Trust by design – the principles

Peter Fettke and Wolfgang Reisig

Deutsches Forschungszentrum für Künstliche Intelligenz
German Research Center for Artificial Intelligence

UNIVERSITÄT DES SAARLANDES

HUMBOLDT-UNIVERSITÄT ZU BERLIN

20th Symposium and Summer School
on Service-oriented Computing (SummerSoC)
Crete, Greece
June 29th – July 4th, 2026

20th anniversary of SummerSOC!

20th anniversary of SummerSOC! And my 5th anniversary at SummerSOC

2022	Systems Mining: from process logs to system models
2023	Next generation modeling in the era of cyber-physical systems
2024	Large process models
2025	Markets: a paradigm for agents in the digital world
2026	Trust by design

Agenda

1. Motivation and background
2. Trust certificates
3. An example
4. Conclusions

dfki ai Deutsches Forschungszentrum für Künstliche Intelligenz
German Research Center for Artificial Intelligence

UNIVERSITÄT DES SAARLANDES

HUMBOLDT-UNIVERSITÄT ZU BERLIN

AUTOBIZ: Pushing the Boundaries of AI-Driven Process Execution and Adaptation, Dagstuhl Seminar 25192, May 4th – 9th, 2025

Agentic Business Process Management: A research manifesto, in: Information Systems (2026)

Information Systems 140 (2026) 102738

Contents lists available at ScienceDirect

Information Systems

journal homepage: www.elsevier.com/locate/is

Agentic Business Process Management: A research manifesto☆,☆☆

Diego Calvanese [a], Angelo Casciani [g], Giuseppe De Giacomo [b], Marlon Dumas [c], Fabiana Fournier [d], Timotheus Kampik [e,f,*], Emanuele La Malfa [b], Lior Limonad [d], Andrea Marrella [g], Andreas Metzger [h], Marco Montali [a], Daniel Amyot [i], Peter Fettke [j,k], Artem Polyvyanyy [l], Stefanie Rinderle-Ma [m], Sebastian Sardiña [n], Niek Tax [o], Barbara Weber [p]

[a] Free University of Bozen-Bolzano, Bozen-Bolzano, Italy
[b] University of Oxford, Oxford, United Kingdom
[c] University of Tartu, Tartu, Estonia
[d] IBM Research, Haifa, Israel
[e] Umeå University, Umeå, Sweden
[f] SAP, Berlin, Germany
[g] Sapienza Università di Roma, Rome, Italy
[h] paluno (Ruhr Institute for Software Technology), University of Duisburg Essen, Essen, Germany
[i] University of Ottawa, Ottawa, Canada
[j] German Research Center for Artificial Intelligence (DFKI), 66123 Saarbrücken, Germany
[k] Saarland University, 66123 Saarbrücken, Germany
[l] The University of Melbourne, Parkville, VIC 3010, Australia
[m] TUM School of Computation, Information, and Technology, TU Munich, Garching, Germany
[n] RMIT University, Melbourne, Australia
[o] Meta, London, United Kingdom
[p] University of St. Gallen, St. Gallen, Switzerland

ARTICLE INFO

Keywords:
Business process management
Autonomous agents
Agentic AI
Framed autonomy
Explainability
Conversational actionability
Self-modification

ABSTRACT

This paper presents a manifesto that articulates the conceptual foundations of *Agentic Business Process Management* (APM), an extension of Business Process Management (BPM) for governing autonomous agents executing processes in organizations. From a management perspective, APM represents a paradigm shift from the traditional view on business processes. This shift is driven by the realization of process awareness by agent-oriented abstractions: software and human agents act as primary functional entities that perceive, reason, and act within explicit process frames. Thus, APM moves away from automation-oriented BPM towards systems in which autonomy is constrained, aligned, and made operational through *process aware* agents.

We introduce the core abstractions and architectural elements required to realize APM systems and elaborate on four key capabilities that agents in APM systems must support: *framed autonomy*, *explainability*, *conversational actionability*, and *self-modification*. These capabilities jointly ensure that agents' goals are aligned with organizational goals and that agents behave in a framed yet proactive manner in pursuing those goals. We discuss the extent to which the capabilities can be realized and identify research challenges whose resolution requires further advances in BPM, AI, and multi-agent systems. The manifesto thus serves as a roadmap for bridging these communities and for guiding the development of APM systems in practice.

☆ This article is part of a Special issue entitled: 'Autonomous Process Execution Systems' published in Information Systems.

☆☆ Given the role as Editor-in-Chief of this journal, Stefanie Rinderle-Ma had no involvement in the peer-review of this article and has no access to information regarding its peer-review. Full responsibility for the editorial process for this article was delegated to another journal editor.

* Corresponding author.
E-mail addresses: diego.calvanese@unibz.it (D. Calvanese), casciani@diag.uniroma1.it (A. Casciani), giuseppe.degiacomo@cs.ox.ac.uk (G. De Giacomo), marlon.dumas@ut.ee (M. Dumas), fabiana@il.ibm.com (F. Fournier), tkampik@cs.umu.se (T. Kampik), emanuele.lamalfa@cs.ox.ac.uk (E.L. Malfa), liorli@il.ibm.com (L. Limonad), marrella@diag.uniroma1.it (A. Marrella), andreas.metzger@paluno.uni-due.de (A. Metzger), marco.montali@unibz.it (M. Montali), damyot@uottawa.ca (D. Amyot), peter.fettke@dfki.de (P. Fettke), artem.polyvyanyy@unimelb.edu.au (A. Polyvyanyy), stefanie.rinderle-ma@tum.de (S. Rinderle-Ma), sebastian.sardina@rmit.edu.au (S. Sardiña), niek@meta.com (N. Tax), barbara.weber@unisg.ch (B. Weber).

https://doi.org/10.1016/j.is.2026.102738
Received 11 December 2025; Received in revised form 8 April 2026; Accepted 8 April 2026
Available online 17 April 2026

Agentic Business Process Management: A research manifesto, in: Information Systems (2026)

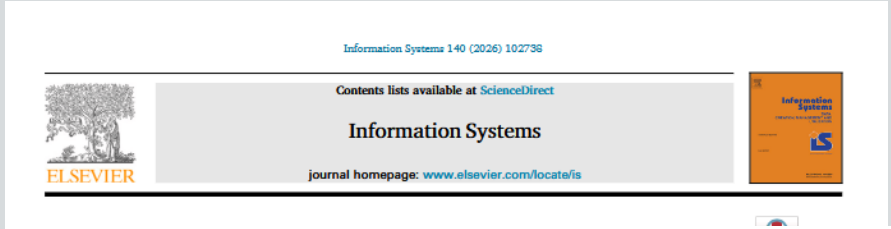
Information Systems 140 (2026) 102738

Contents lists available at ScienceDirect

Information Systems

journal homepage: www.elsevier.com/locate/is

Agentic Business Process Management: A research manifesto☆,☆☆

Diego Calvanese [a], Angelo Casciani [g], Giuseppe De Giacomo [b], Marlon Dumas [c], Fabiana Fournier [d], Timotheus Kampik [e,f,*], Emanuele La Malfa [b], Lior Limonad [d], Andrea Marrella [g], Andreas Metzger [h], Marco Montali [a], Daniel Amyot [i], Peter Fettke [j,k], Artem Polyvyanyy [l], Stefanie Rinderle-Ma [m], Sebastian Sardiña [n], Niek Tax [o], Barbara Weber [p]

[a] Free University of Bozen-Bolzano, Bozen-Bolzano, Italy
[b] University of Oxford, Oxford, United Kingdom
[c] University of Tartu, Tartu, Estonia
[d] IBM Research, Haifa, Israel
[e] Umeå University, Umeå, Sweden
[f] SAP, Berlin, Germany
[g] Sapienza Università di Roma, Rome, Italy
[h] paluno (Ruhr Institute for Software Technology), University of Duisburg Essen, Essen, Germany
[i] University of Ottawa, Ottawa, Canada
[j] German Research Center for Artificial Intelligence (DFKI), 66123 Saarbrücken, Germany
[k] Saarland University, 66123 Saarbrücken, Germany
[l] The University of Melbourne, Parkville, VIC 3010, Australia
[m] TUM School of Computation, Information, and Technology, TU Munich, Garching, Germany
[n] RMIT University, Melbourne, Australia
[o] Meta, London, United Kingdom
[p] University of St. Gallen, St. Gallen, Switzerland

Agentic Business Process Management: A research manifesto, in: Information Systems (2026)

Information Systems 140 (2026) 102738

Contents lists available at ScienceDirect

Information Systems

journal homepage: www.elsevier.com/locate/is

Agentic Business Process Management: A research manifesto☆,☆☆

Diego Calvanese [a], Angelo Casciani [g], Giuseppe De Giacomo [b], Marlon Dumas [c], Fabiana Fournier [d], Timotheus Kampik [e,f,*], Emanuele La Malfa [b], Lior Limonad [d], Andrea Marrella [g], Andreas Metzger [h], Marco Montali [a], Daniel Amyot [i], Peter Fettke [j,k], Artem Polyvyanyy [l], Stefanie Rinderle-Ma [m], Sebastian Sardiña [n], Niek Tax [o], Barbara Weber [p]

[a] Free University of Bozen-Bolzano, Bozen-Bolzano, Italy
[b] University of Oxford, Oxford, United Kingdom
[c] University of Tartu, Tartu, Estonia
[d] IBM Research, Haifa, Israel
[e] Umeå University, Umeå, Sweden
[f] SAP, Berlin, Germany
[g] Sapienza Università di Roma, Rome, Italy
[h] paluno (Ruhr Institute for Software Technology), University of Duisburg Essen, Essen, Germany
[i] University of Ottawa, Ottawa, Canada
[j] German Research Center for Artificial Intelligence (DFKI), 66123 Saarbrücken, Germany
[k] Saarland University, 66123 Saarbrücken, Germany
[l] The University of Melbourne, Parkville, VIC 3010, Australia
[m] TUM School of Computation, Information, and Technology, TU Munich, Garching, Germany
[n] RMIT University, Melbourne, Australia
[o] Meta, London, United Kingdom
[p] University of St. Gallen, St. Gallen, Switzerland

ARTICLE INFO

Keywords:
Business process management
Autonomous agents
Agentic AI
Framed autonomy
Explainability
Conversational actionability
Self-modification

ABSTRACT

This paper presents a manifesto that articulates the conceptual foundations of *Agentic Business Process Management* (APM), an extension of Business Process Management (BPM) for governing autonomous agents executing processes in organizations. From a management perspective, APM represents a paradigm shift from the traditional view on business processes. This shift is driven by the realization of process awareness by agent-oriented abstractions: software and human agents act as primary functional entities that perceive, reason, and act within explicit process frames. Thus, APM moves away from automation-oriented BPM towards systems in which autonomy is constrained, aligned, and made operational through *process aware* agents.

We introduce the core abstractions and architectural elements required to realize APM systems and elaborate on four key capabilities that agents in APM systems must support: *framed autonomy*, *explainability*, *conversational actionability*, and *self-modification*. These capabilities jointly ensure that agents' goals are aligned with organizational goals and that agents behave in a framed yet proactive manner in pursuing those goals. We discuss the extent to which the capabilities can be realized and identify research challenges whose resolution requires further advances in BPM, AI, and multi-agent systems. The manifesto thus serves as a roadmap for bridging these communities and for guiding the development of APM systems in practice.

