

A Taxonomy of 3D Occlusion Management for Visualization

Niklas Elmqvist, *Member, IEEE*, and Philippos Tsigas, *Member, IEEE*

Abstract—While an important factor in depth perception, the occlusion effect in 3D environments also has a detrimental impact on tasks involving discovery, access, and spatial relation of objects in a 3D visualization. A number of interactive techniques have been developed in recent years to directly or indirectly deal with this problem using a wide range of different approaches. In this paper, we build on previous work on mapping out the problem space of 3D occlusion by defining a taxonomy of the design space of occlusion management techniques in an effort to formalize a common terminology and theoretical framework for this class of interactions. We classify a total of 50 different techniques for occlusion management using our taxonomy and then go on to analyze the results, deriving a set of five orthogonal design patterns for effective reduction of 3D occlusion. We also discuss the “gaps” in the design space, areas of the taxonomy not yet populated with existing techniques, and use these to suggest future research directions into occlusion management.

Index Terms—Occlusion reduction, 3D visualization, taxonomy, design patterns, visual cues, depth perception.

I. INTRODUCTION

HUMAN beings employ all manners of visual cues and hints in order to correctly perceive and understand the three-dimensional world surrounding us. Of course, these visual cues can also work against us, fooling our perception into believing things about our environment that are simply not true. In some cases, this is done intentionally through various forms of optical illusions that exploit special characteristics of our minds. A more subtle point, however, is that we can instead choose to directly weaken certain of these visual cues in order to help the human to perceive and understand *more* of her surroundings. In some cases, this selective weakening of visual cues, primarily occlusion, size, and shape, can lead to dramatically increased performance when solving specific tasks in a 3D environment. While this may be difficult to achieve in the real world, it is a perfectly viable approach in a virtual 3D world being visualized on a computer.

In this paper, we explore *occlusion management*, a concept that is applicable to a wide range of domains ranging from general computer graphics to computer vision and robotics. The scope of our work is for visualization and human-computer interaction, where occlusion management techniques modify views, objects, and depth cues in order to increase the spatial awareness of the human user and to facilitate special tasks, such as navigating, searching, or understanding the 3D world.

More specifically, we present a taxonomy consisting of a small set of dimensions describing important characteristics of these techniques, focusing on the purpose, strength, view paradigm, depth cues, interaction model, preserved invariances, and solution

space of each technique. We then go on to classify 50 different methods that have been described previously in the literature into the taxonomy. These classifications form a body of data that we can analyze for trends and the existence of clusters. This analysis yields in turn five orthogonal *design patterns* that characterize current work in the field. The patterns are multiple viewports, virtual X-ray tools, tour planners, volumetric probes, and projection distorters, and we describe the typical uses and characteristics of each pattern. More importantly, the pattern identification process also serves to pinpoint the “gaps” in the taxonomy, i.e. as-of-yet undeveloped techniques that could potentially fulfill a useful role in future research.

Thus, the purpose of this taxonomy is manifold: (i) to provide a common theoretical framework and vocabulary for occlusion management techniques for visualization, giving researchers and practitioners alike a common ground for discussion; (ii) to facilitate qualitative comparison, evaluation, and maybe even benchmarking of different methods for occlusion management; (iii) to suggest a small number of archetypes of design suitable as starting points for implementations and prototypes; and (iv) to inform future directions of research within occlusion management and human perception of 3D space.

This paper is organized as follows: We begin by discussing previous taxonomies in this area and in related areas. We then describe the problem space of occlusion in 3D environments, where the fact that nearby objects occlude more distant ones work against the human perceptual system. This is followed by a presentation of our taxonomy and its dimensions. We also present the full classification of the 50 techniques we have studied in this paper. We then identify and describe the five design patterns, followed by suggestions on future research directions based on unexplored parts of the taxonomy. We finish the paper with some discussions on the design, limitations, and possible future extensions of the taxonomy and our conclusions on the work.

II. RELATED WORK

As mentioned in the introduction, occlusion management is actually an instance of a more general class of visibility problems for computer graphics. As such, some of the aspects discussed in this paper appear in many other contexts: computing occlusion for accelerating real-time rendering, global illumination, or shadow computation, as well as in computational geometry and path planning (such as the art gallery problem). While these topics clearly are outside the visualization scope of this paper, relevant sources include [2], [3], [4].

No previous taxonomy exists in the literature on the class of occlusion management interaction techniques for visualization. More general taxonomies on 3D interaction tend to describe low-level mechanics of manipulative tasks in a morphological fashion, whereas our focus is on high-level aspects of perceptual tasks

N. Elmqvist is with INRIA in France, P. Tsigas is with Chalmers University of Technology in Sweden.

This is an extended version of a paper that appeared at the IEEE Conference on Virtual Reality 2007 held March 10–14 in Charlotte, NC, USA [1].

related to spatial understanding of the 3D environment and its objects. For example, Bowman and Hodges [5] present a general formal framework for 3D interaction in immersive virtual environments (IVEs) based around three tasks: motion control, selection, and manipulation. Bier et al. give a taxonomy of see-through tools [6] for a class of double-handed interaction techniques using transparent sheets called toolglasses that partly inspired this taxonomy. Bowman et al. [7] present a descriptive view of the design space of information display as well as interaction for information visualization within virtual environments.

Although unrelated to the occlusion management area defined here, Pousman and Stasko’s taxonomy of ambient visualization [8] inspired the method employed in this paper for presenting the classifications and deriving the design patterns from the classification data.

We use our taxonomy as a tool for classifying existing techniques and thus validating its generality, but also as a design space. This allows us to identify holes in the taxonomy, akin to Card et al. [9].

Occlusion management deals mainly with 3D visualization, where the user is interested in retrieving information from the environment in order to solve some specific task. The concept of *information-rich virtual environments* (IRVEs), defined by Bowman et al. [7], is of particular interest in this context, because it describes a framework for combining information visualization with 3D virtual environments. Many of the techniques classified in this paper can be applied to managing the challenges [10] presented in their work.

III. PROBLEM SPACE

The occlusion problem space in 3D environments is defined by the intrinsic *properties* of the environment, their interaction with *human cognition*, the *visual tasks* involved, and the ensuing effects caused by the occlusion. When an observer navigates in an environment used for 3D visualization, the environment itself and its geometrical properties causes occlusion of objects. If important target objects are hidden from view, correctness and productivity will suffer.

A. Model

We represent the 3D world U by a Cartesian space $(x, y, z) \in \mathbb{R}^3$. Objects in the set O are volumes within U (i.e. subsets of U) represented by boundary surfaces (typically triangles). The user’s viewpoint $v = (M, P)$ is represented by a view matrix M that includes the position and orientation of the user, as well as a projection matrix P that includes view parameters such as viewport dimensions, focal length, far and near clipping plane, etc.

A line segment r is *blocked* by an object o if it intersects any part of o . An object o is said to be *occluded* from a viewpoint v if there exists no line segment r between v and o such that r is not blocked. Analogously, an object o is said to be *visible* from a viewpoint v if there exists a line segment r between v and o such that r is not blocked. An object o is said to be *partially occluded* from viewpoint v if o is visible, but there exists a line segment r between v and o such that r is blocked.

An object can be flagged either as a *target*, an information-carrying entity, or a *distractor*, an object with no intrinsic information value. Importance flags can be dynamically changed. Occluded distractors pose no threat to any analysis tasks performed

in the environment, whereas partially or fully occluded targets do, potentially causing decreased performance and correctness.




Fig. 1. Basic model for 3D occlusion. Target objects are flagged with “T” and distractors with “D”.

Figure 1 shows a diagram of this basic model. Here we see three objects A , B , and C , the first of which is a distractor and the other two targets. The shaded area represents areas that are invisible to the user from the current view. We can easily see that A is fully visible (but is a distractor), B is partially occluded, and C is occluded.

A set of viewpoints V is said to be *complete* if there exists no object that is occluded in all of the viewpoints v_i . For instance, from the figure it is clear that the set $V = \{v_0, v_1\}$ is complete for the simple environment given in the example (in fact, for this simple situation, it is possible to find a single viewpoint from which all objects are visible).

B. Visual Tasks

The occlusion problem in visualization typically occurs in the following three *visual perception tasks*:

- *target discovery*: finding targets $t \in O$ in the environment;
- *target access*: retrieving graphically encoded information associated with each target; and
- *spatial relation*: relating the spatial properties of a target with other targets and its context.

An example of the target discovery task could be for a fire brigade captain to find all of the water hydrants in a 3D command and control visualization of a burning building. An example of target access might be an engineer retrieving stress measurements in a volumetric visualization of a tunnel construction. Finally, an example of spatial relation could be a power company manager studying the connectivity of the power grid and the relay stations in a 3D visualization in order to reroute around a malfunctioning section of the grid.

Other visual tasks that are of relevance beyond the ones discussed above include object *creation*, *deletion*, and *modification*, etc. In this treatment, however, we consider these to be special cases of discovery and access with regards to inter-object occlusion, and consisting of the same subtasks as these three basic visual tasks.

