



Application Note: Application of KEEL Technology in Clinical Utility Projects (Interpreting Medical Data) (9/23/2003)

Objective:

Analytical validity

Description from notes from Dr. Kirk Hogan (Medical College of Wisconsin - Madison)

Analytical validity stands for the capacity of the test method to detect the presence or absence of a specific DNA sequence variation with precision, and it is supported by data derived from a comparison of the ability of different methodologies to detect the identical alteration. Equivalence is ideally measured in reference to a "gold standard", which itself may be lacking or deficient with rapidly emergent technologies. As described below, oversight of the analytical validity of genetic tests falls well within traditional models of FDA regulation.

Clinical validity

Description from notes from Dr. Kirk Hogan (Medical College of Wisconsin - Madison)

The clinical validity of a genetic test is a measure of its ability to predict a trait, that is, the presence or absence of a clinical condition, predisposition (e.g. drug response), disease or syndrome. A clinically valid correlation between DNA sequence variation and an observed trait may be probabilistic by virtue of near chromosomal proximity of a marker to a causal mutation, or the DNA polymorphism itself may be causal supported by familial and biologic data. Inherent in its name, clinical validity as a measure of how well the test performs in a clinical environment, is both patient- and context-specific, and dependent on such variables as co-existing diseases, intercurrent medications, and personal past medical history for accurate interpretation.

For these reasons, assessment of clinical validity falls away from meaningful FDA scrutiny as a factor intrinsic to the test, relying instead on caregiver competence beyond FDA authority and capacity to regulate. Moreover, consensus on the definition of a trait is often hard to reach and subject to drift over the passage of time.

Clinical utility

Description from notes from Dr. Kirk Hogan (Medical College of Wisconsin - Madison)

The clinical utility of a genetic test is addressed in asking: "Does knowing the genotype have any value to the patient and caregiver as evidenced by improved patient safety and cost-effective care?" Thus, both clinical utility and clinical validity are aspects of test interpretation and indication for use.



SACGT itself was confused on this score stating: "An important distinction in considering the risks and potential benefits of a test is that between the technical aspects of a given test - that is, its clinical validity and utility - versus how it is interpreted by health care providers and the individuals undergoing testing." To the contrary, clinical validity and utility are not intrinsic "technical aspects of a given test" but reflections of the health care provider's interpretation dependent on diagnostic and therapeutic data extrinsic to the test's content.

"A clinically valid test in the hands of a poorly trained health care provider can pose as much risk as a less valid or accurate test that is correctly interpreted."

KEEL technology provides a mechanism for a formal approach for interpreting complex relationships between information. In this manner it provides support for Clinical Utility or a standardized way to interpret medical information.

KEEL technology is not "artificial intelligence". It is just human intelligence packaged in a KEEL technology envelope.

KEEL technology fits in the category of expert systems. However, its rules are not brittle, text-based rules with hard-coupled probabilities that have been developed with inconsistent tests on disconnected sample sets.

KEEL technology incorporates the judgmental values from an expert panel that "explains" relationships between loosely coupled pieces of information.

Because KEEL based systems create explainable decisions, they can be audited and corrected when new information is identified.

The KEEL development model provides an additional benefit, beyond its reasoning capability. It provides a new approach to documenting the diagnostic decision-making model.

The ability to graphically document the relationships between symptoms and other information, allows the model to be created and tested without long delays between concept and evaluation.

Changes or corrections can be entered immediately.

Another advantage of a KEEL based system is its ability to not only explain a decision, but its ability to explain why alternative decisions were not made.

This ability can be helpful in testing and auditing a completed design.



The ultimate goal of any medical analysis is to provide the best care for the patient. A tool embedded with KEEL technology can provide that care in a consistent manner.

When new data is needed to advance the decision-making model, it can easily be incorporated into the design. It can immediately be tested using the features built into the KEEL toolkit.

Since KEEL technology can be deployed in devices, in software applications and in web pages, there is no limit as to how the knowledge is distributed.

Sensor data for gathering symptoms can be locally connected or be delivered over the internet. When the data is ready to be processed the analysis can begin.

Market:

With the recent decoding of the human genome, the amount of data that is /or will be available to the doctor will be overwhelming.

