Dublin Institute of Technology
ARROW@DIT

Articles

Digital Media Centre

2009-01-01

A Generic Approach to Computer-based Clinical Practice Guideling Managment using the ECA Rule Paradigm and Active Databases

Bing Wu
Dublin Institute of Technology, Bing.Wu@dit.ie

Kudakwashe Dube
Dublin Institute of Technology

Follow this and additional works at: http://arrow.dit.ie/dmcart

Part of the Databases and Information Systems Commons, and the Health and Medical Administration Commons

Recommended Citation
Wu, B., Dube, K. (2009) A Generic Approach to Computer-Based Clinical Practice Guideling Managment Using the ECA Rule Paradigm and Active Databases. *International Journal of Technology Management (IJTM)*, Vol 47, Nos 1/2/3, pp. 75 – 95.

This Article is brought to you for free and open access by the Digital Media Centre at ARROW@DIT. It has been accepted for inclusion in Articles by an authorized administrator of ARROW@DIT. For more information, please contact yvonne.desmond@dit.ie, arrow.admin@dit.ie.

This work is licensed under a Creative Commons Attribution-Noncommercial-Share Alike 3.0 License

A generic approach to computer-based Clinical Practice Guideline management using the ECA Rule paradigm and active databases

Kudakwashe Dube* and Bing Wu

Department of Computer Science,
School of Computing, Dublin Institute of Technology,
Kevin Street, Dublin 8, Ireland
E-mail: kudakwashe.dube@dit.ie E-mail: bing.wu@dit.ie
*Corresponding author

Abstract: The increasing demand for reduced cost and improved quality of service in healthcare has prompted the call for better management of medical knowledge. The main emphasis has been on knowledge that is acquired through experience and medical research and then formalised into *Clinical Practice Guidelines* (CPGs). This paper presents a generic approach to CPG information and knowledge management that uses the *Event-Condition-Action* (ECA) rule paradigm and *active databases* within a unified management framework. The paper focuses on an approach for facilitating the use and management of CPGs by clinicians through delivering the CPGs at the point-of-care by a computerised mechanism.

Keywords: clinical guidelines and protocols; event-condition-action; ECA; rule paradigm; active database systems.

Reference to this paper should be made as follows: Dube, K. and Wu, B. (2009) 'A generic approach to computer-based Clinical Practice Guideline management using the ECA Rule paradigm and active databases', *Int. J. Technology Management*, Vol. 47, Nos. 1/2/3, pp.75–95.

Biographical notes: Kudakwashe Dube is a research fellow in the School of Computing at the Dublin Institute of Technology (DIT). He received the PhD Degree from the DIT in 2004. Prior to that, he had completed his Undergraduate studies in Computer Science and Mathematics at the University of Zimbabwe where he obtained his BSc General and BSc Honours degrees. Currently, he holds the Arnold F. Graves Fellowship within the Faculty of Science and works within the Knowledge Camp (K-Camp) Research Group in the Department of Computer Science in the School of Computing at DIT. His research interests include complex information modelling, knowledge management and healthcare informatics.

Bing Wu studied Computer Science at the National University of Defence Technology, China. He did his PhD in Computation at Manchester Institute of Science and Technology, UK, in 1996. He then worked as a research fellow in Computer Science at Trinity College Dublin up to 1998 before joining the DIT. Bing is currently Head of Computer Science in DIT. He is also Head and founder of the K-Camp Research Group. His interests include: knowledge management, discovery and utilisation and applications of advanced database systems, particularly with active, temporal and data mining mechanisms focusing on healthcare and bio-informatics applications.

Copyright © 2009 Inderscience Enterprises Ltd.

1 Introduction

Due to the increasing demand for reduced healthcare costs, minimal clinical practice variation and the optimisation of patient care resource utilisation, calls have been made for the formalisation of medical domain information and knowledge acquired through experience and medical research to create *CPGs*. The US Institute of Medicine has defined a CPG as

> "… a set of systematically developed statements to assist the medical practitioner and the patient in making decisions about appropriate healthcare for specific clinical circumstances." (Field and Lohr, 1992)

Thus, CPGs specify medical knowledge that guides and informs specific activities and interventions that are part of disease management (Gorden et al., 1997). CPGs have also been viewed as "knowledge models of preferred processes of care" (OpenClinical, 2001). Calls have also been made for the incorporation of CPGs into healthcare information systems as an way to avail CPGs at the point-of-care and to positively influence clinician's compliance to these CPGs (Matimer et al., 1992).

ECA rules are specified by an event, a condition and an action whose combined behaviour is such that the event must occur in order for the action to be executed subject to the condition being satisfied (Widom and Ceri, 1996). An *active database* management system is a DBMS that incorporates an ECA rule mechanism and provides ECA rule support facilities that are stipulated in the *Active Database System Manifesto* (Dittrich et al., 1995). Characteristics of ECA rules and their collective behaviour in both relational and object-oriented database systems have been analysed by various researchers in the area of active databases and are now well known (Paton, 1999).

The ECA rule paradigm has been used to support the specification of medical knowledge through the Arden Syntax (Hripscak et al., 1994), which is an established HL7 standard (HL7, 1999). The first application of active databases to computerised CPG management was in HyperCare (Caironi et al., 1997), which uses an active database to implement a hypertension CPG. However, these approaches are lacking in providing generic approaches and methods that can be used to support computerised CPGs within the context of a unified management and implementation framework. This work is part of on-going research work within the Dublin Institute of Technology's School of Computing whose aim is to develop a generic approach to computer-based management of CPG information and knowledge. The ECA rule paradigm is used in formalising CPG knowledge within a unified modelling and implementation framework. This paper aims at presenting our generic and unified framework and method for computer-based CPG management together with a case study, which practices the approach through a prototype implementation system and its use in the management of the Microalbuminuria Protocol (MAP) for diabetes patients.

The rest of this paper is organised as follows: Section 2 presents a review of related work; Section 3 presents a discussion of the representational primitives for specifying CPGs and identifies the ECA rule paradigm as embodying the core representational primitives for CPGs; Section 4 presents the generic approach and the *Specification, Execution and Manipulation* (SpEM) framework for CPG management based on the ECA rule paradigm and active databases; Section 5 presents the method for managing CPGs according to the SpEM framework and approach; Section 6 reports on a case study in which a proof-of-concepts system is developed and used to manage the MAP for the

management of renal complications in diabetes patients; Section 7 presents a discussion of the main issues raised in this paper and future work and Section 8 summarises and concludes this paper.

2 Related work

This section presents a brief review of related work focusing on CPG knowledge management approaches that make use of the ECA rule paradigm. The Arden Syntax is a language for encoding medical knowledge bases that consists of independent modules called the Medical Logic Modules (MLMs). The Arden Syntax is the first approach that made use of the ECA rule paradigm to support computerised medical knowledge management. It is currently the only standard for encoding and sharing medical knowledge between medical institutions (HL7, 1999). This standardisation of the Arden Syntax is an indication of the promise the ECA paradigm has as a pragmatic and viable technology for formalising knowledge in clinical practice. However, since MLMs specifications are stored as individual text files, the Arden Syntax leads to CPGs that can neither be easily queried nor easily manipulated (Jenders et al., 1998). Thus, a limitation of the Arden Syntax is the lack of support for higher-level abstract constructs and modularisation for ECA rules and the support for the manipulation and querying knowledge specification. Hence, the maintenance of the MLMs specifications is difficult.

