

Computer Generation of Process Explanations in Nuclear Astrophysics

SAKIR KOCABAS (UCKOCA@ITU.EDU.TR)

Department of Space Engineering
ITU, 80626 Maslak, Istanbul, Turkey

PAT LANGLEY[◇] (LANGLEY@ISLE.ORG)

Institute for the Study of Learning and Expertise
2164 Staunton Court, Palo Alto, CA 94306 USA

Abstract

In this paper we describe ASTRA, a computational aid for generating process explanations in nuclear astrophysics. The system operates in two stages, the first using knowledge of quantum theory to produce a set of legal reactions among elements and the second searching for pathways of such reactions that explain the construction of some element from lighter ones. ASTRA has found apparently novel reactions that involve proton and neutron capture, as well as novel fusion reactions that produce neutrons and deuterium. The system has also generated reaction pathways for helium, carbon, and oxygen that do not appear in the scientific literature. However, ASTRA also finds many other reaction pathways that are less interesting and that suggest priorities for future research.

Running head: Generating Explanations in Astrophysics

Key words: computational scientific discovery, model formation, reaction pathways,
process explanations

[◇] Also affiliated with the DaimlerChrysler Research and Technology Center, 1510 Page Mill Road, Palo Alto, CA 94304 USA.

1. Introduction

The computational study of scientific discovery has made important strides in its short history. Early research focused on replicating discoveries from the history of science, covering results from disciplines as diverse as physics, chemistry, and biology. The types of discoveries also ranged widely, including numeric laws (e.g., Langley, 1981), qualitative relations (e.g., Jones, 1986), structural models (e.g., Zytlow & Simon, 1986), and process models (e.g., Kulkarni & Simon, 1990). Nevertheless, some critics questioned these results because they involved scientific relations already known to the developers.

In recent years, researchers have turned their efforts toward the computational discovery of new scientific knowledge, with some success (Langley, 1998). For instance, Mitchell, Sleeman, Duffy, Ingram, and Young (1997) report their DAVICCAND system finding a new numeric relation in metallurgy, whereas Buchanan and Lee (1995) describe novel results on whether chemicals cause cancer. Another important example is Valdés-Pérez' (1995) MECHEM system, which has found new reaction pathways in physical chemistry. His progress has encouraged us to examine other branches of science in which reaction pathways occupy a central position.

In this paper, we focus on scientific discovery in the field of astrophysics. In the section that follows, we briefly review this discipline's basic problems and methods. After this, we describe ASTRA, an astrophysical research aid designed to support scientists in explaining the synthesis of elements and their relative abundance in stars. After explaining the inputs, outputs, and operation of the system, we report the reactions and reaction pathways that ASTRA produces for some astrophysical problems. One challenge in research on scientific reasoning is to determine the quality of new discoveries in terms relevant to the field at hand; we devote some attention to this issue, using a careful literature search to evaluate ASTRA's outputs. This analysis reveals some limitations of the system and suggests directions for future research.

Our aim in developing ASTRA is not to automate the scientific process, but rather, as with DAVICCAND and MECHEM, to provide scientists with computational support. Astrophysics has a strong theoretical framework, which it largely shares with particle physics. However, constructing process models for particular phenomena is a tedious task that involves considering many alternative fusion reactions and exploring many potential reaction pathways. We hypothesize that published accounts of such stellar phenomena are not the only such process models that astrophysicists would find acceptable, and expect that scientists will welcome a tool of this sort to help them identify candidate process explanations.

2. Process Explanations in Nuclear Astrophysics

Astrophysics is a branch of science that deals with both the tiniest objects in the universe, elementary particles, and the largest, stars and galaxies. Nuclear astrophysics, one of its subfields, focuses on the formation of chemical elements, through a series of fusion and decay reactions in stars, from hydrogen (H) and helium (4He), thought to have emerged in the early history of the universe. An-

other important problem concerns the distribution of elements, in particular the abundance carbon (^{12}C), nitrogen (^{14}N), and oxygen (^{16}O) relative to lighter elements like lithium (7Li), beryllium (9Be) and boron (^{11}B).

According to the current astrophysical theories, stars traverse several stages in their lifetimes. The first stage, which follows the star's initial formation by the condensation of cosmic clouds and hydrogen gas, involves 'hydrogen burning'. During this period, stars radiate energy emitted by a series of exothermic fusion reactions in which hydrogen is transformed into helium. Astrophysicists have proposed three different pathways (Audouze & Vauclair, 1980, p. 52; Williams, 1991, p. 351) to account for hydrogen burning in stars the size of the sun and smaller. Later stages consist of more complex reactions, typically involving heavier elements.

In their attempt to explain nucleosyntheses, the theorists first select a stellar model in thermal equilibrium that makes certain assumptions about the mass, temperature, density, and distribution of elements in the stellar plasma. Next, they identify the particle and nuclear reactions that, by calculation, are consistent with quantum physics. Finally, they calculate the rates of these reactions using experimental and theoretical knowledge about nuclear cross sections and reactant abundances. In this way, they obtain a set of valid reactions with their rate coefficients. Scientists use reactions with high rates to construct pathways, either by working forward from light elements to the final element or by working backward from the final element to the initial light elements.

