LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

LLNL-TR-651620

Rapid, autonomous analysis of He spectra I: Overview of the RadID program, user experience, and structure

T. B. Gosnell, J. R. Chavez, M. S. Rowland, J. L. Wong

March 11, 2014

Disclaimer

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

Rapid, autonomous analysis of HPGe gamma-ray spectra I:
Overview of the RadID program, user experience, and structure

Thomas B. Gosnell, Joseph R. Chavez, Mark S. Rowland, and James L. Wong

Table of Contents

Abstract

RadID is a new gamma-ray spectrum analysis program for rapid screening of HPGe gamma-ray data to reveal the presence of radionuclide signatures. It is an autonomous, rule-based heuristic system that can identify well over 200 radioactive sources with particular interest in uranium and plutonium characteristics. It executes in about one second. *RadID* does not require knowledge of the detector efficiency, the source-to-detector distance, or the geometry of the inspected radiation source—including any shielding. In this first of a three-document series we sketch the *RadID* program's origin, its minimal requirements, the user experience, and the program operation.

Introduction

Every gamma-ray spectrum tells a story and, like a book, it can be read if you understand the language. A book in good condition with crisp print can be easily read. A book in poor condition, such as a tattered ancient biblical scroll with voids in the document, can only be read with imprecision. Similarly, a gamma-ray spectrum acquired in the field for a brief period of time and smeared by a detector with poor energy resolution is likely to be read with similar imprecision, sometimes resulting in nuclide misidentification [1].

Some gamma-ray spectra are collected by in the field, under suboptimal conditions, for short periods of time, and can result in sparse data that cannot reveal the finest details. Nevertheless, the superior energy resolution of HPGe spectra provides considerably greater information content than is found in commonly employed scintillation detectors, such as NaI(Tl) [1]. For this reason, HPGe excels in the analysis of complex spectra from sources of mixed radionuclides [2] such as plutonium. This is true even for spectra from high-resolution detectors where only the most intense peaks can regularly be observed with confidence.

RadID is a new application built largely on the foundation of an earlier rapid response codebase. In this new cod, rapid execution is retained and still applicable to analysis of spectra obtained in the field. There, data acquisition times are short, in the neighborhood of five minutes and almost always done out of doors. Program execution time also needs to be rapid. For this reason, the application is based on an expert system [3,4]. Once feature extraction is performed, the system traverses a decision tree with a Boolean result at each branch that is accomplished in about one second.

The speed advantage of an expert system is balanced by a disadvantage: an expert system has expertise in a limited domain. If data from outside of that domain is introduced to the system, unpredictable results will likely be obtained. *RadID* was designed to analyze relatively sparse data of the 600-second kind. This is contrary to what most gamma-ray spectrometrists desire. A very long count of a complex spectrum, such as from plutonium for example, to obtain "better" statistics, may reveal peaks with low emission intensities that conflict with those that *RadID* is programed to detect. This can sometimes lead to nuclide misidentification.

RadID origin, requirements, and implementation

Rapid nuclide identification in the field poses unique analytical problems

In the last decade radiation detection instruments were distributed to thousands of first responders for use in case of a radiological emergency. Some of the data acquired are high-resolution gamma-ray spectra.

Because rapid response is required for these events, it was apparent to the first spectrum analysts that the normally excellent software tools for high-resolution gamma-ray spectrum analysis were ill-suited for the rapid nuclide ID task. Information that was normally available in laboratory measurements could be expected to be absent, including the measurement environment (almost always out of doors), the source-to-detector distance, the source geometry, and the packaging around it. Spectra were typically acquired for short acquisition intervals. Energy calibration was frequently challenging and background spectra, when acquired, were typically unreliable.

As a result, new software was created. At LLNL a simple program for rapid HPGe spectrum analysis was written using ladder logic with heuristics to identify 40 radionuclides including uranium, plutonium, and other nuclides that can be commonly found in commerce.

Evolution of a heuristic system

In late 2002 we decided to expand the capability of the original ladder logic application. We developed a graphical user interface, the number of radionuclides that were of interest to identify was considerably expanded, some source characterization capability was added, spectral templates of unshielded radionuclides in the application library were computed to overlay on the measured data to confirm nuclide ID. We had never planned to formally develop an expert system but as the capability of the application grew it began to exhibit the classic elements of a rule-based expert system [4,5].

As our experience with a wide variety of field spectra deepened, we needed to incorporate newly gained knowledge of their characteristics in a relatively painless manner. An expert system differs from traditional programs by having a unique structure. It is modular, divided into two parts: a knowledge base or rule base that can be easily modified or expanded and an inference engine that remains fixed and reasons from the knowledge base in much the same manner as a human expert.

Another characteristic of expert systems is that they are heuristic—like humans they reason with expert judgmental knowledge, rules of thumb, as well as formal established knowledge. As such *RadID* follows the typical reasoning process used by gamma-ray spectrometrists. Here we sketch the elements of *RadID* that are characteristic of expert systems as applied to nuclide ID.

Feature extraction

RadID extracts a number of spectral features such as the presence or absence of key peaks, their net counts, peak-area ratios and integrated counts in selected spectrum energy regions. These features form a working database that is grist for the expert system. These features are examined by a binary decision tree; a rule set (knowledge base) that evaluates several hundred potential analysis results to find those that apply to a particular spectrum.

Fig. 1 is a block diagram showing the relationships between an expert system's elements.

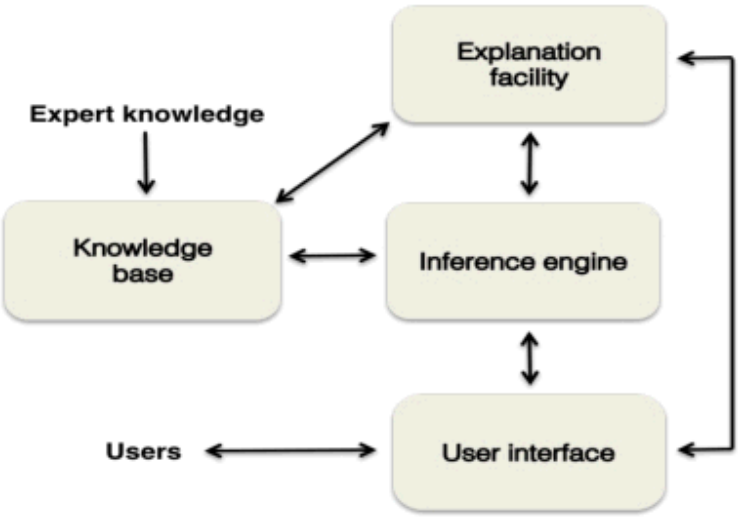

Fig. 1. Expert system block diagram

User interface

RadID employs a graphical user interface (GUI) expressed in three windows, augmented by a fourth utility window for any needed energy calibration adjustment.

Knowledge base

A rule-based system uses a knowledge-based approach rather than a strictly procedural approach where the knowledge base is represented in the form of a set of rules. Rules are an elegant, expressive, straightforward, and flexible means of expressing knowledge]. Facts are presented in IF-THEN-ELSE form to yield Boolean values. Some rules lead directly to a finding. However the Boolean output from most rules become conditions for use by rules that follow in the decision tree. Evaluating the TRUE/FALSE Boolean output of the *RadID* rules provides a swift journey through the decision tree.

