

# Modular Visualization of Distributed Systems

Gustavo A. Guevara S.<sup>1</sup> Jason LaPorte<sup>2</sup> Carlos A. Varela<sup>2</sup>

<sup>1</sup> Universidad Nacional de Colombia, Bogotá, Cundinamarca, COLOMBIA

<sup>2</sup> Rensselaer Polytechnic Institute, Troy, NY 12180, USA

<http://wcl.cs.rpi.edu>

**Abstract.** Effective visualization is critical to developing, analyzing, and optimizing distributed systems. We have developed OverView, a tool for online/offline distributed systems visualization, that enables modular layout mechanisms, so that different distributed system high-level programming abstractions such as actors or processes can be visualized in intuitive ways. OverView uses by default a *hierarchical concentric* layout that distinguishes *entities* from *containers* allowing migration patterns triggered by adaptive middleware to be visualized. In this paper, we develop a *force-directed* layout strategy that connects entities according to their communication patterns in order to directly exhibit the application communication topologies. In force-directed visualization, entities' locations are encoded with different colors to illustrate load balancing. We compare these layouts using quantitative metrics including communication to entity ratio, applied on common distributed application topologies. We conclude that *modular* visualization is necessary to effectively visualize distributed systems since no layout is best for all applications.

## 1 Introduction

Distributed systems promise increased computation power and capabilities for organizations at relatively low prices, although at a cost of increased complexity. This increased complexity rests on the shoulders of distributed application developers to be solved. Consider a scientist who would like to observe the interactions between computers while they run a high performance parallel simulation. Consider a middleware developer who would like to assess the locations of various parts of a distributed program, to ensure that the program's structure on the network is conducive to efficient inter-process communication. Consider visualizing computational activity of BOINC infrastructures[1], such as SETI@Home[2] or MilkyWay@Home [4]. Finally, consider visualizing the structure of a BitTorrent[3] swarm as its activity spikes and dwindles over time. Visualizing distributed systems, whose reconfiguration is not intuitive or easy to keep track of, is critical to assess the efficiency and performance of distributed applications.

The usual approach to distributed system understanding, involves attaching a debugger to each of the locations in which the system is executing, forcing the individual analysis of each of the locations of execution in the system. Progress has been made in terms of the visualization of distributed systems, with the

creation of tools like Pip, WiDS, Jinsight[9], Jive[11], Hy+[5] and OverView[6]. OverView provides a more complete insight to a distributed program’s execution, given it’s capabilities for online/offline visualization and high level abstraction representations, as noted by Desell et al.[6]. Accurate visualization of distributed systems provides scientists and developers with tools for proper assessment of the distributed system’s performance in a large scale, taking all the execution locations into account. More particularly a visualization should provide information as to the proper co-location of actors on different locations.

Given the different network topologies that constitute different distributed systems, it is the purpose of this paper to determine the suitability of a modular approach proposed for OverView[6] to achieve this broad perspective. Two visualization modules are proposed, *hierarchical concentric* and *force-directed*, which are evaluated to find their limitations and to stress the convenience of OverView’s modular architecture to develop suitable visualizations for different types of applications. OverView’s architecture is described, as well as it’s byte-code instrumenting features and visualization module development examples.

The hierarchical concentric module proves effective in representing topologies resulting from structured overlay networks, such as Chord [10] because of it’s ring-like appearance and effective conveyance of the logarithmic number of references. This is further discussed in Section 3. When the network is deployed in a manner that requires actors to send large amounts of messages to each other, visualization becomes more difficult when using the hierarchical concentric module, because the representation becomes cluttered and messages can not be distinguished from each other. The force-directed module increases the amount of actors that can be effectively visualized in some of these topologies with a large communication to entity ratio since space is more efficiently distributed by inferring attraction/repulsion laws from communication between entities as discussed in Section 2. Analytical and experimental analysis of both modules results in the determination of each model’s effectivity at conveying distributed systems in theoretical and real scenarios, as discussed in Section 4.

OverView has been used to visualize various distributed environments, most notably by instrumenting the SALSA programming language [12]<sup>3</sup>. The examination of the existing modules as they are plugged in to OverView, demonstrates the simple nature of using OverView’s event processing approach to develop visualization modules, in order to convey the nature of a distributed application as developers or scientists require.

