NASA/TM 2008-214570

# Testing HyDE on ADAPT

*Adam Sweet*
*NASA Ames Research Center*

National Aeronautics and
Space Administration

Ames Research Center
Moffett Field, California, 94035-1000

**January 2008**

## The NASA STI Program Office . . . in Profile

Since its founding, NASA has been dedicated to the advancement of aeronautics and space science. The NASA Scientific and Technical Information (STI) Program Office plays a key part in helping NASA maintain this important role.

The NASA STI Program Office is operated by Langley Research Center, the Lead Center for NASA's scientific and technical information. The NASA STI Program Office provides access to the NASA STI Database, the largest collection of aeronautical and space science STI in the world. The Program Office is also NASA's institutional mechanism for disseminating the results of its research and development activities. These results are published by NASA in the NASA STI Report Series, which includes the following report types:

- TECHNICAL PUBLICATION. Reports of completed research or a major significant phase of research that present the results of NASA programs and include extensive data or theoretical analysis. Includes compilations of significant scientific and technical data and information deemed to be of continuing reference value. NASA's counterpart of peer-reviewed formal professional papers but has less stringent limitations on manuscript length and extent of graphic presentations.

- TECHNICAL MEMORANDUM. Scientific and technical findings that are preliminary or of specialized interest, e.g., quick release reports, working papers, and bibliographies that contain minimal annotation. Does not contain extensive analysis.

- CONTRACTOR REPORT. Scientific and technical findings by NASA-sponsored contractors and grantees.

- CONFERENCE PUBLICATION. Collected papers from scientific and technical conferences, symposia, seminars, or other meetings sponsored or cosponsored by NASA.

- SPECIAL PUBLICATION. Scientific, technical, or historical information from NASA programs, projects, and missions, often concerned with subjects having substantial public interest.

- TECHNICAL TRANSLATION. English-language translations of foreign scientific and technical material pertinent to NASA's mission.

Specialized services that complement the STI Program Office's diverse offerings include creating custom thesauri, building customized databases, organizing and publishing research results . . . even providing videos.

For more information about the NASA STI Program Office, see the following:

- Access the NASA STI Program Home Page at *http://www.sti.nasa.gov*

- E-mail your question via the Internet to help@sti.nasa.gov

- Fax your question to the NASA Access Help Desk at (301) 621-0134

- Telephone the NASA Access Help Desk at (301) 621-0390

- Write to:
  NASA Access Help Desk
  NASA Center for AeroSpace Information
  7121 Standard Drive
  Hanover, MD 21076-1320

NASA/TM 2008-214570

# Testing HyDE on ADAPT

*Adam Sweet*
*NASA Ames Research Center*

National Aeronautics and
Space Administration

Ames Research Center
Moffett Field, California, 94035-1000

**January 2008**

Available from:

NASA Center for AeroSpace Information
7121 Standard Drive
Hanover, MD 21076-1320
(301) 621-0390

National Technical Information Service
5285 Port Royal Road
Springfield, VA 22161
(703) 487-4650

# Testing HyDE on ADAPT

# September 30, 2007

## 1. Summary

The IVHM Project in the Aviation Safety Program has funded research in electrical power system (EPS) health management. This problem domain contains both discrete and continuous behavior, and thus is directly relevant for the hybrid diagnostic tool HyDE. In FY2007 work was performed to expand the HyDE diagnosis model of the ADAPT system. The work completed resulted in a HyDE model with the capability to diagnose five times the number of ADAPT components previously tested. The expanded diagnosis model passed a corresponding set of new ADAPT fault injection scenario tests with no incorrect faults reported. The time required for the HyDE diagnostic system to isolate the fault varied widely between tests; this variance was reduced by tuning HyDE input parameters. These results and other diagnostic design trade-offs are discussed. Finally, possible future improvements for both the HyDE diagnostic model and HyDE itself are presented.

## 2. Introduction

### 2.1. ADAPT

The Advanced Diagnostic and Prognostic Testbed (ADAPT), located at NASA Ames Research Center, was developed to measure, evaluate, and mature diagnostic and prognostic technologies. It incorporates an electrical power subsystem (EPS) in which faults may be injected by manual or software means. The layout of the ADAPT power system is shown below in Figure 1. The EPS includes elements common to many aerospace applications: power storage, power generation, and power distribution. The power storage consists of three battery modules. Each of the three batteries can be charged by one of the two battery chargers in the power generation element. Finally, any of the three batteries can be used to power either of the two load banks in the power distribution element. This design gives the ADAPT EPS basic redundancy and reconfiguration capability. Note that ADAPT is not a high-fidelity EPS testbed for any particular vehicle system. Rather, the power system serves as a problem domain for testing diagnostic and prognostic applications. These diagnostic systems are referred to as test articles (TAs) for ADAPT.

Test articles are commonly software algorithms, although they may incorporate additional sensing hardware as well. ADAPT has worked with TAs from NASA, academia, and industry. The testing procedure is usually scenario-based, where each scenario may have faults injected into the system. To detect the faults, a test article has access to the telemetry (commands and sensor data) from the ADAPT EPS. The

telemetry and the output of the diagnostic system are saved to a database, and the test article performance is evaluated according to a set of figures of merit.

More information on ADAPT can be found in [1].