Doctors are already being deluged with data as is indicated by the growth in the PACS (Picture Archiving and Communications System) industry.

Hospital networks are being overloaded as a patient's records are moved around the hospital.

Considering the amount of data that is collected, it is no wonder that key pieces of information are overlooked and mistakes are made.

Medical liability is beginning to drive doctors away from the industry.

A new approach is needed to respond to the increasing amounts of data that a doctor has at its disposal in diagnosing disease.

KEEL technology has the potential to cause a paradigm shift in interpreting diagnostic symptoms in the hospital.

The first step will be to offer a second opinion to the doctor.

Later steps may include embedding KEEL engines in medical equipment to assist the doctors in interpreting the results.

KEEL based equipment that incorporate 1. the best medical interpretation, and 2. the skills of the best operator as defined by the equipment manufacturer will change the industry.

Business Drivers

Reduced Engineering Design Cost & Time



Reduced Start-up/Commissioning Cost & Time

Optimization of Product Performance/Life to the Task

Increased Differentiation of Their Products

Reduced Support Cost & Improved Global Support Responsiveness

Easier Compliance to Standards (Certifications)

Reduced Project Risk

Increased Market Coverage

Improved Internal Efficiency - Reduced Fixed Expense

Improved Access to Information

Buying Factors

Ease of Use

Time to Market

Flexibility

Extensibility / Scalability

Desire for Special Features

Performance

Options

Complex Systems

Enterprise Connectivity

Trends

Product Availability

Quality

Prestige

The desire to Learn



Lowest Product Cost

Lowest Installed System Cost

Lowest Lifecycle System Cost

Safety

By-Products

Expectations of the Future

Efficiency

Reduce or Eliminate Other Costs

Form Factor

Explainable / Auditable

Considerations:

The schedule for evaluating the data

When complex cognitive systems are being created, the designer needs to consider the time sensitivity of input and output data.

In human systems, humans may try to postpone decisions until all the relevant data is available. When pressed, however, humans can react to demands that drive them to make instant decisions. In a KEEL system, the designer may need to have the system "adapt" to changing demands.

The flexibility of the KEEL system design allows for normal conservative operation and still adapts to emergency situations.

The system designer needs to consider when the data should be evaluated. Options could be to trigger the evaluation on a scheduled basis; it could be polled by a higher level authority; it could operate on a change of state from one of its sensors; or it could be continuously running.

Human interaction with the solution

In clinical utility applications, the interface between doctor and KEEL engine will be important. In its initial use, KEEL based systems will augment the doctor's diagnosis as a "second opinion". There is likely to be some resistance to having a computer make subjective diagnoses and a "friendly" interaction will be required.



In cases where KEEL engines participate with humans as part of an overall system, consideration should be given to the human interface characteristics of the system.

KEEL engines can trigger the dialog, or the dialog can be initiated by the human.

KEEL engines can provide continuous analog output which can be translated into human readable form or as stimulation to HMI tools.

KEEL engines can accept input from any source, as long as it is translated to the normalized format required by the KEEL engines.

The system designer will consider how and when the dialog between the human and the KEEL engine will be scheduled.

The ability to create complex cognitive solutions in a timely manner
Especially in the early days where KEEL based systems are being evaluated, it will be important to show the ease at which complex solutions can be developed.

Using the KEEL toolkit, it is possible to develop complex cognitive solutions in a short timeframe.

Software features are built into the toolset that allow the cognitive design to be structured in a manner that makes it easy to integrate the KEEL engine(s) into a variety of architectural models. In many cases, these tools allow the glue logic to be created once and be able to support changes to the cognitive model.

The KEEL-FBD tool allows the staged development of complex models in a manner that automatically creates the glue logic. This is another feature that supports rapid development of complex systems.

Architecture Independence

KEEL engines are discrete information processing engines.

In many cases they will be part of a larger system.

Just like humans in a factory, in an army, or on a team, there is interaction between individuals and between devices and equipment. Computer protocols and network architectures have been developed in a manner that somewhat duplicates the development of human language.



People developed different languages to satisfy different needs: voice for people to communicate verbally; sign language for those that cannot hear; telephone and video conferencing for distance communication.

