
Improving a Distributed Software System's Quality of Service via Redeployment

Nenad Medvidovic


(joint work with Sam Malek, Chiyoung Seo, and Marija Mikic-Rakic)

Computer Science Department
University of Southern California
Los Angeles, CA 90089

nen@usc.edu

<http://sunset.usc.edu/~nen/>

Deployment Architecture and QoS



- *Deployment Architecture*: allocation of s/w components to h/w hosts
- h^c deployment architectures are possible for a given system
 - same functionality
 - different qualities of service (QoS)

Problem in a Nutshell

■ Guiding Insight

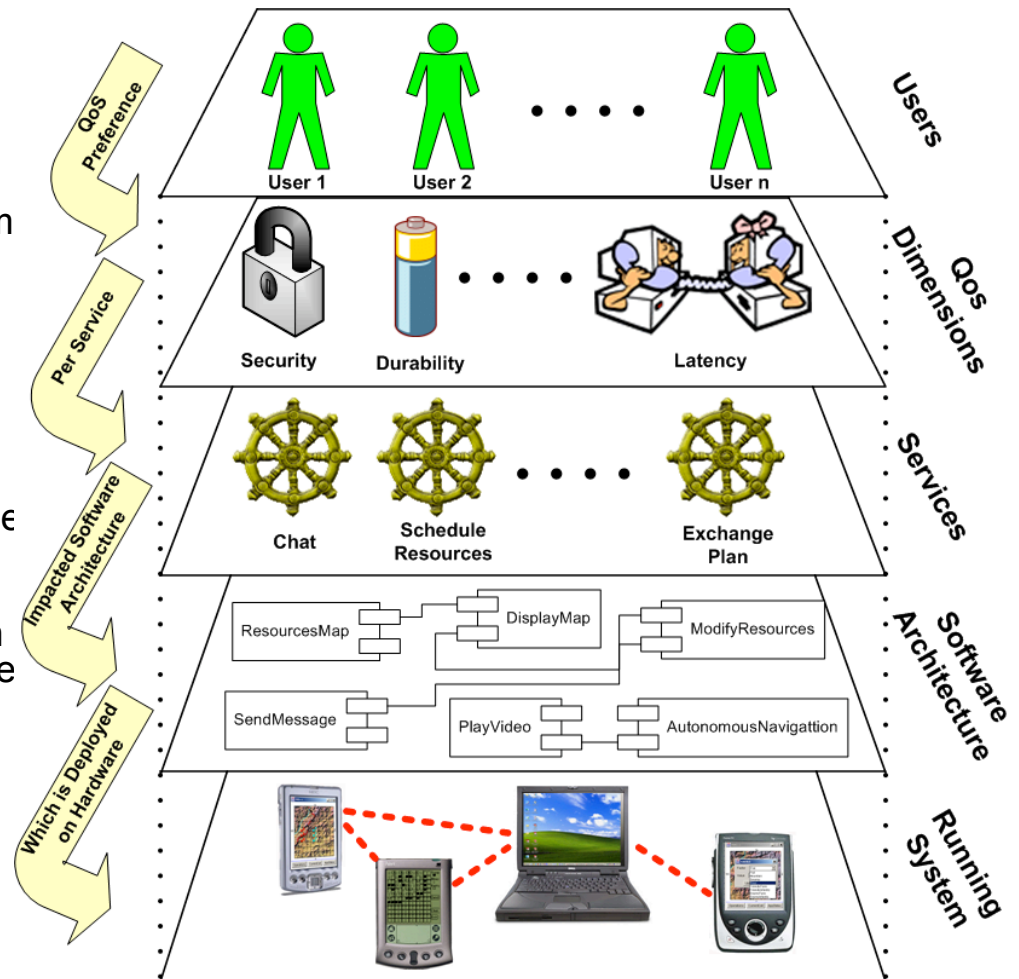
- **System users have varying QoS preferences for the system services they access**
 - Impacts their satisfaction with the system

■ Research Question

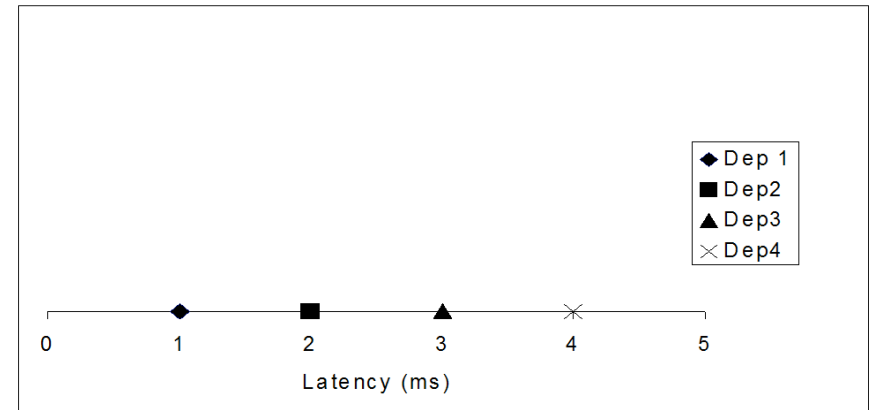
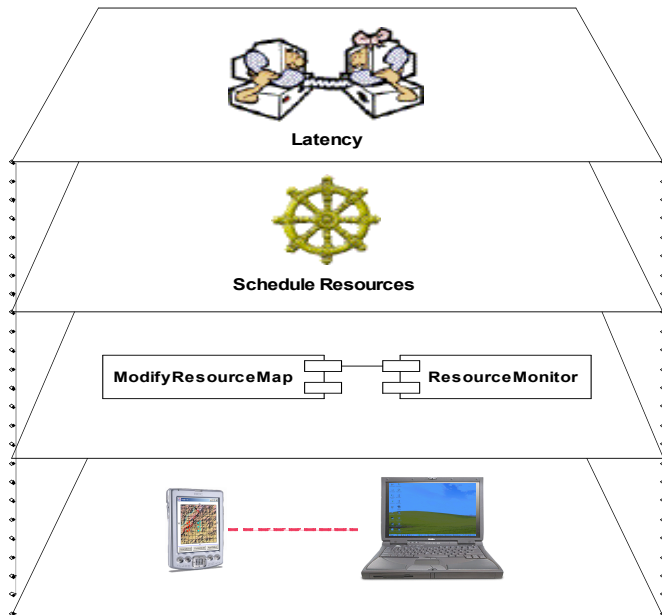
- **How could we improve system's deployment architecture to maximize users' satisfaction?**
 - Where users' satisfaction depends on the system's ability to meet their QoS preferences
 - And where other possible solutions such as caching, hoarding, replication, etc. are not appropriate or ideal