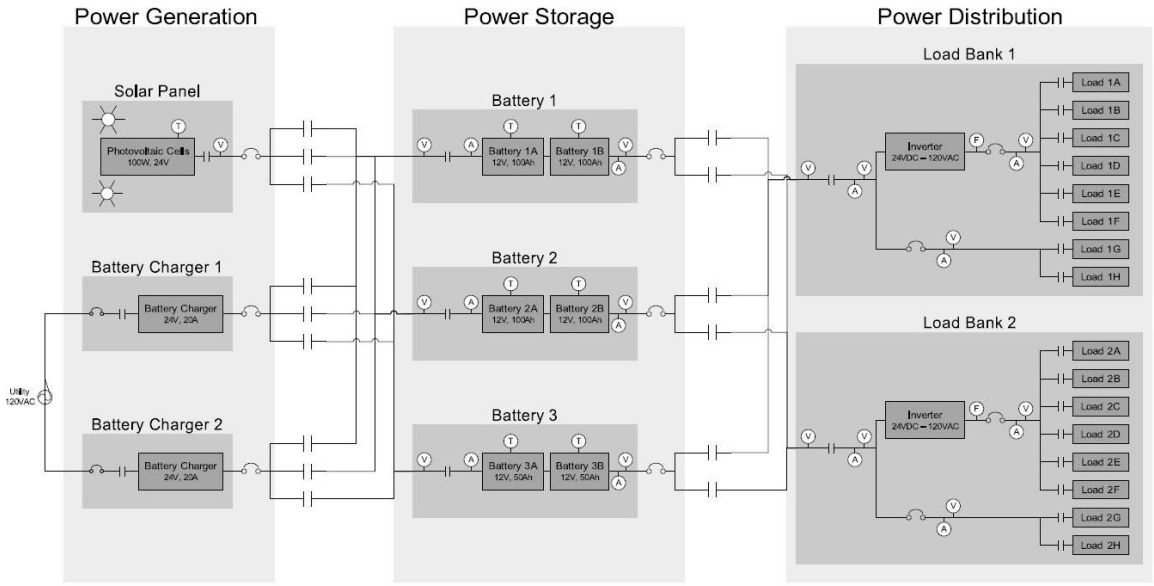


**Figure 1: ADAPT power system**

## 2.2. HyDE

HyDE is a system-monitoring tool developed at NASA Ames Research Center. The general capability of HyDE is to track the state of the target system over time, even as the system progresses through non-observable and fault states. Following from this general capability, HyDE is capable of sensor fusion, fault detection, and fault isolation in the presence of multiple faults. HyDE is a model-based system, meaning it is an inference engine which reasons with a declarative model containing information about a target system. This is in contrast to more traditional approaches to system diagnosis such as rule-based systems, expert systems, and fault trees. The core software and algorithms which make up HyDE are reused across multiple diagnosis applications, and the model is changed in order to adapt HyDE to a particular system.

A HyDE diagnosis consists of one or more candidates, where each candidate is a possible complete state estimate of the target system. A candidate may contain any number of faults, and in general candidates with fewer faults are considered more likely. Note that the candidates contained in a diagnosis are disjuctive: if the first candidate in a diagnosis contains only fault A, and the second candidate contains only fault B, it is incorrect to say that both A and B have failed.

### 2.2.1. Unique HyDE features

There are two main innovations within the HyDE system. The first is the search algorithm used for fault isolation, called conflict-directed search. In this search, the knowledge about what sensor data disagrees with the model's prediction is used to guide the search of the fault space. This conflict-directed search algorithm is inherited from other model-based diagnosis tools. However, HyDE is the first system in which the conflict-directed search algorithm has been used on hybrid systems, which allow models with both discrete and continuous variables and behavior.

The second main innovation of HyDE is its extensible architecture. An interface is defined for users and researchers to create their own types of models and compatible diagnostic algorithms, and use them within the HyDE framework.

HyDE is described in more detail in [2], and the next sections briefly describe HyDE modeling and the HyDE reasoning process.

### 2.2.2. Overview of HyDE modeling

A HyDE model contains system-specific information which HyDE uses to track the target system over time. The elements which make up a HyDE model are components, commands, modes, mode transitions, variables, constants, domains for the variables, and constraints. Logical and basic arithmetic constraints are supported, however more advanced functions such as logarithm and trigonometric functions are not yet supported. First-order differential equations are supported. The syntax of these elements is described in detail in [2].

The system-specific information is gathered from several sources. Most important is a description of the system, such as a schematic. The concept of operations is useful to understand how the system will behave. Finally, a failure modes and effect analysis (FMEA) or similar document is needed to define the diagnostic scope, and give information about the various possible faults to include in the model.

A HyDE model differs from many diagnostic models in two important ways. First, it is predictive. A predictive model means that the model predicts the target system's sensor values according to the current estimated system state. Both nominal and fault modes are modeled in this predictive way. Thus, a predictive model resembles a simulation model. This direction of reasoning is the opposite of many diagnostic systems, which start with sensor data and attempt to classify it into fault signatures.

A second difference in HyDE modeling is that the model is component-based. That means usually only components and their interconnections are modeled, and the HyDE engine determines the behavior of the overall system. It is possible to specify system-level models and constraints in the HyDE model. However, doing so limits the modeler's ability to create reusable component models.

Finally, there is considerable flexibility in the HyDE modeling language. HyDE models may be made with real-valued variables, interval-valued variables, or discrete-valued variables. Systems may be modeled according to the system's physics equations or any abstraction of the system behavior expressible in constraints. These choices will be made by the modeler evaluating trade-offs in the diagnostic system design.

### 2.2.3. Overview of HyDE reasoning process

The HyDE reasoning process is summarized below in Figure 2:

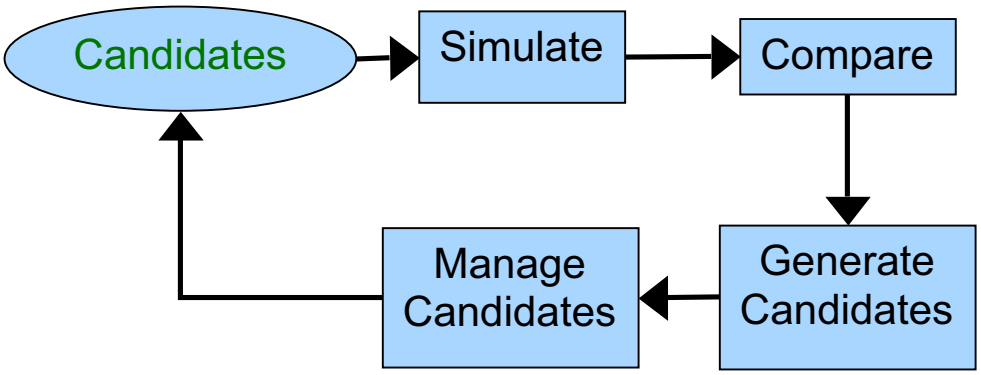


**Figure 2: HyDE reasoning process**

HyDE maintains a set of candidates, which are possible states of the monitored system. To begin monitoring, HyDE is given the initial state of the system, which must not contain any faults. Beginning from that state, HyDE then simulates the system forward in time. When sensor values from the system are presented to HyDE, it compares those values to what the simulation predicts the sensor values to be. If they are consistent, then there is no need to generate and manage candidates: the current candidate(s) are kept and the flow returns to the oval in Figure 2. If the sensor values are inconsistent, then HyDE generates distinct candidates which contain fault modes. It then uses the model constraints to predict the sensor values for each new candidate, and again tests the generated candidate against the sensor information. If the generated candidate is consistent, that candidate is kept. The most likely candidates are generated first, and the generation process continues until the first user-defined search termination parameter is reached.