Computer based protocols have been developed to satisfy different needs: ASCII for simplicity; XML for structure; asynchronous / synchronous / isochronous for flexibility, speed, and time determinism.

KEEL engines can participate with any of these linkages, as long as the data can be normalized to meet the needs of the KEEL engines.

KEEL engines can sit at any point in any architecture. In this way KEEL engines can be integrated into almost any architecture. The simple API enables this valuable attribute.

The "importance" of information

In Clinical Utility applications the importance of information will be a critical attribute of the system. Since KEEL is a rule based system, the importance of information will be heavily discussed.

Over time this information will be tuned to give more accurate results.

A key aspect of any KEEL based system is the dynamically changing importance of information.

In simple systems, there may be no changing importance of information. These types of systems might be built with hard coded solutions or discrete logic.

Cognitive solutions, however, usually deal with complex relationships where information is being interpreted in different ways; in different parts of the same problem.

Cognitive problems often deal with both strategic (future) and tactical (now) problems at the same time. Questions about what to do now and not destroy future opportunities are often addressed. This requires a balancing of information to obtain the best outcome.

Diagnostics and prognostics require that a system adapts its operation to performance variables. If you have difficulty breathing, then breathing becomes more important.

Performance of the system

Different applications require different levels of performance for the KEEL engines.



Some real time control applications may require that information is constantly being evaluated and operated upon.

There will be other times when information changes relatively slowly and therefore does not merit constant evaluation.

There are a variety of techniques to balance performance and system complexity when creating KEEL based solutions.

The evolution of the KEEL engine over time

In Clinical Utility applications, there is limited formal structure for defining the decision making model. KEEL will force a more structured process. As each solution is refined and as new genetic information is integrated into each solution, they will evolve. While this will require constant monitoring, the KEEL Toolkit structure makes this a possible task.

KEEL engines are often developed over time. New pieces of information are added to the design when it is apparent that they contribute to the interpretation of the data provided.

Compsim's KEEL Toolkit provides a number of services to enable enhancements to the engine over time without impacting the glue logic. This saves software development effort.

The KEEL FBD tools allow a complex system to be developed in stages and integrated and tuned as separate components.

The KEEL toolkit incorporates the idea of merge-able objects or decision-making modules.

There are other cases where the model doesn't need to expand, but only needs to be tuned when relationships between information change.

KEEL engines can be created as "classes" in some languages.

The value of "explainable actions"

Because Clinical Utility may be under the direction of the Center for Disease Control (or another government agency), the creation of explainable actions is mandatory. Only with explainable actions, can the process be improved over time.

Many systems can benefit from an engine that creates explainable decisions and actions.

This allows the systems to be audited for their performance in order to tune them over time.



It is possible to have KEEL engines monitor other KEEL engines and potentially provide a feedback mechanism to achieve optimal operation.

There are other cases where there are demands for the creation of explainable actions. This insures that a code of ethics is integrated into the design. Without the ability to decompose every action and every decision, there is the ability to create a system that does not meet the needs of society.

So, KEEL based systems, because they are rule based, can explain why any action was taken or decision was made.

Even though the rules are defined graphically, by inserting a snapshot of the inputs back into the design, the reasoning can be displayed.

Who and how the KEEL actions are monitored

In Clinical Utility applications, an audit committee will be responsible for reviewing the decisions and explanations of the KEEL engines.

KEEL based systems offer several methods for monitoring their performance.

First, because all KEEL actions are visible and explainable in the development environment, they are available for analysis within the tool environment.

Second, because a KEEL engine is a rule based system, it will always respond the same way to an input. An XML schema exists that defines the format for an XML file produced by the real-world device or software applications. If the device or software application logs the input data in this format, it can be read back into the KEEL Toolkit where the reasoning can be reviewed and explained.

The system designer has the responsibility for determining when and how the data is logged and reviewed.

Certain applications may demand more auditing to insure that performance is satisfactory.

It is also likely that new data will enter the application space. This may trigger reviews of performance.

As in human systems, novice operators or players may require closer review than experienced operators or players. The same is true with a KEEL engine.



Consideration for how the KEEL engine(s) fits in the "chain of command"

At least initially, KEEL engines will be observers offering a second opinion and potentially explaining why alternate diagnoses are not applicable.

