SAND2006-5450C

Real Time Consequence Engine[1]

Daniel Briand, Sandia National Laboratories

James E. Campbell, PhD., Sandia National Laboratories

Key Words: consequence analysis, health management, prognostics, PHM

SUMMARY & CONCLUSIONS

This paper presents a brief overview of a PHM system, where consequence analysis performed by the Real Time Consequence Engine (RTCE) fits into a PHM system, and the basic capabilities and architecture of the RTCE. Also, an example of consequence analysis results as applied to a past prototype development effort and the results of ongoing modification efforts to RTCE are presented. By varying operational settings, maintenance schedules, and other parameters of interest, the RTCE can be used to examine the consequences of such actions in terms of common performance metrics; such as, mean time between failures (MTBF), mean time to repair (MTTR), system availability, maintenance cost, downtime cost, etc. In its final form, the RTCE will provide the capability to evaluate the potential cost / benefit of an embedded PHM system, considered to be a precursor to implementing a PHM system; and, provide the capability for real-time consequence analysis, typically viewed as the final step in the development of a complete PHM system.

1 INTRODUCTION

Prognostic and Health Management (PHM) systems are being developed at an increased rate due to potential cost savings and because the DoD is mandating a prognostic capability be implemented in all new major weapon systems. A capable PHM system should be able to predict component failure far enough in advance to allow the modification of operations and maintenance schedules in order to minimize downtime, while obtaining the maximum useful life of a component. The focus of PHM implementation has been in the analysis of the sensor and inspection data necessary to identify impending failure modes and in the development of the data fusion and trend detection algorithms to predict impending failures. However, relatively little work has been accomplished on consequence analysis. Consequence analysis is defined here as the capability to take prognostic information and apply it to the operations and maintenance schedules in an optimal manner so as to maximize the availability of a system while minimizing maintenance and supply costs.

This paper will describe the capabilities of the Real-Time Consequence Engine (RTCE) being developed by Sandia National Laboratories. The purpose of the RTCE is to conduct consequence analysis based on the new information made available by updated failure mode end-of-life predictions. More specifically, the RTCE will predict in real-time the effect that operational strategies, as well as repair / replace / inspect / wait strategies, will have on the modeled system, if that action is taken at the time changes in the reliability of the system are detected. The following paragraphs will first provide a general description of a PHM system. Second, an in-depth description of RTCE will be presented. Third, a past application of the prototype RTCE is presented which evaluates a fixed-wing aircraft's accessory drive gearbox (ADG) time change interval. In addition, a current application, which examines briefly the improvements made to the RTCE in support of an ongoing implementation of a PHM system on the Air Force Airborne Laser (ABL), is reviewed. Finally, a summary with future work is presented.

2 PHM SYSTEM

A PHM system can be broken down into three general areas: data extraction and sensor feature characterization (data analysis), data fusion and system health analysis (evidence analysis), and operational / maintenance impact analysis (consequence analysis) based on updated time-to-failure (TTF) or remaining useful life predictions as shown in Figure 1. Data extraction may involve the collection of sensor data such as vibration, temperature, or load and/or inspection data; such as oil sample analysis, fluid levels, and visual damage. Since the amount of data coming from the sensors may be considerable, sensor feature extraction techniques may be used such as statistical moments, vibration signatures, etc. Sensor feature interpretation techniques, to include determination of failure modes, are being actively researched and may include methods such as sequential probability ratio test (SPRT), multivariate state estimation technique (MSET), neural networks (NN), self organizing maps, etc.

