snap radius.

4.5.2. Center snapping

Many of the ducts, piping, and rebars in an IFC file are often represented as extrusions which naturally gives access to the centerline of the geometry. However, depending on the authoring software (e.g. Tekla, Revit, or Archicad) these types of object might also be represented with explicit geometry, such as a triangular mesh. Due to this a more general algorithm to find an object’s centerline was developed – one that works directly on triangular meshes and finds the centerline by shooting a small number of rays inside the circular geometry. As illustrated in Fig. 10, the algorithm finds the starting point and direction of ray #1 (red) by reversing the intersection normal N (at the position where the user is pointing). The endpoint of ray #1 is then found as the intersection with the circular geometry. From the midpoint of ray #1, a new ray is created (green) by expanding in the direction perpendicular to ray #1. Because the geometry is typically faceted (and not a perfect circle), this process is continued for a fixed number of rays in order to find an approximate center, which is then used as the snapping centerline C . However, in practice, this method also first does a Jordan curve test [33] on the reversed as well as non-reversed normal-ray (N) in order to find out if the user is pointing at the geometry from the outside or from the inside. Because of this, it is equally possible to take exact measurements also from holes (e.g. distance from a wall edge to the center of an

opening hole that should be drilled on-site).

4.6. VR UI and multi-user

The in-VR tools palette is implemented with Dear ImGui, but instead of rendering the UI to the screen, it is rendered to a texture every frame. This texture is then applied on a simple quad rendered above the controller.

The multi-user implementation is based on the Photon Realtime SDK. This is a prototype implementation using no other server infrastructure. All clients load the same model, and then call “JoinOrCreateRoom” (Photon API) with an agreed upon meeting ID. Every modification to the shared environment, such as section planes, 3D-markups, or hiding/showing objects is transferred to all clients with the use of Photon events. These events use the “SendReliable” and “Cached Event” functionality in Photon to make sure that even if a client is connecting much later than the other, that client will still receive all the modification events that have already happened when joining. Position and orientation of all the clients, on the other hand, is using “SendUnreliable” because it is regularly updated anyway. However, in either case, no 3D-data is ever sent over the network, just IDs and transformation matrices. The only exception is 3D markups which are represented as a polyline with 3D coordinates. Voice chat is handled through Discord.

5. Method and evaluation

In order to explore the benefits, use cases, and potential of immersive VR in a construction-oriented context, we evaluated the developed VR system in eight (8) real-world projects with a total of 62 participants from the projects. Being part of a design science research projects, this meant that the evaluations were conducted as multiple iterations distributed over time and also in terms of functionality. The study follows a qualitative approach with empirical data collected by means of a questionnaire together with observations, video recordings, and further discussions with the participants. In addition, a dedicated quantitative rendering performance evaluation has been performed. The following subsections describe the case projects and data collection in more detail.

5.1. Case projects

The VR system has been evaluated at eight different projects – seven buildings and one bridge. These projects were selected primarily because they offered opportunity to get real-world practitioners to use and evaluate the developed technology in their own projects, thereby becoming actual stakeholders and not just general “testers” or “evaluators”. In addition, all projects satisfied a number of basic requirements for the evaluation:

- **BIM-based design for all disciplines:** To support interdisciplinary collaboration but also in order for everybody in the project to be committed to “trust” and rely on the federated model.




Fig. 10. Illustration of center-line-finding-algorithm.

- **IFC-models:** To evaluate the openBIM concept, including metadata, properties, and BCF.
- **Model and project complexity:** To be able to stress-test the developed rendering technology and user interface.

Beyond that, there was an overall strategy (selection criteria) to choose projects that also contributed to variation, mainly in terms of type but also scope and size. Together, the eight selected projects (Fig. 11) represent a diverse set suitable for the evaluation and also satisfying the relevance cycle according to DSR.

These projects were all ongoing, but in different stages (e.g. foundation completed, MEP work has just started, etc.). All of them were considered “BIM-projects” with all of the design and design collaboration done in various BIM systems (e.g. Tekla, Revit, Solibri) and with IFC (IFC2x3) as the collaboration format. One of the projects (G) is a “Total BIM” project where the BIM is the legally binding construction document and no traditional 2D drawings are used [24]. The other projects, however, still hold a combination of (digital) 2D drawings and BIM as the construction documents. All projects feature some degree of parallel work between design and construction, however, *main focus in the study has been from the perspective of construction* and not the architectural design or the client or building end-users perspective. In Table 1, model statistics for the different BIMs is presented. As can be seen, complete BIMs from these types of project tend to become very large and therefore challenging to render in VR in real-time. However, this has not presented itself as a problem using our GPU-driven rendering techniques. All the discipline sub-models in all of the eight projects has been imported directly as IFC-files without any need for further processing or optimization. No issues regarding frame rate or lag has been noted.

5.2. Data collection

In total 62 participants evaluated the VR technology in the construction-oriented context of their current project. Several participants used the VR technology at multiple occasions during their current project. In the majority of cases, we (i.e. the authors) have physically visited the site office in each project to evaluate the VR technology. In some cases, a dedicated “VR room” has already been present at the office (Projects A, B, G), and in other cases we have brought and set up one or more portable VR sets consisting of gaming laptops with NVIDIA GTX 1080 GPU and different HMDs. During the study, HTC Vive (Pro), Oculus Rift S, Oculus Quest 1/2, and HP Reverb G1/G2 have been used. Vive has external tracking sensors suitable for a fixed installation whereas Oculus and HPs HMDs have inside-out tracking that only takes 1–2 min. to setup. Due to the COVID-19 pandemic, three of the evaluation occasions had to take place fully remote and then a supervisor or a BIM/VDC manager at the respective project has set up the VR system. In either case, a brief introduction to the navigation interface has been provided and after that the first participants has often instructed others.

The primary focus during the evaluations has been slightly different

from project to project. In some cases, it has been a more general review of the project as a whole, while in other cases there has been a clear focus on a specific part, for instance a specific floor or a certain installation space. However, important to note is that all of the respondents have been actively involved in each respective project. Projects A, D, F, and H have been visited several times. Still, all of the functionality in the VR system has not been available in all of the evaluation session. For instance, multi-user and BCF-functionality was only available in projects A, D, E, G, and H. However, multiple VR sets have been available during all evaluations sessions. During the evaluation session the researchers asked and noted the participants gender and current work discipline/profession and main trade field connected to the construction project, see Table 2. In addition, data has been collected by means of observations, recordings, as well as open questions and discussions with the participant during the evaluation session.