**Related work** Frishman and Tal[8] have developed a visualization tool which bears a number of similarities to our own. While it is limited to using mobile objects (a single model of distributed systems, and thus is less generic), it takes an interesting approach to scalability. It suggests providing the user a means to select one or more points of focus; that is, parts of the visualization that are of interest. Then, the software visualization will use an algorithm to determine which mobile objects it can filter out as being uninteresting to the user. This

<sup>3</sup> <http://wcl.cs.rpi.edu/salsa/>

stands in contrast to its negative: the approach of selecting the objects one wishes to filter out. Such a means might provide a better-scaling visualization; ultimately, there will always be fewer entities one is interested in than entities one is not.

Another software visualization framework which is similar to OverView is EVolve[13], having the notion of events which describe the runtime behavior of an application being translated into a graphical visualization. It focuses upon object-oriented computations, supporting visualization of method invocations; while its visualizations are less intuitive and generic, they are scalable and provide a large amount of information content, and may even be layered naturally on top of each other to show the user different types of information at once.

## 2 OverView

OverView[6] is a toolkit which permits visualization of Java-based systems; in particular, distributed systems such as those previously described.<sup>4</sup> The toolkit includes three programs, each of which performs a different task (see Figure 1):

- The *OverView Instrumenter*, or OVI, which allows the abstraction of a Java program’s execution into a set of visualizable events by inserting unobtrusive event-sending behavior into existing Java bytecode.
- The *OverView Presenter*, or OVP, which receives and interprets events into a meaningful, interactive, graphical representation of the state of the distributed system. OVP has several *visualization modules*, each of which can display the distributed system in a different layout.
- The *OverView Daemon*, or OVD, which acts as an event relay, collecting events sent by *event sources* (that is, any active instrumented program), and forwarding those events to *event sinks* (that is, any listening visualization program).

Users add event-sending behavior to any existing Java program by writing an *Entity Specification Language* (ESL) file, which uses a simple, declarative syntax to map Java method invocations to OverView events. These events are sent, at run-time, over a network to a listening OVD. ESL’s simple grammar definition (see [6] for details) allows the instrumentation of existing Java bytecode.

OVP can listen for incoming events from multiple sources, both online network connections and offline log files. It will multiplex these events and forward them to a visualization module, in addition to time stamping and logging them, so that they might be played back at a later date if desired. OverView visualization modules are written using the Processing Development Environment<sup>5</sup>.

---

<sup>4</sup> OverView versions 0.1 and 0.2 were developed as a plugin for the Eclipse IDE. OverView’s current version (0.5) is stand-alone to better support future features such as running the visualization from within a web browser as a Java applet.

<sup>5</sup> <http://www.processing.org/>

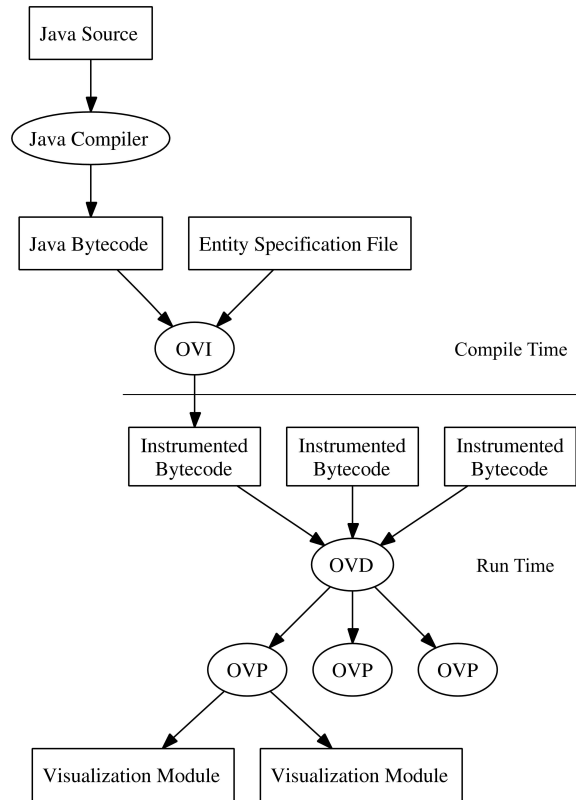


Fig. 1: The OverView framework, showing both compile-time operations and the layout of the run-time architecture.

