

An Architectural Meta-application Model for Coarse Grained Metacomputing

Stephan Kindermann¹ and Torsten Fink²

¹ University of Erlangen-Nuremberg, Germany
snkinder@informatik.uni-erlangen.de

² Free University of Berlin, Germany
tnfink@computer.org

Abstract The emerging infrastructures supporting transparent use of heterogeneous distributed resources enable the design of a new class of applications. These meta-applications are composed of distributed software components. In this paper we describe a new model for component based meta-application design based on a formal architectural description of the gross organization of an application. This structural description is enriched by a formal process algebraic characterization of component behavior. Using this behavioral model we can formally check meta-applications in an early development phase. We present simple architectural styles developed to support data-flow and control-flow driven meta-application design on top of the Amica metacomputing infrastructure.

1 Introduction

There is a growing interest in defining and constructing an infrastructure which gives users the illusion that distributed, heterogeneous (computing and storage) resources constitute one giant transparent environment, a metacomputer [8]. It enables the development of a new class of applications: meta-applications. They are composed of multiple (partly reusable) components. Looking at the current practice of meta-application development there is no agreement on a common programming model for component based application design.

In this paper we describe an abstract, formally concise, and extendable programming model defined to develop meta-applications on top of our metacomputing environment Amica¹ [2]. The overall organization of meta-applications are given in a formal architecture description language (ADL) (see e.g. [6]). Different architectural styles [13] are defined to build up a basic description vocabulary which can be refined to define domain specific extensions. For every element of the vocabulary its behavior is defined by a process algebra term. The overall behavioral model of an application is automatically derived through the appropriate composition of the behavioral description of its components. This model can then be checked against a set of formal requirements (e.g. liveness

¹ Abstract Metacomputing Infrastructure for Coarse Grained Applications

and progress properties). This allows formal correctness checks on Amica meta-applications in an early development phase.

The remainder of this article is as follows. In Sect. 2 we introduce shortly the basic concepts and services of Amica. In Sect. 3 we introduce our new programming model. In Sect. 4 we describe firstly how applications given in our programming model are executed using Amica and secondly how a behavioral model of the application is generated automatically which can be used as input to formal analysis tools. In Sect. 5 we give an overview of related work and discuss the advantages of our approach. Finally, an outlook to future work is given.

2 The Amica Metacomputing Infrastructure

Amica has been designed as a prototypical middleware foundation to investigate the composition of metacomputing applications from reusable components (e.g. legacy systems). Amica provides abstraction of heterogeneous distributed data storage by *data objects* and of computing facilities by *metabricks*. Additional application specific code can be integrated using so called *user bricks*. Data objects (possibly replicated) are related dynamically to real storage resources (data store objects) interconnected by link objects which provide abstraction of the networking infrastructure of the metacomputer. Metabricks are related to the basic computational services (bricks) based on a broker mechanism which looks for appropriate brick factories. The concise meaning of 'appropriate' is given by a cost function which takes into account the current load of computing resources provided by special objects named computation units. Essentially, Amica provides an infrastructure to carry out the instantiation of abstract data storage and computing service requests transparently taking into account the current load within the distributed system. The implementation relies on the standard middleware foundation CORBA and uses in the current version an interpreter based instantiation approach. For a more detailed description of the Amica infrastructure see [2].

3 The Amica Programming Model

A metacomputing application on top of Amica is composed of data storage and computation components which are directly related to the abstract data objects and metabricks of Amica. User-bricks allow additional application specific code to be integrated. A basic set of control flow and data flow connectors is used to specify and control component activation and interaction.

In the following we formally describe the organization of meta-applications as a hierarchical collection of interacting components with well defined properties and interfaces. This structural description is combined with a formal characterization of component behavior based on a process algebra. On this combination we build up a basic vocabulary for meta-application description.