To gather empirical data a questionnaire was used containing both open and closed questions. The questionnaire started with collecting background data about the respondents’ job title and role, age, gender, education, experience in the field. The next four background questions were related to knowledge and use of BIM and what type of BIM software’s they use. The next sub-set of (nine) questions was related to knowledge and use of VR and how they experienced the presented VR system connected to their current work tasks i.e., information they require to perform their work, and how they experienced different aspects of the VR system and its functionality (e.g. measuring tool, multi-user). In total the questionnaire contained 18 questions and the participants conducted the questionnaire after they have used the VR system on their current project. In total 34 participants, (four women), completed the questionnaire, see Table 3.

The dedicated rendering performance evaluation was performed on two different systems: (a) a gaming laptop with Windows 10 equipped with a NVIDIA GTX1080 GPU, Intel i7 CPU, 16GB RAM, and an Oculus Rift S HMD, and (b) a gaming laptop with Windows 11 equipped with a NVIDIA RTX3070Ti GPU, AMD Ryzen 9 CPU, 16GB RAM, and an Oculus Quest 2 HMD connected with the Oculus Link Cable. During this evaluation fly-navigation was used to follow a pre-defined path in two of the buildings – first with occlusion culling activated, and then with only view frustum culling activated – while fetching and recording frame times from the Oculus API. Several round were performed in each mode to see that the performance numbers collected were consistent (i.e. to rule out any effect from temporary Windows background tasks and updates). In order to further analyze and understand the performance results, debug functionality was implemented to be able to “freeze” the culling results and display the view frustum while navigating to other viewpoints (i.e. as illustrated in Fig. 1).

6. Results and discussion

In total 62 persons (seven woman) evaluated the VR technology, many of them at multiple occasions, see Table 2. As already stated they




Fig. 11. Real-world case projects used in this study.

Table 1
Case projects and model statistics.

Project	A	B	C	D	E	F	G	H
Type	Office	Office/ hotel	Facility	Office/ hotel	Bridge	School	Office	Residential building
# of objects	798,059	761,120	227,248	551,654	21,821	158,849	120,585	208,685
# of triangles	111,339,263	85,643,233	28,873,553	48,375,635	2,172,125	3,493,359	10,267,139	29,621,584
Multi-user	X	-	-	X	X	-	X	X

Table 2
No. participants per project (A-H) and design cycle (DC), organized by discipline/field.

Main field/Discipline	A (DC1/DC2/DC3)	B (DC1)	C (DC1)	D (DC1/DC2/DC3)	E (DC3)	F (DC2)	G (DC3)	H (DC3)
MANAGER								
Construction supervisor	3/2/1	1	2	-/-/1	-	2	-	12
Design and Project mngr.	1/-/-	2	-	-/1/1	-	1	2	5
CONSTRUCTION WORKER								
MEP	-	-	2	-/-/3	-	-	-	2
Carpenter	-/1/6	-	-	-	-	1	-	1
Rebar/concrete	-/7/-	-	-	-	-	-	-	1
SPECIALIST/ENGINEER								
MEP/Structural designer	-	-	1	-/-/3	-	2	-	8
VDC/BIM coordinator	1/1/1	1	-	-/2/2	-	1	.	2
Construction engineer	1/1/-	1	-	-/1/-	2	1	-	2

Table 3
Demographics of the participants that completed the questionnaire.

Job title	Total No. participants	Age (Years)	Industry experience in the field (Years)
Construction supervisors/ Foremen	11	Mean = 33,9 (SD = 13,5)	Mean = 11.1 (SD = 7.3)
Design and Project managers	7	Mean = 34,2 (SD = 6.8)	Mean = 11.4 (SD = 7.4)
MEP project leader and engineers	5	Mean = 33.4 (SD = 6.2)	Mean = 11.2 (SD = 7.0)
Carpenter	4	Mean = 33.7 (SD = 5.8)	Mean = 10.5 (SD = 6.2)
Structural engineer	2	33 and 38	14 and 16
VDC-specialist and coordinator	1	35	2
Construction engineer	1	35	16
Rebar/concrete worker	1	22	5
Technical manager	1	39	16
Logistics manager	1	32	6

all had a formal role in any of the projects (e.g. supervisor, MEP fitter, carpenter, etc.) and can therefore be considered actual stakeholders. About half of them, 34 persons, also completed the questionnaire, see Table 3. In the following subsections we present and discuss the results that were captured together via the questionnaire, observations, recordings, open questions, and discussions with the test participants. However, we first present and discuss the results from the rendering performance evaluation.

6.1. Rendering performance

The refresh rate of Rift S and Quest 2 is 80 Hz and 72 Hz, respectively. In practice, this corresponds to a budget of 12.5 ms and 13.9 ms that an application must be able to satisfy in order for everything to be perceived as smooth. The difference between 80 Hz and 72 Hz is not easily noticeable, but an important thing to understand with display technology is that once either system fails to “hold” their target, the refresh rate will automatically be reduced to half (i.e. 40 and 36 Hz, respectively), which becomes a significant drop. However, there are also advanced helper mechanisms, such as Motion Smoothing or

Asynchronous Space-/TimeWarp (ASW/ATW) that allow these performance targets to be relaxed during short periods of time, by reprojecting rendered images as to “simulate” higher refresh rate. Still, failure to meet these targets over longer periods of time (i.e. several seconds) will lower the perceived fidelity and significantly increase the risk of motion sickness among users [7].

To evaluate the rendering efficiency offered by our developed system and to see the benefit and effects of GPU-driven occlusion culling we primarily used project A, simply because it was the most complex in terms of both number of triangles and number of objects. A walkthrough was performed at the 8th floor with and without occlusion culling enabled as to compare it to only using view frustum culling. The path as well as performance graphs showing rendering time is illustrated in Fig. 12. The path is represented with a red markup line, starting at the star symbol in the lower left corner and going in the direction of the arrows. Note that the section plane is for illustrative purposes and in the actual tests no section plane was active. The vertical axis in each chart shows rendering time in milliseconds (lower is better) and the horizontal axis is the number of frames produced during the navigation. The middle chart shows the performance for the GTX1080 system and the right chart shows the RTX3070 system. In both charts red (top curve) shows the render times when only frustum culling is used, and the blue curves show render times when occlusion culling is active.

