

# REX: A Development Platform and Online Learning Approach for Runtime Emergent Software Systems



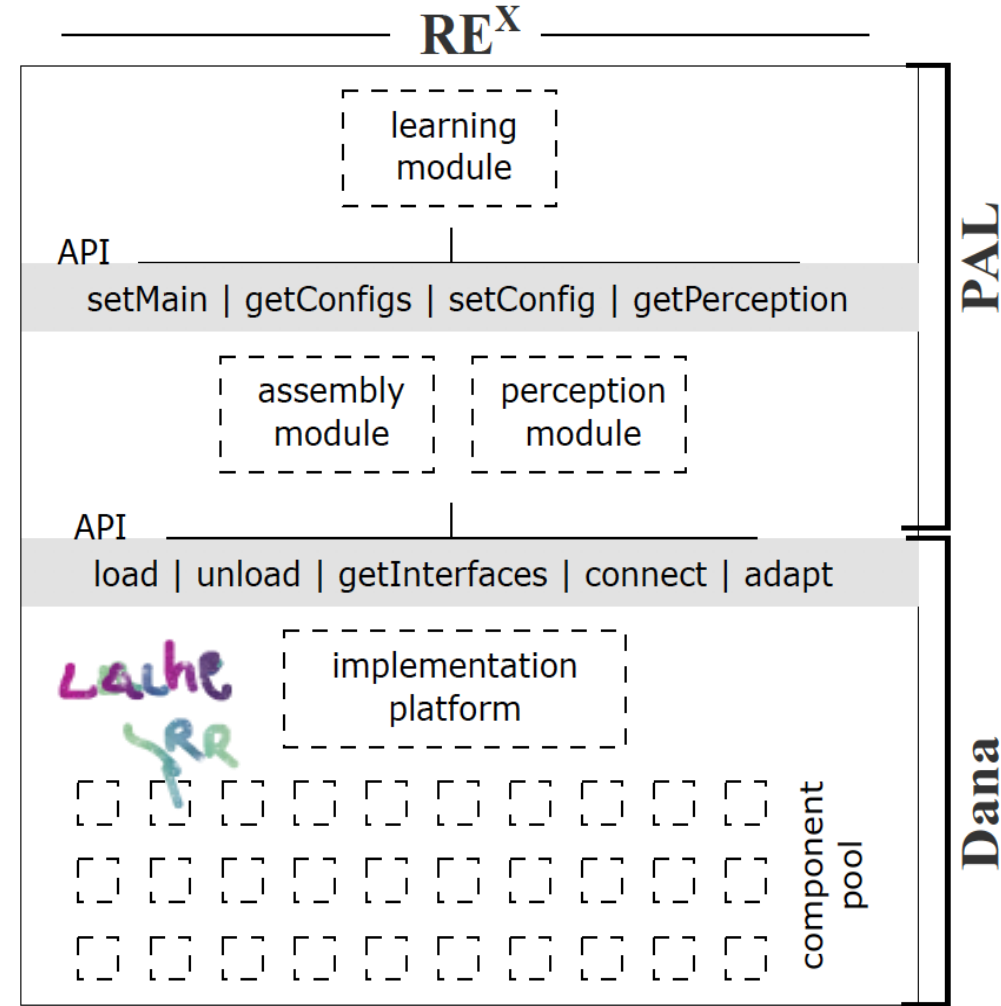
Barry Porter,  
Matthew Grieves,  
Roberto Filho  
David Leslie

# Introduction

- Designing, Analyzing and Maintaining – Millions of LOC:  
Is it sustainable?
- Software development: Models, Policies, and Processes
- Autonomous, Self-Adaptive, and Self Organized Software System
- Emergence of Software System - Autonomously from pool of available building blocks
- Responsive to actual runtime conditions.
- Can show rationale behind the choice

# REX: Development Platform

- Implementation Platform-  
Dana (Dynamic Adaptive Nucleic Architecture)
- PAL Framework
  - Perceive – Internal + External conditions
  - Assemble and Re-assemble modules
  - Learn
- Online Learning – Statistical Linear Bandits, using Thompson Sampling.