☆ This article is part of a Special issue entitled: 'Autonomous Process Execution Systems' published in Information Systems.

☆☆ Given the role as Editor-in-Chief of this journal, Stefanie Rinderle-Ma had no involvement in the peer-review of this article and has no access to information regarding its peer-review. Full responsibility for the editorial process for this article was delegated to another journal editor.

* Corresponding author.
E-mail addresses: diego.calvanese@unibz.it (D. Calvanese), casciani@diag.uniroma1.it (A. Casciani), giuseppe.degiacomo@cs.ox.ac.uk (G. De Giacomo), marlon.dumas@ut.ee (M. Dumas), fabiana@il.ibm.com (F. Fournier), tkampik@cs.umu.se (T. Kampik), emanuele.lamalfa@cs.ox.ac.uk (E.L. Malfa), liorli@il.ibm.com (L. Limonad), marrella@diag.uniroma1.it (A. Marrella), andreas.metzger@paluno.uni-due.de (A. Metzger), marco.montali@unibz.it (M. Montali), damyot@uottawa.ca (D. Amyot), peter.fettke@dfki.de (P. Fettke), artem.polyvyanyy@unimelb.edu.au (A. Polyvyanyy), stefanie.rinderle-ma@tum.de (S. Rinderle-Ma), sebastian.sardina@rmit.edu.au (S. Sardiña), niek@meta.com (N. Tax), barbara.weber@unisg.ch (B. Weber).

https://doi.org/10.1016/j.is.2026.102738
Received 11 December 2025; Received in revised form 8 April 2026; Accepted 8 April 2026
Available online 17 April 2026

LLM Agent

=

LLM + memory + task planning + action + tool use

Agentic Business Process Management: A research manifesto, in: Information Systems (2026)

Information Systems 140 (2026) 102738

Contents lists available at ScienceDirect

Information Systems

journal homepage: www.elsevier.com/locate/is

Agentic Business Process Management: A research manifesto☆,☆☆

Diego Calvanese [a], Angelo Casciani [g], Giuseppe De Giacomo [b], Marlon Dumas [c], Fabiana Fournier [d], Timotheus Kampik [e,f,*], Emanuele La Malfa [b], Lior Limonad [d], Andrea Marrella [g], Andreas Metzger [h], Marco Montali [a], Daniel Amyot [i], Peter Fettke [j,k], Artem Polyvyanyy [l], Stefanie Rinderle-Ma [m], Sebastian Sardiña [n], Niek Tax [o], Barbara Weber [p]

[a] Free University of Bozen-Bolzano, Bozen-Bolzano, Italy
[b] University of Oxford, Oxford, United Kingdom
[c] University of Tartu, Tartu, Estonia
[d] IBM Research, Haifa, Israel
[e] Umeå University, Umeå, Sweden
[f] SAP, Berlin, Germany
[g] Sapienza Università di Roma, Rome, Italy
[h] paluno (Ruhr Institute for Software Technology), University of Duisburg Essen, Essen, Germany
[i] University of Ottawa, Ottawa, Canada
[j] German Research Center for Artificial Intelligence (DFKI), 66123 Saarbrücken, Germany
[k] Saarland University, 66123 Saarbrücken, Germany
[l] The University of Melbourne, Parkville, VIC 3010, Australia
[m] TUM School of Computation, Information, and Technology, TU Munich, Garching, Germany
[n] RMIT University, Melbourne, Australia
[o] Meta, London, United Kingdom
[p] University of St. Gallen, St. Gallen, Switzerland

ARTICLE INFO

Keywords:
Business process management
Autonomous agents
Agentic AI
Framed autonomy
Explainability
Conversational actionability
Self-modification

ABSTRACT

This paper presents a manifesto that articulates the conceptual foundations of *Agentic Business Process Management* (APM), an extension of Business Process Management (BPM) for governing autonomous agents executing processes in organizations. From a management perspective, APM represents a paradigm shift from the traditional view on business processes. This shift is driven by the realization of process awareness by agent-oriented abstractions: software and human agents act as primary functional entities that perceive, reason, and act within explicit process frames. Thus, APM moves away from automation-oriented BPM towards systems in which autonomy is constrained, aligned, and made operational through *process aware* agents.

We introduce the core abstractions and architectural elements required to realize APM systems and elaborate on four key capabilities that agents in APM systems must support: *framed autonomy*, *explainability*, *conversational actionability*, and *self-modification*. These capabilities jointly ensure that agents' goals are aligned with organizational goals and that agents behave in a framed yet proactive manner in pursuing those goals. We discuss the extent to which the capabilities can be realized and identify research challenges whose resolution requires further advances in BPM, AI, and multi-agent systems. The manifesto thus serves as a roadmap for bridging these communities and for guiding the development of APM systems in practice.

☆ This article is part of a Special issue entitled: 'Autonomous Process Execution Systems' published in Information Systems.

☆☆ Given the role as Editor-in-Chief of this journal, Stefanie Rinderle-Ma had no involvement in the peer-review of this article and has no access to information regarding its peer-review. Full responsibility for the editorial process for this article was delegated to another journal editor.

* Corresponding author.
E-mail addresses: diego.calvanese@unibz.it (D. Calvanese), casciani@diag.uniroma1.it (A. Casciani), giuseppe.degiacomo@cs.ox.ac.uk (G. De Giacomo), marlon.dumas@ut.ee (M. Dumas), fabiana@il.ibm.com (F. Fournier), tkampik@cs.umu.se (T. Kampik), emanuele.lamalfa@cs.ox.ac.uk (E.L. Malfa), liorli@il.ibm.com (L. Limonad), marrella@diag.uniroma1.it (A. Marrella), andreas.metzger@paluno.uni-due.de (A. Metzger), marco.montali@unibz.it (M. Montali), damyot@uottawa.ca (D. Amyot), peter.fettke@dfki.de (P. Fettke), artem.polyvyanyy@unimelb.edu.au (A. Polyvyanyy), stefanie.rinderle-ma@tum.de (S. Rinderle-Ma), sebastian.sardina@rmit.edu.au (S. Sardiña), niek@meta.com (N. Tax), barbara.weber@unisg.ch (B. Weber).

https://doi.org/10.1016/j.is.2026.102738
Received 11 December 2025; Received in revised form 8 April 2026; Accepted 8 April 2026
Available online 17 April 2026

LLM Agent

=

LLM + memory + task planning + action + tool use

Two examples

1. organizing a business trip

Agentic Business Process Management: A research manifesto, in: Information Systems (2026)

Information Systems 140 (2026) 102738

Contents lists available at ScienceDirect

Information Systems

journal homepage: www.elsevier.com/locate/is

Agentic Business Process Management: A research manifesto☆,☆☆

Diego Calvanese [a], Angelo Casciani [g], Giuseppe De Giacomo [b], Marlon Dumas [c], Fabiana Fournier [d], Timotheus Kampik [e,f,*], Emanuele La Malfa [b], Lior Limonad [d], Andrea Marrella [g], Andreas Metzger [h], Marco Montali [a], Daniel Amyot [i], Peter Fettke [j,k], Artem Polyvyanyy [l], Stefanie Rinderle-Ma [m], Sebastian Sardiña [n], Niek Tax [o], Barbara Weber [p]

[a] Free University of Bozen-Bolzano, Bozen-Bolzano, Italy
[b] University of Oxford, Oxford, United Kingdom
[c] University of Tartu, Tartu, Estonia
[d] IBM Research, Haifa, Israel
[e] Umeå University, Umeå, Sweden
[f] SAP, Berlin, Germany
[g] Sapienza Università di Roma, Rome, Italy
[h] paluno (Ruhr Institute for Software Technology), University of Duisburg Essen, Essen, Germany
[i] University of Ottawa, Ottawa, Canada
[j] German Research Center for Artificial Intelligence (DFKI), 66123 Saarbrücken, Germany
[k] Saarland University, 66123 Saarbrücken, Germany
[l] The University of Melbourne, Parkville, VIC 3010, Australia
[m] TUM School of Computation, Information, and Technology, TU Munich, Garching, Germany
[n] RMIT University, Melbourne, Australia
[o] Meta, London, United Kingdom
[p] University of St. Gallen, St. Gallen, Switzerland

ARTICLE INFO

Keywords:
Business process management
Autonomous agents
Agentic AI
Framed autonomy
Explainability
Conversational actionability
Self-modification

ABSTRACT

This paper presents a manifesto that articulates the conceptual foundations of *Agentic Business Process Management* (APM), an extension of Business Process Management (BPM) for governing autonomous agents executing processes in organizations. From a management perspective, APM represents a paradigm shift from the traditional view on business processes. This shift is driven by the realization of process awareness by agent-oriented abstractions: software and human agents act as primary functional entities that perceive, reason, and act within explicit process frames. Thus, APM moves away from automation-oriented BPM towards systems in which autonomy is constrained, aligned, and made operational through *process aware* agents.

We introduce the core abstractions and architectural elements required to realize APM systems and elaborate on four key capabilities that agents in APM systems must support: *framed autonomy*, *explainability*, *conversational actionability*, and *self-modification*. These capabilities jointly ensure that agents' goals are aligned with organizational goals and that agents behave in a framed yet proactive manner in pursuing those goals. We discuss the extent to which the capabilities can be realized and identify research challenges whose resolution requires further advances in BPM, AI, and multi-agent systems. The manifesto thus serves as a roadmap for bridging these communities and for guiding the development of APM systems in practice.

☆ This article is part of a Special issue entitled: 'Autonomous Process Execution Systems' published in Information Systems.

☆☆ Given the role as Editor-in-Chief of this journal, Stefanie Rinderle-Ma had no involvement in the peer-review of this article and has no access to information regarding its peer-review. Full responsibility for the editorial process for this article was delegated to another journal editor.

