

5-8-2007

Distributed Diagnosis of Failures in a Three Tier E-Commerce System

Fahad Arshad

Purdue University, faarshad@purdue.edu

Gunjan Khanna

Purdue University, jkhanna@purdue.edu

Ignacio Laguna

Purdue University, ilaguna@purdue.edu

Saurabh Bagchi

Purdue University, sbagchi@purdue.edu

Follow this and additional works at: <http://docs.lib.purdue.edu/ecetr>

Arshad, Fahad; Khanna, Gunjan; Laguna, Ignacio; and Bagchi, Saurabh, "Distributed Diagnosis of Failures in a Three Tier E-Commerce System" (2007). *ECE Technical Reports*. Paper 354.

<http://docs.lib.purdue.edu/ecetr/354>

This document has been made available through Purdue e-Pubs, a service of the Purdue University Libraries. Please contact epubs@purdue.edu for additional information.

Distributed Diagnosis of Failures in a Three Tier E-Commerce System

Gunjan Khanna, Ignacio Laguna, Fahad Arshad, Saurabh Bagchi

Dependable Computing Systems Lab

School of Electrical and Computer Engineering, Purdue University

Email: {gkhanna, ilaguna, faarshad, sbagchi}@purdue.edu

Abstract. For dependability outages in distributed internet infrastructures, it is often not enough to detect a failure, but it is also required to diagnose it, i.e., to identify its source. Complex applications deployed in multi-tier environments, such as the classic three tier e-commerce system, make diagnosis challenging because of fast error propagation, black-box applications, constraints on the diagnosis delay, the amount of states that can be maintained, and imperfect diagnostic tests. Here, we propose a probabilistic diagnosis model for arbitrary failures in components of a distributed application. The monitoring system (the Monitor) passively observes the message exchanges between the components and at runtime, performs a probabilistic diagnosis of the component that was the root cause of a detected failure. The diagnosis model takes into account the possibility of a service failure, link failure, test imperfection, and lack of perfect observability at the monitoring station. We demonstrate the approach by applying it to a J2EE-based e-commerce application called Pet Store exercising a workload of browse-and-buy user transactions. We compare our approach with Pinpoint by quantifying the latency and accuracy of the two systems. The Monitor system outperforms Pinpoint by achieving comparably accurate diagnosis with higher precision in shorter time.

Keywords. *Distributed system diagnosis, runtime monitoring, probabilistic diagnosis, fault injection based evaluation, J2EE e-commerce system.*

I. Introduction

The connected society of today has come to rely heavily on distributed computer infrastructure, be it an ATM network, or the distributed multi-tier applications behind e-commerce sites. The

consequences of downtime of distributed systems may be catastrophic. They range from customer dissatisfaction to financial losses to loss of human lives [40], [41]. There is an increased reliance on Internet services supported by multi-tier applications where the typical three tiers are the web, middleware and database tier. In distributed systems, especially multi-tier applications, the fault in a component may manifest itself as an error, and then propagate to multiple services through the normal communication between the services. The error may remain undetected for an arbitrary length of time causing long error propagation chains. The error can propagate from one component to another and finally manifest itself as a failure. The failure might be detected at a component distant from the originally faulty component. The availability of a system can be quantified as $MTTF/(MTTF+MTTR)$ (MTTF: Mean time to failure, MTTR: Mean time to recovery). There is enormous effort in the fault tolerance community to increase the reliability of components in a distributed system, thus increasing MTTF. There is also a growing number of efforts aimed at reducing MTTR [34]. An important requirement for this is to know which components to recover. This requires tracing back through the chain of errors to determine the component that originated the failure. This serves as the goal for our diagnosis protocol.

For the application and diagnosis model, consider that the application is comprised of multiple services communicating through standard protocols and externally observable messages. Example of such application services are web services and authentication services. The services themselves are comprised of multiple components, e.g. Enterprise Java Beans (EJBs) or servlets, and the interactions between these components are also externally visible. Separate from the application system, we have a monitoring system (the Monitor) that can observe the external interactions of the components but not their internal states. The Monitor initiates diagnosis when a failure is detected through an existing detection system. In this paper we use the existing detection system from our previous work [1].

In practical deployments, the Monitor may not observe the interaction between components of the application perfectly because of congestion or their relative network placement. This is particularly

likely because the application as well as the Monitor is distributed with components spread out among possibly distant hosts. Next, any Monitor will have limited resources and may drop some message interactions from consideration due to exhaustion of its resources (e.g., buffers) during periods of peak load. Third, any diagnostic tests used by the Monitor might not be perfect. Finally, several parameters of the environment are not known deterministically and have to be estimated at runtime. These include the ability of a component to stop the cascade of error propagation (error masking ability) and the unreliability of links between the application components and the Monitor. All these factors necessitate the design of a probabilistic diagnosis protocol, in which the root cause of the failure cannot be deterministically identified.

Our solution implemented in the Monitor rests on three basic techniques. First, the messages between the components are used to build a causal dependency structure between the components. Second, when a failure is detected, the causal structure is traversed (till a well-defined bound) and each component is tested using diagnostic tests. These diagnostic tests are not executed on the components directly but on the component state that had been deduced and stored at the Monitor. We decide against direct tests on the components because the state of the component may have changed since the time it propagated the error and the probing introduces additional stress on the component at a time when failures are already occurring in the system. Third, runtime observations are used to continually estimate the parameters that bear on the possibility of error propagation, such as unreliability of links and error masking capabilities. In our approach, the end goal of the probabilistic diagnosis process is to produce a vector of values called the *Path Probability of Error Propagation* (PPEP). For the diagnosis executed due to a failure detected at component n , PPEP of a component i is the conditional probability that component i is the faulty component that originated the cascaded chain of errors culminating in n .

The basic structuring of an observer and an observed system is not new [1], [2], [3]. The problem of diagnosis of failures in networked environments comprised of black-box entities has also been studied by numerous researchers [28], [29], [35]. Some of these efforts however are aimed at easing

the task of distributed debugging rather than accurate diagnosis of the faulty entity, some are offline approaches, some require accurate prior dependency information between the entities, and yet others need help from the application system through event generation. Our work aims to provide diagnosis of the faulty entities at runtime in a non-intrusive manner to the application.

We apply the diagnosis protocol to a three tier e-commerce system consisting of the Pet Store application deployed on the JBoss application server with the Tomcat web server as the front end and the MySQL database server at the backend. The application supports multiple kinds of browse-and-buy transactions that involve interactions between many components, where components are defined as servlets and EJBs. Through a modification to the JBoss containers, messages between the components are trapped and forwarded to the Monitor. We compare our approach to Pinpoint [29] in terms of accuracy and precision of diagnosis. Pinpoint uses statistical clustering of components with failed transactions to identify the faulty components. We inject errors in the application, where the errors may be due to a single component or interactions between multiple components. Our approach outperforms Pinpoint with the accuracy of the diagnosis improving from 20% to 100% over the Pinpoint algorithm for comparable precision values.

The rest of the paper is organized as follows. Section II, presents the system model. Section III presents the probabilistic diagnosis approach. Section IV presents the implementation and experimental test bed. Section V presents the experimental results. Section VI reviews related work and Section VII concludes the paper.

II. System Model and Background

There are two distinct systems—the Monitor and the application system. The Monitor obtains the protocol messages either through modification to the application’s middleware layer to forward the messages or by a passive snooping mechanism by the Monitor. In either scenario the internal state of the components is not visible to the Monitor and they are treated as black-box for the diagnostic process. The diagnostic process is triggered when a failure is detected.

A. Assumptions

We assume that components can fail arbitrarily, as long as the failures are observable in the message interaction between components. These failures could be caused by incorrect deployment, software bug, security vulnerability or performance problems to name a few. We follow the classical definition of faults being underlying defects that are triggered to become errors and some errors causing end-user visible failures. Errors can propagate from one component to another through the message exchanges between them.

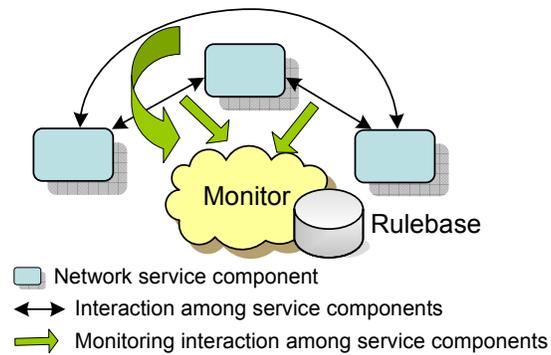


Figure 1. A monitoring system, (the Monitor) verifying the interactions between the service components.