■ Research Objective

- **Devise a solution that is applicable to many classes of application scenarios**

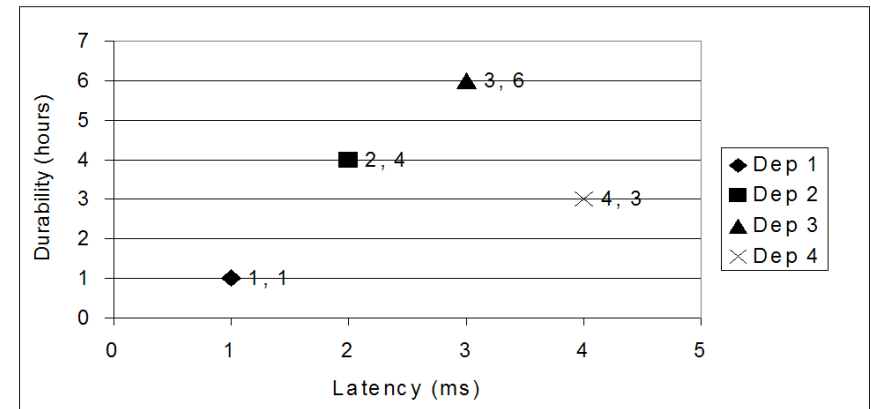
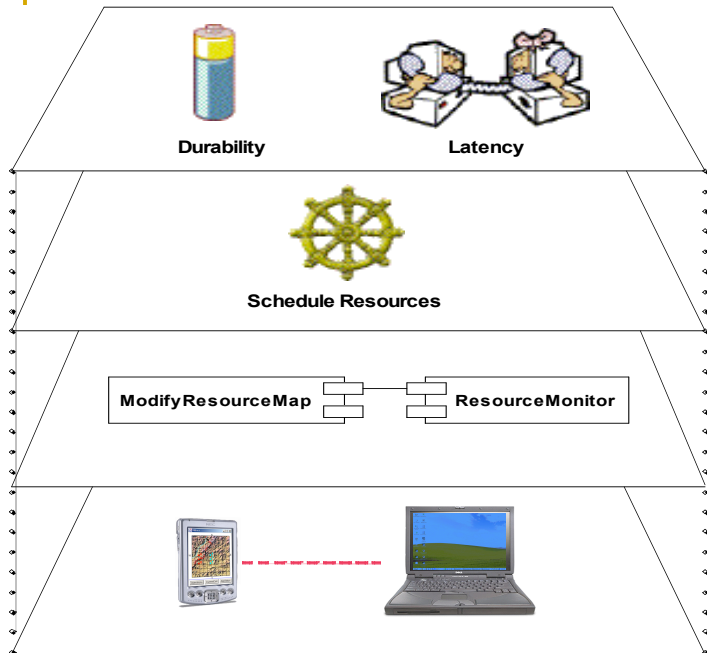


Scenario with a Single QoS Dimension



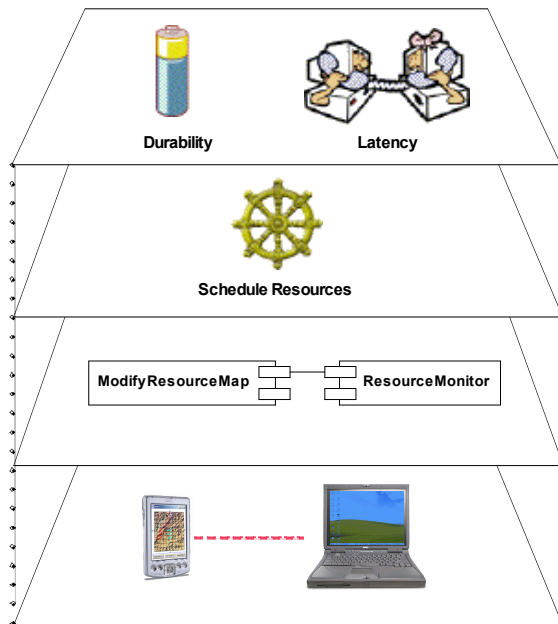
- Objective is to minimize latency
- The *optimal* deployment architecture is deployment 1
- Most all related approaches stop here

Conflicting QoS Dimensions




- Objective is to minimize latency and maximize durability
- There is no optimal deployment architecture!
- Phenomenon known as *Pareto Optimal* in multidimensional optimization


Resolving Trade-Offs between QoS Dimensions



- Explicitly consider
 - system *users*
 - system's *utility* to its users




- A *utility function* denotes a user's preferences for a given rate of improvement in a QoS dimension






- Allows expression of multidimensional optimization in terms of a single scalar value

A Slightly Larger Scenario



- Troop, Latency, Exchange Plan
- Troop, Latency, Schedule Resources
- Troop, Durability, Exchange Plan
- Troop, Durability, Schedule Resources
- Troop, Security, Exchange Plan
- Troop, Security, Schedule Resources
- Commander, Latency, Exchange Plan
- Commander, Latency, Schedule Resources
- Commander, Durability, Exchange Plan
- Commander, Durability, Schedule Resources
- Commander, Security, Exchange Plan
- Commander, Security, Schedule Resources
- Dispatcher, Latency, Exchange Plan
- Dispatcher, Latency, Schedule Resources
- Dispatcher, Durability, Exchange Plan



onsi
fere

Eye-balling the solution quickly becomes impossible!

Proposed Solution

A framework that provides

 an extensible system model

- inclusion of arbitrary system parameters
- definition of QoS dimensions using the parameters
- specification of users' QoS preferences

 multiple QoS improvement algorithms

- different algorithms suited to different classes of systems

 extensible tool support

- deployment, execution, and runtime redeployment
- parameter monitoring and visualization

Proposed Solution

A framework that provides

an extensible system model

- inclusion of arbitrary system parameters
- definition of QoS dimensions using the parameters
- specification of users' QoS preferences

multiple QoS improvement algorithms


- different algorithm suited to different classes of systems

extensible tool support

- deployment, execution, and runtime redeployment
- parameter monitoring and visualization