Again, more details on the HyDE algorithm can be found in [2].

## 3. HyDE model

### 3.1. Expansion over previous HyDE system

HyDE had previously been deployed on ADAPT and undergone a set of acceptance tests. At that time, the diagnostic scope included only the distribution unit of the ADAPT system. This work under the IVHM Project in the Aviation Safety program significantly expanded the diagnostic capability of the HyDE deployment on ADAPT to include elements from the charging and load units as well. The test plans for both are given in Appendix A, and a summary comparison is given below in Table 1:

| | Previous HyDE test | Expanded HyDE test |
|---|---|---|
| Number of fault types | 4 | 8 |
| Number of components | 24 | 155 |

**Table 1: Expansion over previous HyDE system**

The exact scenarios to be tested were chosen at random from a list which included all of the components above. At least one component from all of the fault types was tested.

### 3.2. HyDE model of ADAPT

The current HyDE model of ADAPT is show below in Figure 3:

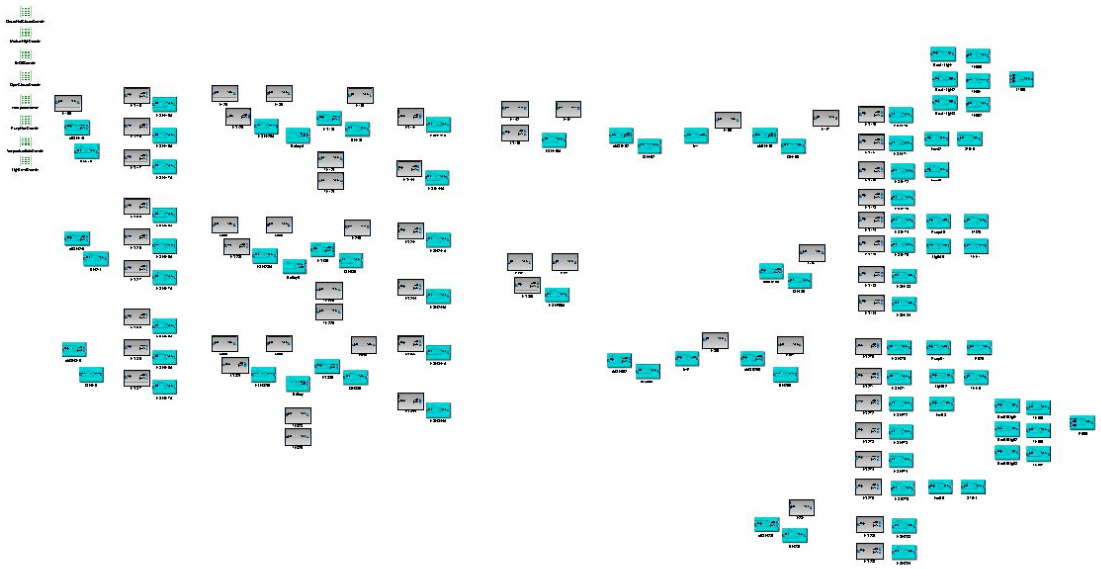

**Figure 3: HyDE model of ADAPT**

The entire model contains components to represent relays, inverters, sensors, circuit breakers, and loads. However, rather than explain the details of the entire model, we will focus on the relay component model as one of the most common components in the ADAPT model.

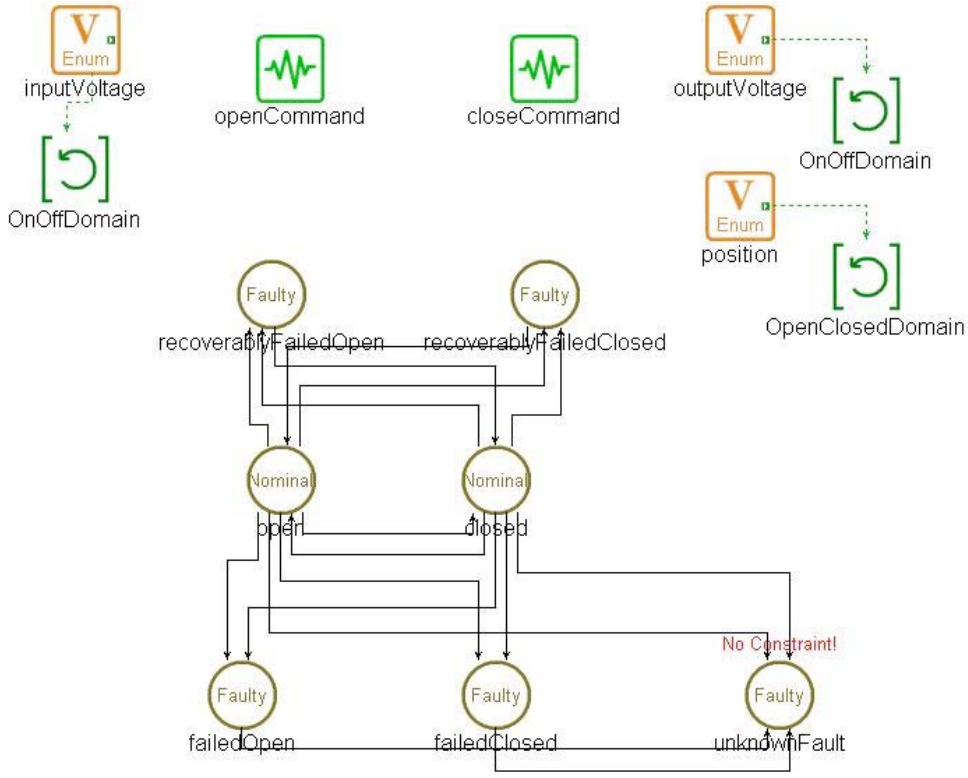


**Figure 4: HyDE model of ADAPT relay component**

In Figure 4 above, there are several parts to the relay component model. First, there are three variables, *inputVoltage*, *outputVoltage*, and *position*. Each of these three is a discrete variable of an enumerated type, and the reference to the domain of the variables is indicated by the green arrow. While these variables are discrete, HyDE variables are allowed to have boolean, discrete, continuous, or interval domains. There are seven modes in the relay model, indicated by the brown circles. The black arrows between the modes are the allowed transitions between the modes. There are two commands, *openCommand* and *closeCommand*, which are used as guards on the mode transitions.