* Corresponding author.
E-mail addresses: diego.calvanese@unibz.it (D. Calvanese), casciani@diag.uniroma1.it (A. Casciani), giuseppe.degiacomo@cs.ox.ac.uk (G. De Giacomo), marlon.dumas@ut.ee (M. Dumas), fabiana@il.ibm.com (F. Fournier), tkampik@cs.umu.se (T. Kampik), emanuele.lamalfa@cs.ox.ac.uk (E.L. Malfa), liorli@il.ibm.com (L. Limonad), marrella@diag.uniroma1.it (A. Marrella), andreas.metzger@paluno.uni-due.de (A. Metzger), marco.montali@unibz.it (M. Montali), damyot@uottawa.ca (D. Amyot), peter.fettke@dfki.de (P. Fettke), artem.polyvyanyy@unimelb.edu.au (A. Polyvyanyy), stefanie.rinderle-ma@tum.de (S. Rinderle-Ma), sebastian.sardina@rmit.edu.au (S. Sardiña), niek@meta.com (N. Tax), barbara.weber@unisg.ch (B. Weber).

https://doi.org/10.1016/j.is.2026.102738
Received 11 December 2025; Received in revised form 8 April 2026; Accepted 8 April 2026
Available online 17 April 2026

LLM Agent

=

LLM + memory + task planning + action + tool use

Two examples

1. organizing a business trip
2. onboarding a new supplier as part of a procurement process

Agentic Business Process Management: A research manifesto, in: Information Systems (2026)

Information Systems 140 (2026) 102738

Contents lists available at ScienceDirect

Information Systems

journal homepage: www.elsevier.com/locate/is

Agentic Business Process Management: A research manifesto☆,☆☆

Diego Calvanese [a], Angelo Casciani [g], Giuseppe De Giacomo [b], Marlon Dumas [c], Fabiana Fournier [d], Timotheus Kampik [e,f,*], Emanuele La Malfa [b], Lior Limonad [d], Andrea Marrella [g], Andreas Metzger [h], Marco Montali [a], Daniel Amyot [i], Peter Fettke [j,k], Artem Polyvyanyy [l], Stefanie Rinderle-Ma [m], Sebastian Sardiña [n], Niek Tax [o], Barbara Weber [p]

[a] Free University of Bozen-Bolzano, Bozen-Bolzano, Italy
[b] University of Oxford, Oxford, United Kingdom
[c] University of Tartu, Tartu, Estonia
[d] IBM Research, Haifa, Israel
[e] Umeå University, Umeå, Sweden
[f] SAP, Berlin, Germany
[g] Sapienza Università di Roma, Rome, Italy
[h] paluno (Ruhr Institute for Software Technology), University of Duisburg Essen, Essen, Germany
[i] University of Ottawa, Ottawa, Canada
[j] German Research Center for Artificial Intelligence (DFKI), 66123 Saarbrücken, Germany
[k] Saarland University, 66123 Saarbrücken, Germany
[l] The University of Melbourne, Parkville, VIC 3010, Australia
[m] TUM School of Computation, Information, and Technology, TU Munich, Garching, Germany
[n] RMIT University, Melbourne, Australia
[o] Meta, London, United Kingdom
[p] University of St. Gallen, St. Gallen, Switzerland

ARTICLE INFO

Keywords:
Business process management
Autonomous agents
Agentic AI
Framed autonomy
Explainability
Conversational actionability
Self-modification

ABSTRACT

This paper presents a manifesto that articulates the conceptual foundations of *Agentic Business Process Management* (APM), an extension of Business Process Management (BPM) for governing autonomous agents executing processes in organizations. From a management perspective, APM represents a paradigm shift from the traditional view on business processes. This shift is driven by the realization of process awareness by agent-oriented abstractions: software and human agents act as primary functional entities that perceive, reason, and act within explicit process frames. Thus, APM moves away from automation-oriented BPM towards systems in which autonomy is constrained, aligned, and made operational through *process aware* agents.

We introduce the core abstractions and architectural elements required to realize APM systems and elaborate on four key capabilities that agents in APM systems must support: *framed autonomy*, *explainability*, *conversational actionability*, and *self-modification*. These capabilities jointly ensure that agents' goals are aligned with organizational goals and that agents behave in a framed yet proactive manner in pursuing those goals. We discuss the extent to which the capabilities can be realized and identify research challenges whose resolution requires further advances in BPM, AI, and multi-agent systems. The manifesto thus serves as a roadmap for bridging these communities and for guiding the development of APM systems in practice.

☆ This article is part of a Special issue entitled: 'Autonomous Process Execution Systems' published in Information Systems.

☆☆ Given the role as Editor-in-Chief of this journal, Stefanie Rinderle-Ma had no involvement in the peer-review of this article and has no access to information regarding its peer-review. Full responsibility for the editorial process for this article was delegated to another journal editor.

* Corresponding author.
E-mail addresses: diego.calvanese@unibz.it (D. Calvanese), casciani@diag.uniroma1.it (A. Casciani), giuseppe.degiacomo@cs.ox.ac.uk (G. De Giacomo), marlon.dumas@ut.ee (M. Dumas), fabiana@il.ibm.com (F. Fournier), tkampik@cs.umu.se (T. Kampik), emanuele.lamalfa@cs.ox.ac.uk (E.L. Malfa), liorli@il.ibm.com (L. Limonad), marrella@diag.uniroma1.it (A. Marrella), andreas.metzger@paluno.uni-due.de (A. Metzger), marco.montali@unibz.it (M. Montali), damyot@uottawa.ca (D. Amyot), peter.fettke@dfki.de (P. Fettke), artem.polyvyanyy@unimelb.edu.au (A. Polyvyanyy), stefanie.rinderle-ma@tum.de (S. Rinderle-Ma), sebastian.sardina@rmit.edu.au (S. Sardiña), niek@meta.com (N. Tax), barbara.weber@unisg.ch (B. Weber).

https://doi.org/10.1016/j.is.2026.102738
Received 11 December 2025; Received in revised form 8 April 2026; Accepted 8 April 2026
Available online 17 April 2026

agentic business process management system (APM) has four main capabilities

1. framing
2. conversational actionability
3. self-modification
4. explainability

Agentic Business Process Management: A research manifesto, in: Information Systems (2026)

Information Systems 140 (2026) 102738

Contents lists available at ScienceDirect

Information Systems

journal homepage: www.elsevier.com/locate/is

Agentic Business Process Management: A research manifesto☆,☆☆

Diego Calvanese [a], Angelo Casciani [g], Giuseppe De Giacomo [b], Marlon Dumas [c], Fabiana Fournier [d], Timotheus Kampik [e,f,*], Emanuele La Malfa [b], Lior Limonad [d], Andrea Marrella [g], Andreas Metzger [h], Marco Montali [a], Daniel Amyot [i], Peter Fettke [j,k], Artem Polyvyanyy [l], Stefanie Rinderle-Ma [m], Sebastian Sardiña [n], Niek Tax [o], Barbara Weber [p]

[a] Free University of Bozen-Bolzano, Bozen-Bolzano, Italy
[b] University of Oxford, Oxford, United Kingdom
[c] University of Tartu, Tartu, Estonia
[d] IBM Research, Haifa, Israel
[e] Umeå University, Umeå, Sweden
[f] SAP, Berlin, Germany
[g] Sapienza Università di Roma, Rome, Italy
[h] paluno (Ruhr Institute for Software Technology), University of Duisburg Essen, Essen, Germany
[i] University of Ottawa, Ottawa, Canada
[j] German Research Center for Artificial Intelligence (DFKI), 66123 Saarbrücken, Germany
[k] Saarland University, 66123 Saarbrücken, Germany
[l] The University of Melbourne, Parkville, VIC 3010, Australia
[m] TUM School of Computation, Information, and Technology, TU Munich, Garching, Germany
[n] RMIT University, Melbourne, Australia
[o] Meta, London, United Kingdom
[p] University of St. Gallen, St. Gallen, Switzerland

ARTICLE INFO

Keywords:
Business process management
Autonomous agents
Agentic AI
Framed autonomy
Explainability
Conversational actionability
Self-modification

ABSTRACT

This paper presents a manifesto that articulates the conceptual foundations of *Agentic Business Process Management* (APM), an extension of Business Process Management (BPM) for governing autonomous agents executing processes in organizations. From a management perspective, APM represents a paradigm shift from the traditional view on business processes. This shift is driven by the realization of process awareness by agent-oriented abstractions: software and human agents act as primary functional entities that perceive, reason, and act within explicit process frames. Thus, APM moves away from automation-oriented BPM towards systems in which autonomy is constrained, aligned, and made operational through *process aware* agents.

We introduce the core abstractions and architectural elements required to realize APM systems and elaborate on four key capabilities that agents in APM systems must support: *framed autonomy*, *explainability*, *conversational actionability*, and *self-modification*. These capabilities jointly ensure that agents' goals are aligned with organizational goals and that agents behave in a framed yet proactive manner in pursuing those goals. We discuss the extent to which the capabilities can be realized and identify research challenges whose resolution requires further advances in BPM, AI, and multi-agent systems. The manifesto thus serves as a roadmap for bridging these communities and for guiding the development of APM systems in practice.

☆ This article is part of a Special issue entitled: 'Autonomous Process Execution Systems' published in Information Systems.

☆☆ Given the role as Editor-in-Chief of this journal, Stefanie Rinderle-Ma had no involvement in the peer-review of this article and has no access to information regarding its peer-review. Full responsibility for the editorial process for this article was delegated to another journal editor.

* Corresponding author.
E-mail addresses: diego.calvanese@unibz.it (D. Calvanese), casciani@diag.uniroma1.it (A. Casciani), giuseppe.degiacomo@cs.ox.ac.uk (G. De Giacomo), marlon.dumas@ut.ee (M. Dumas), fabiana@il.ibm.com (F. Fournier), tkampik@cs.umu.se (T. Kampik), emanuele.lamalfa@cs.ox.ac.uk (E.L. Malfa), liorli@il.ibm.com (L. Limonad), marrella@diag.uniroma1.it (A. Marrella), andreas.metzger@paluno.uni-due.de (A. Metzger), marco.montali@unibz.it (M. Montali), damyot@uottawa.ca (D. Amyot), peter.fettke@dfki.de (P. Fettke), artem.polyvyanyy@unimelb.edu.au (A. Polyvyanyy), stefanie.rinderle-ma@tum.de (S. Rinderle-Ma), sebastian.sardina@rmit.edu.au (S. Sardiña), niek@meta.com (N. Tax), barbara.weber@unisg.ch (B. Weber).

https://doi.org/10.1016/j.is.2026.102738
Received 11 December 2025; Received in revised form 8 April 2026; Accepted 8 April 2026
Available online 17 April 2026

safety cage for an industrial robot

Figure source: https://www.kuk-automation.de/

Agentic Business Process Management: A research manifesto, in: Information Systems (2026)

Information Systems 140 (2026) 102738

Contents lists available at ScienceDirect

Information Systems

journal homepage: www.elsevier.com/locate/is

Agentic Business Process Management: A research manifesto☆,☆☆

Diego Calvanese [a], Angelo Casciani [g], Giuseppe De Giacomo [b], Marlon Dumas [c], Fabiana Fournier [d], Timotheus Kampik [e,f,*], Emanuele La Malfa [b], Lior Limonad [d], Andrea Marrella [g], Andreas Metzger [h], Marco Montali [a], Daniel Amyot [i], Peter Fettke [j,k], Artem Polyvyanyy [l], Stefanie Rinderle-Ma [m], Sebastian Sardiña [n], Niek Tax [o], Barbara Weber [p]

[a] Free University of Bozen-Bolzano, Bozen-Bolzano, Italy
[b] University of Oxford, Oxford, United Kingdom
[c] University of Tartu, Tartu, Estonia
[d] IBM Research, Haifa, Israel
[e] Umeå University, Umeå, Sweden
[f] SAP, Berlin, Germany
[g] Sapienza Università di Roma, Rome, Italy
[h] paluno (Ruhr Institute for Software Technology), University of Duisburg Essen, Essen, Germany
[i] University of Ottawa, Ottawa, Canada
[j] German Research Center for Artificial Intelligence (DFKI), 66123 Saarbrücken, Germany
[k] Saarland University, 66123 Saarbrücken, Germany
[l] The University of Melbourne, Parkville, VIC 3010, Australia
[m] TUM School of Computation, Information, and Technology, TU Munich, Garching, Germany
[n] RMIT University, Melbourne, Australia
[o] Meta, London, United Kingdom
[p] University of St. Gallen, St. Gallen, Switzerland

ARTICLE INFO

Keywords:
Business process management
Autonomous agents
Agentic AI
Framed autonomy
Explainability
Conversational actionability
Self-modification

ABSTRACT

This paper presents a manifesto that articulates the conceptual foundations of *Agentic Business Process Management* (APM), an extension of Business Process Management (BPM) for governing autonomous agents executing processes in organizations. From a management perspective, APM represents a paradigm shift from the traditional view on business processes. This shift is driven by the realization of process awareness by agent-oriented abstractions: software and human agents act as primary functional entities that perceive, reason, and act within explicit process frames. Thus, APM moves away from automation-oriented BPM towards systems in which autonomy is constrained, aligned, and made operational through *process aware* agents.