C. Analysis

It is clear that all visual tasks are severely hampered by the existence of fully occluded objects. For the purposes of target discovery, a fully occluded object will be impossible to discover without the use of some occlusion management strategy, and identifying whether the object is a target never becomes an issue. Analogously for target access, the visual search will fail, and so will the perception of the object's visual properties. As a result, both tasks will affect the efficiency and correctness of users solving tasks using a visualization, but clearly, threats to object discovery are the most serious: if the user is unaware of the existence of an object, she will have no motivation to look for it and access never becomes an issue.

Partial occlusion, on the other hand, has a different effect on these tasks. For target discovery, users may have difficulties distinguishing object identity if too large a portion of the object is occluded. In this situation, the user may either miss the object entirely, count the same object multiple times, or believe different objects are part of the same object. Target access, on the other hand, will succeed in the visual search, although the perception of the object may still fail due to important parts of it being occluded.

Spatial relation, necessary for many complex interactions and visualizations, requires overview of the whole world, and is thus severely affected by both partially and fully occluded objects.

D. Environment Properties

The geometrical properties of the visualization environment are of special interest in this framework because they allow us to characterize the visualization and determine the nature of the occlusion problems that may arise. These properties can also be used to decide which occlusion management strategies are applicable for a specific situation.

In this treatment, we identify three main geometrical properties of the environment that interact to cause inter-object occlusion and influence the three basic visual tasks associated with the environment:

- *object interaction*: spatial interaction of objects in the environment;
- *object density*: amount of objects in the environment with regard to its size; and
- *object complexity*: detail level of individual objects in the environment.

Obviously, these are high-level properties that only generally describe an environment without going into detail on its actual content. Nevertheless, in the following sections we shall see how these property dimensions can serve as powerful reasoning tools for describing a 3D environment and selecting a suitable solution strategy for it.

1) *Object Interaction*: The object interaction property dimension describes how the individual objects in the environment interact spatially with each other, i.e. whether they touch, intersect or merely reside close to each other. There are five ordinal levels to this parameter (see Figure 2 for a visual overview):

- *none*: no spatial interaction between objects (realistically only applicable for singleton environments);
- *proximity*: objects are placed in such close proximity (without intersecting) that they occlude each other from some viewpoint;

- *intersection*: objects intersect in 3D space (without one fully containing another) such that they occlude each other;
- *enclosurement*: one or several objects combine to fully enclose objects (without containing them) such that they are occluded from any viewpoint external to the enclosing objects; and
- *containment*: objects are fully contained in other objects such that they are occluded from any viewpoint.

Examples of these interaction levels exist in all kinds of 3D visualizations: proximity for nodes in 3D node-link diagrams, intersection for visualization of constructive solid geometry (CSG), enclosurement for furniture inside a virtual house, containment for 3D medical volumetric data, etc.




Fig. 2. Object interactions that may cause occlusion in 3D environments.

2) *Object Density*: The object density is a measure of the number of objects inhabiting the 3D environment; it follows naturally that the more objects per volume unit we are dealing with, the greater the chance and impact of occlusion will be. For singleton environments containing a single object, naturally only self-occlusion can occur.

3) *Object Complexity*: The third geometrical property with an impact on the occlusion characteristics of an environment is the complexity of the objects in the environment. With complexity, we refer to the detail level of the 3D objects, i.e. typically the number of triangles (or other 3D primitives, such as quads, lines, and points) that make up the object, but we also include attributes such as color, material, and texture in this parameter. It follows that the more complex an object is, the more information it can potentially encode (and vice versa), and the larger the impact occlusion has on identification and perception of the object.

For simplicity, we can often reduce object complexity by splitting objects into smaller (preferably convex) subobjects. Note that this will often result in an increased object interaction and density index. The same mechanism can be used to handle self-occlusion, i.e. when an object occludes parts of itself.

IV. DESIGN SPACE

In this treatment of the design space of occlusion management for visualization, we will only consider interaction techniques and systems that deal with 3D data and worlds, that aim to visualize the data in some fashion, and that at least roughly fit our concept of occlusion management: making information-carrying targets in 3D environments discoverable, accessible, or relatable despite visibility constraints.

Given these delimitations, we characterize the design space of occlusion management techniques using the following primary dimensions:

- **Primary Purpose (PP)**: Visual task that the technique is primarily targeting. [discovery, access, relation]
- **Disambiguation Strength (DS)**: Maximum object interaction that the technique can handle. [proximity, intersection, enclosurement, containment]

- **Depth Cues (DC):** Strength of depth disambiguation cues for the technique. [very low, low, medium, high, very high]
- **View Paradigm (VP):** View method used for the technique, i.e. the arrangement and layout of the visual substrate. [single view, double separate views, double integrated views, multiple separate views, multiple integrated views]
- **Interaction Model (IM):** Mechanism for target selection as well as operational model for the technique (in terms of user interaction). [passive/online, active/online, passive/offline, active/offline]
- **Target Invariances (A/D/G/L):** Degree of target invariances preserved using the technique. [0–4 aspects: appearance, depth, geometry, location]
- **Solution Space (SS):** Space utilized by the technique to achieve its purpose. [time, image, view, object]

These seven dimensions have been designed to be orthogonal, objective, and capture the full expressivity of the design space of these techniques. In the following sections, we will describe the dimensions in greater detail.

A. Primary Purpose

The purpose of an occlusion management technique describes which particular visual task in the problem space that the technique is primarily targeting (see Section III for the visual tasks). In other words, this dimension can assume any of the values discovery, access, or spatial relation.

More specifically, an interaction technique designed mainly for discovery focuses on making the user aware of the existence of partially or completely occluded targets, not necessarily making retrieval or relation of information from the objects easier.

A technique designed for access, on the other hand, aims not to make users aware of an occluded object, but to allow the user to retrieve the information encoded in the object. Note that supporting access does not necessarily mean that the technique also supports discovery; in some cases, the user may be required to know about the existence of an object in order to access it.

Finally, a technique supporting spatial relation is designed to make not only the object itself but also its surrounding context visible and understandable to the user. This means that it is not possible to simply get rid of the neighboring objects—targets and distractors alike—in the interest of seeing a particular target, since these may carry important information needed to understand the scene. Examples of this includes node-link diagrams, where only a few nodes may be of immediate interest, but the connectivity for the rest of the graph is important for context.

Note that a technique may have more than one purpose; this taxonomy dimension captures the primary purpose of the technique.

- **Domain:** *discovery, access, spatial relation* (nominal)
- **Characteristic Techniques:**
 - *discovery:* image-space dynamic transparency [27], navigation guide avatar [35]
 - *access:* interactive cut-away and break-away views [31], 3D explosion probe [12]
 - *spatial relation:* wayfinder [53], viewpoint selection for intervention planning [48]

B. Disambiguation Strength

Disambiguation strength refers to the maximum degree of object interaction that the technique can handle and still fulfill

its primary purpose. In other words, this is a measure of how complex object interactions the technique can manage using the terminology from the problem space (see Section III-D.1). Note that this metric is unrelated to object density, but that very high object density can confound the situation.

The strength of a technique is an ordinal dimension, and it is generally perceived better for a technique to be able to handle high object interaction. On the other hand, strength is related to other factors of the design space, leading to a trade-off between them. For example, virtual X-ray techniques (see Section V) typically support the highest object interaction (containment), yet are not as scalable as other techniques with more modest strengths.

- **Domain:** *proximity, intersection, enclosure, containment* (ordinal)
- **Characteristic Techniques:**
 - *proximity:* plan-based scene exploration [40], viewpoint-quality-driven scene exploration [49]
 - *intersection:* view-projection animation [47], worlds-in-miniature [54], bird’s eye views [19]
 - *enclosure:* worldlets [55], BalloonProbe [18], SDM [41]
 - *containment:* importance-driven volume rendering [29], view-dependent transparency [46], see-through surfaces [42]

C. Depth Cues

As we hinted at earlier in this paper, actually relaxing some of the visual cues humans rely on for spatial perception will most certainly have a negative impact on the user’s understanding of his or her surroundings, regardless of any advantages gained from doing this. The perception of depth, i.e. the actual 3D component of our vision system, is most vulnerable to this effect, and thus we define a dimension that captures the degree of depth cues that a technique provides.

Depth cues is an ordinal dimension with a five-value scale ranging from very low to very high, signifying the amount of depth cues retained by the technique; very high would mean that in principle all depth cues are preserved (as for normal vision), whereas very low means that practically none are.

There are additional visual cues that help humans perceive their environment and that play a role in the classification of occlusion management techniques, some of which we capture in the “target invariances” dimension below.