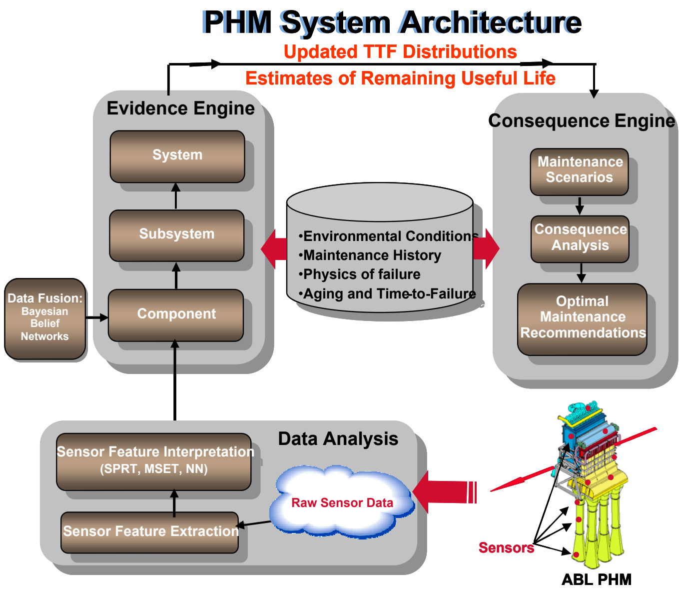

Figure 1. PHM System

Evidence analysis incorporates data fusion techniques such as Bayesian Belief Networks, case-based and model-based reasoning, etc., for each component to provide an overall component level health estimate. These estimates are then rolled up to the next level to provide a subsystem health estimate. Then, subsystem health estimates are rolled up to the system level to provide some form of system health. Finally, updated failure mode time-to-failure distributions and/or remaining useful life predictions are passed on for consequence analysis, the focus of this paper.

3 REAL TIME CONSEQUENCE ENGINE

The purpose of consequence analysis, using a tool such as RTCE, is to look forward in time to help analyze the consequences of various operational and maintenance strategies once a change in a component's time-to-failure or remaining useful life has been detected. The RTCE simulation takes the system health predictions from evidence analysis and develops projections into the future, i.e., what will be the overall impact on the system if a certain impending failure mitigation action is taken. Performance metrics such as mean time between failures (MTBF), mean time to repair (MTTR), availability, maintenance cost, downtime cost, etc., are calculated. Thus, by running the simulation with different operational and maintenance schedules, the consequences of alternative equipment use and maintenance scenarios can be examined.

Possible actions as a result of the consequence analysis might be to modify current operating parameters, shut down immediately and repair the problem, ignore the problem and deal with the failure when it occurs, or schedule maintenance at an appropriate time in the future. For the operator or maintenance personnel to make the best decision, they need to know the consequences of all the possible actions. Consequences might be measured in terms of mission impact, expected downtime, or cost. The intent is to enable this decision making process to occur in real-time, ultimately, as an onboard fully integrated system.

The RTCE can also be used to support PHM cost / benefit assessments such as evaluating the potential effectiveness of a PHM system. PHM systems, no matter how well designed, may either fail to detect a pending problem (false negative) or report a problem when none exists (false positive). False negatives can allow failures to occur that should have been caught. False positives can result in unnecessary (and expensive) maintenance. The end user version of RTCE should help analysts understand the cost / benefit tradeoffs for a PHM system depending on false positive and false negative rates.

3.1 Description

The RTCE, shown in Figure 2, takes as input the updated time-to-failure or remaining useful life predictions, component age and maintenance history, maintenance and operational planned use schedules, and spares availability data. Then, the simulation generates possible maintenance schedules based on predicted failure events which may alter the original maintenance schedules and the operational schedule. Since the failure events are modeled stochastically and maintenance schedules change depending upon predicted failure events, multiple iterations are analyzed and data is collected for calculating the performance measures. An enumeration scheme is then used to develop multiple scenarios in which different aspects of the model, such as repair now or later strategy, repair time duration, combine maintenance event strategies, and spares availability, are varied to provide a collection of all the possible combinations of interest. An optimization scheme is then used to determine the best operational and maintenance strategy based on the user's objectives and the calculated performance measures.

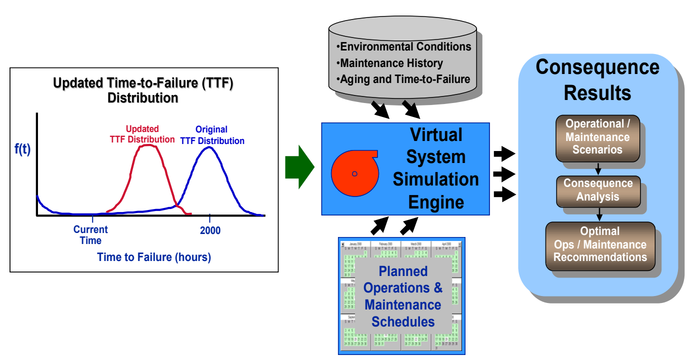

Figure 2. Real Time Consequence Engine

A graphic depicting the current modular architecture of the Virtual System Simulation Engine is shown in Figure 3. Simulation was chosen as the basis for the RTCE because of its flexibility. The simulation mimics the reliability behavior of equipment in terms of simulated equipment failures, repairs, preventive maintenance, and inspections. The simulation is based on user-definable maintenance and inspection schedules and a system reliability model with time-to-failure and time-to-repair distributions for all failure modes. The simulation integrates and drives several modules: schedule, cost, system, and spares.