# Dana

- Component Based Software Paradigm
- All Components – Runtime Replaceable
- Multi-threaded imperative language (what and how)

## Example: From Source code

```
component provides App requires io.Output out{  
    int App:main(AppParam params[]){  
        out.println("Hi! :-)")  
        return 0  
    }  
}
```

```
interface File {  
    transfer char path[]  
    transfer int pos, mode  
  
    File(char path[], int mode)  
    byte[] read(int numBytes)  
    int write(byte data[])  
    bool eof()  
    void close()  
}
```

```
component provides App requires File {  
    int App:main(AppParam args[]) {  
        File ifd = new File(args[0].str, File.READ)  
        File ofd = new File(args[1].str, File.WRITE)  
        while (!ifd.eof()) ofd.write(ifd.read(128))  
        ofd.close()  
        ifd.close()  
        return 0  
    }  
}
```

**Figure 2** – Example interface to open, read and write files (top); and a component that uses this interface to copy a file (bottom).

# Dana: Runtime Adaptation

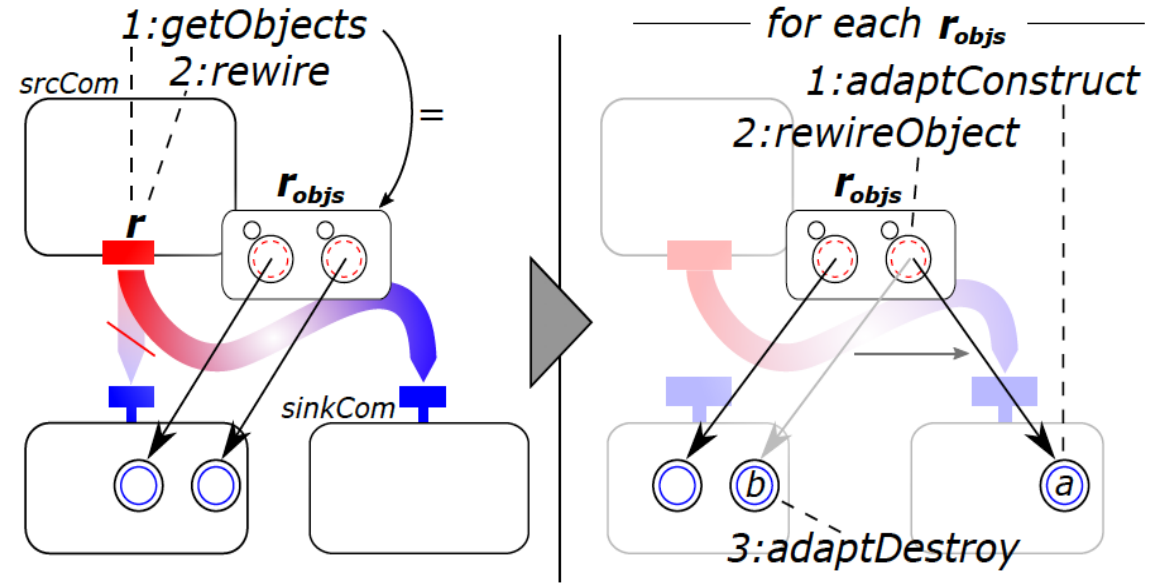
---

**Algorithm 1** Adaptation protocol

---

```
1: srcCom    ▷ Comp. to rewire a required interface of
2: sinkCom   ▷ Comp. with provided interface to wire to
3: intfName  ▷ Interface name being adapted
4: pause(srcCom.intfName)
5: r_objs = getObjects(srcCom.intfName)
6: rewire(srcCom.intfName, sinkCom)
7: resume(srcCom.intfName)
8: for i = 0 to r_objs.arrayLength - 1 do
9:   if pauseObject(r_objs[i]) then
10:    a = adaptConstruct(sinkCom.intfName, r_objs[i])
11:    b = rewireObject(r_objs[i], a)
12:    resumeObject(r_objs[i])
13:    waitForObject(b)
14:    adaptDestroy(b)
15:   end if
16: end for
```

---



**Figure 4** – Adaptation sequence overview. A selected required interface *r* is rewired, followed by each object in the set *r\_objs*.