HyperCare (Caironi et al., 1997) is a prototype system that employs the ECA rule paradigm in the active object-oriented database, Chimera, to capture medical knowledge contained in a hypertension guideline. HyperCare is the first guideline system to use an active database system for guideline management support. HyperCare does not provide a generic protocol specification model and was created specifically to manage a domain- and organisation-specific guideline for essential hypertension. The HyperCare system was designed solely for supporting clinical guideline compliance in the domain of essential hypertension. Other limitations of HyperCare include:

- the difficulty in managing the rules making up the protocol
- the lack of support for dynamic manipulation, querying, versioning and customisation of clinical protocol specifications and instances
- it is an implementation of a specific guideline and does not attempt to provide a generic formalism to support similar protocols.

Up until now, Arden Syntax and HyperCare have been the only approaches that made use of the ECA rule paradigm for managing medical knowledge. The Arden Syntax allows the generic clinical protocols to be specified and executed. Protocol specifications are stored as programming language code in text files. Furthermore, there is no flexible support for the management of both specifications and their instances in the Arden Syntax implementations. HyperCare does not support the creation of generic clinical protocol specifications. Instead, the system was built for a specific clinical protocol, which it implements using ECA rules of an active database system. Both the Arden Syntax and HyperCare do not create patient-specific instances. Instead, all rules in a guideline instance operate at

a global level and have a global scope that covers all patients within a given clinical category.

3 Primitives for clinical guideline representation

A CPG can also be viewed as "a method that identifies actions that are to be performed and that specifies conditions that govern when it is appropriate to perform them" (Pattison-Gordon et al., 1996). Hence, it can be noted that a clinical guideline also includes event monitoring with condition or appropriateness criteria determination. Thus, a CPG embodies the core compositional primitives of the ECA rule paradigm. The recognition of the usefulness of the ECA rule paradigm in supporting the management of information and knowledge in the clinical guideline domain has been confirmed by the development of the Arden.

The core representational primitives for the five main clinical guideline modelling approaches are presented in Table 1. For a more comprehensive review of computerised CPG approaches, the reader is referred elsewhere, especially de Clercq et al. (2004). Four approaches, i.e., GLIF (Ohno-Machado et al., 1998), PROforma (Fox et al., 1998), Asbru (Shahar et al., 1998) and EON (Musen et al., 1996), use control structure primitives (plans, tasks or steps) to represent guidelines. The Arden Syntax (Clayton et al., 1989) is unique in that uses the ECA rule paradigm to model a CPG as an independent modular rule, the MLM. As a result, the Arden Syntax is restricted to modelling simple and highly modular CPGs such as alerts and reminders. From Table 1, it can be noted that primitives that describe conditions and actions are common in all the five approaches while events are directly supported only in the Arden Syntax. The representation of events in the other models is merely implicit except in the EON approach, which explicitly represents a special class of events, termed Adverse_Event.

Besides events, conditions and actions, CPG representations have other primitives. Patient state is modelled only by the EON and GLIF approaches. The Asbru approach stands out from the rest in that it uses knowledge roles, i.e., preferences, intentions, condition and effects, instead of primitives to model guidelines. All approaches support some form of temporal reasoning, with the Arden Syntax and GLIF using three-valued logic (true, false, unknown) to support limited uncertainty.

In CPGs, clinical events are generally detectable happenings that occur to a patient and range from disease progression to what clinicians do to a patient; conditions are checks on patient clinical attributes that are made based on clinical measurements and observations; and actions are clinical interventions that usually follow event occurrences or condition satisfaction or both and can generate events and/or give rise to satisfaction of other conditions. Consequently, it could be concluded that the ECA rule paradigm contains the core representational primitives for CPGs. Although ECA rules, on their own, are not necessarily enough to completely represent all types of CPGs, this work uses them as the core CPG representational primitive for CPGs because they offer a pragmatic way to facilitate immediate incorporation of CPG knowledge and information into existing healthcare systems using facilities already provided with the generally available DBMS technology.

Table 1 CPG representational primitives for the five main clinical guideline modelling approaches

CPG management approach	*Core modelling and representaional primitive*	
Arden Syntax (Clayton et al., 1989)	Medical Logic Module (MLM)	Data slot
		Event slot
		Logic slot (condition)
		Action slot
GLIF (Ohno-Machado et al., 1998)	Guideline steps	Decision
		Patient state
		Branch and synchronisation
		Action
	Medical ontology	
PRO*forma* (Fox et al., 1996)	Task ontology	Root task
		Plan
		Decision
		Action
		Enquiry
Asbru (Shahar et al., 1998)	Preference	
	Intention	
	Condition	
	Effect	
	Plan/action body	
EON (Musen et al., 1996)	Scenario	
	Decision	
	Action	
	Goal	
	Patient data model	Patient
		Qualitative_Entry
		Numeric_Entry
		Adverse_Event
		Condition
		Medication
		Procedure
	Medical-specialty model	

4 Approach and framework to supporting the management of computerised clinical guidelines

This section presents the generic approach and framework based on the ECA rule paradigm and active databases for supporting the management of CPGs. The approach and framework presented in this section addresses the need to:

- manage the clinical guideline knowledge and its execution processes
- consider the clinical situations, which include events and appropriate actions
- to take into account other attributes of the patient mostly contained in the electronic patient record during any intervention.

4.1 The approach to supporting guideline management

The approach adopted here for the management of CPG information and knowledge focuses on supporting monitoring, coordination and generating suggestions and reminders that are relevant to patient care. The underlying principle to this approach is to emphasise the support for monitoring and coordinating clinical interventions while merely providing alerts, reminders and suggestions to the clinicians. The common practice is to make use of advanced AI methods and techniques that strongly emphasise on assisting domain experts with the task of reasoning and/or problem-solving (Miksch, 1999). The approach makes use of a framework that provides three management planes that support specification, execution and manipulation of CPG knowledge and information. The approach also uses the ECA rule paradigm to monitor patient conditions and coordinate clinical interventions as well as suggesting further appropriate clinical interventions such as ordering appropriate clinical laboratory tests whose outcomes are also further monitored. The aim of the approach is to comprehensively support guideline management. Deliberate focus is directed at harnessing the generally available database technology towards providing a tool for assisting domain experts with monitoring and coordinating their activities around the electronic healthcare record.

4.2 The Specification, Execution and Manipulation (SpEM) framework for Clinical Practice Guideline (CPG) management

The SpEM framework for CPG management is made up of the specification, execution, and manipulation planes. Figure 1 illustrates these three planes of the SpEM framework. All aspects of the CPG management process are provided for within the framework. In the specification plane, CPGs are translated into formal specifications which are stored in a suitable form. In the execution plane, the stored CPG specifications are used to create patient-specific CPG instances that are executable by a computer-based execution mechanism. In the manipulation plane, the stored specifications and the executing CPG instances are manipulated using supported operations and queries. It is important to point out here that the active database plays a central role in the SpEM framework. An interesting question relates to the extent to which a modern DBMS with an active rule mechanism can support every process in each of the three planes of the SpEM framework. The next section outlines the role of the active database system in each of the planes of the SpEM framework.

