

Using Runtime Paths for Macro Analysis

Mike Chen, Emre Kıcıman, Anthony Accardi, Armando Fox, Eric Brewer
{mikechen, brewer}@cs.berkeley.edu, {emrek, fox}@cs.stanford.edu, anthony@tellme.com

Abstract

We introduce macro analysis, an approach used to infer the high-level properties of dynamic, distributed systems, and an indispensable tool when faced with tasks where local context and individual component details are insufficient. We present a new methodology, runtime path analysis, where paths are traced through software components and then aggregated to understand global system behavior via statistical inference. Our approach treats components as gray boxes and complements existing micro analysis tools, such as code-level debuggers. We use runtime paths to deduce application state, detect failures, and diagnose problems, all in an application-generic fashion.

1 Introduction

Divide and conquer, layering, and replication are fundamental design principles useful for building large, complex systems, such as Internet services, sensor networks, and peer-to-peer (P2P) systems. Such techniques make building large systems tractable, as they improve availability, increase code reuse, and simplify high-level application structure. Unfortunately, debugging ease and fault monitoring do not scale as well, since global context tends to be dispersed across many small components. Building large, complex systems that are reliable, yet maintainable and extensible, remains a challenge.

Existing debugging techniques make use of various micro analysis tools, such as code-level debuggers and application logs. Such tools tend to provide knowledge limited to component internals, or furnish a thread-level perspective, so that the execution context is lost at the thread boundaries. While they provide valuable, localized knowledge, many of these tools fail to capture aggregate component behavior and macro system properties. By way of analogy, micro analysis allows you to see the details of each honeybee, but macro analysis is needed to understand how the bees interact to keep a beehive functioning.

Various systems have exposed and exploited non-local system context to address performance problems [11, 1]. Macro analysis makes use of non-local context to improve system management and reliability. This is especially im-

portant for large, dynamic systems, where execution context may be distributed across many components.

One key observation we make about dynamic, distributed systems is that most of them have a single system-wide execution path associated with each request that they service. Examples include Internet services that have request/response paths, and P2P systems and sensor networks that have one-way message paths. By tracing these runtime paths, we expose and connect various local contexts dispersed throughout the system. We then use statistics to analyze many of these paths and thereby better understand the system's behavior.

There are several open, challenging problems that can benefit from the high-level system perspective that macro analysis provides:

Deducing system structure: Systems evolve through both changes to their components and changes in how these components interact. Understanding such inter-component relationships enables developers and operators to anticipate potential conflicts and debug problems. Unfortunately, current techniques for tracking changes in these relationships rely on error-prone, manual documentation, which is infeasible for rapidly changing systems. As systems grow and increase in complexity, we desire automated mechanisms for deducing system structure and tracking its evolution.

Detecting application-level failures: Despite our best efforts at unit testing and quality assurance, services still fail. Worse still, many application-level faults are only seen by end users after deployment, even though systems are constantly monitored for signs of failure. One large commercial service has found that such errors take considerably longer to detect than lower-level failures. The difficulty is that broken or misconfigured components or bad component interactions may only exhibit symptoms at the application level, and the global context required to programmatically diagnose such application-level failures is usually not available.

Diagnosing failures: Failures often manifest themselves far from their root cause. In the extreme case, faults are not detected within the system's boundaries at all

and are only visible to external observers. Unfortunately, existing debugging and diagnosis tools have a limited, local view of the system, and thus work best when failures manifest themselves close to the cause. We desire tools that use global failure information to help operators and developers identify the root cause.

The contributions of this paper are:

1. Recognizing that macro analysis is critical when developing, evolving, and maintaining reliable systems as they grow in size and complexity.
2. A path-based macro analysis framework, where we first record the components and resources used to service each request, and then use statistical analysis techniques to deduce system structure, detect application-level failures, and diagnose problems.

The paper is organized as follows: Section 2 develops the runtime path model. Section 3 describes our analysis framework and current status. Section 4 discusses our results addressing the challenging problems above, both in a research setting and as part of a commercial infrastructure at Tellme Networks. We outline future research directions in Section 5 and discuss related work in Section 6.

2 Runtime Paths

We extend the dataflow paths in Scout [11] and Ninja [13]¹ to incorporate runtime properties. A runtime path is the control flow, resources, and performance characteristics associated with servicing a request. Paths can be recorded during runtime by tracing each request through a live system, spanning the system’s layers to access direct component and resource dependencies. Each path then provides a vertical slice of the system from a request’s perspective.

We use the term “request” in a broad sense to mean a unit of work. This includes both requests that require responses (e.g., HTTP) and those that don’t (e.g., one-way messages).