We introduce the core abstractions and architectural elements required to realize APM systems and elaborate on four key capabilities that agents in APM systems must support: *framed autonomy*, *explainability*, *conversational actionability*, and *self-modification*. These capabilities jointly ensure that agents' goals are aligned with organizational goals and that agents behave in a framed yet proactive manner in pursuing those goals. We discuss the extent to which the capabilities can be realized and identify research challenges whose resolution requires further advances in BPM, AI, and multi-agent systems. The manifesto thus serves as a roadmap for bridging these communities and for guiding the development of APM systems in practice.

☆ This article is part of a Special issue entitled: 'Autonomous Process Execution Systems' published in Information Systems.

☆☆ Given the role as Editor-in-Chief of this journal, Stefanie Rinderle-Ma had no involvement in the peer-review of this article and has no access to information regarding its peer-review. Full responsibility for the editorial process for this article was delegated to another journal editor.

* Corresponding author.
E-mail addresses: diego.calvanese@unibz.it (D. Calvanese), casciani@diag.uniroma1.it (A. Casciani), giuseppe.degiacomo@cs.ox.ac.uk (G. De Giacomo), marlon.dumas@ut.ee (M. Dumas), fabiana@il.ibm.com (F. Fournier), tkampik@cs.umu.se (T. Kampik), emanuele.lamalfa@cs.ox.ac.uk (E.L. Malfa), liorli@il.ibm.com (L. Limonad), marrella@diag.uniroma1.it (A. Marrella), andreas.metzger@paluno.uni-due.de (A. Metzger), marco.montali@unibz.it (M. Montali), damyot@uottawa.ca (D. Amyot), peter.fettke@dfki.de (P. Fettke), artem.polyvyanyy@unimelb.edu.au (A. Polyvyanyy), stefanie.rinderle-ma@tum.de (S. Rinderle-Ma), sebastian.sardina@rmit.edu.au (S. Sardiña), niek@meta.com (N. Tax), barbara.weber@unisg.ch (B. Weber).

https://doi.org/10.1016/j.is.2026.102738
Received 11 December 2025; Received in revised form 8 April 2026; Accepted 8 April 2026
Available online 17 April 2026

agentic business process management system (APM) has four main capabilities

1. framing
2. conversational actionability
3. self-modification
4. explainability

Agentic Business Process Management: A research manifesto, in: Information Systems (2026)

Information Systems 140 (2026) 102738

Contents lists available at ScienceDirect

Information Systems

journal homepage: www.elsevier.com/locate/is

Agentic Business Process Management: A research manifesto☆,☆☆

Diego Calvanese [a], Angelo Casciani [g], Giuseppe De Giacomo [b], Marlon Dumas [c], Fabiana Fournier [d], Timotheus Kampik [e,f,*], Emanuele La Malfa [b], Lior Limonad [d], Andrea Marrella [g], Andreas Metzger [h], Marco Montali [a], Daniel Amyot [i], Peter Fettke [j,k], Artem Polyvyanyy [l], Stefanie Rinderle-Ma [m], Sebastian Sardiña [n], Niek Tax [o], Barbara Weber [p]

[a] Free University of Bozen-Bolzano, Bozen-Bolzano, Italy
[b] University of Oxford, Oxford, United Kingdom
[c] University of Tartu, Tartu, Estonia
[d] IBM Research, Haifa, Israel
[e] Umeå University, Umeå, Sweden
[f] SAP, Berlin, Germany
[g] Sapienza Università di Roma, Rome, Italy
[h] paluno (Ruhr Institute for Software Technology), University of Duisburg Essen, Essen, Germany
[i] University of Ottawa, Ottawa, Canada
[j] German Research Center for Artificial Intelligence (DFKI), 66123 Saarbrücken, Germany
[k] Saarland University, 66123 Saarbrücken, Germany
[l] The University of Melbourne, Parkville, VIC 3010, Australia
[m] TUM School of Computation, Information, and Technology, TU Munich, Garching, Germany
[n] RMIT University, Melbourne, Australia
[o] Meta, London, United Kingdom
[p] University of St. Gallen, St. Gallen, Switzerland

ARTICLE INFO

Keywords:
Business process management
Autonomous agents
Agentic AI
Framed autonomy
Explainability
Conversational actionability
Self-modification

ABSTRACT

This paper presents a manifesto that articulates the conceptual foundations of *Agentic Business Process Management* (APM), an extension of Business Process Management (BPM) for governing autonomous agents executing processes in organizations. From a management perspective, APM represents a paradigm shift from the traditional view on business processes. This shift is driven by the realization of process awareness by agent-oriented abstractions: software and human agents act as primary functional entities that perceive, reason, and act within explicit process frames. Thus, APM moves away from automation-oriented BPM towards systems in which autonomy is constrained, aligned, and made operational through *process aware* agents.

We introduce the core abstractions and architectural elements required to realize APM systems and elaborate on four key capabilities that agents in APM systems must support: *framed autonomy*, *explainability*, *conversational actionability*, and *self-modification*. These capabilities jointly ensure that agents' goals are aligned with organizational goals and that agents behave in a framed yet proactive manner in pursuing those goals. We discuss the extent to which the capabilities can be realized and identify research challenges whose resolution requires further advances in BPM, AI, and multi-agent systems. The manifesto thus serves as a roadmap for bridging these communities and for guiding the development of APM systems in practice.

☆ This article is part of a Special issue entitled: 'Autonomous Process Execution Systems' published in Information Systems.

☆☆ Given the role as Editor-in-Chief of this journal, Stefanie Rinderle-Ma had no involvement in the peer-review of this article and has no access to information regarding its peer-review. Full responsibility for the editorial process for this article was delegated to another journal editor.

* Corresponding author.
E-mail addresses: diego.calvanese@unibz.it (D. Calvanese), casciani@diag.uniroma1.it (A. Casciani), giuseppe.degiacomo@cs.ox.ac.uk (G. De Giacomo), marlon.dumas@ut.ee (M. Dumas), fabiana@il.ibm.com (F. Fournier), tkampik@cs.umu.se (T. Kampik), emanuele.lamalfa@cs.ox.ac.uk (E.L. Malfa), liorli@il.ibm.com (L. Limonad), marrella@diag.uniroma1.it (A. Marrella), andreas.metzger@paluno.uni-due.de (A. Metzger), marco.montali@unibz.it (M. Montali), damyot@uottawa.ca (D. Amyot), peter.fettke@dfki.de (P. Fettke), artem.polyvyanyy@unimelb.edu.au (A. Polyvyanyy), stefanie.rinderle-ma@tum.de (S. Rinderle-Ma), sebastian.sardina@rmit.edu.au (S. Sardiña), niek@meta.com (N. Tax), barbara.weber@unisg.ch (B. Weber).

https://doi.org/10.1016/j.is.2026.102738
Received 11 December 2025; Received in revised form 8 April 2026; Accepted 8 April 2026
Available online 17 April 2026

explainability addresses important concerns in an APM system

- accountability
- bias
- compliance
- trust

Agentic Business Process Management: A research manifesto, in: Information Systems (2026)

Information Systems 140 (2026) 102738

Contents lists available at ScienceDirect

Information Systems

journal homepage: www.elsevier.com/locate/is

Agentic Business Process Management: A research manifesto☆,☆☆

Diego Calvanese [a], Angelo Casciani [g], Giuseppe De Giacomo [b], Marlon Dumas [c], Fabiana Fournier [d], Timotheus Kampik [e,f,*], Emanuele La Malfa [b], Lior Limonad [d], Andrea Marrella [g], Andreas Metzger [h], Marco Montali [a], Daniel Amyot [i], Peter Fettke [j,k], Artem Polyvyanyy [l], Stefanie Rinderle-Ma [m], Sebastian Sardiña [n], Niek Tax [o], Barbara Weber [p]

[a] Free University of Bozen-Bolzano, Bozen-Bolzano, Italy
[b] University of Oxford, Oxford, United Kingdom
[c] University of Tartu, Tartu, Estonia
[d] IBM Research, Haifa, Israel
[e] Umeå University, Umeå, Sweden
[f] SAP, Berlin, Germany
[g] Sapienza Università di Roma, Rome, Italy
[h] paluno (Ruhr Institute for Software Technology), University of Duisburg Essen, Essen, Germany
[i] University of Ottawa, Ottawa, Canada
[j] German Research Center for Artificial Intelligence (DFKI), 66123 Saarbrücken, Germany
[k] Saarland University, 66123 Saarbrücken, Germany
[l] The University of Melbourne, Parkville, VIC 3010, Australia
[m] TUM School of Computation, Information, and Technology, TU Munich, Garching, Germany
[n] RMIT University, Melbourne, Australia
[o] Meta, London, United Kingdom
[p] University of St. Gallen, St. Gallen, Switzerland

ARTICLE INFO

Keywords:
Business process management
Autonomous agents
Agentic AI
Framed autonomy
Explainability
Conversational actionability
Self-modification

ABSTRACT

This paper presents a manifesto that articulates the conceptual foundations of *Agentic Business Process Management* (APM), an extension of Business Process Management (BPM) for governing autonomous agents executing processes in organizations. From a management perspective, APM represents a paradigm shift from the traditional view on business processes. This shift is driven by the realization of process awareness by agent-oriented abstractions: software and human agents act as primary functional entities that perceive, reason, and act within explicit process frames. Thus, APM moves away from automation-oriented BPM towards systems in which autonomy is constrained, aligned, and made operational through *process aware* agents.

We introduce the core abstractions and architectural elements required to realize APM systems and elaborate on four key capabilities that agents in APM systems must support: *framed autonomy*, *explainability*, *conversational actionability*, and *self-modification*. These capabilities jointly ensure that agents' goals are aligned with organizational goals and that agents behave in a framed yet proactive manner in pursuing those goals. We discuss the extent to which the capabilities can be realized and identify research challenges whose resolution requires further advances in BPM, AI, and multi-agent systems. The manifesto thus serves as a roadmap for bridging these communities and for guiding the development of APM systems in practice.

☆ This article is part of a Special issue entitled: 'Autonomous Process Execution Systems' published in Information Systems.

☆☆ Given the role as Editor-in-Chief of this journal, Stefanie Rinderle-Ma had no involvement in the peer-review of this article and has no access to information regarding its peer-review. Full responsibility for the editorial process for this article was delegated to another journal editor.