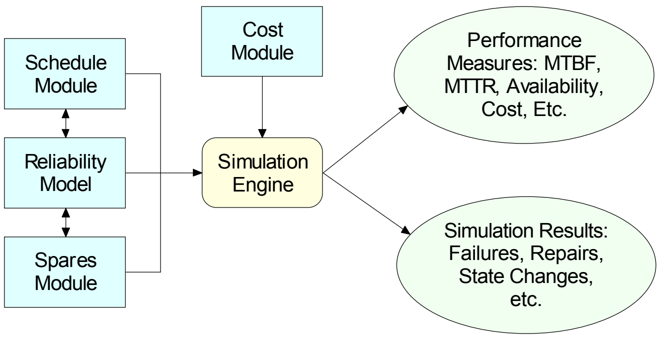

Figure 3. RTCE Simulation Architecture

3.2 Schedule Module

Equipment failures are, by definition, unplanned and are interruptions of the planned scheduled use of the equipment. Simulation of equipment reliability behavior can be viewed as creating a chronology of equipment operational state changes and must begin with the planned use for the equipment. An example of such a chronology is shown in Table 1 where equipment operational states are changing based on the use of the equipment.

Date/Time	**Equipment State**	**Status**
01/01/05 12:00 AM	Other Up Time	Weekend
01/03/05 08:00 AM	Standby	Preparation
01/03/05 10:00 PM	Operational	Normal Ops
01/05/05 01:00 PM	Operational 1	Accelerated Ops
01/06/05 01:00 PM	Operational	Normal Ops
01/07/05 12:00 PM	Preventive Maint.	Weekly Maint.

Table 1. Example Equipment State Chronology

The scheduling module provides the means to specify the planned equipment usage schedule so that component aging, simulated failures, maintenance, etc. can occur in the context of the planned schedule. Setting up an equipment schedule involves the following steps.

1. Identify *special periods*. Special periods are time intervals during which the equipment is scheduled to be in a state other than its default state. Special periods typically model the equipment usage or operational states. Special periods do not include preventive maintenance.
2. Specify *preventive maintenance (PM)* schedules. For each preventive maintenance activity, the failure modes to be addressed are identified and the schedule is specified. Preventive maintenance can be modeled based on calendar time and/or equipment usage time.

3.3 Cost Module

The cost module assumes that the cost of downtime can be characterized by a function that is piecewise constant. The downtime cost function is characterized by a start date, an end date, and the downtime cost per hour which, for example, can represent revenue lost as a result of unplanned equipment failure. In addition to downtime costs, each event (scheduled or unscheduled) can incur additional costs. Each failure mode has an optional cost property including a parts replacement cost. Each scheduled maintenance action includes a cost to perform the maintenance, which is added to the cost to repair any failure modes addressed by the maintenance.

3.4 Reliability Module

The reliability module contains reliability data including time-to-failure and time-to-repair distributions for the failure modes. The model is based on a collection of equipment failure modes. The possibility of redundancy or non-critical system elements is treated through the use of success paths. A success path is a collection of elements (failure modes, components or subsystems) that, if all are operating, determine the operational state of the system. For example, consider the following simple block diagram model in Figure 4.

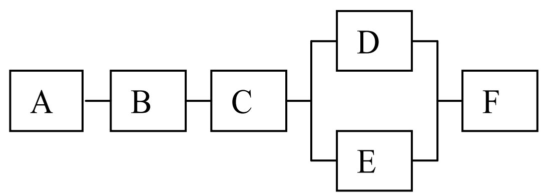

Figure 4. Simple Block Diagram Model

The elements A, B, C and F are in series while D and E are in parallel. If the functionality of the system is unaffected by whether D or E or both are operating, then two success paths would be needed to characterize the system. They are ABCDF and ABCEF, both of which support full functionality. On the other hand, if the functionality of the system in Figure 4 is reduced by the failure of either D or E, then three success paths are needed. Success path ABCDEF supports full functionality while ABCDF and ABCEF support reduced functionality. Of course, series elements do not need to be included in any specific success path since they are, by definition, included in all success paths.