The communication between the components can be asynchronous but the jitter on any given link between component and the Monitor is bounded. We assume that the Monitor maintains a logical clock for each observed component and it is incremented for each event – a send or receive message. The assumption required by the diagnosis protocol is that for an *S(ender)-R(eceiver)* communication, the variation in the latency on the *S-M(onitor)* channel as well as the variation in the sum of the latency in the *S-R* and *R-M* channels is less than a constant Δt , called the *phase*. If messages M_1 and M_2 , corresponding to two send events at *S*, are received at the Monitor at (logical) times t_1 and t_2 , it is guaranteed that the send event M_1 happened before M_2 if $t_2 \geq t_1 + \Delta t$. The communication channel is considered to be unreliable where message duplication, loss or conversion to another correct protocol message may happen.

B. Dependency Information

The Monitor performing diagnosis maintains a causal graph during the times that it is verifying the operation of the application protocol. Let us denote the causal graph at Monitor m by CG_m which is a graph (V, E) where (i) The set V contains all the components verified by m ; (ii) An edge or link e contained in E , between vertices $v1$ and $v2$ (which represent components) indicates interaction between $v1$ and $v2$ and contains state information about all observed message exchanges between $v1$ and $v2$ including the logical clock (LC) at each end. The state information includes a type of interaction and any arguments associated with that interaction. The links in the CG are also time-stamped with the local (physical) time at the Monitor where the link is created. An example of a CG created at the Monitor is given in Figure 2 for the sequence of message exchange events shown with components A, B, C, and D. The number denotes the sequence of the messages. For example, for message '6', the logical clock time at the sender is B.LC4. Since message '2' is assigned a logical time value of B.LC2, it causally precedes message '6'. The LC time stamps helps obtain a partial order over the messages and hence causality. The order of the messages is the order seen by the Monitor which may be different from the order in the application because the communication links are asynchronous.

For a link to be completed in the CG, a matching is required between the sending and the receiving component's messages. The link $A \rightarrow B$ will be matched once both the message sent by A and the corresponding message received by B are seen at the Monitor. The Monitor initially stores the messages in a *Temporary Links* table and moves the matched links to the CG when some trigger is met. As many links as can be matched are transferred to the CG while those that are not matched, but are within the phase from the latest message, are kept in the temporary links. Remaining links in the temporary links table are moved to the CG as unmatched links.

It is imperative to avoid the CG growing in an unbounded manner since this would lead to long delays in traversing the CG during diagnosis leading to high latency in diagnosis. However,

complete purging of the information in the CG can cause inaccuracies during the diagnosis process. We use aggregation of the state information in the CG at specified time points and storing it in an *Aggregate Graph* (AG). The Aggregate Graph contains aggregate information about the protocol behavior averaged over the past. The AG is similar to CG in the structure i.e. a node represents a component and a link represents a communication channel. Unlike the CG there is a single directed link between A and B for all the messages which are sent from A to B. The AG contains some node level information (such as, the node reliability) and some link level information (such as, the reliability of the link in the application system).

C. **Diagnosis Tree**

When a failure occurs, a *Diagnosis Tree* (DT) is constructed using the state information stored in CG. The DT formed for failure F at node D is denoted as DT_{FD} . The tree is rooted at node D and the nodes which have directly sent messages to node D are present at depth 1. Recursively, depth i consists of all the components which have sent messages to nodes at depth $(i-1)$. Since the CG is finite size, the tree is terminated when no causally preceding message is available in the CG after some depth k . The same component might appear several times in the tree at various depths because it might have exchanged messages with various components at different points during the application run. Specifically, a component is represented as many number of times as the number of different states it has been in, while exchanging messages.

D. **Diagnostic Tests**

We assume the existence of *diagnostic tests* which operate on a component and are specific to that component and its state. We impose that the tests should only operate on the information already stored at the Monitor. These tests could be probabilistic in nature, implying that they may not be perfect. The specifics of these tests do not affect our probabilistic model. However, for our implementation, we employ a kind of tests called *causal tests*. A causal test has the following format:

$\langle Type \rangle \langle State1 \rangle \langle Event1 \rangle \langle Count1 \rangle \langle State2 \rangle \langle Event2 \rangle \langle Count2 \rangle$

where, *Type* could be one of {incoming, outgoing, hybrid} depending on the kind of messages being tested. The (*State1*, *Event1*, *Count1*) forms the precondition to be matched, while (*State2*, *Event2*, *Count2*) forms the post-condition that should be satisfied for the node to be deemed correct. The examination of Event 2 is done in an interval of time Δt (a phase) from Event 1. The tuple (S, E, C) refers to the fact that the event E should have been detected in the state S at least count C number of times.

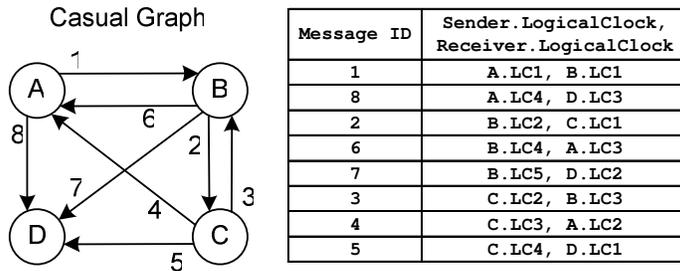


Figure 2.

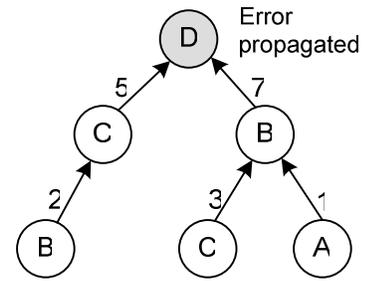


Figure 3.

A sample causal graph. A, B, C and D exchange messages 1-8 among each other. The message ID indicates the causal order, i.e., message 1 precedes the rest of messages **Sample DT for the CG in Fig. 2.**

The correctness rules can be created by examining the *state transition diagram* (STD) of the component and verifying the transitions or by observing some traces of the correct protocol operation. Additionally, rules corresponding to QoS requirements (such as, the number of accesses to the *SignOnEJB* in Pet Store must be restricted to 20 within a 1 sec time window) can be framed by the system administrators. Finally, rules for verifying security properties in the system (such as, the number of logins to Pet Store bounded by a threshold) can be set by security administrators. Rules therefore can be framed through a mix of automated and manual means. This is similar to the situation in all rule based systems, such as intrusion detection systems [4], [5].

III. Probabilistic Diagnosis

The operation of the diagnosis protocol has two logical phases: (1) The diagnostic process that results in a set of nodes being diagnosed as the root cause of failure; (2) Information from the

diagnostic process being used to update the parameters used later diagnoses. The overall process is depicted in Figure 4. Let us first look at the diagnostic process. The goal of the diagnostic process is to calculate the probability of each node in the distributed system being faulty.

A. Path Probability of Error Propagation

The DT forms the basic structure on which the algorithm operates. The path P from any node n_i to the root of the DT constitutes a possible path for error propagation, i.e. a fault present in n_i could have caused the root node to fail during operation. The probability of a path being the chain of error propagation is termed as the *Path Probability of Error Propagation* (PPEP).

A sample DT created from the sample CG in Figure 2 is shown in Figure 3. Here the failure was manifested at node D . The numbers at the links correspond to the message IDs. The root of the tree is the failure node, i.e., D . Depth 1 consists of nodes C and B which have sent messages to D causally before the failure was detected. Here node B is repeated twice because the states of B in which $B \rightarrow C$ and $B \rightarrow D$ communication take place are different.

Definition: $PPEP(N, D)$ is defined as the probability of node N being faulty and causing this error to propagate on the path from N to D , leading to a failure at D . PPEP depends on the following parameters:

(1) **Node reliability** – The node reliability (n_r) is a quantitative measure of the component corresponding to the node being correct. The PPEP for a given node is proportional to $(1 - n_r)$. This node reliability is obtained by running the diagnostic tests on the state of the entity. A set of predetermined tests are performed, each of which yields a ‘0’ (test flags an error) or a ‘1’ (success). If the entire set of tests is denoted by R and a subset of tests which yield ‘1’ be denoted by R' , we define coverage $c(n) = |R'|/|R|$, assuming all tests have equal weights. For the first time the diagnosis is triggered, the node reliability is equal to $c(n)$. During the entire execution of the application, multiple failures cause multiple diagnosis procedures to execute. Each time the

diagnosis is performed, node reliabilities (in AG) corresponding to all of the nodes in the DT are updated. We will explain in Section B how to update the node reliabilities.

