



Formalizing Visibility Characteristics in Hierarchical Systems

Debmalya Biswas, Krishnamurthy Vidyasankar

► To cite this version:

Debmalya Biswas, Krishnamurthy Vidyasankar. Formalizing Visibility Characteristics in Hierarchical Systems. [Research Report] 2007. inria-00156525v1

HAL Id: inria-00156525

<https://inria.hal.science/inria-00156525v1>

Submitted on 21 Jun 2007 (v1), last revised 21 Jun 2007 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

Formalizing Visibility Characteristics in Hierarchical Systems

Debmalya Biswas — Krishnamurthy Vidyasankar

N° ????

June 2007

Thème COM

 ***apport
de recherche***



Formalizing Visibility Characteristics in Hierarchical Systems

Debmalya Biswas*, Krishnamurthy Vidyasankar†

Thème COM — Systèmes communicants
Projets DISTRIBCOM

Rapport de recherche n° 1000 — June 2007 — 34 pages

Abstract: We consider hierarchical systems where nodes represent entities and edges represent binary relationships among them. An example is a hierarchical composition of Web services where the nodes denote services and edges represent the parent-child relationship of a service invoking another service. A fundamental issue to address in such systems is, for two nodes X and Y in the hierarchy whether X can see Y , that is, whether X has visibility over Y . In a general setting, X seeing Y may depend on (i) X wishing to see Y , (ii) Y wishing to be seen by X , and (iii) other nodes not objecting to X seeing Y . The visibility could be with respect to certain attributes like operational details, execution logs, security related issues, etc. In this paper, we develop a generic conceptual model to express visibility. We study two complementary notions: *sphere of visibility* of a node X that includes all the nodes in the hierarchy that X can see; and *sphere of noticeability* of X that includes all the nodes that can see X . We also identify dual properties, *coherence* and *correlation*, that relate the visibility and noticeability notions, and study their variants.

Key-words: Visibility, Noticeability, Hierarchical Systems, Coherence, Correlation, Visibility Policies.

This work is an extended version of [3].

* IRISA, Campus de Beaulieu, Rennes Cedex, France, 35042, dbiswas@irisa.fr

† Dept. of Computer Science, Memorial University of Newfoundland, St. John's, NL, Canada A1B 3X5, vidya@cs.mun.ca

Formalisation de caractéristiques de visibilité dans des systèmes hiérarchiques

Résumé : Nous considérons des systèmes hiérarchiques où les noeuds représentent des entités, et les arêtes des relations binaires entre ces entités; par exemple, une composition hiérarchique de services Web, où les noeuds correspondent aux services, et où les arêtes représentent la relation père-fils d'un service en invoquant un autre. Un problème fondamental à traiter dans de tels systèmes est de savoir, pour deux noeuds X et Y de la hiérarchie, si X peut voir Y , c'est-à-dire si X a une certaine visibilité sur Y . Dans le cas général, le fait que X voie Y peut dépendre (i) du souhait de X de voir Y , (ii) du souhait de Y d'être vu par X , et (iii) de l'absence de refus d'autres noeuds à ce que X voie Y . La visibilité peut être relative à certains attributs comme les détails opérationnels, les logs d'exécution, des problèmes de sécurité, etc. Dans cet article, nous développons un modèle conceptuel générique pour exprimer la visibilité. Nous étudions deux notions complémentaires : la *sphère de visibilité* d'un noeud X , qui inclut tous les noeuds de la hiérarchie que X peut voir, et la *sphère de perceptibilité* de X , qui contient tous les noeuds pouvant voir X . Nous identifions également des propriétés duales, la cohérence et la corrélation, qui relient les notions de visibilité et de perceptibilité, et nous en étudions des variantes. Ces propriétés permettent de définir des motifs intéressants et utiles de visibilité et de perceptibilité, ainsi que leurs représentations.

Mots-clés : visibilité, perceptibilité, systèmes hiérarchiques, cohérence, corrélation, politiques de perception.

1 Introduction

Hierarchical systems are prevalent everywhere. While hierarchical systems provide an elegant mechanism to analyze the system functionality at different levels of abstraction, most of them allow interaction only between adjacent (parent-child) layers. Such restricted means of communication are often not sufficient for real-life scenarios. For example, in a supply chain management system, [5] states the need for visibility across levels as follows: “The information required by downstream entities are mainly material and capacity availability information from their suppliers. The information acquired by an upstream entity is information about customer demand and orders. The depth of information penetration can be specified in various degrees, e.g., isolated, upward one tier, upward two tiers, downward one tier, downward two tiers, and so forth” Also, non-functional aspects such as transactions, monitoring, user-interaction, etc. call for nodes having visibility over their ancestors, descendents and siblings. On the other hand, allowing arbitrary interaction among the hierarchical entities, without any restrictions, may not be an acceptable solution either. In a dynamic and heterogeneous environment, issues such as trust and autonomy force an entity to be selective in the interactions it has with others. For example, [7] considers the hierarchical modeling of hypertexts in the WWW. The WWW can be modeled as a large graph, the pages being the nodes and the hyperlinks being the edges. With respect to a site S , its corresponding graph consists of the pages provided by its departments $d_1, d_2, \dots, d_n$ and the hyperlinks amongst them. Given this, the graph can be hierarchically decomposed as follows: (i) By department: The pages (and corresponding hyperlinks) belonging to a department d are grouped together. (ii) By page structure: Here, we consider a page as an abstraction on its own. This can be useful for an author who wants to change a page while having an overview of its structure. With the above representation, one can also imagine the need for specific nodes to be hidden or visible with respect to a given grouping. The need for such an abstraction has been identified by [7] as follows: “A hierarchical graph data model must allow the representation of this kind of information (hierarchy, visibility), and enforce possible constraints that derive from it (for example, forbidden edges to hidden nodes).” The situation is worse for large scale systems where the number of involved entities may be in hundreds. We encountered the above issues while studying hierarchical Web Services compositions [2]. However, we believe that the proposed solutions are applicable for most hierarchical systems in general.

Thus, we need a model to capture the “visibility” aspect. For a pair of nodes X and Y in the hierarchy, we would like to capture whether X can see Y , that is, whether X has visibility over Y . In a general setting, X has visibility over Y if

- X wishes to see Y : X may be interested in Y due to functional or non-functional requirements.
- Y does not have any objection to X seeing it: As mentioned earlier, security, privacy, confidentiality, etc. issues play an important role in determining the visibility allowed by a provider.

- Remaining nodes in the hierarchy do not have any objections to X seeing Y : Contractual agreements between Y and another node Z may have a bearing on X seeing Y .

The requested visibility is usually with respect to some attribute of the nodes in the hierarchy. Examples of possible attributes (in a Web Services context) are: provider details (URI, physical address), service details (inputs, outputs, pre-conditions and effects) and execution details (execution state, history). Roughly, provider details are required to invoke an operation of the provider, service and execution details for non-functional aspects such as recovery, monitoring, auditing, and others.

The main contribution of this paper is to propose a generic conceptual model to express visibility. Towards this end, we introduce the complementary notions of (i) *Sphere of Visibility* (SoV), and (ii) *Sphere of Noticeability* (SoN). For a node X , SoV of X reflects X 's visibility over others, while SoN of X reflects the visibility others have over X . We also identify dual properties, *coherence* and *correlation*, that relate the visibility and noticeability notions, and study their variants. These properties give rise to interesting and useful visibility and noticeability assignments, and their representations.