As can be seen in Fig. 12, occlusion culling significantly improves the rendering performance, sometimes as much as 7×. Nevertheless, perhaps more important is that view frustum culling alone is not sufficient to provide enough performance on either system. In fact, even on the faster GPU, during several locations along the path the frame time is above 35 ms which corresponds to ~29 Hz. However, due to vsync the effective refresh rate might actually become even lower. Even with motion smoothing or different warping techniques this is close to unacceptable as far as performance goes. In comparison, occlusion culling allows the frame times to be held below the threshold on both systems. The only exception is a very short sequence on the GTX1080 system when the frame times peak at 13 ms. Still, during those brief moments, the ASW/ATW kicks in and the effect for a user is not really noticeable. On the more powerful RTX3070-system the rendering time is consistently below 10 ms and therefore has surprisingly much room left before performance becomes an issue. Furthermore, as seen on both systems, the occlusion culling mode is also offering rather consistent frame times, without the strong peaks and fluctuations seen in the frustum culling




Fig. 12. Navigation path (left), and rendering time for occlusion culling (blue) and frustum culling (red) in project A for the GTX1080 GPU (middle) and RTX3070 GPU (right). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

only alternative.

To better understand the reasons for the huge differences in performance between using only frustum culling and occlusion culling, Fig. 13 shows the number of objects actually rendered in each mode at a certain point along the walkthrough path. On the left side the actual viewpoint is seen (top) together with the location along the path and view-cone (bottom). On the right side the view-cone is shown together with those objects that are rendered when only frustum culling is enabled. By only using view frustum culling the system has no knowledge about inter-object occlusion and therefore has to render every object that is inside the field of view, as illustrated by the view-cone. In the middle we instead show only those objects which are identified as *potentially visible* according to the occlusion culling algorithm and therefore the only ones that are rendered. Basically, this comparison shows the simple reason for why performance is much better with occlusion culling: *significantly less number of objects are rendered, and therefore the performance increases.*

However, although the difference is not always this strong, frustum culling tend to have much more variance in the general case. Fig. 14 shows the results from a similar walkthrough from project H. As illustrated by the red markup line the path follows the shared balcony access, before entering an apartment and finally ends up at the private balcony looking across the courtyard (left). The maximum difference in frame time between occlusion culling and frustum culling is $\sim 2\times$, but while OC is having a fairly consistent performance around 8 ms, FC shows much more fluctuating characteristics ranging from ~ 5 to ~ 18 ms. Due to OC having a slight overhead cost, FC can actually be faster in certain viewpoints where there are few objects within the field of view. Still, such viewpoints are less likely to stress the performance and in the general case utilizing OC is a net win. The two rightmost screenshots show rendered objects for FC and OC at the end of the walkthrough and also here it becomes obvious why performance is better – *fewer objects are rendered when utilizing OC.* Nevertheless, it's important to note that none of these cases are a worst-case scenario for view frustum culling, and it would be easy to find viewpoints – such as when the mini-model is

fully visible – where the difference compared to occlusion culling would be much greater. Still, we wanted to give a few representative “normal” cases.

6.2. Navigation, overview, and dimensioning

Overview and orientation of AEC projects in VR is often highlighted as an issue in the literature [50]. The purpose of the mini-model is mainly to allow participant to get an overview of the project and ability to “jump into” a specific location in the building. The concept is not new and has almost become a de-facto standard in recent years [42,69]. Still, for very tall buildings it can be difficult to pinpoint a specific floor from the mini-model perspective, which was expressed at the first evaluation at Project A. Especially construction workers express that a lot of the discussions and planning is in relation to the different floors (e.g. “next week the ventilation subcontractor will start at the seventh floor.”). To handle this situation all the floors (i.e. IfcBuildingStorey) are directly extracted from the IFC-files, and exposed as a drop down list in the tools palette. Selecting a specific floor will then section the building 1.2 m above the floor with the purpose being that it would be easy to jump into a specific location (e.g. Room) at a specific floor. For the taller buildings this was seen as a huge improvement compared to guessing or visually counting floors. This is a good example on how non-geometrical, spatial objects in a BIM – that cannot be transferred in general 3D-file formats – add value in a visualization session. However, a bit surprisingly, this also turned out to be a powerful “2D drawing mode”, allowing participants to discuss and review larger areas, such as the main MEP “corridors”, in a similar way that is typically done on a blueprint. Fig. 15 illustrates such an example from Project D, which was performed fully remote.

Another aspect seen as an issue in previous research is the ability to easily take accurate measurements from the BIM in VR [47]. In fact, even in the context of non-VR BIM-viewers this is still considered one of the main obstacles for going fully “drawingless” on the construction site [24]. Most of the time, it is sufficient to use the perpendicular




Fig. 13. Illustration of rendered objects with frustum culling (right) or occlusion culling (middle) for the specific viewpoint (left) in project A.




Fig. 14. Render times (middle) and rendered objects with frustum or occlusion culling (rightmost two) for the specific path and viewpoint (left) in project H.




Fig. 15. Mini-model used for navigation and overview (From: remote multi-user meeting in Project D).

measurement (which shoots a dimension ray perpendicular to the selected surface), which is a quick way to get the height to the ceiling or width of a corridor. However, once slightly more complicated dimensions are needed, such as shortest distance between two edges, snapping becomes very important in order for it to be user-friendly. For the MEP-workers, the center snapping was actually seen as something completely novel, even in a non-VR context. Many of them were familiar with Solibri or BIMCollab Zoom, which does have edge, vertex, and plane snapping, but not the ability to use center snap. However, as they typically start by placing the hangers, they need the distance from a wall or reference plane to the center of the pipe or duct, something that they now could easily extract, as seen in Fig. 16.

6.3. Information and understanding

Previous studies have highlighted how directly and intuitively people seem to perceive various situations in a VR environment [82,84]. This is strongly confirmed by our data. Above all, it is the fact that the model is viewed on a 1:1 scale which is highlighted and several comments are linked to how size, spaces, and details, are given a completely different understanding and feeling compared to when the model is (re) viewed on a regular screen. Furthermore, it becomes clear from the

observations that review in VR is similar to how you look or inspect something in the real world. In several cases, the VR controllers were used to represent different tools to see if there was enough space in the model to perform a certain set of work steps. On several occasions, users also asked for the ability to replace the controller models with that of tools, such as a screw driver or wrench, which should be considered in future studies. In general, size-related issues, such as height of railings or dimensions of openings, are primarily discovered by viewing the model in a 1:1 scale and that the users instinctively relate to their own, egocentric perspective (e.g.: “that feels like it’s a little low” [Supervisor] or “that feels a bit tight” [Construction worker (MEP)]) before using the measurement tool. To better illustrate the respondents’ thoughts, we give some representative examples of what they answered to the question “What information do you find (in VR) and want to use?”. Also here it becomes clear that it is mainly the visual and geometric representation that conveys a large part of the information they seek and find. In fact, almost 70% of the answers to this question relate to the visual or geometric representation of the model.