Success paths are defined by a collection of references to failure modes in the reliability model and a reference to the equipment operational state that results from the success path. When a success path is active, the system is in an operational state (note that the user must define these operational states and the groups of failure modes that lead to various operational states). Operational states may differ from the default operational state in terms of the effect on the system being simulated. For example, a helicopter maneuvering at low altitude is subject to more stress than one flying straight and level at higher altitude. During intervals of increased (or decreased) operational stress, selected components and subsystems may effectively age more or less rapidly than normal. RTCE allows alternate equipment operational states that provide a means to treat such effects. Operational states are characterized by a collection of affected failure modes and an "acceleration factor" which causes the failure mode to age more or less rapidly than normal during the alternate operational state.

3.5 Spares Module

The spares module provides the impact of a spares inventory on the system being simulated. If a part fails during the simulation and requires replacement, the spares inventory is queried to see if a spare part that fixes that component exists, and if it does, how long it takes to acquire that spare. The spares inventory is a collection of spare parts each of which has properties such as restock time, withdraw time, purchase cost, storage cost, reorder level, usage rate, etc. A spares model is included since the availability of spares may be a major factor in decision-making when a pending failure is identified.

3.6 RTCE Simulation Engine

The RTCE simulation logic represents a discrete-event simulation which executes an evolving schedule of preventive maintenance actions and failure and repair events. Failure and repair events are generated by drawing from time-to-failure and time-to-repair distributions. A "master clock" looks at each event, calculates the state the system is in given that event, and then finds the next event that will happen. The first step in preparing the simulation is to set up the planned schedule. The simulation event's start and end dates define the period to be simulated while the planned schedule might include preventive maintenance (PM), scheduled shut down periods, etc. The simulation maintains the planned schedule (which can be delayed somewhat during the simulation), the actual schedule (which is the result of events being added during the simulation), and a collection of current events (which are events that started prior to the current date but have not yet ended). Events are allowed to overlap and can be modeled to reflect efficiencies in combining maintenance actions.

At any time during the simulation there are three possibilities for the next system state change: 1) the end of a current event, 2) the occurrence of a failure mode, or 3) the beginning of a planned event. The simulation engine determines which of these occurs next and takes the appropriate action. If the next event is a failure mode or a planned event, the event is added to the actual schedule and to the current events collection. If the end of a current event causes the next state change, that event is removed from the current events collection. Time is then advanced to the next state change and the process continues. At the end of the simulation, statistics on system performance measures such as system MTBF, availability, downtime, and cost are calculated and presented to the user. If desired, the user can view the history of events in the simulation.

4 APPLICATIONS

The following paragraphs describe two applications of the RTCE to help further demonstrate its potential for use in evaluating PHM systems and optimally modifying operations and maintenance schedules.

4.1 Application 1

The first application illustrates the potential effectiveness of a PHM system based on a time-change interval optimization analysis performed on a typical Accessory Drive Gearbox (ADG) for a fixed wing aircraft. The ADG provides a main engine starting capability on the ground or in flight and provides power from the aircraft engine to the accessories such as hydraulic pumps and generators. It was assumed that the ADG had a 5% probability of burn-in failure during the first 100 hours of operation, a 10% probability of random failure occurring between burn-in and onset of end-of life, and the end-of-life characterized by a normal distribution with a mean of 4000 hours and a standard deviation of 500 hours.

The Consequence Engine (an earlier prototype of the RTCE) simulates inspections in significant detail. The simulation looks ahead to the next scheduled inspection and generates a probability that the component will fail before the next inspection. A failure event is predicted for that component based on its projected usage, its time to failure distribution, and a random number draw. If a failure is predicted before the next inspection, there is a probability that the current inspection will not detect this impending failure. This is considered a "false negative" probability, which in some sense constitutes a failure of the PHM system. If a failure will not occur before the next inspection, there is a probability that the current inspection will indicate an impending failure. This is considered a "false positive' probability, meaning that the inspection indicated a pending failure when there was none, and the component was replaced prematurely. In this analysis, the false negative input value was varied while the false positive rate was set at 1%.