- **Domain:** *very low, low, medium, high, very high* (ordinal)
- **Characteristic Techniques:**
 - *very low:* artistic multiprojection [16]
 - *low:* ghosting (IBIS) [23]
 - *medium:* visual access distortion [51] blueprints [20]
 - *high:* 3D-zoom [15], magic mirror [34]
 - *very high:* StyleCam [43], standard 3D navigation [57]

D. View Paradigm

Different occlusion management techniques utilize the view and the view space in different ways; this dimension captures the paradigm employed for managing the visual substrate. Typically, interaction techniques are either based on a single view, double views, or a large number of views (multiple); similarly, for the case when there are additional views beyond the main one, they may either be separate windows in an overview+detail approach,

Technique	PP	DS	DC	VP	IM	A/D/G/L	SS
3D cutting planes [11]	access	contain	med	single	act/on	N/N/Y/Y	object
3D explosion probe [12]	access	enclose	med	single	act/on	Y/Y/Y/N	object
3D view management [13]	access	contain	med	single	pass/on	Y/Y/Y/Y	view
3D virtual X-ray magic lenses [14]	access	contain	low	single	act/on	N/N/Y/Y	object
3D-zoom [15]	access	intersect	high	single	act/on	Y/Y/N/Y	object
artistic multiprojection [16]	access	intersect	v. low	mult/int	pass/off	Y/Y/N/N	view
automatic tour generation [17]	relation	intersect	v. high	single	act/off	Y/Y/Y/Y	time
BalloonProbe [18]	access	enclose	med	single	act/on	Y/Y/Y/N	object
bird's eye views [19]	discovery	intersect	v. high	double	act/on	Y/Y/Y/Y	view
blueprints [20]	access	contain	low	single	pass/on	N/N/Y/Y	image
consistent 3D labels [21]	access	contain	v. low	single	pass/on	Y/Y/Y/Y	image
CubicalPath [22]	access	intersect	v. high	single	act/off	Y/Y/Y/Y	time
cutaways (IBIS) [23]	access	contain	v. low	single	pass/on	Y/N/Y/Y	image
deformation-based volume explosion [24]	access	contain	low	single	pass/on	Y/Y/N/N	object
general transfer-function lenses [25]	access	contain	med	single	act/on	N/Y/N/Y	object
ghosting (IBIS) [23]	access	contain	low	single	pass/on	N/N/Y/Y	image
illustrative shadows [26]	access	contain	v. low	single	pass/on	Y/N/Y/N	image
image-space dynamic transparency [27]	discovery	contain	med	single	pass/on	N/N/Y/Y	image
importance-driven focus of attention [28]	access	contain	low	single	pass/on	N/N/Y/Y	object
importance-driven volume rendering [29]	access	contain	med	single	pass/on	N/N/Y/Y	object
intelligent multi-shot visualization [30]	access	intersect	v. high	multiple	pass/on	Y/Y/Y/Y	time
interactive break-away views [31]	access	contain	v. low	single	pass/on	N/N/Y/Y	image
interactive cut-away views [31]	access	contain	v. low	single	pass/on	N/N/Y/Y	image
interactive figure captions [32]	access	contain	v. low	single	pass/on	Y/Y/Y/Y	image
looking glass [33]	access	contain	low	single	act/on	N/N/Y/Y	image
magic mirror [34]	access	intersect	high	double/int	act/on	Y/Y/N/Y	view
multiple inset views (IBIS) [23]	access	intersect	v. high	multiple	pass/on	Y/Y/Y/Y	view
navigation guide avatar [35]	discovery	intersect	v. high	single	act/off	Y/Y/Y/Y	time
object removal (IBIS) [23]	access	contain	v. low	single	pass/on	Y/N/Y/Y	object
occlusion-free route animation [36]	access	intersect	med	single	pass/on	Y/N/N/N	object
orthotumble [37]	access	intersect	low	double/int	act/on	Y/Y/N/N	view
path-planning navigation [38]	discovery	intersect	v. high	single	act/off	Y/Y/Y/Y	time
perspective cutouts [39]	access	enclose	v. low	single	act/on	Y/N/Y/Y	image
plan-based scene exploration [40]	discovery	proximity	v. high	single	pass/on	Y/Y/Y/Y	time
SDM [41]	access	enclose	med	single	act/on	Y/Y/N/N	object
see-through surfaces [42]	discovery	contain	med	single	act/on	Y/N/Y/Y	object
StyleCam [43]	access	intersect	v. high	single	act/off	Y/Y/Y/Y	time
task-level camera control [44]	discovery	intersect	single	v. high	pass/on	Y/Y/Y/Y	time
temporal non-linear projections [45]	access	intersect	low	single	act/on	Y/Y/N/N	view
view-dependent transparency [46]	access	contain	low	single	pass/on	N/N/Y/Y	image
view-projection animation [47]	discovery	intersect	low	double/int	act/on	Y/Y/N/N	view
viewpoint selection [48]	relation	intersect	v. high	single	pass/on	Y/Y/Y/Y	time
viewpoint-quality-driven exploration [49]	discovery	proximity	v. high	single	pass/off	Y/Y/Y/Y	time
virtual multiprojection cameras [50]	access	intersect	v. low	mult/int	pass/off	Y/Y/N/N	view
visual access distortion [51]	access	contain	med	single	act/on	Y/N/N/Y	object
volumetric rendering transfer functions [52]	access	contain	med	single	pass/on	N/N/Y/Y	object
wayfinder [53]	relation	intersect	v. high	single	pass/off	Y/Y/Y/Y	time
world-in-miniature [54]	discovery	intersect	v. high	double/int	act/on	Y/Y/Y/Y	view
worldlets [55]	discovery	enclose	v. high	multiple	act/off	Y/Y/Y/Y	view
X-ray tunnel [56]	discovery	contain	low	double	act/on	N/N/Y/Y	image

TABLE I

CLASSIFICATION OF THE 50 TECHNIQUES INTO THE TAXONOMY (ALPHABETICAL ORDER).

or integrated in the same image in a focus+context [58] way. The view paradigm dimension is used to classify techniques according to a combination of these two metrics.

The degree of integration can sometimes be tricky to assess—

for example, in the case of the worlds-in-miniature (WIM) [54] technique, there is a very obvious second view, i.e. a miniature version of the world, yet since it is a first-class object in the environment, we classify it as being integrated. For bird's eye

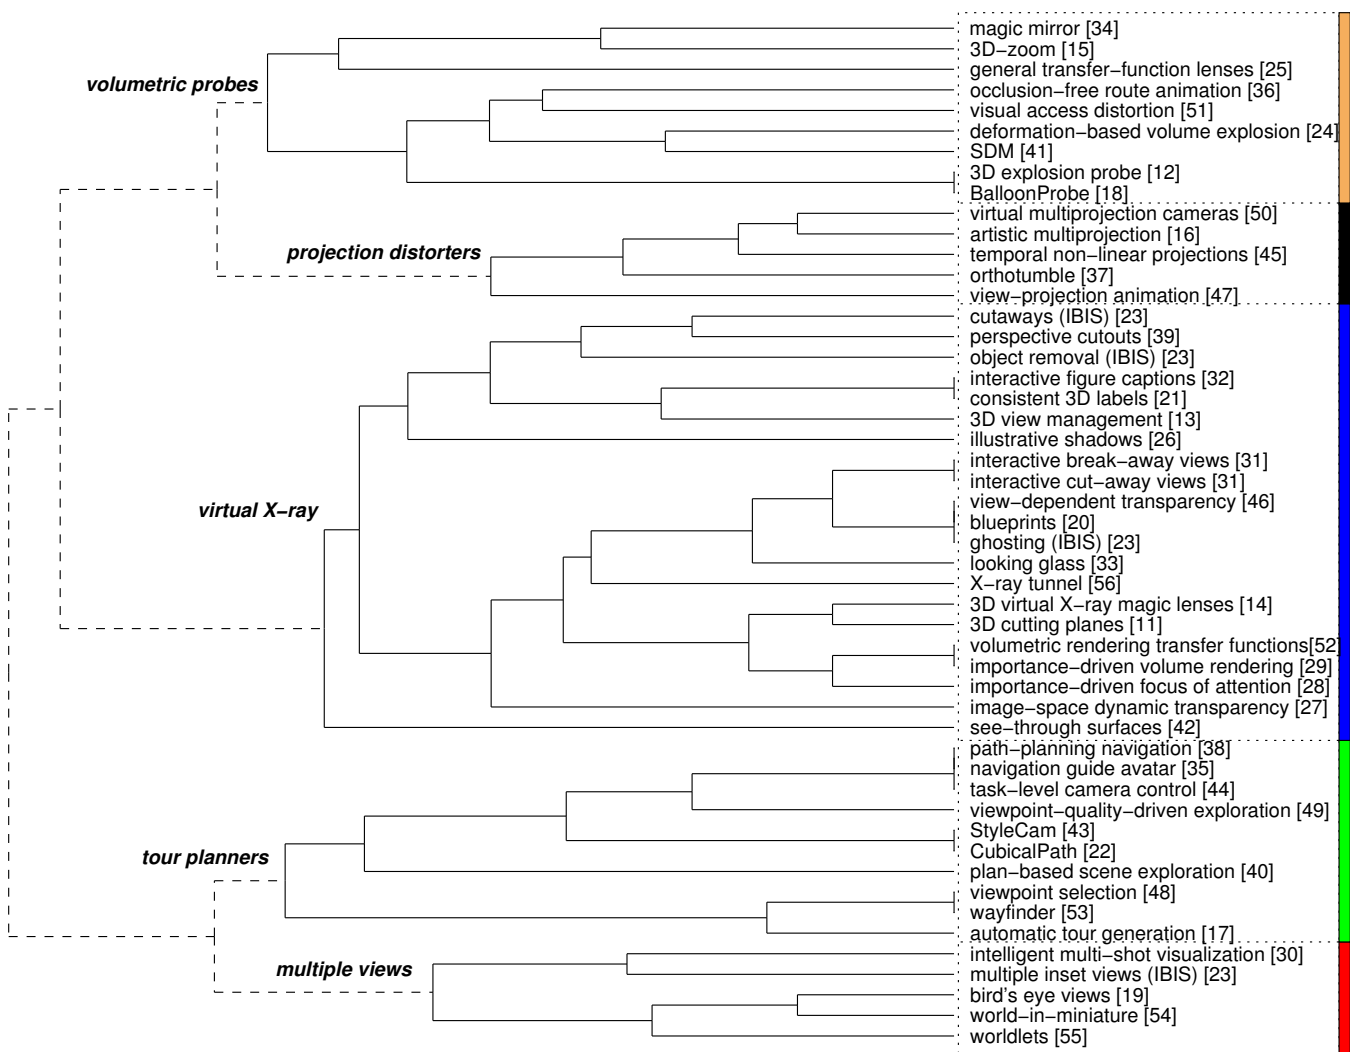


Fig. 3. Dendrogram showing the hierarchical clustering of the 50 techniques classified into the taxonomy. The five design patterns are drawn as solid lines whereas the higher-level hierarchy is shown using dashed lines.

views [19], on the other hand, the secondary view is in a separate window, and is thus classified as having double separate views. This factor is also the reason why separating the number of views from their integration is difficult. In the case of Singh's multiprojection techniques [50], the single view actually consists of multiple different cameras, non-linearly combined into one.