## 2.1 Events

OverView’s visualization model is based on two fundamental units, called *entities* and *containers*. An entity embodies the concept of a discrete unit of computation, which could refer to an object, an actor, an ambient, a process, or even a virtual machine. A container refers to the environment in which an entity exists. Every OverView visualization is composed of some aggregation of these two basic elements.

- **Position/1:** Tells OverView to create a particular entity, or to move one if it already exists, outside of any container.
- **Position/2:** Tells OverView to create a particular entity inside a particular container. Again, if the entity already exists, it is moved instead of being created. If the container does not yet exist, it is created and placed outside any other container.
- **Deletion/1:** Tells OverView to delete a particular entity. If such an entity does not exist, nothing happens.

- **Communication/2:** Tells OverView that two entities have communicated (for example, via method invocation or message passing).

## 2.2 Visualization Modules

A visualization module for OverView can be developed to suit the application to be visualized, since every OVP allows the connection of new modules that render information to the scientist or developer’s particular preferences. The approach involves to implement the `Renderer` interface provided by OverView which will allow to implement `handle` and `unhandle` methods, as well as processing’s `draw` method. `RingVisualization.java` which can be found in OverView’s website<sup>6</sup> is a good example for the creation of visualization modules.

**Hierarchical Concentric** The first module, called *hierarchical concentric*, is a visualization in which all top level entities and containers are arranged in a ring around the center of the screen, each scaling to fit if necessary. Those which are containers are differentiated from entities by being drawn as a square, rather than a circle; furthermore, any entities or containers it contains will be arranged in a circle within it. Upon communication a line is drawn between two communicating entities, disappearing momentarily. Following is a description of the actions triggered by each of the received events:

- **Entity Creation:** After a check for its existence an entity will be created only if it does not exist inside or outside a container.
- **Entity Communication:** A line is drawn between two entities that exist, the origin located on the message emitter and the destination on the message receiver.
- **Entity Migration:** Upon removal from its current location the entity will be created in another location. To represent this transition, the interface displays the entity as it moves along the interface from one location to its destination, this is represented through the use of animation easing.
- **Entity Deletion:** The entity is removed from the interface.

**Force-directed** In *force-directed* visualization, entities and containers are distinguished through color-coding. While only entities are drawn, containers are distinguished by the entity’s color. A color convention is provided on the screen. Entities are not positionally constrained, they move depending on three external forces.

- **Entity Creation:** After a check for its existence an entity will be drawn in a color that depends on its container.
- **Entity Communication:** A line is drawn between two entities that exist, the origin located on the message emitter and the destination on the message receiver. This line represents the exertion of attraction force described in detail later.

<sup>6</sup> <http://wcl.cs.rpi.edu/overview>

- **Entity Migration:** Upon removal from its current location the entity will be created in another location. To represent this transition, the interface displays the entity as it changes its color to that assigned randomly at visualization initialization to its new container.
- **Entity Deletion:** The entity is removed from the interface along with its connections to other entities.

*Entity Dynamics:* There is a force of repulsion between entities. They repel each other following an inverse square rule, like Coulomb’s law for electrostatics, as follows:

$$F_r = K_r \frac{1}{r^2} \quad (1)$$

where  $r$  represents the distance between two given entities and  $K_r$  represents a repulsion coefficient. There is attraction between entities modeling communication as a force, rendering communicating entities as connected by a spring, which follows Hooke’s Law:

$$F_a = -K_a r \quad (2)$$

where similarly,  $r$  represents the distance between two entities and  $K_a$  represents the attraction coefficient. The space through which they move provides a constant friction which acts as a damper. This results in a particle system that displays the entities in a co-location scheme influenced directly by communication between entities or lack thereof.