Inference engine

The inference engine is coding that decides which facts are germane to the analysis, prioritizes their execution order, determines which rules are satisfied by these facts, and prioritizes the results. Typically, the presence of a nuclide will be reported when key line combinations for that nuclide are present and, usually, some further conditions that must also be satisfied to prevent nuclide misidentification and to report additional findings beyond simple nuclide ID.

Findings

The inference engine prepares a list of findings (sometimes oddly called an agenda), prioritized by a figure of merit. The rules that generate a findings list must be satisfied by the features in the working database. Once the inference engine establishes that all of the conditions in a rule and its antecedents are satisfied, it then adds the finding to its appropriate rank in the findings list reported once the journey through the rule set reaches its conclusion.

Explanation facility

The user interface provides graphical and textual information that aids the analyst in determining the validity of the *RadID* findings. Of greatest importance, computed spectral templates are available for comparison to candidate nuclides or complex sources. The templates are not models of the measured data but are used for characteristic line comparison. Additionally, the knowledge domain of *RadID* encompasses some rarely encountered radionuclides that may be unfamiliar to the analyst. Because spectral information is rarely context free. The user interface provides radionuclide usage information that, combined with contextual information provided with the spectrum, can provide analytical confidence. For example, if a radiation source appears to be within a person who claims to have had medical radiation treatment and *RadID* identifies it as ^{67}Ga, the GUI usage pane provides the following information: *3.2612 d half-life, Used in humans for imaging of abdominal infections, (…) Can be detected for 2–3 weeks*

Advantages and disadvantages of expert systems

Expert systems are designed to operate within a narrow, well-defined domain of knowledge. Their advantages include:

- Creation of efficiencies and reduction of the time needed to analyze a spectrum,
- A modular database of expert knowledge that is independent of the inference engine and *is more easily modifiable than conventional programs*,
- Within the domain of the knowledge base, a different problem can be solved using the same program without reprogramming efforts,
- The ability to reason heuristically. The human mental process is internal, and it is too complex to be represented as an algorithm. However, most experts are capable of expressing their knowledge in the form of rules for problem solving.

Expert system disadvantages include:

- Lack of common sense. For example, *expert systems have difficulty in recognizing domain boundaries.* When given a task different from the typical problems (for example when a complex and unusual unknown spectrum, such as perhaps mixed nuclear waste, is presented to *RadID*), an expert system might attempt to solve it and fail in rather unpredictable ways,
- Lack of creativity. Human experts can respond creatively to unusual situations, expert systems cannot,
- Problems with excessive scope expansion of the knowledge base: for example adding new nuclides in *RadID,* can create conflicts with existing rules, increasing the difficulty of knowledge-base maintenance,
- To err is human. Even a brilliant expert is only a human and thus can make mistakes. This suggests that an expert system built to perform at a human expert level should also be expected to make occasional mistakes. Errors and omissions in the knowledge base can lead to incorrect findings. Expanding and debugging the *RadID* knowledge base has been in progress for ten years. It is now valued by its users but still remains a work in progress. Remaining errors in its conclusions have proven to be readily detectable by its users—experienced gamma-ray spectrometrists—experts themselves. As such, *RadID* is a tool to *assist* experienced gamma-ray spectrometrists, not replace them. It provides rapid preliminary results

for the analyst but is not a tool intended for use by those unfamiliar with gamma-ray spectrometry.

The user experience with basic nuclide identification

The user's viewpoint

Launching *RadID* reveals its primary window. An example of weapon-grade uranium is shown in Fig. 2.

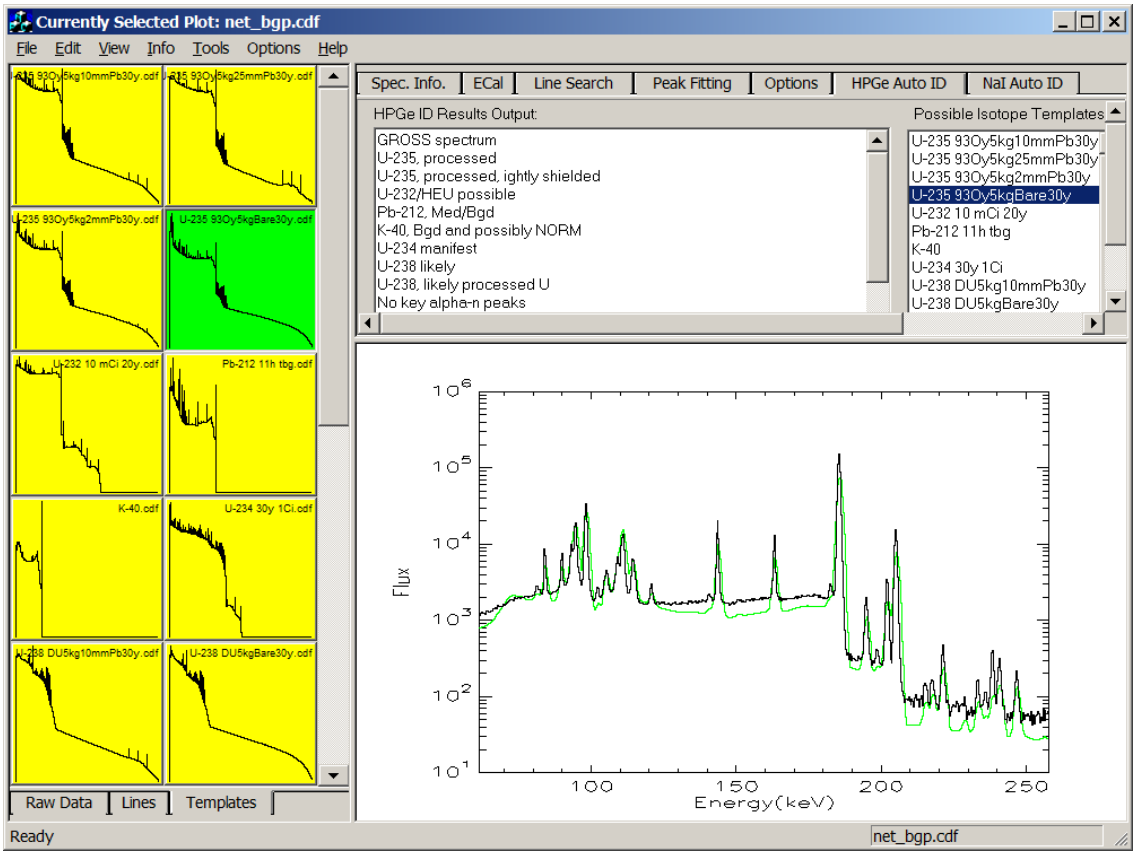

Fig. 2. The *RadID* primary window. In this figure, the user has used the File menu in the upper left hand corner of the window to read in a foreground spectral data file in the pulse-height analyzer's proprietary format. RadID translated the proprietary format into a common internal format and displayed it in the primary viewer pane. Abbreviated and prioritized result findings appeared in the upper middle pane. These events unfolded within about one second. A number of candidate spectral templates that might be displayed for the user to compare with the displayed spectrum data have appeared in the upper right-hand pane. Spectral thumbnails of the candidates are revealed on the left-hand side by clicking on the templates tab in the lower left-hand corner. The user has since zoomed the primary view to the first 300–600-keV of the spectrum and sampled a variety of candidate templates for comparison to the measured data, finally settling on a template of bare 30year-old ^{235}U, shown in green.