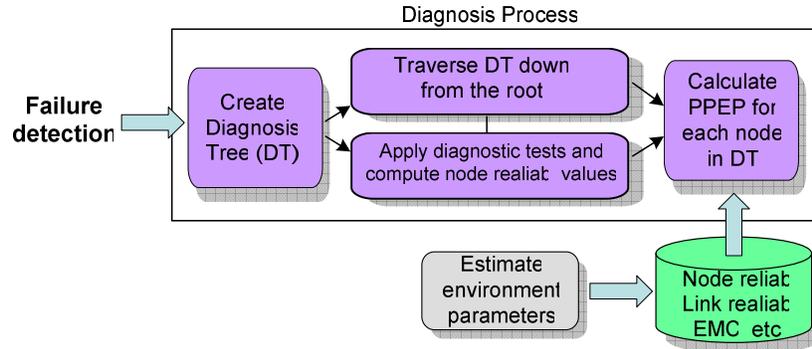


Figure 4. Schematic with the overall process flow of the diagnostic process

(2) **Link Reliability** – In simple terms the link reliability between two nodes n_i and n_j denoted as $l_{r(i,j)}$ and measures the number of received packets by the receiver over the number of packets actually sent. The PPEP for a given node is proportional to the link reliability, because high link reliability increases the probability of the path being used for propagating the error. The link reliability is maintained for each edge in the AG. Note that since the Monitor is only observing the system, the errors within the Monitor in observing the messages also may erroneously affect the link reliability. These errors cannot be distinguished in our framework. Monitor deduces the link reliability through observing the fraction of successful message transmissions over a particular link.

(3) **Error masking capability (EMC)** – The error masking capability (e_m) quantifies the ability of a node to mask an error and not propagate it to the subsequent link on the path in the DT toward the root. The EMC of an entity depends on the type of error, e.g., syntactical or semantic errors. Additionally, a node may have different error masking capabilities depending on the type of message being processed and forwarded, e.g., if there is an *off-by-one* bug in the counter check on the number of simultaneous JDBC connections, it will mask the errors when the number of JDBC connections is one more than the maximum threshold. The EMC of node C in Figure 3 is denoted by $e_m(C)$ and is a function of *message type* and *error type*. The PPEP for a given node is inversely

proportional to the EMC values of nodes in the path since the intermediate nodes are less likely to have propagated the error to the root node.

With the DT in Figure 3, $PPEP(C, D) = (1-n_r(C)) \cdot l_{r(C,D)}$, $PPEP(B, D) = (1-n_r(B)) \cdot l_{r(B,C)} \cdot (1-e_m(C)) \cdot l_{r(C,D)}$. For a general path P consisting of nodes $n_1, n_2 \dots n_k$ with link $l_{r(i,j)}$ between nodes n_i and n_j , the $PPEP(n_1, n_k)$ for a failure detected at node n_k (root node in the corresponding DT) is given by

$$PPEP(n_1, n_k) = (1-n_r(n_1)) \cdot l_{r(1,2)} \cdot (1-e_m(n_2)) \cdot l_{r(2,3)} \dots (1-n_r(n_i)) \cdot l_{r(i,i+1)} \cdot (1-e_m(n_{i+1})) \cdot l_{r(i+1,i+2)} \dots (1-e_m(n_{k-1})) \cdot l_{r(k-1,k)}$$

Consider a far simpler alternative that may at first glance seem sufficient for the problem of diagnosis. The system simply looks at all messages in the system at runtime and checks them using the causal tests. The Monitor could detect the error very quickly after the earliest message that failed a test. However, this scheme would be incorrect. First, the system cannot check the messages as they come in. The rules are specified in terms of logical time and not absolute time (a corollary of the asynchronous nature of the system). Additionally, all rules cannot be matched at runtime because that would impose unnecessary overhead and would not be useful in most executions (when there is no failure). The first component that failed a test does not necessarily implicate the component that sent it, other factors are to be considered, such as, how reliable were the links between that component and the component at which the failure was ultimately detected. This is because the distance from the root is not a matter simply of the number of links on the DT. Also, the tests are not perfect and cannot therefore be trusted to indict a component by itself.

For an autonomous system, the parameters used in the diagnosis process should be automatically updated during the lifetime of the system as more failures and message interactions are observed and this forms the topic of our discussion next.

B. Updating Node reliability

The node reliabilities are updated based on the results of the diagnostic tests. Let the set of tests that can be applied to the node i (based on the event and the state) be T_i . This set is partitioned into two

sub-sets A' and A respectively depending upon the test outcome, i.e., the test returned a value of '1' or '0'. If the weight of test j in T_i is $w_{i,j}$ then the current coverage of node i is given by

$$c(i) = \sum_{j \in A'} w_{i,j} / \sum_{j \in T_i} w_{i,j} .$$

The weight of a test is proportional to (1) the frequency of invocations ($w^{(f)}$) where the test gave the correct result; and (2) whether the test examines state for a period of time greater than the length of transients in the system ($w^{(r)}$). The overall weight is calculated as

$$w_{i,j} = w_{i,j}^{(f)} \cdot w_{i,j}^{(r)} .$$

The first factor is quite self-explanatory—higher the number of times the test has been seen to operate correctly, the higher is the confidence placed on it. The second factor argues that violation of a tested property for less than the transient period in the system is not indicative of a fault. The value $c(i)$ is the current coverage of the node and does not consider the past history of node present in AG. The final node reliability is updated using exponential weighting as $n_r(A) = (1 - \rho)c(i) + \rho n_r(A)$. The value $n_r(A)$ is maintained in the AG. The weight can be adjusted depending upon the particular scenario.

C. Updating Link Reliability

Link reliability of the edge from A to B is calculated as $l_r(A,B) = n_m / n_t + n_m$, where n_m is the number of matched edges for A-to-B communication and n_t is the number of unmatched edges from A to B . This formulation is based on the observation that if the sender notifies the Monitor of a sent message and the receiver does not report the corresponding packet, this is indication of a lossy link. Overall $l_r(A,B)$ is updated with the link reliability for the current round ($l_r(c)$) using exponential weighting as $l_r(A,B) = (1 - \rho)l_r(c) + \rho l_r(A,B)$. Notice that in the PPEP calculation, the link reliabilities of adjoining links are multiplied though the events are actually not independent. Consider a linear chain of communication from A to B and B to C . The probability of a successful communication from A to C is $P(A \rightarrow B \text{ is successful}) \cdot P(B \rightarrow C \text{ is successful} \mid A \rightarrow B \text{ is successful})$. In the link reliability calculation, the dependence is implicitly taken into account since the matched and

unmatched link count on the $B \rightarrow C$ link is affected by the events on the $A \rightarrow B$ link and the measurement of the matched and unmatched links is used in the link reliability computation.

D. Updating Error Masking Capability

Consider in the DT of Figure 3, nodes C and A at the same depth 2, are both faulty but $PPEP(C)$ is the highest and $PPEP(A)$ is low. This will cause node C to be diagnosed as faulty. In this case node B is taken to have masked the error originating from node A and not propagated it if the following three conditions are satisfied: (i) Running diagnostic tests on B yields a high coverage; (ii) Running diagnostic tests on A yields a low coverage; and (iii) Link reliabilities $l_r(A,B)$ and $l_r(B,D)$ are high (to ensure that the error must have propagated).

The incremental change $\Delta(EMC)$ is thus given by

$$\Delta(EMC) = EMC_{prev} \frac{(1 - c(A)) \cdot l_{r(A,B)} \cdot l_{r(B,D)}}{(1 - c(B))}$$

$$EMC_{new} = EMC_{old} + \Delta(EMC)$$

We decrease the EMC for every intermediate node residing on a path which is finally diagnosed to have caused the error propagation.

E. Distributed PPEP

The nodes in the Diagnosis Tree may be present in several networks and even organizational boundaries and be verified by different Monitors each of which constructs a part of the AG and the CG obtained from its local information. During diagnosis it is quite likely that the DT contains nodes which are verified by some other Local Monitors. We leverage the fact that due to the multiplicative form of the PPEP computation, the PPEP value can be computed incrementally by each Monitor for the part of the DT under its verification domain. Each Local Monitor runs the diagnosis algorithm to obtain the PPEP values for the nodes under its verification domain. This information is aggregated at a higher level Monitor (such as an Intermediate Monitor) to determine the faulty entity.

IV. Experimental Testbed

A. Application

We use for our evaluation PetStore (version 1.4), a sample J2EE application developed under the Java BluePrints program at Sun Microsystems [37]. It runs on top of the JBoss application server [38] with MySQL database [39] as the back-end providing an example of a 3-tier environment. Figure 5 depicts the application topology for the experiments. The PetStore application is driven by a web client emulator which generates client transactions based on sample traces. The web client emulator is written in Perl using lynx as the web browser. For the mix of client transactions, we mimic the TPC-WIPSo [36] distribution with equal percentage of browse and buy interactions. The servlets and the EJBs are considered as components in our experiments and these serve as the granularity level at which diagnosis is done. This design choice is based partly on the fact that in JBoss a faulty servlet or an EJB can be switched out at runtime for a correct one. We identify a total of 56 components in the application.