Naturally, there are many possible reactions among substances, and an even greater number of reaction pathways that could account for the nucleosynthesis of a given element. Astrophysicists deal with this large space by focusing their attention on only a small set of reactions, relying on heuristics to constrain the generation of explanatory hypotheses, such as preferring pathways that involve high reaction rates. This behavior, heuristic search guided by domain-specific constraints, is consistent with other findings about human cognition on nonroutine problems.

However, the fact that astrophysicists find effective ways to limit their search does not mean they find the best solutions to their problems. We hypothesize that more systematic searches of the same space might uncover alternative process models, and that computational tools could aid this search process, in the generation of possible reactions within a selected energy band and in the construction of pathways from a given set of reactions.

In the remaining pages, we examine ASTRA, a system designed as a computational aid for developing process explanations in nuclear astrophysics. After describing the system's operation, we report results on processes that involve hydrogen and helium burning, the formation of carbon, nitrogen, and oxygen through these processes, and on anomalies in the relative abundance of light elements. Our treatment draws repeatedly on publications in nuclear astrophysics, including Audouze and Vauclair (1980), Clayton (1983), Fowler (1986), Fowler et al. (1967), Fowler et al. (1975), Harris et al. (1983), Cujec and Fowler (1980), Kippenhahn and Weigert (1994), Lang (1974), and Williams (1991). In closing, we discuss the implications of our results with ASTRA, related research on computational scientific discovery, and directions for future work.

3. The ASTRA Discovery System

Before we describe our application of ASTRA to nuclear astrophysics, we should first consider the two main stages of the system's operation. The first module generates all reactions that are valid theoretically, whereas the second produces reaction chains that constitute process explanations for the nucleosynthesis of elements. In each case, we examine the module in terms of its inputs, outputs, and search mechanisms.

3.1 Generating Nuclear Reactions


The first stage of ASTRA takes as input descriptions for a set of elements and isotopes. Each entity is characterized in terms of five quantum properties – rest mass (in MeV/c^2), electric charge, spin count, lepton count, and baryon count – based on information taken mainly from the *CRC Handbook of Chemistry and Physics* (Weast & Astle, 1981). We also give ASTRA theoretical knowledge about conservation relations over these quantum properties that hold in reactions among the elements and isotopes. We can also constrain the system to consider only exothermic reactions, or we can extend the energy band to include endothermic reactions (e.g., to -3.0MeV), as some endothermic reactions also contribute to stellar nucleosynthesis.¹

Based on this information, ASTRA systematically generates all reactions among these elements that obey the conservation laws and that have inputs of the form A , $A + B$, or $A + B + C$, and outputs of the form D , $D + E$, or $D + E + F$. This gives nine possible forms, including three for decay reactions. The algorithm instantiates each form in all possible ways and retains each proposed reaction only if it conserves all the relevant properties. ASTRA also calculates the energy emissions, or Q values, of each reaction in terms of mega-electron volts (MeV). The basic operation at this stage is equivalent to one module of the BR-4 system, which we have described elsewhere (Kocabas & Langley, 1995). Note that this mechanism generates not only fusion reactions, which is considered the mainstay of nuclear astrophysics, but also decay reactions, in which one element decomposes into other substances.

For each set of reacting entities, ASTRA constructs a set L of possible resultant entities using a simple domain heuristic: The total rest mass of the resultant entities cannot be much greater than that of the reactants. For example, the reaction products of $H + H \rightarrow \dots$ cannot include elements heavier than tritium, because the next heavier element would be ${}^3\text{He}$, which has a rest mass of $2809.4 \text{ MeV}/c^2$ that far exceeds the total rest mass of the two hydrogens (1877.56). The program builds sets of one, two, and three combinations of entities from the set L as possible resultant entities, then tests each combination for the conservation of quantum properties. The reactions that pass all the tests are added to the knowledge base as plausible reactions, along with their assigned Q values.

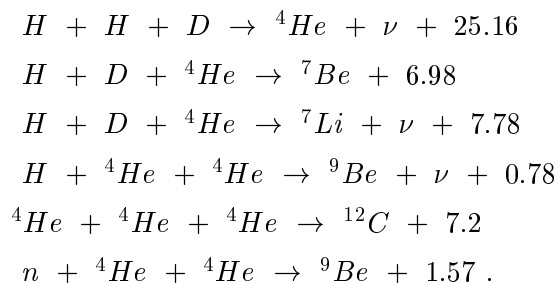
1. Endothermic reactions play an important role in helium burning, but the majority of stellar nucleosyntheses are exothermic.

For instance, this ASTRA module generates three distinct reactions² that involve the elements hydrogen (H) and lithium (6Li):



In each example, hydrogen and lithium (on the left-hand side) combine to form one or more new substances (on the right-hand side), along with the energy, specified in MeV's.