To date, most user analyses have been for benign sources with commonly encountered radionuclides such as medical ^{99m}Tc or Naturally Occurring Radioactive Material (NORM and TENORM (Technologically Enhanced NORM) [5,6,7]). For these spectra, the results in the Primary window may well suffice. For unusual or complex data, additional information can be

obtained from the Details window chosen from the menu bar in the primary window. Fig. 3 shows the Details window.

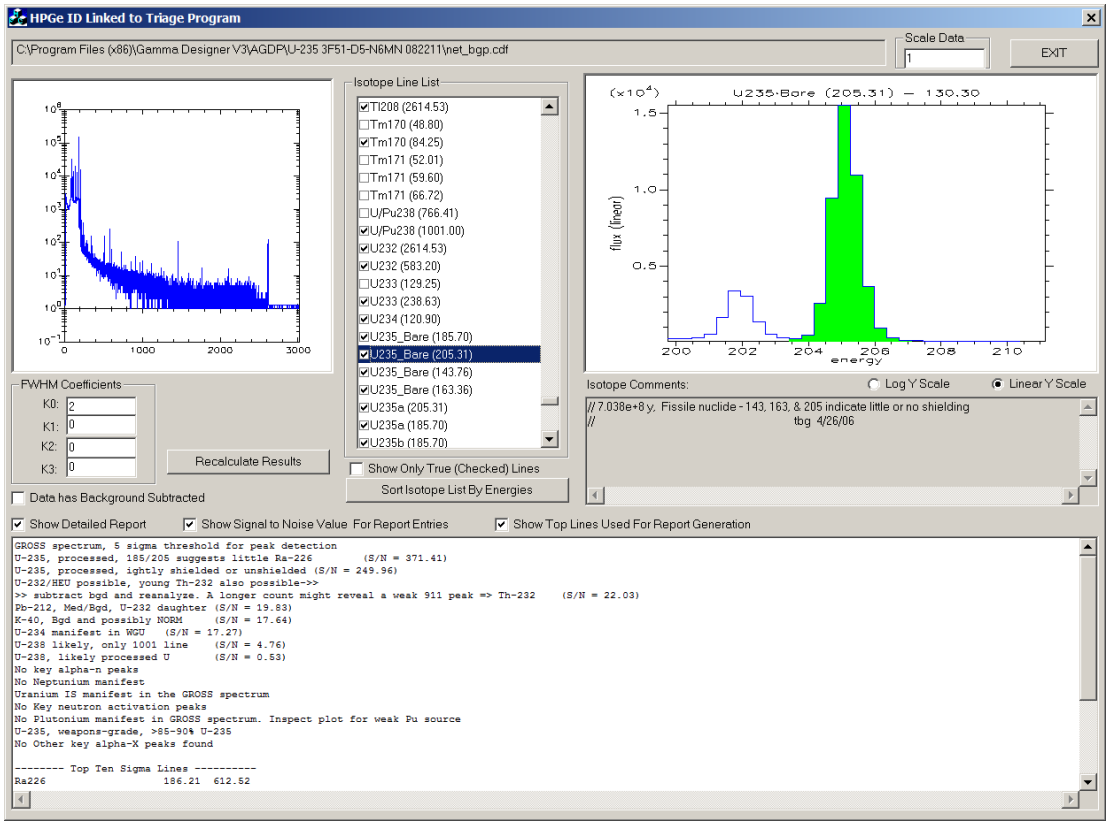

Fig. 3. The Details window. This figure includes expanded textual results from the primary window. They are prioritized by a descending figure of merit, S/N. They are followed by summary results indicating the presence or absence of uranium, plutonium, neptunium, common neutron activation activity, common alpha-n activity, and other common alpha-x activity. A thumbnail display of the entire spectrum appears in the upper left pane. The upper center window reveals the 600+ lines that *RadID* searches for and lists them alphabetically by associated element and mass number. Small boxes to the left of the line descriptions contain checks if the line is an initial candidate for ID. Below this window is a button that will re-sort the lines by gamma-ray energy. A close-up single-line view is shown in the upper right pane for the 205-keV ^{235}U line, chosen by the user from the isotope line list. The green shading indicates the spectral region used for ID of this peak. It is determined by the energy calibration and indicates the expected width of the peak at its base on the underlying continuum. Below the line-view pane is a description of the nuclide's provisional ID, its half-life, categorization (e.g. medical, industrial, fissile, research, or impurity) and, usually, its uses/applications or other raison d'être.

Qualitative descriptions of uncertainty

Problem solving concerns the search for a solution. The principal role of the inference engine is to search for the most appropriate item of knowledge at a given moment [4]. For *RadID,* achieving search efficiency necessitated the reliance on heuristics as commonly used in expert systems.

> *"A heuristic is a rule of thumb, strategy, trick, simplification, or any other kind of device which drastically limits search for solutions in large search spaces.* [8]*."*

When we report out the findings, some of the rules can produce results of dubious certainty or, on occasion, incorrect results. *RadID* is now sufficiently mature that unrecognized incorrect results have become increasingly rare for the nature of the data that have been presented to *RadID* to date. When dubious results *do* occur and *are* recognized by the inference engine, we indicate a subjective judgment of uncertainty in the rule findings report. This may be expressed as a word or one or more question marks:

- No caveat $\Rightarrow$ near certainty
- Likely
- Probable
- ? $\Rightarrow$ Possible
- ?? $\Rightarrow$ Doubtful
- *???* $\Rightarrow$ Very doubtful

Findings accompanied with the ?? symbol are typically associated with multiplets. Because RadID has no multiplet resolution capability, this symbol is often associated with attempts to identify the presence of characteristic X-rays. If these X-rays are of particular interest to the user, a careful visual inspection will usually resolve uncertainty.

Figure of merit

The figure of merit, S/N, shown in the results pane in the Details window is a vestigial indication of relative signal strength found in all incarnations of this application since its infancy. A review of the S/N algorithm revealed that it lacks some statistical rigor as an estimate of signal-to-noise ratio. Nevertheless, it has been in use for nearly 20 years has proven to be a useful and reliable indicator of signal strength. With few exceptions, to be worthy of reporting, findings must be accompanied with a S/N value that exceeds 5. Experience has shown that S/N values approaching or exceeding 10 are of high statistical certainty but do not exclude the possibility of systematic error. This is particularly true for the nuclides that emit only one strong gamma ray.

There has been insufficient experience with the new background compensation capability in RadID to evaluate the figure of merit efficacy in this circumstance.