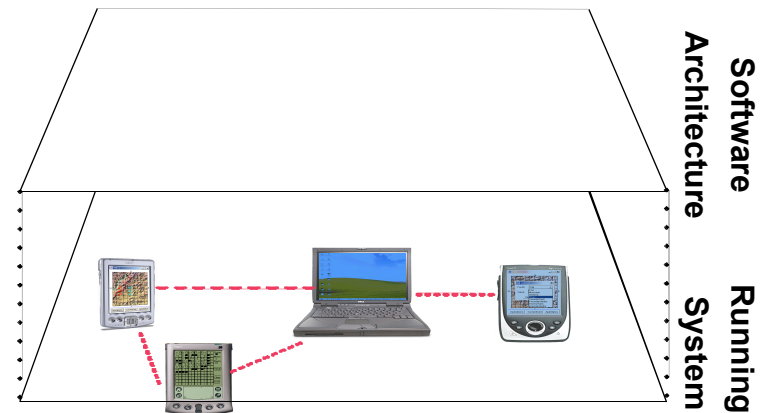
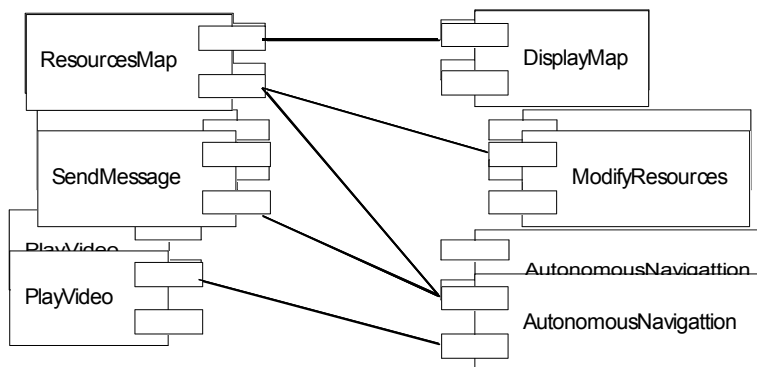
Model of the Hardware System

- A set H of hardware nodes
 - $H=\{PDA1, PDA2, PDA3, Laptop\}$
- A set HP of host parameters
 - $HP=\{memory, battery\}$
- A function $hParam:H\times HP\rightarrow R$
 - $hParam(PDA1, memory)=20MB$
- A set N of network links
 - $N=\{Link1, Link2, Link3, Link4\}$
- A set NP of network link parameters
 - $NP=\{reliability, bandwidth\}$
- A function $nParam:N\times NP\rightarrow R$
 - $nParam(Link1, bandwidth)=256kb/s$



Model of the Software Architecture

- A set C of software components
 - $C = \{ResourcesMap, DisplayMap, \dots\}$
- A set CP of component parameters
 - $CP = \{size, CPU\ usage\}$
- A function $cParam: C \times CP \rightarrow R$
 - $cParam(DisplayMap, size) = 50Kb$
- A set I of logical links
 - $I = \{renderMap, updateMap, \dots\}$
- A set IP of logical link parameters
 - $IP = \{frequency, average\ event\ size, \dots\}$
- A function $IParam: I \times IP \rightarrow R$
 - $IParam(renderMap, frequency) = 20$
- A set $DepSpace = \{d1, d2, \dots\}$ of all possible deployment mappings



Model of the System Services

- A set S of service
 - $S = \{Chat, Scheduler Resources, Exchange Plan\}$
- A function $sParam: S \times \{H \cup C \cup N \cup I\} \times \{HP \cup CP \cup NP \cup IP\} \rightarrow R$ of values for service-specific system parameters
 - $sParam(Schedule Resources, renderMap, frequency of execution) = 3$



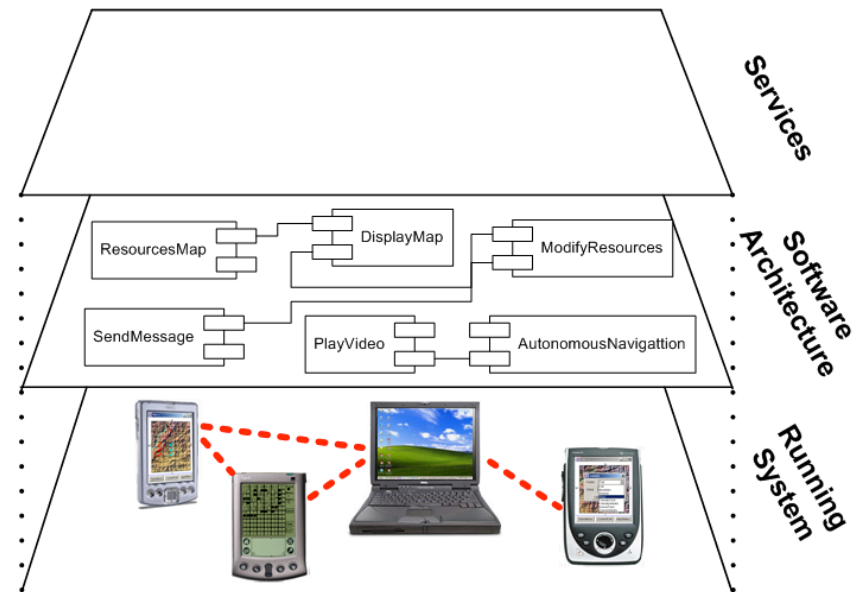
Chat



Schedule Resources



Exchange Plan



Model of the QoS Dimensions


- A set Q of quality of service dimensions
 - $Q = \{\text{security, durability, latency}\}$
- A function $qValue: S \times Q \times DepSpace \rightarrow R$ that quantifies the achieved level of QoS
 - $qValue(\text{chat, latency, } d1) = 5ms$
- A function $qType: Q \rightarrow \{-1, 1\}$
 - -1 denotes it is desirable to minimize the QoS
 - 1 denotes it is desirable to maximize the QoS



