High confidence design and implementation of cyber-physical systems (CPSs) is widely recognized as the enabler for a large number of exciting applications —in healthcare, energy-distribution, and industrial automation— with enormous societal impact. It is equally recognized that the current design and verification methodologies fall short of what is required to design these systems in a robust, yet cost-effective manner.

Current system design and verification techniques take a 0-1 view of the world: a property either holds or it doesn't, and a synthesized system provides no guarantees if the environment deviates in any way from the model. In practice, this view can be overly restrictive. First, CPSs need to operate for extended periods of time in environments that are either unknown or difficult to describe and predict at design time. For example, sensors and actuators may have noise, there can be mismatches between the dynamics of the physical world and its model, software scheduling strategies can change dynamically. Thus, asking for a model that encompasses all possible scenarios places an undue burden on the programmer, and the detailed book-keeping of every deviation from nominal behavior renders the specifications difficult to understand and maintain. Second, even when certain assumptions are violated at run-time, we would expect the system to behave in a *robust* way: either by continuing to guarantee correct behavior or by ensuring that the resulting behavior only deviates modestly from the desired behavior upon the influence of small perturbations. Unfortunately, current design methodologies for CPSs fall short in this respect: the Boolean view cannot specify or guarantee that small changes in the physical world, in the software world, or in their interaction, still results in acceptable behavior. In this poster, we sketch a theory of robustness for cyber- physical systems. Our starting point is the observation that a notion of robustness and associated design tools have been successfully developed in continuous control theory. There, the

control designer designs the system for the nominal case, while bounding the effects of uncertainties and errors on the performance of the system. Our goal is to provide a similar theory and algorithmic tools in the presence of both discrete and continuous changes on the one hand, and in the presence of more complex temporal specifications —given, for example, in linear temporal logic (LTL) or as ω-automata— on the other hand. We do this in three steps.

First, robustness is a topological concept. In order to define it, we need to give meaning to the words "closeness" or "distance." For this, we define a metric on the system states. Second, instead of directly modeling the effect of every disturbance, we model a *nominal* system (the case with no disturbance), together with a set of (unmodeled) disturbances whose effect can be bounded in the metric. That is, while making no assumption on the nature or origin of disturbances, we assume that the disturbances can only push a state to another one within a distance γ. Third, under these three assumptions, we show how we can derive strategies for ω-regular objectives that are robust in that the deviation from nominal behavior can be bounded as a function of the disturbance and the parameters of the system.