The rest of the paper is organized as follows: In section 2, we provide an informal introduction to SoV with the help of an e-shopping scenario. The formal definitions of SoV, coherence and correlation properties, and their variants, are given in Section 3. Section 4 discusses global visibility assignment policies. Section 5 introduces SoN and brings out the relationship between the visibility and noticeability notions. Section 6 considers a special visibility assignment. Section 7 discusses some related work and section 8 concludes the paper.

2 An Informal Introduction to SoV

The Sphere of Visibility (SoV) of a node X (SoV_X) consists of nodes visible to X in the hierarchy. The visibility is with respect to some attribute A such as provider, service and execution details (discussed earlier). The visibilities of a provider corresponding to different attributes are independent, that is, a provider Y might have visibility over X 's provider details only, service details only, execution details only or any combination of the above. The visibility of each provider may be determined based on (a) the functional and non-functional requirements of the provider, (b) security, privacy and anonymity characteristics of the providers whose visibility is sought, and (c) the global policies and constraints imposed by the environment.

For example, let us consider an e-shopping scenario (Fig. 1). A customer U orders a few goods from a store S . S splits the order into two parts and sends them to suppliers $S-A$ and $S-B$. Supplier $S-B$ uses supplier $S-C$ to fulfill part of the order. $S-A$ and $S-B$ use courier companies $C-A$ and $C-B$ respectively to ship the goods to the customer. The store uses a financial service P for processing payment for the goods. This involves charging a credit card, by the credit card company H , and awarding bonus air miles, by another service B . The store also uses a monitor/auditor M to keep track of the service execution.





Figure 1: A hierarchical composition graph H corresponding to an e-shopping scenario.

Taking the attribute A as service details, the visibilities of some of the providers over other providers in the hierarchy of Fig. 1 may be as follows, shown in Fig. 2. In the illustration of visibility of X over Y , X is represented in double ovals, Y is represented in thick oval, and the other nodes, if any, are represented in thin ovals.

1. The store S has visibility over its parent and all its children. It does not have visibility over the next level descendants (Fig. 2(a)).

2. The bonus air miles processing unit B has visibility over only the credit card company H and the customer U (Fig. 2(b)). It is only concerned with the customer's credit card number and the purchase amount *without any need to know the context*, namely the goods purchased and the store. We call the visibilities of B over H and U as *weak visibilities* (or *weak references*), whereas the visibilities of S over U , $S-A$, $S-B$, M and P , described above, are referred to as *strong visibilities* (or *strong references*), meaning that the “structures” of the nodes $S-A$, $S-B$, M and P , relative to S , in the hierarchy are also visible to S .
3. The courier company $C-A$ has strong visibility over supplier $S-A$, and weak visibility over customer U , to get the address details (Fig. 2(c)).
4. The courier company $C-B$ has strong visibility over $S-B$, S and U (Fig. 2(d)).
5. The visibilities of U are described in Fig. 2(e): (i) the visibility over S is strong; (ii) the visibility over B is weak; and (iii) the visibilities over $S-A$ and H are of “intermediate strength”; we call this *partially strong visibility* (or *reference*). The partially strong visibility of U over H may be interpreted, for the service details attribute, as: U gets the service details of H directly (by weak reference to H), or from P (by weak reference to P). Note that U does *not* have visibility over P . Thus, U can get the service details of H from P , but not the service details of P itself.
6. Some non-hierarchical (horizontal) visibilities are illustrated in Fig. 2(b) (B over H) and Fig. 2(f) ($S-A$ over $S-B$, $S-C$, and M).
7. Visibilities need not be symmetric. For example, S does not have visibility over $C-B$ (Fig. 2(a)), perhaps $S-B$ wishing to hide the details of $C-B$ from S , whereas $C-B$ has visibility over S (Fig. 2(d)).
8. Visibilities of the providers in the hierarchy need not be related. For example, U has weak visibility over B (Fig. 2(e)) to ensure that the bonus air miles are credited, but U 's child S does not have visibility over B (Fig. 2(a)).

Figure 2: SoV's of some of the providers in hierarchy H of Fig. 1.

3 Formal SoV

In the sphere of visibility of X , we identify the nodes visible to X and their “type” of visibility. First, we introduce some terminology. We consider a hierarchy H as an undirected tree. For any graph G , the set of nodes in G is denoted as $V(G)$, and the set of edges as $E(G)$. For nodes X and Y in H , $H[X, Y]$ denotes the subgraph of H consisting of the nodes and edges in the path from X to Y . The (unique) path from X to Y in the hierarchy is denoted (X, Y) , and also by the sequence of nodes from X to Y , $(X, \dots, Y)$, where such an expanded form is required. Throughout this paper, we refer to a generic *visibility assignment* $\mathcal{V}$ in H with respect to an attribute A . $\mathcal{V}$ consists of the set of subgraphs $\mathcal{V}[X, Y]$, for all pairs X, Y of nodes in H , defined as follows: $\mathcal{V}[X, Y]$ is either (i) a connected subgraph of $H[X, Y]$ that contains Y , or (ii) the null graph, meaning that X does not have visibility over Y . $\mathcal{V}[X, Y]$ denotes the *type*, also *strength*, of visibility X has over Y . We assume that $\mathcal{V}[X, X]$, for every X , is the graph containing just the node X .

Definition 3.1 (Sphere of Visibility) The *Sphere of Visibility* of a node X in hierarchy H , denoted SoV_X , is the set of nonnull subgraphs $\mathcal{V}[X, Y]$, for Y in H . The set of nodes Y for which $\mathcal{V}[X, Y]$ is nonnull is denoted as V_X .

We say that X has a *weak* reference to any node Y that is visible to X . If $\mathcal{V}[X, Y]$ has some edges then we say that X has a *partially strong* reference to Y . If $\mathcal{V}[X, Y]$ is $H[X, Y]$, that is, it has all the nodes and edges in the path from X to Y in H , then we say that X has a *strong* reference to Y . The weak, partially strong and strong references are also denoted as *weak*, *partially strong* and *strong visibilities*, respectively. We use the following notation to represent the strength of visibility. If $\mathcal{V}[X, Y]$ is $H[Y_1, Y]$, for Y_1 in the path from X to Y , then it is denoted as (Y_1, Y) . Then strong reference will be (X, Y) , and weak reference will be (Y) .

In the following, we define two special properties for visibility assignments: coherence and correlation. We first illustrate these properties with a strong reference $\mathcal{V}[X, Y]$, in Fig. 3 and Fig. 4. Here, $H[X, Y] = (X, Y_3, Y_2, Y_1, Y)$ is the path from X to Y .

- Coherence refers to the property that X has a strong reference to Y implies X has a strong reference to every node in the path from X to Y . That is, $\mathcal{V}[X, Y]$ is (X, Y) implies (i) $\mathcal{V}[X, Y_1]$ is (X, Y_1) , (ii) $\mathcal{V}[X, Y_2]$ is (X, Y_2) , and (iii) $\mathcal{V}[X, Y_3]$ is (X, Y_3) .
- Correlation refers to the property that X has a strong reference to Y implies every node in the path from X to Y has a strong reference to Y . That is, $\mathcal{V}[X, Y]$ is (X, Y) implies (i) $\mathcal{V}[Y_3, Y]$ is (Y_3, Y) , (ii) $\mathcal{V}[Y_2, Y]$ is (Y_2, Y) , and (iii) $\mathcal{V}[Y_1, Y]$ is (Y_1, Y) .