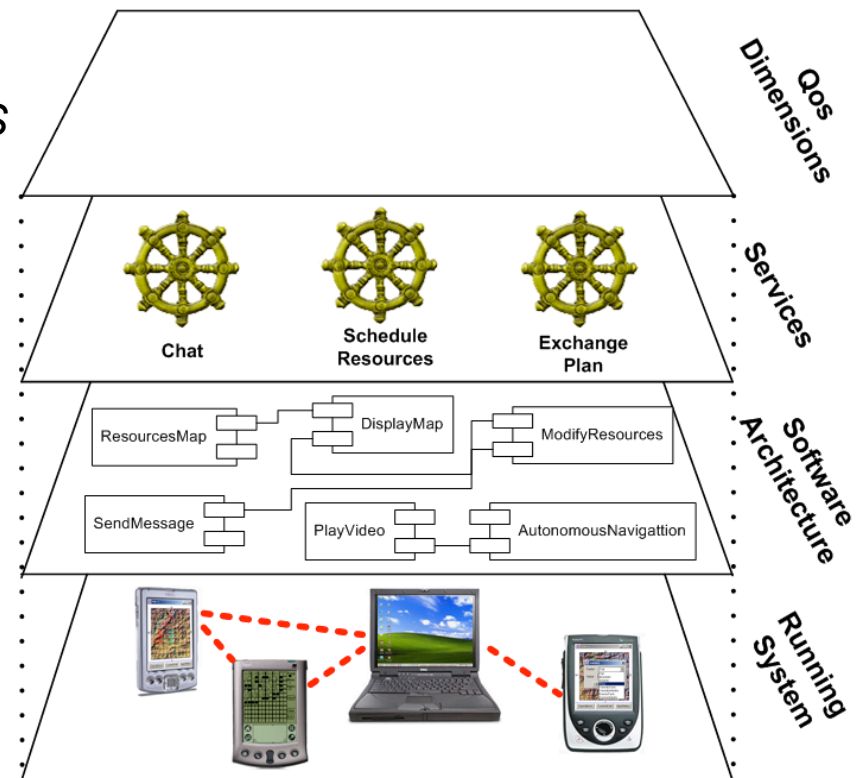
Security



Durability

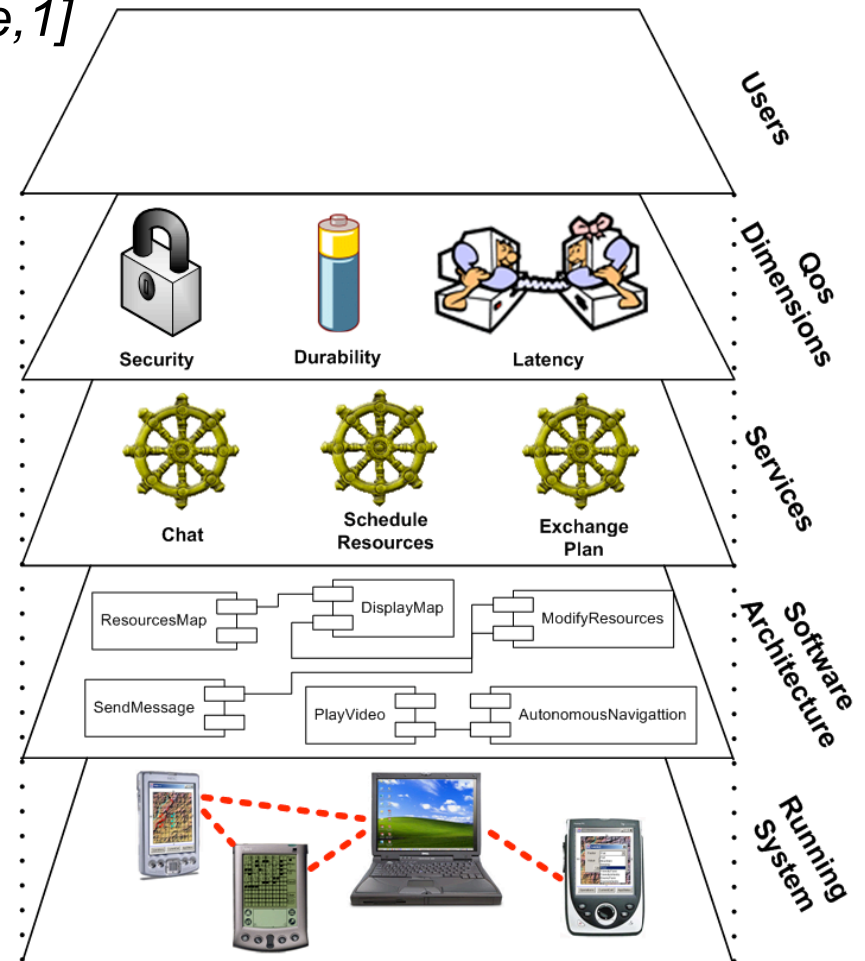


Latency



Model of the System Users

- A set U of users
 - $U = \{Troop, Commander, Dispatcher\}$
- A function $qosRate: U \times S \times Q \rightarrow [MinRate, 1]$
 - represents the rate of change in QoS
- A complementary function $qosUtil: U \times S \times Q \rightarrow [0, MaxUtil]$
 - represents the utility for that rate of change
- A user's priority can be expressed as the ratio of $MaxUtil$ to $MinRate$