- **Domain:** *single view, double separate views, double integrated, multiple separate, multiple integrated* (nominal)
- **Characteristic Techniques:**
 - *single view:* 3D view management [13], 3D cutting planes [11]
 - *double separate views:* bird's eye views [19]
 - *double integrated views:* magic mirror [34], view-projection animation [47]
 - *multiple separate views:* intelligent multi-shot visualization [30], multiple inset views (IBIS) [23]
 - *multiple integrated views:* virtual multiprojection cameras [50]

E. Interaction Model

We are also interested in capturing the interaction model employed by each technique. This dimension is a cross of both

interactivity and the operation model of technique, i.e. whether the technique allows for an active or passive interaction mode when discovering targets, as well as whether it operates in an online or offline fashion when changing the target selection.

Active and online techniques require direct manipulation by the user to select targets [39], dynamically animate the viewing transform [37], or move a magic mirror to show a second view of the scene [34]. No offline recomputation is needed as users perform these operations.

Active and offline techniques, on the other hand, allow for some direct manipulation control (such as of the viewpoint), but changing the target selection requires an offline recomputation. Automatic tour generation [17], for example, lets users partially control their movement in the 3D environment but requires a precomputation step for building the grand tour based on the targets to visit.

Passive and online techniques require no explicit input from the user to expose hidden targets. Examples include image-space dynamic transparency [27] and ghosting [23], where in both cases targets are exposed using transparency whenever they are hidden from the view of the user without the user having to do anything. Changing the target selection is possible, and does not require

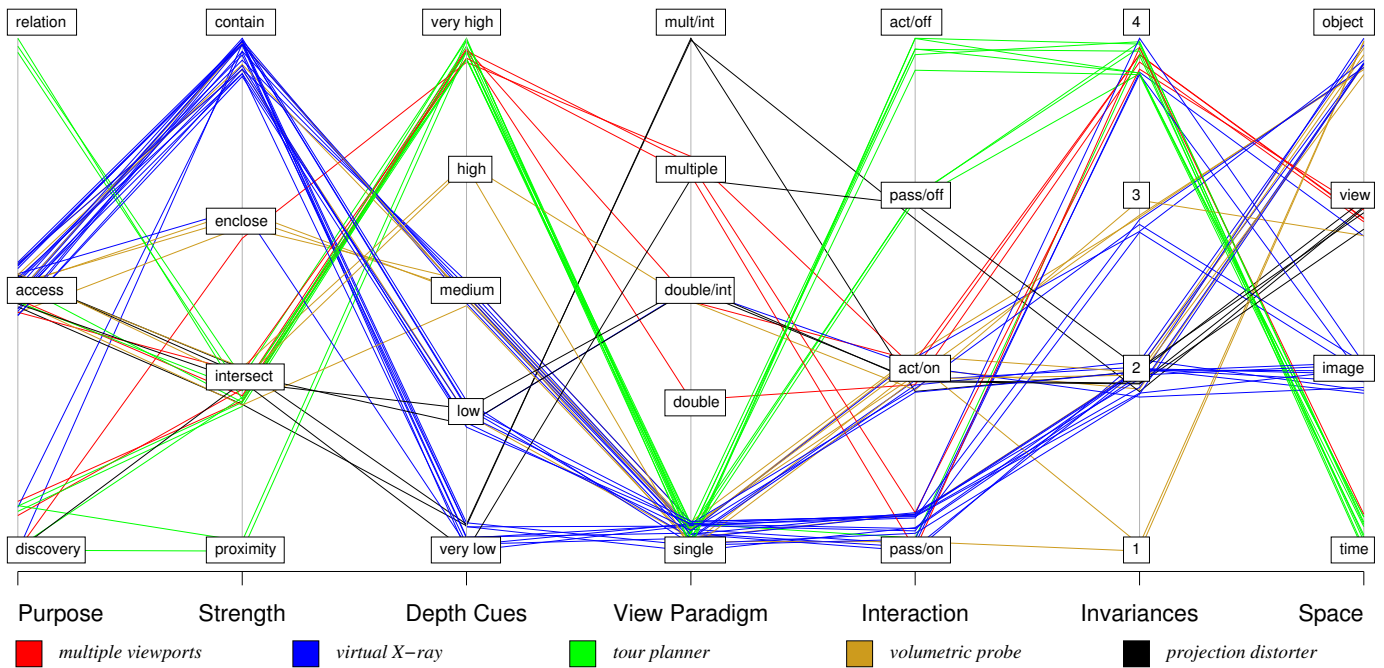


Fig. 4. Classification of 50 different occlusion management techniques using the taxonomy (points have been jittered to show distribution).

offline recomputation.

Finally, passive and offline techniques allow no direct manipulation control, and changing the target selection triggers a non-interactive recomputation step. By way of example, artistic multiprojection [16] lets the user compose views of several different viewpoints into a single image in a special phase, but the output image allows for no direct interaction (essentially because artistic multiprojection is not primarily an occlusion management technique, but rather for artistic composition of images).

One drawback of passive modes is that a technique employing such an interaction model typically must have prior semantic knowledge about the targets the user considers important. One example is the importance-driven volume rendering technique by Viola et al. [29], where each individual target in the volumetric dataset has an associated importance value. Active mode, on the other hand, puts these decisions in the hands of the user, providing for more flexible interaction. However, active control requires the user to manually discover hidden targets, which may cause them to miss important information.

- **Domain:** *passive/online*, *active/online*, *passive/offline*, *active/offline* (nominal)
- **Characteristic Techniques:**
 - *passive/online*: image-space dynamic transparency [27], object removal (IBIS) [23]
 - *active/online*: 3D virtual X-ray magic lens [14], general transfer-function volumetric lenses [25]
 - *passive/offline*: viewpoint-quality-driven scene exploration [49], wayfinder [53]
 - *active/offline*: navigation guide avatar [35], Cubical-Path [22]

F. Target Invariances

The sixth primary dimension of our taxonomy describes the number of invariances preserved by the technique. A complement

to the depth cues parameter above, target invariances describes how many of the following properties of the **targets** (not necessarily distractors) in the environment are retained:

- **Appearance (A):** Color, texture and material of the target.
- **Depth (D):** Depth information for the target (related to the depth cues dimension).
- **Geometry (G):** Shape and size of the target.
- **Location (L):** Position and orientation of the target.

All of the above properties are all more or less important for visualization applications in 3D environments. For instance, in a simple 3D scatterplot, the location of each data point is vital for the data to be interpreted correctly, so an occlusion management technique designed for use with such data should definitely preserve this property. In a color-coded 3D tree-representation of a file system, it might make sense to displace location (as long as connectivity information is retained) but the appearance should not be altered.

The higher number of invariances a technique retains, the better it is, and so this is an ordinal dimension. However, as discussed in the introduction, our normal visual cues are often at an odds with understanding various properties of an environment (e.g. seeing all the targets despite occlusion), and thus this is an example of a classical trade-off decision specific to each technique. Often, designers can gain certain attractive properties by relaxing others, all depending on the particular application area of the technique.

- **Domain:** *0-4: appearance, depth, geometry, location* (ordinal)
- **Characteristic Techniques:**
 - *appearance*: preserve: BalloonProbe [18]; discard: volumetric rendering transfer functions [52]
 - *depth*: preserve: interactive figure captions [32]; discard: X-ray tunnel [56]
 - *geometry*: preserve: importance-driven focus of attention [28]; discard: non-linear projection [16], [45]

- *location*: preserve: interactive cut-away and break-away views [31]; discard: 3D explosion probe [12]

G. Solution Space

The solution space of an occlusion management technique captures which mechanism the technique utilizes for achieving its purpose. In essence, this is the canvas used for solving the occlusion problem. While some techniques may involve combinations of different solution spaces, we are interested in the primary space for purposes of classification.