The program generates three-particle reactions in the same manner as those involving two reactants. Some examples of three-particle reactions that ASTRA produces are:



For the runs described in this paper, we provided ASTRA with the elements from hydrogen to oxygen, their isotopes, and a few elementary particles like the electron, proton, neutron, and the neutrino, as well as their antiparticles, giving a total of 36 distinct entities. From these, the system generated some 400 different reactions, but some were minor variations on one another, such as ${}^3He + {}^9Be \rightarrow {}^{12}C + e + \bar{e}$ and ${}^3He + {}^9Be \rightarrow {}^{12}C + \nu + \bar{\nu}$. We eliminated such near repetitions manually, leaving 276 reactions that included 262 fusion reactions and 14 decays.

3.2 Generating Reaction Pathways

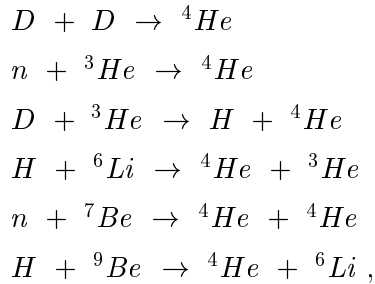
ASTRA's second stage takes as input the primitive reactions generated by the first module, along with an element E whose synthesis the user wants explained and the basic elements/isotopes $\{B\}$ that he assumes as given (typically hydrogen and deuterium). In response, the system generates all reaction chains that lead from the starting elements to the final element through the various reactions identified in the first stage.

The system uses a depth-first, backward-chaining search to construct these process explanations. On the first step, ASTRA finds those reactions that give as an output the final element E . Upon selecting one of these reactions, R , it recursively finds those reactions that give as an output one of R 's input elements. The algorithm continues this process, halting its recursion when it finds a


2. The reaction formulations of ASTRA are based on neutral atoms. For this reason, our notation has minor differences from that found in textbooks, as in the second hydrogen-lithium reaction, for which the textbook version is $H + {}^6Li \rightarrow {}^7Li + \bar{e} + \nu$ instead of $H + {}^6Li \rightarrow {}^7Li + \nu$.

reaction chain for which all the reacting elements are in $\{B\}$ or when it cannot find a reaction off which to chain. ASTRA generates all possible reaction chains in this systematic manner.


To clarify this procedure, let us partly follow ASTRA's construction of reaction chains for the synthesis of ${}^4\text{He}$ through a few steps. The system begins by retrieving reactions that produce the target element ${}^4\text{He}$, which include




as well as many others. The module next looks for reactions that generate items on the left of these candidates for final steps, such as



and other reactions that generate deuterons (D), neutrons (n), and ${}^3\text{He}$. The system also considers the reactions



and similar processes that create varieties of lithium (${}^6\text{Li}$ and ${}^7\text{Li}$), as well as



and other reactions that produce various isotopes of beryllium (${}^7\text{Be}$, ${}^8\text{Be}$, and ${}^9\text{Be}$). This backward chaining continues until each reaction pathway arrives at a starting element, in this case hydrogen, that produces each of the intermediate terms along the pathway, or until the search exceeds a depth limit. The above example suggests a breadth-first search, but in fact the system uses depth-first search to create one reaction chain at a time, backtracking each time it exhausts its options at a given level.

In this manner, ASTRA produces all possible exothermic fusion reactions and decays, including those we have found in the astrophysics literature. This basic process is similar to that used in Valdés-Pérez' MECHEM to discover reaction pathways in physical chemistry. Both systems find chains of reactions that explain how one set of entities transforms into another set, and both use extensive search through the space of pathways, constrained by knowledge of legal reactions. One difference is that our system also generates its reactions from a deeper background theory.³ This provides another layer of reasoning at which scientists may have overlooked explanatory processes, and thus gives ASTRA additional opportunities for novel discoveries, as we can now illustrate.

4. Experimental Results in Astrophysics


In this section we report on our experience with ASTRA, which we organize around conceptual distinctions in the literature on nuclear astrophysics. We first address two broad classes of reactions that astrophysicists believe play an important role in stellar nucleosyntheses, then turn to reaction pathways that explain the generation of heavier elements from lighter ones.

4.1 Proton, Electron, and Neutron Captures


Astrophysicists explain the synthesis of chemical elements from hydrogen and helium in stellar systems through a series of fusion and decay reactions. Two of the main processes are proton and neutron captures, in which a nucleus reacts with a proton or a neutron to give a heavier element or isotope. Electron captures, in which an atomic nucleus absorbs an orbital electron and emits a neutrino, play a relatively minor role in the nucleosyntheses.

Proton captures are an important class of exothermic reactions that contribute to hydrogen-burning processes, in which hydrogen atoms are transformed into helium atoms. We have found 33 examples of proton captures given in astrophysics literature (e.g., Fowler et al., 1967, 1975, 1983) for elements from hydrogen (H) to oxygen (${}^{16}\text{O}$).