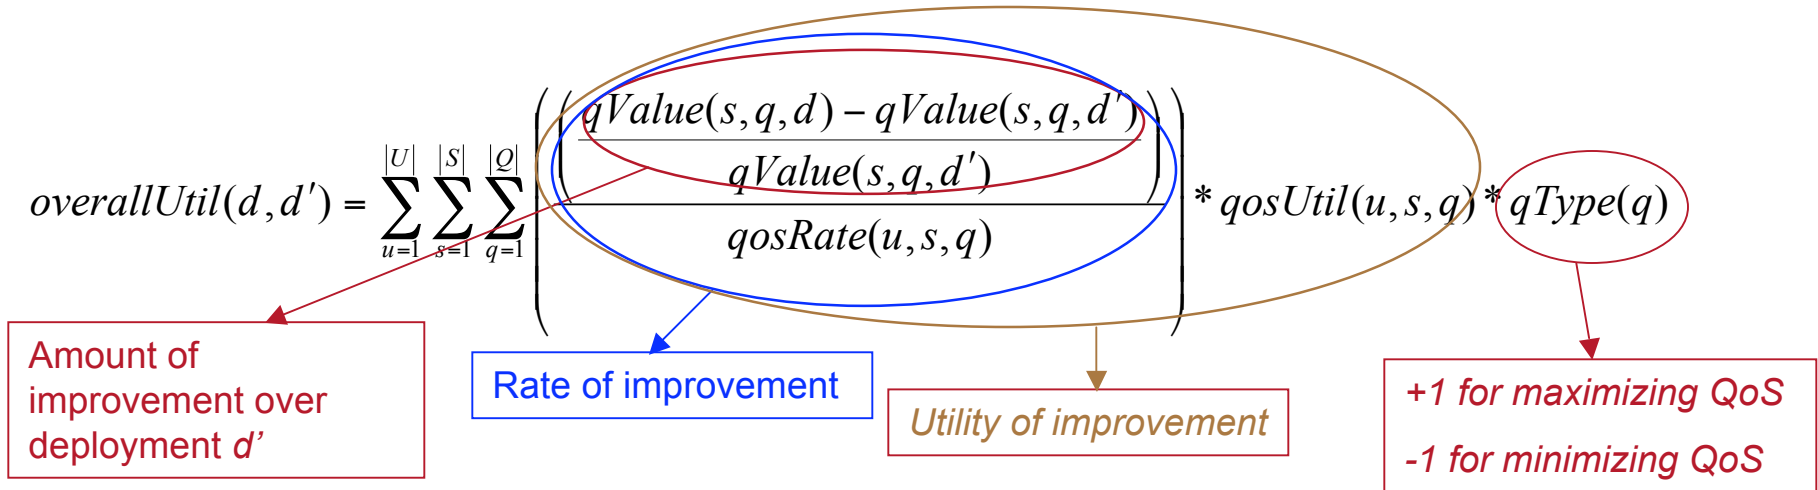


Model of the Constraints

- A set PC of parameter constraints
 - $PC = \{memory, bandwidth, \dots\}$
- A function $pcSatisfied: PC \times DepSpace \rightarrow [0, 1]$
 - 1 if constraint is satisfied
 - 0 if constraint is not satisfied
- Functions that restrict locations of software components
 - $loc: C \times H \rightarrow [0, 1]$
 - $loc(c, h) = 1$ if c can be deployed on h
 - $loc(c, h) = 0$ if c cannot be deployed on h
 - $colloc: C \times C \rightarrow [-1, 1]$
 - $colloc(c1, c2) = 1$ if $c1$ has to be on the same host as $c2$
 - $colloc(c1, c2) = -1$ if $c1$ cannot be on the same host as $c2$
 - $colloc(c1, c2) = 0$ if there are no restrictions
- ...

Problem Definition

Given the current deployment of the system d' , find an improved deployment d such that the users' overall utility defined as the function



is maximized and specific conditions are satisfied:

- $\forall c \in C, loc(c, H_c) = 1$ \longrightarrow All location constraints are satisfied
- $\forall c1 \in C, \forall c2 \in C,$
 if $(colloc(c1, c2) = 1) \rightarrow (H_{c1} = H_{c2}),$
 if $(colloc(c1, c2) = -1) \rightarrow (H_{c1} \neq H_{c2})$ \longrightarrow All collocation constraints are satisfied
- $\forall constr \in PC, pcSatisfied(constr, d) = 1$ \longrightarrow All system parameter constraints are satisfied
- ...

Framework Instantiation

- The engineer needs to specify the “loosely” defined elements of the model

 Define the pertinent properties of the application scenario

 Define QoS dimensions in terms of system properties

$$qValue(s, availability, d) = \sum_{c1=1}^{C_s} \sum_{c2=1}^{C_s} sParam(s, I_{c1,c2}, freq) * nParam(N_{H_{c1}, H_{c2}}, rel)$$

 Define system parameter constraints

Proposed Solution

A framework that provides

 an extensible system model

- inclusion of arbitrary system parameters
- definition of QoS dimensions using the parameters
- specification of users' QoS preferences






 **multiple QoS improvement algorithms**

- different algorithm suited to different classes of systems

 extensible tool support

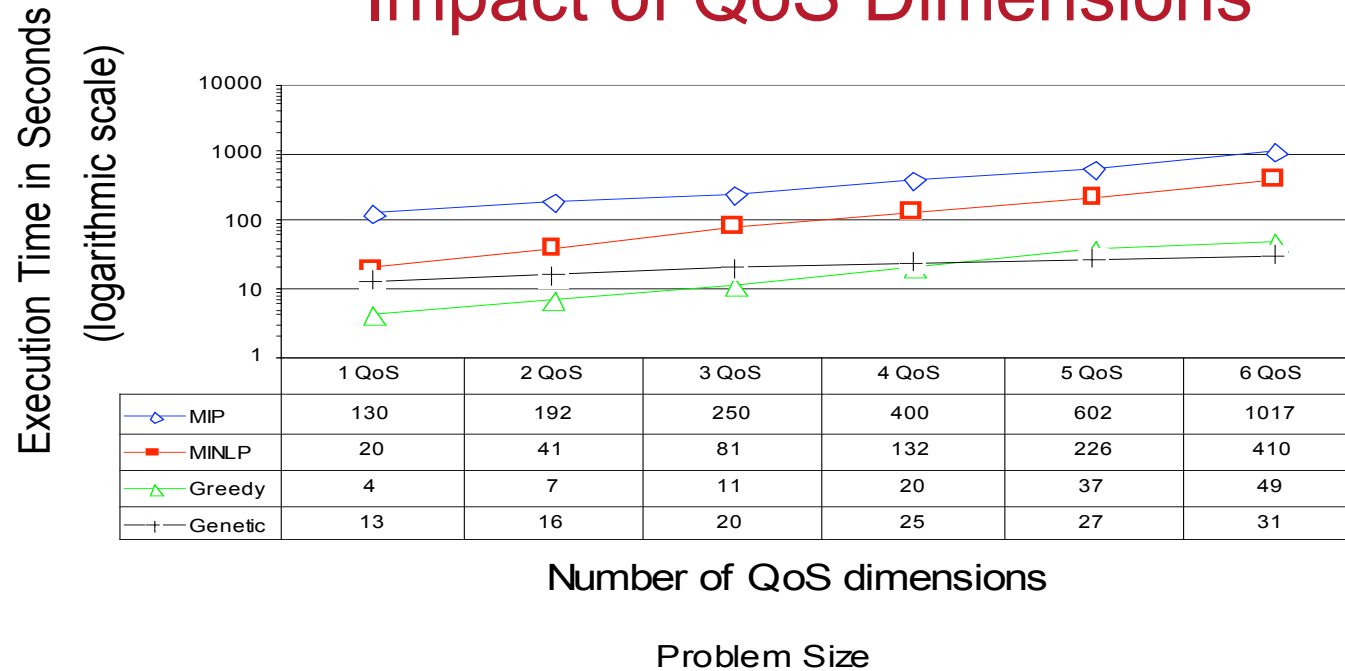
- deployment, execution, and runtime redeployment
- parameter monitoring and visualization

Algorithms

-  **MINLP – polynomial (?)**
 - Represented the problem as a set of (non-)linear constraint functions
 - Does not guarantee the optimal solution or convergence
-  **MIP – exponential: $O(2^{|H|} |C|^2)$**
 - Devised an approach to transform our MINLP problem to MIP
 - Developed heuristics to decrease complexity to $O(|H|^{|C|})$
-  **Greedy – polynomial: $O(|S|^3 (|C| |U| |Q|)^2)$**
 - An iterative algorithm that leverages several heuristics for
 - Ranking elements of our problem (services, hosts, components, ...)
 - Assigning software components to hardware hosts
 - Makes local decisions that often maximize the global objective
-  **Genetic – linear: $O(\#populations \times \#evolutions \times \#individuals \times |S| |U| |Q|)$**
 - An individual represents a solution composed of a sequence of genes
 - A population contains a pool of individuals which are evolved via cross-overs and mutations
 - The accuracy on the representation depends on the ability to promote “good” genes
 - Bad representation does not promote “good” genes → random search
-  **Market-Based (Auctioning)**
 - Under development and evaluation