Some techniques manage occlusion by sequentially showing different parts or views of a scene in an animation or interactive exploration—these techniques make use of temporal space as the solution space. Some perform image-based composition or special effects—they employ the image space. Some manipulate the viewing transform or other properties of the virtual camera to make as many targets as possible visible to the user—the view space. Finally, still some techniques manipulate the 3D world itself, changing the geometry, positions, and even appearances of 3D objects to optimize viewing and visibility—they operate in the object space.

- **Domain:** *temporal space, image space, view space, object space* (nominal)
- **Characteristic Techniques:**
 - *temporal space*: path-planning navigation [38], task-level camera control [44]
 - *image space*: image-space dynamic transparency [27], blueprints [20]
 - *view space*: orthotumble [37], virtual multiprojection cameras [50]
 - *object space*: visual access distortion [51], SDM [41]

V. DESIGN PATTERNS

We have classified the 50 techniques involved in our survey using our taxonomy; see Table I. We then study this body of classifications to see patterns and trends, using hierarchical clustering (normalizing the data column-by-column and clustering using average linkage, see Figure 3 for a dendrogram). This analysis yields five distinct and orthogonal archetypes of design, or *design patterns* [59], i.e. a generic and reusable solution to a commonly occurring problem within a specific context. The result can be summarized in the parallel coordinate plot in Figure 4. The five patterns we have identified we call Multiple Viewports, Virtual X-Ray, Tour Planner, Volumetric Probe, and Projection Distorter. We will describe these in the following sections.

According to pattern lore, a design pattern has four essential elements: a name, a problem (occlusion management, the same for all patterns), a solution, and the consequences of using the pattern. We use these elements in our discussion of each pattern. We also show the distribution of techniques implementing the pattern on the design space and give an example picture of an instance of the pattern.

A. Multiple Viewports

The Multiple Viewports pattern (red in Figure 4) is characterized by a view paradigm based on two or more separate views, resulting in an overview+detail kind of layout. Instances of this pattern also tend to preserve most, if not all, invariances—the trick lies in the placement of the additional cameras, not

manipulating the image seen from them. It is most effective for 3D environments that lend themselves to overviews, such as landscapes and structured buildings. Furthermore, the interaction model tends to be active; no existing technique performs the automatic placement of cameras that would be necessary for passive interaction.

The pattern is widely used in 3D CAD applications to simultaneously show an object under construction from several directions. Figure 5 shows an example of the Multiple Viewports pattern in action for Blender3D, a standard 3D modeling application. See Baldonado et al. [60] for design principles of multiple view visualizations. Figure 6 shows the distribution of multiple viewport classifications in the taxonomy.

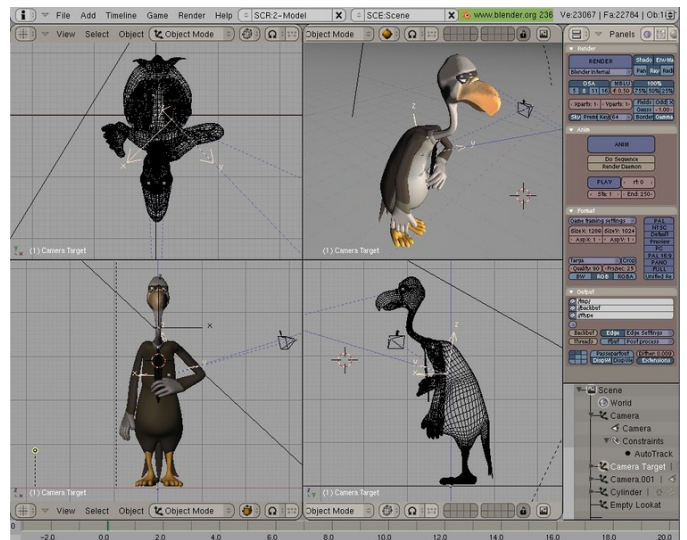


Fig. 5. Multiple viewports in the Blender3D (<http://www.blender.org/>) open source modeling program showing different views of the same 3D object.




Fig. 6. Classification distribution of multiple viewports techniques.

1) *Solution*: Manage discovery and access of targets by providing several alternate (often separate) viewports of the 3D environment. Typically, one viewport is designated as the main viewport, with the other viewports as secondary and generally smaller. Accordingly, the main viewport is often used for detail or first-person views, whereas the alternate views give either static or dynamic overviews of the environment (such as an overhead map).

2) *Consequences*: The use of the Multiple Viewports pattern trades screen estate and user attention for increased discovery and access; the user will have a smaller main visualization window than otherwise, and may have to split his or her attention

across all of the viewpoints. Furthermore, in some situations, it is not clear what constitutes an overview, and thus introducing additional viewpoints may have diminishing returns. However, this is a very powerful approach for certain types of environments (such as environments that are essentially 2D in nature and lends themselves to overhead maps).

3) *Examples:* Worlds-in-miniature [54], worldlets [55], bird's eye views [19].

B. Virtual X-Ray

The Virtual X-Ray pattern (blue in Figure 4) is (mostly) based on an image-space approach where occlusion can be easily detected and sometimes even delegated to programmable fragment shaders. While we only concern ourselves with 3D in our treatment, this pattern is also used in dynamic transparency techniques for 2D windowing systems, such as the free-space transparency [61] and multiblending [62] techniques. Typically, example techniques have very high disambiguation strength. Furthermore, there is a clear division between two types of Virtual X-Ray techniques; active ones, where the user controls a "searchlight" on the 2D view, and passive ones, where semantic information allows the system to automatically uncover targets.

Figure 7 depicts the Virtual X-Ray pattern in the guise of image-space dynamic transparency [27] showing the engine of a jeep through a semi-transparent and interactive breakaway in its hood. Figure 8 shows the distribution of virtual X-Ray classifications in the taxonomy.




Fig. 7. Image-space dynamic transparency [27], a Virtual X-Ray technique, uncovering the engine inside a jeep.




Fig. 8. Classification distribution of virtual X-Ray techniques.

1) *Solution:* Make targets visible through intervening distractors by turning occluding surfaces invisible or semi-transparent. The method for distractor removal is characteristic: some techniques are view-dependent whereas others are static; some eliminate distractors (or parts of distractors), others merely make distractors semi-transparent. Active interaction facilitates exploration whereas passive interaction requires target information but yields a potentially higher correctness.

2) *Consequences:* The Virtual X-Ray pattern makes discovery trivial and facilitates access by selectively removing distractors occluding the targets. However, this is a direct weakening of occlusion depth cues, causing a decrease in depth perception and making spatial relation more difficult. The use of semi-transparency also results in high visual complexity and imposes a high cognitive load on the user. Finally, Virtual X-Ray can make visibility computations for rendering optimization useless.

3) *Examples:* Perspective cutouts (active) [39], X-ray tunnel [56], image-space dynamic transparency (passive) [27], IBIS cutaways (passive) [23].

C. Tour Planner

The family of Tour Planner techniques (green in Figure 4) is characterized by a hybrid interaction model consisting of an offline and an online phase where paths first are defined or computed and then interactively shown in the environment itself. Typically no distortion is imposed on the view (a temporal canvas is used), so all invariances are usually retained.

Figure 9 shows the offline phase of automatic tour generation [17], an instance of the Tour Planner pattern, and Figure 10 shows the distribution of tour planner classifications in the taxonomy.




Fig. 9. Offline voxelization process for automatic tour generation [17], a Tour Planner technique, computing a grand tour that visits all landmarks in a given 3D environment.




Fig. 10. Classification distribution of tour planner techniques.

1) *Solution*: Present all targets in an environment by constructing a complete (i.e. all targets are visible in at least one point) path (or a number of paths) through it. It should also conform to a number of additional constraints (such as short or optimal length, closed, uniform visual complexity, etc). Often realized in an offline precomputation or specification step (automatic or interactively built) followed by an interactive exploration phase where the user is guided by the computed path.

2) *Consequences*: The Tour Planner pattern is non-invasive and thus will not modify the environment itself and will typically retain all invariances. This however means that the pattern's disambiguation strength is generally low. The path computation step can sometimes be costly in terms of computation time, and intractable to dynamically changing situations. In case the paths are interactively built by a human designer, there may be no completeness guarantee.

3) *Examples*: Wayfinder [53], StyleCam [43], viewpoint selection for intervention planning [48].

D. Volumetric Probe

Volumetric Probes (golden in Figure 4) manage occlusion in the object space through active user interaction in a direct manipulation approach. The probe itself is volumetric and is thus a first-class object in the environment. Typically, techniques operate by performing some kind of transformation on the contents of the probe and affecting some of its invariances. For instance, an "interactive explosion" metaphor can be adopted, meaning that target location is not retained.

Figure 13 depicts a Volumetric Probe in the shape of a spherical BalloonProbe [18] being used to separate objects in two different kinds of environments. Figure 11 shows the distribution of volumetric probe classifications in the taxonomy.




Fig. 11. Classification distribution of probe techniques.

1) *Solution*: Provide a user-controlled distortion probe that locally transforms objects to manage occlusion. The approach is based either on (i) removing distractors or (ii) separating targets; in the former case, we want to eliminate objects that get in the way, whereas in the latter, we instead want to disambiguate between several targets that share the same space. The actual transformation used depends on the technique and can range from changes to appearance, geometry, and position. The interaction is active and under direct user control in the object space.