Each of the modes contains constraints on the variables which define the behavior of the component in that mode. When the relay is estimated to be in a mode, HyDE will activate that mode's constraints on the variables. As the text inside the brown circle indicates, some modes are considered nominal and some are considered faulty. The faulty modes are distinguished by the transitions leading to them – if a guardless transition leads into a mode, that mode is defined as faulty. The faulty modes contain the same constraints as the equivalent nominal mode, but are different in the transitions to and

from that mode. The reason for this is that when the relay is open, it has the same behavior regardless if it opened in response to a command or due to a failure. However, the *open* mode will transition to *closed* if a *closeCommand* is received, but the *failedOpen* mode will not transition if a *closeCommand* is received. The *recoverablyFailedOpen* mode represents a case where sending a command again allows the system to recover back to nominal operations, and *unknownFault* is a special catch-all fault mode with no constraints on the variables.

## 4. Testing Procedures

### 4.1. Test Plans

The HyDE parameters used for testing are defined in Appendix A. The test plans for both the original HyDE testing and the expanded testing are attached in Appendix B and Appendix C. These define how the tests are to be done and which ADAPT faults will possibly be used in the testing.

Both test plans define sets of fault types and lists of particular components which could exhibit that fault. The current testing included a scenario from each of the fault types with the actual component chosen randomly from the component list in the test plan.

### 4.2. Figures of Merit

The two figures of merit for which HyDE was tested are defined below:

- **Correctness of fault isolation** – HyDE's output is correct if at least one candidate returned is the injected fault, and all candidates returned are possible explanations of the fault signature.

- **Time to fault isolation** – The time elapsed from when the fault was injected to when the diagnostic system isolates the fault.

Some ADAPT tests define a third metric, the time to detection. HyDE currently doesn't provide a separate fault detection notification, so this metric is not given in the results. It is equal to the time to fault isolation.

## 5. Test Results

The locations of the injected faults which were randomly chosen according to the test plan are shown in Figure 5.

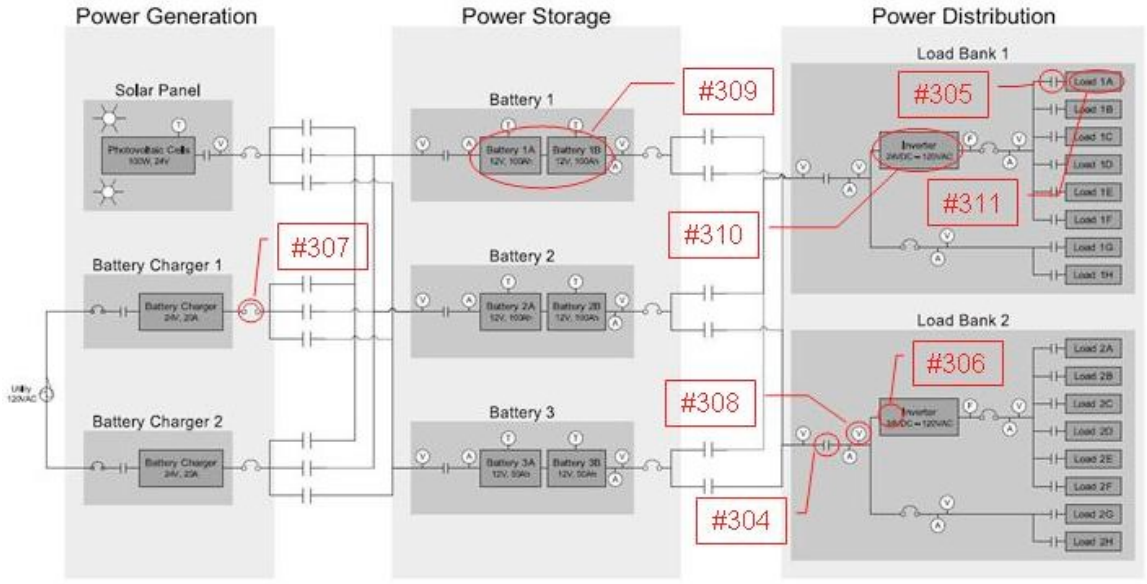


**Figure 5: Locations of injected faults, with experiment number**

The results of the HyDE tests are given below in Table 2:

| ADAPT Experiment ID | Injected Fault and Location | HyDE Diagnosis | Correct Fault Isolation? | Time (s) |
|---|---|---|---|---|
| 317 | None | No faults reported | yes | - |
| 304 | Relay Failed Open, EY260 | EY260.recoverablyFailedOpen<br>EY260.failedOpen<br>EY260.recoverablyFailedOpen ESH260A.unknownFault | yes | 35.7 |
| 305 | Relay Feedback Sensor Failed, ESH175 | EY175.failedOpen<br>EY175.recoverablyFailedOpen<br>ESH175.unknownFault | yes | 8.4 |
| 306 | Circuit Breaker Tripped, cbISH262 | cbISH262.tripped<br>cbISH262.failedOpen<br>Inv2.unknownFault ISH262.unknownFault | yes | 89.5 |
| 307 | Circuit Breaker Feedback Sensor Failed, ISH210 | ISH210.unknownFault<br>cbISH210.tripped<br>cbISH210.failedOpen | yes | 11.9 |
| 308 | Voltage Sensor Failed, E261 | E261.unknownFault<br>E261.unknownFault<br>E261.unknownFault | yes | 16.6 |
| 309 | Battery Low | BatteryA.lowVoltage | yes | 93.1 |

| | | | | |
|---|---|---|---|---|
| | Voltage, BATT1 | EY136.failedOpen<br>ISH136.unknownFault<br>E135.unknownFault | | |
| | | EY136.tripped<br>ISH136.unknownFault<br>E135.unknownFault | | |
| 310 | Inverter Failed Off, INV1 | Inv1.unknownFault | yes | 12.2 |
| | | Inv1.unknownFault | | |
| | | Inv1.unknownFault | | |
| 311 | Load Sensor Failed, LT500 | LT500.unknownFault | yes | 11.6 |
| | | LT500.unknownFault | | |
| | | LT500.unknownFault | | |

**Table 2: Results of HyDE diagnostic testing on ADAPT**

In all eight tests, HyDE's fault isolation was correct. The average time for fault isolation was 34.9 seconds, but the time to isolation for the scenarios varied widely. These tests were performed on a Intel 2.8GHz Xeon processor running Windows XP. HyDE was compiled with Microsoft Visual Studio 2005 Express in release mode, which includes compiler optimization.