* Corresponding author.
E-mail addresses: diego.calvanese@unibz.it (D. Calvanese), casciani@diag.uniroma1.it (A. Casciani), giuseppe.degiacomo@cs.ox.ac.uk (G. De Giacomo), marlon.dumas@ut.ee (M. Dumas), fabiana@il.ibm.com (F. Fournier), tkampik@cs.umu.se (T. Kampik), emanuele.lamalfa@cs.ox.ac.uk (E.L. Malfa), liorli@il.ibm.com (L. Limonad), marrella@diag.uniroma1.it (A. Marrella), andreas.metzger@paluno.uni-due.de (A. Metzger), marco.montali@unibz.it (M. Montali), damyot@uottawa.ca (D. Amyot), peter.fettke@dfki.de (P. Fettke), artem.polyvyanyy@unimelb.edu.au (A. Polyvyanyy), stefanie.rinderle-ma@tum.de (S. Rinderle-Ma), sebastian.sardina@rmit.edu.au (S. Sardiña), niek@meta.com (N. Tax), barbara.weber@unisg.ch (B. Weber).

https://doi.org/10.1016/j.is.2026.102738
Received 11 December 2025; Received in revised form 8 April 2026; Accepted 8 April 2026
Available online 17 April 2026

explainability addresses important concerns in an APM system

- accountability
- bias
- compliance
- **trust**

"issues may arise among stakeholders – including process owners, business analysists, end users, and customers – who may hesitate to rely on agentic process recommendations or automated agent decisions when the underlying rationale is unclear." (p. 6)

Trust in science

- Scientific theory contains a set of sentences (propositions, statements etc.)
- Problem: How can we trust a sentence? How can we prove the truth of a sentence?
- Example: “All swans are white.”

“All swans are white.”

“All swans are white.”
A black swan?

“All swans are white.”
A black swan: improbable but possible!

black swan (cygnus atratus) in Tasmania, Australia

source: brockhaus.de

“All swans are white, some are black.”

“All swans are white, some are black.”
A green-and-pink-striped swan?

Highly improbable but possible!

Trust in mathematics

- All sentences are deduced from premises.
- Example: $a^2 + b^2 = c^2$ (Pythagoras's theorem)

"The remainder of the book shows **370** different proofs [of Pythagoras's theorem], whose origins range from 900 B.C. to 1940 A.D." (cp. abstract)

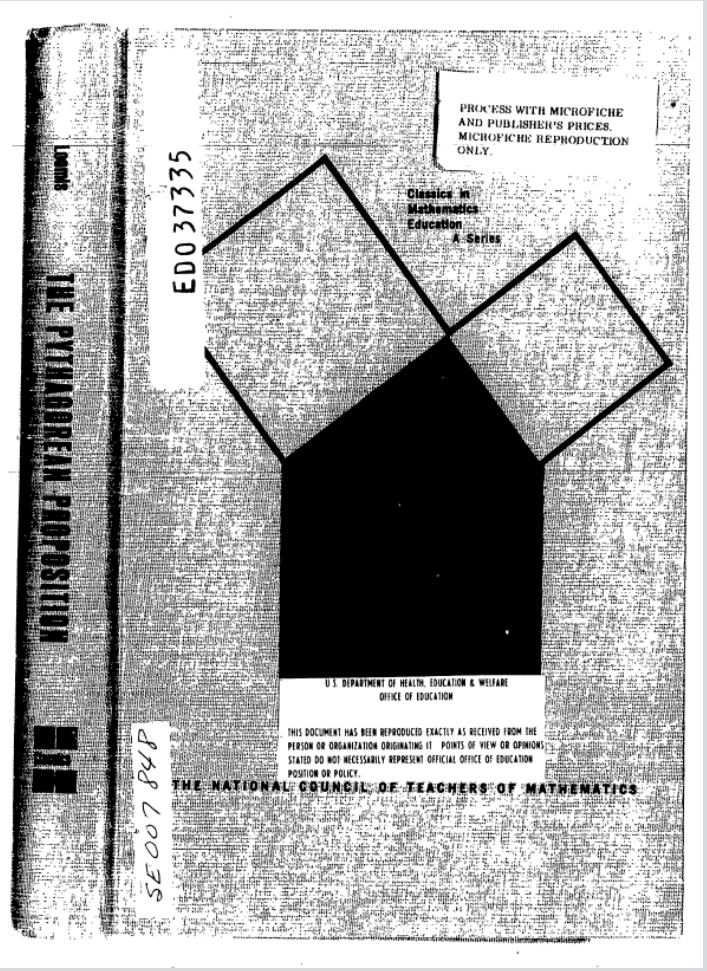

E. S. Loomis: The Pythagorean Proposition, National Council of Teachers of Mathematics, 1940

Trust in mathematics

Fermat's Last Theorem

$a^n + b^n = c^n$ has no non-zero integer solutions for any integer exponent $n > 2$

It took the mathematical community more than three centuries to officially (!) prove the theorem!

Proof by Andrew Wiles published in 1995 is 109 (!) pages long

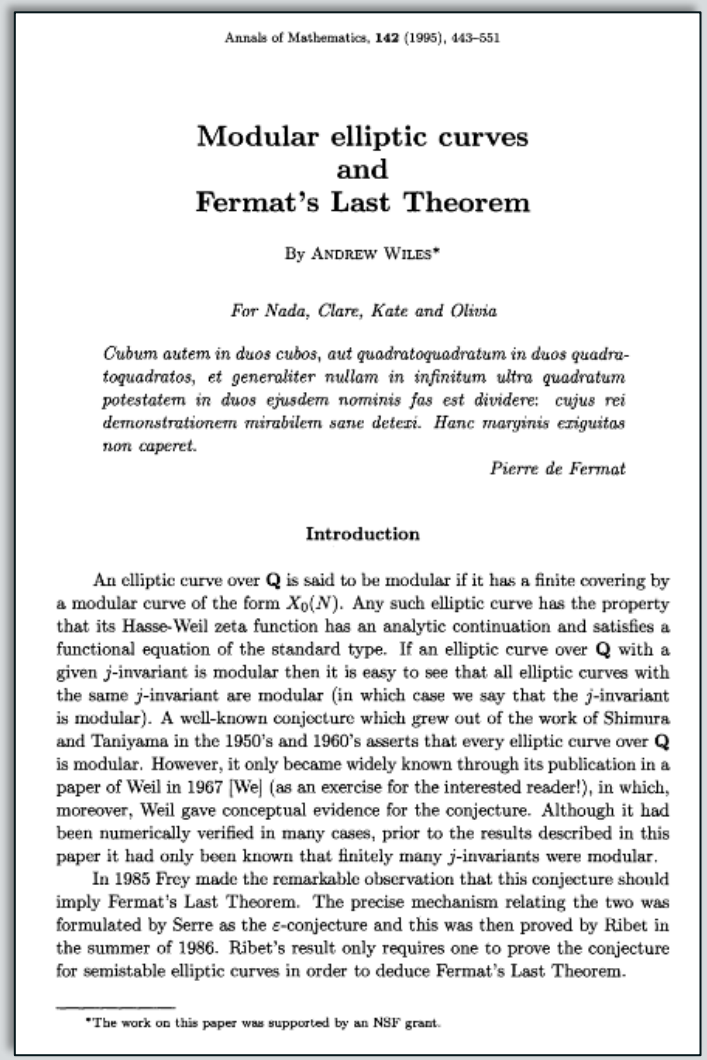

Annals of Mathematics, **142** (1995), 443–551

Modular elliptic curves and Fermat's Last Theorem

By ANDREW WILES*

For Nada, Clare, Kate and Olivia

Cubum autem in duos cubos, aut quadratoquadratum in duos quadratoquadratos, et generaliter nullam in infinitum ultra quadratum potestatem in duos ejusdem nominis fas est dividere: cujus rei demonstrationem mirabilem sane detexi. Hanc marginis exiguitas non caperet.

Pierre de Fermat

Introduction

An elliptic curve over **Q** is said to be modular if it has a finite covering by a modular curve of the form $X_0(N)$. Any such elliptic curve has the property that its Hasse-Weil zeta function has an analytic continuation and satisfies a functional equation of the standard type. If an elliptic curve over **Q** with a given j-invariant is modular then it is easy to see that all elliptic curves with the same j-invariant are modular (in which case we say that the j-invariant is modular). A well-known conjecture which grew out of the work of Shimura and Taniyama in the 1950's and 1960's asserts that every elliptic curve over **Q** is modular. However, it only became widely known through its publication in a paper of Weil in 1967 [We] (as an exercise for the interested reader!), in which, moreover, Weil gave conceptual evidence for the conjecture. Although it had been numerically verified in many cases, prior to the results described in this paper it had only been known that finitely many j-invariants were modular.

In 1985 Frey made the remarkable observation that this conjecture should imply Fermat's Last Theorem. The precise mechanism relating the two was formulated by Serre as the ε-conjecture and this was then proved by Ribet in the summer of 1986. Ribet's result only requires one to prove the conjecture for semistable elliptic curves in order to deduce Fermat's Last Theorem.

*The work on this paper was supported by an NSF grant.

We have to trust in the proof!

Is mathematics just a social process?

Trust by design

Disclaimer

We do not present a final answer!

We like to point to important questions and present some preliminary results!

Agenda

1. Motivation and background
2. Trust certificates
3. An example
4. Conclusions

dfki ai Deutsches Forschungszentrum für Künstliche Intelligenz
German Research Center for Artificial Intelligence

UNIVERSITÄT DES SAARLANDES

HUMBOLDT-UNIVERSITÄT ZU BERLIN

Example: bipartite graph

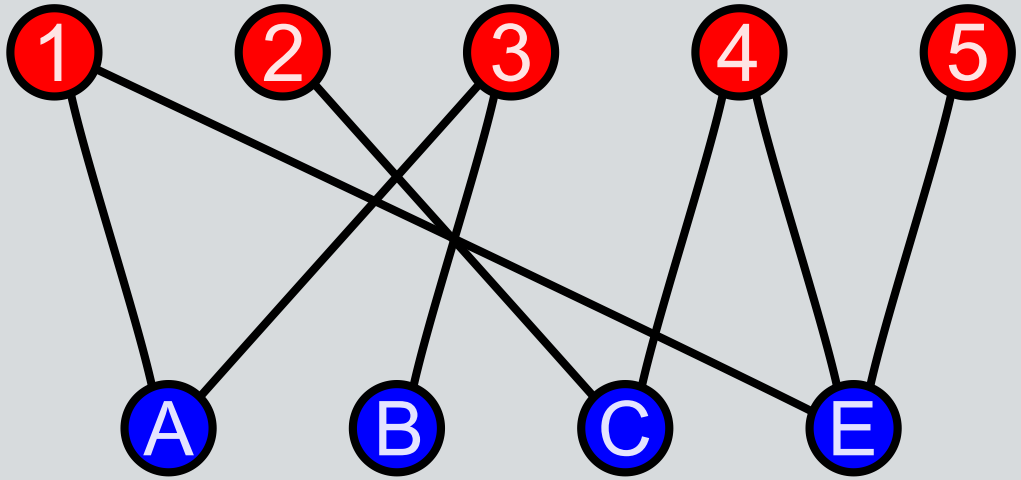

A graph is bipartite if its vertices can be partitioned into two sets in such a way that no edge joins two vertices in the same set.

Example: Is this graph bipartite?