3.1 Architecture Description

Despite the variety of existing software architecture description languages (ADLs) there is a considerable agreement about the role of structure in architecture description. Our description of metacomputing applications is based on ACME [5] which emerged from a joint effort of the architecture research community to provide a common intermediate representation for different ADLs.

Meta-applications are given as collections of components interconnected by connectors in a meta-application graph. The structural description is enriched by a behavioral characterization of components and connectors using the process algebraic description language Lotos [10].

Definition 1. A meta-application graph (MG) is given as a bipartite graph, which is characterized by a quintuple $G = (Nodes, I, E, Beh)$.

- The *Nodes* of the graph consist of a set of components *Comp* and a disjoint set of connectors *Conn*: $Nodes = Comp \cup Conn, Comp \cap Conn = \{\}$.
- The function $I : Nodes \rightarrow 2^{IP}$ associates each node with its set of interconnection points ($\in IP$). Interconnection points are subdivided into ports (for components) and roles (for connectors). ($IP = Ports \cup Roles$).
- The interconnection of components and connectors via their ports and roles is given by the relation $E \subseteq (Comp \times Ports) \times (Conn \times Roles)$.
- The function $Beh : Nodes \rightarrow LotosTerm$ associates each node to its behavioral description in form of a Lotos process algebra term.

For a node c in a meta-application graph with an interconnection point $p1$ and an associated set of actions $\{a_1, \dots, a_n\}$ (e.g. services needed or provided or events emitted at the interconnection point) the associated process term $Beh(c)$ defines a Lotos process c with parameter list $[p1_a_1, \dots, p1_a_n]$. This list is extended accordingly if multiple interconnection points are defined for a node. Interaction with other (node-) processes is exclusively possible via synchronization with these externally observable actions.

3.2 An Architectural Style for Amica Meta-applications

Each node and each interconnection point in a meta-application graph is an instance of a type from a set of predefined type definitions. These types are used to build up a basic vocabulary *Voc* to describe the architecture of a meta-application. This vocabulary along with a set of constraints is often called an architecture style [13].

Definition 2. A meta-application vocabulary *Voc* to build up (behavioral) meta-application graphs is given as a quintuple (NT, IPT, SC, BC) where *NT* is a set of node (component or connector) types, *IPT* is a set of interconnection point (port or role) types, *SC* is a set of structural constraints, and *BC* is a set of behavioral constraints.




Fig.1. Basic meta-application vocabulary

The description of metacomputing applications on top of Amica is currently based on a simple architectural vocabulary which is illustrated in Fig. 1.

It contains component and connector types to characterize the flow of control and the flow of data, to describe farm parallelism, and to provide access to data storage and computing resources. For brevity we omit the structural and semantic constraints. In general, structural constraints are given by the associations in the class diagram and first order logic predicates. Semantic constraints currently are given in an action based temporal logic.

Data is stored in instances of the type `DataObjectCpT`. Its definition includes two port types for read-write access and two port types to create and delete object instances. The access to data object components is done via connectors of the type `DFlowCnT`. It provides roles for connection to data objects and to objects needing data access. Each component which is wired in control flow is an instance of type `CFlowCpT`. This type defines two interconnection points of type `CinPT`, `CoutPT`, characterizing incoming and outgoing control flow. Control flow components are connected to control flow connectors (type `CFlowCnT`). Different subtypes characterize connectors which split and combine control flow.

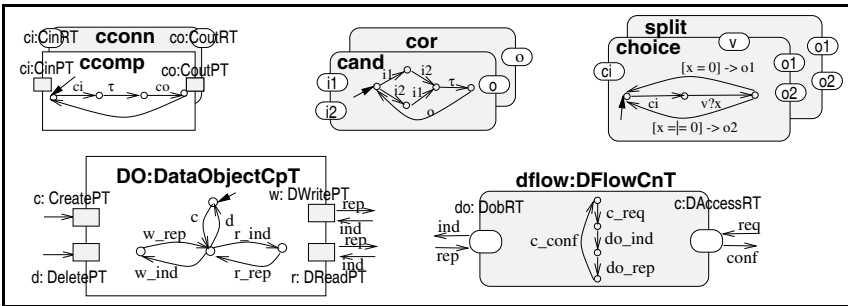


Fig.2. Basic data and control flow components and connectors

Examples of basic data and control flow components are given in Fig. 2. Their basic behavior is illustrated in the form of labeled transition systems. Thus control flow is simply propagated and data flow is based on a simple request–confirm protocol.