# PAL Framework: Perception & Assembly

## ➤ Perception

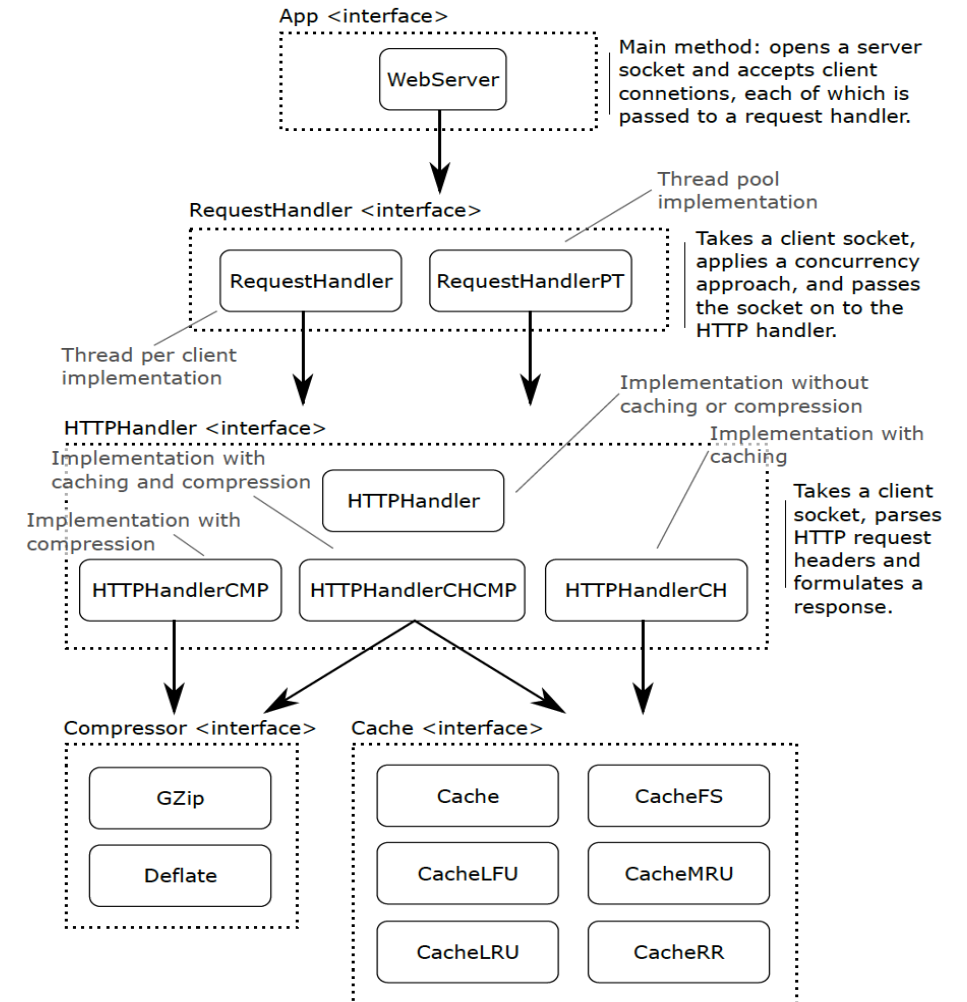
- Implemented using Recorder Interface
- Data – Event and Metrics (Name, Value, Flag)

## ➤ Assembly

- Starts with main component of target system
- Read Required components (recursively)
- Search interfaces in resources directory and their potential implementation
- Example – Interface (io.File) → io (Implementation Directory)
- Create a list of configurations
- Use Adaptation protocol to reassemble


# Sample Implementation: Webserver

- Number of components in system = 30
- File System, String Parsers,
- Number of configurations –  $2*3*(2+5) = 42$
- Request Handler – Avg response time
- HTTPHandler – Events for requested resource & their size



# Exploration Vs Exploitation:

REGRET (LOWER IS BETTER)



- Upper Confidence Bound Action Selection

$$A_t \doteq \arg \max_a \left[ Q_t(a) + c \sqrt{\frac{\ln t}{N_t(a)}} \right]$$

- Greedy – Exploit current knowledge to maximize immediate reward
- Posterior Sampling- Thompson Sampling
  - Estimate posterior distribution using prior distribution