Based on the large variation in the initial timing results, a second round of testing was done on the scenarios which took HyDE more than 20 seconds to isolate the fault. These were ADAPT experiments #304, #306, and #309. The "minimum candidate probability" parameter of HyDE was changed in order to speed up the search. This parameter is defined and discussed in section 6.4. The results from these second tests are shown below in Table 3:

| ADAPT Experiment ID | Injected Fault and Location | HyDE Diagnosis | Correct Fault Isolation? | Time (s) | Time Decrease |
|---|---|---|---|---|---|
| 314 | Battery Low Voltage, BATT1 | BatteryA.lowVoltage | yes | 19.4 | 79.2% |
| 315 | Circuit Breaker Tripped, cbISH262 | cbISH262.tripped | yes | 38.8 | 56.6% |
| | | cbISH262.failedOpen | | | |
| 316 | Relay Failed Open, EY260 | EY260.recoverablyFailedOpen | yes | 17.5 | 51.0% |
| | | EY260.failedOpen | | | |

**Table 3: Results of 2nd tests with changed HyDE parameters**

With the parameter change, the time to fault isolation was greatly reduced in all retested scenarios. This speed increase is due to fewer candidates being returned in the HyDE diagnosis. HyDE still isolates the fault correctly in all scenarios. Finally, there still is a large variance in the time to fault isolation between the scenarios. These issues are discussed in the next section.

## 6. Discussion and Diagnostic Trade-offs

### 6.1. Expanded Scope vs. Timing

The current testing of HyDE on ADAPT reported longer times to fault isolation than the previous testing. The reason is that several components in the expanded diagnostic scope have much longer settling times than the other components. The AC inverters have a six-second delay from when power is applied to when they will produce output. Several of the loads have settling times as well: the fans take 4.5 seconds after a transition to achieve their steady-state value, and the pumps take three seconds.

The HyDE diagnostic system currently deployed on ADAPT waits for the system's settling time to pass before HyDE begins processing the sensor data. This is necessary because HyDE looks at data from all sensors in creating a diagnosis‘ and the current model doesn't include the transient periods of the system. It handles these by waiting for a specified timeout after a command is sent before attempting a diagnosis. The timeout can be specified for each command, and the worst-case timeout is 11.1 seconds (slightly larger than the sum of the inverter and fan timeouts listed above).

Therefore, the actual computation time of HyDE is less than that reported in the time-to-isolation metric. In general, any system which waits for a mode transition to complete and the transients to settle out will have this system latency incorporated in the diagnostic latency. Systems which are able to model the transient regions could potentially isolate faults more quickly.

### 6.2. Trade-offs with level of modeling abstraction

Different knowledge-based diagnostic systems use different levels of abstraction to capture the knowledge of a target system. High levels of abstraction include the diagnostic logic based on constant threshold values on data from a single sensor; detailed levels of modeling abstraction can include full-system simulations incorporating fast system dynamics and transient effects.

#### 6.2.1. Capability

In general, diagnostic systems with a very detailed level of modeling abstraction will have greater capability to detect and isolate faults, and be able to do that with fewer

sensors. As noted in section 6.1, the current system waits for transient periods to end; if transient periods were modeled, the time to isolate faults would be reduced.

#### 6.2.2. Time to isolation

The capability that comes with a detailed level of abstraction comes at a cost. Often the cost manifests as an increase in the computational requirements for diagnosis, which for a fixed amount of processing power will result in a longer time to isolation. This happens particularly if the model explicitly includes fast transients that must be characterized or simulated.

#### 6.2.3. Time/expense of implementation

Another cost of detailed system modeling is the time required to implement the model. Usually the implementation effort required goes up as the system modeling gets more detailed.

#### 6.2.4. Robustness to system changes

Finally, the more detailed the system model, the less robust the model is to system changes. These system changes could be the result of design changes, recalibrations, or degradation of the system over time. Often, detailed models require more complex characterization of the target system. In these cases, when the system behavior changes, the characterization must be performed again.

### 6.3. HyDE diagnostic system considerations

#### 6.3.1. Flexibility of modeling

An advantage of HyDE is its flexibility in system modeling. A HyDE model may be discrete, continuous, or a combination of the two. This gives the modeler the capability to model the system at various levels of abstraction even within the same model. If the system is well-characterized, then more detailed modeling may be done; if not, then diagnostic information may still be included at a more abstract level.

#### 6.3.2. Requires initial system mode

One disadvantage of HyDE is that it currently requires the initial system mode to begin tracking the system. This constraint arises from HyDE's ability to track hidden system modes. In general, the complete state of a system cannot be determined from the

sensor data. This is especially true in aerospace systems, which tend to make minimal use of sensors to reduce weight. This is a disadvantage if the system is in an unknown state and the software must be restarted.

To mitigate this issue the HyDE development team is investigating allowing multiple candidates even in the initial state. This would likely include a no-fault assumption, but even with that assumption there could be uncertainty in the initial state of the system, depending on the number of hidden states. However, once HyDE had determined an initial set of candidates, the operation of HyDE would continue as usual.

### 6.4. HyDE Parameters Affecting Results

HyDE has several parameters that allow a user to customize the diagnostic search strategy. In this testing, two parameters had a significant impact on the length of time required for HyDE to return a diagnosis: the history length and the minimum candidate probability.

The history length is the amount of time for which HyDE will store telemetry from the system. Reasoning across time allows HyDE the capability to diagnose faults that manifest slowly over time, and also allows HyDE to revise its diagnosis based on later sensor information. To prevent HyDE's memory usage from growing unboundedly with time, HyDE deletes any telemetry older than the history length. In general, a larger history length increases the time to isolation, because it increases the size of the search space in which HyDE is looking for faults. The user thus must set the history length parameter to be large enough to capture a target system's slow-manifesting faults, but small enough so that the time to fault isolation is still acceptable.

The minimum candidate probability may also affect the size of the fault search space, and thus the time to fault isolation. HyDE will continue to search for candidates until the probability of a candidate is smaller than this number (or some other parameter is met and terminates the search). If this parameter is very small, HyDE may spend an increasing amount of time searching for less likely candidates. This is seen in the test results: in experiments #304, #306, and #309, the time to fault isolation is much larger than the other scenarios because HyDE is searching for the very unlikely double- and triple-fault candidates. When those scenarios were retested with the minimum candidate probability parameter raised, the time to fault isolation is reduced to be comparable to the other scenarios. Here the design trade-off is as follows: the minimum candidate probability must be small enough to capture all fault modes of interest, but large enough so that the time to fault isolation is still acceptable.