Instances of the special connector type **FarmCnT** are used to express and control "bag of task" like parallel computations. Farm connectors are used only in a well defined cooperation with data object and worker components, see Fig. 3. After started the farm reads in a bag containing the tasks to be distributed (using **bi:DobRT**) and starts a number of workers over **ws:CoutRT** (this number is given as a property value of the connector). Then the tasks are distributed to the workers using role **do:DAccessRT**. Thereafter the results are collected over **di:DAccessRT** and stored in a result bag data object attached to role **bo**. This description of the behavior corresponds directly to the Lotos description given as skeleton in Fig. 3.

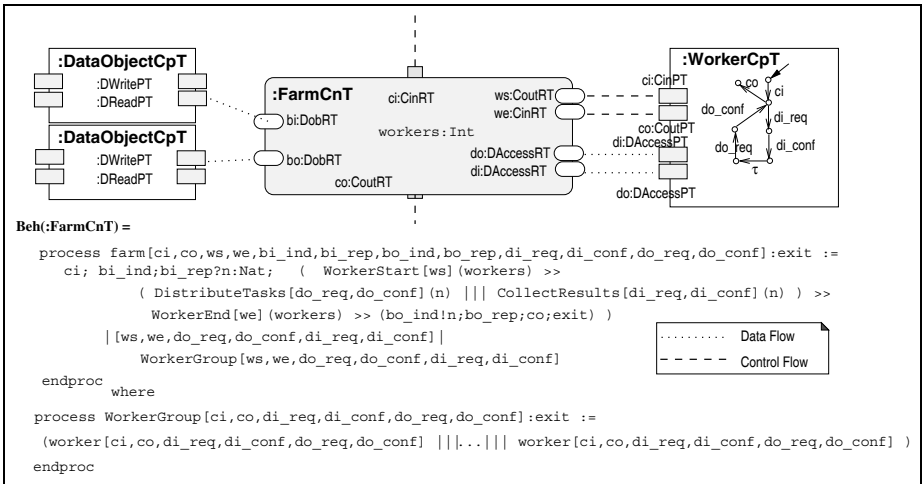


Fig.3. Bag of task like parallelism based on farm connector and worker component

3.3 A Small Example

A simple exemplary composition based on our vocabulary defining a ray tracing meta-application is given in Fig. 4. Two data objects are involved, **scen** stores the three dimensional scenario and **pic** stores the generated picture.

In **init** these data objects are created and initialised. Then, the remote computation is started by a metabrick. In parallel the user can monitor the current state of the picture by a specialized user brick. Using the graphical

