Synchronized Phasor Measurements for Response-Based One-Shot Control

S. M. Rovnyak, *Member, IEEE*, and K. Mei, *Student Member, IEEE*

***Abstract*--This paper presents some new ideas about response-based transient stability control using apparent resistance R and its rate of change Rdot. These ideas were developed while teaching a special topics course on one shot control at IUPUI. The reader is referred to previous publications for background information. The potential application to synchronized phasor measurement units (PMUs) is then explained. Some of the new ideas are important for the extension to PMUs.**

***Index Terms*-- Stability control, decision trees, transient stability.**

I. INTRODUCTION

THIS paper describes an algorithm to automate the design of wide area stability controls using pattern recognition tools such as decision trees (DTs). The method is applied to the phase-plane design of an out-of-step relay as in [1]. The resistance R is the real part of V / I measured around the center of an intertie. R can become small very quickly during short circuits. When there is loss of synchronism R becomes small relatively slowly. These features were used to detect loss of synchronism using R and Rdot in [2], [3].

The DT pattern recognition methodology uses large amounts of simulation data to develop the controls. The resulting control scheme is not necessarily optimal. Our goal is only to demonstrate net improvement in the power system because of the control.

One shot refers to a control that operates in a discrete action such as opening a circuit breaker, tripping generators or shedding load. The term also refers to a step change in a continuous quantity such as HVDC power transfer. One shot controls that unstress the power system are probably going to be less sensitive to communication time delays than feedback controls.

One shot controls are easier to simulate because it is not necessary to imbed the control algorithm in the power system simulation software. It is possible to run the simulation once without control, analyze the output and rerun the simulation with control. Simulation programs often have a file that specifies faults and line tripping and so on to occur during a simulation. One shot controls such as tripping generators and switching capacitors and shedding loads can be represented easily in this file.

This material is based upon work supported by the National Science Foundation under Grant No. ECS-0426189. Kejun Mei received travel reimbursement from Carnegie Mellon University.

S. M. Rovnyak is with the Department of Electrical and Computer Engineering, Indiana University - Purdue University Indianapolis, Indianapolis, IN 46202 USA (e-mail: srovnyak@iupui.edu).

K. Mei is with the School of Electrical and Computer Engineering, Purdue University, West Lafayette, IN 47907 USA (e-mail: kmei@purdue.edu).

II. EVENT AND RESPONSE-BASED CONTROLS

Event-based one-shot controls are very common. They are often called remedial action schemes (RAS) or special protection systems (SPS). These controls are predetermined for specific events through off-line simulation. They typically consist of generator tripping and reactive element switching.

Event-based controls are helpful for meeting reliability criteria. The well known N-1 criterion means that the power system should withstand the loss of any one major component. When there is an operating point for which the loss of one major component causes instability or the violation of some operating constraint, it is necessary to design controls that are automatically triggered in response to the loss of that particular component. These event-based controls are typically designed to prevent violation of an operating constraint for the loss of one particular component.

Response-based controls are not automatically triggered by the loss of some particular component. Response-based controls are proposed to operate based on measured and usually continuous quantities such as real power flow or apparent resistance, for example. The R-Rdot relay is response-based like most relays. When the R-Rdot relay is used to trigger one shot controls then it is performing response-based control. Response-based controls can potentially operate for a much larger number of events than event-based controls.

III. DESIGN FOR RESPONSE-BASED CONTROL

First, the methodology detailed in [1] can be used to train DTs to detect loss of synchronism across the Pacific AC Intertie (PACI). Simulation output was converted to input-output pairs for DT training as follows. In the table below, the PACI angle refers to the phase angle difference between generators on the opposite ends of the PACI. The DT input vector consists of { R , Rdot }. The target refers to the output the DT should produce for the given input vector. In this example the target is set to 1 (Action) if the PACI angle is larger than 120 degrees.

Time	R	Rdot	PACI angle	Target
4.833	17.64	-26.77	112.10	0 (None)
4.850	16.89	-44.70	114.16	0 (None)
4.867	16.16	-44.02	116.32	0 (None)
4.883	15.40	-45.51	118.59	0 (None)
4.900	14.58	-48.93	120.99	1 (Action)
4.917	13.84	-44.50	123.51	1 (Action)
4.933	12.72	-64.78	126.16	1 (Action)

This paper proposes a simpler method of converting simulations to input-output pairs. The previous method required finding the 120 degree threshold by hand. The method in this paper uses DT software CART [4] to find such a threshold separating stable from unstable swings. If any angle difference exceeds 360 degrees in this model then the simulation is definitely unstable. All the samples from an unstable simulation including those with PACI angle less than 120 degrees are defined to have target = Action. Fig. 1 shows how the training set looks according to this definition. Points with target = Action are marked 'x' ; points with target = None are marked 'o'.