ASTRA's first stage predicts that all elements from hydrogen to nitrogen (${}^{15}\text{N}$), with the exception of ${}^4\text{He}$, participate in proton capture. The program produces 46 such reactions, including all 33 examples we have found in texts, but also 13 others, some of which are:




3. However, in more recent work, Zeigarnik et al. (1997) have extended MECHEM to accept component reactions from an external source, thus reducing the difference between our two approaches.




We have not seen any of these reactions in any of the astrophysics texts that we examined. Thus, they are candidates that scientists might accept as theoretically possible but that they have overlooked due to the nonsystematic nature of their generation strategies.

Electron capture reactions are weak interactions in which an electron is absorbed by the atomic nucleus to be transformed into one with a smaller atomic number. An important example in the literature, to which we will return, is the reaction




In this context, ASTRA's first stage produces six electron capture reactions, but only the one above appears in astrophysics texts. Two other electron-capture reactions that ASTRA generates are:



The first reaction seems unlikely to play a role in stars, because ${}^9\text{B}$ has a very short lifetime ($2 \times 8^{-19}\text{sec}$), but the second is a more serious candidate.


In fusion reactions that involve neutron capture, an element combines with a neutron to form a heavier isotope of the same element. We found 17 neutron captures for light elements in the literature, while ASTRA predicts 59 such reactions that are theoretically possible for the same elements. These include five reactions that we did not see in the texts:



The third reaction may play an important role in stellar reaction pathways, which we will consider shortly.

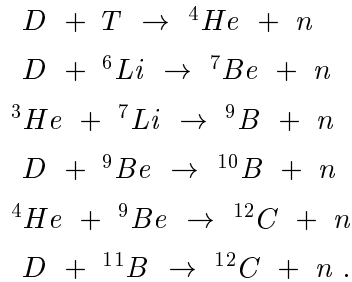
4.2 Neutron and Deuteron Production

Neutron capture requires a continuous supply of neutrons in the stellar plasma, so that it relies on some neutron-producing reaction. Audouze and Vauclair (1980, p. 86) suggest that



which combines deuterons (an isotope of hydrogen), is the only reaction that releases neutrons

in the hydrogen-burning stage of main-sequence stars. Yet ASTRA also predicts six additional reactions that produce neutrons:




The first two of these reactions appear likely in main-sequence stars, as ${}^6\text{Li}$, D , and T (tritium, another isotope of hydrogen) all exist in them. The second reaction seems important, as both D and ${}^6\text{Li}$ are stable isotopes, and thus could play a role in stellar reaction chains. However, astrophysicists would ignore most of these candidates, given the low abundance of their reactants in the stellar plasma.

Most of the neutron-producing reactions rely on a deuteron as one of their inputs. The best known deuteron source is the reaction



in which two hydrogens react, and in astrophysics texts we have found two other reactions that produce deuterium ($T + {}^3\text{He} \rightarrow {}^4\text{He} + D$ and $H + {}^9\text{Be} \rightarrow {}^8\text{Be} + D$). However, the first stage of ASTRA predicts 15 other reactions of this sort. These include:



The first two of these reactions should take place in main-sequence stars, as ${}^6\text{Li}$ and ${}^7\text{Li}$ are known to exist there, yet we have not found either reaction in the literature that we examined. Again, we anticipate that astrophysicists would reject most of these candidates because their reactants occur with low abundance.

4.3 Helium Synthesis through Hydrogen Burning

Astrophysicists hypothesize that the transformation of hydrogen into helium serves as the principal source of energy in main-sequence stars. The standard reaction pathways for helium synthesis given in astrophysics texts (e.g., Audouze & Vauclair, 1980, p. 52; Williams, 1991, p. 351) are the hydrogen-burning processes called ‘proton-proton’ chains. The first such pathway is

- a. $H + H \rightarrow D + \bar{e} + \nu$
- b. $D + H \rightarrow {}^3\text{He}$
- c. ${}^3\text{He} + {}^3\text{He} \rightarrow {}^4\text{He} + H + H$.

The net effect of this chain, when reaction *a* occurs twice, is $4 H \rightarrow {}^4\text{He} + 2 \nu + 26.72 \text{ Mev}$. Another hypothesized pathway, known as the ‘alpha-catalyzed chain’, is

- d. ${}^3\text{He} + {}^4\text{He} \rightarrow {}^7\text{Be}$
- e. ${}^7\text{Be} + e \rightarrow {}^7\text{Li} + \nu$
- f. $H + {}^7\text{Li} \rightarrow {}^8\text{Be}$
- g. ${}^8\text{Be} \rightarrow {}^4\text{He} + {}^4\text{He}$,

in which reactions *b* and *c* provide both the ${}^3\text{He}$ and the ${}^4\text{He}$ needed by reaction *d*. An alternative pathway, which also appears in texts, replaces reaction *e* with $H + {}^7\text{Be} \rightarrow {}^8\text{B}$ and *f* with ${}^8\text{B} \rightarrow {}^8\text{Be} + \bar{e} + \nu$, which produce the ${}^8\text{Be}$ needed by the final reaction through a different mechanism. Astrophysicists refer to these three pathways as the *pp1*, *pp2*, and *pp3* chains, respectively.