3.1 Coherence and correlation

Definition 3.2 (Coherent Visibility) A visibility assignment $\mathcal{V}$ is *coherent* if for each pair of nodes X and Z , and every node Y in the path from X to Z in H , Y not equal to X or Z , $\mathcal{V}[X, Z] \cap H[X, Y]$ is a subgraph of $\mathcal{V}[X, Y]$.




Figure 3: Coherence: (a) implies (b), (c) and (d).




Figure 4: Correlation: (a) implies (b), (c) and (d).




Figure 5: Sample path.




Figure 6: Coherent and non-coherent pairs.

Nodes X , Y and Z are illustrated in Fig. 5. We use this figure to comprehend and relate the different properties in this paper. Informally, coherence means that the strength of visibility of X over Y is *at least* as much as the strength used for visibility of X over Z : $\mathcal{V}[X, Z] \cap H[X, Y]$ refers to the strength of visibility of X over Y “used” for visibility over Z , whereas $\mathcal{V}[X, Y]$ is simply the strength of visibility of X over Y . Fig. 6 describes several pairs of $\mathcal{V}[X, Z]$ and $\mathcal{V}[X, Y]$ in coherent and non-coherent visibility assignments.




Figure 7: Correlation.

Correlation property captures the intuitive notion that, referring to Fig. 5 again, the strength of visibility of Y over Z is *at least* as much as the strength of visibility of X over Z restricted to $H[Y, Z]$.

Definition 3.3 (Correlated Visibility) A visibility assignment $\mathcal{V}$ is *correlated* if for each pair of nodes X and Z , and every node Y in the path from X to Z in H , Y not equal to X or Z , $\mathcal{V}[X, Z] \cap H[Y, Z]$ is a subgraph of $\mathcal{V}[Y, Z]$.

Fig. 7 illustrates correlation. Informally, along the path from Z to X , the strength of visibility over Z may first be increasing, and then remain the same, and finally be decreasing. At some stage, Z may not be visible and, if so, it remains invisible to all the other nodes along the path.

We note that coherence and correlation are orthogonal properties. Fig. 8 shows a visibility assignment which is coherent but not correlational: X_1 is visible to X_4 , but not to its parent X_3 . Fig. 9 shows a visibility assignment which is correlated but not coherent: $\mathcal{V}[X_4, X_1] \cap H[X_4, X_2]$ is not a subgraph of $\mathcal{V}[X_4, X_2]$.




Figure 8: Illustration of coherent, but not correlated, SoV's.




Figure 9: Illustration of correlated, but not coherent, SoV's.




Figure 10: Inverse coherence.

3.2 Related properties

In this section, we define some variants of the coherence and correlation properties that will be useful in some application environments. First we consider coherence. By replacing “subgraph” by “supergraph” in the definition of coherence, we get inverse coherence. For graphs G and G' , we call G a *supergraph* of G' if G' is a subgraph of G .

Definition 3.4 (Inversely Coherent Visibility) A visibility assignment $\mathcal{V}$ is *inversely coherent* if for each pair of nodes X and Z , and every node Y in the path from X to Z in H , Y not equal to X or Z , $\mathcal{V}[X, Z] \cap H[X, Y]$ is a supergraph of $\mathcal{V}[X, Y]$.

Informally, inverse coherence means that the strength of visibility of X over Y is *at most* as much as the strength used for visibility of X over Z . Fig. 10 shows a visibility assignment where adjacent nodes are weakly visible, while non-adjacent nodes are strongly visible. Here (Fig. 10), inverse coherence is satisfied.

Another variant of the coherence property is the following.

Definition 3.5 (Uniformly Coherent Visibility) A visibility assignment $\mathcal{V}$ is *uniformly coherent* if for each pair of nodes X and Z , and every node Y in the path from X to Z in H , Y not equal to X or Z , $\mathcal{V}[X, Z] \cap H[X, Y]$ is equal to $\mathcal{V}[X, Y]$.

Note that uniform coherence is the same as uniform inverse coherence. Fig. 11 illustrates the notion of uniform coherence.

Uniform coherence implies that when $\mathcal{V}[X, Z] \cap H[X, Y]$ is null, $\mathcal{V}[X, Y]$ must also be null. This must be so irrespective of whether $\mathcal{V}[X, Z]$ itself is null or nonnull. We envisage that, in some applications, it is reasonable to have nonnull $\mathcal{V}[X, Y]$ when $\mathcal{V}[X, Z]$ is nonnull




Figure 11: Uniform coherence.

but $\mathcal{V}[X, Z] \cap H[X, Y]$ is null. Note that this possibility is included in the coherence property definition because the null graph is a subgraph of any graph. Also, this possibility is contrary to the intuitive notion of inverse coherence, namely, the visibility of distant nodes is stronger than that of nearby nodes. We define a special kind of uniform coherence that allows this possibility:

Definition 3.6 (Nonnull Uniformly Coherent Visibility) A visibility assignment $\mathcal{V}$ is *nonnull uniformly coherent* if: for each pair of nodes X and Z , and every node Y in the path from X to Z in H , Y not equal to X or Z , if $\mathcal{V}[X, Z] \cap H[X, Y]$ is nonnull, then it is equal to $\mathcal{V}[X, Y]$.

Uniform coherence facilitates representing SoV_X , of a node X , in a single graph. The graph is obtained just by merging all the graphs in SoV_X . We denote this graph as $\mathcal{V}_X$. We show the validity of such representation in the following. The connected component of $\mathcal{V}_X \cap H[X, Y]$ that contains Y is denoted $\mathcal{V}_X[X, Y]$. In this paper, we use “ $\subseteq$ ” to denote the subgraph relationship. By an *extension* of a path (X, Y) we mean a path (X, Z) , for some Z , such that (X, Y) is a subgraph of (X, Z) .

Proposition 3.1 For a uniformly coherent visibility assignment $\mathcal{V}$, $\mathcal{V}_X$ represents SoV_X in a “lossless” fashion, that is, $\mathcal{V}_X[X, Y]$ equals $\mathcal{V}[X, Y]$, for every Y .

Proof: For any Y , clearly, $\mathcal{V}[X, Y] \subseteq \mathcal{V}_X[X, Y]$. We show the equality. We need to consider only the case where $\mathcal{V}_X[X, Y]$ is nonnull. We first show that $\mathcal{V}[X, Y]$ is nonnull. Assume the contrary. Then, for some node Z in the extension of the path (X, Y) , $\mathcal{V}[X, Z]$ must contain the path (Y, Z) . (Note that Z could be Y or some other node.) Then,




Figure 12: Illustration for the proof of Proposition 3.1.

$\mathcal{V}[X, Z] \cap H[X, Y]$ is nonnull (containing at least the node Y), and it should be equal to $\mathcal{V}[X, Y]$, by uniform coherence. Thus, $\mathcal{V}[X, Y]$ cannot be null.