Figure 1 The generic framework, SpEM, for managing domain knowledge and information such as Clinical Practice Guideline (CPGs)

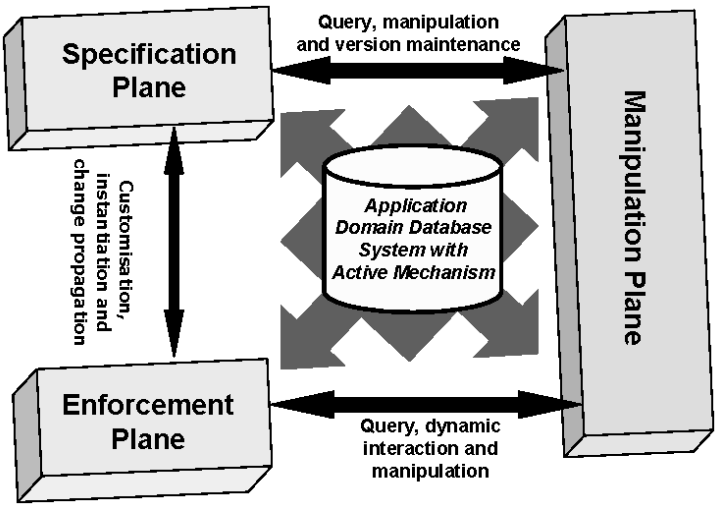

4.3 Existing support for the Specification, Execution and Manipulation (SpEM) framework

Table 2 summarises the literature review findings on the support for the SpEM framework and the computational formalisms employed in major existing CPG approaches and systems. Guideline support approaches and systems for the domain of diagnosis and therapy planning provide advanced and comprehensive modelling concepts, frameworks and computational formalisms. However, these guideline support approaches provide guideline support mainly in terms of the specification and execution of guideline knowledge and not much attention has been directed towards:

- the comprehensive support for the manipulation, i.e., performing operations and issuing queries on aspects of the computerised guideline knowledge and information
- the exploitation of database features for managing information and knowledge.

The manipulation of guideline knowledge and the information about its execution state is important to allow flexibility and the ease-of-use of guideline management support systems. Flexibility and ease-of-use are the major determining factors in the acceptability of guideline systems by clinicians. In terms of the SpEM framework, the guideline systems and models reviewed in this section support mainly the specification and execution planes. With the probable exception of the Asgaard/Asbru guideline system (Shahar et al., 1998), most systems do not provide support for the manipulation plane. Our work, which is illustrated in the shaded row in the table, places emphasis on covering all aspects that make up the SpEM framework.

Table 2 Support for the SpEM framework in existing CPG management systems

		*SpEM framework support key: √ – full support, * – weakly supported, X – no support*				
				Manipulation		*DB feature use*
Guideline/protocol system	*Computational formalism employed*	*Spec*	*Exec*	*Operations*	*Query*	
DILEMMA/PRESTIGE (Gordon et al., 1994; Gordon and Veloso, 1996)	*Network-based*: network of components, state-transition model of action execution	√	√	X	X	X
EON/Dharmma (Musen et al., 1996; Tu and Musen, 2001)	*Hybrid*: network-based core model, Boolean criteria, temporal patterns and selected formalisms for suitable for each task components	√	√	X	X	X
PROforma (Fox et al., 1996)	*Hybrid*: network of plans and procedures, declarative formal logic	√	√	X	X	X
SIEGFRIED (Lobach et al., 1997)	*Hybrid*: Structured and procedural representation with a relational data model	√	√	*	*	√
GLIF (Ohno-Machado et al., 1998)	*Network-based*: flowchart of structured actions and decisions	√	*	X	X	X
Asgaard/Asbru (Shahar et al., 1998)	*Hybrid*: hierarchical skeletal planners with a library of various problem-solving methods	√	√	*	*	X
GUIDE (also Pavia Model) (Quaglini et al., 2000)	*Network-based*: flowcharts based on Petri Nets	√	√	X	X	X
PROGIDY (Johnson et al., 1999)	*Network-based*: augmented transitions of patient states and decisions	√	√	X	X	X
GASTON (de Clercq et al., 2001)	*Hybrid*: frame-based model with flowcharts and production rules	√	√	*	X	X
GLARE (Terenziani et al., 2001)	*Network-based*: a control network of actions and their composites	√	√	X	X	X
SpEM/TOPS (this work)	*Hybrid*: state transition model, Structured and procedural representation using ECA rule paradigm with a active database mechanism and a database model	√	√	√	√	√

4.4 The active database within the Specification, Execution and Manipulation (SpEM) framework

The active database plays a crucial role in the framework and approach by providing the basis for a readily available execution engine for CPGs. Modern relational database systems already support, in a very basic way, the mechanism for monitoring and coordination in the form of the rudimentary active rule mechanism generally referred to as the trigger mechanism. The second role of the active database in the framework and approach is that it provides excellent facilities for manipulating information to support the tasks of monitoring and coordination. The electronic patient record could be held in the database together with guideline information, knowledge and executing processes. This makes available uniform operation and query facilities for both the patient record and guideline information. The third role played by the active database in the framework and approach is that of guaranteeing future sharing of information through the generic nature of databases and the standard language, the SQL. Furthermore, tools already exist to map data from databases to XML for information exchange between systems. Although the implementation of trigger mechanisms in modern DBMS's is currently not standardised, the SQL standard already specify standard SQL trigger specifications, which form the basis for future DBMS compliance.

5 Managing computerised Clinical Practice Guidelines (CPGs) within the Specification, Execution and Manipulation (SpEM) framework

To comprehensively support the management of computerised CPGs, several aspects need to be incorporated as components of the management process. Domain knowledge that exists in the form of expertise and literature on recent advances and discoveries in medical knowledge is the source of CPGs. The translation of this domain knowledge into CPGs is done by clinicians and is outside the scope of this paper. The formal representation of CPGs and the creation of formal specifications and their subsequent storage is an important aspect of the computerisation of CPGs. The instantiation and execution of computerised CPGs with respect to specific individual patient is also a vital component of the management of computerised CPGs. The manipulation of both the formal specifications and the enforcement process consists of the two aspects: querying; and performing manipulation operations on the CPGs.

The process illustrated in Figure 2 allows CPGs to be formally specified, stored, executed with respect to a given patient, and manipulated through querying and operations. This process is comprehensive and will ensure that most aspects of the computerised CPGs are easy to manage. Figure 3 further illustrates the enabling technologies that are used in the method of supporting the management of CPGs. These enabling technologies provide the following features and functionality:

- *The CPG specification model and language*. In order to address the requirements of the Specification Plane, a declarative language, the clinical Protocol LANguage, PLAN, together with its model, were developed (Wu, 1998). PLAN uses the ECA paradigm as the core representation construct for specifying clinical protocols.

  The storage of the PLAN specifications is achieved by the use of the relational database. PLAN is a declarative ECA paradigm-based language. Figure 3 illustrates

the general syntax of PLAN using the Backus- Nuar Form (BNF). A PLAN specification consists of:

- a descriptive header
- a set of schedules
- the protocol rule set.

Just like a protocol, a schedule is named and consists of some *entry criteria* and a *list of rules* each of which is either a dynamic or static rule. For a more detailed discussion of PLAN, the reader is referred to an earlier paper (Wu and Dube, 2001).