When asked to generate reaction chains from hydrogen to helium, the ASTRA system finds all of these reaction chains, including another well-known helium-producing pathway, called the ‘CNO cycle’, that involves carbon, nitrogen, and oxygen atoms, and that takes place in stars more massive than the sun. However, ASTRA also formulates another reaction chain:

- $H + H \rightarrow D + \nu$
- $D + H \rightarrow {}^3\text{He}$
- ${}^3\text{He} + {}^6\text{Li} \rightarrow {}^9\text{Be} + \nu$
- $H + {}^9\text{Be} \rightarrow {}^4\text{He} + {}^6\text{Li}$,

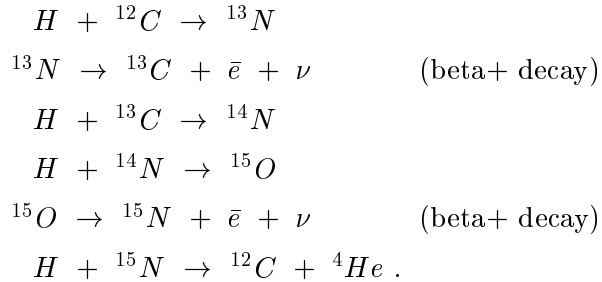
which has the same net effect, provided that there exists some ${}^6\text{Li}$ in the plasma, but we have seen no record of this pathway in the literature. In addition, the program finds 44 other processes of helium synthesis that differ in their last link of the chains. Many of these would be disregarded by astrophysicists because they have small cross sections, including the chain:

- $H + H \rightarrow D + \nu$
- $D + {}^4\text{He} \rightarrow {}^6\text{Li}$
- $H + {}^6\text{Li} \rightarrow {}^7\text{Li} + \nu$
- $H + {}^7\text{Li} \rightarrow {}^4\text{He} + {}^4\text{He}$,

as *D* is believed to be quickly destroyed by the reaction $D + H \rightarrow {}^3\text{He}$. Both Cujec and Fowler (1980) and Harris, Fowler, Caughlan, and Zimmerman (1983) argue that reactions involving *D* are unlikely due to its low abundance. However, Clayton (1983, pp. 371–2) notes that the density

of deuterium in the interstellar medium and the sun remains unknown, and suggests that the substance might be more common than usually believed.


In the CNO chain of reactions, hydrogen is transformed into helium through a series of hydrogen capture reactions involving the isotopes of carbon, nitrogen, and oxygen. The standard account of the first CNO chain, as given in the texts, is:




ASTRA produces a variant of this process using the electron capture of ${}^{13}\text{N}$ in the reaction $e + {}^{13}\text{N} \rightarrow {}^{13}\text{C} + e + \nu$ mentioned earlier, in place of the slow beta+ decay of ${}^{13}\text{N}$ into ${}^{13}\text{C}$ (9.97 minutes) in the second reaction above. The rest of the reactions follow as in the standard account. This constitutes a CNO analogue of the pp2 chain and, if the cross section of $e + {}^{13}\text{N}$ is sufficiently high, the reaction would be competitive with the current explanation. Again, we have not see this alternative chain in the texts.

4.4 Generation of Carbon and Oxygen


The origin and relative abundance of carbon and oxygen have long been central concerns of astrophysics. The standard account (e.g., Fowler, 1986, pp. 5–6) relies on the process of ‘helium burning’, in which helium nuclei react with one another, and with other light elements, to form carbon and oxygen in the steps:



However, there are theoretical problems with this account, in that the first reaction is endothermic and the lifetime of ${}^8\text{Be}$ is very short ($2 \times 10^{-16} \text{sec}$). Later calculations showed that ${}^8\text{Be}$ resonances are sufficiently stable to allow the reaction with an alpha particle to produce carbon, as in the second reaction. In this run, ASTRA does not formulate the reaction ${}^4\text{He} + {}^4\text{He} \rightarrow {}^8\text{Be}$ because it is slightly endothermic, but the system finds 20 other reactions that produce ${}^8\text{Be}$, such as




Once ${}^8\text{Be}$ is available, ${}^4\text{He} + {}^8\text{Be} \rightarrow {}^{12}\text{C}$ can take place exothermically, so ASTRA formulates this reaction. The system produces 24 additional chains that differ in their final steps to ${}^{12}\text{C}$. These include:



which relies on the existence of neutrons in the stellar medium. Briefly, if ${}^8\text{Be}$ captures a neutron before it decays, then it transforms into its stable isotope ${}^9\text{Be}$. This in turn produces carbon by reacting with ${}^4\text{He}$, where the emitted neutron from the latter reaction can combine with another ${}^8\text{Be}$. Most astrophysicists believe this process can compete with the standard one only in explosive stars that produce many neutrons.

Once ${}^{12}\text{C}$ has formed, in whatever manner, it can react with ${}^4\text{He}$ exothermically to produce oxygen through the reaction




In summary, ASTRA finds a number of reaction pathways to carbon and oxygen that do not appear in astrophysics literature, as it does in the case of the lighter element helium. All of these pathways are theoretically possible, but final judgement about their scientific value requires further evaluation, as we discuss next.