2) *Consequences*: Using a Volumetric Probe can help disambiguate even very difficult situations, but the very nature of the pattern means that many invariances are not preserved. The pattern is best suited for discovery or access. The local influence model means that there may be a problem of reach in a virtual environment.

3) *Examples*: 3D explosion probe [12], deformation-based volume explosion [24], 3D magic lenses [14].

E. Projection Distorter

This pattern (black in Figure 4) is signified by a view-space approach presented using two or more integrated views. Since non-linear projections are typically employed to pack as many of the targets as possible into a single view, few invariances are retained. Thus, this pattern is often best used for discovery, rarely for access, and almost never for relation.

Figure 14 shows parallel and perspective views of the same 3D scene. View-projection animation [47], an instance of the Projection Distorter pattern, provides an interactive animation between these two views to help users discover occluded targets. Figure 12 shows the distribution of projection distorter classifications in the taxonomy.




Fig. 12. Classification distribution of projection distorter techniques.

1) *Solution*: Integrate several different views of targets into a single view in order to maximize discovery. The solution is then often reminiscent of a focus+context technique with one focus per view. Individual view selection is often actively controlled by the user in an online or offline manner. In one case, a hybrid approach is employed where target semantic information is extracted from previous user explorations using data mining techniques and then used to inform the technique [45].

2) *Consequences*: The use of the Projection Distorter pattern affects only the view projection code of an application and is thus relatively easy to integrate into existing code. On the other hand, the resulting visual displays can often become disconcerting and disorienting to the user. Few object properties are retained.

3) *Examples*: Artistic multiprojection [16], view-projection animation [47], orthotumble [37].

VI. DISCUSSION

In this section, we discuss the design of the taxonomy, our criteria for selecting and classifying techniques into the taxonomy, and some of its limitations. We also study the current state of the art in order to identify possible future research directions.

A. Taxonomy Design

The taxonomy presented in this paper has been designed to be orthogonal and objective, with no dimension being reducible to another and having a minimum of coupling to the other dimensions. In our dimension selection process, we strived to find a minimal set of characteristic dimensions through discussions among ourselves and with external researchers, studies of the

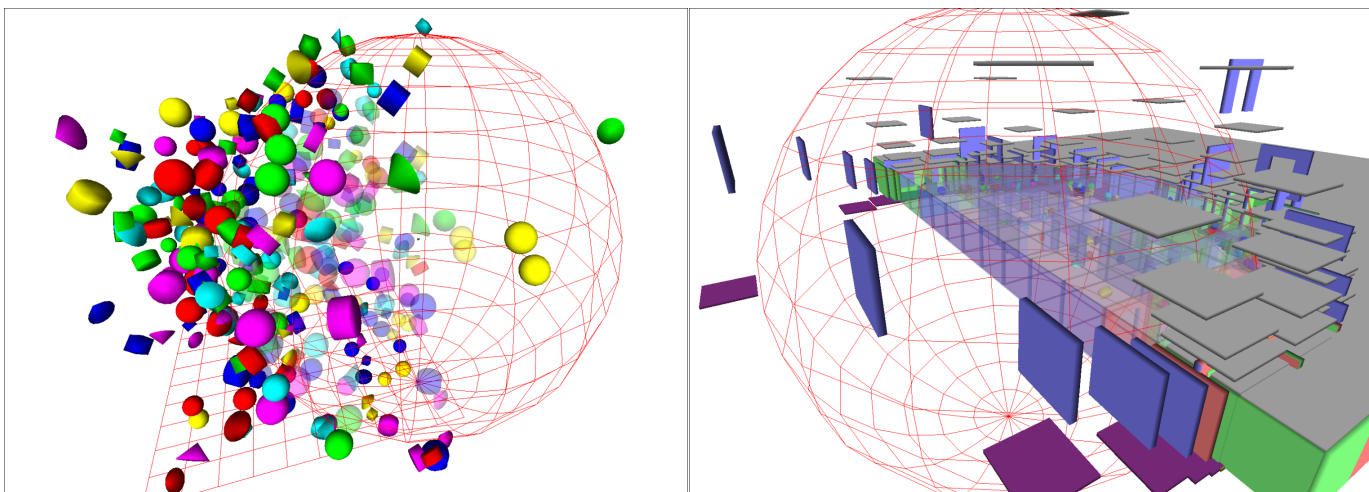


Fig. 13. Spherical BalloonProbes [18], a Volumetric Probe technique, separating objects in abstract and building walkthrough environments.

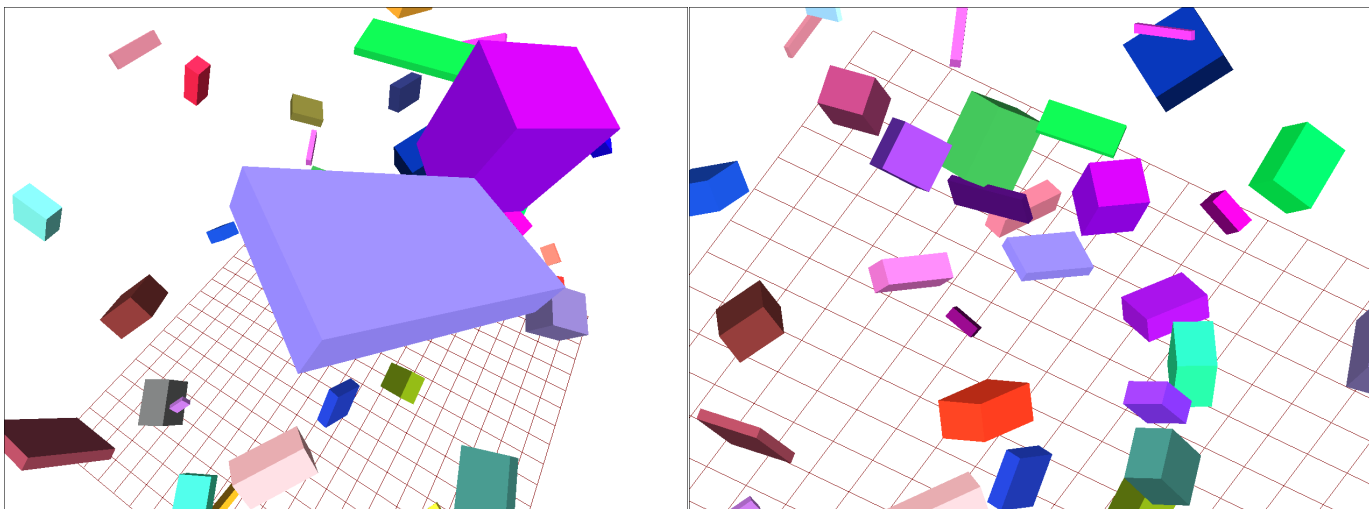


Fig. 14. Perspective and parallel views of the same scene for view-projection animation [47], a Projection Distorter technique.

literature, as well as through analysis. In particular, the Solution Space dimension was added after the conference version of this paper because clustering showed that this dimension seemed to separate the patterns in the original classification. Furthermore, the Interaction Model dimension was substantially revised following feedback from the research community.

Regardless, it is always possible to debate the inclusion or exclusion of specific property dimensions to a taxonomy. We believe this to be a valid one, and the successful classification of 50 different techniques using it confirms this claim.

Nevertheless, property dimensions that were excluded for various reasons include scalability (the amount of object density the technique can handle), influence level (i.e. whether the interaction technique operates on a local, regional, or global level), and dimensionality (2D, 2.5D, 3D, etc). Scalability has been pinpointed by some people to be a particularly useful property that would surely help designers in choosing an appropriate technique given a specific problem. The problem here was the lack of a neutral benchmark for such an objective metric such as scalability, not to mention that we did not have access to actual implementations of each technique, and in the end this caused us to omit it from the taxonomy. This lack of implementations is also what prevented

us from measuring performance metrics for the techniques.

B. Technique Selection and Delimitations

Some 200 scientific papers were perused while performing the classification process, and out of these techniques, we selected only 50 to classify into the taxonomy presented in this paper. The selection was guided by the delimitations stated in Section IV of this paper; we included only techniques that had the following properties:

- deals with 3D data;
- has an interactive component;
- aims at visualization of data (at least in some sense); and
- is concerned with making information-carrying targets discoverable, accessible, and/or relatable despite 3D visibility constraints.

While we believe that our selection reflects these properties, there may be instances that can be argued to be excluded or included in the classification. For the former case, we attempted to include some techniques that were not strictly concerned with occlusion management with the purpose of making the classification a little broader and open to non-traditional approaches. For

the latter, no taxonomy or classification is ever complete, and we encourage others to follow up on our work and help us fill in the missing data.

In particular, a number of 3D navigation techniques were indeed included in the classification in this paper. These were selected due to them being on the borderline of what constitutes an occlusion management technique, or representative for a specific class of techniques. Many other 3D navigation techniques in the literature were excluded from this classification; the line had to be drawn somewhere.

C. Limitations of the Taxonomy

In line with the statement expressed above, that being that no taxonomy is ever complete, there may also be structural problems with the taxonomy that could be modified. In particular, the taxonomy is not able to capture some aspects of an occlusion management technique, such as the means of selecting views for a multiple viewports technique, or whether a tour planner technique makes use of automatic or user-built paths. Furthermore, looking at the dendrogram in Figure 3 or the classification in Table I, we can see that a few techniques are classified exactly the same, when they certainly are different in at least some aspects. Addressing these points should be one focus of future improvements to the taxonomy.