- *The CPG execution model and mechanism*. For each patient, the relevant PLAN specification is customised and installed as an instance that executes within the ECA rule mechanism of a database system. The execution of the guideline instance proceeds according to the ECA rule mechanism which monitors events in the patient record held within the database and the time points of interest to the CPG. Thus, the same database where protocol specifications and the patient record are held is also used as the execution engine for the specified CPGs.
- *Providing for CPG manipulation*. Provision is made to perform operations and to issue queries against the CPG specifications and the instance's execution process. A query and manipulation language that is based on the SQL can be used to achieve the purpose of the Manipulation Plane.

Figure 2 The management support process supporting the SpEM framework

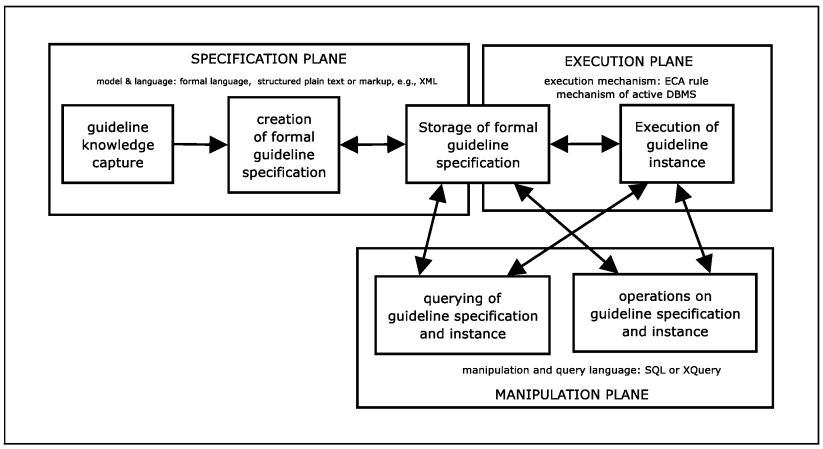

Figure 3 The high-level BNF syntax of the language, PLAN

```
<protocol> ::= PROTOCOL<protocol_body>END PROTOCOL.
<protocol_body> ::= protocol_header>SCHEDULE_SET<schedule_list>END
        SCHEDULE_SET <protocol_rule_set>
<protocol_header> ::= protocol_name>;<description>;<creator>;<category>;
<schedule-list> ::= <schedule> | <schedule>,<schedule-list>
<schedule> ::= SCHEDULE<schedule_body>END SCHEDULE
<schedule_body> ::= schedule_name>;<entry_criteria>;<schedule_rule_list>;
<schedule_rule_list> ::= <schedule_rule>|<schedule_rule>;<schedule_rule_list>
<schedule_rule> ::= <static_rule>|<dynamic_rule>
```

6 Case Study: the proof-of-concepts system and the management of the Microalbuminuria Protocol (MAP)

The next section presents the design of TOPS the prototype system that follows the SpEM framework by providing the functionality that is required in the three planes for the management of CPG knowledge and information. This section also presents a demonstration of the use of the approach presented in this paper to the management of the Microalbuminuria (MA) guideline for diabetes patients.

6.1 TOPS: the proof-of-concepts system

An advanced active database application system, called TOPS, has been developed as a proof-of-concepts implementation of the SpEM framework and approach. A high level architecture of TOPS is illustrated in Figure 4. The architecture of TOPS, as illustrated in Figure 4, has three layers:

- external to TOPS, are users and external systems
- the top layer is the clinical protocol management functionality that allows users to specify, store, execute manipulate and query clinical protocols and external systems to supply and receive information from the system
- the middle layer provides services that extend the ECA rule execution mechanism of the underlying database system handle connections to the database
- the bottom layer is the ECA rule mechanism in a modern database system, which currently uses the Oracle9i active mechanism.

Figure 4 The architecture of the proof-of-concepts system, TOPS, which has been implemented using the Oracle RDBMS

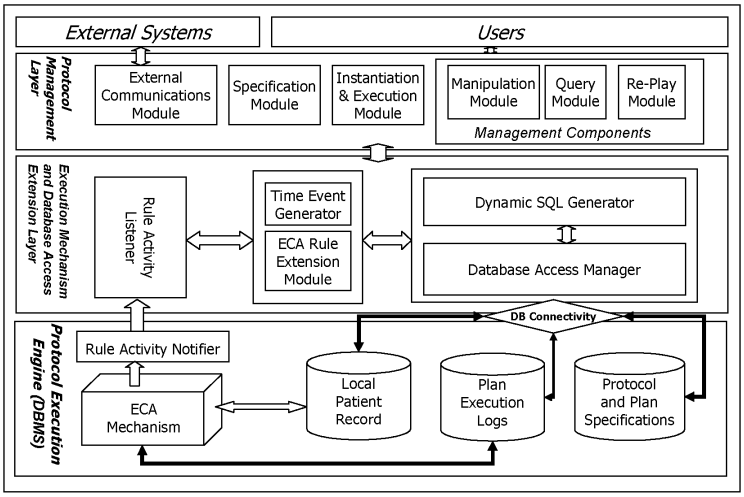

The TOPS implementation architecture in Figure 6 was realised using the Oracle database system and fully utilises the Oracle trigger mechanism. The architecture provides the management functionality to allow operations to be performed and queries

to be issued dynamically at any time during the execution of patient-specific CPG instances. The architecture provides support for the three planes in the SpEM framework. Issues of concurrency and efficiency in rule execution are handled by the trigger mechanism within the Oracle9i DBMS with the exception of ECA rule extensions that are implemented externally.

6.2 Guideline rules and SQL triggers in TOPS

The TOPS protocol execution engine uses the active rule mechanism of the DBMS, the trigger mechanism, and programmatic extensions implemented in Java as its kernel. PLAN specifications are automatically translated to SQL triggers and added to the protocol execution engine. Figure 5 illustrates the process used by TOPS to automatically translate a PLAN rule to one or move SQL triggers for an underlying DBMS. One PLAN rule may be translated to one or more SQL triggers. During the translation process, a PLAN rule is split into its event, condition and action components. Each of these components is translated and formatted to the corresponding SQL trigger code segment, which is then assembled to create code for an SQL trigger. The resulting code for the SQL trigger is DBMS-specific. In the current version of TOPS, the SQL trigger code generated is specific to the Oracle 9i DBMS. SQL trigger support for other DBMS environment could be added through a plugin mechanism. It is possible for a rule component to be translated to a fully independent trigger. The translation of a PLAN rule to one on more SQL triggers may require more information than what a PLAN specification may contain. Hence it may be necessary for the system to prompt the user for more information at some points during the translation process.