5. Discussion of Results

We have carefully compared ASTRA's outputs, at both the reaction and pathway level, to those available in astrophysics texts (Audouze & Vauclair, 1980; Clayton, 1983; Cujec & Fowler, 1980; Fowler et al., 1967, 1975, 1983; Kippenhahn & Weigert, 1994). Of course, an exhaustive comparison to these fields' literatures would be intractable, but a more detailed evaluation of the system's results is needed (preferably with the help of a domain expert) before we can make claims of originality.

We can evaluate our results with ASTRA on two main fronts. The first concerns false negatives or errors of omission. Here the system generally fares well, in that it has found nearly all basic reactions we have seen in texts, along with all pathways for proton, electron, and neutron capture, for neutron and deuteron creation, and for helium, carbon, and oxygen production. The few omissions all concern endothermic (energy-using) reactions, which we forbid in most runs in order to focus our efforts on the more common exothermic reactions.

Nevertheless, some endothermic reactions play an important role in the nucleosynthesis of heavier elements, particularly when they lead to highly exothermic reactions. For instance, astrophysicists hypothesize that the reaction ${}^4\text{He} + {}^4\text{He} \rightarrow {}^8\text{B} + -0.1 \text{ MeV}$ enables the carbon-producing reaction ${}^4\text{He} + {}^8\text{Be} \rightarrow {}^{12}\text{C}$. As we have noted, ASTRA can formulate reactions in any selected energy band, and runs that let the system accept reactions with Q values to -3.0 MeV have produced:



However, we want ASTRA to consider plausible endothermic reactions and to ignore implausible ones.⁴ Fortunately, the abundance of such reactions depends on the energy produced by colliding particles and nuclei in the stellar medium, which in turn depends on the star’s size and chemical composition. Given this energy level, we can constrain future runs to consider only those endothermic reactions that level would support.

The second main issue concerns false positives or errors of commission, and here issues of evaluation become more complex. ASTRA generates only those reactions consistent with quantum theory but, as we have noted, astrophysicists do not find all possible reactions equally plausible, apparently using the low rates of many component reactions to rule out the vast majority of possible pathways. Without these rate constraints, the system generates orders of magnitude more reaction pathways than appear in astrophysical papers, the vast majority of which would certainly be unacceptable to scientists.

The current version of ASTRA can use reaction rates to rule out candidates, retaining only those reactions with the highest rates to construct reaction pathways. But the rate for each reaction must be given by the user, as the current version cannot calculate them. Preliminary studies with hydrogen-burning reactions show that, taking this approach, ASTRA generates the *pp1* chain described in Section 4.3, but rules out the many other pathways it normally produces. Still, the existing system requires manual entry of reaction rates, and giving ASTRA the ability to calculate realistic reaction rates should be our highest priority for future work. Such calculations will require selecting a certain stellar model, with assumptions about temperature distribution and abundances of elements in the plasma, and we are pursuing collaboration with an expert in nuclear astrophysics to incorporate the knowledge needed to automate this process.

Despite ASTRA’s current limitations, we remain convinced that a computational aid of this sort could aid astrophysicists significantly, since they cannot generate systematically, on their own, all reactions that satisfy their criteria for plausibility. For example, Fowler et al. (1967, 1970, 1983) mention 88 reactions for elements from *H* to ¹⁶*O* in their research, whereas ASTRA generates 276 such reactions. Within these 88 reactions, the authors cite 33 hydrogen captures, 17 neutron captures, and 8 deuteron fusion reactions, whereas the system formulates 46 hydrogen captures, 59 neutron captures, and 75 deuteron fusions for the same range of elements.

Similarly, ASTRA can search a much larger space of reaction pathways than can human scientists and, although the processes for hydrogen and helium burning have been dealt with extensively in the literature, there may still be room for contributions, especially concerning the synthesis of

4. Astrophysicists usually obtain endothermic reactions by reversing known exothermic reactions, but ASTRA can directly generate such reactions without missing any possible candidates.

heavier elements from oxygen to iron, which also take place in stellar environments. This seems a likely area in which a systematic discovery system like ASTRA could aid the process of explanation generation. We intend to invoke the extended system in these relatively unexplored regions and to recruit astrophysicists to evaluate the plausibility of the resulting process models, with the ultimate aim of contributing to the scientific literature in this intriguing field.

6. Related Research on Computational Scientific Discovery

Our approach to computational discovery draws many ideas from previous work on the topic. As we noted earlier, our system shares many goals and techniques with Valdés-Pérez' MECHEM, so it seems worth examining their similarities and differences in some depth. In terms of their tasks, both MECHEM and ASTRA are designed to generate candidate reactions and reaction mechanisms in combinatorial spaces, but the first system focuses on the domains of organometallic chemistry and catalysis, whereas our own addresses nuclear astrophysics. Both are designed as computational aids for scientists, but MECHEM includes a graphical interface that lets users influence system behavior, whereas ASTRA has a less flexible non-graphical interface.