### 3 Visualizing Common Communication Patterns

A measure of efficient use of space can be proposed by setting a fixed *communication/entity* ratio and measure how many entities can be displayed without communications giving out false information. (*e.g.*, communications between A and B passing through uninvolved entity C or a communication being crossed at more than 2 points). These highly intolerant thresholds are set in order to provide a general basis of comparison, a user could define his/her own *clarity thresholds*.

$$r = \frac{\text{Communications}}{\text{Entities}} \quad (3)$$

$C_x$  represents the amount of points along a communication’s trajectory through which it is crossed by another.

$$C_x < 3 \quad (4)$$

$E_x$  represents the amount of entities that are crossed over by a communication.

$$E_x = 0 \quad (5)$$

Force-directed visualization is also evaluated by measuring the amount of time it will take to reach a stable point, that is, when the velocity of the system is below an arbitrary value that tends to zero referred to as *stability threshold* which is determined in terms of (*pixels/cycle*). A cycle is measured to be  $1,781 \cdot 10^{-2} s$

(using a 2000 cycle average measurement), the stability threshold was defined at  $1 \cdot 10^{-5}$  pixels/cycle which is considered a system velocity low enough to allow for comprehensive reading of the visualization. This is done to measure the cost of increased clarity in those representations where force-directed visualization performs significantly better in terms of the clarity threshold.

Herein, we examine the performance of the hierarchic concentric and force-directed visualizations based on their meeting of the clarity thresholds in different communication/entity ratio settings, the latter provided by different topologies.

### 3.1 Linear and Grid-based Topologies:

Linear and grid-based topologies are characterized by having a number of nodes set up in some  $n$ -dimensional lattice, with each node connected to and communicating with its neighbors. The most common class of programs that use such a structure are physical simulation programs, such as heat distribution or fluid dynamics applications, where each node represents an area in space (see Figure 2.)

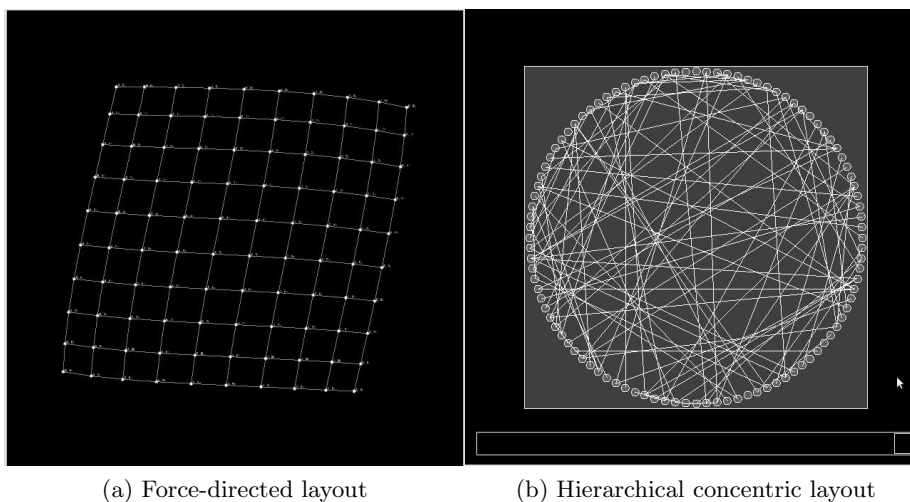


Fig. 2: Snapshot of a grid simulation containing 100 entities using OverView.

*Performance:* Grids are better represented by force-directed visualizations because they do not display cluttered entities and the communication links are clear, the *clarity threshold* is met since there are no crossing lines or crossed entities. Hierarchic visualizations, highly dependent on the order of entity creation do not meet the clarity threshold. The stability threshold was reached by the force-directed approach in an average 80.523 seconds.

### 3.2 Ring-based Topologies

Ring-based topologies find a common application in peer-to-peer networks (see e.g., Figure 3.) Chord[10] increases the number of neighbor links from a constant number to a logarithmic number (based on the size of the network), which decreases the search time from linear to logarithmic as well; this makes it possible to construct robust and high-performance peer-to-peer networks (see Figure 4.).