If $\mathcal{V}[X, Y]$ is strong, that is, (X, Y) , then clearly, $\mathcal{V}_X[X, Y]$ is also (X, Y) . Now suppose $\mathcal{V}[X, Y]$ is a proper subgraph of $H[X, Y]$, say (Y_1, Y) for some Y_1 in the path from X to Y . We claim that $\mathcal{V}_X[X, Y]$ does not contain any more edges than in (Y_1, Y) . Suppose, on the contrary, that for node Y_2 preceding Y_1 in the path from X to Y , $\mathcal{V}_X[X, Y]$ contains the edge (Y_2, Y_1) also. Then, for some Y' in the extension of the path (Y_2, Y_1) , $\mathcal{V}[X, Y']$ must contain (Y_2, Y') . Now, Y' may be in (X, Y) or in the extension of (X, Y) , or may be in a completely different path from Y_1 . (An instance of the last case is illustrated in Fig. 12.) In all cases, $\mathcal{V}[X, Y'] \cap H[X, Y_1]$ should equal $\mathcal{V}[X, Y_1]$ by uniform coherence and hence contain (Y_2, Y_1) . However, again by uniform coherence, $\mathcal{V}[X, Y] \cap H[X, Y_1]$ should equal $\mathcal{V}[X, Y_1]$ and so $\mathcal{V}[X, Y]$ should include (Y_2, Y_1) , a contradiction to the assumption otherwise. The proof is complete.

We describe some characteristics of $\mathcal{V}_X$, for a uniformly coherent visibility $\mathcal{V}$, with Fig. 13: Essentially, if a node T is visible to X with strength (T_1, T) , then every node T' in the extension of (X, T) is visible to X with strength (T_1, T') . In Fig. 13:

- Y_i , for each i , is visible with strength (X, Y_i) ;
- Z_3 , Z_4 and Z_5 are visible with strengths (Z_3) , (Z_3, Z_4) , and (Z_3, Z_5) , respectively.
- Z_8 is visible with strength (Z_8) ;
- V_1 is visible with strength (V_1) ; and
- W_1 and W_2 are not visible.




Figure 13: Single graph representation - Uniformly coherent visibility.

The above proposition is applicable for nonnull uniform coherency also. Fig. 14 illustrates nonnull uniformly coherent visibility. Here, in addition to the nodes visible according to Fig. 13, we have:

- Z_1 and Z_6 are visible with strengths (X, Z_1) and (X, Z_6) , respectively; and
- W_1 is visible with strength (X, W_1) .

Fig. 15(a) illustrates a nonnull uniformly coherent SoV of the user U in Fig. 1. Note the change in $\mathcal{V}[U, S-A]$ (for uniform coherence) and the addition of $\mathcal{V}[U, P]$ (for coherence), compared to the SoV of U shown in Fig. 2(e). Fig. 15(a) is depicted in a single graph, as shown in Fig. 15(b). Fig. 16 shows a nonnull uniformly coherent SoV of the supplier $S-A$ (Fig. 1). The partially strong visibility of $S-A$ over M in Fig. 2(f) has been changed to a strong visibility in Fig. 16.

We now consider variants of the correlation property. The definitions are similar to those of coherence.

Definition 3.7 (Inversely Correlated Visibility) A visibility assignment $\mathcal{V}$ is *inversely correlated* if for each pair of nodes X and Z , and every node Y in the path from X to Z in H , Y not equal to X or Z , $\mathcal{V}[X, Z] \cap H[Y, Z]$ is a supergraph of $\mathcal{V}[Y, Z]$.

It is easy to check that the visibility assignment in Fig. 10, where adjacent nodes are weakly visible and non-adjacent nodes are strongly visible, is also inversely correlated (in addition to being inversely coherent).





Figure 14: Single graph representation - Nonnull uniformly coherent visibility.

Figure 15: Nonnull uniformly coherent SoV of the user U in Fig. 1.




Figure 16: Nonnull uniformly coherent SoV of the supplier $S-A$ in Fig. 1.

Definition 3.8 (Uniformly Correlated Visibility) A visibility assignment $\mathcal{V}$ is *uniformly correlated* if for each pair of nodes X and Z , and every node Y in the path from X to Z in H , Y not equal to X or Z , $\mathcal{V}[X, Z] \cap H[Y, Z]$ is equal to $\mathcal{V}[Y, Z]$.

Definition 3.9 (Nonnull Uniformly Correlated Visibility) A visibility assignment $\mathcal{V}$ is *nonnull uniformly correlated* if: for each pair of nodes X and Z , and every node Y in the path from X to Z in H , Y not equal to X or Z , if $\mathcal{V}[X, Z] \cap H[Y, Z]$ is nonnull, then it is equal to $\mathcal{V}[Y, Z]$.

In Fig. 7, for uniform correlation, $\mathcal{V}[X_2, Z]$, $\mathcal{V}[X_1, Z]$, etc. will be the same as $\mathcal{V}[X_3, Z]$. That is, the decreasing phase will not be there; the strength may increase, then remain the same. With nonnull uniform correlation, the visibility may become null at some point. If so, it will stay null from thereon. Replacing $\mathcal{V}[X_2, Z]$ with null in that figure gives nonnull uniformly correlated visibility. (Note that if $\mathcal{V}[X, Z]$ is nonnull, then $\mathcal{V}[X, Z] \cap H[X_1, Z]$ is also nonnull and so $\mathcal{V}[X_1, Z]$ must be nonnull too, by uniform correlation.)

As with uniform coherence, uniform correlation can also be depicted in a single graph. We show this in the next section.

4 Visibility Policies

In general, a visibility assignment can be done by prescribing $\mathcal{V}[X, Y]$ for each pair of nodes X and Y individually. However, in many applications, a global visibility policy could be formulated and the individual assignments based on this policy. In addition to simplifying the logic behind the individual assignments, this facilitates defining visibility characteristics even for nodes that have not joined the hierarchy yet, in a dynamic environment. In this section, we illustrate a few, very simple examples of, global policies. We start with a policy that is already illustrated in Fig. 10.

Policy 1: Weak visibility to adjacent nodes and strong visibility to all other nodes.

As mentioned earlier, this visibility has inverse coherence and inverse correlation properties.

We illustrate the rest of the policies with the path
 $(X = Y_5, Y_4, Y_3, Y_2, Y_1, Y_0 = Y)$.

Policy 2: (Fig. 17) Strong visibility to nodes at distance 1 or 2, and partially strong visibility consisting of a path with two edges to all other nodes. That is, $\mathcal{V}[X, Y]$ is:

- (X, Y) if Y is adjacent or at distance 2 from X , and
- (Y_2, Y) if Y is at distance greater than 2 from X , where Y_2 is the third last node in the path from X to Y .

This visibility assignment is coherent and uniformly correlated.

Example of coherence is,

- $\mathcal{V}[X, Y] \cap H[X, Y_2]$ is (Y_2) , and is a subgraph of
- $\mathcal{V}[X, Y_2]$ which is (Y_4, Y_2) .

An example of uniform correlation is,

- $\mathcal{V}[X, Y] \cap H[Y_3, Y]$ is (Y_2, Y) , and is equal to
- $\mathcal{V}[Y_3, Y]$ which is also (Y_2, Y) .

Policy 3: (A slight variant of Policy 2) (Fig. 18) Strong visibility to nodes at distance 1 or 2, partially strong visibility consisting of a path with two edges to nodes at distance 3 or 4, and no visibility of all other nodes. That is, $\mathcal{V}[X, Y]$ is:

- (X, Y) if Y is adjacent or at distance 2 from X ,
- (Y_2, Y) if Y is at distance 3 or 4 from X , where Y_2 is the third last node in the path from X to Y , and




Figure 17: Illustration for visibility policy 2.




Figure 18: Illustration for visibility policy 3.




Figure 19: Illustration for visibility policy 4.

- null if Y is at distance 5 or more.

This visibility assignment is coherent and nonnull uniformly correlated. Here, uniform correlation is satisfied only when $\mathcal{V}[X, Y]$ is not null.

Policy 4: (Fig. 19) No visibility to adjacent nodes and weak visibility to all other nodes.