Figure 5 The process of translating plan rules into SQL triggers for a target DBMS

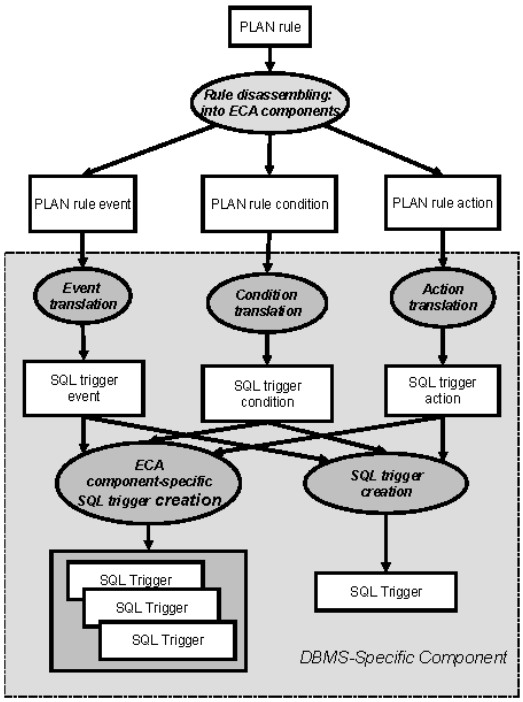

6.3 Managing the Microalbuminuria Protocol (MAP) using TOPS

The framework and approach for CPG management presented in this paper will be demonstrated by using a case study involving the CPG for the diagnosis and management of MA in diabetes mellitus as interpreted by a practicing clinician at the local Diabetes Day Clinic in Dublin.

- *Modelling the microalbuminunia protocol*. Figure 6 illustrates the state chart for the MAP.The renal screening process starts with the annual screening of blood and leucocytes in urine using the Dipstick Urine Test (DUT). If the DUT is positive, i.e., blood and leucocytes are present in urine, then screening for other infections is done before a patient can be referred to a nephrologist. If the DUT is negative, i.e., blood and leucocytes are absent from urine, then the patient is screened for MA, which involves three measurements of urine albumin using the Albumin-Creatinine Ratio (ACR) test over a period of six months. If ACR is less than 20 mg/l at any point, then the patient is cleared of MA and becomes subject to the annual DUT. If ACR is greater than 200 mg/l, then the patient is referred to the nephrologists. If ACR is in the range 20–200 mg/l in two of the three measurements taken over six months, then the patient is diagnosed with MA. This diagnosis is confirming with the 24-hour creatinine clearance and protein loss measurements. If MA is confirmed, then treatment and monitoring of MA commences. At any point during the treatment of MA, the patient is referred to the nephrologist if ACR is greater than 200 mg/l. The patient is also placed on annual screening if ACR drops to less than 20 mg/l. A more detailed version of the state chart illustrated in Figure 6 is used to generate ECA rules that implement the logic of the protocol. For each state and its associated transitions, rules are designed to handle the following:
  - perform what must be done when the patient enters the state
  - perform what must be done during the patient's stay in the state; Perform what must be done when a patient exits from the state
  - monitor the conditions that cause the patient to be moved from one state to another, i.e., conditions for state transitions.
- *Specifying the MAP*. Figure 7 presents the outline structure for the MAP. The schedule and protocol rule sets are designed by following a few simple guidelines that will allow the systematic creation of a CPG specification based on the state chart.

  By applying the guidelines to the rules obtained with the aid of the state chart for the MAP, the specification for the MAP with the outline structure and content presented in Figure 7 is obtained. Due to space limitations, the complete PLAN specification for the MAP is not presented here. The MAP specification, which is expressed in the language PLAN, is parsed by the PLAN parser after which the specification is stored in the database where it can be managed. Once stored in the database; the MAP specification can be queried and manipulated at various levels down to the ECA rule components. The CPG specifications could also be converted to XML for sharing.

Figure 6 State chart for the Microalbuminuria Protocol (MAP)

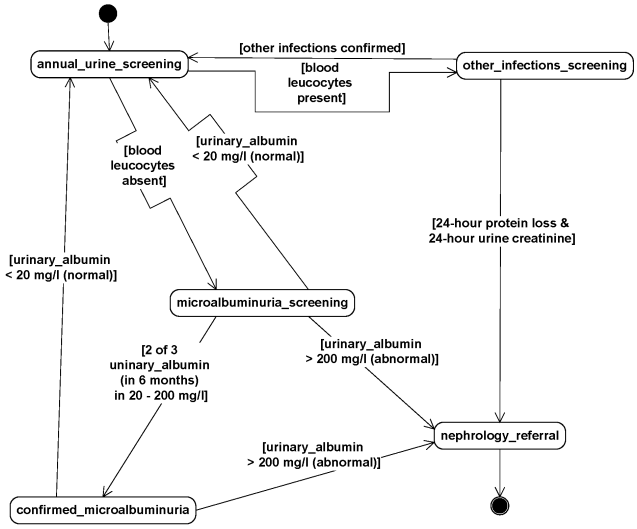

Figure 7 A high-level structure of the specification of the Microalbuminuria Protocol (MAP)

```
@PROTOCOL@ MAP;
DESCRIPTION: This is a protocol for the diagnosis and
management of microalbuminuria in diabetes patients;
CREATOR: DR JOHN NOLAN;
CATEGORY: DIABETIC_NEPHROPATHY;
#SCHEDULE_SET#
^SCHEDULE^ AUS, <AUS_rules> ^END SCHEDULE^
^SCHEDULE^ OIS, <OIS_rules> ^END SCHEDULE^
^SCHEDULE^ MAS, <MAS_rules> ^END SCHEDULE^
^SCHEDULE^ CMA, <CMA_rules> ^END SCHEDULE^
^SCHEDULE^ NPH, <NPH_rules> ^END SCHEDULE^
#END SCHEDULE_SET#
~RULE_SET~
STATIC_RULE AUS1,<body of rule AUS1>;
RULE OIS1,<body of rule OIS1>;
RULE MAS1,<body of rule MAS1>;
RULE CMA1,<body of rule CMA1>;
RULE CMA2,<body of rule CMA2>;
RULE CMA3,<body of rule CMA3>;
RULE CMA4,<body of rule CMA4>;
RULE NPH1, <body of rule NPH1>;
~END RULE_SET~
@END PROTOCOL@
```

- *Customising and instantiating the MAP*. Figure 8(a) illustrates the specification of the ECA rule, MAS5, as it appeared in the PLAN specification of the protocol, MAP. The specification of the rule MAS5 after parsing the MAP specification in TOPS is illustrated in Figure 8(b). The attributes of the rule at this stage are held in a Java object. This rule specification is returned by the toString() method of the PDRule() class, which is illustrated in Figure 8(b). Figure 8(c) illustrates the Oracle database trigger SQL code for the rule, MAS5, generated by TOPS during the creation of a patient plan. This translation of MAS5 to a database trigger is done automatically by TOPS and may involve user input. The trigger has a number of customisations. As can be seen in Figure 8(c), the rule name has been translated from just MAS5 to PL$81$1$MAS5 where 81 is the patient's ID and one is the category ID. This ensures that the rule name is unique within the database, which is a requirement imposed by the Oracle DBMS and useful for the management of patient plans in TOPS.