There are two main requirements for a system to support runtime paths. First, it must be possible to associate a unique path with each distinct request. For example, if the same request is handled by different components in different, possibly distributed processes, we must be able to make this connection, perhaps by using a unique ID that travels with the request. Second, it must be possible to report *observations* and associate them with the components that made them. For a pipelined system, a logging mechanism together with knowledge of component request entry and exit would be sufficient.

The components make local observations, from which global context is obtained by stringing them together along

¹Scout: “a logical channel through a multi-layered system over which I/O data flows within a single host”. Ninja: “a flow of typed data through multiple services across the wide area”.




Figure 1. Analysis Framework.

a runtime path. Each observation contains information about some active component, such as its name, location, timestamp, latency, and arguments. The tracing should be extensible to allow for integration with micro analysis techniques. For example, identifiers could be included in each observation to provide a link with standard application logs.

3 Analysis Framework

The analysis framework consists of five major modules, as illustrated in Figure 1. The *Tracer* tracks requests through the target system, reporting any observations made. Although *Tracer* is platform-specific, it can be application-generic for platforms that host application components by monitoring requests that enter and exit the components.

The *Aggregator* receives these observations and reconstructs the runtime paths, which the *Path Repository* stores. The *Statistical Declarative Query Engine* handles the data management complexity. It enables us to transparently optimize data storage, and evaluate and update different analysis algorithms. Monitoring and debugging tools should be built on top of this engine. The *Visualization* module helps users understand system behavior. Paths have a natural graph representation: nodes are observations and edges indicate request propagation.

We implemented an extensible *Aggregator*, a *Path Repository*, and a *Query Engine*, with plugins for data clustering and structural anomaly detection. We are currently experimenting with different analysis algorithms.

We built a *Tracer* for a web server, Jetty, and a clustered J2EE application server, JBoss [9]. The web server inserts a unique request ID into the HTTP request header. This ID is placed in thread local storage for intra-thread component calls, and passed via a modified RMI library for inter-thread calls. The total code impact was 428 lines in 10 files.

We have been running *PetStore*, an e-commerce application, and *ECperf* [15], an industry-standard benchmark for J2EE application servers. Both have a 3-tier architec-

ture consisting of a web server, application components, and a database. The tracing instrumentation is application-generic, so no application changes were necessary.

4 Applying Macro Analysis

4.1 Deducing System Structure

Understanding a service’s structure, including the relationship between external requests and the service’s internal components and state, enables developers and operators to anticipate potential problems before they upgrade the system. Knowing how components and shared state are used is critical when debugging failed requests. Dependency models have been proposed to improve system reliability and availability [7], but there are few techniques that generate such models automatically.

Key idea: Paths directly capture application structure.

Runtime paths record how a system services real requests, which compares favorably with error-prone, human-generated models and static analysis that predict how the system *might* service such requests. Automatically generated models help developers and operators understand the *actual* behavior of systems under investigation, and can be used as input to recovery mechanisms, such as recursive restarts [3], to reduce mean time to repair.

The Magpie project is using macro analysis techniques to generate detailed and accurate models of distributed system workloads [10]. While currently used for capacity planning, these models may also be applied to performance debugging, system tuning, and fault diagnosis.

Key idea: Paths associate requests with internal state.

Internet services typically store persistent state in a database to allow for easier front-end scaling. Different requests often share persistent state, but the components handling each request may be unaware of any such sharing. For example, although checkout and login requests may seem like independent HTTP requests, they may share a user profile. A bug during checkout could corrupt this shared state and cause subsequent login failures. Such bugs are difficult to diagnose without an understanding of how various dispersed local contexts depend on each other.

By tracing runtime paths from the web servers, through the application components, and to the databases, we can easily determine how state is shared across requests.

Table 1 shows the mapping between request types and their reads and writes to database tables in PetStore, our desired level of state granularity.

4.2 Detecting Failures using Anomalies

Key idea: Paths often behave differently in failure modes. Hence we can detect failures via changes in path behavior.

Request Type	Database Tables				
	Product	Signon	Account	Banner	Inventory
verifysignin	R	R	R		
cart	R			R	R/W
commitorder	R				W
category	R				
search	R			R	
productdetails	R				R/W
newaccount		R	R		
checkout					W

Table 1. An automatically generated partial state dependency table for PetStore. To determine which request types share state, group the rows by common entry under the desired column. For example, the checkout request only writes to the Inventory table, and shares state with three other requests: cart, commitorder, and productdetails.

Application-level failure detection remains a major challenge today. In practice, quality assurance testing mainly catches simple bugs. Many complex bugs exist in deployed software because of difficulties accurately simulating the workload of a production environment, difficulties modeling the production environment itself, incomplete test coverage, and economic factors. Detecting these bugs in a live system can be difficult, since many bug symptoms are only evident to an end-user, such as incorrect text on a web page.