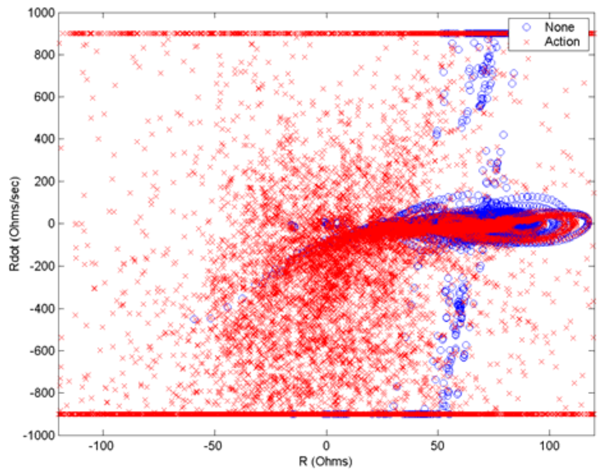

Fig. 1. Training data for DT classifier. 'x' = Action. 'o' = None.

IV. Threshold Between Action and None

The threshold of 120 degrees in [1] was found by hand. In higher dimensions it will be necessary to use software tools to find such thresholds. The discussion in this section about the threshold between Action and None might be generalized to the situation with higher dimensions.

It is safe to assume if the PACI angle exceeds 360 degrees then there is loss of synchronism. If the PACI angle never exceeds 360 degrees then the simulation is probably but not necessarily stable. The set of simulations with the PACI angle less than 360 may include a few that are beginning to lose synchronism near the end of the simulations. We suspect this set does not contain many simulations losing synchronism one or two seconds before the end. Early loss of synchronism would likely result in an angle greater than 360 before the end of the simulation. Results below show no simulations losing synchronism at one second before the end.

We calculated three numerical distributions of samples as functions of the PACI angle as shown in Fig. 2. There were 385 six-second simulations. Data for all three distributions comes from the set of simulations with the PACI angle always less than 360 degrees. The bottom line was obtained by including only the first 5 seconds of those simulations. The middle line counts samples in the first 5.5 seconds. The top line represents all six seconds of simulation data. The graph zooms in on the part where the counts of samples go to zero.

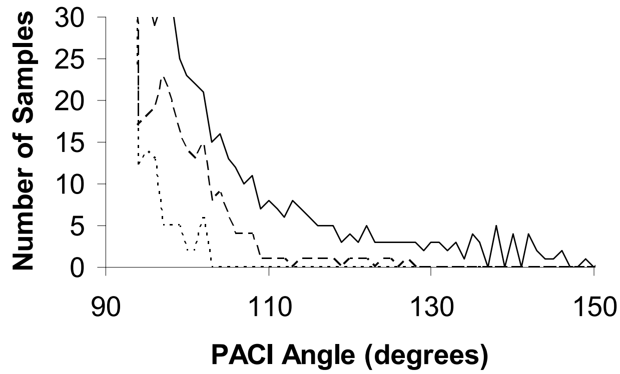

Fig. 2. Numerical distribution of samples as functions of PACI angle. Tails in the top two distributions represent a few events beginning to lose synchronism at the end of a simulation.

Consider the data represented by the lower line that is plotted. Fig. 2 shows this collection contains no sample with the PACI angle greater than about 103. In the collection represented by the middle line it looks like there is just one simulation that goes over 110 degrees. It appears the criterion of 120 degrees derived by hand was conservative. We could have defined the target for a data sample to be Action if the PACI angle is greater than 103 degrees.

The DT software CART can be used to find the 103 degree threshold from the data [4]. First, let us introduce a very useful parameter in the DT building process called relative misclassification cost. The misclassification cost affects where the decision boundary is drawn when there is overlapping data with different target values. For example, the one node DTs in Fig. 3 are obtained using different misclassification costs. The DTs are trained from the data in Fig. 1 and they perform R-Rdot relaying as in [1]. When the DT is forced to be one node then it is just an R relay.

The pre-fault equilibrium value of R is 72 ohms. The DT in Fig. 3 with lower relative misclassification cost will output Action when R = 36.8 ohms. The other DT will output Action when R = 22.6 ohms. The apparent R is moving relatively slowly when loss of synchronism begins to occur. The DT with lower relative misclassification cost will output Action sooner.