Figure 8 The process of translating rule MAS5 from the MAP specification to the Oracle SQL trigger in TOPS

```
RULE MAS5,
DESCRIPTION: if ACR > 200 mg/l then
refer patient to nephrologist for
possible proteinuria,
ON: RESULT_ARRIVAL('ACR'),
IF: ACR%RESULT$DATABASE$T_RESULTS >
200%DOUBLE,
DO:
PATIENT_STATE('nephrology_referral'
);
```

(a) Rule MAS5 as it appears in the MAP specification

```
RULE_NAME: MAS5;
DESCRIPTION:
if ACR > 200 mg/l then refer patient to nephrologist
for possible proteinuria;
ON: RESULT_ARRIVAL('ACR');
IF: ACR > 200.0;
DO: PATIENT_STATE('nephrology_referral');
```

(b) Rule MAS5 after its has been parsed into a TOPS rule object (the following is an output of `PDRule.toString()`)

Object DDL of TOPS.PL$81$1$MAS5

```
CREATE OR REPLACE TRIGGER "TOPS"."PL$81$1$MAS5" AFTER
INSERT ON "T_RESULTS" FOR EACH ROW WHEN (NEW.RESULT_VALUE > 200.0 AND NEW.TEST_ID = 9) DECLARE
P_ID NUMBER(10,0);
BEGIN
SELECT PATIENT_ID INTO P_ID FROM T_ORDERED_TESTS WHERE ID = :NEW.ORDER_ID;
IF P_ID = 81 THEN
PATIENT_STATE('PL$81$1$MAS5','ibmt20',81,'nephrology_referral');
END IF;
END;
```

Save As... Close Help

(c) The rule MAS5 as an SQL trigger in the Oracle active database mechanism

The event, result_arrival('ACR'), has been translated into two parts: the first part is the database triggering event INSERT ON T_RESULTS and the second part is the condition, NEW.TEST_ID = 9, where nine is the TOPS ID for the ACR test. Thus the rule is now able to monitor the arrival of ACR results. A further customisation has been done to ensure that the rule performs a change in the state of the specific patient to whom the result belongs. The rule is now more specific than what it was in Figure 8(a). The MAS5 rule action invokes an Oracle stored procedure, PATIENT_STATE(), an Oracle Java call specification, which is an interface to a Java stored procedure within the DBMS.

- *Executing the MAP*. Figure 9 illustrates a TOPS session that shows the execution process for a MAP instance associated with a fictitious patient named Ben Ferguson. Such a session is started when a TOPS user requests the creation of a MAP instance for the specified patient. The MAP specification is retrieved and used to create a MAP instance for the patient. Execution then proceeds in an event-driven manner according to the set of ECA rules making up the MAP instance. The ECA rules monitor the patient record and react appropriately as specified in the MAP.

- *Querying the MAP*. In the query illustrated in Figure 10, a patient-specific protocol instance's snapshot at a given time or interval is retrieved. Thus, the protocol instance snapshot refers to the instance's rule composition and the status of its rules at that time or interval. In executing the above query, TOPS first determines if at least one patient plan snapshot exists within the interval specified in the query.

  If the patient plan snapshot does not exist, the query returns the plan's snapshot at the time this query is being processed. The TOPS query illustrated in the TOPS session illustrated in Figure 11 provides information on what tests where ordered with respect to the specified patient during the given time interval. The query target is the order while the source is the patient. The target condition is a time interval, which means that the orders of interest must first belong to the patient with ID 61

and must fall within this time interval, [2004-7-16 17:48:30, 2004-7-16 17:51:25]. The term order in the query can be generalised to rule-action so that one can obtain information on rule actions that have been performed during the specified time interval.

Figure 9 The TOPS execution log for a Microalbuminuria (MA) CPG instance for a fictitious patient named Ben Ferguson

```
[2004-07-19 12:23:53.565] : TOPS System Execution Log: Session Starting at 2004-07-19 12:23:52.243
[2004-07-19 12:23:53.675] : ----------------------------------------
[2004-07-19 12:24:19.683] : Getting confirmation to create the TOPS database objects.
[2004-07-19 12:24:22.927] : Listener active ...
[2004-07-19 12:24:24.7] : Waiting ...
[2004-07-19 12:27:26.922] : Analysing command: create ...
[2004-07-19 12:27:26.932] : Executing command: create(plan)
[2004-07-19 12:28:49.731] : Retrieving the protocol spec ...
[2004-07-19 12:29:21.777] : Creating plan for patient [ Ben Ferguson, ID: 81 ]
[2004-07-19  12:29:25.082]  :  [start_time:  2004-07-19  12:29:21.887,  end_time:  2004-07-19  12:30:21.887,
interval: 60]
[2004-07-19 12:29:33.234] : [Rule: AUS1; State changed to: READY]
...
[2004-07-19 12:30:46.94] : [Rule: PL$81$1$main$AUS1; State changed to: READY]
[2004-07-19 12:30:49.644] : [Rule: PL$81$1$main$AUS1; State changed to: ACTIVE]
...
 [2004-07-19 12:31:13.338] : [PL$81$1$main$AUS1] executing ...TEST_ORDER (81|DSU,|)
[2004-07-19 12:31:14.569] : [Rule: PL$81$1$main$AUS1; State changed to: FINISHED]
...
[2004-07-19  12:31:25.996]  :  [ORDER_EVENT_TRIGGER]  is  executing  TEST_ORDER_EVENT  (81|81|8|2004-07-19
12:31:23.0|)
...
[2004-07-19 12:31:33.396] : Received |-->[PL$81$1$OIS1%TEST_ORDER*81|UTI,|*%]
...
[2004-07-19 12:31:33.747] : Received |-->[PL$61$1$OIS1%TEST_ORDER*61|UTI,|*%]
...
[2004-07-19 12:31:33.927] : Received |-->[PL$81$1$AUS2%PATIENT_STATE*81,other_infections_screening,*%]
...
```

Figure 10 A query for a snapshot of the composition of a CPG instance in TOPS

```
TOPS:\> query
QUERY:\>    --->SELECT  PLAN  FROM  SNAPSHOT  WHERE
TARGET:2004-7-19 01:55:02,2004-7-19 01:55:58;
No plan snapshot in given interval.
Providing current snapshot, instead:
…
PLAN [ PL$25$1$ ] SNAPSHOT @[2004-07-19 22:30:26.91]
[rule 1]--->[ 72, PL$25$1$AUS2, DYNAMIC, READY ]
[rule 2]--->[ 73, PL$25$1$AUS3, DYNAMIC, READY ]
[rule 3]--->[ 74, PL$25$1$OIS2, DYNAMIC, READY ]
[rule 4]--->[ 75, PL$25$1$OIS3, DYNAMIC, READY ]
[rule 5]--->[ 76, PL$25$1$OIS4, DYNAMIC, READY ]
[rule 6]--->[ 77, PL$25$1$OIS5, DYNAMIC, READY ]
[rule 7]--->[ 78, PL$25$1$MAS2, DYNAMIC, READY ]
[rule 8]--->[ 79, PL$25$1$MAS3, DYNAMIC, READY ]
[rule 9]--->[ 80, PL$25$1$MAS4, DYNAMIC, READY ]
[rule 10]--->[ 81, PL$25$1$MAS5, DYNAMIC, READY ]
[rule 11]--->[ 82, PL$25$1$CMA5, DYNAMIC, READY ]
[rule 12]--->[ 83, PL$25$1$CMA6, DYNAMIC, READY ]
[rule 13]--->[ 84, PL$25$1$OIS1, DYNAMIC, READY ]
[rule 14]--->[ 85, PL$25$1$MAS1a, DYNAMIC, READY ]
[rule 15]--->[ 86, PL$25$1$MAS1b, DYNAMIC, READY ]
[rule 16]--->[ 87, PL$25$1$CMA1, DYNAMIC, READY ]
[rule 17]--->[ 88, PL$25$1$CMA2, DYNAMIC, READY ]
[rule 18]--->[ 89, PL$25$1$CMA3, DYNAMIC, READY ]
[rule 19]--->[ 90, PL$25$1$CMA4a, DYNAMIC, READY ]
[rule 20]--->[ 91, PL$25$1$CMA4b, DYNAMIC, READY ]
[rule 21]--->[ 92, PL$25$1$NPH1, DYNAMIC, READY ]
[rule   22]--->[   93,   PL$25$1$main$AUS1,   STATIC,
EXECUTING ]
END SNAPSHOT FOR PLAN PL$25$1$.

QUERY:\>  --->exit

TOPS:\>
```