Some observations about background interference in field measurements

Background radiation is ever present in field spectra and can interfere with the signal of interest from a measured radiation source. If the signal of interest is especially weak, background can even mask the source presence entirely. To mitigate this difficulty, in typical radioanalytical laboratory measurements, samples are prepared by a radiochemist in a standard geometry and counted in a standard geometry in a heavily shielded counting chamber in a laboratory. This dramatically reduces background intensity in the chamber and allows for a trustworthy background measurement when the source is absent.

In contrast, it can be extremely difficult to measure a trustworthy background spectrum in the field. Moving or reorienting the detector to obtain a source-free background measurement will frequently produce a measurement that will likely differ significantly, and sometimes even exceeding, the true background at the source position. It is not at all unusual for environmental background to vary significantly over distances of a few meters. Furthermore, the possibility of small instrumental shifts between foreground and background measurements would produce a distorted net spectrum upon background subtraction. For this reason, all earlier versions of

RadID's preceding incarnations purposely excluded a background subtraction capability as an unnecessary added complexity that would likely be more harmful than useful.

However recently another application area was identified in which background subtraction was needed and trustworthy background spectra could be obtained. . In early 2014, a background subtraction capability was added to form a new application—*RadID*.

RadID—Adding a background subtraction capability

If the user wishes to subtract a trusted time-scaled background spectrum from the foreground spectrum, the usual method proceeds as follows. As in the basic example of nuclide ID, the user will first read in the gross spectrum, show in Fig. 4 as an analysis of weapons-grade plutonium.

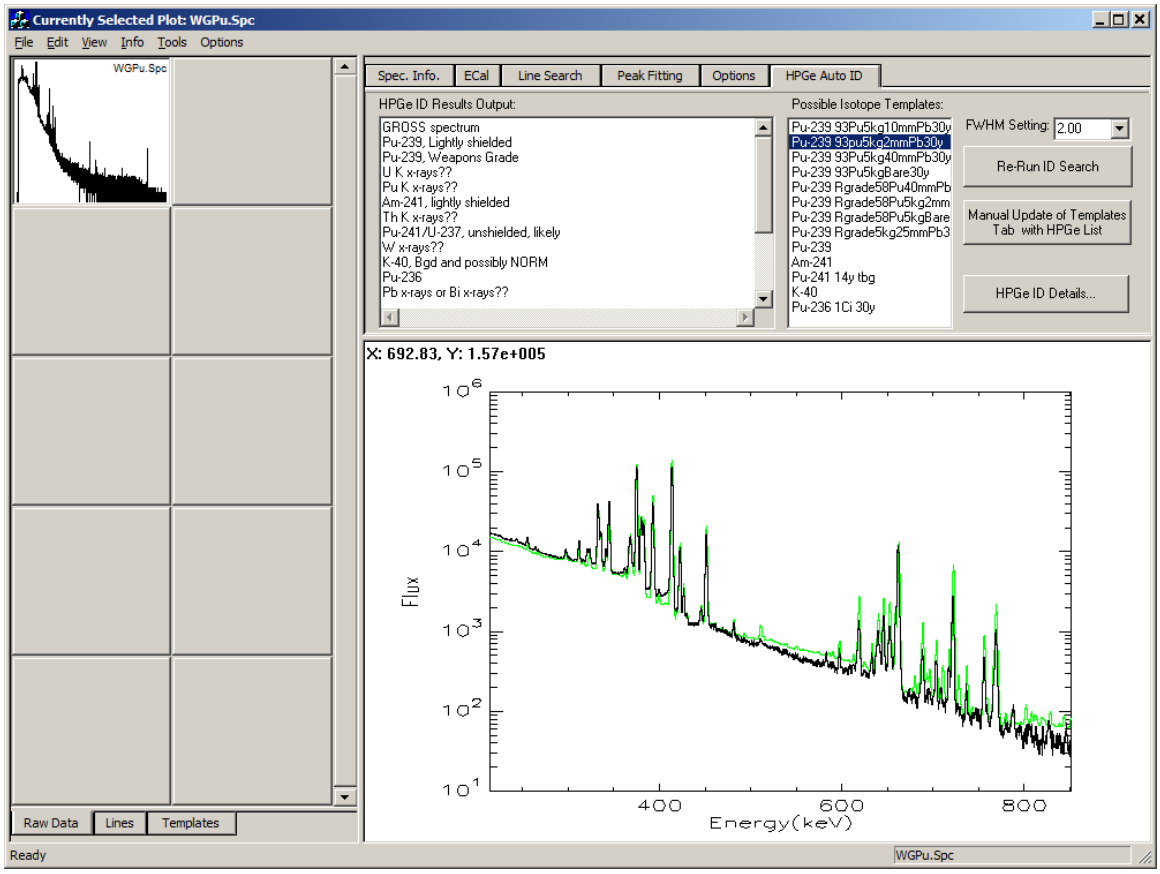

Fig. 4. The user has read in a spectrum from weapons-grade plutonium and zoomed into 300- to 800-keVregion to examine the most of the most defining portion of ^{239}Pu signature. Multiple templates for ^{239}Pu are available for comparison and a template shielded by 2-mm of lead is closest to the measured data.

When the user wishes to subtract a trusted time-scaled background spectrum from the foreground spectrum, the background spectrum can be read in using the File menu. It is shown in Fig. 5.