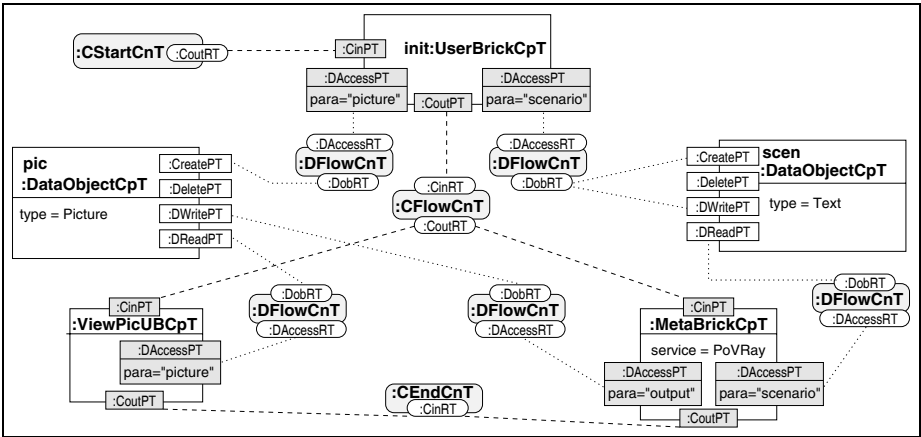


Fig.4. An exemplary application

front-end for ACME this application can be intuitively defined by drag-and-drop nodes out of our meta-application vocabulary. Another simple application to parallel simulation of mobile communication systems is described in [2].

4 Meta-application Execution and Formal Analysis

A meta-application graph is dynamically interpreted and mapped to the Amica metacomputer. Also automatically a global behavioral model of the meta-application is generated in a compositional way. This model can be analyzed and checked for design errors (e.g. resulting from component composition mismatches). These two steps, namely interpretation and behavioral model generation are discussed in the following section in more detail.

A simple meta-application interpreter is used to map our basic component and connector vocabulary to the services provided by Amica: Components of type `DataObjectCpT` and `MetaBrickCpT` are directly correlated to the data object and metabrick abstractions of basic data storage and computing services provided by Amica. The data store network and the broker mechanism of Amica relate these dynamically to the distributed data objects and bricks (created by brick factories). Data flow connectors and control flow connectors are interpreted to control component interaction and activation. In the special case of a farm connector (type `FarmCnT`) a specified number of workers is instantiated, the tasks contained in a bag are distributed and the results are collected.

The structure information given in the meta-application graph is used to compose the individual node behaviors to the combined overall system behavior. Composition is based on appropriate synchronization of component and connector actions. To correlate these actions we have to define a renaming operator

$\setminus\{(old_1 \setminus new_1), \dots, (old_n \setminus new_n)\}$ which replaces all actions old_i_act over interconnection point names old_i to new_i_act .

Given a behavioral meta-application graph $G = (Nodes, I, E, Beh)$ with $Nodes = Comp \cup Conn = \{CP_1, \dots, CP_N\} \cup \{CN_1, \dots, CN_M\}$ the associated overall system behavior is given by the parallel (full interleaving $|||$) composition of all component instances and also all connector instances. These two groups synchronize over all actions of their interconnection points ($||$). Thus the general structure of the overall system behavior description is given as the following process term scheme:

$$(NCP_1 ||| \dots ||| NCP_N) || (NCN_1 ||| \dots ||| NCN_M)$$

where

$$\begin{aligned} NCP_i &= Beh(CP_i) \setminus \{(old \setminus new) \mid old \in I(CP_i) \wedge new = CN_j_newp \\ &\quad \wedge attachedCN(CP_i, old) = (CN_j, newp)\} \\ NCN_i &= Beh(CN_i) \setminus \{(old \setminus new) \mid old \in I(CN_i) \wedge new = CN_i_old\} \end{aligned}$$

The function *attachedCN* used above gives exactly one attached (connector, role) pair for a given component and port, that is attachments of multiple roles to a port are disallowed. In the case of multiple attachments of ports to a role, the above scheme implies an or semantics (the role-action synchronizes with an action associated with one of the attached ports). We have extended this to generally allow n-port to one role attachments with different semantics (e.g. and) given by the type of the role. This generation scheme was implemented in Java and we apply powerful formal analysis tools (the CADP tool set [4] and the model checker XTL [11]) for abstract meta-application property checks.

5 Related Work and Conclusion

Different approaches to build component based meta-applications are used in literature, ranging from simple data-flow models to general component frameworks. In WebFlow [9] (restricted) data flow graph models are proposed for component composition. In [14] a distributed component architecture toolkit is described for meta-application design on top of the Globus metacomputing infrastructure [3]. Also scripting languages can be used for meta-application description, see e.g. [12]. In contrast to these approaches having no or only a very restricted semantic foundation we use a more flexible and general formal process algebraic model combined with an architectural description of meta-applications using a well defined (and extendable) set of components and connectors.