“A lot about planning work steps” [Supervisor]
 “Mainly geometries and clashes between models. It’s like Solibri on steroids.” [Structural engineer]




Fig. 16. Sectioning, measurement, and dimensioning created directly in VR (From: Project B).

“Complex solutions that becomes clearer in a larger 3D model. Better understanding of how everything is connected.” [Design manager]

“A VR model creates a better understanding than a normal, traditional BIM models” [Construction engineer]

“Dimensions, object information, details and sections” [Construction engineer]

“Mainly using VR to see what space is available to be able to plan and build in the right order” [Carpenter]

“Clashes between installations, issues around accessibility” [Carpenter]

“The same object information as in BIM” [VDC coordinator]

6.4. Multi-user

The multi-user functionality was evaluated on five (5) occasions, where three (3) of them was performed fully remote. Interestingly, out of all the questions in the questionnaire, the one about added value provided by multi-user functionality was rated highest and most consistent. From the observations, the main difference and benefit compared to single-user VR (with multiple participants) is that with multi-user it is possible for all participants to share a common frame of reference and perspective at the same time. Together with the possibility to also point and use markups, etc. to communicate both verbally and “(meta-)physically”, it is clear that this leads to a better understanding among participants. In the case with multiple users and single HMD, it is possible for the VR-user to verbally explain and show and refer to locations in VR (i.e. as the other participants can look at what is happening on the screen). However, people “outside” can only verbally communicate with the person using the HMD. Typical - and recurring - examples that led to some frustration during the sessions with single-HMD-multiple-participants was when non-VR participants addressed the VR participant while pointing at locations and referring to something on the screen (which, obviously, was not possible to see for the VR-user). Also, even if non-VR participants can see the same things as the VR-user, the scale and perspective is not at all the same, which has led to several confusions and misunderstandings during these sessions when only single-user VR was used. When considering collaboration and understanding among participants, multi-user has a clear advantage.

In one of the cases, it concerned a job planning session with a total of eight participants who took turns to use three different VR headsets that were all available in the same “VR room” at the site office. An additional screen was connected to each headset so that those standing outside could see what the different people were doing in the VR model, see Fig. 17. This session had also been preceded by previous constructability reviews and job planning sessions in VR, however on other parts of the project and without multi-user functionality.

The focus in this case was framework supplement on a single floor and the responsible supervisor started with a review (in VR) regarding how progress was planned on this floor and in what order the different disciplines would go and which moments were considered the most challenging. This then turned into more discussions about specific places and other participants took turns to be inside the model at the same time. In order to exemplify the type of discussion that was going on

between the participants, below is an excerpt from the recorded conversation during the meeting (#1 = supervisor, #2-3 = carpenters):

#1: Here we will have to put a screen on the outside because we will never get access here with the drywalling. I see that now...

#2: Yeah, it's really tight there...

#1: [takes measurements in the model] ...there's not even enough space for a screwdriver so that's the way we'll have to do it...

#1: ...here we also need to build screens first, because this is insanely tight. Here we have to go first and THEN the installers. Just have to make sure they don't sneak start here...

#2: Ohh, yeah, it's tight all the way here!

#1: Yeah, it's very tight.

#3: There's so much installations here...

#3: What if we could go in with a rail before they even cast...? Nah, the ceiling is too high here...

To what extent the use of VR – and multi-user VR – in job planning generally improves the end result is difficult to give a definitive answer to, as it is possible that a “traditional” session using 2D drawings or BIM would have worked just as well. However, many discussions during the meeting are about size and space, and in that aspect we know that VR contributes to a better understanding. From the questionnaire, it is also a very positive rating that is given regarding *understanding of the project, details, as well as understanding the work of other professional groups and disciplines*. Furthermore, there are several examples in this particular project of design and constructability issues first discovered in VR (discussed more in subsection 6.7). Together, all these factors point to added value in using VR during job planning and design and constructability review. What can be said with certainty, however, is that the possibility of being several people in the same VR model is considered a clear added value compared to only being able to be one person at a time. Or, as the supervisor (#1 in above conversation) expressed it:

“I thought this was awesome, that we could be several people in the model at the same time!”

6.5. Remote multi-user

Three of the projects took advantage of the possibility to perform multi-user meetings fully remote. The first was a bridge project (E) which was reviewed among development and construction engineers, and although this was mainly considered to be a technology test, some minor issues and clashes regarding the reinforcement was discovered. These issues could then be sent back to the design team as a BCF-file, see Fig. 18. However, being a bridge project it allowed the VR technology to be evaluated in an infrastructure-context, both in terms of functionality but also in terms of what types of models and information content that can be expected. Notable observations here is that teleport navigation have less usage when doing constructability review for a bridge and also that sectioning-by-floor doesn't make sense in this context. Instead, it is expected that sectioning along the road- or alignment curve would be the corresponding use for infrastructure project with IFC4x3 in the future.




Fig. 17. Job planning with several users simultaneously in the VR model (From: Project A).



Fig. 18. BCF of design issue created in VR and sent back to the design team (From: Project E).

The second occasion with multi-user was an informative meeting around the production sequencing in project G. The meeting was planned to start with testing and evaluate the possibilities of using VR as a tool during the project. However, as the VR multi-user worked so well the originally planned meeting (i.e. Zoom meeting) didn't happen and instead the sequencing was discussed only in VR. As the federated model was very well structured, with properties in each object indicating the sub-contract (e.g. ContractID), it was possible to use the filtering functionality to color-code and turn them on/off in VR. This essentially allowed for a simplified 4D model, albeit without animations, see Fig. 19. During this meeting one of the HMDs was an Oculus Quest 2 that was connected to the computer using Oculus AirLink, which makes it untethered. Although no detail performance analysis was made it was observed that it work surprisingly well. There were occasional "lag spikes" that made the user aware that the device was not connected by cable (i.e. as these spikes never occur otherwise), but other than that no major differences, which makes it interesting to investigate further.

In project D, design and construction were ongoing (in parallel) and the VR-session involved two MEP-designers connecting from their office, two MEP- and VDC-coordinators at the site office, and two researchers. In addition to this, one VDC-specialist connected to the meeting in "spectator" mode, i.e. without a HMD. The primary focus was MEP-coordination at a single floor where ordinary clash detection and coordination had just been completed, but during the meeting several other parts of the project, such as the main installation room, were also reviewed, see Fig. 20. Below is an excerpt that shows the dynamics that arise during the meeting, despite everyone being in different physical locations (#1–2 = MEP designers, #3 = MEP coordinator, #4 = VDC coordinator):

- #1: So, then we can sneak into the main installation room right here. It's probably a bit interesting...
- #2: Damn, it's going to be a high ceiling here...
- #1: This was insanely cool to fly around in here!
- #2: Crazy big! In the model, it doesn't look like that otherwise. I mean, when you look at a normal screen...
- #3: I found here, here it seems to be a big issue! I can gather you here...
- #3: It feels like a lot of things are colliding here. If you look down here...
- #4: Yeah, we haven't really done any [design] collaboration in this area. Still a couple of months away...