### 6.5. Future possible diagnostic expansions

There are several possible future expansions or changes to the current HyDE diagnostic system as deployed on ADAPT, discussed in sections below.

### 6.5.1. Data acquisition system (DAQ) faults

The data acquisition system of ADAPT is not currently modeled. This would be a useful model addition in several ways. First, while individual sensors are currently modeled and sensor faults can be diagnosed, all sensor faults are assumed to be independent. If a data acquisition module were to fail, HyDE would currently require very long computation time to find an extremely unlikely candidate in which all affected sensors failed. Secondly, this would demonstrate HyDE's ability to reason over multiple off-nominal sensor readings to isolate faults. Finally, the addition of data acquisition modules should be straightforward to add to the HyDE model.

### 6.5.2. Real-valued model

Converting the current discrete model to a real-valued model would also be interesting future work. Some ADAPT components which are more continuous by nature are left out of the current discrete model, and they could be more readily modeled with these model types. Also, a major strength of HyDE is its ability to model hybrid systems, and the current discrete model is not making use of that capability. Therefore the full capability of HyDE was not being utilized in the current model. Finally, a real-valued model could incorporate models of the transient periods of components, which could obviate the need for lengthy timeout periods due to settling time. Converting to a real-valued model would be the most logical next step for the HyDE model.

## 6.6. HyDE Recommendations

The results of this test lead to two specific suggestions for improving HyDE. First is to allow a separate notification for fault detection. As implemented in HyDE, the fault detection step is relatively fast, followed by a much more intensive fault isolation. Having a fast notification after the fault detection step that a fault has occurred would be valuable information to users of the target system.

The second suggestion is for HyDE to return candidates individually, rather than all at once. This could let users of a HyDE diagnostic system begin to evaluate the HyDE output while the search is proceeding for more candidates. Another refinement of this suggestion is to allow more interactive control of HyDE's fault isolation search: as candidates are found by HyDE, the user could look at those candidates and decide if HyDE should continue the search.

# 7. Conclusion

The work completed in FY2007 resulted in a HyDE model with the capability to diagnose five times the number of ADAPT components than was previously tested. The expanded diagnosis model passed a corresponding set of new ADAPT fault injection scenario tests with no incorrect faults reported. The average time from fault injection to receipt of the diagnosis was 34.9 seconds. The time required for the HyDE diagnostic system to isolate the fault varied widely between tests. This was partially alleviated by tuning HyDE's minimum candidate probability parameter to limit HyDE from spending excessive CPU time looking for unlikely candidates. The work did illustrate several suggestions for HyDE: allowing a separate notice for fault detection, and allowing a user to control the candidate search more directly, such as returning candidates individually rather then all at once.

## 8. Acknowledgments

Thanks to the HyDE development team, Sriram Narasimhan, Lee Brownston, and David Hall for their support in this activity. Also, thanks to the ADAPT team, and particularly to Justin Yu for his help in conducting the test runs.

## 9. References

[1] S. Poll, A. Patterson-Hine, J. Camisa, D. Garcia, D. Hall, C. Lee, O. Mengshoel, C. Neukom, D. Nishikawa, J. Ossenfort, A. Sweet, S. Yentus, I. Roychoudhury, M. Daigle, G. Biswas, and X. Koutsoukos. Advanced Diagnostics and Prognostics Testbed. 18th International Workshop on Principles of Diagnosis, pp. 178-185, May 2007.

[2] S. Narasimhan, L. Brownston. HyDE - A General Framework for Stochastic and Hybrid Model-based Diagnosis. 18th International Workshop on Principles of Diagnosis, pp. 162-169, May 2007.

# Appendix A. HyDE parameters used for testing

The testing was done with the following values for the HyDE parameters:

```
CommandsToBackTrackAcross UNBOUNDED
FilterType constraints
HistoryTime 5
SystemLocationChangesToBackTrackAcross UNBOUNDED
MaximumCandidateCount 3
MaximumCandidateSize UNBOUNDED
MaximumCandidatesToTry UNBOUNDED
MinimumCandidateProbability 0
NoiseModel Percentage
PreferNewerCandidates true
maximumcandidategenerationtime UNBOUNDED
UseDependencyGraphs false
```

As described in the main body, for the second re-tested scenarios the only changed parameter was

```
MinimumCandidateProbability 0.05
```

# Appendix B. Original HyDE test plan

## Advanced Diagnostics and Prognostics Testbed (ADAPT)
## HyDE Test Plan and Acceptance Criteria

1.0 Purpose

HyDE will be tested in the ADAPT system to determine its maturity as an ISHM reasoning system, and to determine its suitability to be used as a proven test article in future possible ADAPT work. This future work may involve demonstrations of the ADAPT system, and it may involve collaborations of ADAPT with other areas in higher-level work such as mission operations utilizing ISHM techniques.

2.0 Communications testing

This is a test of HyDE's ability to interface to the ADAPT software systems. HyDE must implement the test article interface defined in the "ADAPT Software Interface and Requirements Document." The test consists of three parts listed below.

2.1 Sensor data message test

To test HyDE's ability to receive sensor data messages, the ADAPT software system (or an ADAPT-provided test application) will send sensor data messages as specified in the Interface and Requirements document. The data will be sent at 5Hz (the same rate as expected in ADAPT operations), for a period of 5 minutes. To pass the test, HyDE must show the following:

1) Show that HyDE received at least 99% of the sensor messages (maximum message loss rate of 1 in 100)
2) The sensor messages received must contain the same data as the messages sent by ADAPT.

2.2 Command message test

To test HyDE's ability to receive command messages, the ADAPT software system (or an ADAPT-provided test application) will send command messages as specified in the Interface and Requirements document. Commands will be sent at random time intervals for a period of 5 minutes, but not more quickly than once every 5 seconds. To pass the test, HyDE must show the following:

1) Show that HyDE received at least 99.9% of the command messages (maximum message loss rate of 1 in 1000)
2) The command messages received must contain the same data as the messages sent by ADAPT.

2.3 Sending diagnosis messages

To test HyDE's ability to send diagnosis messages, HyDE (or a HyDE-provided test application) must send diagnosis messages to the ADAPT software

system. A total of 100 diagnosis messages must be sent within a 5-minute time period. To pass the test, HyDE must show the following:

1) Show that ADAPT received at least 99% of the diagnosis messages (maximum message loss rate of 1 in 100)
2) The diagnosis messages received by ADAPT must be in the format specified in the Interface and Requirements document.
3) The diagnosis messages received by ADAPT must contain the same data as the messages sent by HyDE.