Here, for any Y_i , $i < 5$, $\mathcal{V}[X, Y] \cap H[X, Y_i]$ is null. Therefore, we have nonnull uniform coherence (which is also nonnull uniform inverse coherence). For any i , $\mathcal{V}[X, Y] \cap H[Y_i, Y]$ is (Y) ; it is equal to $\mathcal{V}[Y_i, Y]$, for $i > 1$, and null for $i = 1$. Hence, we have inverse correlation.

Policy 5: (Fig. 20) Strong visibility to nodes at distance 1 and 2, weak visibility to nodes at distance 3 or 4, and strong visibility to all other nodes.

It is not clear where such visibility assignment will be useful. This assignment does not have any of the properties we have discussed. We are stating this policy just to illustrate that a visibility assignment need not satisfy any properties.

Example of coherence is, with X , Y_3 and Y_2 ,

- $\mathcal{V}[X, Y_2]$ is (Y_2) .
- $\mathcal{V}[X, Y_2] \cap H[X, Y_3]$ is null, and is a subgraph of




Figure 20: Illustration for visibility policy 5.

- $\mathcal{V}[X, Y_3]$ which is (X, Y_3) .

Example of inverse coherence is, with X , Y_1 and Y ,

- $\mathcal{V}[X, Y]$ is (X, Y) .
- $\mathcal{V}[X, Y] \cap H[X, Y_1]$ is (X, Y_1) , and is a supergraph of
- $\mathcal{V}[X, Y_1]$ which is (Y_1) .

An example of inverse correlation is, with X , Y_3 and Y ,

- $\mathcal{V}[X, Y]$ is (X, Y) .
- $\mathcal{V}[X, Y] \cap H[Y_3, Y]$ is (Y_3, Y) which is a supergraph of
- $\mathcal{V}[Y_3, Y]$ which is (Y) .

An example of correlation is, with X , Y_3 and Y_1 ,

- $\mathcal{V}[X, Y_1]$ is (Y_1) .
- $\mathcal{V}[X, Y_1] \cap H[Y_3, Y_1]$ is (Y_1) which is a subgraph of
- $\mathcal{V}[Y_3, Y_1]$ which is (Y_3, Y_1) .

5 Sphere of Noticeability

For a node X , the Sphere of Noticeability notion is intended to capture: (i) which nodes have visibility over X ; and (ii) what type of visibility they have of X . First we define noticeability independent of, but in a way analogous to the definition of, visibility. We refer to a general *noticeability assignment* $\mathcal{N}$ in H with respect to an attribute A . $\mathcal{N}$ consists of a set of subgraphs $\mathcal{N}[X, Y]$, for all pairs X, Y of nodes in H , defined as follows: $\mathcal{N}[X, Y]$ is either (i) a connected subgraph of $H[X, Y]$ that contains X , or (ii) the null graph. $\mathcal{N}[X, Y]$ denotes the *type*, also *strength*, of noticeability, that is, the type of visibility Y has over X . In the last case, X is not noticed by Y . We assume that $\mathcal{N}[X, X]$, for every X , is the graph containing just the node X . Note that since visibility and noticeability notions are complementary, $\mathcal{V}$ and $\mathcal{N}$ definitions are also complementary. That is, for $\mathcal{N}$ that “corresponds to” a $\mathcal{V}$, for X and Y , $\mathcal{N}[X, Y]$ is the same as $\mathcal{V}[Y, X]$. We use $\mathcal{N}[X, Y]$ most of the time in the definitions and discussions in this section, though $\mathcal{N}[Y, X]$ could also be used instead. In the illustration of noticeability of X by Y , X is represented in an oval inside a rectangle, Y is represented by a thick oval, and other nodes, if any, are represented in thin ovals.

Definition 5.1 (Sphere of Noticeability) The *Sphere of Noticeability* of a node X in hierarchy H , denoted SoN_X , is the set of non-null subgraphs $\mathcal{N}[X, Y]$, for Y in H . The set of nodes Y for which $\mathcal{N}[X, Y]$ is non-null is denoted as N_X .

Note that, for a specific node X , SoV_X is the set of $\mathcal{V}[X, Y]$ ’s for different Y ’s, whereas SoN_X is the set of $\mathcal{V}[Y, X]$ ’s for different Y ’s.

An obvious application of SoN is for change management. An example is a provider X notifying the providers, who have visibility over X , when there is some change in the provider URI (provider details), metrics used to compute the service (service details), log format (execution details), etc. An interpretation of the relationship between SoV and SoN using the e-shopping scenario introduced earlier (Fig. 1) follows. For a node X , SoV_X can be considered as the nodes from which some information (input) is expected. SoN_X can be considered as the nodes to which some information (output) is to be sent. In both cases, the type of visibility reflects how the information may be received or sent. For example (Figures 1 and 2), for the air miles provider B , SoV_B conveys that B is expecting the credit charge information from H and the air miles account details from U . On the other hand, SoN_B may contain U , P and H , reflecting that B should send confirmation of the air miles reward to U , P and H .

Coherence and correlation properties for noticeability assignments can be defined analogous to those for visibility assignments.

Definition 5.2 (Coherence Properties of Noticeability) A noticeability assignment $\mathcal{N}$ is *coherent* (respectively, *inversely coherent*, *uniformly coherent*), if for each pair of nodes Z and X , and every node Y in the path from Z to X in H , Y not equal to Z or X , $\mathcal{N}[Z, X] \cap H[Z, Y]$ is a subgraph of (respectively, a supergraph of, equal to) $\mathcal{N}[Z, Y]$.




Figure 21: Illustration for the proof of Proposition 5.1.

Definition 5.3 (Nonnull Uniformly Coherent Noticeability) A noticeability assignment $\mathcal{N}$ is *nonnull uniformly coherent* if: for each pair of nodes Z and X , and every node Y in the path from Z to X in H , Y not equal to Z or X , if $\mathcal{N}[Z, X] \cap H[Z, Y]$ is nonnull, then it is equal to $\mathcal{N}[Z, Y]$.

Here also, both uniform and nonnull uniform coherence facilitate merging all the subgraphs of SoN_X , of a node X , into a single graph. The graph is denoted $\mathcal{N}_X$. The connected component of $\mathcal{N}_X \cap H[X, Y]$ that contains X is denoted $\mathcal{N}_X[X, Y]$.

Proposition 5.1 For a uniformly coherent noticeability assignment $\mathcal{N}$, $\mathcal{N}_X$ represents SoN_X in a “lossless” fashion, that is, $\mathcal{N}_X[X, Y]$ equals $\mathcal{N}[X, Y]$, for every Y .

Proof: For any Y , clearly, $\mathcal{N}[X, Y] \subseteq \mathcal{N}_X[X, Y]$. We show the equality. We need to consider only the case where $\mathcal{N}_X[X, Y]$ is nonnull. We first show that $\mathcal{N}[X, Y]$ is nonnull. Assume the contrary. Then, for some node Z in (i) $H[X, Y]$ or (ii) its extension, $\mathcal{N}[X, Z]$ must contain at least the node X , that is, the path (X) . In case (i), $\mathcal{N}[X, Y] \cap H[X, Z]$ is null and so not equal to $\mathcal{N}[X, Z]$, violating uniform coherence. In case (ii), $\mathcal{N}[X, Z] \cap H[X, Y]$ is nonnull and should be equal to $\mathcal{N}[X, Y]$ by uniform coherence, contrary to the assumption.