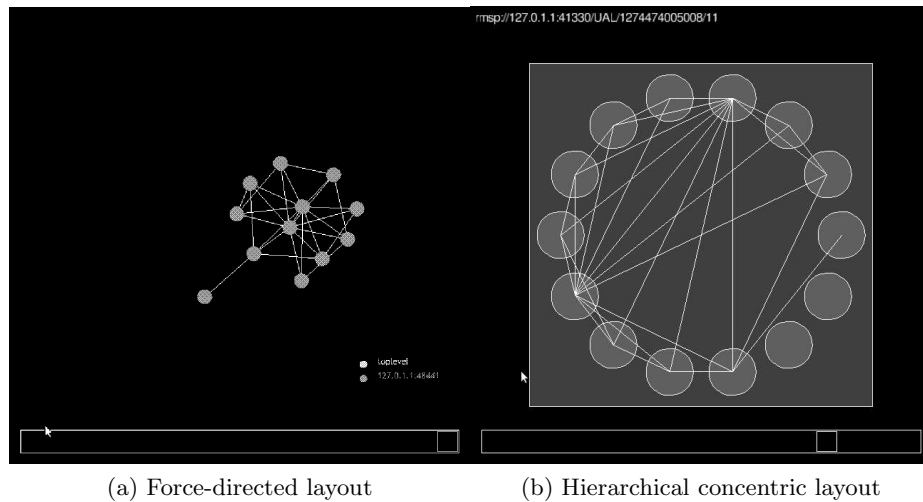


Fig. 3: Flavius Josephus problem where each third man in the circle commits suicide, the solution for the "safe" position is visualized using OverView.


*Performance:* Both models fail to meet the clarity threshold, however once the threshold is lowered, the first model to reach compliance of it, is the hierarchical concentric. The dense communication pattern for Chord is not efficiently viewed using force-directed approaches. The stability threshold is reached at an average 4.202 seconds.

### 3.3 Cube and Hypercube-based Topologies

Cubes, hypercubes, and other  $n$ -dimensional cubes have a number of notable properties useful to distributed computation; among them is the upper bound on hops between any two nodes on the network, while maintaining only a small number of edges between nodes. Hypercubes are also notable for having a very structured high-dimensional organization, which is difficult to map to 2 or 3 dimensions well (see Figure 5.)

*Performance:* This topology does not meet the clarity threshold on either visualization, a less objective evaluation would imply that if the scale is small





(a) A force-directed representation. (b) A hierarchical concentric representation.

Fig. 4: Chord networks visualized using OverView show that the clarity threshold in the force-directed approach is violated by far, while the hierarchical concentric approach provides an effective means of visualization.

enough, a Force Directed approach could be intelligible. This topology lies beyond the scope of both visualizations.

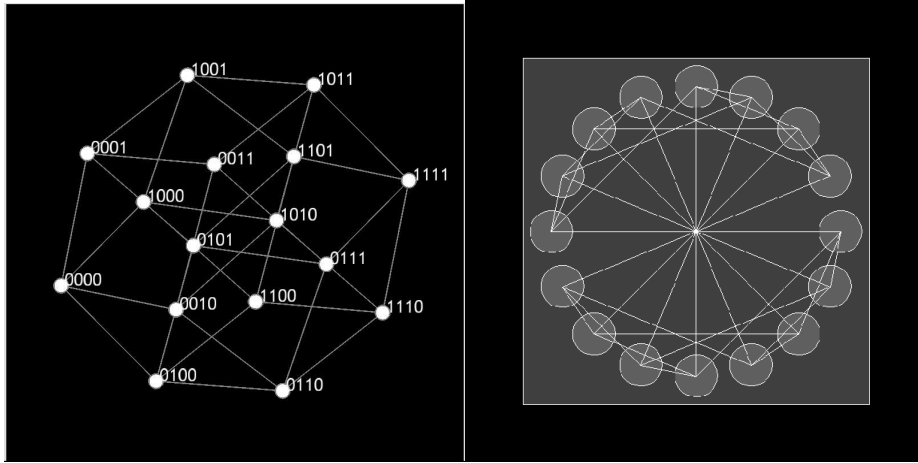
### 3.4 Recursive Topologies

Recursive computation is as important in distributed computing as it is in sequential computing; many divide-and-conquer recursive algorithms can be made to run efficiently in a distributed setting (see Figure 6.)