The results of this analysis are shown in Figure 5. The probability of a false negative is the X-axis with the larger numbers meaning that the inspection (i.e. prognostics) is not that good at detecting onset of failure. The probability of ADG failure per year, per aircraft is on the Y-axis. The different lines show the results of the analysis, parametrically varying the time-change interval Tc from 2500 hours to 4000 hours.

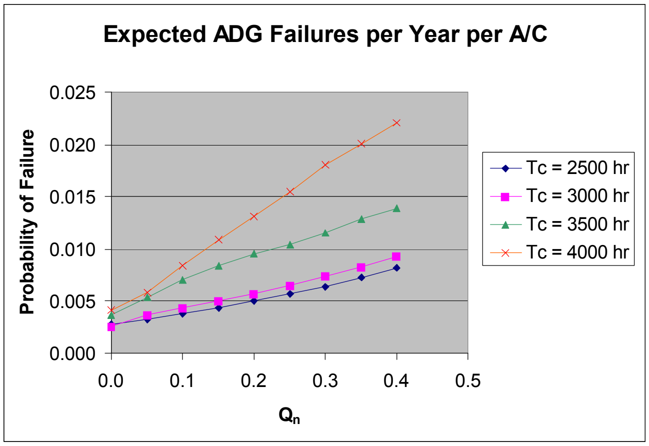

Figure 5. Annual ADG Failures per Year per A/C

The bottom line shows that with the current replacement interval of about 2500 hours, there is approximately a 0.3% chance of failure per aircraft per year, assuming a perfect inspection (this is the baseline failure rate primarily due to infant mortality, or failures that occur before an inspection is scheduled.) Even with a false negative rate of 40%, the probability of ADG failure is still under 1% because the time change is at the 3σ limit on the wear out portion of the assumed time-to-failure distribution (4000-3*500). As the time change interval is lengthened, the probability of ADG failure increases; however, the assumption that many of the failures are caught at the 300 hour inspections means that overall, the ADG failure rate does not increase significantly. In the worst case scenario, with a 40% false negative rate and a time change at 4000 hours, the failure probability is around 2.2% per aircraft per year.

This graph also shows how the RTCE can be used to determine the required precision of a prognostic system to maintain a certain reliability level. If, for example, an annual failure probability for the ADG per aircraft was desired to be 0.5% or less, the accuracy of a prognostic system with a 3500 hour time change interval would require a false negative probability of about 3% or less. If an annual failure probability for the ADG per aircraft was desired to be 1% or less, the accuracy of a prognostic system with a 4000 hour time change interval would require a false negative probability of no more than 13%.

4.2 Application 2

In the second application, the RTCE is part of a complete PHM system that is being prototyped in support of the ABL's chemical-oxygen iodine laser (COIL). ACTA Corporation is developing the evidence engine that will interface with Sandia's RTCE to demonstrate the feasibility and capability for a real time, on-board, complete PHM system [2]. The prototype is being applied to the iodine fluid flow system, which is considered to be problematic. It will support PHM for the ABL by recommending an appropriate change in scheduled operations and/or maintenance actions. Possible actions might be to modify current operating parameters, shut down immediately and repair the problem, ignore the problem and deal with the failure when it occurs, or schedule maintenance at an appropriate time in the future. The intent is to enable this decision making process to occur in real-time, ultimately, as an on-board fully integrated system to model multiple ABL platforms.

As part of this effort, several improvements are being made. The most significant improvement is the new user interface. With multiple screens, the user can input operations and maintenance schedules, update component failure information, etc., much quicker than in the original data input file. In addition, a schedule viewer screen allows the user to visualize the current operations schedule, calendar-based preventive maintenance schedule, the predicted use-based maintenance schedule based on the current operations schedule, and predicted failures based on the operations, calendar-based maintenance and use-based maintenance schedules. Once the RTCE is run, the schedule viewer screen provides the resulting or optimized schedule on the same screen as the original as shown in Figure 6. Now the user can clearly determine what failures are about to occur, what changes in both the calendar-based and use-based maintenance schedules should be considered, and what changes in the operations and maintenance schedules may provide the maximum availability for the minimum cost.