Algorithms' Performance


Impact of QoS Dimensions



Impact of Heuristics

Mapping in Genetic

Objective Function Value



Problem Size

—◇— Genetic without mapping	13	57	73	65	102
—■— Genetic with mapping	48	150	130	192	285
—△— Genetic with mapping and three parallel executing populations	51	156	132	210	294

Algorithms in Practice

QoS	MIP				MINLP				Greedy				Genetic			
	Avail.	Latency	Comm. Security	Energy Cons.	Avail.	Latency	Comm. Security	Energy Cons.	Avail.	Latency	Comm. Security	Energy Cons.	Avail.	Latency	Comm. Security	Energy Cons.
service 1	56%	-8%	18%	-8%	33%	2%	-5%	14%	24%	-8%	4%	-4%	16%	-2%	18%	-8%
service 2	93%	94%	97%	24%	91%	41%	32%	24%	83%	91%	62%	15%	93%	84%	35%	18%
service 3	39%	30%	22%	49%	32%	38%	11%	69%	39%	30%	22%	49%	19%	30%	22%	49%
service 4	215%	97%	302%	7%	215%	97%	302%	7%	165%	50%	220%	12%	180%	91%	150%	10%
service 5	59%	7%	25%	26%	23%	5%	39%	21%	43%	7%	19%	18%	29%	5%	35%	33%
service 6	99%	55%	37%	44%	83%	35%	45%	32%	99%	55%	37%	44%	99%	55%	37%	44%
service 7	91%	57%	20%	47%	97%	29%	44%	25%	91%	37%	14%	23%	91%	43%	4%	49%
service 8	43%	22%	7%	56%	41%	11%	-5%	72%	32%	21%	-10%	58%	13%	51%	7%	72%
Average	86%	44%	66%	30%	76%	32%	57%	33%	72%	35%	46%	26%	67%	44%	38%	33%
overallUtil	64				57				55				52			

- Results of running the algorithms on an example scenario of 12 Comps, 5 Hosts, 8 Services, and 8 Users
- Significant improvements for all the four QoS dimensions by all the algorithms
- The more important QoS dimensions of services have improved significantly more than others

Algorithmic Trade-Offs

- Architectural style
 - E.g., Client-Server vs. Peer-to-Peer
 - MIP algorithm for very constrained architectures
 - One of the optimization algorithms for flexible and large architectures
- Large number of QoS dimensions
 - Genetic outperforms the greedy
 - Genetic is only linearly affected by the number of QoS dimensions
- Stable vs. unstable systems
 - For small and stable systems, MIP algorithm is worth the time and resources required to compute a solution
 - For large and unstable systems, genetic or greedy is more applicable
- Resource constrained systems
 - Genetic algorithm can execute in parallel on multiple devices
 - Sharing the overhead among many hosts
- Centralized vs. decentralized systems
 - Market-based algorithms could also be leveraged in a decentralized setting

Proposed Solution

A framework that provides

 an extensible system model

- inclusion of arbitrary system parameters
- definition of QoS dimensions using the parameters
- specification of users' QoS preferences

 multiple QoS improvement algorithms

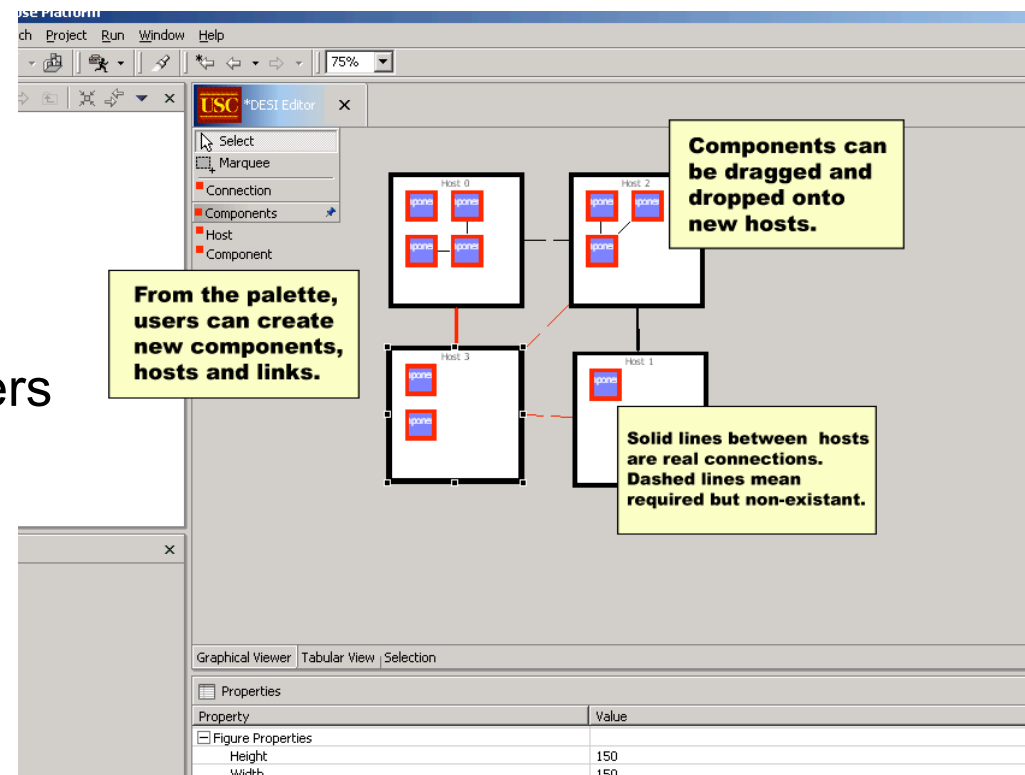
- different algorithm suited to different classes of systems

 **extensible tool support**

- deployment, execution, and runtime redeployment
- parameter monitoring and visualization

Modeling and Analysis Support – DeSi

- DeSi is a visual environment for analyzing deployment architectures
- It allows for modeling a distributed system in terms of four basic elements
 - Software components
 - Hardware devices
 - Network links
 - Logical (interaction) links
- Each of these elements has an associated set of parameters
 - Accessed via property sheets
- DeSi is extensible
 - Allows for modeling of new parameters and properties
 - Views are completely separated from the model