Figure 11 A query to find out what clinical test orders where made during the specified time interval

```
TOPS started.
Initialising ...
Creating log files ...
…
TOPS Command Line Facility, version 1.0.
Copyright(C) 2000-2003, K. Dube
Computer Science Department, School of Computing,
DIT, Ireland.
Started at: Wed Jul 21 05:28:13 BST 2004
TOPS:\>QUERY
Executing command: QUERY()
command code: 13
QUERY:\>     --->SELECT ORDER FOR PATIENT WHERE
TARGET:2004-7-16     17:48:30,2004-7-16     17:51:25;
SOURCE:PATIENT_ID=61
processing query ...
launching specialised query handler ...
processing ORDER query [ 2004-7-16 17:48:30,2004-7-
16 17:51:25 ] for [ PATIENT ] ...

Tests ordered for [PATIENT_ID=61] during time
interval [2004-7-16 17:48:30,2004-7-16 17:51:25]
Dip_stick_urine Profile, DSU, 2004-07-16 17:49:28.0
Urinary_Tract_Infection Profile, UTI, 2004-07-16
17:50:06.0
Urinary_Tract_Infection Profile, UTI, 2004-07-16
17:50:06.0
-------------------
End test listing.

QUERY:\>  --->EXIT
TOPS:\>
```

6.4 Case study review and conclusion

In this case study, the use of the highly intuitive state chart makes it easy to communicate with domain experts during CPG information and knowledge elicitation, capture and specification. The use of the UML state chart also makes the subsequent extraction of the relevant ECA rules easier since the state chart naturally supports the ECA rule paradigm (Calestam, 1999) and is easily understood by domain experts. The TOPS protocol specification parser, which uses an object-based mapping between the PLAN specification and the underlying relational database for storing protocol specifications, proved to be efficient and effective as a simple tool for mapping the ECA rule-based protocol specifications into the database. The MAP specification was stored in the Oracle relational database. A single protocol specification in the database consisted of components that were spread over several relations or tables. This offered a simple way to visualise specification information using the familiar tabular format. The relational database model was found to offer a uniform and flexible way to access, manipulate and query all information from specification, to executing process state, to data in the patient record. Flexibility was guaranteed by the SQL, which allows queries that combine data on attributes from several entities subject to constraints within the database.

The generation of SQL trigger code that implement the ECA rules of the MAP was automatically supported by TOPS and required no user intervention. This makes it easy for application domain experts to use TOPS with no knowledge of the SQL trigger specification language. However, domain experts still needed to be familiar with the protocol specification language, PLAN, which should ideally be closer to their domain language than the SQL.

The execution of the rule actions is subject to the availability of the appropriate software module that implements the action. In other words, rule actions in the MAP needed to be predefined and any new action required by the protocol requires that the module to implement such an action be developed. However, the rule actions could be designed such that they were generic and re-usable by other rules in possibly different protocols.

This case study has demonstrated the applicability and effectiveness of the SpEM framework and the ECA rule paradigm in enabling the support for the management of the MAP for diabetes patients. It has been shown that the MAP knowledge can be modelled and specified by using the ECA rule paradigm guided by the state chart. The functionality provided by the three management planes then made available by using a system such as the prototype system, TOPS, for application to the protocol. The specification language, PLAN, was used to specify the resulting protocol specification. Once a PLAN specification is obtained, it is stored in the database for effective management. This makes it possible to execute, perform operations and query various aspects of the MAP using a manipulation language based on the SQL.

7 Discussion and future directions

The work presented in this paper differs from research works in literature in two major respects. First, instead of following the purely decision-support approach, it follows the ECA rule paradigm and uses a generic management framework and an associated specification language for supporting computerised clinical guidelines. The central focus of the chosen approach is on a pragmatic way to integrate guidelines with the electronic healthcare record and support both process-oriented as well as highly modular and fine-grained computerised guidelines that detect errors in test orders, abnormal test orders and results, and issue alerts, reminders and pagers based on information in the electronic healthcare record.

Second, emphasis is placed not only on the creation and execution, but also on the management of the computerised guidelines for each clinical category of patients on a higher level, and the instances of guidelines for each individual patient on a lower level. A clear line has been drawn between the high level static and generic aspects of a guideline for each category of patients, and the lower level dynamic and patient-specific guideline that is instantiated and executed for an individual patient. Most other works are mainly concentrated on developing systems that support specification and execution of the computerised guideline with little emphasis on support for the management, i.e., manipulation, query and replay, of guideline information and knowledge and easy integration with the electronic healthcare record. As a consequence of this and other factors, guideline systems are not in widespread use in practice despite having over a decade of research results.

Our future work will focus on improving our framework though investigating enriched specification, execution and query models as well as making PLAN language more expressive and XML-enabled. Investigations will also look into more efficient methods of exploiting and enhancing the ECA mechanism as well as investigating hybridising it with other paradigms. We have also started investigating useful novel concepts and methods for information scenario manipulation, querying and replay within the context of the SpEM framework. These investigations will enrich the on-going improvements on the prototype system, TOPS, for which a more user-friendly interface will be developed. Future work will also include the deployment of the improved prototype system into the real-life setting within the patient care domain where clinicians will be able to evaluate it more comprehensively. Deployment currently faces challenges associated with patient privacy, confidentiality and security concerns within the patient care practice setting as well as licensing restrictions on access to the database schema and API for the existing hospital systems.

8 Summary and conclusion

This paper has presented a unified and generic framework, SpEM, consisting of the three planes for Specification, Execution and Manipulation for the management of knowledge and information for supporting computerised CPGs. The paper has also presented a model and language, PLAN, for guideline specification and the prototype system, TOPS, which implements the SpEM framework through a generic method that harnesses the ECA rule paradigm and the active database to provide support for all aspects of the SpEM framework. The overall approach exploits the active database, which combines the ECA rule paradigm with the data management functionality of a DBMS to present a promising environment for supporting CPGs and their integration with the electronic healthcare record and clinical workflow. The paper has also presented a case study that demonstrated the technical feasibility of the framework and approach by computerising a real protocol for managing MA in diabetes patients.

Most other research works on managing computerised CPGs have tended to focus only on their specification and execution without providing for comprehensive manipulation, which would improve on the flexibility of resulting CPG systems. This work contributes a generic approach with the framework, SpEM, that unify the core CPG management dimensions, which include manipulation (operations, query, replay, etc.) and an active database method for computational support. The benefits of the approach are:

- the flexibility based on the information management dimension of the manipulation plane within SpEM framework
- the ease-of-integration of CPGs with electronic healthcare records and clinical workflows due to the exploitation of the active database features of data management and ECA rule execution
- the ease-of-incorporation of CPG management system into the healthcare systems due to ubiquity of database systems within most institutions.