We consider a *web interaction* to be a complete cycle of communication between the client emulator and the application, as it is defined by the TPC Benchmark WIPSo specification [36]. This cycle starts when the client emulator initiates a web request and it is completed when the last byte of data from the response page has been received by the client emulator. Examples of web interactions could be entering the Welcome page or executing a Search. A *transaction* is a sequence of web interactions. An example of a transaction by a user who is searching and viewing information about a particular product is: *Welcome page* → *Search* → *View Item details*. For our experiments we created a total of 55 different transactions. A *round* is a permutation of these 55 transactions modeling different user activities that occur on the web store. Within a round, transactions are executed one at a time. Two transactions are considered to be non-unique if they use exactly the same components, neglecting the order in which the components are used. Thus, a transaction that comprises: *Welcome, Search, Search* is not unique with respect to another that

comprises: *Welcome*, *Search*. There are 41 unique transactions in the set of 55 transactions that we use. Although 55 is not an exhaustive set of possible transactions in the application, the chosen set exercised a wide variety of web-interactions and between them, touched all the components of PetStore. We note that the results presented here depend on the exact set of transactions used to exercise the system.

We instrumented the JBoss application server to snoop over the message communication between PetStore components. JBoss has a layered architecture and each communication traverses multiple interceptors. We modify the *SecurityInterceptor* to forward messages to the Monitor for updating the causal graph. Thus, the PetStore application is left unchanged.

B. Monitor configuration

The diagnosis algorithm in the Monitor is implemented in Java. The Monitor is provided an input of state transition diagrams for the verified components and causal tests used during calculation of PPEP values. The size of the causal graph is bounded at 100 links.

Figure 6 shows an example STD for *CreditCardEJB* used by the Monitor in our experiments. A start state S_0 signifies a no request state. If a request for processing is received from another component, the state of the EJB moves from S_0 accordingly. With the STD, we have some simple causal tests which can be derived from the STD itself. As explained in section II.D, causal tests are dependent on the state and event of the component. For example, if the EJB is requested for *getData()* then in state S_1 there must be a return from *getData()* to ensure correct operation of the EJB. This is verified using the first rule in Figure 6. The exhaustive list of STDs and rules used for the experiments here is provided in [42].

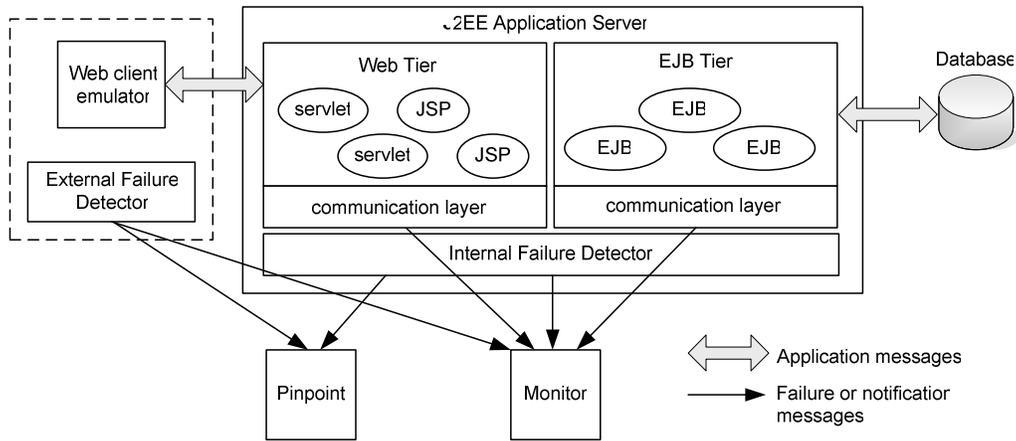


Figure 5. Logical Topology of the Client and Server for the Experiments

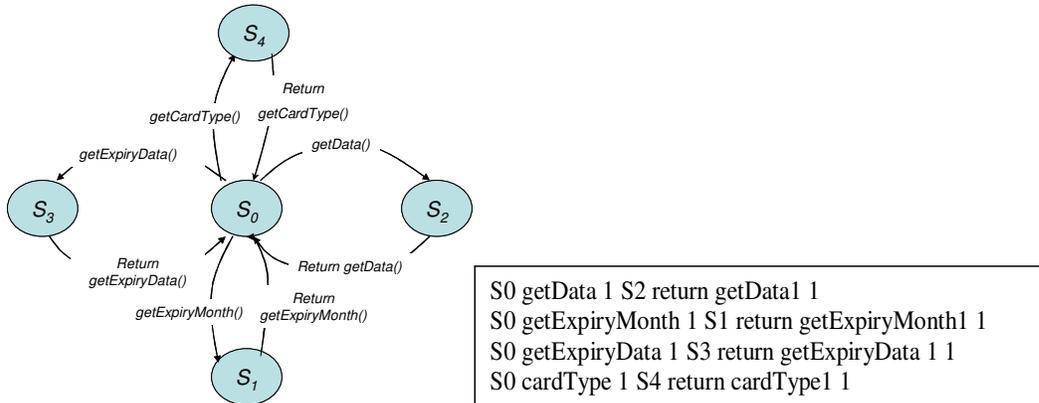


Figure 6. An example STD for CreditCardEJB along with some illustration of Causal Tests.

C. Pinpoint Implementation

Pinpoint serves as a valid point of comparison with the Monitor since both systems have the same focused goal (diagnosis, as opposed to say performance debugging as in [28] with diagnosis being a side issue) and have the same target application model (black-box or gray-box application and passive observation of the application for diagnosis). Importantly, Pinpoint represents a recent state-of-the-art development ([29], [35]) and has been well explained and demonstrated on an open source application (compare to Magpie [30] where the application is not available to us), and its algorithms are not dependent on a huge set of parameters whose settings are left mysterious in the publication (compare to the machine learning approach in [22][23] where several statistical distributions would have to be assumed).

We implement the Pinpoint algorithm (as explained in [29]) for comparison with our Monitor’s diagnosis approach. Pinpoint diagnosis algorithm requires as input a dependency table—a mapping of which components each transaction depends on. This is in contrast to the Monitor approach, where such dependency information does not have to be determined *a priori* and fed into the system before execution. Instead the Monitor deduces the dependencies through runtime observations as described in Section II.B. For Pinpoint, when transactions are executed, their failure status is determined by the failure detectors. A table (called the input matrix) is then created with the rows being the transactions, the first column being the failure status, and the other columns being the different components. If a cell $T(i, 1)$ is 1, it indicates transaction i has failed. If a cell $T(i, j)$ ($j \neq 1$) is 1, it indicates transaction i uses the component j . Pinpoint correlates the failures of transactions to the components that are most likely to be the cause of the failure. The input matrix is fed as input to a data clustering engine. The transpose of this binary input matrix is used by the data analysis engine. The data analyzer computes the dissimilarity between the rows of transposed matrix, which is represented by Jaccard’s distance. This matrix containing the distances between the components is fed to a clustering algorithm called the Unweighted Pair Group Method with Arithmetic Mean (UPGMA). The components that fall in the cluster having the failure row are diagnosed to be faulty.

A crucial point for the accurate operation of Pinpoint is that the transactions should be diverse enough, i.e., use distinct non-overlapping components. Two transactions T_1 and T_2 are called distinct with respect to a set of components $\{C_1, C_2, \dots, C_k\}$ if there is no overlap between these columns for T_1 and T_2 , i.e., when T_1 ’s row has a 1 in any of these columns, T_2 ’s row has a zero, and vice-versa. Pinpoint as described by the authors in [29] is an offline approach. For comparison with the Monitor, we convert it into an online protocol. We incrementally feed the transactions and their corresponding failure status as they occur in the application, rather than waiting for all the transactions in a round to be completed before executing Pinpoint. To provide a comparable platform between the Monitor and Pinpoint, we keep the testbed identical to that in [29]—same client, web server, application server (with identical components), and database server. Since the

performance of the Monitor and Pinpoint are sensitive to the transactions used, we would have liked to use the same set of transactions as used by Pinpoint in [29]. However, the paper is silent on the issue—it does not even provide the total number of transactions used. We contacted the authors of Pinpoint but they were unable to provide us with the transactions either (see personal communication in [42]).

D. Detectors

We create the same detectors as in [29]. An internal and an external failure detector are built which provide failure status of transactions to Pinpoint and the Monitor. The external detector detects failures that will be visible to the user, such as application-specific failures, machine crashes or complete service failures. We implemented this external detector as part of the client emulator. It examines the output error log of lynx and flags a failure if an HTTP error is observed. Alternately, if a transaction does not complete within 20 seconds, timeout occurs and the detector flags a failure. An internal detector is used to detect a failure that may not immediately manifest itself to users. The internal detector is built to catch Java declared and undeclared exceptions in the application and is embedded in each component.