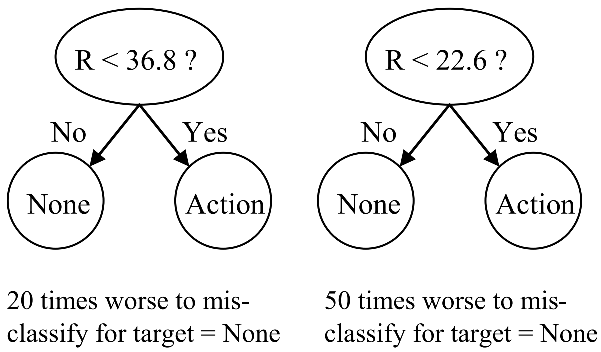

Fig. 3. Different one-node relays resulting from different relative misclassification costs.

We set the Target = Action for all samples of all simulations that had any PACI angle greater than 360 degrees. We set the Target = None for samples from the first five out of six seconds of simulations that had the PACI angle always less than 360. The samples with Target = Action are scattered all over the place in Fig. 1. Most of the samples with Target = None have R greater than about 20 Ohms. We set the cost of misclassifying a sample with Target = None to be 1000 times more than cost of misclassifying a sample with Target = Action. We construct DTs from data using CART. Pruning the tree back to one node, the result is the PACI angle < 103.0 degrees. The resulting tree classifies as Target = Action almost everywhere except where points with Target = None exist. The CART software has a complexity parameter that might be used to obtain a good fit for higher dimensional problems. A one node DT may not be a good solution in higher dimensions.

Another useful subset of the training data includes only a relatively short segment of each unstable trajectory around the time loss of synchronism begins to occur. Long after loss of synchronism occurs the R-Rdot trajectory can go all over the place as seen in Fig. 1. This widespread data impairs the performance of a DT designed to detect loss of synchronism. A previous method of handling this data was to set Target = None everywhere except the general area where we expected the boundary to be drawn [1]. The data shown in Fig. 4 contains only a segment of each unstable simulation during which the PACI angle was between 100 and 200 degrees. The method described in this paper may generalize better to higher dimensions than the previous method.

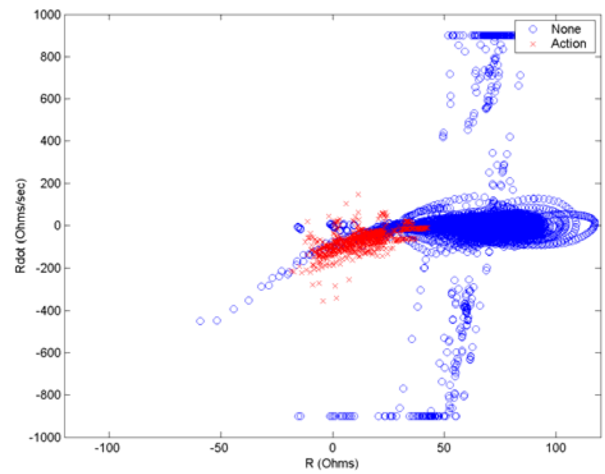

Fig. 4. Improved training data. Only those samples with 100 degrees < PACI angle < 200 degrees are included from the unstable simulations. Data from all six seconds of stable simulations are included.

V. Design for Response-Based Control Continued

This section summarizes the methodology for designing response-based control in [5].

Step 1: Train DT to Detect or Predict Stability.

Described above in Section 4 and in [1], [5].

Step 2: Find a Good Combination of Controls by Trial and Error.

Choose one specific combination of one shot controls. Control combinations that reduce steady state phase angle differences are good candidates. Change the file that specifies the faults and one shot controls for every training set simulation that has two consecutive DT outputs of Action. Control should be simulated to occur following a short actuation delay such as 0.1 second after the second consecutive DT output of Action. Rerun all the simulations in the training set that have one shot control initiated by the DT and evaluate the results. Choose a different combination of controls and repeat.

Reference [6] describes how to calculate through simulation the sensitivity of an index to various controls and how to use this information to search for an optimal combination of one shot controls for one simulation. The method is a form of gradient descent so it can get stuck in local minima. It is still able to find combinations of controls that stabilize some extremely unstable events [6]. A similar approach might be tried for one shot control using R-Rdot. The index could be summed over all the training set simulations that have control action in the R-Rdot application.

There are two HVDC lines in our 176 bus model. HVDC fast power changes are inexpensive compared with tripping generation and load. The results discussed further below were obtained by first performing an exhaustive search of all feasible combinations of HVDC fast power change amounts.