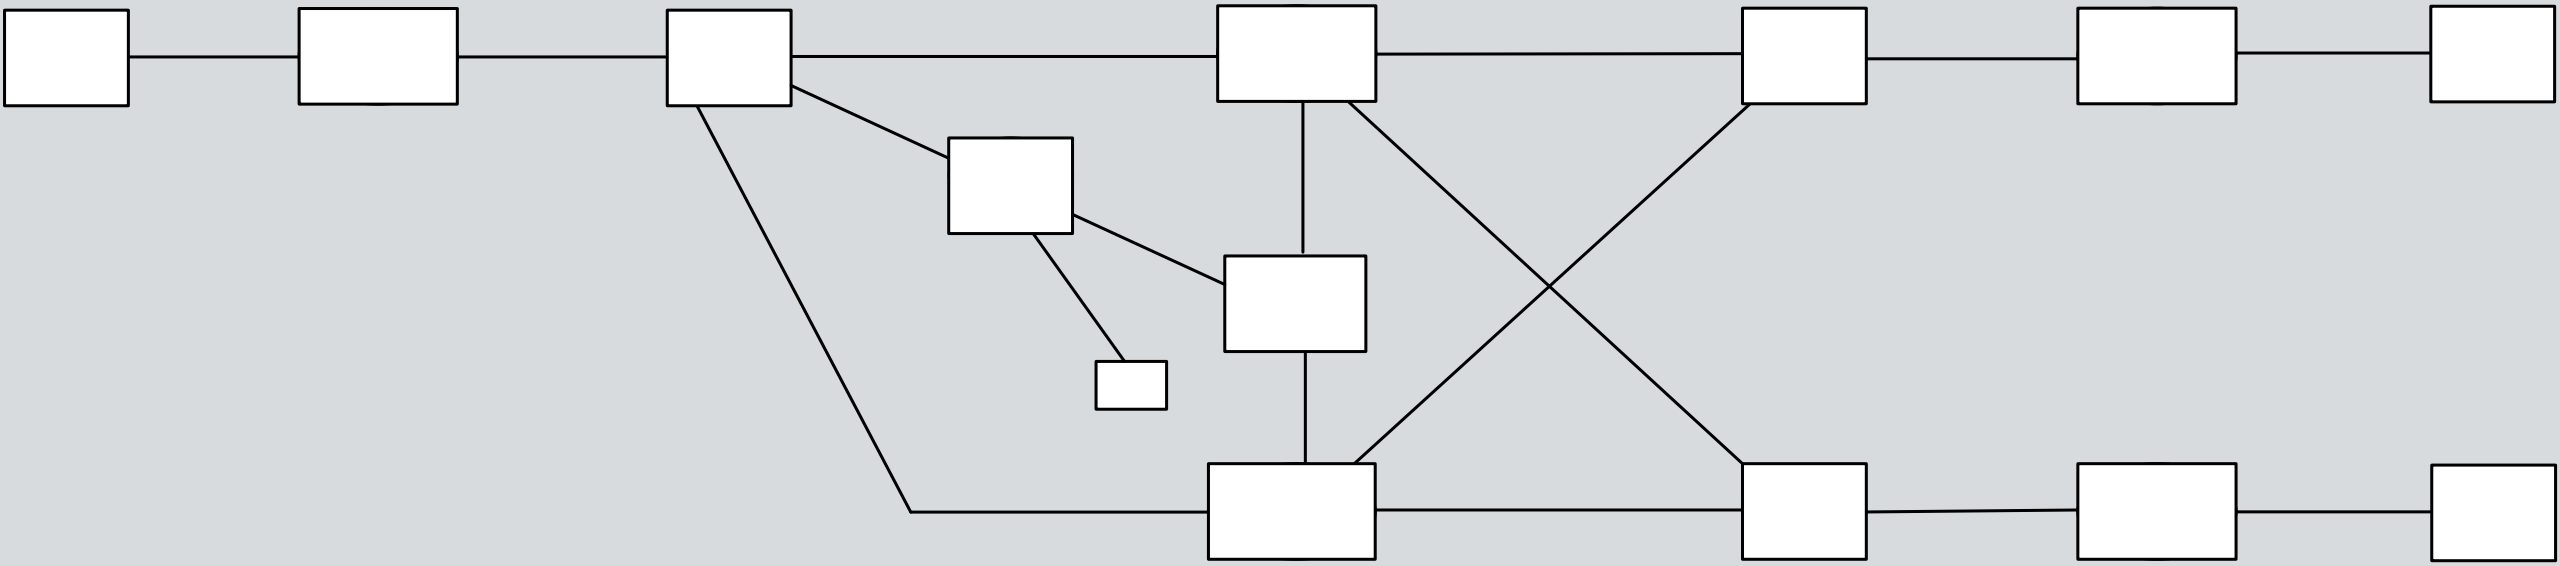

Example: Is this graph bipartite?
Yes! But how to prove it?

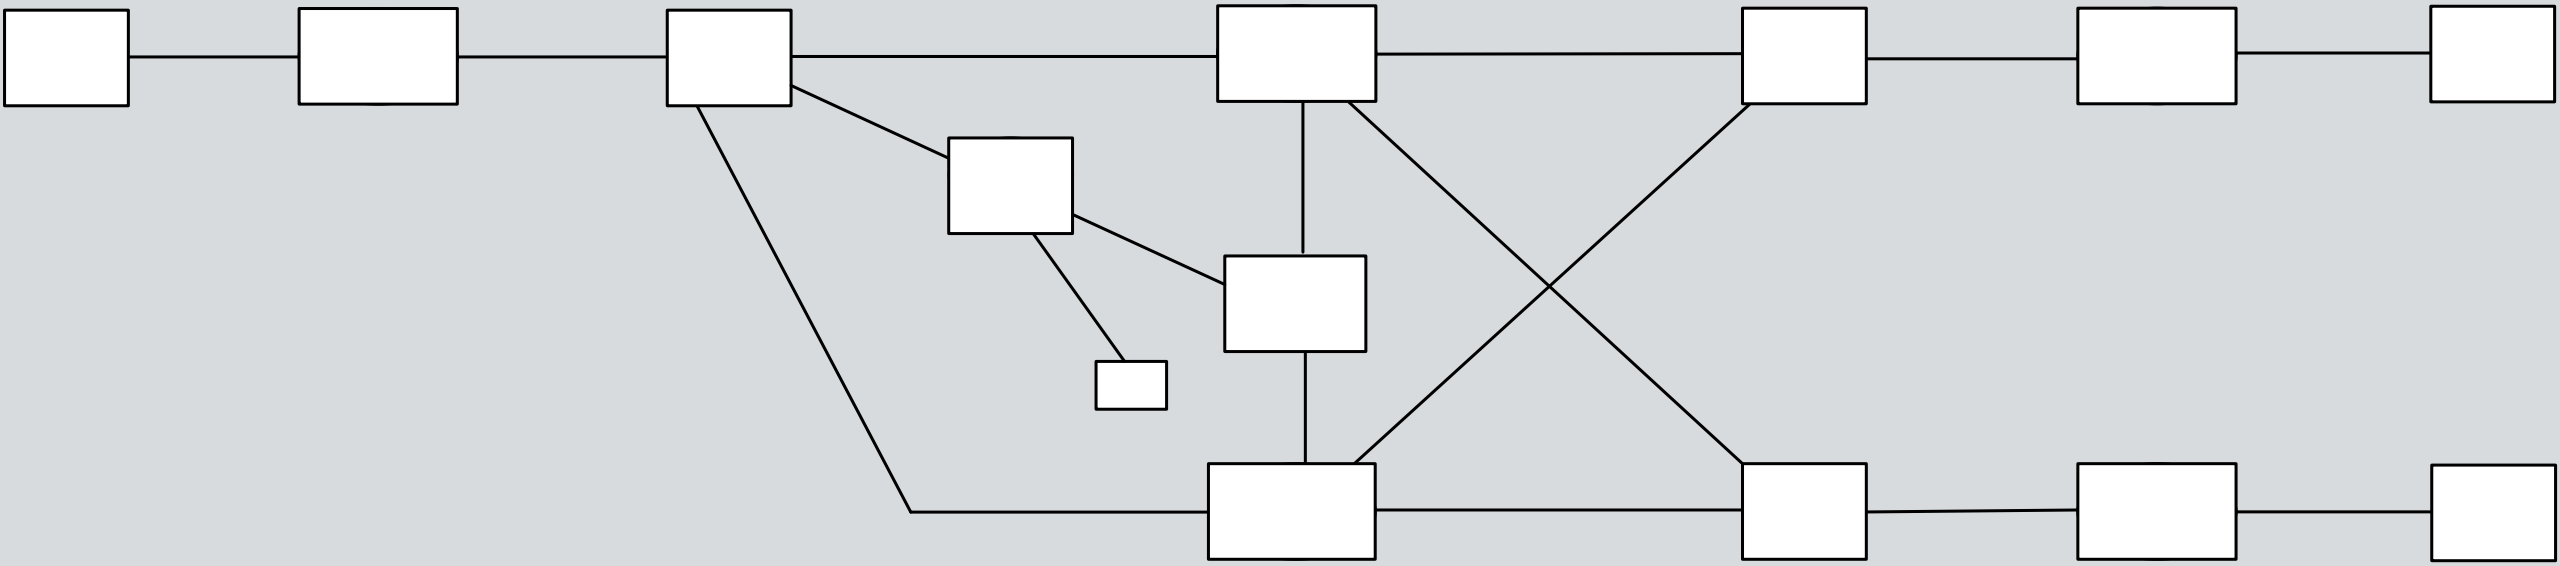

Example: Is this graph bipartite?
Yes! But how to prove it?