The two systems also differ somewhat in terms of their inputs and outputs. MECHEM requires the inputs and outputs of some reaction, for which it generates all the simplest reaction pathways that meet known constraints. The program also accepts a network of chemical reactions generated by CHEMNET (Bruck et al., 1998), as well as optional constraints specified by the user. Analogously, ASTRA requires information, in terms of five quantum numbers, about the elementary particles and isotopes of elements from hydrogen to oxygen, along with the initial and final elements; from these it generates all possible reactions that satisfy the quantum constraints, along with all acceptable reaction pathways. The system also accepts constraints from the user about known reaction rates and acceptable energy bands for reactions.

The differences between the two systems' internal representations are more significant. MECHEM represents molecules as graphs and denotes reactions as ordered lists of reactants and products that specify the quantity for each substance and, when known, the associated energy values. Reaction pathways are lists of reaction steps, which in practice have at most two reactants and two products. In contrast, ASTRA represents elements and particles as atomic symbols, reactions as ordered lists with associated energy emissions, and pathways as a series of reaction links for each element that takes part in the chain. Reactions can involve from one to three substances in both the reactants and the results.

Both systems carry out search through the space of reaction pathways, but they proceed by quite different steps. MECHEM uses a breadth-first approach, iterating through the number of substances S and, for a given S , iterating through the number of reaction steps R , hypothesizing a new unseen substance each time it increments S . For each combination of S and R , the system considers all ways the substances might occur in the posited reactions, using the constraints implied by the reactions to solve for the formulae of unseen substances. MECHEM tests each such combination to determine whether it satisfies additional constraints, continuing to search until finding one or more pathways that explains how the reactants are transformed into the results. In contrast, ASTRA carries out a depth-first search, starting from the resulting products and chaining backward until it

finds a pathway that links them, using known reactions, to the initial elements. The system has no explicit bias toward simpler hypotheses, continuing its search until finds all reaction chains under a specified length, but it can take advantage of user-specified rates to ignore any unlikely reactions.

Another close relative of ASTRA is Hendrickson's (1995) SYNGEN, which addresses the task of chemical synthesis, where one must determine not only the reaction paths but also the starting molecules. The constraints for this domain are more similar to those used by MECHEM, but the system's search framework is more closely akin to our own. In particular, both ASTRA and SYNGEN generate candidate pathways by chaining backward from the final products, using known reactions, until reaching acceptable starting substances.

Our system differs from SYNGEN in its focus on astrophysics and in its ability to generate the basic reactions from the elements involved and the principles of quantum physics. ASTRA inherits this latter ability from our previous BR-4 system (Kocabas & Langley, 1995), which carries out theory revision in particle physics, much like its predecessor BR-3 (Kocabas, 1991). The BR-3 system in turn descends directly from STAHL (Zytkow & Simon, 1986) and STAHLp (Rose & Langley, 1986), which modeled qualitative discovery in chemistry. Unlike its ancestors, BR-4 includes a module that predicts reactions from its current theory, which forms the basis for ASTRA's capacity along these lines.

Other discovery systems that formulate process theories seem less closely related. Kulkarni and Simon's (1990) KEKADA generates reaction pathways but relies on experimental data to determine intermediate steps. Karp (1990), Rajamoney (1990), and O'Rorke, Morris, and Schulenberg (1990) describe systems that produce process explanations in biology, physics, and chemistry, respectively, but all start with given accounts and revise them in response to unexpected observations. Our work has focused on generating process models from a deeper theory, rather than on their revision.

7. Concluding Remarks

In this paper we described ASTRA, which we designed to serve as an aid for astrophysical research. Given a set of elements, isotopes, and particles, the system determines all valid reactions among these entities that are consistent with quantum theory. We found that ASTRA generates all reactions we have seen in the astrophysics literature involving proton, electron, and neutron captures, as well as neutron and deuteron production. Moreover, given an element observed within stars and the likely starting elements, the algorithm generates all reaction pathways that can explain the latter's transformation into the former. Again, studies reveal that ASTRA reproduces all reaction pathways that, to our knowledge, scientists have proposed for the creation of helium, carbon, and oxygen.

However, the system also generates many reactions and pathways that we have not found in the scientific literature. Discussions with an expert in astrophysics indicate that some of these results hold theoretical interest, in that they may provide alternatives, in certain stellar conditions, to generally accepted accounts. But the vast majority of generated reactions and pathways appear to be implausible due to low reaction rates. This suggests that ASTRA needs to make better use of constraints about these factors, and we intend to incorporate them in future versions and thus make the system a more useful research tool for astrophysicists.

Acknowledgements

We thank Raul Valdés-Pérez for discussions about the relation between ASTRA and MECHEM, as well as the anonymous reviewers for comments on an earlier draft. An earlier version of this paper appeared in the *Proceedings of the ECAI-98 Workshop on Machine Discovery*. Brighton, UK.