If $\mathcal{N}[X, Y]$ is strong, that is, (X, Y) , then clearly, $\mathcal{N}_X[X, Y]$ is also (X, Y) . Now suppose $\mathcal{N}[X, Y]$ is a proper subgraph of $H[X, Y]$, say (X, Y_1) for some Y_1 in the path from X to Y . We claim that $\mathcal{N}_X[X, Y]$ does not contain any more edges than in (X, Y_1) . Suppose, on the contrary, that for a node Y_2 following Y_1 in the path from X to Y , $\mathcal{N}_X[X, Y]$ contains the edge (Y_1, Y_2) . Then, for some Y' in the extension of the path (Y_1, Y_2) , $\mathcal{N}[X, Y']$ must contain (Y_1, Y_2) . Now Y' may be in (X, Y) or in the extension of (X, Y) , or may be in a completely different path from Y_2 . The last case is illustrated in Fig. 21. In all cases, $\mathcal{N}[X, Y'] \cap H[X, Y_2]$ should equal $\mathcal{N}[X, Y_2]$ by uniform coherence, and hence contain (Y_1, Y_2) . However, again by uniform coherence, $\mathcal{N}[X, Y] \cap H[X, Y_2]$ should equal $\mathcal{N}[X, Y_2]$ and so $\mathcal{N}[X, Y]$ must contain the edge (Y_1, Y_2) also, a contradiction to the assumption otherwise. The proof is complete.

With uniformly coherent noticeability, if X is noticed by a node T with strength (X, T_1) , where T_1 is not T , then X is noticed by every T' in the extension of (X, T) with the same strength (X, T_1) . Fig. 22 illustrates $\mathcal{N}_X$ for a uniformly coherent noticeability.




Figure 22: Single noticeability graph - Uniformly coherent noticeability.

X is noticeable to:

- Y_1 with strength (X, Y_1) ;
- each Y_i , for $i > 1$, with strength (X, Y_2) ;
- V_1 with strength (X) ;
- Z_j , for j equal to 6, 7 and 8, with strength (X, Z_6) ;
- Z_k , for k from 1 to 5, with strength (X, Z_1) ; and
- X is not noticeable to W_1 and W_2 .

Again, the above proposition is applicable to nonnull uniform coherence also. Fig. 23 illustrates a nonnull uniformly coherent noticeability. Note that X is not noticeable to Z_3 , Z_4 and Z_5 . Here, X may be noticeable for some nodes along a path (X, Y) and may become unnoticeable for Y . If that happens, then X will remain unnoticeable for every node Z in the extension of (X, Y) .

A nonnull uniformly coherent SoN of the user U (Fig. 1) is shown in Fig. 24. It is shown as a single graph in Fig. 25. A nonnull uniformly coherent SoN of the supplier $S-A$ (Fig. 1) and its representation as a single graph, are shown in Fig. 26 and Fig. 27, respectively.

We now define the correlation property and its variants.

Definition 5.4 (Correlation Properties of Noticeability) A noticeability assignment $\mathcal{N}$ is *correlated* (respectively, *inversely correlated*, *uniformly correlated*) if for each pair of





Figure 23: Single noticeability graph - Nonnull uniformly coherent noticeability.

Figure 24: Nonnull uniformly coherent SoN of the user U in Fig. 1.




Figure 25: Nonnull uniformly coherent SoN of the user U in Fig. 1 (represented as a single graph).




Figure 26: Nonnull uniformly coherent SoN of the supplier $S-A$ in Fig. 1.




Figure 27: Nonnull uniformly coherent SoN of the supplier $S-A$ in Fig. 1 (represented as a single graph).

nodes Z and X , and every node Y in the path from Z to X in H , Y not equal to Z or X , $\mathcal{N}[Z, X] \cap H[Y, X]$ is a subgraph of (respectively, a supergraph of, equal to) $\mathcal{N}[Y, X]$.

Definition 5.5 (Nonnull uniformly Correlated Noticeability) A noticeability assignment $\mathcal{N}$ is *nonnull uniformly correlated* if: for each pair of nodes Z and X , and every node Y in the path from Z to X in H , Y not equal to Z or X , if $\mathcal{N}[Z, X] \cap H[Y, X]$ is nonnull, then it is equal to $\mathcal{N}[Y, X]$.

We now show that the coherence and correlation properties are indeed related.

Proposition 5.2 In a hierarchy H , a visibility assignment $\mathcal{V}$ is coherent (respectively, inversely coherent, uniformly coherent, nonnull uniformly coherent) if and only if the corresponding noticeability assignment $\mathcal{N}$ is correlated (respectively, inversely correlated, uniformly correlated, nonnull uniformly correlated), and vice versa.

Proof: The proof follows from the complementary relationship between visibility and noticeability. Recall that, for $\mathcal{N}$ that “corresponds to” a $\mathcal{V}$, for X and Y , $\mathcal{N}[X, Y]$ is the same as $\mathcal{V}[Y, X]$. Applying the above transformation to the coherent visibility definition, we get the correlated noticeability definition and vice versa. The same applies to the other equivalences.

We note that due to the equivalence of uniformly correlated visibility and uniformly coherent noticeability, and the property that uniformly coherent noticeability can be represented in a single graph, it follows that uniformly correlated visibility can be represented in a single noticeability graph. As stated earlier, uniformly coherent visibility can be represented in a single visibility graph.

6 A Special Case

We have shown that uniformly coherent visibility facilitates representing the SoV of any node X in a single graph $\mathcal{V}_X$. For different X 's, the graphs $\mathcal{V}_X$'s will also be different. In this section, we show that under certain special conditions, the SoV's of all nodes in H can be represented in a single graph.

Definition 6.1 (sv-symmetry) A visibility assignment $\mathcal{V}$ is *strong visibility symmetric*, abbreviated *sv-symmetric*, if for every pair of nodes X and Y , if $\mathcal{V}[X, Y]$ is strong (that is, equals (X, Y)), then $\mathcal{V}[Y, X]$ is also strong (that is, equals (Y, X)).

We note that, under uniformly coherent visibility assignment, the single graph representing $\mathcal{V}_X$ is implicitly a directed tree rooted at X , with edges directed towards X . Similarly, for a node Y different from X , $\mathcal{V}_Y$ is a directed tree rooted at Y . When we try to represent both $\mathcal{V}_X$ and $\mathcal{V}_Y$ in a single graph, the edges in the path from X to Y have to be represented in both directions. The sv-symmetry property enables representing the edges in undirected form, in a single graph representation.

Definition 6.2 (Harmonious Visibility and Noticeability) A visibility assignment $\mathcal{V}$ is *harmonious* if it satisfies the following properties:

- nonnull uniform coherence;
- nonnull uniform correlation;
- sv-symmetry; and
- if a node is visible to some node in H , then it is visible to every node in H .

A noticeability assignment $\mathcal{N}$ is *harmonious* if its corresponding visibility assignment is harmonious.

It turns out that *for a harmonious visibility assignment, the visibility graph $\mathcal{V}_X$, for any node X , is the visibility graph for all nodes in H .*

Proposition 6.1 For a harmonious visibility assignment $\mathcal{V}$, $\mathcal{V}_X$, for any node X , represents $\mathcal{V}$ in a “lossless” fashion. That is, $\mathcal{V}_X[Y, Z]$ equals $\mathcal{V}[Y, Z]$, for any nodes Y and Z .