# Multi-armed Bandit

- Arm – One configuration of webserver
- Action – choose one config and deploy

$$\begin{aligned} &\beta_0 + \beta_1 \mathbb{I}_{\text{RequestPT}} \\ &+ \beta_2 \mathbb{I}_{\text{HttpCMP}} + \beta_3 \mathbb{I}_{\text{HttpCH}} + \beta_4 \mathbb{I}_{\text{HttpCHCMP}} \\ &+ \beta_5 \mathbb{I}_{\text{Deflate}} \\ &+ \beta_6 \mathbb{I}_{\text{CacheFS}} + \beta_7 \mathbb{I}_{\text{CacheLFU}} + \beta_8 \mathbb{I}_{\text{CacheMRU}} \\ &+ \beta_9 \mathbb{I}_{\text{CacheLRU}} + \beta_{10} \mathbb{I}_{\text{CacheRR}} \end{aligned} \quad (1)$$

$$x_{\text{conf}} = (1, \mathbb{I}_{\text{RequestPT}}, \mathbb{I}_{\text{HttpCMP}}, \mathbb{I}_{\text{HttpCH}}, \mathbb{I}_{\text{HttpCHCMP}}, \mathbb{I}_{\text{Deflate}}, \mathbb{I}_{\text{CacheFS}}, \mathbb{I}_{\text{CacheLFU}}, \mathbb{I}_{\text{CacheMRU}}, \mathbb{I}_{\text{CacheLRU}}, \mathbb{I}_{\text{CacheRR}}),^3$$

$$y = x_{\text{conf}} \beta + \varepsilon,$$

---

**Algorithm 2** Learning Algorithm

---

```
1: //matrix of all available  $x_{\text{conf}}$  vectors (configurations)
2:  $actionMatrix = assembly.getConfigs()$ 
3:  $X = new Matrix()$  //list of observed  $x_{\text{conf}}$ 's to date
4:  $y = new Vector()$  //list of rewards seen for each X
5:  $n = 0$ 
6: while running do
7:   //do linear regression & sample from posterior
8:    $\Lambda = X^T X + \Lambda_0$ 
9:    $\beta = \Lambda^{-1}(\Lambda_0 \tilde{\beta} + X^T y)$ 
10:   $a = a_0 + (n/2)$ 
11:   $b = b_0 + (y^T y + \tilde{\beta}^T \Lambda_0 \tilde{\beta} - \beta^T \Lambda \beta) \times 0.5$ 
12:   $\sigma^2 = new InverseGamma(a, b).sample()$ 
13:   $sample = new Normal(\beta, \sigma^2 \Lambda^{-1}).sample()$ 
14:
15:  //select the new configuration to use
16:   $i = \arg \max(actionMatrix * sample)$ 
17:   $assembly.setConfig(i)$ 
18:
19:  //wait for 10 seconds, then record observations
20:   $result = 1/perception.getAverageMetric()$ 
21:  add row  $i$  of  $actionMatrix$  as new row of  $X$ 
22:  add  $result$  as new element of  $y$ 
23:   $n++$ 
24: end while
```

---

# Handling Environment Changes

- Entropy –
  - ✓ High => Request for different resources
  - ✓ zero => Single resource requested repeatedly
- Text Volume – Highly Compressible  
Example – HTML, CSS Files
- High entropy interval – more than 50% request of high entropy
- 7 Extra Regression Coefficients
- Total number of configurations =  $42 * 4 = 168$

$$(1, \mathbb{I}_{\text{RequestPT}}, \mathbb{I}_{\text{HiEnt}}, \mathbb{I}_{\text{HiTxt}}, \mathbb{I}_{\text{HttpCMP}(\text{LowTxt})}, \\ \mathbb{I}_{\text{HttpCMP}(\text{HiTxt})}, \mathbb{I}_{\text{HttpCH}(\text{LowEnt})}, \mathbb{I}_{\text{HttpCH}(\text{HiEnt})}, \\ \mathbb{I}_{\text{HttpCHCMP}(\text{LowTxt}, \text{LowEnt})}, \mathbb{I}_{\text{HttpCHCMP}(\text{HiTxt}, \text{LowEnt})}, \\ \mathbb{I}_{\text{HttpCHCMP}(\text{LowTxt}, \text{HiEnt})}, \mathbb{I}_{\text{HttpCHCMP}(\text{HiTxt}, \text{HiEnt})}, \\ \mathbb{I}_{\text{Deflate}}, \mathbb{I}_{\text{CacheFS}}, \mathbb{I}_{\text{CacheLFU}}, \mathbb{I}_{\text{CacheMRU}}, \\ \mathbb{I}_{\text{CacheLRU}}, \mathbb{I}_{\text{CacheRR}}).$$

# Results – Runtime Adaptation

	Average	Maximum	Minimum
setConfig (idle)	509.60 ms	615.00 ms	397.00 ms
setConfig (busy)	1350.32 ms	5811.00 ms	510.00 ms
pause/resume (idle)	8.50 µs	9.94 µs	7.81 µs
pause/resume (busy)	13.22 µs	31.21 µs	8.51 µs
pauseObject/resumeObject (idle)	4.51 µs	5.34 µs	3.84 µs
pauseObject/resumeObject (busy)	28.54 µs	387.17 µs	4.35 µs
components adapted in setConfig()	1.22	3.00	1.00

**Table 1** – Adaptation speed measured in different ways, from full configuration changes to individual component adaptations.