E. Fault Injection

We perform fault injection into the components of the PetStore application (i.e., Servlets and EJBs). We choose a set of 9 components called *target components* consisting of six EJBs and three servlets for fault injection. The names of the components are *AddressEJB*, *AsyncSenderEJB*, *CatalogEJB*, *CreditCardEJB*, *ContactInfoEJB*, *SupplierClientLocalFacadeEJB*, *enter_order_information.screen*, *order.do*, and *item.screen*. We use four different kinds of fault injection as in [29]:

- Declared Exception: We inject *IOException* as the representative declared exception.
- Undeclared Exception: This is a Runtime Exception not caught in the application.
- Endless call: The target component has an infinite *while* loop.
- Null call: Instead of returning the appropriate value, a method returns a null object.

The internal detector is more likely to detect the declared and the undeclared exceptions, and the null calls while the external detector is more likely to detect the endless call. For a given round only one target component is injected. We use 1-component, 2-component and 3-component triggers. In a 1-component trigger, every time the target component is touched by a transaction, the fault is injected in that component. In a 2-component trigger, a sequence of 2-components is determined and whenever the sequence is touched during a transaction, the last component in the transaction is injected. This mimics an interaction fault between two components, and, in the correct operation of a diagnosis protocol, both components should be flagged as faulty. The 3-component fault is defined similarly.

V. Results

A. Performance Metrics

We use *precision* and *accuracy* as output metrics as in the Pinpoint work to enable a comparison. A result is accurate when all components causing a fault are correctly identified. For example, if two components, A and B, are interacting to cause a failure, identifying both would be accurate. Identifying only one or neither would not be accurate. However, if the predicted fault set (by the diagnosis algorithm) is {A, B, C, D, E} and the fault was in components {A, B}, then the accuracy is still 100%. Precision captures the non-idealness in this case. Precision is the ratio of the number of faulty components to the total number of entities in the predicted fault set. In the above example, the precision is 40%. Components {C, D, E} are false positives. Lower precision implies high false positives. There is a tension between accuracy and precision in most diagnosis algorithms. When the algorithm is sensitive, it generates highly accurate results, but also causes a large number of false alerts reducing precision. Pinpoint uses the UPGMA clustering algorithm and varying the size of the faulty cluster varies the precision and accuracy. In the Monitor, after the diagnosis algorithm terminates, an ordered list of components is produced in decreasing order of PPEP. We define the

predicted fault set as the top k components in the ordered output list. We vary k to obtain different accuracy and precision values.

B. Single Component Faults

In single component faults, the fault injection trigger consists of a single component. If a transaction touches the target component then one of the four kinds of faults (chosen randomly) is injected and the injection remains permanent for the remainder of the round. First, let us consider the effect of varying cluster size on the performance of Pinpoint. The total number of injections for these results is 36—9 target components for injection and all 4 types of injection done on each component. The averaged results for accuracy and precision are plotted in Figure 7 (the bars show 90% confidence interval). As the size of the cluster increases, we see an increase in the accuracy which is intuitive because at some point the failure cluster includes all the components that are actually faulty. Beyond that, increase in cluster size does not impact the accuracy. As the cluster size increases, the precision increases to a maximum value and then decreases thereafter. The increase occurs till all the faulty components are included in the failure cluster. Thereafter, increasing the cluster size includes other non-faulty components and thus brings down the precision. The maximum value of precision occurs when all the faulty components are included in the failure cluster. However the precision is still poor (less than 10%). This is explained by the observation that for the transactions in the application, there is tight coupling between multiple components. Whenever the entire set of tightly coupled components does *not* appear together as a fault trigger, which is the overwhelming majority of the injections, the precision suffers. The amount of tight coupling between the components is quantified through the experiment in Section F. We emphasize that if we were to hand pick transactions such that they are distinguishable with respect to the target components, then the performance of Pinpoint would improve. Two transactions T_i and T_j are indistinguishable with respect to a set of components $\{C_1, C_2, \dots, C_k\}$ if the columns of T_i in the input matrix corresponding to these components are identical to that of T_j . Figure 7(a) shows the variation of

Accuracy with False Positives for Pinpoint and the Monitor. This is averaged across the 36 injections for the presented results. For 1-component faults, Pinpoint has high false positives rates but the accuracy eventually reaches 1. In contrast the Monitor has a much higher accuracy keeping a low false positive rate. Monitor’s accuracy also reaches 1 but at a much lower value of false positives (0.6) as compared to Pinpoint (> 0.9). The latency of detection in our system is very low. Thus, the faulty component is often at the root of the Diagnosis Tree in the Monitor. Since error propagation is thus minimized, the PPEP value for the faulty entity is high causing it to be diagnosed by the Monitor. This explains the high accuracy for the Monitor. However, Pinpoint’s algorithm does not take advantage of the temporal information—the temporal proximity between the component where detection occurs and the component that is faulty. As a consequence its accuracy suffers relative to that of the Monitor.

Notice that in Pinpoint, for a given value of false positives, two different accuracy values are achieved since a given precision value is achieved for two different cluster sizes (Figure 7(b)). Since accuracy is a monotonically increasing plot with cluster size (Figure 7(a)), the different cluster sizes give two different accuracy values.

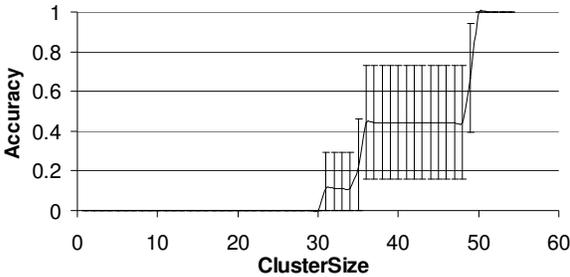


Fig 7(a)

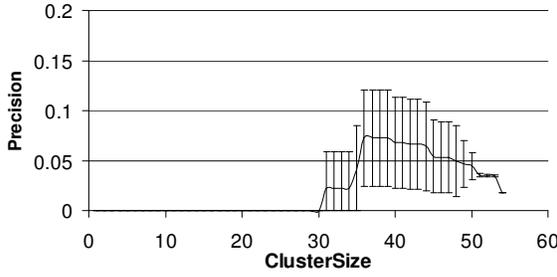


Fig 7(b)

Figure 7. 1-component fault injection: Variation of Accuracy and precision with cluster size in Pinpoint.

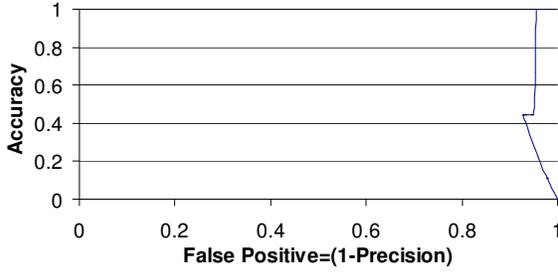


Fig 8(a)

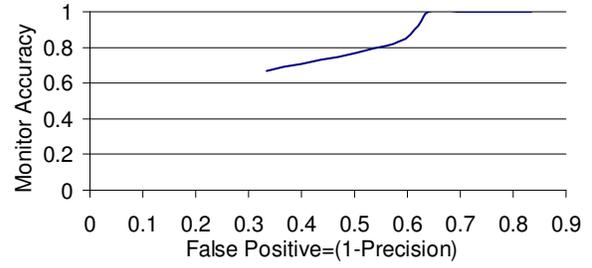


Fig 8(b)

Figure 8. Single component fault injection: Performance of Pinpoint and Monitor. Both can achieve high accuracy but Pinpoint suffers from high false positive rates.

C. Two Component Faults

The 2-component fault injection results are shown in Figure 9. Pinpoint results improve in terms of the false positives implying higher precision. This is attributed to the fact that Pinpoint’s clustering method works better if the failing transactions are better distinguishable from the successful transactions. Recollect distinguishable is discussed in the context of components. A 2-component fault includes two components as the trigger and going from one component to two components increases the distinguish-ability of transactions. Consider transaction T_1 and T_2 both of which use component C_1 (the trigger in a single component fault injection). However, for a two component fault injection with trigger as $\{C_1, C_2\}$, the transactions T_1 and T_2 will be distinguishable as long as both T_1 and T_2 do not use C_2 . Thus, say T_1 uses $\{C_1, C_2\}$ and T_2 does not use C_2 . Then only T_1 will fail and T_2 will not, leading to the diagnosis (considering simplistically that these are the only transactions and components) of C_1 - C_2 as the faulty entities.

In contrast, the Monitor results although, still significantly better than Pinpoint, suffer in the 2-component fault injection. One can see that accuracy reaches a maximum of only 0.83 compared to 1.00 in 1-component injection. The number of times in a round the trigger for the 2-component fault is hit is lower than for the single component fault. Each detection causes an execution of the diagnosis process and each execution of the diagnosis process updates the parameters of the causal graph away from an arbitrary initial setting toward an accurate set of values. Thus, for the 2-component faults, the Monitor gets less opportunity for refining the parameter values and

consequently the PPEP calculation is not as accurate as for the single component faults. This explains the decline in performance of the Monitor for the 2-component faults.