DeSi – Control Panel

The screenshot shows the DeSi Editor Eclipse Platform interface. The main window is titled "Resource - DESI Editor - Eclipse Platform". The interface is divided into several sections:

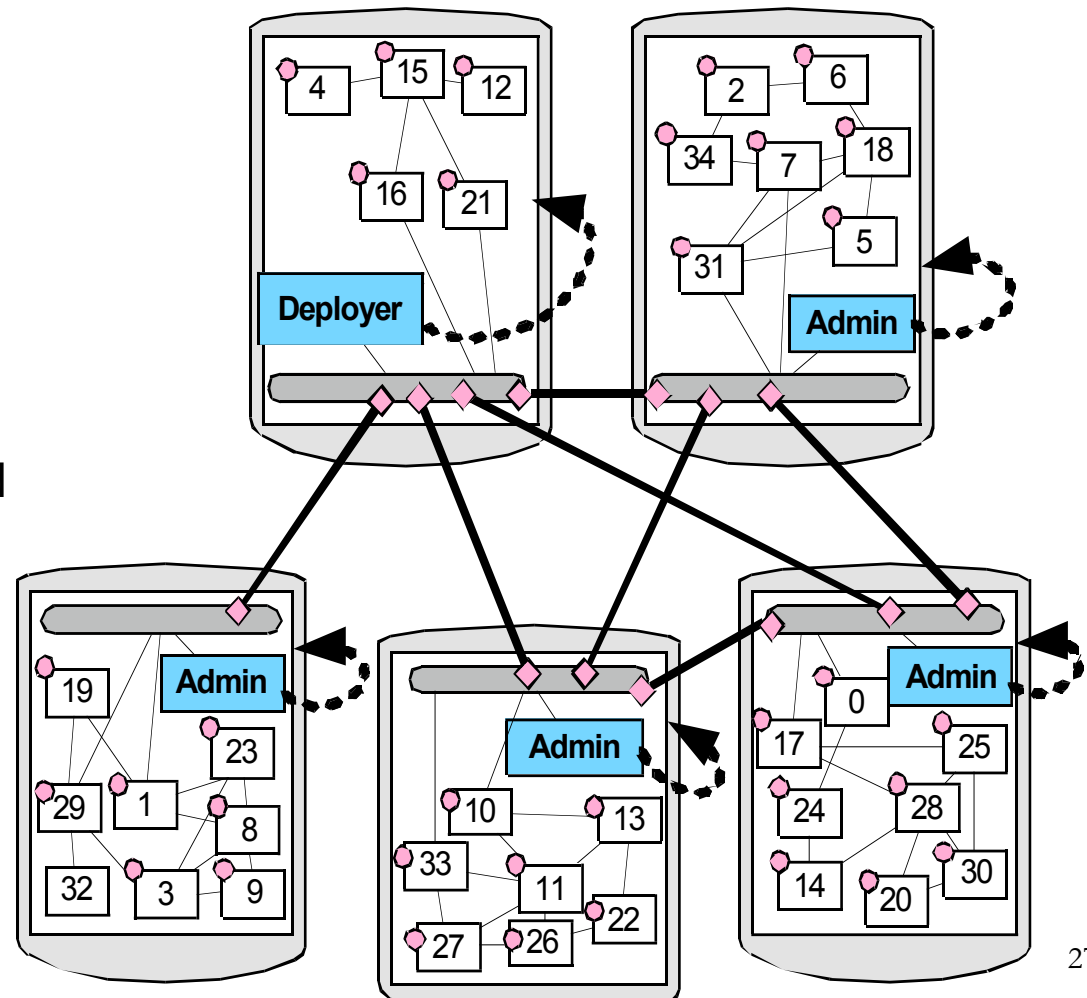
- Input:** A list of configuration parameters such as "Number of components: 10", "Number of hosts: 4", "Minimum comp. memory (in KB): 2", etc.
- Constraints:** A table mapping components to hosts. A yellow callout box states: "Components can be collocated or restricted to/from certain hosts." Buttons for "On the same host", "Not on the same host", and "Fix to host" are present.
- Algorithms:** A section for selecting algorithms. A dropdown menu is open, showing options: "Just run", "Just run and preview", "Run and effect", and "Unbiased Stochastic". Other options include "Biased Stochastic", "Greedy Approximation", "Clustered", "Level of awareness", "Decentralized Iterative", and "Benchmark". A yellow callout box states: "Here algorithms are run and their performance can be benchmarked." A "Revert to current deployment" button is at the bottom.
- Availability:** A large yellow callout box states: "New systems can be generated according to specified properties." Below it, the availability is shown as "Availability: 0.491368".
- Tables of parameters:** A section with tabs for "Hosts: reliability and memory", "Comps: frequency and memory", "Hosts: bandwidth", and "Comps: avg. e". The "Hosts: reliability and memory" tab is active, showing a table:

Host/Host	0	1	2	3
0	1.0	0.0	0.0	0.808
1	0.0	1.0	0.889	0.0
2	0.0	0.889	1.0	0.0
3	0.808	0.0	0.0	1.0