- Webserver is actually paused
- pauseObject – busy waiting for new function call
- pause – prevent new objects from being instantiated

# Results: Divergent Systems

Entropy\Text	Low	High
Low	Cache	Cache & Compress
High	Default (Due to hash collision)	Cache & compression


Request pattern	File size (b) [GZ]	Default	Caching	Caching & compression
Text <i>low entropy</i>	156,983 [12,757]	11.94 ms	9.56 ms	<b>0.70 ms</b>
Text <i>low entropy</i>	82,628 [11,949]	4.05 ms	<b>0.60 ms</b>	0.66 ms
Text <i>low entropy</i>	3,869 [1,930]	1.18 ms	<b>0.59 ms</b>	0.63 ms
Image <i>low entropy</i>	1,671,167 [1,667,464]	160.81 ms	<b>150.72 ms</b>	154.42 ms
Image <i>low entropy</i>	84,760 [66,914]	4.02 ms	<b>0.66 ms</b>	0.74 ms
Image <i>low entropy</i>	4,001 [3,895]	1.22 ms	<b>0.55 ms</b>	0.62 ms
Text <i>high entropy</i>	156,983 [12,757]	19.27 ms	19.66 ms	<b>3.04 ms</b>
Text <i>high entropy</i>	82,628 [11,949]	4.61 ms	3.27 ms	<b>3.07 ms</b>
Text <i>high entropy</i>	3,869 [1,930]	<b>1.25 ms</b>	2.93 ms	2.52 ms
Image <i>high entropy</i>	1,671,167 [1,667,464]	<b>156.50 ms</b>	156.64 ms	157.66 ms
Image <i>high entropy</i>	84,760 [66,914]	4.48 ms	3.19 ms	<b>2.94 ms</b>
Image <i>high entropy</i>	4,001 [3,895]	<b>1.30 ms</b>	2.90 ms	2.67 ms

**Table 2** – Results of different configurations under different request patterns, showing average response times. The standard deviation throughout these results is low, at around 0.2.


# Results:

- 1 test iteration = 10 second (1000 experiments)
- Large File => Less training samples