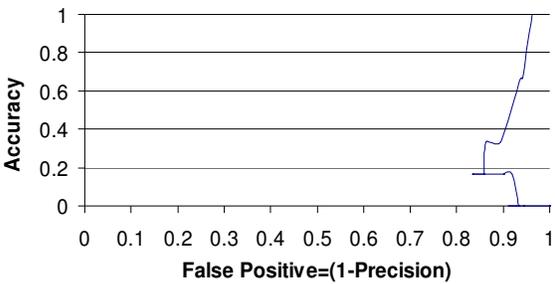


Fig. 9(a)

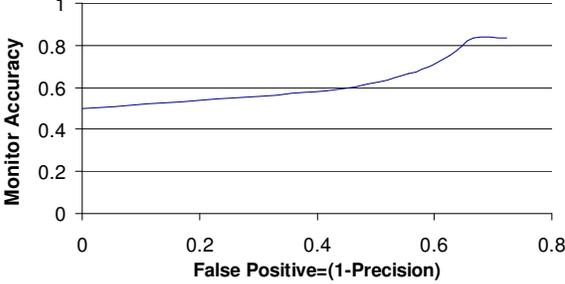


Fig. 9(b)

Figure 9. 2-component fault injection: Performance of Pinpoint and Monitor. Performance of Monitor declines and Pinpoint improves from the single component fault, but Monitor still outperforms Pinpoint.

D. Three Component Faults

The 3-component fault injections show even better results for Pinpoint with the maximum average precision value touching 27%. This is again attributed to the fact that more number of components causes selected transactions to fail leading to a better performance by the clustering algorithm. The Monitor again outperforms Pinpoint by achieving higher accuracy at much lower false positives. The Monitor’s performance again declines compared to the 2-component faults due to the same reason pointed in the previous section (the number of diagnoses for the 3-component trigger is less than that for the 2-component trigger).

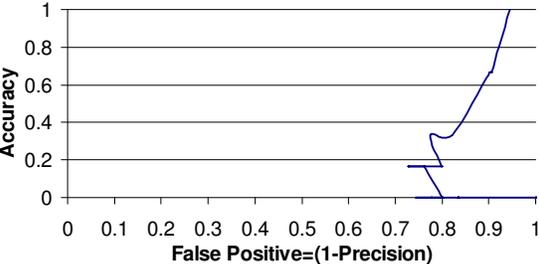


Fig. 10(a)

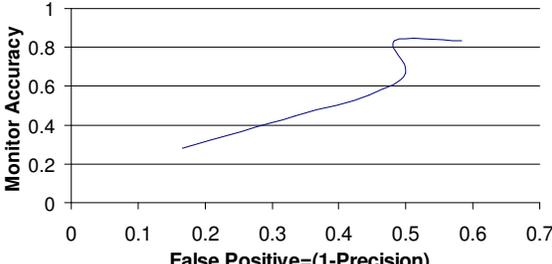


Fig. 10(b)

Figure 10. 3-component fault injection: Performance of Pinpoint and Monitor. Performance of Monitor declines and Pinpoint improves from the single and two component fault, but Monitor still outperforms Pinpoint.

E. Latency

In its online incarnation, Pinpoint takes as input the transactions and corresponding failure status every 30 seconds during a round. It runs the diagnosis for each of these snapshots taken at 30 second intervals, terminating when the round is complete and Pinpoint executes on the entire input matrix corresponding to all the 55 transactions. Pinpoint’s performance only becomes reasonable at 3.5 minutes and above and hence we report only this part of the plot. Arguably this is a subjective decision, but we find the meaningful insights are only possible when Pinpoint has data worth 3.5 minutes or more. The latency plots show that after 3.5 minutes the accuracy and precision increase monotonically with latency.

We define the latency of diagnosis for the Monitor as the time delay from the receipt of the detector alert which marks the beginning of the diagnosis till the PPEP ordered list is generated. The Monitor has an average latency of 58.32 ms with a variance of 14.35 ms, aggregated across all three fault injection campaigns.

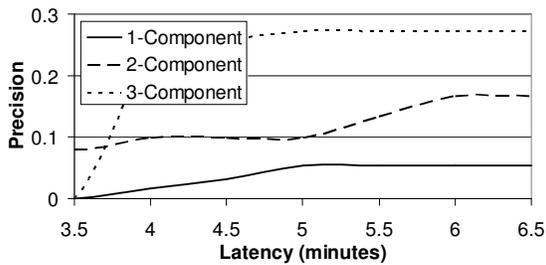


Fig. 11(a)

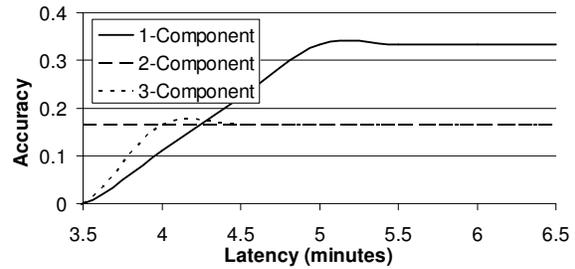


Fig. 11(b)

Figure 11. Single component fault injection: Variation of accuracy and precision with latency for Pinpoint. Higher latency means higher number of transaction data points and Pinpoint’s performance improves monotonically.

F. Behavior of Components

The PetStore application has some components which are tightly coupled, i.e., they tend to be invoked together for the different transactions supported by the application. We have noted earlier that tight coupling negatively impacts Pinpoint’s clustering algorithm. For our experiments, we

inject failures in 9 components and here we consider how tightly coupled these components are with the other components in PetStore. *AddressEJB* is tightly coupled with 4 components implying that it always occurs with these 4 components in all the 55 transactions in our experimental setup. Pinpoint cannot distinguish between sets of components that are tightly coupled and thus reports all the tightly coupled components as faulty even though in reality only a subset of these may be faulty. This is the fundamental reason why its precision is found to be low in all our experiments. To counter this problem, one can synthetically create transactions that independently use different components (as noted by the authors themselves in [29]). However, for an application like PetStore, components are naturally tightly coupled and thus generating such synthetic transactions is a difficult task. Also even if we could devise such “unnatural” transactions that would make components distinguishable, it cannot be assumed that such transactions will be created by regular users in the system.

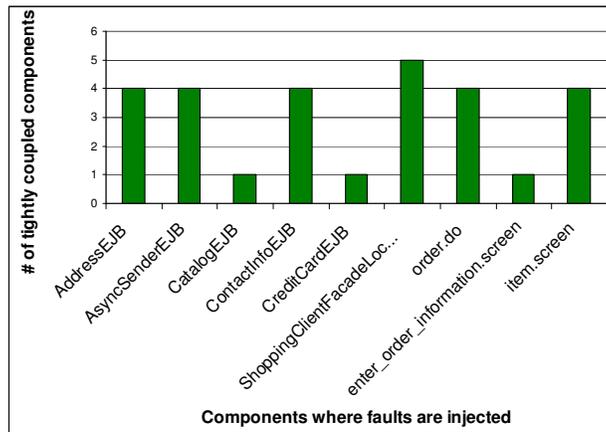


Figure 12. Tightly connected components

VI. Related Work

White box systems: The problem of diagnosis in distributed systems can be classified according to the nature of the application system being monitored – white box where the system is observable and, optionally, controllable; and black box where the system is neither. The Monitor system falls in

the latter class. White box diagnostic systems often use event correlation where every managed device is instrumented to emit an alarm when its state changes [8][18]-[20]. By correlating the received alarms, a centralized manager is able to diagnose the problem. Obviously, this depends on access to the internals of the application components. Also it raises the concern whether a failing component's embedded detector can generate the alert. This model does not fit our problem description since the target system for the Monitor comprises of COTS components, which *have to be* treated as black-box. A number of white box diagnosis systems that correlate alarms have been proposed in the intrusion detection area [43][44]. An alternative diagnostic approach is to use end-to-end probing [21]-[23]. A probe is a test transaction whose outcome depends on some of the system's components; diagnosis is performed by appropriately selecting the probes and analyzing the results. Probe selection is typically an offline, inexact, and computationally heavy process. Probing is an intrusive mechanism because it stresses the system with new requests. Also it is not guaranteed that the state of the system with respect to the failure being diagnosed has stayed constant till the time of the probe.