*Performance:* When nodes are created with a particular ordering, the hierarchical concentric visualization meets the clarity threshold on binary trees which are common in recursive algorithms. A force directed approach does not rely on the order of creation to meet the clarity threshold, and allows to recognize the fractal nature of the calculation. No entity crossing or communication crossing occurs once the system is at a stable position. The stability threshold is reached at average in 27.962 seconds.

## 4 Discussion on Visualization Models

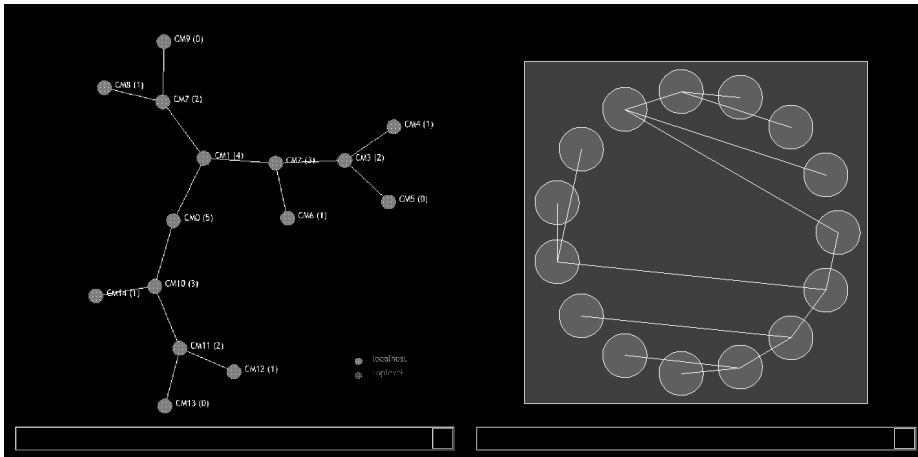
Figure 7 shows a table containing the data obtained for each of the topologies evaluated in Section 3.



(a) A force-directed layout.

(b) A hierarchical concentric layout.

Fig. 5: A Hypercube network is visualized with OverView.



(a) A force-directed layout.

(b) A hierarchical concentric layout

Fig. 6: The fractal nature of the recursive calculation is observed using OverView.

	<i>Grid</i>	<i>Ring – based</i>	<i>Hypercube</i>	<i>Recursive</i>
Entities	100	45	32	15
Communications	180	150	16	14
Communication/Entity ratio	1.8	3.3	2.0	0.93
Reached stability threshold	80.523	4.202	6.845	27.962
Clarity thresholds cleared (hierarchical)	No	No	No	Yes
Clarity thresholds (force-directed)	Yes	No	No	Yes

Fig. 7: Performance of visualization modules on different application classes.

#### 4.1 Hierarchical Concentric Visualization

Hierarchical concentric places entities around a hierarchical series of rings. This placement mechanism was selected for several reasons: firstly, it minimizes ambiguous communications—that is, communication lines overlapping multiple entities, so one cannot determine at a glance which entities are referred to; this is a commonly-arising problem in grid-based visualizations, since many communications will either be purely horizontal or purely vertical and reach across entire rows or columns. Secondly, it can display a hierarchy of elements intuitively, even with naïve placement of new elements at the end of the list, allowing it to support recursive computations with reasonable effectiveness. This hierarchical, concentric structure is both intuitive and generic, mapping naturally both to flat systems (such as actors within theaters) and nested systems (such as mobile ambients).

Hierarchical concentric’s main drawback is in scalability. While it is effective for small numbers of entities, it rapidly breaks down and becomes confusing with as few as one hundred entities, with many communications filling the rings. Furthermore, logically grouped entities may be placed far from each other by the visualization, and since one cannot zoom in upon both sides of the ring at once, it can become difficult or impossible to understand what is occurring in parts of a system. However, despite these weaknesses, the visualization can be very effective for some application topologies, such as linear and circular topologies, as illustrated in Section 3.