## 3.0 Scenario testing

The scenarios will test the diagnosis capability of HyDE. Each scenario will begin with the ADAPT system in the same default configuration, and will be commanded by a human user(s) to execute the scenarios. Faults will be injected in many of the scenarios, and HyDE must detect and isolate the fault to the component level (within a set of candidates), and must do so in a limited period of time.

### 3.1 Nominal operation

No faults will be injected into this scenario test. This test may be run multiple times, and individual runs will be passed if HyDE does not report any faults through the entire run. The results of individual runs will be recorded, and when HyDE has achieved 90% pass rate over all of the runs, this test will be passed. Only runs performed after the last change made to the HyDE diagnostic system will be used in determining the success.

The scenario will be:
System begins in default "off" configuration
System is commanded to drive any one of the loads
Load is left on for a minimum of 1 min.
System is returned to default "off" configuration

### 3.2 Relay Failed Open

The fault injected into this test will be chosen by the ADAPT personnel running the test from one of the following relays on the ADAPT system: EY141, EY144, EY241, EY244, EY341, EY344, EY160, or EY260. The relay chosen should result in a change in the system state at the time of the fault injection: that is, the fault should not be injected on a relay which is already open.

This test may be run multiple times, and individual runs will be passed if HyDE shows the following:

1) The injected fault is included as at least one of the candidates.
2) The diagnosis is reported within a time window of 10 seconds from the time of fault injection.
3) No fault diagnosis is reported outside of the time window.

The results of individual runs will be recorded, and when HyDE has achieved a 90% pass rate over all of the runs, this test will be passed. Only runs performed after the last change made to the HyDE diagnostic system will be used in determining the success.

The scenario will be:

System begins in default configuration

System is commanded to drive any one of the loads

Load is left on for a minimum of 1 min. At any point in that time, the fault will be injected to the system.

After the fault injection, the system will remain in that state for at least the length of the time window allowed to HyDE for fault detection and isolation

If necessary, system commanded to restore power to the load

System is returned to default "off" configuration

3.3 Voltage Sensor Failed

The fault injected into this test will be chosen by the ADAPT personnel running the test from one of the following voltage sensors on the ADAPT system: EI135, EI140, EI142, EI161, EI235, EI240, EI242, EI261, EI335, or EI340. The voltage sensor fault injected should result in a marked change of the sensor value, at least 15 volts from the previous value.

This test may be run multiple times, and individual runs will be passed if HyDE shows the following:

4) The injected fault is included as at least one of the candidates.
5) The diagnosis is reported within a time window of 10 seconds from the time of fault injection.
6) No fault diagnosis is reported outside of the time window.

The results of individual runs will be recorded, and when HyDE has achieved a 90% pass rate over all of the runs, this test will be passed. Only runs performed after the last change made to the HyDE diagnostic system will be used in determining the success.

The scenario will be:

System begins in default configuration

System is commanded to drive any one of the loads

Load is left on for a minimum of 1 min. At any point in that time, the fault will be injected to the system.

After the fault injection, the system will remain in that state for at least the length of the time window allowed to HyDE for fault detection and isolation

If necessary, system commanded to restore power to the load

System is returned to default "off" configuration

3.4 Circuit Breaker Tripped

The fault injected into this test will be chosen by the ADAPT personnel running the test from one of the following circuit breakers on the ADAPT system: EY136, EY236, or EY336. The circuit breaker should not be in the tripped position before the fault is injected. This test may be run multiple times, and individual runs will be passed if HyDE shows the following:

7) The injected fault is included as at least one of the candidates.
8) The diagnosis is reported within a time window of 10 seconds from the time of fault injection.
9) No fault diagnosis is reported outside of the time window.

The results of individual runs will be recorded, and when HyDE has achieved a 90% pass rate over all of the runs, this test will be passed. Only runs performed after the last change made to the HyDE diagnostic system will be used in determining the success.

The scenario will be:

System begins in default configuration

System is commanded to drive any one of the loads

Load is left on for a minimum of 1 min. At any point in that time, the fault will be injected to the system.

After the fault injection, the system will remain in that state for at least the length of the time window allowed to HyDE for fault detection and isolation

If necessary, system commanded to restore power to the load

System is returned to default "off" configuration

3.5 Battery Overheating

The fault injected into this test will be chosen by the ADAPT personnel running the test from one of the following batteries on the ADAPT system:BATT1A, BATT1B, BATT2A, BATT2B, BATT3A, or BATT3B.

This test may be run multiple times, and individual runs will be passed if HyDE shows the following:

10) The injected fault is included as at least one of the candidates.
11) The diagnosis is reported within a time window of 1 min. from the time of fault injection. Note the longer time period allowed here as a result of the slower thermal characteristics of this fault.
12) No fault diagnosis is reported outside of the time window.

The results of individual runs will be recorded, and when HyDE has achieved a 90% pass rate over all of the runs, this test will be passed. Only runs performed after the last change made to the HyDE diagnostic system will be used in determining the success.

The scenario will be:

System begins in default configuration

System is commanded to drive any one of the loads

Load is left on for a minimum of 1 min. At any point in that time, the fault will be injected to the system.

After the fault injection, the system will remain in that state for at least the length of the time window allowed to HyDE for fault detection and isolation

If necessary, system commanded to restore power to the load

System is returned to default "off" configuration

4.0 Acceptance Criteria

To be accepted, HyDE must pass all 3 of the communication tests, and must pass 3 of the 5 scenario tests.

# Appendix C. Expanded HyDE test plan

## Advanced Diagnostics and Prognostics Testbed (ADAPT)
## HyDE Expanded Test Plan

1.0 Purpose

HyDE successfully passed the success criteria outlined in the first HyDE Test Plan. This plan for expanded testing was created to test the expansions of the diagnostic scope of the HyDE deployment on ADAPT.

2.0 Scenario testing

The scenarios will test the diagnosis capability of HyDE. Each scenario will begin with the ADAPT system in the same default configuration, and will be commanded by a human user(s) to execute the scenarios. Faults will be injected in many of the scenarios, and HyDE must detect and isolate the fault to the component level (within a set of candidates), and must do so in a limited period of time.

2.1 Scenario procedures

2.1.1 Fault manifests immediately

The injected fault must begin to have an effect immediately in the system. For example, a relay may not be failed open if it is already open. This is required to easily score the timing figures of merit.