Multiprocessor system diagnosis: The traditional field of diagnosis has developed around the target of multiprocessor systems, first addressed in a seminal paper by Preparata *et al.* [11] known as the PMC method. The PMC approach, along with several other deterministic models [7], assumed tests to be perfect and mandated that each entity be tested a fixed number of times. Probabilistic diagnosis, on the other hand, can only diagnose faulty nodes with a high probability but can relax assumptions about the nature of the fault (intermittent faulty nodes can be diagnosed) and the structure of the testing graph [14], [15]. Follow up work focused on multiple syndrome testing [13] where multiple syndromes were generated for the same node proceeding in multiple lock steps. Both use the comparison based testing approach whereby a test workload is executed by multiple nodes and a difference indicates suspicion of failure. The authors in [24] propose a fully distributed algorithm that allows every fault-free node to achieve diagnosis in, at most, $(\log N)^2$ testing rounds. More recently, in [25] the authors extend traditional multiprocessor diagnosis to

handle change of failure state during the diagnostic process. All of these approaches are fundamentally different from ours since there is no separation between the payload and the monitor system. This implies the application system has to be observable as well as controllable (to generate the tests and analyze them).

Embedded system diagnosis: There has also been considerable work in the area of diagnosis of embedded systems, particularly in automotive electronic systems as in [26] (detection and shut down of faulty actuators in automotive systems) and [27] (use assertions to correlate anomalies from multiple components). The papers in this domain do not consider imperfect observability of the sensor input or the actuator output, possibly because of tight coupling between the payload and the monitor systems. They are focused on scheduling monitor processes under tight computational resource constraints while we do not have such constraints.

Debugging in distributed applications: There has been a spurt of work in providing tools for debugging problems in distributed applications – performance problems [28]-[30], misconfigurations [31], unexpected behavior [31], etc. The general flavor of the approaches in this domain is that the tool collects trace information at different levels of granularity (line of code to process) and the collected traces are automatically analyzed, often offline, to determine the possible root causes of the problem. For example, in [28], the debugging system performs analysis of message traces to determine the causes of long latencies. The goal of these efforts is to deduce dependencies in distributed applications and flag possible root causes to aid the programmer in the manual debug process, and not to produce automated diagnosis. Moreover tools like Pip [45], debug distributed applications by comparing the actual behaviour to the expected behaviour of an application. They give an expectation language that allows the programmer to express in a declarative language the expectations about the system. Also they provide tools for gathering events, checking behavior and visualizing valid and invalid behaviors which aid in tracking the bugs.

More recent work has produced powerful tools for debugging of distributed applications. In [45], the authors present a tool called liblog that aids in recreating the events that occurred prior to and

during failure. The replay can be done offline at a different site. The tool guarantees that the event state in its log will be consistent, i.e., no message is received before it has been sent. This work stops short of automated diagnosis. Some other mechanism, not described in the paper, is responsible for taking the replayed events and determining the root cause. There are several other offline tools that aid diagnosis, such as tools for data slicing [46], backtracking [47], and deterministic replay [48][49], but they all require manual effort in diagnosing the faulty components.

Online tools focus on data collection [48][50][51] rather than on automated diagnosis. A recent work [52] performs online failure diagnosis through a novel concept. When a failure occurs, some failure triggering events are captured right then. The system can roll back to a checkpoint in case of failure and when re-executing, heavy duty instrumentation is applied. Inspired by manual debugging, the Triage Diagnosis Protocol (TDP) provides a top-down failure-guided diagnosis. Based on the failure symptom and previous diagnosis results, TDP automatically decides the next step of instrumentation or analysis. The initial diagnosis happens at the user site, preserving the privacy of any user data, and then the detailed diagnosis happens at the programmer's site. This work relies on perfect checkpointing and the assumption that the state will not change between the original failing execution and the re-execution. It also assumes the application is white-box and can be instrumented. The work has to deal with the challenge of the user not trusting the debugging system completely. This problem is assumed away in our work—the Monitor can observe any interaction between the application components.

Network diagnosis: Diagnosis in IP networks is addressed in Shrink [54]. This tool used for root cause analysis of network faults models the diagnosis problem as a Bayesian network. It specifically diagnoses inaccurate mappings between IP and optical layers. The work in [55] studies the effectiveness and practicality of Tree-Augmented Naive Bayesian Networks (or TANs) as a basis for performing offline diagnosis and forecasting from system-level instrumentation in a three-tier network service. The TAN models are studied to select combinations of metrics and thresholds values that correlate with performance states of the systems (compliance with Service Level

Objectives). This approach differs from the Monitor approach in the sense that it relies on monitoring performance metrics rather than diagnosing the origin of the problem over a set of possible components.

Automated diagnosis in COTS systems: Automated diagnosis for black-box distributed COTS components is addressed in [32], [33]. The system model has replicated COTS application components, whose outputs are voted on and the replicas which differ from the majority are considered suspect. This work takes the restricted view that all application components are replicated and failures manifest as divergences from the majority. In [17], the authors present a combined model for automated detection, diagnosis, and recovery with the goal of automating the recovery process. However, the failures are all fail-silent and no error propagation happens in the system, the results of any test can be instantaneously observed, and the monitor accuracy is predictable.

In none of the existing work that we are aware of is there a rigorous treatment of the impact of the Monitor's constraints and limited observability on the accuracy of the diagnosis process. There are sometimes statements made on this without supporting reasoning – for example, in [28], it is mentioned that drop rates up to 5% do not affect accuracy of the diagnosis.

VII. Conclusion

In this paper we presented an online diagnosis system called the Monitor for arbitrary failures in distributed applications. The Monitor passively observes the message exchanges between the components of the application and at runtime, performs a probabilistic diagnosis of the component that was the root cause of a detected failure. The Monitor is compared to the state-of-the-art diagnosis framework called Pinpoint.

We tested the two systems on a 3-tier Java-based e-commerce system called PetStore. Extensive fault injection experiments were performed to evaluate the accuracy and precision of the two schemes. The Monitor outperformed Pinpoint particularly in precision, though its advantage

narrowed for interaction faults. As part of future work we are looking at diagnosis in high throughput network streams. In these streams, the Monitor may have to decide to drop some parts of a stream. We are looking into intelligent decision making to maintain a high accuracy. We are also investigating machine learning based diagnosis in the presence of uncertain information.

References

- [1] G. Khanna, P. Varadharajan, and S. Bagchi, "Self Checking Network Protocols: A Monitor Based Approach," In Proceedings of the 23rd IEEE Symposium on Reliable Distributed Systems (SRDS '04), pp. 18-30, October 2004.
- [2] M. Diaz, G. Juanole, and J.-P. Courtiat, "Observer-A Concept for Formal On-Line Validation of Distributed Systems," IEEE Trans. on Software Engineering, vol. 20, no. 12, pp. 900-913, Dec 1994.
- [3] M. Zulkernine and R. E. Seviora, "A Compositional Approach to Monitoring Distributed Systems," IEEE International Conference on Dependable Systems and Networks (DSN'02), pp. 763-772, Jun 2002.
- [4] K. Bhargavan, S. Chandra, P. J. McCann, and C. A. Gunter, "What Packets May Come: Automata for Network Monitoring," In ACM SIGPLAN Notices, vol. 36, no. 3, pp. 206-219, 2001.
- [5] "Snort Flexible Response Add-On," Available at: http://cerberus.sourceforge.com/~jeff/archives/snort/sp_respond2/
- [6] C. A. Carver, J. M.D. Hill, and U. W. Pooch, "Limiting Uncertainty in Intrusion Response," In Proceedings of the 2001 IEEE Workshop on Information Assurance and Security United States Military Academy, West Point, NY, 5-6 June, 2001.
- [7] R. Buskens and R. Bianchini Jr., "Distributed on-line diagnosis in the presence of arbitrary faults," In 23rd International Symposium on Fault-Tolerant Computing (FTCS-23), 1993.
- [8] I. Katzela and M. Schwartz, "Schemes for Fault Identification in Communication Networks," In IEEE/ACM Trans. On Networking, vol. 3 , no. 6 , pp. 753-764, Dec 1995.
- [9] M. A. Hiltunen, "Membership and System Diagnosis," In 14th IEEE Symposium on Reliable Distributed Systems (SRDS '95), pp. 208-217, Sep 1995.
- [10] C. Dwork, N. Lynch, and L. Stockmeyer, "Consensus in the presence of partial synchrony", In the Journal of the Association for Computing Machinery, Vol 35, No. 2, April 1988, pp. 288 – 323.
- [11] F.P. Preparata, G. Metze, R.T. Chien. "On the Connection Assignment Problem of Diagnosable Systems", In IEEE Transactions on Electronic Computers, Vol. EC-16, No. 6, December 1967,, pp. 848-854.
- [12] S. Maheshwari and S. Hakimi, "On models for diagnosable systems and probabilistic fault diagnosis," IEEE Transactions on Computers, C-25:228–236, Mar 1976.
- [13] D. Fussel and S. Rangarajan, "Probabilistic Diagnosis of Multiprocessor Systems with Arbitrary Connectivity," 19th Int. IEEE Symp. on Fault-Tolerant Computing, pp. 560-565, 1989.
- [14] Sungu Lee and Kang Shin, "Probabilistic diagnosis of multiprocessor systems," ACM Computing Surveys (CSUR), Volume 26 , Issue 1, 1994.
- [15] A. T. Dahbura, "System-level diagnosis: A perspective for the third decade," In Concurrent Computation: Algorithms, Architectures, Technologies Plenum, New York, 1988.