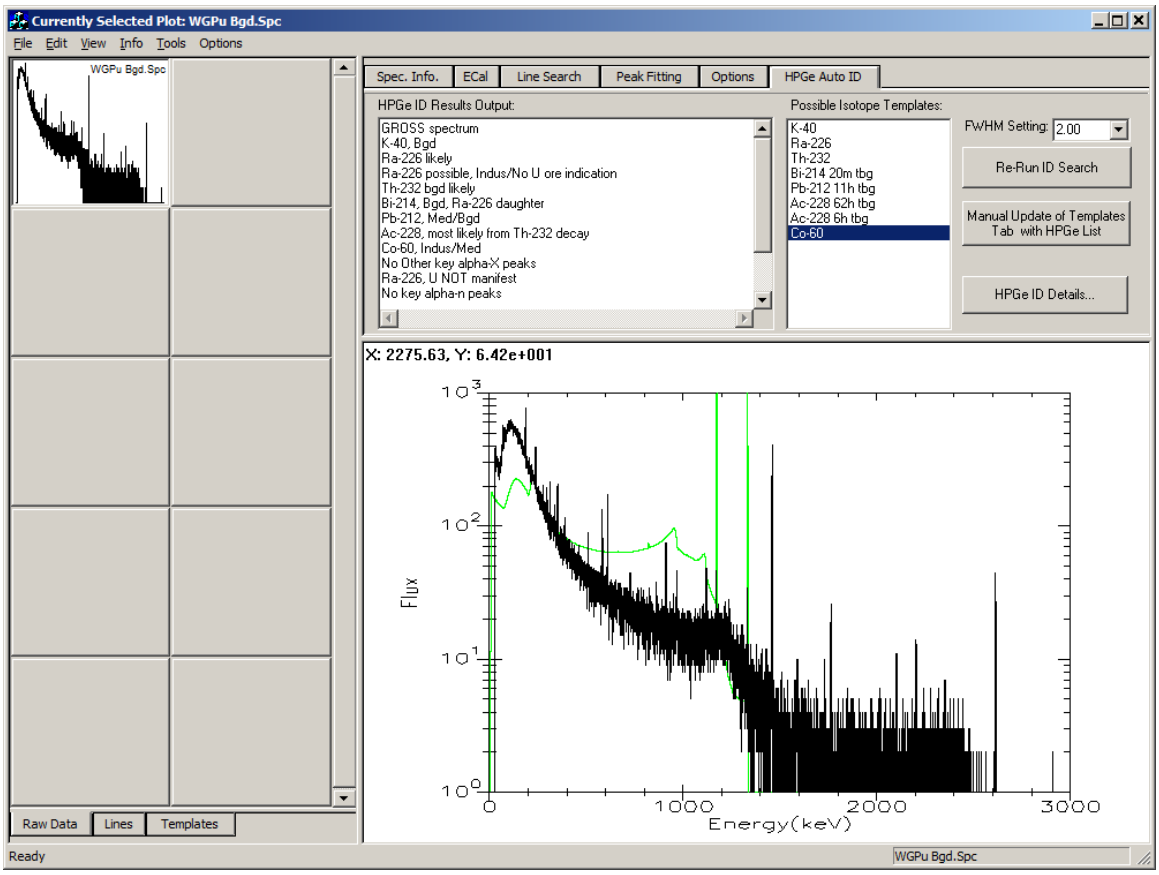

Fig. 5. The user has returned to the File menu and chosen an appropriate background spectrum for the foreground spectrum. The most prominent peaks in background spectra at 1461-keV from primordial ^{40}K and at 2615-keV for the distant ^{208}Tl daughter of primordial ^{232}Th. Distant daughters of primordial ^{238}U (^{226}Ra and ^{214}Bi) are identified. Other distant daughters of primordial ^{232}Th (^{228}Ac and ^{212}Pb) are also identified in the results output pane. All of the sources identified in the pane are from background activity save for a spurious source of ^{60}Co, present in the measurement environment that is also identified and its 1173 and 1333-keV peak signature is verified with a green template chosen by the user.

Once the foreground spectrum is read in a second time, the background spectrum subtracted automatically. Fig. 6 shows the net spectrum formed by channel-by-channel subtraction of the time-scaled background from the foreground spectrum. In this figure, we notice that zero and fractional values are evident in the net spectrum because some values are plotted below unity. This artifact results from low values with large Poisson uncertainty in the background spectrum being subtracted from similarly large uncertainty values in the foreground spectrum. If full spectrum background subtraction is done, this artifact is unavoidable.

Another artifact that we can also observe is the effect of imprecise registration of the foreground and background spectra in the net spectrum. It is clearly noticeable in the 1461-keV ^{40}K peak. This peak is entirely from background radiation and, if the registration is precise, it should vanish in the net spectrum. It does not. It is greatly reduced but it does not vanish. The energy resolution

of HPGe detectors is so high that if the tiniest instrumental drift occurs between foreground and background measurements, some noticeable misregistration will often occur.

We might expect a similar misregistration on the 2615-keV ^{208}Tl peak. The peak is reduced by about a factor of two but does not largely disappear, nor should it. In the foreground spectrum this peak arises from the distant ^{208}Tl daughter of both background ^{232}Th and the trace, but very active, 8^{236}Pu impurity in the WGPU (see Appendix A). Once background has been removed, the contribution of ^{236}Pu is revealed.

If the degree in misregistration observed is not objectionable, then the analyst can proceed to the Details window. However, if the misregistration is troubling some gain tweaking can be done via the ECal tab. This tab can be accessed for spectrum calibration adjustment and possibly achieve system gain alignment of foreground and background.

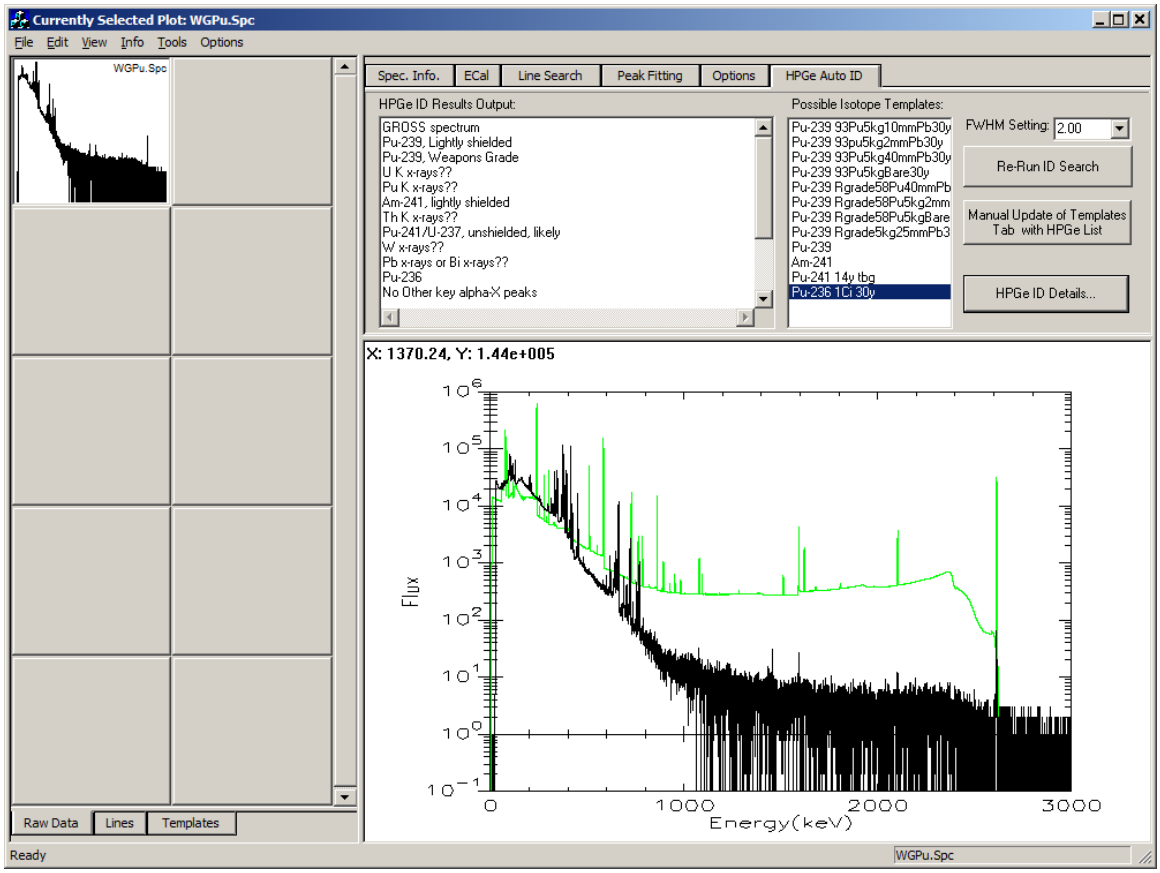

Fig. 6. Net spectrum. In this figure, to obtain the net spectrum the user has once again had to return to the File menu and reload the foreground file. Because loading a background file preceded this action, when the foreground file is read in, the background data have been automatically time-scaled and subtracted on a channel-by-channel basis to produce a net spectrum. A tiny residual of background-only ^{40}K is present at 1461 –keV from slight misregistration of the foreground and background spectra. Residual counts in the ^{208}Tl peak at 2615-keV is from the trace but highly active ^{236}Pu that is always present in weapons-grade Pu. A green template for unshielded ^{236}Pu provides confidence in this interpretation.