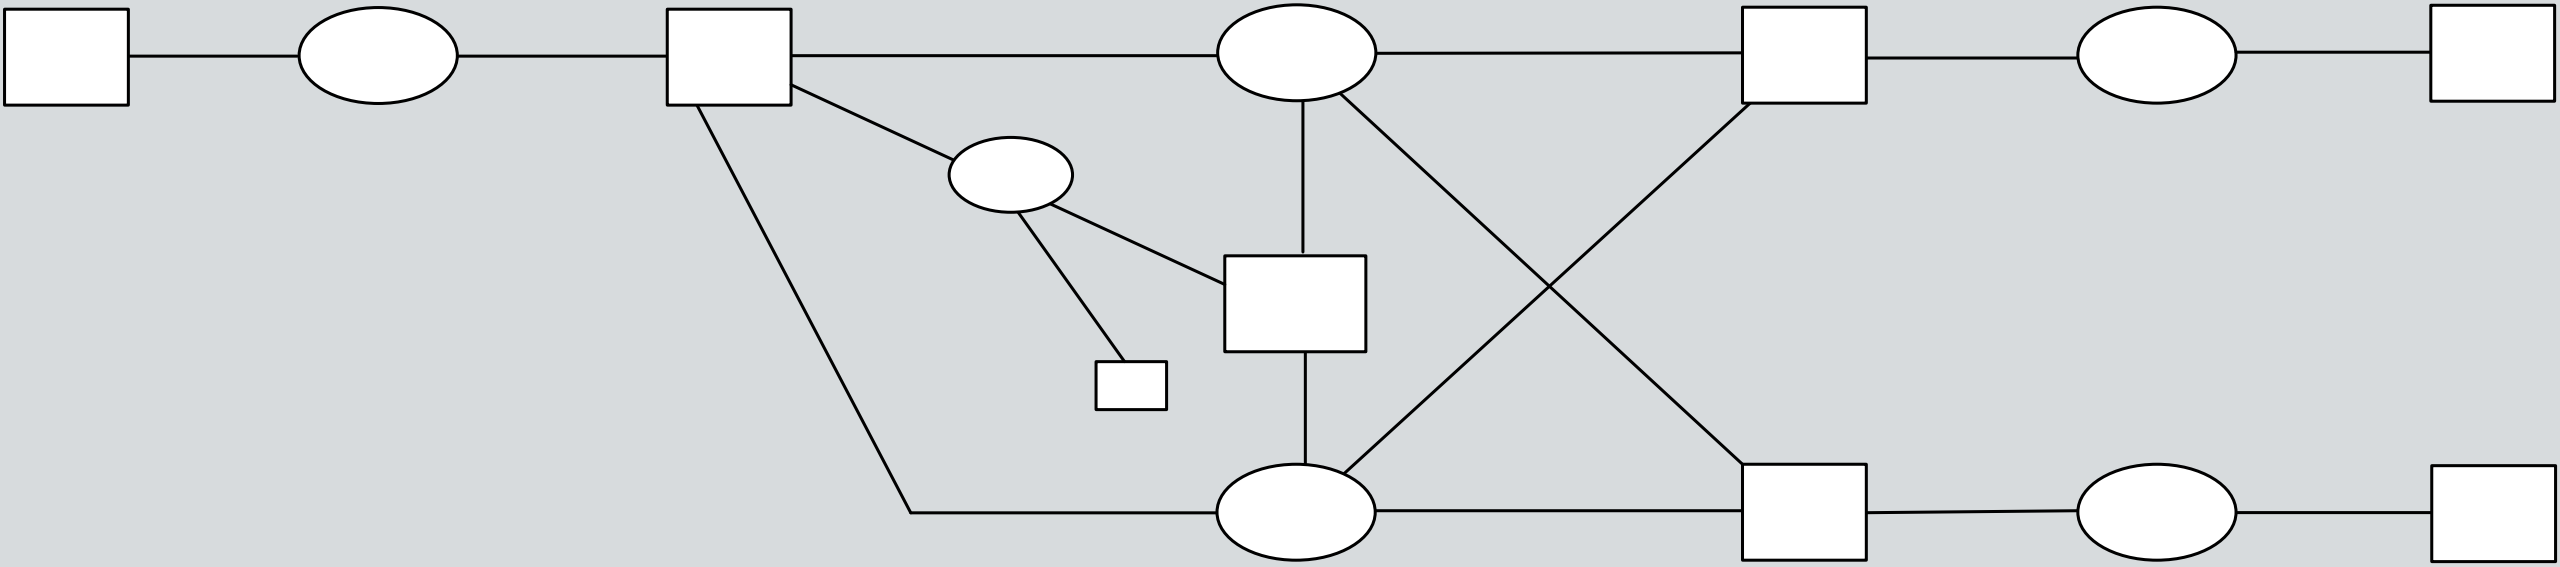

Example: Is this graph bipartite?

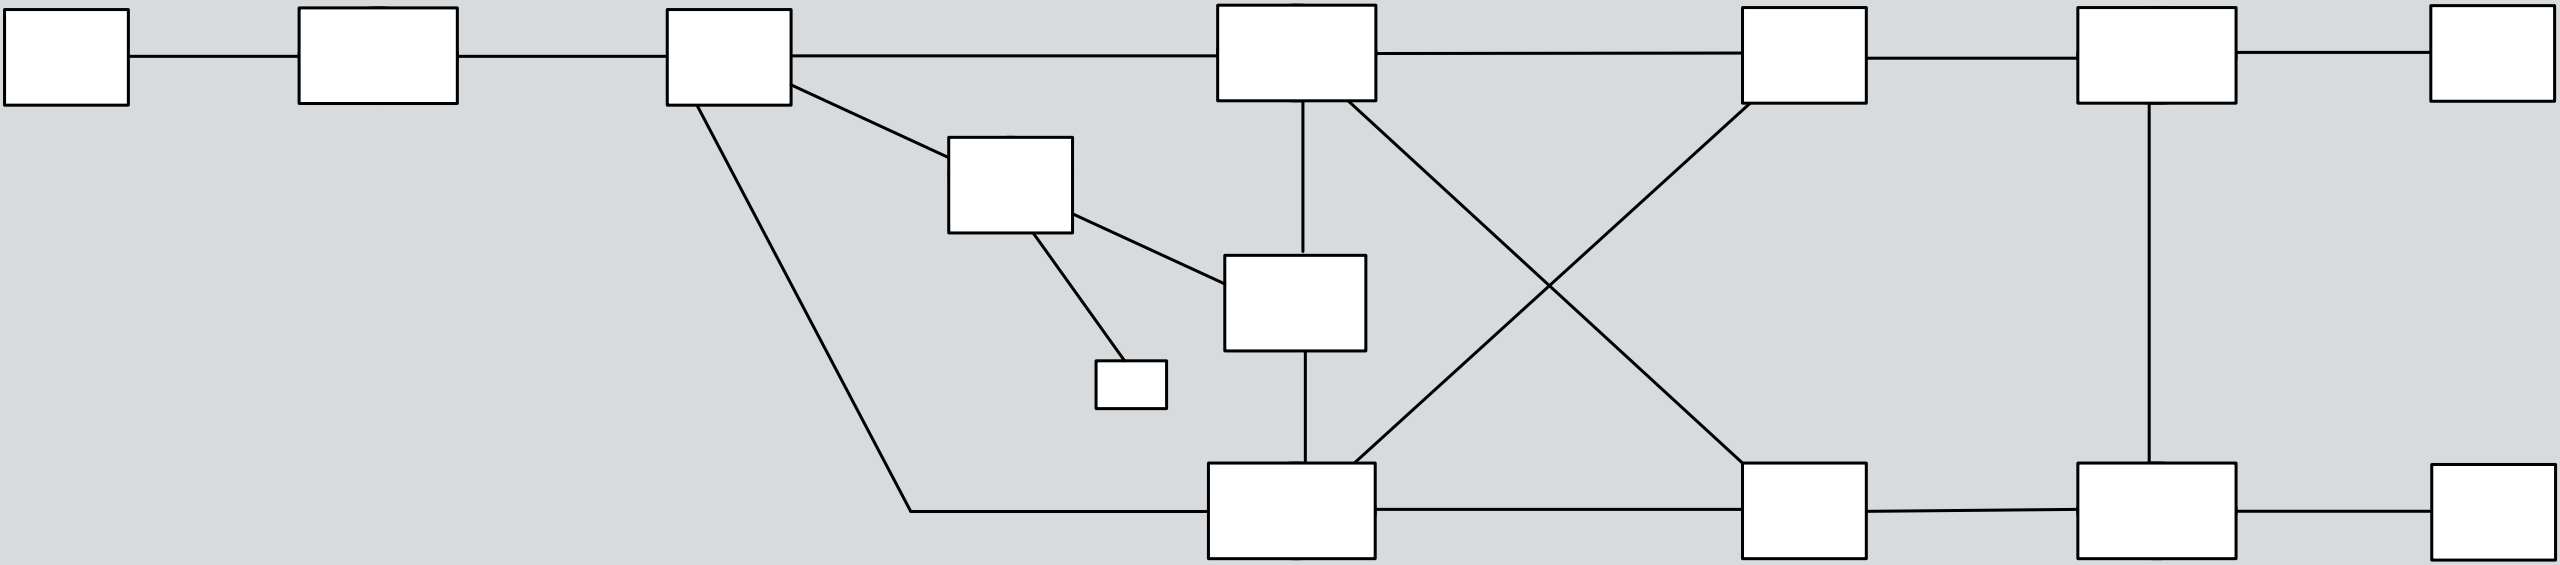

Example: Is this graph bipartite?
No! But how to prove it?

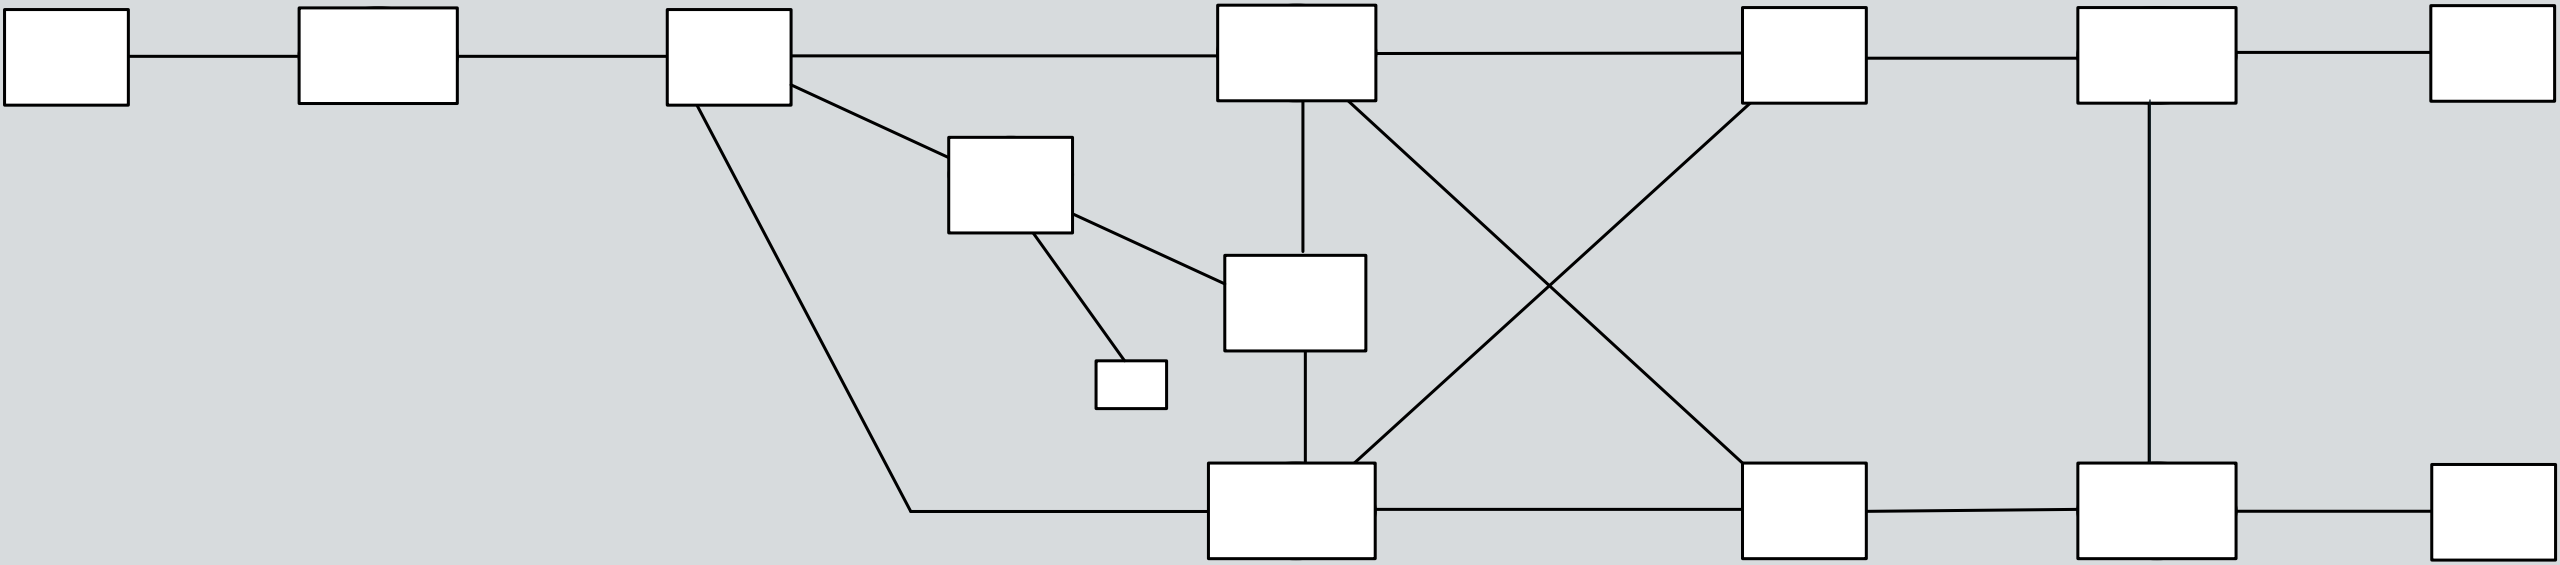

Example: Is this graph bipartite?
No! But how to prove it?