However, on the floor that was the primary focus, no immediate problems with the installations were discovered. Still, this session

clearly shows the possibilities of having people from both production and design jointly review and discuss the project in VR without even having to leave their own workplace. Also, compared to the co-located sessions, we see no major differences in behavior among participants, either in questionnaire response, or from the observations.

Overall, it can be stated that the possibility of being several people at the same time in the VR model provided a clear added value and contributed to increased understanding and communication between the participants. Compared to when several people take turns putting on the VR glasses and entering the model, communication and understanding is facilitated by being able to see where the others are, what they are looking at and pointing on, and scaling 1:1 and stereo vision also means that "everyone sees the same thing". Even in cases where all participants connected from widely different locations, observations together with comments and reactions from the participants show that they act very much like standing next to each other and looking at the same things in the model.

Previous research on multi-user VR has highlighted the importance and effects of realistic avatars on user experience [60]. However, the data from our study doesn't identify that as particular important. Nothing was ever mentioned negatively regarding the simplicity of the avatars. In fact, the only thing explicitly commented about the avatars was a question regarding why one of them had ordinary glasses (i.e. spectator mode), and that some participants expressed it was a cool thing that a user inspecting the mini-model will appear as a "giant" to those being in scale 1:1. Other than that, participants reacted surprisingly normal in the interactions with each other and with primary focus on the tasks (e.g. design review, planning). Typically they gathered around a specific issue, like a clash and then discussed verbally while also using pointing, selection, or markups to further aid the discussion. It is thus suggested that request and need for realism is instead strongly connected to the application and main purpose with the use of the technology. With the focus in this context being on use of VR for design and constructability review, planning, work preparation, and general information about the project, things like photorealism is of less importance. In fact, material and colors are instead asked to be more symbolic or schematic to illustrate discipline, or functionality (e.g. fresh air inlet). With that focus it is perhaps logical that the avatars are of more symbolic character. Also, site personnel in Scandinavia today are getting more and more used to the concept of BIM and the BIM-tools that are often used on site are more "engineering software" with focus on geometry and data [24]. That is, high level of detail appears to be more




Fig. 19. Illustration of sequencing using color-coding and filtering from IFC-properties (From: Project G).




Fig. 20. Constructability review with multiple users performed fully remote (From: Project D).

important than high level of photorealism. It is still possible that more realistic, human-like, avatars would improve the user experience, level of presence, communication, and social interaction [60], however, it is clear that the current approach is “good enough” and does not negatively impact the multi-user experience.

Still, one thing that was identified as a potential issue with remote meetings was the number of participants, mainly when the participants don’t know each other or have met before. With many users it might be difficult to always understand who is talking, especially when the participants have never met before and can’t easily recall each other’s voices. Therefore it is perhaps good to have a limit of around 5–6 participants. Important to note here is that this was not something mentioned by any of the participants (that had all met before in real-life), but instead observed by one of the researchers (that had not met most of the participants in real-life) during the D project meeting.

6.6. openBIM, properties, and filtering

In Scandinavia, BIM-projects are typically synonymous with IFC-files from the contractor perspective. In fact, only in rare cases do contractors and the on-site organization even have access to the native BIM models. There are multiple reasons for this, such as contractual and the need for a neutral collaboration format, but also that almost all clash detection and model review for non-infrastructure projects is done in Solibri or BIMCollab Zoom, using IFC. As such, it is clear that any VR system that needs to be used on a regular basis in a construction-context needs to have support for IFC. Furthermore, it is clear that the concept of properties and metadata is starting to become common knowledge throughout the on-site organization, much thanks to the use of software like Solibri, Dalux, and BIM360. Features found in those systems, such as filtering, are then requested, or even expected to be found in the VR system. As already discussed, in these contexts, realism, materials, and textures appears to be of less importance, while instead geometric detail and color-coding are features that are considered important and often asked for. Given the focus on constructability and planning, for instance, the preference for high geometric detail and schematic coloring is definitely logical. From the observations (please see the examples from project A and C in section 6.7) and comments in the questionnaire we can understand the benefits of detailed components for constructability review, and also schematic coloring as it helpful to illustrate responsibilities (e.g. contractor ID in project G, new vs. existing in project C). Still, this might also – at least to some degree – be connected to expectations, as they are all informed that they are about to enter and experience “the BIM” in VR. At least for people on-site, from their perspective “BIM” and “IFC” is often analogous with Solibri, BIMCollab Zoom, Dalux, StreamBIM, or the BIM360 viewer, and probably what they are expecting to see. If told that they were about to experience a VR-visualization of the project, or a real-time rendering from Revit or something from a game engine, expectations might instead be more towards photorealism. However, it should also be noted that all of this is from the perspective of production, and not from the client or architect.

Either way, there is definitely potential for future research around this aspect.

From a data management perspective, issues (e.g. design issues) are handled differently in all the projects, ranging from simply exchanging BCF-files to using more integrated cloud-based solutions (e.g. BIMcollab Cloud). However, the common denominator is that BCF can be used as a transfer format between different systems, and the possibility to take snapshots in VR and export as BCF was considered very useful and also important in order for a smooth integration with all the other BIM-systems. Still, several participants also asked for the possibility to import BCF-issues from other systems into VR for further inspection:

“Very often we find potential issues and clashes in Solibri, that would be very interesting to look further into in a VR environment.” [Design manager]

Regarding interoperability, we have had very few issues with the xBIM Toolkit and the IFC-import procedure. With xBIM creating an in-memory representation of the whole IFC-file, all of the metadata contained in it (e.g. properties, levels, relations, etc.) is available “by definition” as we are not reading another representation of the file, but, in fact, *the* file. The main challenge with IFC-files is instead the geometry, as it is often an implicit geometry representation (e.g. CSG geometry) that has to be converted to an explicit representation (i.e. triangular mesh) in order for the GPU to be able to rasterize it. Still, the only issue encountered in terms of geometry is that several objects in the structural prefab model in Project A were “inverted” the first time we tested that model. BIMXplorer has a “recompute normals on selected objects” functionality that we used at the time to fix the problem. However, after some searching in the xBIM Github forum we got the tip to check the volume of the object – using a *tetrahedron volume computation* – and if negative we simply flip all triangles in the model. Hopefully this will be implemented and fixed in the xBIM code as well in the future. Still, as this is a very fast calculation we do that during import on all objects with a solid geometry, and have had no problem since. In other projects (i.e. outside of this study) we sometimes encounter geometry in Tekla models that is not imported correctly (but is imported correctly in Solibri), but it is a rare case, something like a single object every tenth project we evaluate. To summarize, xBIM is a mature, stable, and fast IFC Toolkit.