References

- Adouze, J., & Vauclair, S. (1980). *An introduction to nuclear astrophysics*. Holland: D. Riedel.
- Bruk, L. G., Gorodskii, S., Zeigarnik, A. V. Valdes-Perez, R. E., & Temkin, O. N. (1998). Oxydative carbonylation of phenylacetylene catalyzed by Pd(II) and Cu(I): Experimental tests of forty-one computer generated mechanistic hypotheses. *Journal of Molecular Catalysis A: Chemical*, *130*, 29–40.
- Clayton, D. D. (1983). *Principles of stellar evolution and nucleosynthesis*. Chicago: University of Chicago Press.
- Cujec, B., & Fowler, W. A. (1980). Neglect of D , T , and ${}^3\text{He}$ in advanced stellar evolution. *The Astrophysical Journal*, *236*, 658–660.
- Fowler, W. A. (1986). The synthesis of the chemical elements carbon and oxygen. In S. L. Shapiro & S. A. Teukolsky (Eds.), *Highlights of modern astrophysics*. New York: John Wiley & Sons.
- Fowler, W. A., Caughlan, G. R., & Zimmermann, B. (1967). Thermonuclear reaction rates. *Annual Review of Astronomy and Astrophysics*, *5*, 525–570.
- Fowler, W. A., Caughlan, G. R., & Zimmermann, B. (1975). Thermonuclear reaction rates: II. *Annual Review of Astronomy and Astrophysics*, *13*, 69–112.
- Harris, M. J., Fowler, W. A., Caughlan, G. R., & Zimmerman, B. (1983). Thermonuclear reaction rates: III. *Annual Review of Astronomy and Astrophysics*, *21*, 165–176.
- Hendrickson, J. B. (1995). Systematic synthesis design: The SYNGEN program. *Working Notes of the AAAI Spring Symposium on Systematic Methods of Scientific Discovery* (pp. 13–17). Stanford, CA: AAAI Press.
- Jones, R. (1986). Generating predictions to aid the scientific discovery process. *Proceedings of the Fifth National Conference of on Artificial Intelligence* (pp. 513–517). Philadelphia: Morgan Kaufmann.
- Karp, P. D. (1990). Hypothesis formation as design. In J. Shrager & P. Langley (Eds.), *Computational models of scientific discovery and theory formation*. San Mateo, CA: Morgan Kaufmann.
- Kippenhahn, R., & Weigert, A. (1994). *Stellar structure and evolution*. London: Springer-Verlag.
- Kocabas, S. (1991). Conflict resolution as discovery in particle physics. *Machine Learning*, *6*, 277–309.
- Kocabas, S., & Langley, P. (1995). Integration of research tasks for modeling discoveries in particle physics. *Working notes of the AAAI Spring Symposium on Systematic Methods of Scientific Discovery* (pp. 87–92). Stanford, CA: AAAI Press.

- Kulkarni, D., & Simon, H. A. (1990). Experimentation in machine discovery. In J. Shrager & P. Langley (Eds.), *Computational models of scientific discovery and theory formation*. San Mateo, CA: Morgan Kaufmann.
- Lang, K. R. (1974). *Astrophysical formulae: A compendium for physicists and astrophysicists*. New York: Springer-Verlag.
- Langley, P. (1981). Data-driven discovery of physical laws. *Cognitive Science*, *5*, 31–54.
- Langley, P. (1998). The computer-aided discovery of scientific knowledge. *Proceedings of the First International Conference on Discovery Science*. Fukuoka, Japan: Springer.
- Langley, P., Simon, H. A., Bradshaw, G. L., & Zytkow, J. M. (1987). *Scientific discovery: Computational explorations of the creative processes*. Cambridge, MA: MIT Press.
- O'Rourke, P., Morris, S., & Schulenberg, D. (1990). Theory formation by abduction: A case study based on the chemical revolution. In J. Shrager & P. Langley (Eds.), *Computational models of scientific discovery and theory formation*. San Mateo, CA: Morgan Kaufmann.
- Rajamoney, S. (1990). A computational approach to theory revision. In J. Shrager & P. Langley (Eds.), *Computational models of scientific discovery and theory formation*. San Mateo, CA: Morgan Kaufmann.
- Rose, D., & Langley, P. (1986). Chemical discovery as belief revision. *Machine Learning*, *1*, 423–451.
- Valdés-Pérez, R. E. (1995). Machine discovery in chemistry: New results. *Artificial Intelligence*, *74*, 191–201.
- Weast, R. C., & Astle, M. J. (Eds.). (1981). *CRC handbook of chemistry and physics* (62nd ed.). Florida: CRC Press.
- Williams, W. S. C. (1991). *Nuclear and particle physics*. Oxford: Clarendon Press.
- Zeigarnik, A. V., Valdés-Pérez, R. E., Temkin, O. N., Bruk, L. G., & Shalgunov, S. I. (1997). Computer-aided mechanism elucidation of acetylene hydrocarboxylation to acrylic acid based on a novel union of empirical and formal methods. *Organometallics*, *16*, 3114–3127.
- Zytkow, J. M., & Simon, H. A. (1986). A theory of historical discovery: The construction of componential models. *Machine Learning*, *1*, 107–137.