In some cases KEEL engines will participate as components of a larger system.

They can perform administrative roles by interpreting information in a consistent manner and responding with consistent command decisions.

They can perform subservient roles by accepting direction from above and adapting the commands to modify their actions.

They can sit in the middle of a chain, by accepting commands from above and reviewing status from below. They can modify their own strategy according to the rules provided to them and the information they observe on their own.

They can make requests to humans and devices above them in the chain-of-command, and can deliver commands to humans and devices below. They can collaborate with their peers according to the rules that dictate responsibility.

Should it be appropriate, the KEEL engines can develop their own levels of trust in collaborative environments. The system designer will determine the communication protocols and the flexibility of the system.

Normalizing the data

KEEL engines operate on normalized data (0 to 100) values that can be either integer or floating point, as defined by the KEEL project.

While this suggests a linear range of normalized input data, the inputs can drive a curve which allows the data to be interpreted in almost any way.

In this manner a single normalized input value can be interpreted according to any number of independent curve relationships.

Much of the development work in architecting a KEEL solution is spent defining the relationships between information. Because this is all done graphically, there is no need to write "code" to see the results of the analysis.

The level of "trust" attached to input information



In Clinical Utility cases, trust may not be a primary factor. However, there may be cases where trust may be associated with the timeliness of the data.

Many judgmental decisions are made by including a level of trust to validate the information. When the level of trust is diminished, then the information may be given a lower level of importance in the overall solution.

KEEL engines can include the level of "trust" as an input to the system. How this is interpreted, is the responsibility of the system designer.

The concept of risk associated with the decisions and actions associated with the system

In Clinical Utility applications, risk can be assigned as a negative factor when choosing between several alternative treatments. It can also be assigned to impact quality of life.

Many cognitive decisions need to incorporate risk into the decision-making model.

Risk can be included as an input to a KEEL engine.

How risk participates in the decision-making model, is the responsibility of the system designer.

The responsibility for the overall system remains with the solution architect

The system architect is responsible for the overall architecture of the system.

This will include segmentation of the system, determining when and how the KEEL engines will be scheduled.

The cognitive model for interpreting the input data and causing decisions and actions to be promoted is also the responsibility of the system designer.

How time and space relationships might contribute to the solution

Time and space often impact the importance of information when making cognitive decisions. KEEL supports these concepts with its "clipper" features.

This allows decisions and actions to be tuned for different times and locations.



In cases where an optimal solution is being targeted, such as the time to send a message or the time to shoot at a moving target, then tools to support these kinds of decisions are required. They are built into the KEEL toolkit and available to the system designer.

System architecture

The system architecture is the definition of the relationships between all system components.

The cognitive segment is usually just part of the system. The system architecture defines the layout for performance, flexibility, extensibility, cost, resources, etc.

The system architecture is often the result of a balancing act: balancing time to market, resources, performance, and cost. The system architecture is driven with an objective where the features are defined. The objective is addressed by identifying potential solutions: collections of components and methods of tying them together.

KEEL technology can be integrated into the architecture from the beginning, or it can be an "add-on". Because individual KEEL engines are architecture neutral, they can be integrated into an overall architecture at different times; even after a program is completed.

The potential for autonomous operation

In Clinical Utility cases, autonomous operation will probably not be the primary goal; at least not initially. However, as the technology is integrated into medical equipment, this may become a primary goal. There may be cases where there is a lack of medical talent where autonomous operation would be better than no support at all.

Because KEEL engines can interpret information in a human-like manner, and direct relative actions to be taken based on that interpretation, KEEL based systems have the potential to operate without human intervention.

Alternatively, KEEL based systems can operate as either backups to human operators, giving advice or recommendations, or they can provide the primary decision-making engines that are backed up by humans.

The outputs from the system

The outputs from a KEEL engine are normalized values between 0 and 100. This information may have to be transformed into other forms for use by the external controls or monitoring equipment.

KEEL could generate control signals.



KEEL engines interpret information and provide outputs that represent a balancing of the inputs. These values can be used to generate complex commands in the form of control signals to other equipment.

KEEL could generate information for other KEEL engines.