6.7. Design, constructability, and planning review

In project C, the focus was a large installation space that involved both new and existing installations. Here, rule-base color-coding was applied in order to clarify between *existing-to-keep*, *existing-to-remove*, and *new installations*. The work in this area was to begin the following week and the MEP subcontractor fitter reviewed the VR model for about an hour together with the project manager. They went through the planned execution in detail and discovered several places in need of modifications. Above all, they discovered that in several places it was not an optimal design from a production perspective, and the installation fitter then asked if it was possible to move or draw new pipes and

ducts directly in the VR model. This was not possible at the time in BIMXplorer, but the fitter was then creative on his own and used the markup tool with an excessively “thick” pen to simply draw the alternative piping in 3D, see Fig. 21. After this was drawn up, a number of screenshots were taken of the model and these were then used as print-outs the following week during the actual execution – a clear example of how the use of immersive VR allows better solutions to be found than initially designed and planned for.

This scenario is interesting for a number of reasons; First, it is an actual example of *Production-Oriented Views* [47] created fully in VR, by the installation fitter, and then used on-site; Second, it shows the potential – and also the need – of bringing construction-knowledge into the design process. Ideally, this VR review session should have been held already during the design stage; Third, it highlights the importance of being able to color-code objects on-demand based on properties, in this case illustrated by the need to easily distinguish between different installations; Fourth, it appears that VR, as a user-friendly interface to the BIM, “encourage” users to actually edit or change the model in case of errors. This last point is also brought up several times in project H, from different MEP disciplines, where they basically ask “*Can I move this?*”

The multi-user VR job planning session in project A (section 6.4) had been preceded by similar planning and review sessions, however without the multi-user functionality. Also then, it was a lot about sequencing and whether lack of space meant that the planned execution order had to be changed. In addition, the following major issues were detected in the model, which later had to be updated in the design:

- Non-optimal placement of stairs in the climbing formwork system (logistical issues)
- Too small openings in prefabricated wall elements (logistical/transport problems otherwise)
- Too low railings (safety issues), see Fig. 22
- Wrong type of radiators on one floor (106 pcs), see Fig. 22

Important to note here is that conventional BIM model review and clash detection in Solibri had already been performed without these issues being detected. Here, the radiators are perhaps the most interesting. Being a high-rise office building, the sill height is the same on all floors, except for one. However, design had assumed the same sill height and just copied and used the same radiator layout on all floors, which wouldn't work with a lower sill height. From a 2D plan view this is impossible to spot, and clash detection will not find it either. Instead, one of the construction workers reviewed this floor in VR and instinctively thought “*I can't mount the brackets behind those radiators*”.

Similar examples that have been found in other projects during this study include wrong components (e.g. fire door instead of normal door), wrong placement of components (e.g. smoke detectors), lack of space (e.g. spaces around doors, stairs ending too close to a wall or door), and constructability issues and general design errors (e.g. wrong connections between walls and roof). Beyond that, a lot of (physical) clashes have been identified. However, these clashes are often in places where there

hasn't yet been a “final” clash control. Overall, it appears that clash control (e.g. Solibri) work well in most cases and that it is not primarily to detect “physical” collisions that VR should be used for – a similar conclusion that Haahr & Knak [31] landed at. However, as the above examples show, a clash-free model does not automatically guarantee that there are no problems or that the design is sound from a constructability perspective. Ideally, clash detection should therefore run in parallel with constructability reviews and in this context VR is highlighted as a suitable tool by the respondents.

7. Research limitations

Due to the strong focus on projects with a relatively high level of BIM-use, the participants were generally interested in new technology, had knowledge about BIM and also experience from working with digitalized, model-based construction. On the one hand, they therefore represent a modern, digitalized branch within the construction industry that we are likely to see much more of in the future, but on the other hand they may not be reflecting the typical or average situation today. As such, our results and conclusions might – at least to some degree – exhibit pro-innovation bias, which we acknowledge is a limitation of the research. However, at the same time the results are forward looking.

When instead looking purely at the rendering technology, almost the opposite can be said. We showed that the developed technology and a GPU-driven occlusion culling algorithm allowed large federated BIMs from real-world projects to be rendered in real-time on modern HMDs. Still, all the HMDs used during the evaluations were connected to a (gaming) computer, which, considering the overall focus today on developing untethered, standalone devices, might not be a true reflection of the hardware we are likely to see in the future [19]. Even as we see impressive performance increase from generation to generation on mobile devices, they are currently not near that of the dedicated GPUs that has been used in the evaluations in this research project. As such, our technological contribution can be considered missing a forward looking component, which we acknowledge can be seen as a limitation. However, there are multiple sides also to this. With streaming approaches for standalone VR (e.g. remote, server-side rendering across 5G networks) the developed technology is, in fact, directly applicable as these solutions do use dedicated GPUs. Similarly, if we look enough forward in time, also mobile devices will have the same performance that dedicated GPUs have today, which will then make the presented algorithms equally applicable.

8. Conclusions

In this paper we have presented the design, development, and evaluation of a multi-user VR system for immersive visualization of large and complex BIMs. The VR system address several technical challenges and barriers for efficient integration of BIM and VR such as interactivity and rendering performance, as well as interoperability and data management through openBIM. The evaluation of this system has focused on



Fig. 21. Design review and sketch of more production-adapted MEP installation (From: Project C).




Fig. 22. Errors found in VR; too low railings (left) and wrong type of radiators (right) (From: Project A).

identifying and understanding benefits and use cases of both single and multi-user VR in a construction-oriented context. As gathered from eight real-world projects, our data and analysis show that there are great benefits and opportunities by letting construction personnel use VR technology for design and constructability review, sequencing, and job planning. By involving staff with knowledge and experience from construction production, many issues have been identified and resolved before reaching the actual production stage. In some cases, it has been pure design errors or constructability issues, while in other cases it has been about changing the sequencing order between disciplines, or identifying alternative solutions that would be better from a productivity perspective. The primary benefit of VR compared to non-immersive alternatives is found in the 1:1 scale and stereo vision, which gives a much better understanding of size, space, and details – this at the same time as the miniature model with by-floor sectioning still can offer easy overview and orientation similar to a traditional 2D plan view. In addition, our observations suggest that participants inspect immersive VR-models much like they inspect real environments, using egocentric cues and interactions.

