REGULATORY ASPECTS OF DIGITAL TWINS

BRIAN M. GOLCHERT
Manager, Focused R&D,
WESTINGHOUSE ELECTRIC COMPANY

SCOTT E. SIDENER
Chief Engineer, Digital,
WESTINGHOUSE ELECTRIC COMPANY

Across the nuclear industry, many applications of digital twins are being considered for existing plants and systems as well as for new reactor designs. Many definitions of digital twins exist. For this discussion, digital twins are defined as combinations of multiple types of analysis coupled with data to give a real-time virtual analytic representation of a process, component, or system with potential, informed feedback to the operating process, component, or system. A digital twin may combine multiple types of analysis (such as fluids analysis, structural analysis, reactor physics calculations, and statistical analysis) coupled with live and/or historical data.

To be clear, digital twins are not a 'new' idea. Coupling data to numerical analyses has been done for decades with quite a bit of regulatory success. Instrumentation and Control (I&C) systems may have numerical analyses built into their logic to provide warning and/or control thus making them forms of digital twins. Control room simulators are essentially entire plant digital twins using pre-computed numerical analyses. What typically defines the 'modern' digital twin is speed and therefore the implication of real-time or as-needed real-time analytics.

Figure 1 shows a schematic of the potential components of a digital twin. In this case, a digital twin for a reactor coolant pump seal is depicted. The orange box on the left represents the physical operating component while the blue box on the right represents the digital counterpart to the physical pump seal. It contains the computational models that virtually rep-

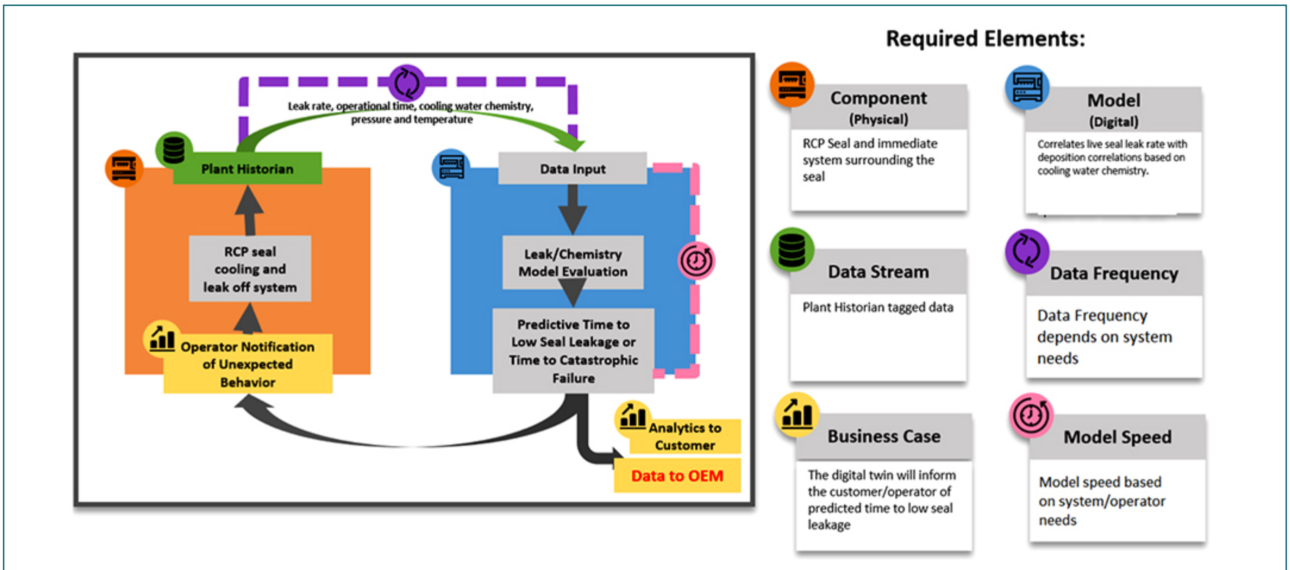

Figure 1. Schematic of Components of an Example Digital Twin.

resent the operating physical component. The arrow on the top that connects the orange box to the blue box represents the transmission of data from the physical component to the digital domain (models). The arrow at the bottom connecting the blue box to the orange box represents any feedback (control) from the model to the physical component. The arrow at the bottom right represents the transmission of the model outputs to the customer for their records and to the manufacturer of the seal. The information harvested from the seal by the manufacturer could be used to enhance the design of the component as well as improve the digital computational models.