A common practice is to segment a system into multiple KEEL engines. It is likely that one KEEL engine will provide data to the input(s) of other KEEL engines.

The KEEL FBD tool assists in connecting KEEL engines into a single solution.

KEEL engines could also be distributed across a network or in multiple tasks where messaging or data sharing could provide the mechanism for one KEEL engine to feed others.

KEEL could provide inputs to other systems (non-KEEL).

The outputs from a KEEL engine could be supplied to other non-KEEL subsystems for further processing.

KEEL outputs could be part of a local feedback loop.

KEEL engines can be part of a feedback loop where the output of the system is connect back to the input through some other circuitry.

KEEL outputs could be part of a distributed feedback loop.

KEEL outputs could be fed to other external devices which, in turn, feed data back to the input of the KEEL engine. The other devices could be local or remote to the KEEL engine.

Warning messages could be triggered.

The outputs from the KEEL engine could be used to trigger warning messages. The warning messages could use other outputs to describe the warning in relative terms.

Information messages could be triggered to indicate status.

Analog values could be included to explain subjective interpretation.

KEEL engines can used to supply variable information in the form of informational messages. Complex messages could be structured from multiple variable output signals.

Commands to operators could be generated.

KEEL engines interpret information and provide outputs that represent a balancing of the inputs. These values can be used to generate complex commands to an operator.



Logging of information could be triggered. KEEL could log its own decisions or it could log other inputs and outputs.

The outputs from KEEL engines could be logged for historical records or for audits.

The inputs to the system could also be logged. If the log format is in XML compliant with the KEEL Input Schema, then the data could be used in the development environment to recreate the decision-making model for exact interpretation.

KEEL could cause state changes of the system to take place based on subjective evaluations.

KEEL engines interpret input information and drive outputs. These outputs could drive an external state machine that could cause the equipment holding the KEEL engine (or any other system with or without the KEEL engine) to change state. In this case, the KEEL engine is supplying inputs to the state machine.

KEEL could generate diagnostic interpretations.

Beyond just generating processed / interpreted information, KEEL engines can interpret diagnostic information and explain the interpretation in detail.

In addition to explaining the interpretation, it can provide the information to explain why other diagnostic interpretations are not considered.

The sources of input data

KEEL engines can accept inputs from almost any type of data source, as long as the data can be transformed into the normalized data format required by the KEEL engine.

When the input is textual or verbal, it will have to be transformed into the normalized format.

Human operator

In systems where the human operator is part of the system there is the potential that the operator will be providing input data to the system.

For example, the operator could be supplying input data to the system as part of the job function. The operator could be reading values or recording physical observations about characteristics of the problem domain. This could be a doctor that records physical symptoms of the patient or of the environment that may contribute to the symptoms.

In other cases, the operator could be directed to take specific measurements. An example might be an automotive service



technician that could be directed to take readings in an automotive electrical system to try and isolate the problem.

Inputs from human operators are commonly gathered through some kind of Human Interface Device that will transform the information from human terms to formats more conducive to digital processing. This might be via a keypad, a pushbutton, or in some cases it might be voice input. It could also be in some form of visual form where information is generated by physical movement. It is possible that any of the human senses could trigger inputs to a KEEL system.

Sensors

Sensors of all types can be used as inputs to a KEEL system. As long as the information can be transformed into one or more normalized input values, it can be interpreted by a KEEL engine.

Sensors exist to detect and measure almost all physical states. For example: time, temperature, pressure, torque, speed, acceleration, distance, density, color, edges, shapes, counts, volume, etc. There are probably sensors to measure anything for which a value can be assigned to it.

Collections of sensors can also detect and measure non-physical information: stress, truthfulness, pain. These values are determined with some algorithm that synthesizes the information.

Databases

Databases are used to store historic and synthesized data. This data can be manipulated by any number of mathematical processes to provide running averages, identify trends, detect shifts, etc.

The result of database queries can generate numeric information that can be used as inputs to KEEL engines.

Databases that are constantly updated have the ability to send evolving data to KEEL engines and thus tune the KEEL engines with new data.

External Data Sources

In addition to databases, external data sources can be any device or software application that generates information.

For example, machine tools may have counters embedded in them that count completed operations or completed orders. This information is gathered as the equipment operates and can provide input information to KEEL engines.