The Details window, chosen by clicking on the HPGeID details button, above the main spectrum viewer widow, is shown in Fig. 7.

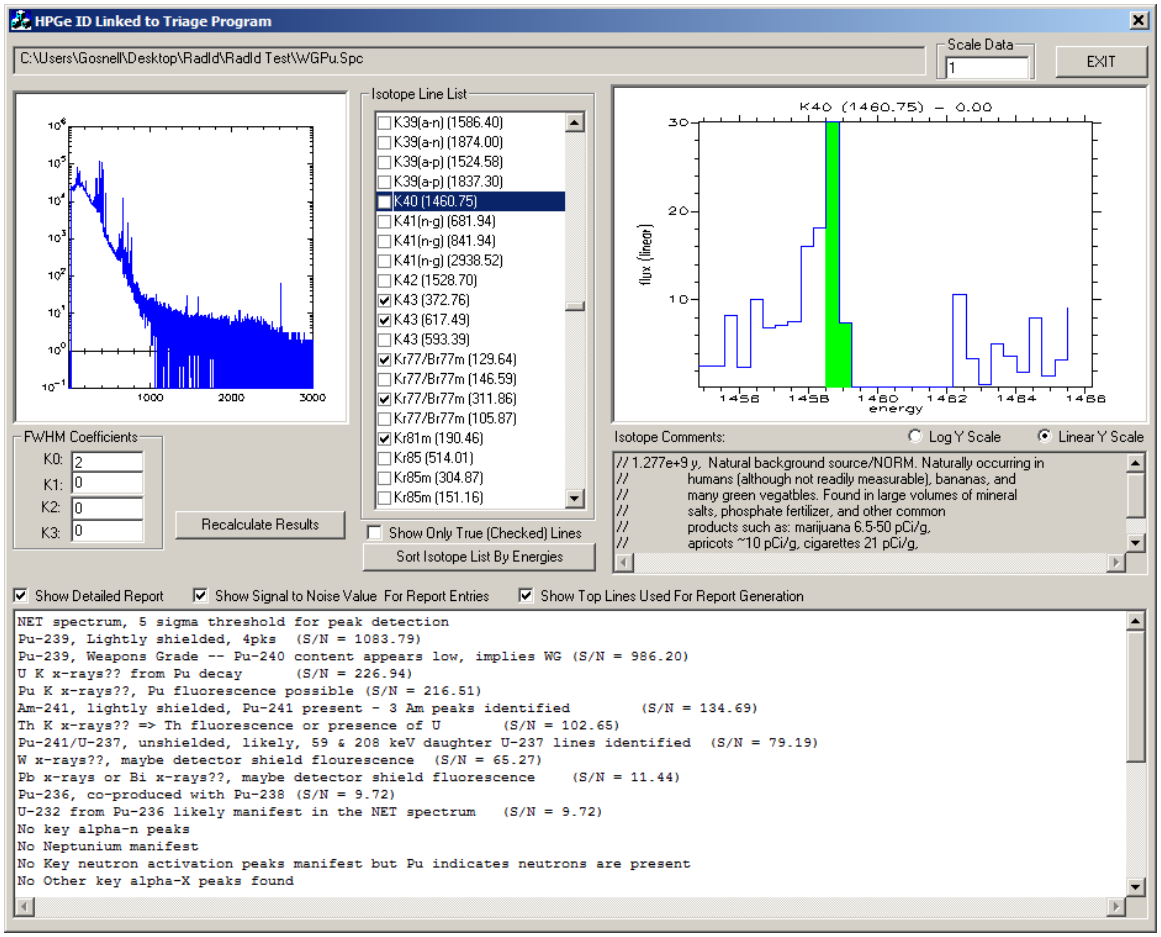

Fig. 7. The Details window. In this figure these include expanded textual results from the net spectrum window (Fig. 6). Of particular interest is the single line window that displays the small, malformed residual of the net 1461-keV peak.

The single-line view in Fig. 7 shows the small mal-formed residual for the 1461-keV peak after background subtraction. The green band likely indicates residual from the low-energy side of the background peak. To the right of that, the body of he peak is absent.

Program operation

Feature extraction

In *RadID* feature extraction searches the gamma-ray spectrum to find features to populate a working database that is the input to rule-based system for analysis.

Counts in selected regions of interest (ROI)

Counts in selected regions of interest are integrated for detecting Doppler-broadened peaks, detection of possible fission gammas from neutron multiplication, and determining the presence of weapons-grade plutonium.

Characteristic gamma-ray lines for nuclide ID

In HPGe spectra peaks are the most prominent and most important feature. Typically in HPGe spectrum analysis programs, a search is made over the entire spectrum to find all peaks all peaks of significant strength. In its original application, *RadID* had to run on a very slow processor that took too long to perform the traditional peak search. Instead we identified key peaks that could be used to identify nuclides of interest, examined the spectrum at the expected peak locations, and tested for peak presence. Because, in the original and most current *RadID* applications counting intervals are expected to be short, the radiation source can be shielded from the detector, or the radiation source may be a weak emitter, we search for only the most intense peaks emitted by each radionuclide of interest.

RadID advances serially to the energy positions where each of the key peaks could be found if present. For each peak, its width on the underlying continuum is defined by the detector resolution width at that energy. We determine the ratio of net counts in a full-energy-peak to the underlying continuum. If the net counts exceed the figure of merit threshold, typically set at five, it is given a Boolean value of TRUE, and the net counts are reserved for subsequent use.

Obviously good spectral energy calibration is essential for *RadID* to produce accurate peak location results. Application of *RadID* to a poorly calibrated spectra will invariably produce a hodgepodge of fantasy identification results. Although *RadID* possesses an energy calibration capability, chosen by the *ECal* tab shown in Fig. 2, highly precise energy calibration over a wide energy range cannot be expected of the spectra presented to *RadID.* Therefore *RadID* is necessarily forgiving of small errors in energy calibration, even though such errors exacerbate the peak identification search.

At this point peak identification is considered provisional. Misidentification of a peak can occur when the excursion above the continuum is not due to a photopeak. This is not commonplace but occurs from time to time particularly at the lower-level discriminator edge. Recalling that *RadID* is forgiving of small deviations of energy calibration accuracy, provisional peak ID may occur for an imposter peak that emits a gamma ray with energy close to the search energy. With the large number of radionuclides that *RadID* attempts to identify, this possibility is commonplace. Most peak misidentification is resolved in the following section of the code that determines the provisional presence of radionuclides

Ratios of peak counts and/or ROI counts for nuclide ID

Nearly 100 ratios of peak areas are used for nuclide discrimination, , presence of source shielding, and to determine source characteristics such as estimates of minimum uranium enrichment. Ratios of counts in regions of interest are used to identify the presence of Doppler broadened peaks and to detect the presence of weapons-grade plutonium.