Proof: If $\mathcal{V}[Y, Z]$ is null, then $\mathcal{V}[X, Z]$ is also null, by the last property in the statement of the proposition. Therefore Z will not be in $\mathcal{V}_X$. In the following, we consider the case where $\mathcal{V}[Y, Z]$ is nonnull. Let $\mathcal{V}_X[Y, Z]$ be (Z_1, Z) , where Z_1 is a node in the path from Y to Z . Note that Z_1 could be Y , Z , or some other node. We will show that $\mathcal{V}[Y, Z]$ is also (Z_1, Z) .

We consider the various possibilities of the relative positions of X , Y and Z in H , illustrated in Fig. 28.

(a) X is in an extension of the path (Z, Y) .

Here, $\mathcal{V}[X, Z] \cap H[Y, Z]$ is also (Z_1, Z) and this equals $\mathcal{V}[Y, Z]$, by nonnull uniform correlation.

(b) X is in the extension of the path (Y, Z) .

Here, $\mathcal{V}[X, Z_1]$ contains (Z, Z_1) . By nonnull uniform correlation, $\mathcal{V}[X, Z_1] \cap H[Z, Z_1]$ equals $\mathcal{V}[Z, Z_1]$. By sv-symmetry, this equals $\mathcal{V}[Z_1, Z]$. Now, by nonnull uniform correlation, $\mathcal{V}[Y, Z] \cap H[Z_1, Z]$ equals $\mathcal{V}[Z_1, Z]$. Therefore, $\mathcal{V}[Y, Z]$ includes (Z_1, Z) . We show that it indeed equals (Z_1, Z) . Suppose on the contrary that $\mathcal{V}[Y, Z]$ has Z_2 (preceding Z_1) and the edge (Z_2, Z_1) , that is, it includes (Z_2, Z) . Then, by uniform correlation $\mathcal{V}[Z_2, Z]$ must be (Z_2, Z) ; by sv-symmetry, $\mathcal{V}[Z, Z_2]$ must be (Z, Z_2) ; by uniform correlation, $\mathcal{V}[X, Z_2] \cap H[Z, Z_2]$ must be (Z, Z_2) ; and so $\mathcal{V}[X, Z_2]$ must include (Z_2, Z) , and hence the edge (Z_2, Z_1) . Then, $\mathcal{V}_X[Y, Z]$ will also have (Z_2, Z_1) , contradicting the assumption otherwise.

(c) X is in the path (Y, Z) .

Three subcases arise:

(i) Z_1 is in the path (X, Z) and is different from X . Then, by nonnull uniform correlation, $\mathcal{V}[Y, Z] \cap H[X, Z]$ equals $\mathcal{V}[X, Z]$ which is (Z_1, Z) . Therefore, $\mathcal{V}[Y, Z]$ is (Z_1, Z) .

(ii) Z_1 is in the path (Y, X) and is different from X . Here, $\mathcal{V}[X, Z_1]$ must be (X, Z_1) . By sv-symmetry, $\mathcal{V}[Z_1, X]$ is (Z_1, X) . Also, $\mathcal{V}[Y, Z] \cap H[X, Z]$ equals (X, Z) , and therefore $\mathcal{V}[Z_1, Z]$ is (Z_1, Z) . By nonnull uniform correlation, $\mathcal{V}[Y, Z] \cap H[Z_1, Z]$ is $\mathcal{V}[Z_1, Z]$. Thus, $\mathcal{V}[Y, Z]$ includes (Z_1, Z) . Suppose $\mathcal{V}[Y, Z]$ includes another edge (Z_2, Z_1) , for Z_2 preceding Z_1 , as in the previous case. Then, by nonnull uniform correlation, $\mathcal{V}[Y, Z] \cap H[Z_2, Z]$ is (Z_2, Z) ; by nonnull uniform coherence, $\mathcal{V}[Z_2, Z] \cap H[Z_2, X]$ is (Z_2, X) ; by sv-symmetry, $\mathcal{V}[X, Z_2]$ is (X, Z_2) . Therefore, the edge (Z_2, Z_1) must also be in $\mathcal{V}_X$, in contradiction to the assumption otherwise.

(iii) Z_1 is X . Here also, $\mathcal{V}[Y, Z] \cap H[Z_1, Z]$ is $\mathcal{V}[Z_1, Z]$. Thus, $\mathcal{V}[Y, Z]$ includes (Z_1, Z) . The proof that $\mathcal{V}[Y, Z]$ does not contain additional edges follows as in subcase (ii) above.

(d) X is neither in (Y, Z) nor in any of its extensions.

Then let W be the node in (Y, Z) common to the paths (X, Y) and (X, Z) . Two subcases arise.

(i) Z_1 is in the path (W, Z) , and is different from W . Then, $\mathcal{V}[X, Z]$ must be (Z_1, Z) . By nonnull uniform correlation, $\mathcal{V}[X, Z] \cap H[W, Z]$ is $\mathcal{V}[W, Z]$ and is equal to (Z_1, Z) ; by nonnull uniform correlation, $\mathcal{V}[Y, Z] \cap H[W, Z]$ is (Z_1, Z) . Therefore, $\mathcal{V}[Y, Z]$ is (Z_1, Z) .

(ii) Z_1 is W or is in the path (Y, W) . In this case, by nonnull uniform coherence, $\mathcal{V}[Y, Z] \cap H[Y, W]$ is $\mathcal{V}[Y, W]$. Therefore, we need to show only that $\mathcal{V}[Y, W]$ is (Z_1, W) . This follows from case (b), by substituting W in place of Z .

The proof is complete.

The above proposition asserts that when the visibility assignment is harmonious, a single graph can represent the visibility graph of every node, including those which are not visible




Figure 28: Cases for the proof of Proposition 6.1.

to any node at all. The latter part follows from the fact that the proof of the proposition did not make use of whether X itself is visible (to every node) or not, and similarly with the visibility of Y . We also note that the same graph, with the addition of all invisible nodes, is the noticeability graph for all the visible nodes. The noticeability graph for each invisible node will be the graph containing just that node.

7 Related Works

The notion of spheres of visibility and noticeability is based on the concept of Spheres of Control (SoC) [1]. A Sphere of Control encapsulates entities sharing a similar set of properties or having a dependency relation. The dependency relations considered in [1] are atomicity, commitment, resource allocation, recovery, auditing, consistency, etc. SoV and SoN logically group the nodes (and their attributes) visible to another node in a hierarchy. Also, [1] considers homogeneous and non-autonomous systems where visibility is not an issue. Thus, our work can be considered as complementary to the work in [1] to heterogeneous and autonomous systems.

Later works have extended the initial concept of SoC to Spheres of Atomicity [6] and Commitment [12]. [6] utilizes the properties of the processes (pivot, compensatable and retrievable) in a Sphere of Atomicity to determine if the sphere, as a whole, guarantees atomicity. [12] applies the concept of SoC to Multi-Agent Systems (MAS) to structure agents based on their commitment guarantees. However, the above works are not directly related to the work presented in this paper and we mention them for the sake of completeness.

No other work (that we are aware of) has attempted to formalize the visibility aspect for hierarchical systems. Some of the works which have touched upon this aspect are the following: [8] identifies real-life scenarios where there might be a need to deviate from the inheritance of access rights upwards through the hierarchy in a role-based access control. [13] discusses the visibility aspect with respect to the visibility of the results of a subtransaction in a nested transactional system. Basically, [13] advocates the provision to be able to expose the results to a particular ancestor to improve performance. [9] proposes a formal model based on π -calculus to capture the behaviour of nested long running transactions in a Web services context. Of particular interest are the different modes of failure propagation: up-propagation, down-propagation, down-specific propagation and non-propagation. [4] studies the problem of determining the minimum number of nodes which need to be visible such that the system is still compensable. [10] considers the visibility aspect with respect to sending publish/subscribe notifications for event based systems. [11] models P2P systems, where a node has knowledge of (visibility over) only a subset of the remaining nodes, abstracted as P2P communities. However, none of the above works considers the effect of *relationships* which might exist among the visibilities of different entities of a system, such as, coherence, correlation, etc.