2.1.2 Scenario chronology

1) System begins in default configuration of all relays open, all circuit breakers closed
2) System is commanded on to charge or power loads, depending on fault location
3) Fault injected into system at any point up to 1 min after reaching desired ADAPT configuration
4) After the fault injection, the system will remain in that state for enough time for all faults to become observable, and a maximum of 1 min thereafter.

2.1.3 Scenario success criteria

1) The injected fault is included as at least one of the candidates reported in the HyDE diagnosis
2) No fault diagnosis is reported before a fault is injected

The results of individual runs will be recorded, and when HyDE has achieved a 90% pass rate over all of the runs, this test will be passed. Only runs performed after the last change made to the HyDE diagnostic system will be used in determining the success.

2.2 Nominal operation

No faults will be injected into this scenario test.

2.3 Relay Failed Open

The fault injected into this test will be chosen at random by the ADAPT personnel running the test from one of the following relays on the ADAPT system:

Charging: EY115, EY116, EY117, EY215, EY216, EY217, EY315, EY316, EY317, EY126, EY226, EY326

Distribution: EY141, EY144, EY241, EY244, EY341, EY344, EY160, EY260

Load: EY170, EY171, EY172, EY173, EY174, EY175, EY183, EY184, EY270, EY271, EY272, EY273, EY274, EY275, EY283, EY284

2.4 Relay Position Sensor Failed

The fault injected into this test will be chosen at random by the ADAPT personnel running the test from one of the following relay sensors on the ADAPT system:

Charging: ESH115, ESH116, ESH117, ESH215, ESH216, ESH217, ESH315, ESH316, ESH317, ESH126, ESH226, ESH326

Distribution: ESH141, ESH144, ESH241, ESH244, ESH341, ESH344, ESH160, ESH260

Load: ESH170, ESH171, ESH172, ESH173, ESH174, ESH175, ESH183, ESH184, ESH170, ESH271, ESH272, ESH273, ESH274, ESH275, ESH283, ESH284

2.5 Voltage Sensor Failed

The fault injected into this test will be chosen at random by the ADAPT personnel running the test from one of the following voltage sensors on the ADAPT system:

Charging: E125, E135, E225, E235, E325, E335

Distribution: E140, E142, E161, E240, E242, E261, E340, E342

Loads: E165, E167, E265, E267

2.6 Circuit Breaker Tripped

The fault injected into this test will be chosen at random by the ADAPT personnel running the test from one of the following circuit breakers on the ADAPT system:

Charging: cbISH110, cbISH210, cbISH310

Distribution: EY136, EY236, or EY336

Loads: cbISH162, cbISH166, cbISH262, cbISH266

2.7 Circuit Breaker Position Sensor Failed

The fault injected into this test will be chosen at random by the ADAPT personnel running the test from one of the following circuit breakers on the ADAPT system:

Charging: ISH110, ISH210, ISH310

Distribution: ISH136, ISH236, or ISH336

Loads: ISH162, ISH166, ISH262, ISH266

2.8 Battery Low Voltage

The fault injected into this test will be chosen at random by the ADAPT personnel running the test from one of the following batteries on the ADAPT system:

BATT1, BATT2, BATT3

2.9 Inverter Failed Off

The fault injected into this test will be chosen at random by the ADAPT personnel running the test from one of the following inverters on the ADAPT system:

INV1, INV2

2.10 Load Sensor Failed

The fault injected into this test will be chosen at random by the ADAPT personnel running the test from one of the following load sensors on the ADAPT system:

LT500, LT505, ST515, FT525, FT520, ST516

# REPORT DOCUMENTATION PAGE

*Form Approved*
*OMB No. 0704-0188*



| 1. REPORT DATE *(DD-MM-YYYY)* | 2. REPORT TYPE | 3. DATES COVERED *(From - To)* |
|---|---|---|
| 18-01-2008 | NASA STI Technical Memorandum | Oct 2006 - Sept 2007 |

**4. TITLE AND SUBTITLE**

Testing HyDE on ADAPT

**5a. CONTRACT NUMBER**

**5b. GRANT NUMBER**

**5c. PROGRAM ELEMENT NUMBER**

ARMD/AVSP/IVHM 645846.02.07.01.01

**6. AUTHOR(S)**

Adam Sweet

**5d. PROJECT NUMBER**

**5e. TASK NUMBER**

**5f. WORK UNIT NUMBER**

**7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)**

NASA Ames Research Center,
Intelligent Systems Division
Moffett Field, CA 94035-1000

**8. PERFORMING ORGANIZATION REPORT NUMBER**

NASA/TM-2008-214570

**9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)**

National Aeronautics and Space Administration
Washington, DC 20546-0001

**10. SPONSORING/MONITOR'S ACRONYM(S)**

NASA

**11. SPONSORING/MONITORING REPORT NUMBER**

NASA/TM-2008-214570

**12. DISTRIBUTION/AVAILABILITY STATEMENT**

Unclassified -- Unlimited
Subject Category: 18 Distribution: Standard
Availability: NASA CASI (301) 621-0390

**13. SUPPLEMENTARY NOTES**

**14. ABSTRACT**

The IVHM Project in the Aviation Safety Program has funded research in electrical power system (EPS) health management. This problem domain contains both discrete and continuous behavior, and thus is directly relevant for the hybrid diagnostic tool HyDE. In FY2007 work was performed to expand the HyDE diagnosis model of the ADAPT system. The work completed resulted in a HyDE model with the capability to diagnose five times the number of ADAPT components previously tested. The expanded diagnosis model passed a corresponding set of new ADAPT fault injection scenario tests with no incorrect faults reported. The time required for the HyDE diagnostic system to isolate the fault varied widely between tests; this variance was reduced by tuning HyDE input parameters. These results and other diagnostic design trade-offs are discussed. Finally, possible future improvements for both the HyDE diagnostic model and HyDE itself are presented.

**15. SUBJECT TERMS**

hybrid diagnosis, model-based reasoning, electrical power system health management, integrated vehicle health management

| 16. SECURITY CLASSIFICATION OF: | | | 17. LIMITATION OF ABSTRACT | 18. NUMBER OF PAGES | 19a. NAME OF RESPONSIBLE PERSON |
|---|---|---|---|---|---|
| a. REPORT | b. ABSTRACT | c. THIS PAGE | | | 19b. TELEPHONE NUMBER *(Include area code)* |
| U | U | U | UU | 27 | |