Regarding multi-user VR, this feature was rated very high by all of the participants, and it is clear from our results that it improves communication, understanding, and collaboration. Even in remote sessions, there appears to be a surprisingly high sense of presence among participants, mainly facilitated by them being able to see where the other users are, what they are looking at and pointing at. Also, multi-user interactions do not seem to require realistic avatars in order to be efficient.

Furthermore, with increased on-site use of BIM-viewers like Dalux, BIM360, and StreamBIM, site personnel are getting more used to the concept of BIM, federation, and properties, and are now starting to request features like BIM-properties, filtering, and color-coding also in VR. Such features can be provided by directly supporting openBIM and IFC which additionally solves many problems with BIM and VR reported in previous research, such as interoperability and data management issues. In relation to this we also see that high level of detail in combination with schematic coloring appears to be more important than realism and textures when used in a construction-oriented context. Still, directly supporting federated, high-detail BIMs can be a challenging task for a graphics application to manage in real-time, especially with stereo rendering and VR. However, as we have shown with BIMXplorer, this can be solved by occlusion culling and modern, GPU-driven rendering techniques.

For future work it would be interesting to explore the use of VR technology with a construction-oriented focus already during the design stage by bringing in knowledge and experience from production. Although we see in this study that the VR technology allows design issues – or even better solutions – to be found before reaching the actual production stage, it would be preferable to identify it already at the design stage. Given that the general use of desktop computers, design

software, and BIM-tools are typically higher during this phase it is probably even easier to integrate the VR technology already during design. Here, multi-user capabilities can also make it easier to gather the required competence without the need for everyone to meet at the same physical location. Another, related, interesting topic would be investigation around the relation and importance of geometric detail and suitable representation and level of realism for various stages and disciplines.

Regarding technology it is reasonable to expect that with untethered and standalone VR devices it would be easier to integrate VR as tool for daily use in construction, simply because then we can exclude the need of a PC. Still, as we have presented in this paper, real-world models from real-world projects tend to be extremely large and complex, which makes this a less easy task in practice. The next logical step would therefore be to focus on developing strategies and algorithms to support large model rendering even on standalone devices like Oculus Quest. Even if GPU-driven, indirect rendering and occlusion culling is expected to be useful also here, these devices will still not be able to render nearly as many visible triangles as dedicated GPUs, which means that also level-of-detail or prioritized rendering needs to be considered.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Some data will be made available on request (all models and data cannot be shared, however)

Acknowledgements

This is an extended and updated version of one of the best papers presented initially at CONVR 2022 organized and led by Prof. Chansik Park of Chung Ang University in Seoul, Korea. This work is funded and part of the Digital Twin Cities Centre supported by Sweden's Innovation Agency Vinnova under Grant No. 2019-00041 and by SBUF Grant No. 13763 (Development Fund of the Swedish Construction Industry).

References

- [1] A. Abbas, M. Choi, J. Seo, S.H. Cha, H. Li, Effectiveness of immersive virtual reality-based communication for construction projects, *KSCSE J. Civ. Eng.* 23 (2019) 4972–4983, <https://doi.org/10.1007/s12205-019-0898-0>.
- [2] M. Afzal, M.T. Shafiq, Evaluating 4D-BIM and VR for effective safety communication and training: a case study of multilingual construction job-site crew, *Buildings* 11 (8) (2021) 319, <https://doi.org/10.3390/buildings11080319>.