#### 4.2 Force-directed Visualization

Force-directed visualization was primarily developed to remedy some of the shortcomings of hierarchical concentric. Unlike it, there is no hierarchical placement; entities are only placed upon the primary visualization surface. The only reference as to an entity’s belonging to a particular container is the color used to draw it. Furthermore, if like colors are well clustered, it provides good information on whether entities that communicate frequently are near each other physically, which is a very natural notion for the visualization to convey. An outlier to this affirmation is the case of farmer-worker topologies, where such information is not accurately conveyed since there is no communication between

workers (see Figure 8a.) Given this scenario, a user can interact with the visualization by click-n-drag actions, allowing for manual configuration of entities. We note, however, that the force-directed layout provides a very intuitive visualization in most cases, especially in reference to the application communication topology.

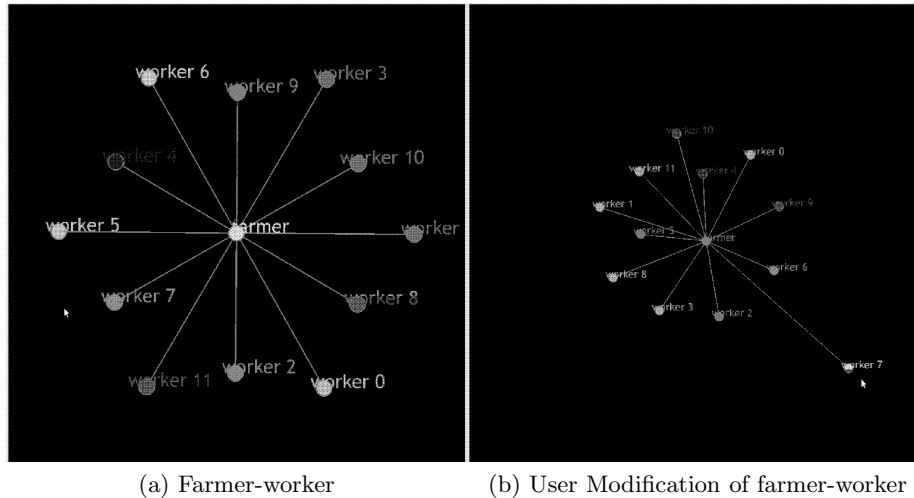


Fig. 8: A farmer-worker topology does not involve workers communicating with each other, on these conditions force-directed approaches are not advisable for visualization, on 8b the user interacts with the visualization by dragging one of the workers.

The force-directed layout lacks support for nesting which is one of its most significant weaknesses, however it allows to visualize the appearance of clustering between entities that are tightly coupled in terms of communication which is a good indicator of proper application of load-balancing policies.

## 5 Conclusions

Our experiences with these visualizations indicate clearly and beyond a doubt that a modular approach to OverView is the correct choice, since different applications have different visualization needs. Furthermore, our experience developing these modules for OverView leads us to believe that a developer or scientist can better use a customized visualization module. This possibility of visualization renders the subjective nature of visualization clarity and effectiveness to a simple matter of instrumenting bytecode according to the events that are intended to be visualized and if the existing modules do not convey information in a satisfactory manner, the task of creating a visualization module is relatively simple.

OverView’s hierarchical concentric module provides a way to convey the nesting of entities in a clear fashion, finding its boundaries when communication patterns are dense. It approaches the issue by having the representation of a communication event fade over time, which could lead to misrepresented persistent links. The force-directed module focuses its attention on the communication topology of an application at a given point, which allows to evaluate the load balancing policies of adaptive middleware or the load balancing algorithms implemented within the application. An approach to migration representation is proposed within this module with the color coding of containers, however this technique is not as effective since there are topologies which do not present communication between co-located entities.

The powerful nature of OverView’s architecture is put to the test with the development of the force-directed approach that looked to the problem that arises from the hierarchical concentric module not conveying coupling. The combined usage of these two visualization modules provides the means for comprehensive evaluation of an application’s performance. Furthermore, it paves the way for more visualization modules to be developed in order to gain perspective on other particular characteristics of a distributed system.

## 5.1 Future Work

The most obvious next step to take in the direction of distributed systems visualization using OverView is to produce modules that focus on other characteristics of distributed systems whose visualization could prove useful. For instance, a module that focuses on visualizing the different policies available in adaptive middleware to distribute autonomous actors. Such visualizations would permit users to fully comprehend the behavior induced by the policies selected for the system at hand (*e.g.* energy management, high performance, task management, etc.).