Ideally, the digital twin should execute at near real time or as close to the time needed for the process. Most processes in the nuclear industry do not change rapidly so 'real time' may be over a relatively large time scale. As an example, a digital twin of the reactor that monitors/predicts control rod movement during steady-state operations is on a very slow time scale (control rod positions change very slowly as a function of time). On the other hand, a digital twin built for accident analyses must have the capability of running in 'faster than real time' to predict/minimize negative reactions due to an accident.

A frequent barrier to real-time operation comes in the execution time of a digital twin that has high fidelity models coupled to live data. Often, these high-fidelity physics-based models generate huge amounts of data (sometimes generating terabytes of data) and it is not practical to execute or use all of the data from these models due to their magnitude. A typical method for reducing the execution time of the digital twin is to replace slow high-fidelity models with high-speed surrogate representations or Reduced Order Models (ROMs). Often the creation ROMs include the application of machine learning (ML) methods. These ROMs are expected to capture the required details of the high-fidelity models in a functional form from the immense amount of high-fidelity analysis data thus allowing the quicker execution of the twin.

The nuclear industry has established guidelines for submitting individual types of analyses listed above (for example, NUREG-2152[1] provides guidelines for computational fluid dynamics analysis of spent fuel storage devices). However, currently, there does not exist an established guideline for documenting the creation, performance, Verification and Validation (V&V), and credibility of a digital twin, though there is currently some regulation in draft form for comment. Additionally, there is no clear guidance for the creation and V&V of the ROMs which may be used in digital twins. It is crucial that the ROMs accurately represent the pertinent portions of the physical process being modeled, an established process to create, assess, and V&V these ROMs is a necessity particularly for their increased use in advanced reactor analyses.

In addition, the statistical variation of these ROMs, when compared to either the results of high-fidelity simulations or real streaming data, warrants Uncertainty Quantification (UQ) to ensure the credibility of the resultant predictions or outputs. This is analogous to the UQ that is required by regulation for many of the potential components of a digital twin. For complex digital twin applications, it is quite possible that multiple ROMs created for multiple analyses are required. It is thus necessary to evaluate and quantify the uncertainties of the entire digital twin by propagating through the combined ROMs and propagate uncertainty through the model hierarchy when coupled together into a digital twin. A consistent methodology used to create the ROM, quantify, and propagate the uncertainties associated with inclusion of ROM(s) into a digital twin would be extraordinarily helpful in future advanced, and existing, reactor digital twin regulatory submissions. A methodology for this predictive Verification, Validation, and Uncertainty Quantification (VVUQ) is required to ensure credibility of the ROMs for use in design applications & analyses.

For any proposed digital twin guidelines, it is expected that existing regulatory and code guidelines would be leveraged as well incorporating advanced VVUQ methods. However, any proposed guidelines need to keep cost in mind otherwise regulatory requirements would make the building, use, and regulation of a digital twin financially untenable. Without practical guidance where 'what-if' scenarios' for the regulatory process are reasonably limited, it is quite possible that the regulatory requirements would become onerous to the point where digital twins would not be practical. A process to create and approve a digital twin that cost-effectively ensures safety would be greatly desired and is suggested here.

The topic of a 'what-if' scenario needs further explanation. For any single analysis submitted to the regulator, one generally needs to perform a best estimate case and also perform uncertainty analyses to determine a confidence or tolerance interval around the best estimate case. These extra cases alone could require a significant time and financial expenditure to perform. Consider a simple digital twin that integrates two computational analyses. From a regulatory perspective, the combined effect of UQ through both analyses to the result is needed. Each of the analyses would have their own base case and uncertainty quantification. It should not be required as part of any regulatory guidance to exercise both analyses simultaneously with all of the differing conditions to produce the combined effect of uncertainty. As can be imagined, as the number of different computational analyses in the digital twin increases, the combinations of evaluated conditions increase dramatically. So, rather than performing all of these 'what-if' scenarios, it is suggested that efficient methods for statistical combination of the uncertainties for all the individual analyses be utilized when determining the uncertainty of the digital twin. As a suggestion, using a Square Root Sum of Squares (SRSS) combination of all the pertinent uncertainties would be appropriate assuming that the uncertainty from each analysis is independent from the other analyses, as often is the case.