- [55] C. Morse, Gaming engines: Unity, unreal, and interactive 3D spaces, *Technology | Architecture + Design* 5 (2) (2021) 246–249, <https://doi.org/10.1080/24751448.2021.1967068>.
- [56] A.A. Muhammad, I. Yitmen, S. Alizadehsalehi, T. Celik, Adoption of virtual reality (VR) for site layout optimization of construction projects, *Teknik Dergi* 31 (2) (2019) 9833–9850, <https://doi.org/10.18400/tekderg.423448>.
- [57] J. Ojala, J. Selin, T. Partala, M. Rossi, Virtual construction: Interactive tools for collaboration in virtual reality, in: *Advances in Information and Communication. FICC 2020. Advances in Intelligent Systems and Computing* vol. 1130, Springer, Cham, 2020, https://doi.org/10.1007/978-3-030-39442-4_26.
- [58] G. Ozcan-Deniz, Expanding applications of virtual reality in construction industry: a multiple case study approach, *J. Construct. Eng. Manag. Innov.* 2 (2) (2019) 48–66, <https://doi.org/10.31462/jcemi.2019.02048066>.
- [59] D. Paes, J. Irizarry, D. Pujoni, An evidence of cognitive benefits from immersive design review: comparing three-dimensional perception and presence between immersive and non-immersive virtual environments, *Autom. Constr.* 130 (2021), 103849, <https://doi.org/10.1016/j.autcon.2021.103849>.
- [60] M. Pakanen, P. Alavesas, N. van Berkel, T. Koskela, T. Ojala, “Nice to see you virtually”: thoughtful design and evaluation of virtual avatar of the other user in AR and VR based telepresence systems, *Entertain. Comput.* 40 (2022), 100457, <https://doi.org/10.1016/j.entcom.2021.100457>.
- [61] A. Prabhakaran, A.M. Mahamadu, L. Mahdjoubi, Understanding the challenges of immersive technology use in the architecture and construction industry: a systematic review, *Autom. Constr.* 137 (2022), 104228, <https://doi.org/10.1016/j.autcon.2022.104228>.
- [62] A. Prabhakaran, A.M. Mahamadu, L. Mahdjoubi, P. Boguslawski, BIM-based immersive collaborative environment for furniture, fixture and equipment design, *Autom. Constr.* 142 (2022), 104489, <https://doi.org/10.1016/j.autcon.2022.104489>.
- [63] L. Potseluyko, F.P. Rahimian, N. Dawood, F. Elghaish, A. Hajirasouli, Game-like interactive environment using BIM-based virtual reality for the timber frame self-build housing sector, *Autom. Constr.* 142 (2022), 104496, <https://doi.org/10.1016/j.autcon.2022.104496>.
- [64] F.P. Rahimian, V. Chavdarova, S. Oliver, F. Chamo, L.P. Amobi, OpenBIM-Tango integrated virtual showroom for offsite manufactured production of self-build housing, *Autom. Constr.* 102 (2019) 1–16, <https://doi.org/10.1016/j.autcon.2019.02.009>.
- [65] A. Rashidi, W.Y. Yong, D. Maxwell, Y. Fang, Construction planning through 4D BIM-based virtual reality for light steel framing building projects. *Smart and Sustainable Built Environment*, (ahead-of-print), 2022, <https://doi.org/10.1108/SASBE-06-2022-0127>.
- [66] J.F. Fernández Rodríguez, Implementation of BIM virtual models in industry for the graphical coordination of engineering and architecture projects, *Buildings* 13 (3) (2023) 743, <https://doi.org/10.3390/buildings13030743>.
- [67] M. Roupé, M. Johansson, M. Viklund Tallgren, F. Jörnebrant, P.A. Tomsa, Immersive visualization of building information models, in: *Proceedings of the 21st International Conference of the Association for Computer-Aided Architectural Design Research in Asia (CAADRIA 2016)*, 2016, pp. 673–682, <https://doi.org/10.52842/conf.caadria.2016.673>.
- [68] R. Sacks, C. Eastman, G. Lee, P. Teicholz, *BIM Handbook: A Guide to Building Information Modeling for Owners, Designers, Engineers, Contractors, and Facility Managers*, John Wiley & Sons, 2018 (ISBN: 9781119287537).
- [69] S. Sateei, M. Roupé, M. Johansson, Collaborative design review sessions in virtual reality: multi-scale and multi-user, in: *Proceedings of the 27th CAADRIA Conference (CAADRIA 2022)*, 2022, pp. 29–38, <https://doi.org/10.52842/conf.caadria.2022.1.029>.
- [70] J. Scholtz, M.P. Steves, A framework for real-world software system evaluations, in: *Proceedings of the 2004 ACM Conference on Computer Supported Cooperative Work*, 2004, pp. 600–603, <https://doi.org/10.1145/1031607.1031710>.
- [71] Y. Shi, J. Du, S. Lavy, D. Zhao, A multiuser shared virtual environment for facility management, *Proc. Eng.* 145 (2016) 120–127, <https://doi.org/10.1016/j.proeng.2016.04.029>.
- [72] Y. Shi, J. Du, E. Ragan, K. Choi, S. Ma, Social influence on construction safety behaviors: a multi-user virtual reality experiment, in: *Construction Research Congress 2018*, 2018, pp. 174–183, <https://doi.org/10.1061/9780784481288.018>.
- [73] M.F. Shiratuddin, W. Thabet, D. Bowman, Evaluating the effectiveness of virtual environment displays for reviewing construction 3D models, in: *4th Conference on Construction Applications of Virtual Reality (CONVR 2004)*, Pp, 2004, pp. 87–98, <https://researchportal.murdoch.edu.au/esploro/outputs/conferencePaper/Evaluating-the-effectiveness-of-virtual-environment/991005539578007891> (accessed Oct. 23, 2023).
- [74] V. Schwind, K. Wolf, N. Henze, Avoiding the uncanny valley in virtual character design, *Interactions* 25 (5) (2018) 45–49, <https://doi.org/10.1145/3236673>.
- [75] O.V. Smyslova, A.E. Voiskounsky, Usability studies: to meet or not to meet intrinsic motivation, *Psychology J.* 7 (3) (2009) 303–324, [http://www.psychology.org/File/PNJ7\(3\)/PSYCHOLOGY_JOURNAL_7_3_SMYSLOVA.pdf](http://www.psychology.org/File/PNJ7(3)/PSYCHOLOGY_JOURNAL_7_3_SMYSLOVA.pdf) (accessed Oct. 23, 2023).
- [76] M.V. Tallgren, M. Roupé, M. Johansson, 4D modelling using virtual collaborative planning and scheduling, *J. Inform. Technol. Construct. (ITcon)* 26 (42) (2021) 763–782, <https://doi.org/10.36680/j.itcon.2021.042>.
- [77] S. Tea, K. Panuwatwanich, R. Ruthankoon, M. Kaewmoracharoen, Multiuser immersive virtual reality application for real-time remote collaboration to enhance design review process in the social distancing era, *J. Eng. Design Technol.* 20 (1) (2022) 281–298, <https://doi.org/10.1108/JEDT-12-2020-0500>.
- [78] P. Truong, K. Hölttä-Otto, P. Beceril, R. Turtiainen, S. Siltanen, Multi-user virtual reality for remote collaboration in construction projects: a case study with high-rise elevator machine room planning, *Electronics* 10 (22) (2021) 2806, <https://doi.org/10.3390/electronics10222806>.
- [79] M. Umair, A. Sharafat, D.E. Lee, J. Seo, Impact of virtual reality-based design review system on user’s performance and cognitive behavior for building design review tasks, *Appl. Sci.* 12 (14) (2022) 7249, <https://doi.org/10.3390/app12147249>.
- [80] D. Ververidis, S. Nikolopoulos, I. Kompatsiaris, A review of collaborative virtual reality systems for the architecture, engineering, and construction industry, *Architecture* 2 (3) (2022) 476–496, <https://doi.org/10.3390/architecture2030027>.
- [81] J. Wang, R. Shi, W. Zheng, W. Xie, D. Kao, H.N. Liang, Effect of frame rate on user experience, performance, and simulator sickness in virtual reality, *IEEE Trans. Vis. Comput. Graph.* 29 (5) (2023) 2478–2488, <https://doi.org/10.1109/TVCG.2023.3247057>.
- [82] J. Wolfartsberger, Analyzing the potential of virtual reality for engineering design review, *Autom. Constr.* 104 (2019) 27–37, <https://doi.org/10.1016/j.autcon.2019.03.018>.
- [83] S.E. Yoon, E. Gobetti, D. Kasik, D. Manocha, Real-Time Massive Model Rendering, *Springer Nature*, 2008, <https://doi.org/10.1007/978-3-031-79531-2>.
- [84] R. Zaker, E. Coloma, Virtual reality-integrated workflow in BIM-enabled projects collaboration and design review: a case study, *Vis. Eng.* 6 (1) (2018) 1–15, <https://doi.org/10.1186/s40327-018-0065-6>.
- [85] Z. Zou, L. Arruda, S. Ergun, Characteristics of models that impact transformation of BIMs to virtual environments to support facility management operations, *J. Civil Eng. Manag.* 24 (6) (2018) 481–498, <https://doi.org/10.3846/jcem.2018.5689>.
- [86] Y. Collet, M. Kucherawy, Zstandard Compression and the Application/Zstd Media Type.. Technical Report RFC8478. RFC Editor., 2018, <https://doi.org/10.17487/RFC8478>.