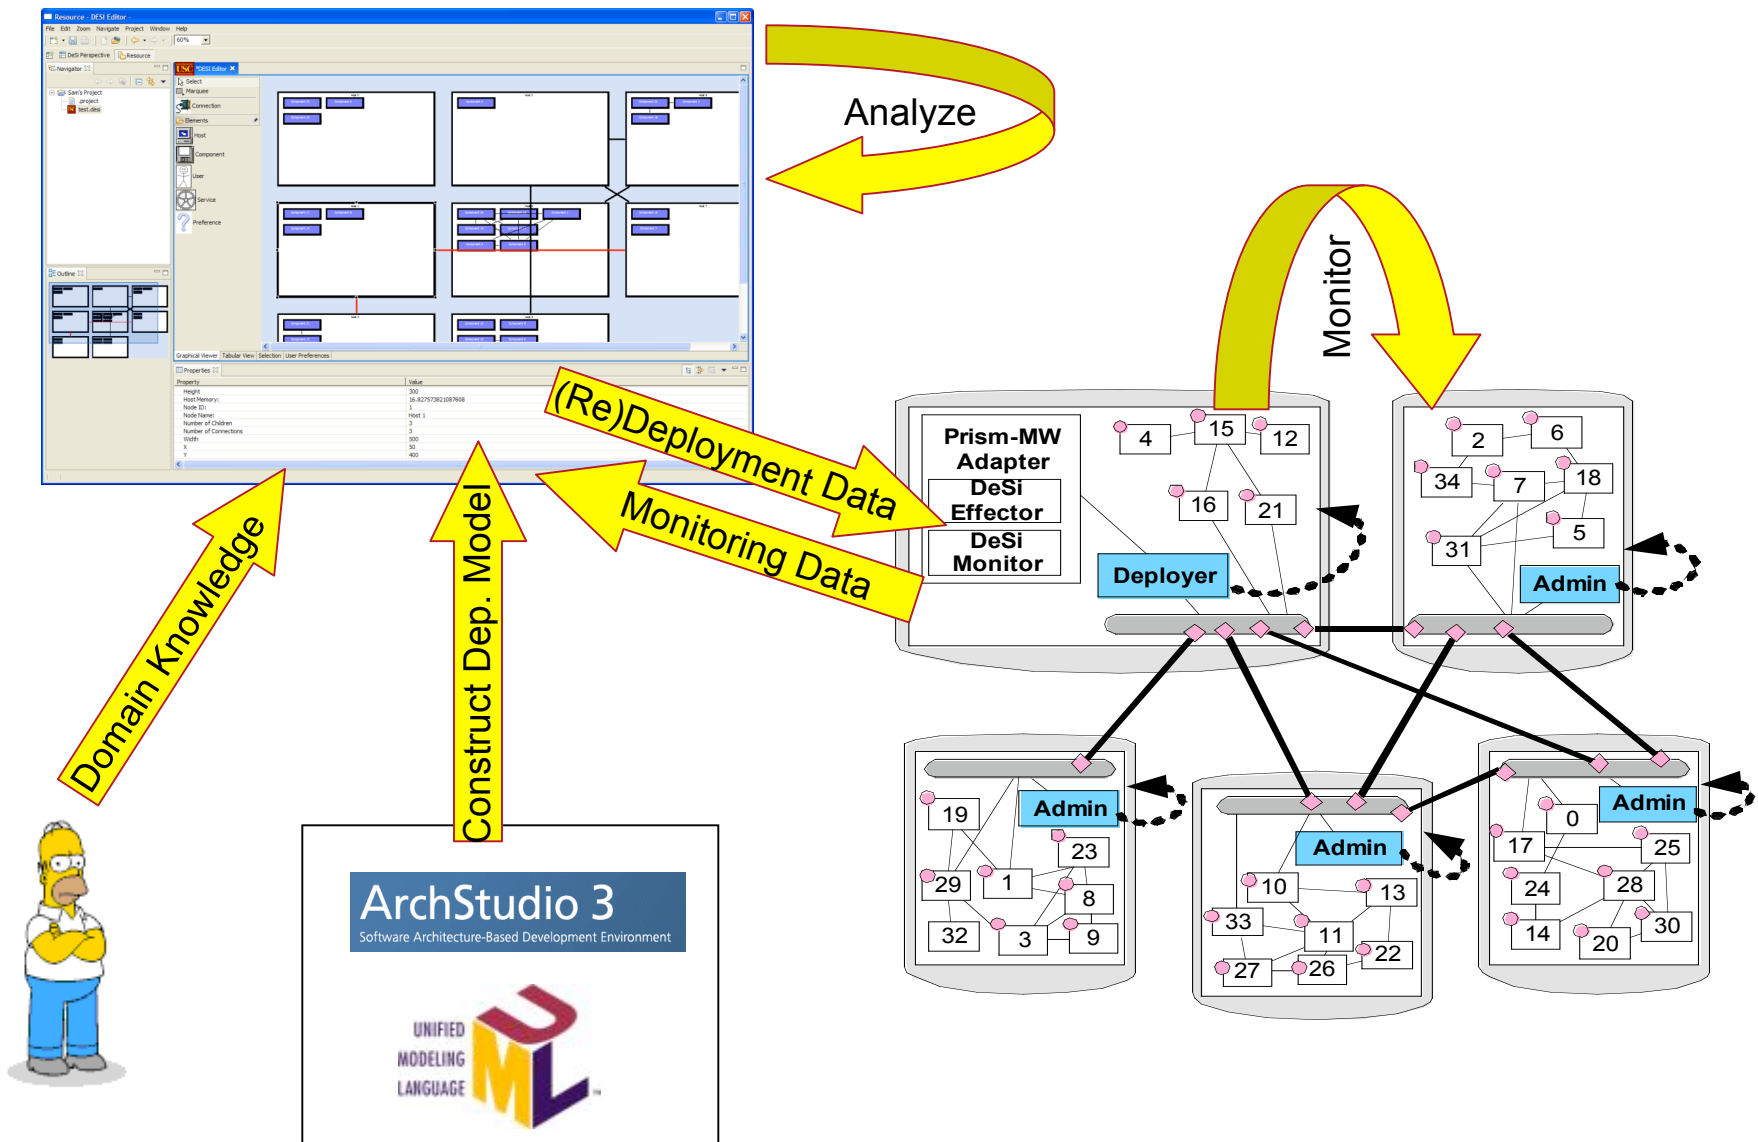
- Results:** A table showing the results of the simulation. A yellow callout box states: "Here we can see the details of our system and the overall system availability." The table has columns: "Component", "Initial d...", "Exact", "Unbias...", "Biased ...", "Greedy", "Decent...", and "Clustered".

Implementation and Execution Support – Prism-MW

- Prism-MW is an extensible architectural middleware
 - PL-level constructs architectural concepts
 - components
 - connectors
 - ports, etc.
 - Facilities for monitoring and (re)deployment of a distributed system
 - Allows for the addition of new monitoring and deployment facilities



Tool Suite Integration



Contributions

- Address system deployment as a multidimensional optimization problem
 - Leverages users' preferences to resolve inherent trade-offs in conflicting QoS dimensions
- Explicitly consider system's high-level services and their internal architecture
- An extensible modeling approach that can be leveraged across different application scenarios
 - Specify arbitrary system parameters
 - Define arbitrary QoS dimensions in terms of system parameters
- A suite of generic multidimensional optimization algorithms
 - Operate on top of an instantiated model of a system
- A suite of customizable tools
 - A number of extension points are leveraged to configure the tools to the application scenario at hand
 - Promotes reuse and cross-evaluation of solutions to this problem

On-Going Work

- Further profiling of the algorithms
 - Determine which algorithms are suitable to what classes of systems
- Several on-going enhancements to DeSi
 - Addition of new modeling elements: users, user preferences, services, etc.
- Complete the integration of Prism-MW, DeSi, and ArchStudio
- Develop the support for autonomically selecting appropriate redeployment algorithms
- Evaluate the approach on real distributed systems
 - Troops Deployment System (TDS)
 - Midas

Questions