Other approaches exist which are promoting the general idea of combining an architectural description with a formal behavioral model in the more general context of distributed software design. The ADL Darwin is combined with labeled transition systems in [7] to facilitate a compositional reachability analysis and in [1] CSP is used within the ADL Wright.

To conclude, we have described a formal model supporting component based meta-application development. A formal architecture description language was

combined with a process algebraic behavioral description to define a basic extendable vocabulary of component types for meta-application development. This vocabulary has been implemented on top of the Amica metacomputing infrastructure. Component compositions can automatically be checked based on well known state space analysis methods (e.g. model checking). Next steps include an extension of our vocabulary (including e.g. event propagation and handling) and an application to more complex problems. As we want to check not only functional but also performance properties of meta-applications we plan to use a stochastic process algebra description of component behavior. Also a behavioral model of the Amica infrastructure itself will be integrated in the future.

References

- [1] R. Allen and D. Garlan. A formal basis for architectural connection. *ACM TOSEM*, 6(3):213–249, July 1997.
- [2] T. Fink and S. Kindermann. First steps in metacomputing with Amica. In *Euromicro-PDP 2000*, pages 197–204. IEEE Computer Society, 2000.
- [3] I. Foster and C. Kesselman. Globus: A metacomputing infrastructure toolkit. *The International Journal of Supercomputer Applications and High Performance Computing*, 11(2):115–128, 1997.
- [4] H. Garavel, M. Jorgensen, R. Mateescu, C. Pecheur, M. Sighireanu, and B. Vivien. CADP'97 - status, applications, and perspectives. In *Proceedings of 2nd COST 247 Int. Workshop on Applied Formal Methods in System Design*, 1997.
- [5] D. Garlan, R.T. Monroe, and D. Wile. ACME: An architecture description interchange language. In *Proceedings of CASCON '97*, November 1997.
- [6] D. Garlan and M. Shaw. *Software Architecture: Perspectives on an emerging Discipline*. Prentice Hall, April 1996.
- [7] D. Giannakopoulou, J. Kramer, and S.C. Cheung. Behaviour analysis of distributed systems using Tracta. *Journal of Automated Software Engineering*, 6(1):7–35, January 1999. R. Cleaveland and D. Jackson, Eds.
- [8] A. Grimshaw, A. Ferrari, G. Lindahl, and K. Holcomb. Metasystems. *Communications of the ACM*, 41(11), 1998.
- [9] T. Haupt, E. Akarsu, and G. Fox. Webflow: a framework for web based metacomputing. In *HPCN Europe '99*, April 1999.
- [10] ISO/IEC. Lotos — a formal description technique based on the temporal ordering of observational behaviour. International Standard 8807, ISO — Information Processing Systems — OSI, Genève, September 1988.
- [11] R. Mateescu and H. Garavel. Xtl: A meta-language and tool for temporal logic model-checking. In Tiziana Margaria, editor, *STTT'98 (Denmark)*, July 1998.
- [12] R.P. Mc Cormack, J.E. Koontz, and J. Devaney. Seamless computing with Web-Submit. *Concurrency: Practice and Experience*, 11(15):946–963, 1999.
- [13] M. Shaw and P. Clements. A field guide to boxology: Preliminary classification of architectural styles for software systems. In *Proceedings COMPSAC*, 1997.
- [14] J. Villacis, M. Govindaraju, D. Stern, A. Withaker, F. Berg, P. Deuskar, T. Benjamin, D. Gannon, and R. Bramley. Cat: A high performance, distributed component architecture toolkit for the grid. In *Proceedings of the High Performance Distributed Computing Conference*, 1999.