- [16] X. Y. Wang, D. S. Reeves and S. F. Wu, "Tracing Based Active Intrusion Response," In *Journal of Information Warfare*, Vol. 1 No. 1, September 2001.
- [17] K. R. Joshi, W. H. Sanders, M. A. Hiltunen, R. D. Schlichting, "Automatic Model-Driven Recovery in Distributed Systems," At the 24th IEEE Symposium on Reliable Distributed Systems (SRDS'05), pp. 25-38, 2005.
- [18] A. T. Bouloutas, S. Calo, and A. Finkel. Alarm correlation and fault identification in communication networks. *IEEE Transactions on Communications*, 42(2/3/4):523--533, 1994.
- [19] B. Gruschke, "Integrated Event Management: Event Correlation Using Dependency Graphs," at the 10th IFIP/IEEE International Workshop on Distributed Systems: Operations and Management (DSOM), pp. 130-141, 1998.
- [20] S. Kliger, S. Yemini, Y. Yemini, D. Ohsie, and S. Stolfo, "A coding approach to event correlation," *Intelligent Network Management*, no., pp. 266-277, 1997.
- [21] A. Frenkie and H. Lee, "EPP: A framework for measuring the end-to-end performance of distributed applications," in *Proc. Performance Engineering Best Practices Conference*, no., 1999.
- [22] I. Rish, M. Brodie, and S. Ma, "Intelligent probing: A cost-efficient approach to fault diagnosis in computer networks," *IBM Systems Journal*, vol. 41, no. 3, pp. 372-385, 2002.
- [23] I. Rish, M. Brodie, M. Sheng, N. Odintsova, A. Beygelzimer, G. Grabarnik, and K. Hernandez, "Adaptive diagnosis in distributed systems," *IEEE Transactions on Neural Networks*, vol. 16, no. 5, pp. 1088-1109, 2005.
- [24] E. P. Duarte, Jr. and T. Nanya, "A hierarchical adaptive distributed system-level diagnosis algorithm," *IEEE Transactions on Computers*, vol. 47, no. 1, pp. 34-45, 1998.
- [25] S. Arun and D. M. Blough, "Distributed diagnosis in dynamic fault environments," *IEEE Transactions on Parallel and Distributed Systems*, vol. 15, no. 5, pp. 453-467, 2004.
- [26] N. Kandasamy, J. P. Hayes, and B. T. Murray, "Time-constrained failure diagnosis in distributed embedded systems," at the *International Conference on Dependable Systems and Networks (DSN)* pp. 449-458, 2002.
- [27] P. Peti, R. Obermaisser, and H. Kopetz, "Out-of-norm assertions [diagnostic mechanism]," at the 11th IEEE Real Time and Embedded Technology and Applications Symposium (RTAS) pp. 280-291, 2005.
- [28] M. K. Aguilera, J. C. Mogul, J. L. Wiener, P. Reynolds, and A. Muthitacharoen, "Performance debugging for distributed systems of black boxes," at the *Proceedings of the nineteenth ACM symposium on Operating systems principles*, Bolton Landing, NY, USA, pp. 74-89, 2003.
- [29] M. Y. Chen, E. Kiciman, E. Fratkin, A. Fox, and E. Brewer, "Pinpoint: problem determination in large, dynamic Internet services," at the *International Conference on Dependable Systems and Networks (DSN)*, pp. 595-604, 2002.
- [30] P. Barham, R. Isaacs, R. Mortier, and D. Narayanan, "Magpie: online modelling and performance-aware systems " at the 9th Workshop on Hot Topics in Operating Systems (HotOS IX), pp. 85-90, 2003.
- [31] H. J. Wang, J. Platt, Y. Chen, R. Zhang, and Y.-M. Wang, "PeerPressure for automatic troubleshooting," at the *Proceedings of the joint international conference on Measurement and modeling of computer systems*, New York, NY, USA, pp. 398-399, 2004.
- [32] A. Bondavalli, S. Chiaradonna, D. Cotroneo, and L. Romano, "Effective fault treatment for improving the dependability of COTS and legacy-based applications," *Dependable and Secure Computing*, *IEEE Transactions on*, vol. 1, no. 4, pp. 223-237, 2004.
- [33] L. Romano, A. Bondavalli, S. Chiaradonna, and D. Cotroneo, "Implementation of threshold-based diagnostic mechanisms for COTS-based applications," at the 21st IEEE Symposium on Reliable Distributed Systems, , pp. 296-303, 2002.

- [34] A. Brown and D. A. Patterson, "Embracing Failure: A Case for Recovery-Oriented Computing (ROC)," 2001 High Performance Transaction Processing Symposium, Asilomar, CA, October 2001.
- [35] G. Candea, E. Kıcıman, S. Kawamoto, and A. Fox, "Autonomous Recovery in Componentized Internet Applications," Cluster Computing Journal, Vol. 9, Number 1 (February 2006).
- [36] TPC- Benchmark: <http://www.tpc.org/tpcw>.
- [37] PetStore J2EE Application: <http://java.sun.com/blueprints/code/index.html>.
- [38] JBoss Application Server: <http://labs.jboss.com>.
- [39] MySQL: Open Source Database: www.mysql.com.
- [40] META Group, Inc., "Quantifying Performance Loss: IT Performance Engineering and Measurement Strategies", November 22, 2000. Available at: <http://www.metagroup.com/cgi-bin/inetcgi/jsp/displayArticle.do?oid=18750>
- [41] FIND/SVP, 1993, "Costs of Computer Downtime to American Businesses," At: www.findsvp.com.
- [42] Addendum to SRDS '07 submission-data files and communication: www.ece.purdue.edu/~sbagchi/Papers/srds07_add.pdf
- [43] F. Cuppens and A. Mieke, "Alert correlation in a cooperative intrusion detection framework," Proceedings of the 2002 IEEE Symposium on Security and Privacy, May 12-15, 2002.
- [44] H. Debar and A. Wespi, "Aggregation and Correlation of Intrusion Detection Alerts," Proceedings of the 4th Symposium on Recent Advances in Intrusion Detection (RAID 2001), Davis, CA, USA, October 2001. Springer LNCS 2212, pages 85-103.
- [45] D. Geels, G. Altekar, S. Shenker, and I. Stoica, "Replay Debugging for Distributed Applications," USENIX Annual Technical Conference, pp. 289-300, 2006.
- [46] X. Zhang, R. Gupta, and Y. Zhang, "Precise dynamic slicing algorithms," ICSE, pp. 319-329, 2003.
- [47] S. T. King, G. W. Dunlap, and P. M. Chen, "Debugging operating systems with time-traveling virtual machines," USENIX Annual Technical Conference, pp. 1-15, 2005.
- [48] S. Narayanasamy, G. Pokam, and B. Calder, "BugNet: Continuously recording program execution for deterministic replay debugging," ISCA, pp. 284-295, 2005.
- [49] S. M. Srinivasan, S. Kandula, C. R. Andrews, and Y. Zhou, "Flashback: A lightweight extension for rollback and deterministic replay for software debugging," USENIX Annual Technical Conference, pp. 29-44, 2004.
- [50] A. Ayers, R. Schooler, C. Metcalf, A. Agarwal, J. Rhee, and E. Witchel, "TraceBack: First fault diagnosis by reconstruction of distributed control flow," PLDI, pp. 201-212, 2005.
- [51] D. Brumley, J. Newsome, D. Song, H. Wang, and S. Jha, "Towards automatic generation of vulnerability-based signatures," In IEEE Symposium on Security and Privacy, pp. 2-16, 2006.
- [52] J. Tucek, S. Lu, C. Huang, S. Xanthos and Y. Zhou, "Automatic Online Failure Diagnosis at the End-User Site", Second Workshop on Hot Topics in System Dependability, HotDep'06.
- [53] P. Reynolds, C. Killian, J. Wiener, J. Mogul, M. Shah, and A. Vahdat, "Pip: Detecting the Unexpected in Distributed Systems," In Proceedings of NSDI, San Jose, CA, May 2006.
- [54] S. Kandula, D. Katabi, J. Vasseur, "Shrink: A Tool for Failure Diagnosis in IP Networks," ACM SIGCOMM Workshop on mining network data (MineNet-05), Philadelphia, PA, August 2005.
- [55] I. Cohen, M. Goldszmidt, T. Kelly, S. Julie, J. Chase, "Correlating instrumentation data to system states: A building block for automated diagnosis and control", Operating System Design and Implementation (OSDI), Dec 2004.