Heuristic rule set for nuclide ID and characterization

Provisional nuclide ID and nuclide uses

Provisional radionuclide ID

Provisional nuclide ID occurs with a series of rules, at least one for each radionuclide. Each rule has a unique label followed by one or more conditions that are the Boolean values assigned to each of one to four key peaks used for the provisional peak ID. This rule examines each of the key lines to assess its Boolean value. If they are all found to be TRUE, a logical AND, then the provisional nuclide ID is declared TRUE.

Currently we have specified search for over 600 peaks for identifying more than 200 sources. In spite of the high resolution of HPGe, multiple peak energies, especially at low energies, can fall within the resolution width of the detector. For example, in Fig. 3 the green shading at the base of the peak in the peak window defines the resolution width for the 205-keV peak from ^{235}U. Not shown in the figure is the tentative ID of the 204.1-keV peak from medical ^{255}Fm as well as the 204.12-keV peak from fission product and medical ^{95}Nb. Moreover, complex spectra, such as from plutonium, will have gamma-ray lines that are sufficiently close to many of the 600+ key lines, that they can become imposter candidate peaks. Bogus candidate peaks begin to be winnowed out with the search for candidate nuclides that usually require the presence of two or more candidate peaks.

Further winnowing is typically required to achieve final nuclide ID by attaching additional conditions, such as appropriate ratios of the areas of winnowed peaks. Attributes associated with nuclides, for example uranium enrichment estimates, require nuclide ID followed by additional conditions.

Radionuclide uses/applications

If an unexpected nuclide ID of a rarely encountered nuclide occurs, the user can refer to the radionuclide *Isotope Comments* pane in the Details window to determine if the presence of this nuclide is consistent with the measurement context. If not, a misidentification is likely and should be reported for *RadID* correction.

Findings—Final nuclide ID and source characterization

The findings section comprises approximately 1000 rules that attempt to wrest as much information as is available from the extracted features. There can be multiple final ID findings for a particular nuclide that report nuclide presence and/or some particular characteristic. Each rule has a unique label that is followed by the provisional nuclide ID and then followed by zero to many conditions to finalize the nuclide ID and/or to determine nuclide characteristics. Examples of multiple findings are illustrated in all of the preceding user interface figures.

RadID findings are determined in the following categories.

- Findings based entirely on ratios
- Presence of background indicators
- Alpha particle activation indicators
- Neutron activation indicators
- Uranium presence findings
- Findings of possible HEU masking

- Uranium enrichment estimates
- Neptunium findings
- Plutonium reports
    - Presence of absence of weapons grade plutonium
    - Presence of absence of reactor-grade plutonium
    - Presence of absence of heat-source plutonium
- Americium-241 findings
- PET imaging sources
- Presence of other, principally medical and/or industrial, nuclides
- Key findings summary that includes:
    - Presence of absence of manifest uranium
    - Presence of absence of manifest plutonium
    - Presence of absence of manifest neptunium
    - Presence of absence of key alpha-n peaks
    - Presence of absence of other alpha-X peaks
    - Presence of absence of neutron activation peaks

Debugging the *RadID* knowledge base

Most errors in RadID findings are misidentification of nuclides not present in the spectrum. For misidentification, the debugging proceeds generally as follows:

- An experienced analyst using *RadID* determines that a finding is suspect and forwards the spectrum being analyzed with a description of the misidentification to the R*adID* knowledge base expert
- The *RadID* expert:
    - Uses *RadID* knowledge base with the problem spectrum to duplicate the suspect finding
    - Examines the key gamma-ray lines used to determine the preliminary ID
    - Examines the explanation template for the suspect nuclide
    - Determines if the ID is actually correct
    - If the finding is incorrect, determines where the feature overlap occurs. This may sometimes require referral to data compilations such as The Table of Isotopes [9] and NuDat [10]
    - Adjusts the finding rule and issues a new version of the *RadID* knowledge base to the *RadID* programmer for distribution in the next maintenance update of *RadID*

In uncommon cases, an analyst observes a signature or feature of interest in the spectrum that is not reported in the findings. For missed identification, the debugging proceeds generally as follows:

- The analyst alerts the *RadID* expert of the problem and supplies the spectrum of interest
- The *RadID* expert determines if the nuclide or feature in question is indeed in the spectrum
    - The *RadID* expert determines if the nuclide in question *is not* included among the *RadID* library of nuclides of interest

        - Ignore missed ID report as a spurious incident
    - The *RadID* expert determines if the nuclide in question *is* included in the *RadID* nuclides of interest
        - I*f so,* the expert determines which branch in the decision tree was erroneously negotiated and corrects the rule to provide an ID finding
        - *If not* and the nuclide is of some interest or if it is liable to occur frequently enough to become a nuisance
            - If the nuclide is not of particular interest but *is* likely to occur frequently then add it to the rule set to prevent future missed ID reports
            - If the nuclide is not of particular interest but *is not* likely to occur frequently then it can be ignored to avoid library clutter

Conclusion

RadID is a gamma-ray spectrum analysis program originally written to assist in the detection of illicit movement of nuclear material. It is specific to the rapid analysis of HPGe gamma-ray data to reveal the radionuclide signatures of interest that may be present in the spectra. While some users have found *RadID* to be useful for other problems, the program retains much of its original character. In this document, we have sketched the program's origin, its minimal requirements, the user's experience, and program operation.

The program execution proceeds through two types of analysis. The first analysis is feature extraction, a quantitative analysis, that locates key peaks of interest, determines their net counts above and signal strength above the continuum. It then computes selected ratios of peak areas. Finally it integrates the counts in a few small regions of interest and computes selected ratios of integrated counts.

In the second part of the analysis we is an application based on heuristics and employs the classic elements of a rule-based expert system. In addition to simple nuclide identification, *RadID* includes some focused examination of radiation source characteristics such as estimates of uranium enrichment and plutonium grades and reports these findings.

RadID differs from its previous incarnation by adding the capability to subtract a trustworthy background spectrum from the foreground spectrum on a channe-by-channel basis to obtain a net spectrum.

Appendix A: The thorium decay series

A major source of structure in background spectra is from the thorium decay series shown in Fig. A-1. The figure also includes a collateral series beginning with ^{236}Pu that joins the thorium series at ^{228}Th.