If we analyze a live system using macro analysis techniques, however, we can often see secondary effects of failures. Many errors cause runtime paths to end prematurely, while others send paths to less-often used error handlers. Still others, such as fail-stutter faults, simply cause deviations in the latencies of particular path components.

One macro analysis technique for detecting these changes in path behavior is to search for asymmetries in the interactions among replicated, load-balanced components. Consider an Internet service where all but one of the middle tier nodes are sending queries to a database. Since the middle tier nodes are replicas of one another, it is likely that the one node’s database inactivity is a symptom of some problem. We can extend this idea to treat excessively heavy or light component interaction volume as a sign of failure.

Another technique is to group paths by request type and search for significant deviations in latency or structure.

4.3 Diagnosing Failures through Correlations

Key idea: Runtime paths make the root cause’s interaction with a failed request apparent, so that we can quickly explain a failure by a bad component-level or systemic behavior, and can also quickly assess the user visible impact (and hence the priority) of the problem.

When something fails, we want to know why. The challenge here is in tracing externally observed (application-level) failures back to a system fault or root cause.

In practice, we start with sets of suspected failed requests (e.g., reported by users or discovered via anomaly detection) and successful requests, and face the task of identifying the runtime path features underlying any real failures.

The diagnosis task can be cast as a data mining problem, by using data clustering to group the components associated with each failed request’s path. Using a previous prototype, Pinpoint [4], we showed that for single component failures the data clustering approach provides a tradeoff between accuracy, at 70-90%, and false positives, at 20-40%. This compares favorably with direct fault detection and other automatic analysis methods, which either offer 40% accuracy or many (almost 90%) false positives. Also, we found that paths are vital when dealing with multi-component faults.

Alternatively, the task can be cast as a classification problem in the machine learning domain. Here, the root causes would be the strongest classification rules.

An important benefit of our failure diagnosis approach is that the runtime path data links the logically separate tasks of failure detection and diagnosis, which enables an understanding of mechanisms where the connection between problem causes and symptoms is not otherwise apparent.

4.4 A Commercial Example

Tellme Networks has developed a path-based macro analysis infrastructure, its *observation logging system*, to help ensure the high reliability of the Tellme Network. Tellme runs voice applications; for the purpose of our current discussion, the system is a telephony network with an Internet back-end, and serviced requests include an audio response to a telephone caller’s voice query.

After explaining how runtime paths behave in this example, we will discuss a failure and describe how we used paths to first diagnose the problem, then deduce runtime system structure, and ultimately craft a detection algorithm.

An actual, although simplified, voice response path is illustrated in Figure 2. There are 31 observations in this runtime path, plotted by the relative time at which each observation was made. We call attention to 4 of these, indicating when the user begins and finishes speaking an utterance and when we start and stop sending audio in response.

The latency profile of the runtime path shown in Figure 2 is close to ideal for this system; most of the time is spent listening to or talking to the user, and the logic connecting the request with the response is rapidly executed. An important metric is the latency between end of speech and start of playback. This 1760 ms delay is perceived by the user, and indicates how quickly the system appears to respond.

Now consider a real network failure. In this case, the process supplying the audio to the machine on the telephone




Figure 2. A runtime path from the Tellme Network.

network experienced an internal fault, so that it produced very short waveforms instead of the desired ones. This failure is not catastrophic at the system level; an occasional skipped waveform may go unnoticed at best or confuse the user at worst. The problem is therefore difficult to detect via low-level system monitoring and requires significant application knowledge to handle effectively.

Despite this challenge, we quickly diagnosed this problem using global context in the observation logs. A caller, familiar with a particular application that exhibited the symptoms, provided details of his phone call that enabled us to quickly locate the relevant runtime path. Once visualized as in Figure 2, a short 20 ms playback time suggested a skipped waveform, the preceding observations confirmed that the audio server process thought it had successfully serviced the audio request, and observation context information pointed to logs for the suspect process, which made the root cause evident. All this information was found in this one path.

Once we understood the runtime path characteristics for this failure, we were able to query the observation logs to detect all similar occurrences throughout the Tellme Network. We deduced enough system state to know which components affected which applications, so we could isolate the failing components and assess application impact.

Finally, we were able to craft a monitor to detect failures upon seeing an anomalous proportion of suspect paths with structure similar to that of the path for the failed request.

5 Future Directions

In addition to refining our framework and exploring different algorithms, we are applying our path-based macro analysis methodology to sensor networks and P2P systems.