Acknowledgements

The authors would like to thank the Arnold F Graves Trustees and the Office of Postgraduate Studies and Research of the Dublin Institute of Technology who are sponsoring this work.

References

Caironi, P.V.C., Portoni, L., Combi, C., Pinciroli, F. and Ceri, S. (1997) 'HyperCare: a prototype of an active database for compliance with essential hypertension therapy guidelines', *AMIA Ann Fall Symposium*, Hanley and Belfus, Philadelphia, PA, pp.288–292.

Calestam, B. (1999) 'OMT-A: an extension of OMT to model active rules', *Computer Science*, University of Skovde, Skovde, p.133.

Clayton, P.D., Pryor, T.A., Wgertz, O.B. and Hripcsak, G. (1989) 'Issues and structures for sharing medical knowledge among decision-making systems: the 1989 Arden homestead retreat', *Annu. Symp. Comput. Appl. Med. Care*, pp.116–121.

de Clercq, P.A., Blom, J.A., Korsten, H.H. and Hasman, A. (2004) 'Approaches for creating computer-interpretable guidelines that facilitate decision support', *Artif. Intell. Med.*, Vol. 31, No. 1, pp.1–27.

de Clercq, P.A., Hasman, A., Blom, J.A. and Korsten, H.H.M. (2001) 'Design and implementation of a framework to support the development of clinical guidelines', *Int. J. Med. Inf.*, Vol. 64, Nos. 2–3, pp.285–318.

Dittrich, K.R., Gatziu, S. and Geppert, A. (1995) 'The active database management system manifesto: a rulebase of ADBMS features', *2nd Workshop on Rules in Databases (RIDS)*, Springer, Athens, Greece.

Field, M.J. and Lohr, K.N. (1992) *Guidelines for Clinical Practice: From Development to Use*, National Academy Press, Washington DC.

Fox, J., Johns, N., Rahmanzadeh, A. and Thomson, R. (1996) *PROforma: A Method and Language for Specifying Clinical Guidelines and Protocols*, Medical Informatics Europe, Amsterdam.

Fox, J., Johns, N. and Rahmanzadeh, A. (1998) 'Disseminating medical knowledge: the PROforma approach', *Artificial Intelligence in Medicine*, Vol.14, No. 1, pp.157–182.

Gordon, C. and Veloso, M. (1996) 'The PRESTIGE project: implementing guidelines in healthcare', *Medical Informatics Europe 96*, IOS Press, pp.887–891.

Gordon, C., Herbert, I., Johnson, P., Nicklin, P., Pitty, D. and Reeves, P. (1997) 'Telematics for clinical guidelines: a conceptual modelling approach', *Studies in Health Technology and Informatics*, Vol. 43, Pt A, pp.314–318.

Gordon, C., Jackson-Smale, A. and Thomson, R. (1994) 'DILEMMA: logic engineering in primary care, shared care and oncology (AIM Project A2005)', *Computer Methods and Programs in Biomedicine*, Vol. 45, pp.37–39.

HL7 (1999) *Arden Syntax for Medical Logic Modules*, Standards of the Health Level 7, USA.

Hripscak, G., Luderman, P., Pryor, T.A., Wigertz, O.B. and Clayton, P.D. (1994) 'Rationale for the Arden Syntax', *Computers and Biomedical Research*, Vol. 27, pp.291–324.

Jenders, R.A., Huang, H., Hripcsak, G. and Clayton, P.D. (1998) 'Evolution of a knowledge base for a clinical decision support system encoded in the Arden Syntax', *AMIA Symp.*, pp.558–562.

Johnson, P., Tu, S.W., Booth, N., Sugden, B. and Purves, I. (1999) *A Guideline Model for Chronic Disease Management in Primary Care*, Report SMI-1999-790, Stanford Medical Informatics, Stanford University School of Medicine.

Lobach, D.F., Gadd, C.S. and Hales, J.W. (1997) 'Structuring clinical practice guidelines in a relational database model for decision support on the internet', *AMIA Annual Fall Symposium*, pp.158–162.

Matimer, D., McCauley, B., Nightingale, P., Ryan, M., Peters, M. and Neuberger, J. (1992) 'Computerised protocols for laboratory investigation and their effect on use of medical time and resources', *Journal of Clinical Pathology*, Vol. 45, pp.572–574.

Miksch, S. (1999) 'Plan management in the medical domain', *AI Communications*, Vol. 12, No. 4, pp.209–235.

Musen, M.A., Tu, S.W., Das, A.K. and Shahar, Y. (1996) 'EON: a component-based approach to automation of protocol-directed therapy', *JAMIA*, Vol. 3, No. 6. pp.367–388.

Ohno-Machado, L., Gennari, J.H., Murphy, S., Jain, N.H., Tu, S.W., Oliver, D.E., Pattison-Gordon, E., Greenes, R.A., Shortliffe, E.H. and Barnett, G.O. (1998) 'The GuideLine interchange format: a model for representing guidelines', *JAMIA*, Vol. 5, No. 4, pp.357–372.

OpenClinical (2001) *The Medical Knowledge Crisis and its Solution through Knowledge Management*, OpenClinical.org.

Paton, N.W. (Ed.) (1999) 'Active rules in database systems', *Monographs in Computer Science*, Springer-Verlag, New York.

Pattison-Gordon, E., Cimino, J.J., Hripcsak, G., Tu, S.W., Gennari, J.H., Jain, N.L. and Greenes, R.A. (1996) *Requirements of a Sharable Guideline Representation for Computer Applications*, Stanford University, Stanford, California, USA.

Quaglini, S., Stefanelli, M., Cavallini, A., Micieli, G., Fassino, C. and Mossa, C. (2000) 'Guideline-based careflow systems', *Artif. Intell. Med.*, Vol. 20, No. 1, pp.5–22.

Shahar, Y., Miksch, S. and Johnson, P. (1998) 'The Asgaard project: a task-specific framework for the application and critiquing of time-oriented clinical guidelines', *Artificial Intelligence in Medicine*, Vol. 14, pp.29–51.

Terenziani, P., Molino, G. and Torchio, M. (2001) 'A modular approach for representing and executing clinical guidelines', *Artif. Intell. Med.*, Vol. 23, No. 3, pp.249–276.

Tu, S.W. and Musen, M.A. (2001) 'Representation formalisms and computational methods for modeling guideline-based patient care', in Heller, B., Loffler, M., Musen, M. and Stefanelli, M. (Eds.): *Computer-Based Support for Clinical Guidelines and Protocols-Proceedings of EWGLP 2000*, IOS Press, Amsterdam, Vol. 83, pp.115–132.

Widom, J. and Ceri, S. (Eds.) (1996) *Active Database Systems: Triggers and Rules for Advanced Database Processing*, Morgan Kaufmann, San Francisco, USA.

Wu, B. (1998) *PLAN: The Framework and its Language for Modelling Clinical Test Request Protocols*, Department of Computer Science, DIT, Dublin.

Wu, B. and Dube, K. (2001) 'PLAN: a framework and specification language with an event-condition-action (ECA) mechanism for clinical test request protocols', *34th Hawaii International Conference on System Sciences (HICSS-34): The Mini-Track in Information Technology in Healthcare*, Maui, Hawaii, IEEE Computer Society, Los Alamitos, California.