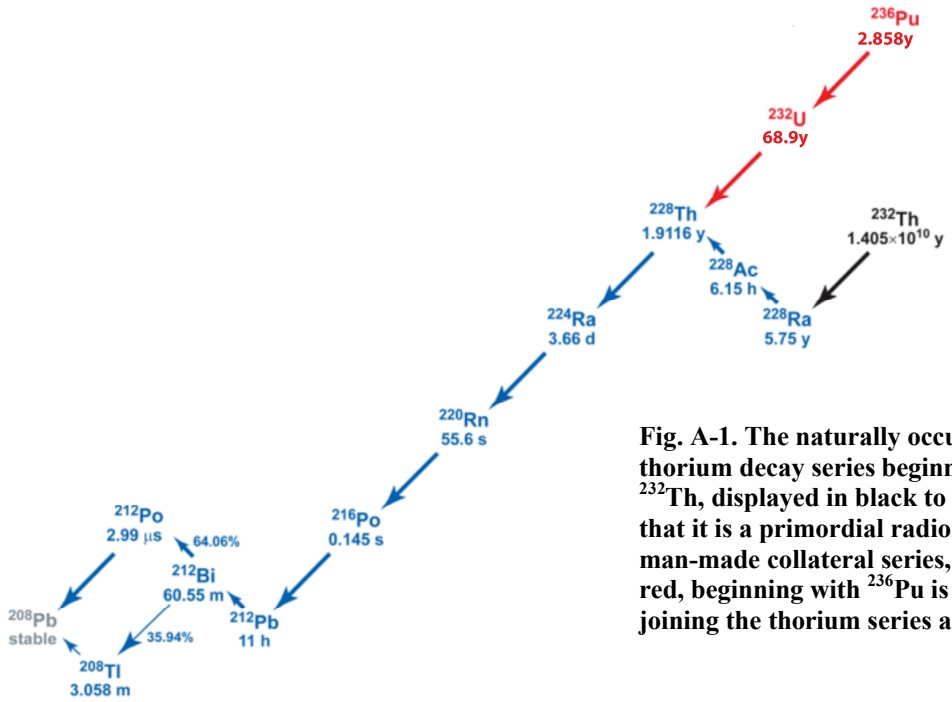

Fig. A-1. The naturally occurring thorium decay series beginning with ^{232}Th, displayed in black to indicate that it is a primordial radionuclide. A man-made collateral series, shown in red, beginning with ^{236}Pu is shown joining the thorium series at ^{228}Th.

Thorium-228: Nexus nuclide—a ubiquitous and ambiguous signature

The natural thorium decay series arises from the decay of the 1.4x10^{10}-y primordial radionuclide ^{232}Th. ^{228}Th, with a 1.9-y half-life, is the great granddaughter of ^{232}Th. It decays to seven successive radioactive daughters. These daughters are so short-lived that they require the presence of the longer-lived ^{228}Th to sustain their existence for more than a few weeks. One of these daughters, ^{208}Tl, provides one of the most recognizable signatures in background radiation with strong gamma rays at 583-, 861-, and 2615-keV. Observation of ^{208}Tl, therefore, indicates the presence of ^{228}Th. A computer-simulated point source ^{228}Th spectrum is shown in Fig. A-2 illustrates its gamma-ray signature.

The importance of ^{228}Th is that it is the nexus of two decay series with the nuclides ^{232}Th and ^{236}Pu as their progenitors (Fig. A-1). The presence of ^{228}Th in a gamma-ray spectrum is therefore not necessarily unique to a particular parent.

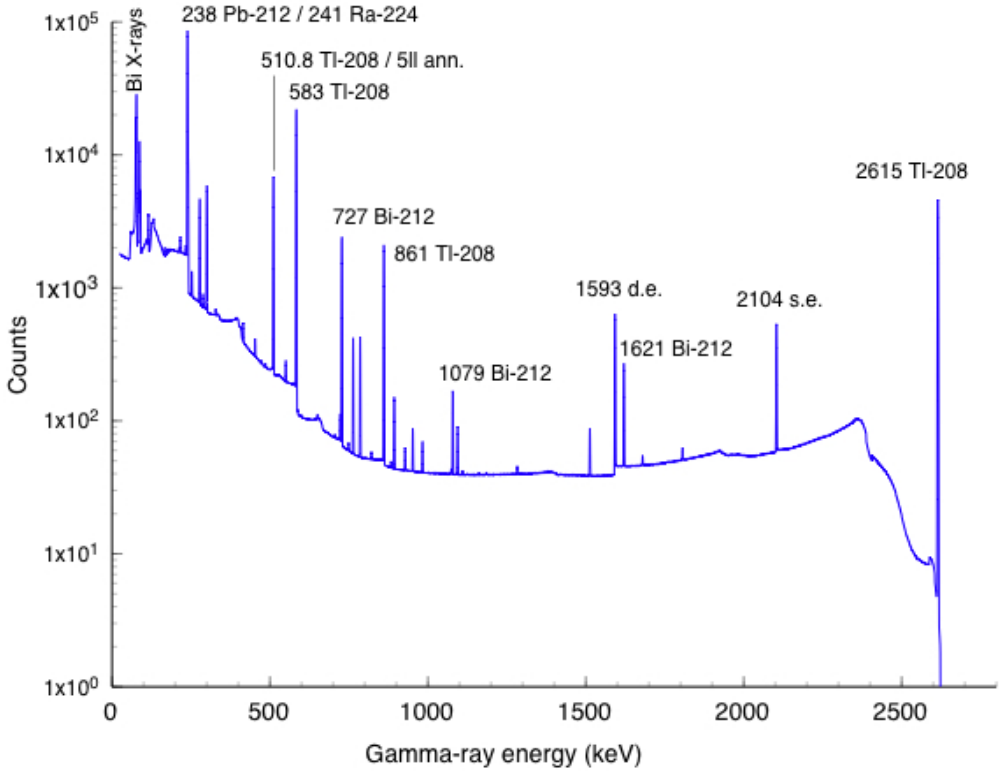

Fig. A-2. Illustration of the ^{228}Th signature with a computer-simulated HPGe spectrum from a year-old ^{228}Th point source

References

1. Karl E. Nelson, Thomas B. Gosnell, David A. Knapp, *The effect of energy resolution on the extraction of information content from gamma-ray spectra*, Nucl. Instr. Meth. Phys. Research A **659** (2011) 207.
2. Glenn F. Knoll, *Radiation Detection and Measurement,* Third Ed., Wiley, New York 2000, p. 428.
3. Artificial Intelligence, *Rule-based expert systems,* http: intelligence.worldofcomputing.net/expert-systems-articles/rule-based-expert-systems.html
4. Adrian A Hopgood, *Intelligent Systems for Engineers and Scientists*, 2nd Ed., CRC Press, pp 1–47, New York, 2001.
5. U.S. Environmental Protection Agency, *About TENORM*, http://www.epa.gov/rpdweb00/tenorm/about.html
6. U. S. Environmental protection agency, *Technologically-Enhanced, Naturally-Occurring Radioactive Materials,* http://www.epa.gov/radiation/tenorm/
7. U. S. Environmental protection agency, *TENORM in consumer products*, http://www.epa.gov/radiation/tenorm/consumer.html
8. A. Barr and E. A. Eigenbaum, *The Handbook of Artificial Intelligence,* vol. 1, Addison-Wesley, 1986.
9. Lawrence Berkeley National Laboratory, *Exploring the Table of Isotopes, http://ie.lbl.gov/education/isotopes.htm*
10. Brookhaven National Laboratory, *NuDat 2.6,* http://www.nndc.bnl.gov/nudat2/