A clock or calendar is another example of an external device that can generate information for a KEEL engine. These devices generate time related information.

A communication network might generate traffic information.

Other KEEL Engines

KEEL engines may be components of a larger cognitive system. In these cases it is likely that one KEEL engine will feed other KEEL engines.

The KEEL FBD tool provides a mechanism for integrating multiple KEEL engines in a single application.

A more loosely coupled solution would be to connect KEEL engines across a network or some other connectivity approach.

Locally accumulated data

A KEEL engine will be embedded in a device or software application.

It is likely that the device or software application will be performing functions in addition to the cognitive process associated with the KEEL engine. In these cases other values generated by the application may be used as inputs to the KEEL engine. Certainly diagnostic and prognostic data generated by a device might be used to drive a KEEL engine.

Preprocessed data

Preprocessed data can exist anywhere in a system. This preprocessed data could have gone through a validation process or a transformation process. It could carry with it confidence data or some other biasing information that could be used in conjunction with the preprocessed data. It could be accumulated locally or it could exist anywhere in the system where it could be move to the KEEL engine for interpretation and processing.

Control Signals from Other Devices

Control Signals from pieces of equipment or software applications can be used as inputs to KEEL engines.

In some cases KEEL engines are part of autonomous devices. They react to their surroundings and decide what to do for themselves. In some of these cases, the KEEL engine could intercept control signals from another device that are directed to perform operations for that other device. This information could provide intelligence for the device containing the KEEL engine.



Connectivity

The ability to install KEEL engines at any point in a system makes this a valuable attribute where components of the system are likely to be distributed across different pieces of equipment in different locations.

KEEL engines can be connected by any media and in any format as long as the data is converted to the normalized format before triggering the KEEL engine to process it.

Directly wired to source

In its simplest form a sensor can be directly wired to an input pin on a microprocessor where the signal is transformed to a normalized data format used by the KEEL engine.

The same is true for the output. In its simplest form the normalized data output from the KEEL engine is transformed before sending it out a pin on the microprocessor where a wire carries the signal to an actuator (control point).

A direct wire is the simplest form of network.

Network connected - any topology

Connectivity to and from KEEL engines can be accomplished with any type of network with any topology.

Some wired networks might be termed point to point, multi-drop, token passing, star, web, loop, trunk, etc.

Messaging techniques might include: Broadcast, Store and Forward, Directly Addressable, Group Addressable, All-call.

Any message packaging technique might be used: structured or unstructured, packed, block mode, etc.

Any character encoding can be used: ASCII, Async, Bisync, Isochronous, or any other.

The data can be encrypted or non-encrypted.

The choice of network connectivity is left to the system designer.

Infrared connection

The KEEL engine is not restricted to any specific connectivity to its inputs and outputs. Infrared links can be used.

Supplied by the same processor running the KEEL engine



Because some data sources and data sinks for KEEL engines will be within the same microprocessor, input and output data can be generated and consumed locally.

Radio Frequency

KEEL engines are independent of the connectivity between inputs and outputs and the KEEL engine. Radio frequency connectivity is appropriate for some applications where wired and infrared connections are not appropriate.

Development Process:

Determine where KEEL engines might be located in the system and what information will be exchanged

This is where the number of KEEL engines will be identified to satisfy the system need.

This is also where the information model will be defined: What input information is needed at each KEEL engine and how that information will flow through the system.

Define the information hierarchy which includes how and when information is accumulated

Some systems will fit in a "chain-of-command". The concept of authority needs to be considered in these types of systems.

This will dictate who is a supplier of information and who is a receiver of information.

This is where the concept of trust needs to be included in the model.

The concept of risk may also be considered.

Plan for staged introduction

In systems that might require multiple KEEL engines, consideration for a staged introduction might be appropriate.

Since KEEL engines operate autonomously, this is primarily an issue that is driven by other parts of the system.

If cognitive technology is to be included in the overall system, then it should be scheduled for introduction when it can be tested with the availability of its inputs and outputs.

Evaluate KEEL segmentation

KEEL engine segmentation is an issue that is required with distributed systems. It should also be considered with complex cognitive situations.



The KEEL FBD tools provide support for the integration of multiple KEEL engines into a single compile for a localized solution.