In a previous work [2], we had introduced the SoV concept with respect to hierarchical Web Services compositions, and shown its application in the context of performing compensation under visibility constraints. However, [2] only considered vertical visibility (that is, visibility over ancestors and descendents) as compared to the more generalized notion of visibility presented in [3] (visibility over siblings, uncles, cousins, etc.). The notions of noticeability, coherence and correlation were also introduced in [3]. In this paper, we further extend the dual properties of coherence and correlation, to study their inverse and uniform variations. The discussions with respect to visibility policies and single visibility (noticeability) graph representations are also novel contributions of this paper.

8 Discussion and Conclusion

Consider a node X of a static hierarchy H . Initially, X has some visibility requirements and noticeability restrictions. Note that the above reflects only X 's point of view, that is, visibility and noticeability that X would like. As such, the finally assigned SoV_X may be quite different from its initial expectations due to conflicts with the noticeability restrictions of nodes Y over which X would like to have visibility. The same applies for SoN_X also. That is, let us assume that X would like to have visibility over Y , but Y does not want to be visible to X . Given this, there are two possibilities: (i) Y 's restriction cannot be overruled and so X is not allowed to have visibility over Y , and (ii) X 's requirement has higher priority leading to negotiation with Y , and X gets visibility over Y . In addition, X 's visibility (noticeability) over (of) Y may be affected by the requirements and restrictions of the remaining nodes in H , global coherence/correlated properties, etc. Often (in a Web Services context), it is not the service provider itself but some higher level logical entity or an agent acting on behalf of the provider which is responsible for regulating the visibility of a provider. For example, with reference to the e-shopping scenario in Fig. 1, let us assume that $S-A$ would like to have visibility over the courier companies (such as, $C-B$) used by other suppliers to find the cheapest option. On the same lines, $S-B$ might like to keep the details of its courier company $C-B$ hidden due to competitive reasons, and should be in a position to reject $S-A$'s request for visibility over $C-B$. Resolving such conflicts is application/domain dependent and beyond the scope of this paper. Basically, we assume the existence of a nonconflicting set of visibility requirements for H .

The main contribution of this paper starts with a formalization of the visibility notion. The SoV definition is intuitive and encompasses "strong", "weak" and a variety of partially strong notions. The SoN definition follows naturally from that of SoV. We have identified (i) the dual properties of coherence and correlation, and its inverse and uniform variations, for both visibility and noticeability assignments, and (ii) how these properties relate the two spheres. We have discussed policies which exploit this interrelationship in different ways to associate meaningful visibility (noticeability) characteristics to the nodes of a hierarchy. In general, the SoV (and similarly, SoN) is represented by a set of graphs. We have identified some special cases where (i) the SoV of a node, (ii) the SoN of a node, and (iii) the SoV and SoN of all the nodes can be represented in a single graph.

It will be interesting to look for other nice visibility properties that can be applied to all hierarchical systems or at least to special cases. From an implementation perspective, we would also like to consider dynamic hierarchies where the structure of the hierarchy itself may change. We are also considering its practical applications with respect to non-functional aspects, such as, transactions, monitoring, security, etc.

References

- [1] C. Davies, Jr, Data processing spheres of control, *IBM Systems Journal*, 17(2) (1978) 179-198.
- [2] D. Biswas and K. Vidyasankar, Spheres of Visibility, in: *Proc. 3rd IEEE European Conference on Web Services (ECOWS '05)* (IEEE Computer Society Press) 2-13.
- [3] D. Biswas and K. Vidyasankar, Modeling Visibility in Hierarchical Systems, in: *Proc. 25th International Conference on Conceptual Modeling (ER '06)*, *Lecture Notes in Computer Science*, vol. 4215 (Springer, Berlin, 1985) 155-167.
- [4] D. Biswas and B. Genest, Compensation for Sequential Databases and Decomposition, Submitted Apr 2007.
- [5] F.-R. Lin, G. W. Tan and M.J. Shaw, Modeling Supply-Chain Networks by a Multi-Agent System, in: *Proc. 31st Annual Hawaii International Conference on System Science (HICSS '98)* (IEEE Computer Society Press) 105-114.
- [6] G. Alonso and C. Hagen, Exception Handling in Workflow Management Systems, *IEEE Transactions on Software Engineering* 26(10) (2000) 943-958.
- [7] G. Busatto, H.-J. Kreowski and S. Kuske, Abstract hierarchical graph transformation, *Mathematical Structures in Computer Science* 15 (2005) 773-819.
- [8] J. D. Moffet, Control principles and role hierarchies, in: *Proc. 3rd ACM Workshop on Role-Based Access Control (RBAC '98)* (ACM Press) 63-69.
- [9] L. Bocchi, Compositional Nested Long Running Transactions, in: *Proc. 7th International Conference on Fundamental Approaches to Software Engineering (FASE '04)*, *Lecture Notes in Computer Science*, vol. 2984 (Springer, Berlin, 1985) 194-208.
- [10] L. Fiege, Visibility in Event-Based Systems, Ph.D. Thesis, Department of Computer Science, Darmstadt University of Technology, Darmstadt, Germany, 2005.
- [11] M. Khambatti, K. Ryu and P. Dasgupta, Peer-to-peer Communities: Formation and Discovery, in: *Proc. International Conference on Parallel and Distributed Computing Systems (PDCS '02)* (ACTA Press) 161-166.
- [12] P. Yolum and M. P. Singh, Commitment Machines, in: *Proc. 8th International Workshop on Intelligent Agents (ATAL '01)*, *Lecture Notes in Computer Science*, vol. 2333 (Springer, Berlin, 1985) 235-247.
- [13] Q. Chen and U. Dayal, A Transactional Nested Process Management System, in: *Proc. 12th International Conference on Data Engineering (ICDE '96)* (IEEE Computer Society Press) 566-573.



Unité de recherche INRIA Rennes
IRISA, Campus universitaire de Beaulieu - 35042 Rennes Cedex (France)

Unité de recherche INRIA Futurs : Parc Club Orsay Université - ZAC des Vignes
4, rue Jacques Monod - 91893 ORSAY Cedex (France)

Unité de recherche INRIA Lorraine : LORIA, Technopôle de Nancy-Brabois - Campus scientifique
615, rue du Jardin Botanique - BP 101 - 54602 Villers-lès-Nancy Cedex (France)

Unité de recherche INRIA Rhône-Alpes : 655, avenue de l'Europe - 38334 Montbonnot Saint-Ismier (France)

Unité de recherche INRIA Rocquencourt : Domaine de Voluceau - Rocquencourt - BP 105 - 78153 Le Chesnay Cedex (France)

Unité de recherche INRIA Sophia Antipolis : 2004, route des Lucioles - BP 93 - 06902 Sophia Antipolis Cedex (France)

Éditeur
INRIA - Domaine de Voluceau - Rocquencourt, BP 105 - 78153 Le Chesnay Cedex (France)
<http://www.inria.fr>
ISSN 0249-6399