Then tripping different generators were simulated to occur in addition to the HVDC power changes. Typically one or two generator trips were selected for inclusion in the final combination of one-shot controls. An objective function approach was used in which one point was added for each simulation stabilized by the control and three points were subtracted for each simulation destabilized by the control.

All the experimentation in Step 2 is done on the training set. A Final Control Combination (FCC) having the best objective function score is selected after all the iteration is done.

Step 3: Evaluate DT to Trigger the FCC on New Simulations.

Run a test set of simulations different than the training set. Trigger the FCC the first time during any simulation that a set of measurements results in two consecutive DT outputs of Action. Evaluate the results over all the test simulations.

VI. Simulation Study

Reference [5] describes a simulation study involving a 176-bus simplified model of the WECC with 29 detailed generator models. There was data from 385 simulations in the training set and from 1600 simulations in the test set. The events covered a wide variety of contingencies including double outages.

The simulated R and Rdot were monitored near the middle of the PACI as in [1], [5]. The FCC in [5] consisted of three simultaneous one-shot controls. There were two HVDC fast power changes and one generator tripping. There were two consecutive DT outputs of Action in 116 of the 385 training simulations and in 491 of the 1600 test simulations.

Reference [5] considers two options for assessing stability. One way is to declare a simulation stable if there is no loss of synchronism across the PACI. The other way is to declare a simulation stable if there is no loss of synchronism anywhere in the network. The former definition turns out to be much better suited for selecting the Final Control Combination (FCC) in the R-Rdot application. Tables I and II show the performance of response-based control that is trained to maintain synchronism across the PACI only. In other words, loss of synchronism across the PACI is the criterion for whether control is counted as stabilizing a simulation in the objective function used in training. Tables I and II evaluate the performance of the same control according to different criteria. Table I shows how many simulations were stabilized according to whether synchronism is maintained across the PACI. In Table II the definition of stability is whether there is synchronism across the entire network.

TABLE I

Performance of DT-triggered one-shot control

	Train	Test
Stabilized By the Control	46	253
Stable With or No Control	61	185
Unstable With or No Control	9	53
Destabilized By the Control	0	0

TABLE II

Performance of same control. Stability now refers to maintaining network-wide synchronism.

	Train	Test
Stabilized By the Control	5	19
Stable With or No Control	59	103
Unstable With or No Control	52	369
Destabilized By the Control	0	0

The first line from each of the preceding tables is summarized in Table III as a guide for interpreting the data in Table IV, which shows the performance of a DT trained to maintain synchronism across the entire network. The performance of this DT according to the network-wide criterion on the training set is the only number that is not many times worse than the previous DT. An event can be stabilized network-wide but not according to PACI angle because the thresholds used to judge their stability are different.

TABLE III

Summary performance of DT trained for PACI synchronism

	Train	Test
Stabilized Over PACI	46	253
Stable Network-Wide	5	19

TABLE IV

Summary performance of DT trained for network-wide synchronism

	Train	Test
Stabilized Over PACI	10	74
Stable Network-Wide	17	1

These DTs that use R and Rdot effectively observe the PACI angle. The performance of the DT shown in Table III generalizes well from the training set to the test set. The test set numbers are roughly proportional to the training set numbers. The DT in Table IV does not generalize well from the training set to the test set.

VII. Extension to Synchronized Phasor Measurements

A similar design that uses synchronized phasor measurement units (PMUs) could potentially be trained to maintain stability network-wide and have good generalization from the training set to the test set. Similar to the technique described above we can define the target = Action for all samples of a simulation that has any generator angle difference greater than 360 degrees. This definition corresponds to the network-wide stability criterion. Then we will set a large cost on misclassifying a case that has target = None. We may experiment with a complexity cost parameter that can be used to reduce the size of a DT in CART. There is a need to force the DT to have a small number of nodes to prevent over training.

The DT input vector will contain phase angle measurements and their rates of change similar to [7]. The last second of data will be removed from any training data having target = None. The result is expected to be similar to the data represented by the lowest line in Fig. 2. Samples in the training data with target = None will be in the normal range for stable simulations. The resulting DT is expected to reflect the boundary in feature space between stable and unstable simulations. The DT thresholds are expected to be significantly less than 360 degrees.

The DT will be trained to output Action only when phase angle differences are outside the stable operating region. The target = None samples have been designed to exclude data from the end of simulations that are beginning to lose synchronism. If these samples are not excluded then the data will contain a few samples with angle differences up to 360 degrees. If the relative misclassification cost is set very high, the resulting DT will have corresponding thresholds closer to 360 degrees, which would not be as useful. A good DT will trip early for unstable simulations and should not trip unintentionally for stable simulations.