There is also the potential for integrating multiple independent segments into a single KEEL engine. In situations where one is dealing with simple cognitive processes, it might be appropriate to include several of them in the same engine.

For each KEEL engine, define the expectations (outputs) from the system

The outputs of the system can drive external controls, provide information or logging data, drive external logic, provide warnings or triggers, or drive other processing engines (like other KEEL engines) or hard coded logic.

The outputs of the system define what the KEEL engine is trying to accomplish.

For each KEEL engine, identify the sources of information

The sources of information are the inputs to the system as viewed from the outside of the KEEL engine. They can come from sensors, databases, human operators, other KEEL engines, other control systems, or internal logic (like self diagnostics and prognostics or internal calculations).

KEEL General:

This section identifies the general benefits that could be derived from a KEEL based solution.

Judgmental decisions by trained operators are potentially "tricked" into overlooking critical attributes

There are some applications where there is an attempt to trick humans.

Magic shows attempt to redirect attention while the trick is performed.

Some markets like airport security target people that may try to trick the system. A terrorist is not going to advertise he is a terrorist.

When people are trained for airport security, they are trained to look for specific attributes. These are the rules.

KEEL engines are rule based systems. The rules identify the trigger points.

While the perfect airport security agent will never miss a critical attribute, human nature makes them susceptible to structured trickery.

A KEEL engine will perform according to its rules.



Human experts take too long to make judgmental decisions

Human experts are trained to perform a task. This works well when the task operates in normal ways.

However, when humans are "surprised", it takes them some amount of time to react.

Adrenaline is used to supercharge the human in special circumstances. This accelerates their response mechanism and makes them more aware of their surroundings.

Even with adrenaline, however, reaction time is based on the complexity of the task.

Computers running structured code react to emergencies (events) in a structured way.

Computers, however, cannot watch out for everything. The code to develop this type of system would be too complex.

The adaptable nature of a KEEL engine may provide a middle ground: faster than a human and not so complex that one would have to write too much code to respond.

The dynamic nature of a KEEL engine enables it to handle complex situations in adaptable ways.

Applications where the judgmental decisions must be explained

While not all decisions or actions need to be explained, there are a number of reasons for having an explainable system.

First, it is difficult to enhance a system when you don't know why it is doing what it is doing. With a KEEL system, if one takes a snapshot of the critical inputs, those inputs can be loaded into a KEEL engine where the reasoning can be investigated.

Some systems need an explainable solution for legal reasons.

A system that creates explainable actions may be more marketable and may be more likely to receive market acceptance.

Complex situations where it is not economical to develop and maintain straight line code (IF, THEN, ELSE)

It is very expensive to develop conventional software. Complex software is prone to logic and typing errors.



Large complex systems are brittle and break with unexpected, untested situations.

KEEL systems allow the development of complex systems without writing code. The decision-making model is developed graphically and the code is created automatically. The model can be tested from within the toolkit, before it is translated to source code.

The same cognitive model can be translated into multiple languages at the click of a button.

Situations where the environment is dynamic and the importance of information changes. The system must react to change

A key attribute of KEEL systems is that the dynamic importance of information is built into the KEEL technology.

KEEL was designed to accept and adapt to the changing importance of information.

Information changes in importance when it is being evaluated from different time and space aspects. (For example, the temperature is important because you are close to the source. It is not important if you are far away from the source. What is NOW is more important than what is FAR IN THE FUTURE if you are thinking about a problem now. It is less important now if you are addressing a strategic problem.)

Situations where there is an advantage to be able to create one design and execute it on multiple platforms: device, software simulation, web

In some circumstances there is a large value for being able to demonstrate the same model in a device, in a simulator, and on a web page.

Using the KEEL Toolkit, the same design can be implemented in C, C#, VB, VB.NET, Flash, Java, PLC Structured Text.

When it is important to demonstrate the cognitive technology in multiple environments, KEEL can provide a significant advantage.

Where architectural issues may prohibit other solutions (KEEL technology is architecture independent: localized, distributed, web based, multiprocessor, etc.)

KEEL engines are implemented as two or three subroutines in any of a variety of languages.

One subroutine is called at the initialization of the system to load the arrays with data.