Implementation of the RTCE on the ABL would proceed as follows. Prior to a mission, anticipated operations, scheduled maintenance, and use-based maintenance events would be updated in the RTCE. In addition, time to failure distributions and component age would be updated depending upon parts maintenance and / or replacements that occurred prior to mission start. An evidence engine would monitor the condition of critical components and as failure / wear-out indications are prognosed, pass the updated time-to-failure distributions to the RTCE. The RTCE would use the updated time-to-failure distributions to determine changes in the operations, scheduled maintenance, and use-based maintenance schedules and provide the onboard operator / technician an updated recommended schedule that maximizes the mission's effectiveness. Mission effectiveness may be maximized by implementing such options as: continue operating as planned, operate the laser at a reduced readiness or output, shorten the mission, abort the mission, etc. Monitoring multiple ABL platforms simultaneously with a PHM system that included both an evidence engine and a RTCE would provide the greatest flexibility for operations while maximizing overall mission effectiveness.

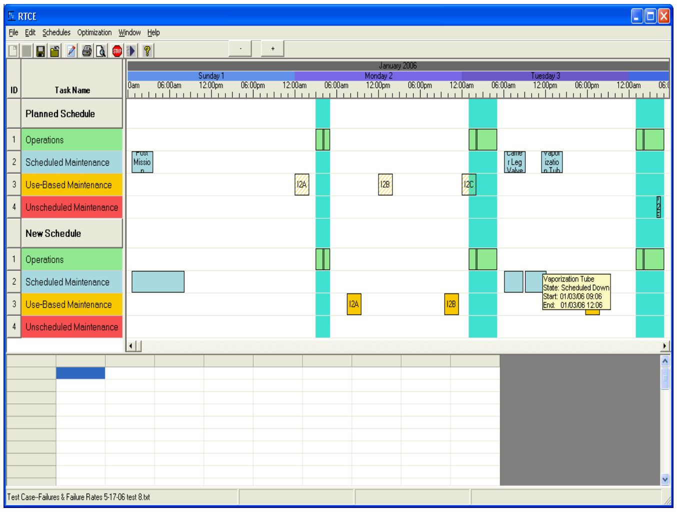

Figure 6. RTCE Schedule Viewer Screen

5 NEXT STEP

In reality, there is an additional step in the process which has yet to be implemented. This step will have the consequence analysis provide direct feedback to the operating system to allow automated changes to component / system use in response to changes in component health. For example, when evidence analysis determines that a power supply is beginning to malfunction, the consequence analysis will evaluate different options such as shut down the power supply and terminate the operations immediately, operate the power supply at a lower power output to allow the system to "limp home", or continue to operate as originally scheduled, but shorten operations commensurate with failure predictions. This is the future direction for the RTCE.

REFERENCES

1. Multiple portions of the information contained in this paper were obtained from: Swiler, Laura P., James E. Campbell, et al, Algorithm Development for Prognostics and Health Management (PHM), SAND Report, SAND2003-3820, October, 2003.

2. Briand, Daniel, Project Manager. Data Driven Prognostics: PHM for Airborne Laser Systems, WFO Proposal Package, Proposal 152040915, 6 June, 2005.

BIOGRAPHIES

Daniel Briand
P.O. Box 5800 MS 1011
Sandia National Laboratories
Albuquerque, NM 87185-1011
Sandia National Laboratories

e-mail: dbriand@sandia.gov

Daniel Briand is a Senior Member of the Technical Staff at Sandia National Labs focusing on reliability and prognostic issues for system sustainment and readiness programs. His background is in operations analysis and aircraft operations & training, the majority of which stems from having served 20 years in the US Air Force as an operations research analyst and helicopter pilot. Dan is a graduate of the US Air Force Academy, has a MS in Operations Research from the Air Force Institute of Technology, and is currently pursuing a Ph.D. in Statistics from the University of New Mexico

James E. Campbell, Ph. D.
CompP.O. Box 5800 MS 1011
Sandia National Laboratories
Albuquerque, NM 87185-1011
Sandia National Laboratories

e-mail: jecampb@sandia.gov

James Campbell is a Distinguished Member of the Technical Staff at Sandia National Labs. Prior to joining Sandia, he was a Vice President at Intera Inc. in Austin, Texas. His areas of expertise include reliability analysis, decision analysis, analysis of complex systems of systems, and software design and engineering. He has bachelor's degrees in mathematics and physics and a Ph.D. in physics.