Implicit in the above discussion is that all best estimate analyses are based on the same approximate conditions. The demonstration of consistent and accurate combination of uncertainties from the multiple analyses should be one of the criteria in any future regulatory guidance. To be sure, it might not always be possible to achieve consistency for all analyses in a digital twin. If there are cases where the best estimate cases differ, the submitter should be required to provide documentation detailing a more sophisticated error analysis than SRSS.

---

[1]NUREG-2152, "Computational Fluid Dynamics Best Practice Guidelines for Dry Cask Applications: Final Report," March 2013.

To be clear, any uncertainty propagation analysis should only be done after defining target parameters (Key outputs of the digital twin to be used). As an example, peak clad temperature or maximum void fraction could be chosen as a target parameter. Any uncertainty propagation analysis should only include those inputs and analyses that affect those target parameters. Limiting scope and simplifying VVUQ reinforces the idea of cost-effectiveness, it is strongly suggested that a limited number of critical target parameters be required for uncertainty propagation to avoid excessive and non-productive cost/effort. The list of target parameters really depends on the functionality of the digital twin; effectively what is needed to be calculated or predicted for the intended application

Many numerical analyses produce a tremendous amount of data, and these large amounts of data can impede a digital twin running near real-time. Even with a fast supercomputer, it is highly unlikely to execute a digital twin near real-time with all this data. Fortunately, a large portion of the calculated numerical analysis data may not be 'needed' for the execution of a digital twin. Using the pre-selected target parameters, a limited set of input parameters that significantly affect the target parameter will need to be identified. After this selection process, a restricted number of numerical simulations that cover the expected operating conditions are run. For conservatism, a few cases that exceed normal expected operating parameters should also be run. All cases that covered the expected operating conditions provide data that would 'feed' into a reduced order model.

A frequent advantageous benefit of using ROMs is the simplification of a massive amount of data into a tractable functional representation with usually fewer required inputs than the original model. This reduction of inputs required to produce an analytic result can simplify the scale of the VVUQ. This alone, is of great value when considering the engineering costs of performing the analysis.

Ultimately, a reduced order model functionally represents the dependency of the target parameter on the key input parameters. It can be a single variable dependency, but it might likely be a multivariable dependency. Using this approach changes the numerical results from the analysis to a function that can be used in a digital twin allowing it to run much faster than a data lookup. Each analysis may have multiple target variable dependencies, but the process described below should be consistently applied to all ROM.

This type of data reduction is not new since it has been used in many industries (including nuclear) but due to the algorithms used to encode many forms of ROM's they can often be seen as being a 'black box'. To avoid regulatory issues with a black box approach, all of these considerations need to be included in the digital twin VVUQ. It is suggested that standard machine learning algorithms be used to create the reduced order model. However, that raises the question of what 'standard' and which model is applicable to the data set. To avoid these concerns, it is suggested that a regulatory submission apply multiple machine learning algorithms to the data and then decide which ROM is most credible, or 'best.' The definition of best needs to fully detailed in any regulatory submission (it could be something as simple as an R2 value). This approach somewhat eliminates bias towards the choice of specific machine learning algorithms. The final selected algorithm is then considered to be the best estimate ROM.

Deciding which ROM is best, or if any is good enough is the key challenge. A metric like R2 is based on a comparison to the selected validation points from the higher fidelity model. What the regulator really needs to know is if the ROM can be trusted to provide credible predictions over the intended application space or domain. In addition, the regulators will check to see if the ROM is well behaved and its uncertainty can be quantified in the operating space at points not exactly on those training and validation points. This requires an advanced parametric analysis, sensitivity analysis, and uncertainty analysis which should be included in any regulatory guidance for digital twins.

When training multiple ROMs with different ML algorithms to describe the same set of data, the variation between those proposed ROM's can be used to quantify a model form, or algorithm, uncertainty. This model uncertainty can be combined in the digital twin UQ. Since the data used to generate the ROM covers the entire expected operating range, it is assumed that this approach would provide a reasonable engineering evaluation of the final chosen ROM uncertainty.

