Research Brief
PRiME run smarter, live longer
Power-efficient, Reliable, Many-core Embedded systems www.prime-project.org

# Modelling for parameters with order graphs

***PRiME** is a five year EPSRC funded research programme (2013-2018), in which four UK universities address the challenges of **power consumption** and **reliability** of future high-performance embedded systems utilising many-core processors.*

## Structures vs. parameters in modelling

Modelling is an important part of complex system design. Traditional hierarchical modelling methods tend to have layers of abstraction corresponding to the naturally existing layers of multi-level systems. Although this is logical, it is not always the best approach for effective analysis and design. For instance, parts of a system in the same logical and structural layer may not contribute to the same degree on a metric such as system power consumption.

## Unnecessary model complexity

As the complexity of systems keep increasing – for instance the ever increasing number of cores in multi/many core systems – models are also becoming more complex. These complex models are more difficult to derive and debug, demanding increasing amounts of designer effort and providing decreasing degrees of confidence.

Structural-centric models, when used for studying a single (or a set of) parameters, inevitably include wasteful complexity. Ideally, less parametrically significant parts of the system should be modelled at higher levels of abstraction and more significant ones with more detail. This approach then focuses more on optimal parametric fidelity than on logical intuition.

## Order graphs

PRiME has studied ways to overcome modelling complexity in both structural and behavioural aspects:

1) Structural:

> Representations of systems as resource configurations that reflect the parameters of interest, such as power.
>
> Ways of zooming through these representations - with the purpose of finding the right balance in exploring the system at the right depth, whilst avoiding getting into unnecessary levels of detail.

2) Behavioural:

> Ways of capturing self-similarities, such as the behaviour of components that have little difference in their impact on the given parameters of interest (e.g. power), and hence do not need to be differentiated.
>
> Ways of compressing the behavioural differentiation by effectively "zooming out" from discerning deterministic trajectories to less discriminating stochastic patterns.

To achieve this, PRiME has introduced a new modelling formalism called **order graphs.** These have a clear hierarchical structure, whilst providing straightforward vertical zooming across multiple layers (orders) of model abstraction. This can be done independently in different regions of a model, resulting in the isolation of cuts that may run through different orders (for reasons important for the designer).

A significant innovation of order graphs is that an *arc* at a higher order (level of abstraction) is a *node* at the next lower order, as shown in Figure 1.

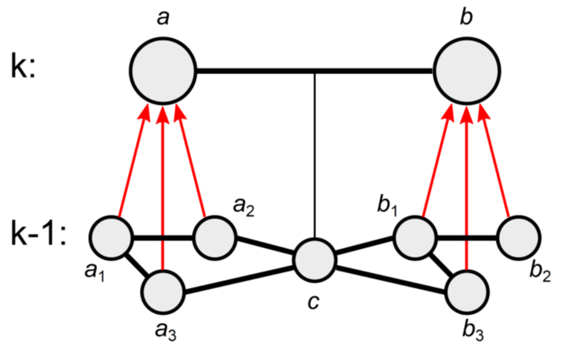


Figure 1. An arc becomes a node.

PRiME run smarter, live longer

EPSRC Engineering and Physical Sciences Research Council

With order graphs, it is easy to represent system resources with *nodes* and the dependencies between resources as *arcs*. A *node* at a higher order, therefore, can open up to an entire graph at a lower order, as we zoom in to look at the details of a resource. A dependency at a higher order is actually a resource when looked at in detail.

This concept can be understood by looking at Figure 1 as a model of a client-server system, with order k showing client *a* connected to server *b* – *nodes a* and *b* are resources and the *arc* between them representing their cross-dependency. This dependency at a more detailed level of abstraction – order k-1 – is exposed as a resource itself, namely *c*. Further orders below can show the details of *c*, perhaps a network.

## Parametric-proportional modelling

Figure 1 follows the conventional approach used by designers, with levels of abstraction following system structural and logical complexity levels. However, consider the case when studying power consumption and it is known that the client consumes a relatively small amount of power compared to the server, which in turn is less power hungry than the network. To have a proper power model accurate to a certain amount of wattage, we need to model the network down to a very complex level of detail. At that order, maybe k-3, the models of the server and especially the client would be unnecessarily complex. Each node would be contributing very little to the overall system power consumption.

Order graphs allow us to easily generate parametric-proportional cross-layer cuts across a system model. For instance we can keep the client model at order k with a single node, dig deeper into order k-1 for the server, and go down still deeper to k-3 for the network. The order graph method and tools help generate this multi-order cross layer cut, which can be used to study system behaviour focusing on the particular parameter or set of parameters (here power consumption) for which nodes across the cut make approximately the same contribution. In other words, at no place in the cut is the parameter under-modelled or over-modelled.

## Studying real systems

Experiments were conducted on an Odroid XU3 system. This is based on an embedded ARM big.LITTLE processor architecture with four ARM Cortex A7 cores for low power operations and four ARM Cortex A15 cores for high performance.

An order graph model of this system was constructed to a structural detail at the core level – i.e. each A7/A15 core was represented by a single node in the lowest order. Subsequent system characterisation produced data showing that each A7 core in typical operation uses less than ¼ the power of an A15 core. The order graph method then generated a cross-layer cut so that all four A7 cores are together represented by a single node and each A15 core is represented by a separate node (Figure 2). In analysis, this reduces the model state space by a factor of over 4000 without reducing power precision.

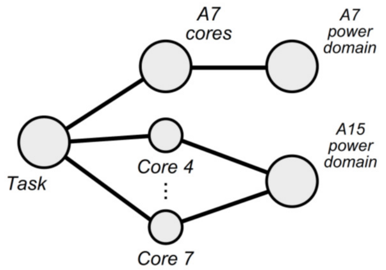


Figure 2. Odroid system model cross-layer cut for power.

## Future work: order graphs in RTM design toolkit

Runtime management (RTM) is a major concern of PRiME and this parametric-proportional modelling approach via order graphs will be fully adopted in the RTM design process, as well as for the models used within the RTM itself.

## More information

Visit the PRiME programme web site, including the opportunity to sign-up for programme updates, or contact the PRiME Collaboration Manager (including industry liaison):

**Gerry Scott**
Email: gerry.scott@soton.ac.uk
Tel: +44 (0)23 8059 2749

www.prime-project.org

Imperial College London | Newcastle University | MANCHESTER 1824 The University of Manchester | UNIVERSITY OF Southampton