The second subroutine is scheduled when appropriate. Before this routine is called, the KEEL arrays are loaded with input data. Then the second routine is called repetitively until a flag is set indicating that a stable system has been determined. After that the outputs from the system are distributed to the appropriate places.

This simple API and small subroutine set allows KEEL engines to be integrated into a variety of system architectures.

KEEL engines don't care where the input data comes from or where the outputs are distributed.

This makes the KEEL engines appropriate for almost any system configuration.

The architecture selection is left to the system architect.

Human experts are required to interpret information to make the best decisions or take the most appropriate actions

KEEL engines interpret information in a human like manner. In this light, KEEL technology is an expert system that makes decisions or takes actions using rules established by the designer.

Like human decision-making, KEEL engines utilize a linkage of information to solve problems. An input may impact a number of problem domains; each in a different way. A single problem domain may impact a number of other problem domains in a variety of ways.

The linkages defined for the KEEL engine can be compared to the various ways a human might link pieces of information.

The KEEL toolkit provides a unique way to document the human decision-making model for specific applications.

Devices must operate autonomously and make judgmental decisions on their own

Each KEEL engine has its set of inputs and outputs.

When the KEEL engine is scheduled to evaluate the inputs, it iterates until a stable answer to all of its problems is achieved.

Each KEEL engine is a stand-alone function. If a complete system is incorporated into a single engine, then a complete stand-alone solution is developed.



A system can be built with one engine or it can be built with several engines that are integrated into a single autonomous solution by using the KEEL FBD tool. Services provided within the KEEL Toolkit support the software engineer's need to compile multiple engines with a single compile.

KEEL technology has been developed to specifically address the needs of adding cognitive technology to devices in order to provide autonomous operation.

KEEL technology supports addressing the judgmental decisions required when there is a need to make subjective decisions in a human-like manner.

Devices can make control decisions when human operators are not present

Because KEEL engines can operate autonomously, they are able to provide backup to human operators when the human operators are unavailable.

KEEL engines can make the same judgmental decisions as their human operators. These judgmental decisions are appropriate for making control decisions.

These control decisions can be binary (on/off) or relative (analog). Just like a human operator, a KEEL system can take observed data (inputs) and apply cognitive judgmental rules that are defined graphically and make the control decisions.

When the small memory footprint of a KEEL engine is an advantage
KEEL technology is implemented in two or three small subroutines. The subroutines do not expand as the size or complexity of the system increases.

The remainder of the system is made up of tables. These take much less space than "code".

This allows KEEL engines to run in 8 bit microprocessors and above.

This is important in many "device" applications.

Repetitive judgmental decisions are prone to error

Some human activities or occupations require that those humans are trained to observe and react to situations. This could be an airport security guard or it could be a medical diagnostician.



The airport security guard is asked to look for things out of the ordinary. This triggers increased scrutiny of passengers and of luggage.

The medical diagnostician is trained to look at diagnostic tests.

Both of these activities are repetitive. There are potentially large amounts of data (characteristics) to observe.

The decisions or observations by the reviewer are prone to error, just by the number of repetitive judgments that are made.

The airport guard that hasn't identified a problem passenger in five years may get passive as time goes by.

The medical diagnostician that never finds a disease may miss the one critical test that identifies a particular diagnosis.

Summary:

KEEL technology is used to interpret information. This is the function provided by the doctors as they interpret diagnostic information.

Once defined, KEEL engines can perform the same interpretation as the doctors. And KEEL engines can provide an explainable answer which can be used to validate the doctor.

At some time in the future, the KEEL engine will make the diagnosis and the doctor will validate the diagnosis.

There are some cases where groups of people have to live without direct access to human medical personnel. A KEEL based system could perform the same diagnosis as if there was a live doctor involved.

In manufacturing, we have come to expect that products built with automated equipment will out perform those manufactured by hand. Eventually medical diagnosis will be performed primarily by machines and the consumer will have the same level of expectations.

Disclaimer

This application note suggests the potential for KEEL technology to respond to certain market needs. The end users are totally responsible for assuring that the technology performs as expected.

The application note may also assume that certain external technology exists to support the KEEL engine in an effective manner. This may or may not be accurate in all cases.