Another key challenge is ensuring credible predictions over the intended application space. Quite often the application domain is a multi-dimensional space on the order of the number of key inputs to the digital twin. This typically high-dimension application space is almost never a uniformly distributed space. It can be a very sparse space with clusters of covariant data points. It may also be so large that it is impossible to cover the space with validation data. If high dimensional ROMs are included in a regulatory submission, it is suggested that an accurate mapping of the application space be performed.

So, the creation of a digital twin starts with a variety of detailed analyses, identifies target variables, and then creates ROM to determine the functional relationship between the target variables and the key inputs. These ROM are then combined to describe the component or system thus creating a digital twin. Since the digital twin is now made up of only functional relationships, it can run very fast.

Another conceptual regulatory question regarding a constructed digital twin is likely to be: Where does the digital twin 'live' as an entity? Commercial software vendors are offering their software to integrate the representative ROM of components, but this raises a regulatory challenge. How is it known that this commercial software provides 'real' validated results under all application? By, 'real' it is meant that the software is under configuration control so it will produce qualified results. Historically, software qualification has been part of regulatory submittals and if a digital twin is submitted using a commercial code, it should be expected that the software be qualified and some demonstration/proof that this software is more than adequate to perform its function would be required.

However, as was mentioned earlier, it might be possible to embed a digital twin into an I&C system. For submittals of this kind, it is suggested that normal protocols for I&C systems be applied with the inclusion of additional information on ROM generation and uncertainty analysis previously discussed. But one open ended question (which will not be answered here) regarding I&C is around autonomous control. It is conceivable that a digital twin of an entire reactor could be built that would allow remote autonomous control based on feedback from sensors in the system. In many ways, that is the holy grail of digital twins, but it is a topic for another day. The details to consider for regulation of digital twins with those characteristics is not covered here, however, one should always consider where the technology is going when considering what solutions are needed today to address regulatory needs.

Finally, all the previous discussions have been focused on the numerical analyses that create ROM that help build a digital twin. The value of a digital twin rests in the incorporation, interpretation, and evaluation of and intelligence derived from both live and historical data. It is this data that drives the performance of a digital twin. An example of applying a digital twin is found in component condition monitoring. In this situation, historical data is used in combination with live data to determine the current 'state' of the component or system. In some designs, this combination of historical and live data is used for Bayesian statistical and/or ML to predict expected near-term behavior using evolving or continuously re-trained models. From a regulatory perspective, the historical data quality and its acquisition can be dealt with using existing regulatory guidance. However, the interpretation of the historical data and utilization of the live data needs to be detailed in any submission. The historical data can be considered in the same vein as the ROM analysis from a regulatory perspective. Multiple machine learning models can be applied to the historical data to develop ROM of the expected behavior of that signal. This can all be detailed in the regulatory submission.

Live data needs a slightly different approach and there is no clear single path forward. A submission could include using a qualified ROM to interpret the live data (such as indication as to failure probability). This approach would not require additional regulatory oversight since the details of the ROM would have already been provided. However, it is expected that the usage of live data would be more sophisticated. The Bayesian approach would effectively continuously update the ROM. In a regulatory submittal, it would be expected that the technical details of the Bayesian or re-training methodology as well as some form of validation of the methodology be provided in the submittal.

In summary, modern digital twins provide compelling opportunities for operational value never before attainable, however, capturing that value will require resolving some key regulatory considerations that are associated with this technology. A digital twin is the 'glue' that holds a bunch of models (analyses) and data together with the physical operating systems. From a regulatory perspective, it is suggested that existing guidance be applied to creation these individual models as the foundational building block for the application of digital twins. However, there are special considerations that need to be defined and addressed specific to this technology. Suggestions presented here have a bias towards cost-effective approaches to VVUQ of computational digital twins such that efforts in digital twin submittal maintain business viability. It is expected that this suggested approach to regulatory guidance for digital twins will increase the adaption of digital twins throughout the industry. These suggestions represent discussion points for thoughtful consideration. They are not complete by any means, and the questions herein are likely to lead to additional questions. To be sure, the regulatory consideration for digital twins is a quickly advancing/evolving topic.■