VIII. Conclusions

This paper has demonstrated some new techniques for converting simulation data to input-output pairs for classifier training. Setting target = Action for all the samples from unstable simulations is demonstrated for an R-Rdot relay. The possible extension to phasor measurement will rely on a complexity cost parameter or human involvement to get a DT of appropriate size. The DT will have to be larger than one node to be effective for judging network-wide stability.

After the system trajectory crosses the basin boundary the rate of angle separation increases. After this point it appears to take less than one second for the angle difference to exceed 360 degrees in this model. A technique of omitting the last second of data from each "stable" simulation causes the DT boundary to be closer to the basin boundary. A "stable" simulation means that no angle difference exceeds 360 degrees although loss of synchronism could be starting to occur at the end of the simulation. Data from all the "stable" simulations will have target = None in the training set. Instability at the end is rare but in a large and varied data set it is likely to occur. When it does happen it generates data with angles greater than those of stable simulations. If the relative misclassification cost is set too high then the DT thresholds would be closer to 360 degrees and less useful.

The method described in the PowerPoint slides for this conference is slightly different. The slides suggest setting target = "Action" when the maximum angle difference in the network exceeds 360 degrees at the time the sample is taken. The proposal in the slides is expected to have some DT thresholds closer to 360 degrees and to not be as useful.

IX. Acknowledgment

The authors are thankful to all the organizers for this conference and for travel funds for graduate student Kejun Mei. The citation for this paper is provided in [8].

X. References

[1] S. M. Rovnyak, C. W. Taylor, and Y. Sheng, "Decision Trees Using Apparent Resistance to Detect Impending Loss of Synchronism," *IEEE Trans. on Power Delivery*, Vol. 15, No. 4, pp. 1157-1162, October 2000.

[2] C. W. Taylor, J. M. Haner, L. A. Hill, W. A. Mittelstadt, and R. L. Cresap, "A New Out-of-Step Relay with Rate of Change of Apparent Resistance Augmentation," *IEEE Trans. on Power Apparatus and Systems*, Vol. PAS-102, pp. 631-639, March 1983.

[3] J. M. Haner, T. D. Laughlin, and C. W. Taylor, "Experience with the R-Rdot Out-of-Step Relay," *IEEE Trans. on Power Delivery*, Vol. PWRD-1, No. 2, pp. 35-39, April 1986.

[4] [Online] www.salford-systems.com

[5] K. Mei, and S.M. Rovnyak, "Response-based decision trees to trigger one-shot stabilizing control," *IEEE Trans. on Power Systems*, pp. 531-537, February 2004

[6] G. Li, and S.M. Rovnyak, "Integral Square Generator Angle Index for Stability Ranking and Control," *IEEE Trans. on Power Systems*, Vol. 20, No. 2, pp. 926-934, May 2005.

[7] S.M. Rovnyak, S.E. Kretsinger, J.S. Thorp, and D.E. Brown, "Decision Trees for Real-Time Transient Stability Prediction," *IEEE Trans. on Power Systems*, Vol. 9, No. 3, pp. 1417-1426, August 1994.

[8] [This paper] S. M. Rovnyak and K. Mei, "Synchronized Phasor Measurements for Response-Based One-Shot Control," Second Carnegie Mellon Conference in Electric Power Systems, Pittsburgh, PA, Jan. 11-12, 2006. http://www.ece.cmu.edu/~electriconf/

XI. Biographies

Steven M. Rovnyak (S'89-M'95) was born in Lafayette, IN, on July 4, 1966. He received the B.A. degree in mathematics in 1988 and the B.S., M.S., and Ph.D. degrees in electrical engineering in 1988, 1990, and 1994, respectively from Cornell University, Ithaca, NY.

Currently, he is Assistant Professor of Electrical and Computer Engineering at Indiana University - Purdue University Indianapolis (IUPUI). His research interests include the development of pattern recognition methodologies for protective relaying and for one-shot stability controls. He is also working on protective relaying for distributed generators.

Kejun Mei (S'01) received the B.S. degree in computer engineering from Wuhan University of Hydraulic and Electric Engineering, Wuhan, China, in 1990, and the M.S. degree in computer engineering from China Electric Power Research Institute, Beijing, China, in 1993.

He is currently pursuing the Ph.D. degree in computer engineering at Purdue University, West Lafayette, IN. His research interests are in data mining, digital signal processing, pattern recognition and the application of artificial intelligence techniques to the monitoring and control of large systems and industrial processes.