Small HTML Files with low latency



Large HTML Files with low latency



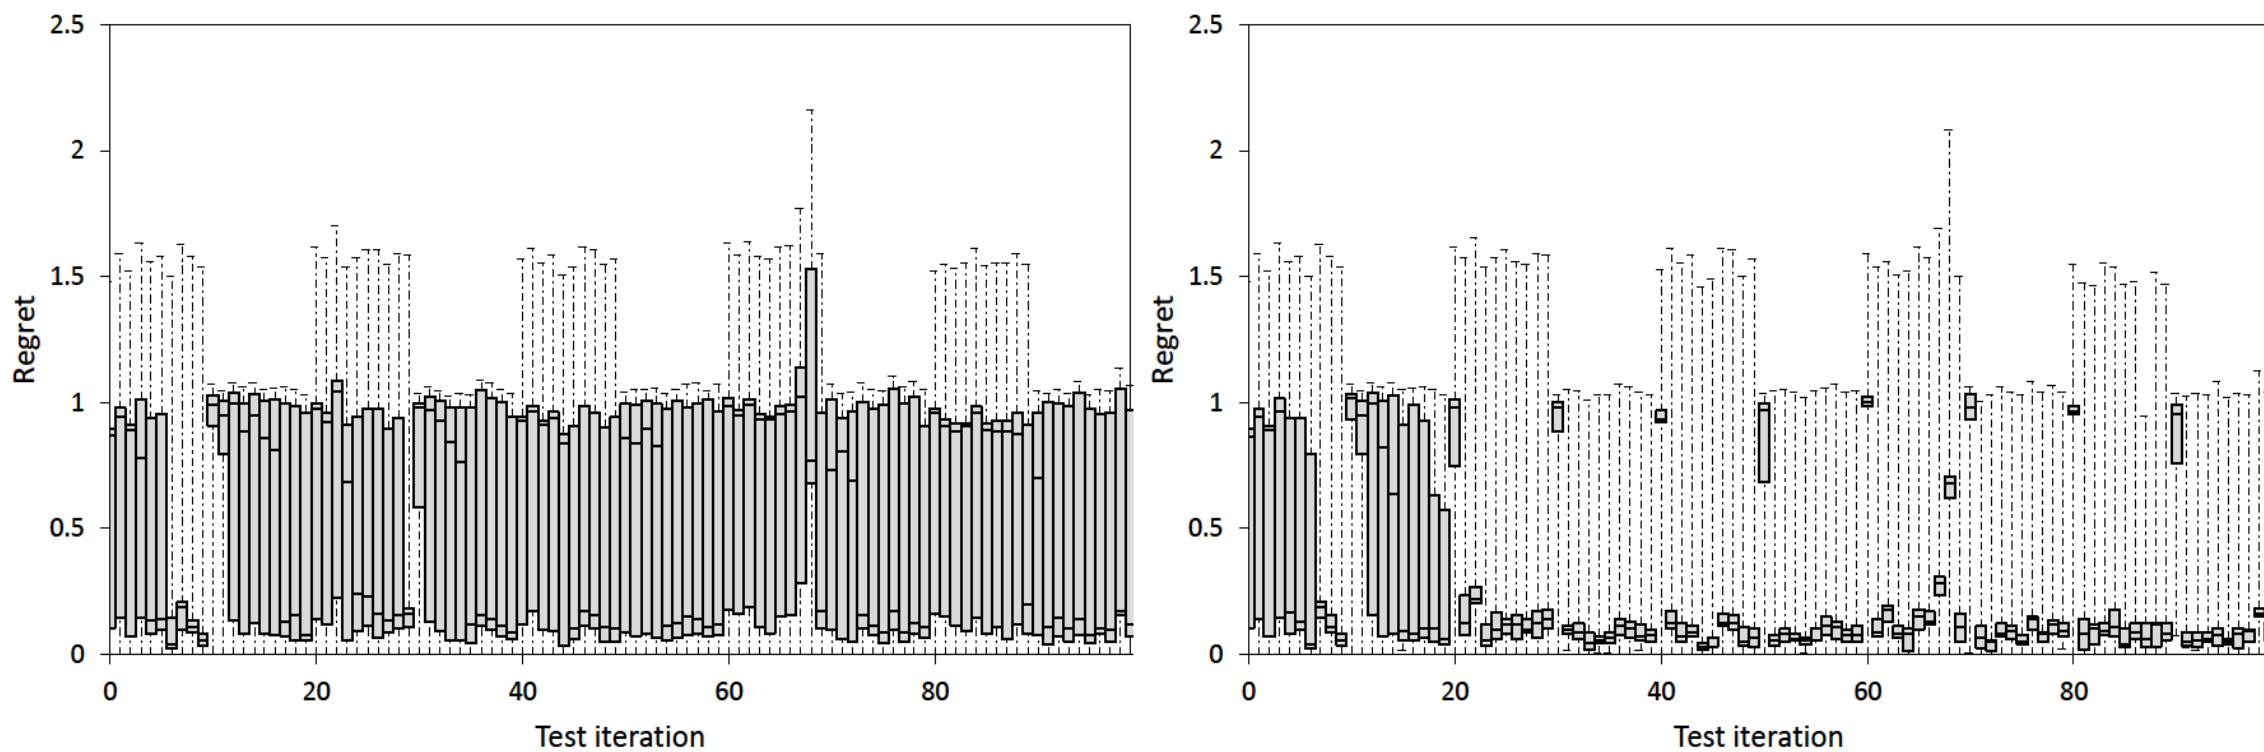
**Figure 7** – Learning using response times to large text files, with adjusted prior values for  $\hat{\beta}_0$  and  $b_0$ .



**Figure 9** – Learning using response times to a realistic (and highly varying) request pattern, using the NASA server trace [2].

# Results: Alternating Request Pattern

➤ Left - Constantly forget and re-learn



**Figure 8** – Learning without (left) and with (right) categorization on a request pattern that changes every ten iterations.

# Thoughts

- Will the adaptation be computationally expensive as number of components and metrics increase ? – Scalable?
- Impact on QoS during transition
- Ease of adding access patterns in model
- Overhead of providing various implementation for a component Vs Simple Knob Tuning
- Extra overhead of module loading for a large system

Questions ?