Extensive work on the matter has spawned a generation of middleware applications that specialize on load balancing[7] and the evaluation of such software is critical. Upon deployment of middleware, the task of load balancing the system becomes transparent to the developer, which in turn raises the question of whether the middleware selected as well as its policies are optimal given the application-level objectives.

Additional OverView visualization modules that can clearly separate physical-layer network aspects such as inter-processor latencies and bandwidths for particular clusters, from application-layer aspects would enable a better assessment of middleware resource management policies.

## Acknowledgements

This work has been partially supported by the following grants: NSF CAREER CNS Award No. 0448407, IBM Innovation Awards.

## References

1. D. P. Anderson, *Boinc: a system for public-resource computing and storage*, Proceedings of the Fifth IEEE/ACM International Workshop on Grid Computing., IEEE Computer Society, 2004, pp. 4–10.
2. David P. Anderson, Jeff Cobb, Eric Korpela, Matt Lebofsky, and Dan Werthimer, *Seti@home: an experiment in public-resource computing*, Commun. ACM **45** (2002), no. 11, 56–61.
3. Bram Cohen, *Incentives build robustness in BitTorrent*, Tech. report, May 2003.
4. Nathan Cole, Travis Desell, Daniel Lombranaa Gonzalez, Francisco Fernandez de Vega, Malik Magdon-Ismail, Heidi Newberg, , Boleslaw K. Szymanski, and Carlos A. Varela, *Parallel and computational intelligence*, ch. Evolutionary Algorithms on Volunteer Computing Platforms: The MilkyWay@Home Project, pp. 64–90, Springer-Verlag Berlin Heidelberg, 2010.
5. Mariano P. Consens, Masum Z. Hasan, and Alberto O. Mendelzon, *Using hy+ for network management and distributed debugging*, CASCON '93: Proceedings of the 1993 conference of the Centre for Advanced Studies on Collaborative research, IBM Press, 1993, pp. 450–471.
6. Travis Desell, H. Iyer, Abe Stephens, and Carlos A. Varela, *OverView: A framework for generic online visualization of distributed systems*, Proceedings of the European Joint Conferences on Theory and Practice of Software (ETAPS 2004), eclipse Technology eXchange (eTX) Workshop (Barcelona, Spain), March 2004.
7. Travis Desell, Kaoutar El Maghraoui, and Carlos A. Varela, *Load balancing of autonomous actors over dynamic networks*, Proceedings of the Hawaii International Conference on System Sciences, HICSS-37 Software Technology Track, January 2004, pp. 1–10.
8. Yaniv Frishman and Ayellet Tal, *Visualization of mobile object environments*, SoftVis '05: Proceedings of the 2005 ACM symposium on Software visualization (New York, NY, USA), ACM Press, 2005, pp. 145–154.
9. Wim De Pauw and John M. Vlissides, *Visualizing object-oriented programs with jinsight*, ECOOP '98: Workshop on Object-Oriented Technology (London, UK), Springer-Verlag, 1998, pp. 541–542.
10. Ion Stoica, Robert Morris, David Karger, M. Frans Kaashoek, and Hari Balakrishnan, *Chord: A scalable peer-to-peer lookup service for internet applications*, Proceedings of the 2001 ACM SIGCOMM Conference, 2001, pp. 149–160.
11. James T. Streib and Takako Soma, *Using contour diagrams and jive to illustrate object-oriented semantics in the java programming language*, SIGCSE '10: Proceedings of the 41st ACM technical symposium on Computer science education (New York, NY, USA), ACM, 2010, pp. 510–514.
12. Carlos A. Varela and Gul Agha, *Programming dynamically reconfigurable open systems with SALSAs*, ACM SIGPLAN Notices. OOPSLA'2001 Intriguing Technology Track Proceedings **36** (2001), no. 12, 20–34.
13. Qin Wang, Wei Wang, Rhodes Brown, Karel Driesen, Bruno Dufour, Laurie Hendren, and Clark Verbrugge, *Evolve: an open extensible software visualization framework*, SoftVis '03: Proceedings of the 2003 ACM symposium on Software visualization (New York, NY, USA), ACM Press, 2003, pp. 37–ff.