Key idea: Violations of macro invariants are signs of system intrusion or buggy implementations. Macro analysis can help discover invariants, detect violations, and pinpoint the offending components.

Because of the highly distributed and dynamic nature of these systems, many domain-specific macro invariants are difficult to validate using micro analysis or static analysis [6]. Consider an upper bound on the number of hops made during message delivery in a P2P system as a macro invariant. By applying root cause analysis, we can identify peers that incorrectly route messages. In this example, although each node may detect an invariant violation, a diagnosis is difficult without the context contained in a runtime path.

6 Related Work

We consider both micro and macro analysis work, as well as hybrid approaches.

Macro analysis Magpie [10] profiles web sites to observe the processing state machine for each HTTP request and to measure request resource consumption (CPU, disk, and network usage) at each stage. The focus is on building probabilistic models of the workload suitable for performance prediction, tuning, and diagnosis.

Micro analysis: Anomaly detection has been used to identify intrusions from resource usage [5], system calls [8], and network packets [12]. Paths provide non-local context and may make the detection of a new class of intrusions possible.

Hybrid: There are several recent commercial request tracing systems of note. PerformaSure [14] and AppAssure [2] focus on performance diagnosis. IntegriTea [16] focuses on capturing and replaying failure conditions. These systems work with isolated requests, while we aggregate multiple paths and use statistical techniques to infer collective system behavior.

7 Conclusion

Macro analysis satisfies a need when monitoring and debugging large, complex systems where local context is of insufficient use. To this end, we have presented a runtime path-based approach along with a family of macro analysis tools that are proving effective in addressing several challenging and important problems. Our method involves dynamically tracing runtime paths through live systems, and recording local observations about performance, interacting components, and resource usage along the way. We subsequently apply data mining techniques to statistically infer aggregate system behavior. This approach is applicable to a large variety of systems, and complements existing micro analysis tools that provide additional insight into individual components.

Our results with distributed Internet systems demonstrate promising progress in 1) deducing system structure and dependencies, 2) detecting failures via path anomalies, and 3) diagnosing problems. We plan to extend our methodology

to peer-to-peer systems and sensor networks by validating macro invariants.

References

- [1] M. B. Abbott and L. L. Peterson. Increasing Network Throughput by Integrating Protocol Layers. *IEEE/ACM Transactions on Networking*, 1(5):600–610, 1993.
- [2] Alignment Software. AppAssure, 2002. <http://www.alignmentsoftware.com/>.
- [3] G. Candea and A. Fox. Recursive Restartability: Turning the Reboot Sledgehammer into a Scalpel. In *HotOS VIII*, 2001.
- [4] M. Chen, E. Kıcıman, E. Fratkin, E. Brewer, and A. Fox. Pinpoint: Problem Determination in Large, Dynamic Internet Services. In *Symposium on Dependable Networks and Systems (IPDS Track)*, 2002.
- [5] G. W. Dunlap, S. T. King, S. Cinar, M. Basrai, and P. M. Chen. ReVirt: Enabling Intrusion Analysis through Virtual-Machine Logging and Replay. In *OSDI*, 2002.
- [6] D. Engler, B. Chelf, A. Chou, and S. Hallem. Checking System Rules Using SystemSpecific, Programmer-Written Compiler Extensions. In *OSDI*, 2000.
- [7] B. Gruschke. A New Approach for Event Correlation based on Dependency Graphs. In *5th Workshop of the OpenView University Association*, 1998.
- [8] S. A. Hofmeyr, S. Forrest, and A. Somayaji. Intrusion Detection Using Sequences of System Calls. *Journal of Computer Security*, 6(3):151–180, 1998.
- [9] JBoss.org. JBoss J2EE Application Server, 2001. <http://www.jboss.org>.
- [10] Microsoft Research. Magpie, 2003. <http://research.microsoft.com/projects/magpie/>.
- [11] D. Mosberger and L. L. Peterson. Making Paths Explicit in the Scout Operating System. In *OSDI*, 1996.
- [12] V. Paxson. Bro: A System for Detecting Network Intruders in Real-time. *Computer Networks (Amsterdam, Netherlands: 1999)*, 31(23–24):2435–2463, 1999.
- [13] S. D. Gribble et al. The Ninja Architecture for Robust Internet-scale Systems and Services. *Computer Networks*, 35(4):473–497, 2001.
- [14] Sitraka. PerformaSure, 2002. <http://www.sitraka.com/software/performasure/>.
- [15] Sun Microsystems. ECperf J2EE benchmark, 2001. <http://java.sun.com/j2ee/ecperf/>.
- [16] TeaLeaf Technology. Integritea, 2002. <http://www.tealeaf.com/>.