The clustering of techniques is mechanical, but stems directly from their individual classification. It is certainly possible to find problems or oddities that may need future attention. For instance, the virtual X-ray pattern currently includes at least three labeling techniques (3D view management [13], consistent 3D labels [21], and interactive figure captions [32]) that do not really fit under this pattern. Other singularities include the camera-control techniques in the tour planner and multiple views patterns (StyleCam [43] and intelligent multi-shot visualization [30]). While it may be argued that these problems stem from wrong classification or should be excluded entirely, they may also be indications of limited expressivity of the taxonomy in some aspects.

D. Future Research Directions

Besides identifying existing design archetypes in the literature, we can also extract possible future research opportunities from our taxonomy by studying the as-of-yet unexplored parts of the design space. As can be seen from Figure 15, most of the design space has been covered by the existing techniques. However, the question is if we can find combinations of properties that would be particularly fruitful for future research.

One such combination would be techniques that make users aware of occluded content without compromising visual quality and imposing a high cognitive load on the user. Retaining a high degree of depth cues is important for complex visual tasks such as spatial relation. Another interesting area to explore is hybrid-interaction methods where the user's own actions are used to inform the target selection. This approach may help solve the trade-off between the precision that a passive interaction model provides as opposed to the more general nature of active user interaction.

Combinations of patterns could be profitable ways of utilizing the strong points of two different methods while at the same time making up for the weak ones. For example, a multiple-viewport technique could be augmented with virtual X-Ray support in one

or several of the views. A tour planner could be paired with a volumetric probe to help disambiguate in difficult situations of locally high target congestion.

The field is also open for using this taxonomy to introduce variations to already established and well-researched design patterns. In particular, the multiple viewports design pattern is well-known and extensively used in 3D visualization. Looking at its distribution (Figure 6), we can see that there is a fairly well-defined notion of a multiple viewports technique is. However, perhaps there may be ways of improving their disambiguation strength by clever (automatic or user-guided) selection of *which* viewports to include.

There is one additional point to consider: General 3D navigation has been shown to be a task with a very high cognitive load; for every traveled world unit, the user runs the risk of becoming disoriented, totally lost, or even nauseous. For the longer term, it can be noted that the ultimate goal of occlusion management techniques should be to help minimize the need for 3D navigation in general. Perhaps the class of interaction techniques described in this paper can help short-circuit excessive navigation in the first place.

VII. CONCLUSIONS

Occlusion management for visualization is a subset of 3D interaction techniques concerned with improving human perception for specialized visual tasks through manipulation of visual cues such as occlusion, size, and shape. In this paper, we have presented five archetypical design patterns for occlusion management based on a classification of existing interaction techniques. The patterns include multiple viewports, virtual X-ray, tour planners, volumetric probes, and projection distorters. The underlying taxonomy used for this classification is based on seven characteristic properties of occlusion management techniques. Analysis of this taxonomy also yields additional missing patterns, such as primarily techniques for target awareness and hybrid-interaction approaches with an emphasis on retaining a high degree of depth cues and supporting spatial relation.

ACKNOWLEDGMENT

The authors would like to thank the anonymous reviewers and participants of the IEEE Virtual Reality 2007 conference for their feedback and comments on this taxonomy. Many thanks to the reviewers who offered feedback on the extended journal version of this article. Thank you also to the numerous people who have offered their insight and opinions on this work during the course of the project.

REFERENCES

- [1] N. Elmqvist and P. Tsigas, "A taxonomy of 3D occlusion management techniques," in *Proceedings of the IEEE Conference on Virtual Reality*, 2007, pp. 51–58.
- [2] D. Cohen-Or, Y. L. Chrysanthou, C. T. Silva, and F. Durand, "A survey of visibility for walkthrough applications," *IEEE Transactions on Visualization and Computer Graphics*, vol. 9, no. 3, pp. 412–431, July/Sept. 2003.
- [3] J. Bittner and P. Wonka, "Visibility in computer graphics," *Environment and Planning B: Planning and Design*, vol. 30, no. 5, pp. 729–756, Sept. 2003.
- [4] M. Berg, M. Kreveld, M. Overmars, and O. Schwarzkopf, *Computational Geometry: Algorithms and Applications*. Springer-Verlag, 1997.




Fig. 15. Distribution of all of the design patterns in the taxonomy. Only outlines are shown to prevent 2D occlusion.

- [5] D. A. Bowman and L. F. Hodges, "Formalizing the design, evaluation, and application of interaction techniques for immersive virtual environments," *Journal of Visual Languages and Computing*, vol. 10, no. 1, pp. 37–53, 1999.
- [6] E. Bier, M. Stone, K. Fishkin, W. Buxton, and T. Baudel, "A taxonomy of see-through tools," in *Proceedings of the ACM CHI'94 Conference on Human Factors in Computing Systems*, 1994, pp. 358–364.
- [7] D. A. Bowman, C. North, J. Chen, N. F. Polys, P. S. Pyla, and U. Yilmaz, "Information-rich virtual environments: theory, tools, and research agenda," in *Proceedings of the ACM Symposium on Virtual Reality Software and Technology*, 2003, pp. 81–90.
- [8] Z. Pousman and J. Stasko, "A taxonomy of ambient information systems: Four patterns of design," in *Proceedings of the ACM Conference on Advanced Visual Interfaces*, 2006, pp. 67–74.
- [9] S. Card, J. Mackinlay, and G. Robertson, "The design space of input devices," in *Proceedings of the ACM CHI'90 Conference on Human Factors in Computing Systems*, 1990, pp. 117–124.
- [10] N. F. Polys and D. A. Bowman, "Design and display of enhancing information in desktop information-rich virtual environments: challenges and techniques," *Virtual Reality*, vol. 8, no. 1, pp. 41–54, 2004.
- [11] K. Hinckley, R. Pausch, J. C. Goble, and N. F. Kassell, "Passive real-world interface props for neurosurgical visualization," in *Proceedings of the ACM CHI'94 Conference on Human Factors in Computing Systems*, 1994, pp. 452–458.
- [12] H. Sonnet, M. S. T. Carpendale, and T. Strothotte, "Integrating expanding annotations with a 3D explosion probe," in *Proceedings of the ACM Conference on Advanced Visual Interfaces*, 2004, pp. 63–70.
- [13] B. Bell, S. Feiner, and T. Höllerer, "View management for virtual and augmented reality," in *Proceedings of the ACM Symposium on User Interface Software and Technology*, 2001, pp. 101–110.
- [14] J. Viega, M. J. Conway, G. Williams, and R. Pausch, "3D magic lenses," in *Proceedings of the ACM Symposium on User Interface Software and Technology*, 1996, pp. 51–58.
- [15] A. Raab and M. Rüger, "3D-Zoom: Interactive Visualisation of Structures and Relations in Complex Graphics," in *3D Image Analysis and Synthesis*, B. Girod, H. Niemann, and H.-P. Seidel, Eds. infix-Verlag, 1996, pp. 125–132.
- [16] M. Agrawala, D. Zorin, and T. Munzner, "Artistic multiprojection rendering," in *Proceedings of the Eurographics Workshop on Rendering Techniques*, 2000, pp. 125–136.
- [17] N. Elmqvist, M. E. Tudoreanu, and P. Tsigas, "Tour generation for exploration of 3D virtual environments," in *Proceedings of the ACM Symposium on Virtual Reality Software and Technology*, 2007, to appear.
- [18] N. Elmqvist and M. E. Tudoreanu, "Occlusion management in immersive and desktop 3d virtual environments: Theory and evaluation," *International Journal of Virtual Reality*, vol. 6, pp. 21–32, 2007.
- [19] S. Fukatsu, Y. Kitamura, T. Masaki, and F. Kishino, "Intuitive control of "bird's eye" overview images for navigation in an enormous virtual environment," in *Proceedings of the ACM Symposium on Virtual Reality Software and Technology*, 1998, pp. 67–76.
- [20] M. Nienhaus and J. Döllner, "Blueprints: Illustrating architecture and technical parts using hardware-accelerated non-photorealistic rendering," in *Proceedings of Graphics Interface*, 2004, pp. 49–56.
- [21] B. Preim, A. Ritter, T. Strothotte, T. Pohle, D. R. Forsey, and L. Bartram, "Consistency of Rendered Images and Their Textual Labels," in *Proceedings of CompuGraphics*, 1995, pp. 201–210.
- [22] S. Beckhaus, F. Ritter, and T. Strothotte, "Guided exploration with dynamic potential fields: the cubical path system," *Computer Graphics Forum*, vol. 20, no. 4, pp. 201–210, Dec. 2001.
- [23] S. Feiner and D. D. Seligmann, "Cutaways and ghosting: satisfying visibility constraints in dynamic 3D illustrations," *The Visual Computer*, vol. 8, no. 5-6, pp. 292–302, 1992.
- [24] M. J. McGuffin, L. Tancu, and R. Balakrishnan, "Using deformations for browsing volumetric data," in *Proceedings of the IEEE Conference on Visualization*, 2003, pp. 401–408.
- [25] T. Ropinski and K. H. Hinrichs, "Interactive volume visualization techniques for subsurface data," in *Proceedings of the International Conference on Visual Information Systems*, 2005, pp. 121–131.
- [26] F. Ritter, H. Sonnet, K. Hartmann, and T. Strothotte, "Illustrative Shadows: Integrating 3D and 2D Information Displays," in *Proceedings of the International Conference on Intelligent User Interfaces*, 2003, pp. 166–173.
- [27] N. Elmqvist, U. Assarsson, and P. Tsigas, "Employing dynamic transparency for 3D occlusion management: Design issues and evaluation," in *Proceedings of INTERACT*, ser. LNCS, C. Baranauskas, P. Palanque, J. Abascal, and S. D. J. Barbosa, Eds., vol. 4662. Springer, 2007, pp. 532–545.
- [28] I. Viola, M. Feixas, M. Sbert, and M. E. Gröller, "Importance-driven focus of attention," *IEEE Transactions on Visualization and Computer Graphics*, vol. 12, no. 5, pp. 933–940, 2006.
- [29] I. Viola, A. Kanitsar, and E. Gröller, "Importance-driven volume rendering," in *Proceedings of the IEEE Conference on Visualization*, 2004, pp. 139–145.
- [30] W. H. Bares and J. C. Lester, "Intelligent multi-shot visualization interfaces for dynamic 3D worlds," in *Proceedings of the International Conference on Intelligent User Interfaces*, 1999, pp. 119–126.
- [31] J. Diepstraten, D. Weiskopf, and T. Ertl, "Interactive cutaway rendering," in *Proceedings of Eurographics*, 2003, pp. 523–532.
- [32] B. Preim, R. Michel, K. Hartmann, and T. Strothotte, "Figure Captions in Visual Interfaces," in *Proceedings of the ACM Workshop on Advanced Visual Interfaces*, 1998, pp. 235–246.
- [33] J. Looser, M. Billingham, and A. Cockburn, "Through the looking glass:

- the use of lenses as an interface tool for augmented reality interfaces,” in *Proceedings of GRAPHITE*, 2004, pp. 204–211.
- [34] J. Grosjean and S. Coquillart, “The magic mirror: A metaphor for assisting the exploration of virtual worlds,” in *Proceedings of the Spring Conference on Computer Graphics*, 1999, pp. 125–129.
- [35] L. Chittaro, R. Ranon, and L. Ieronutti, “Guiding visitors of Web3D worlds through automatically generated tours,” in *Proceedings of the ACM Conference on 3D Web Technology*, 2003, pp. 27–38.
- [36] S. Takahashi, K. Yoshida, K. Shimada, and T. Nishita, “Occlusion-free animation of driving routes for car navigation systems,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 12, no. 5, pp. 1141–1148, 2006.
- [37] T. Grossman, R. Balakrishnan, G. Kurtenbach, G. Fitzmaurice, A. Khan, and B. Buxton, “Creating principal 3D curves with digital tape drawing,” in *Proceedings of the ACM CHI 2002 Conference on Human Factors in Computing Systems*, 2002, pp. 121–128.
- [38] B. Salomon, M. Garber, M. C. Lin, and D. Manocha, “Interactive navigation in complex environments using path planning,” in *Proceedings of the ACM Symposium on Interactive 3D Graphics*, 2003, pp. 41–50.
- [39] C. Coffin and T. Höllerer, “Interactive perspective cut-away views for general 3D scenes,” in *Proceedings of the IEEE Symposium on 3D User Interfaces*, 2006, pp. 25–28.
- [40] S. Desroche, V. Jolivet, and D. Plemenos, “Towards plan-based automatic exploration of virtual worlds,” in *Proceedings of the International WSCG Conference*, 2007.
- [41] M. C. Chuah, S. F. Roth, J. Mattis, and J. Kolojehchick, “SDM: Selective dynamic manipulation of visualizations,” in *Proceedings of the ACM Symposium on User Interface Software and Technology*, 1995, pp. 61–70.
- [42] L. Chittaro and I. Scagnetto, “Is semitransparency useful for navigating virtual environments?” in *Proceedings of the ACM Symposium on Virtual Reality Software and Technology*, 2001, pp. 159–166.
- [43] N. Burtnyk, A. Khan, G. Fitzmaurice, R. Balakrishnan, and G. Kurtenbach, “StyleCam: interactive stylized 3D navigation using integrated spatial & temporal controls,” in *Proceedings of the ACM Symposium on User Interface Software and Technology*, 2002, pp. 101–110.
- [44] S. M. Drucker and D. Zeltzer, “Intelligent camera control in a virtual environment,” in *Proceedings of Graphics Interface*, 1994, pp. 190–199.
- [45] K. Singh and R. Balakrishnan, “Visualizing 3D scenes using non-linear projections and data mining of previous camera movements,” in *Proceedings of AFRIGRAPH*, 2004, pp. 41–48.
- [46] J. Diepstraten, D. Weiskopf, and T. Ertl, “Transparency in interactive technical illustrations,” *Computer Graphics Forum*, vol. 21, no. 3, pp. 317–325, 2002.
- [47] N. Elmqvist and P. Tsigas, “View-projection animation for 3D occlusion management,” *Computers and Graphics*, 2007, to appear.
- [48] K. Muehler, M. Neugebauer, C. Tietjen, and B. Preim, “Viewpoint selection for intervention planning,” in *Proceedings of the IEEE/Eurographics Symposium on Visualization*, 2007, pp. 267–274.
- [49] D. Sokolov, D. Plemenos, and K. Tamine, “Viewpoint quality and global scene exploration strategies,” in *Proceedings of the International Conference on Computer Graphics Theory and Applications*, 2006, pp. 184–191.
- [50] K. Singh, “A fresh perspective,” in *Proceedings of Graphics Interface*, 2002, pp. 17–24.
- [51] M. S. T. Carpendale, D. J. Cowperthwaite, and F. D. Fracchia, “Distortion viewing techniques for 3D data,” in *Proceedings of the IEEE Symposium on Information Visualization*, 1996, pp. 46–53.
- [52] R. A. Drebin, L. Carpenter, and P. Hanrahan, “Volume rendering,” in *Computer Graphics (SIGGRAPH '88 Proceedings)*, vol. 22, Aug. 1988, pp. 65–74.
- [53] C. Andújar, P.-P. Vázquez, and M. Fairén, “Way-finder: guided tours through complex walkthrough models,” in *Proceedings of Eurographics*, 2004, pp. 499–508.
- [54] R. Stoakley, M. J. Conway, and R. Pausch, “Virtual Reality on a WIM: Interactive worlds in miniature,” in *Proceedings of the ACM CHI '95 Conference on Human Factors in Computing Systems*, 1995, pp. 265–272.
- [55] T. T. Elvins, D. R. Nadeau, and D. Kirsh, “Worldlets – 3D thumbnails for wayfinding in virtual environments,” in *Proceedings of the ACM Symposium on User Interface Software and Technology*, 1997, pp. 21–30.
- [56] R. Bane and T. Höllerer, “Interactive tools for virtual X-ray vision in mobile augmented reality,” in *Proceedings of the IEEE and ACM International Symposium on Mixed and Augmented Reality*, 2004, pp. 231–239.
- [57] C. Ware and S. Osborne, “Exploration and virtual camera control in virtual three dimensional environments,” *Computer Graphics*, vol. 24, no. 2, pp. 175–183, 1990.
- [58] G. W. Furnas, “Generalized fisheye views,” in *Proceedings of the ACM CHI '86 Conference on Human Factors in Computer Systems*, 1986, pp. 16–23.
- [59] C. Alexander, S. Ishikawa, M. Silverstein, M. Jacobson, I. Fiksdahl-King, and S. Angel, *A Pattern Language*. Oxford University Press, 1977.
- [60] M. Q. W. Baldonado, A. Woodruff, and A. Kuchinsky, “Guidelines for using multiple views in information visualization,” in *Proceedings of the ACM Conference on Advanced Visual Interfaces*, 2000, pp. 110–119.
- [61] E. W. Ishak and S. K. Feiner, “Interacting with hidden content using content-aware free-space transparency,” in *Proceedings of the ACM Symposium on User Interface Software and Technology*, 2004, pp. 189–192.
- [62] P. Baudisch and C. Gutwin, “Multiblending: displaying overlapping windows simultaneously without the drawbacks of alpha blending,” in *Proceedings of the ACM CHI 2004 Conference on Human Factors in Computing Systems*, 2004, pp. 367–374.



Niklas Elmqvist is a postdoctoral researcher in the AVIZ team of INRIA Saclay - Île-de-France located at Université Paris-Sud in Paris, France. Previously he was a Ph.D. student at the Department of Computer Science and Engineering at Chalmers University of Technology in Gothenburg, Sweden, where he received his Ph.D. in December 2006. His research specialization is in information visualization, human-computer interaction, and visual analytics. He is a member of the IEEE and the IEEE Computer Society.



Philippos Tsigas is an Associate Professor in the Department of Computer Science and Engineering at Chalmers University of Technology in Göteborg, Sweden. He is the scientific leader of the Distributed Computing and System research group and his research interests include distributed and parallel computing and systems, as well as general information visualization. He is a member of the IEEE and the IEEE Computer Society.