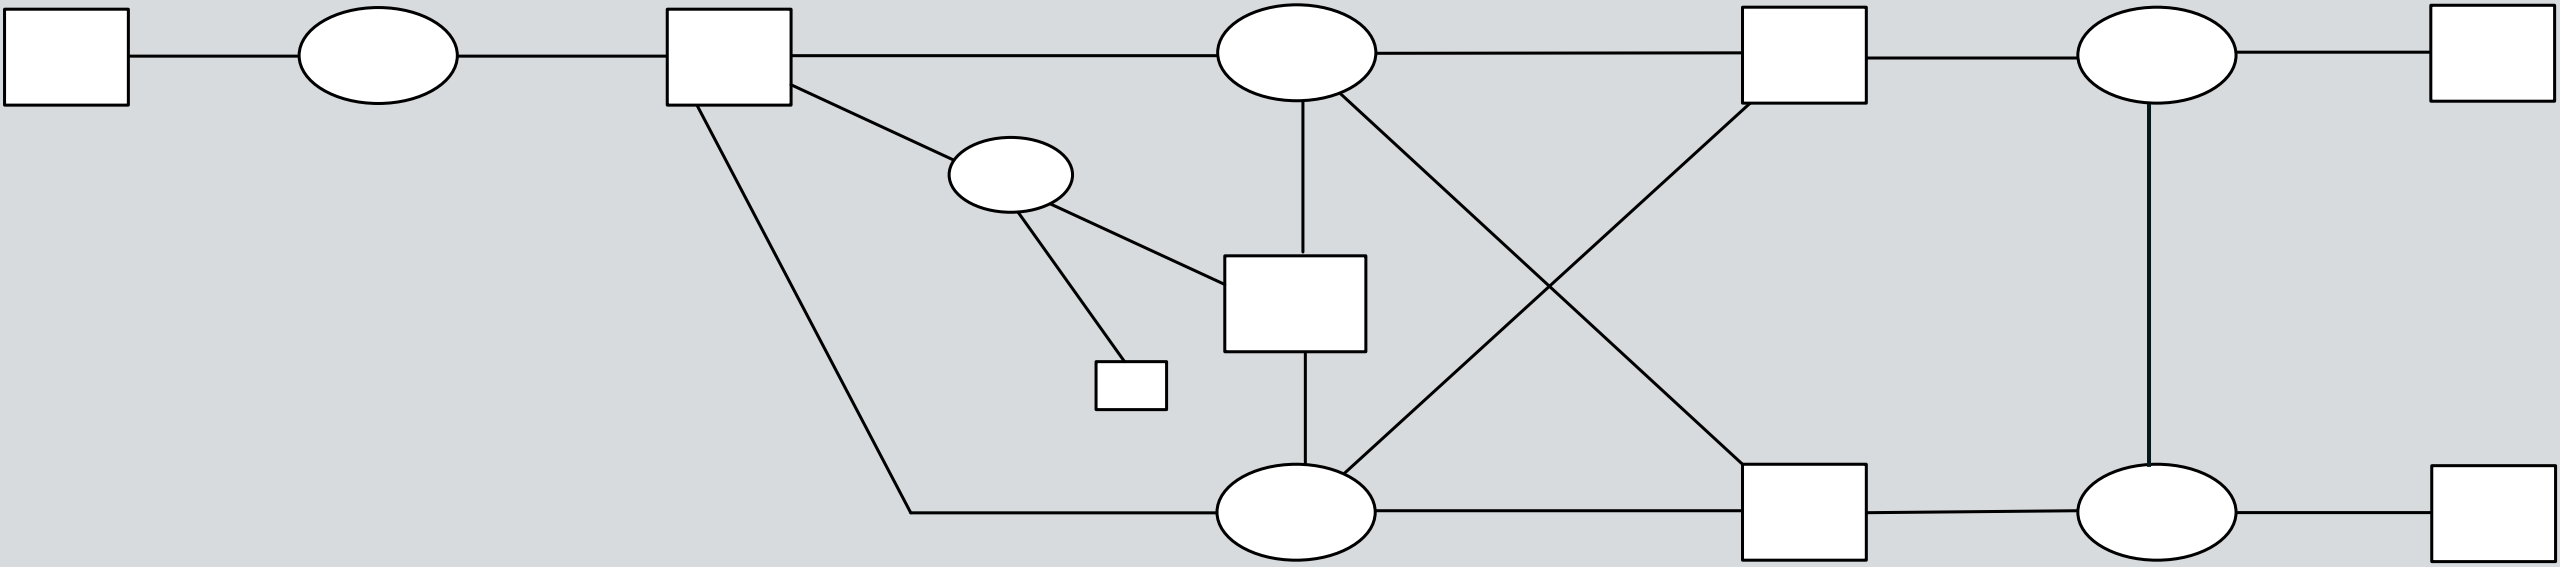

Example: Is this graph bipartite?
No! But how to prove it?

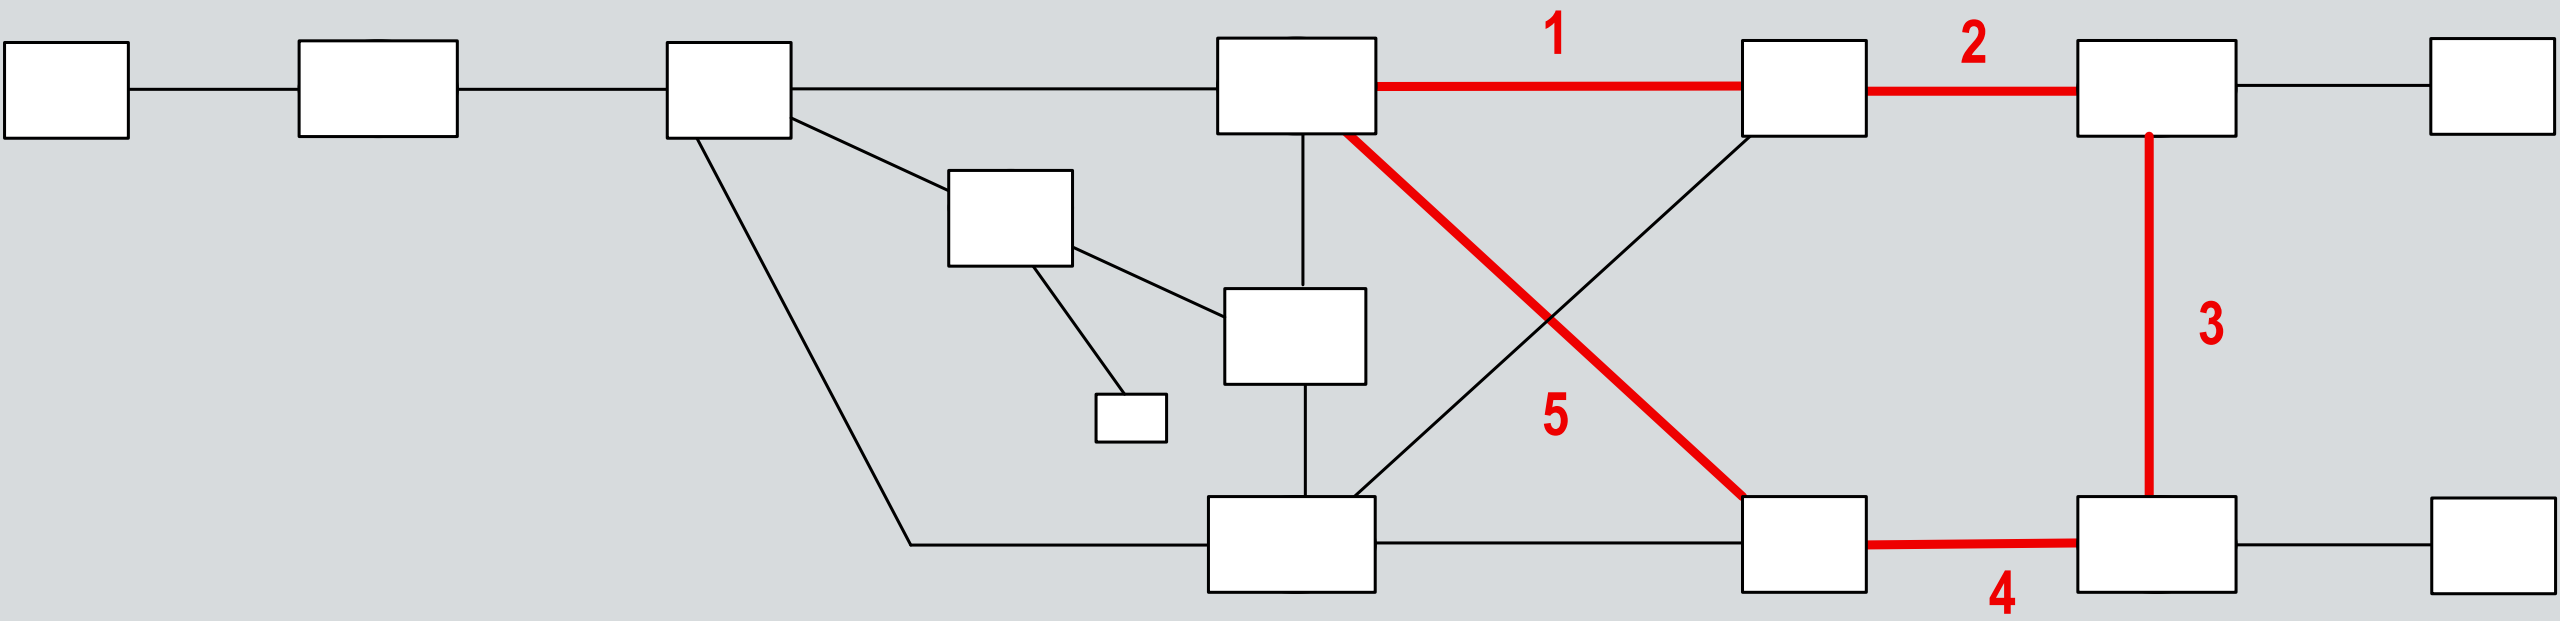

Example: Is this graph bipartite? No! But how to prove it?

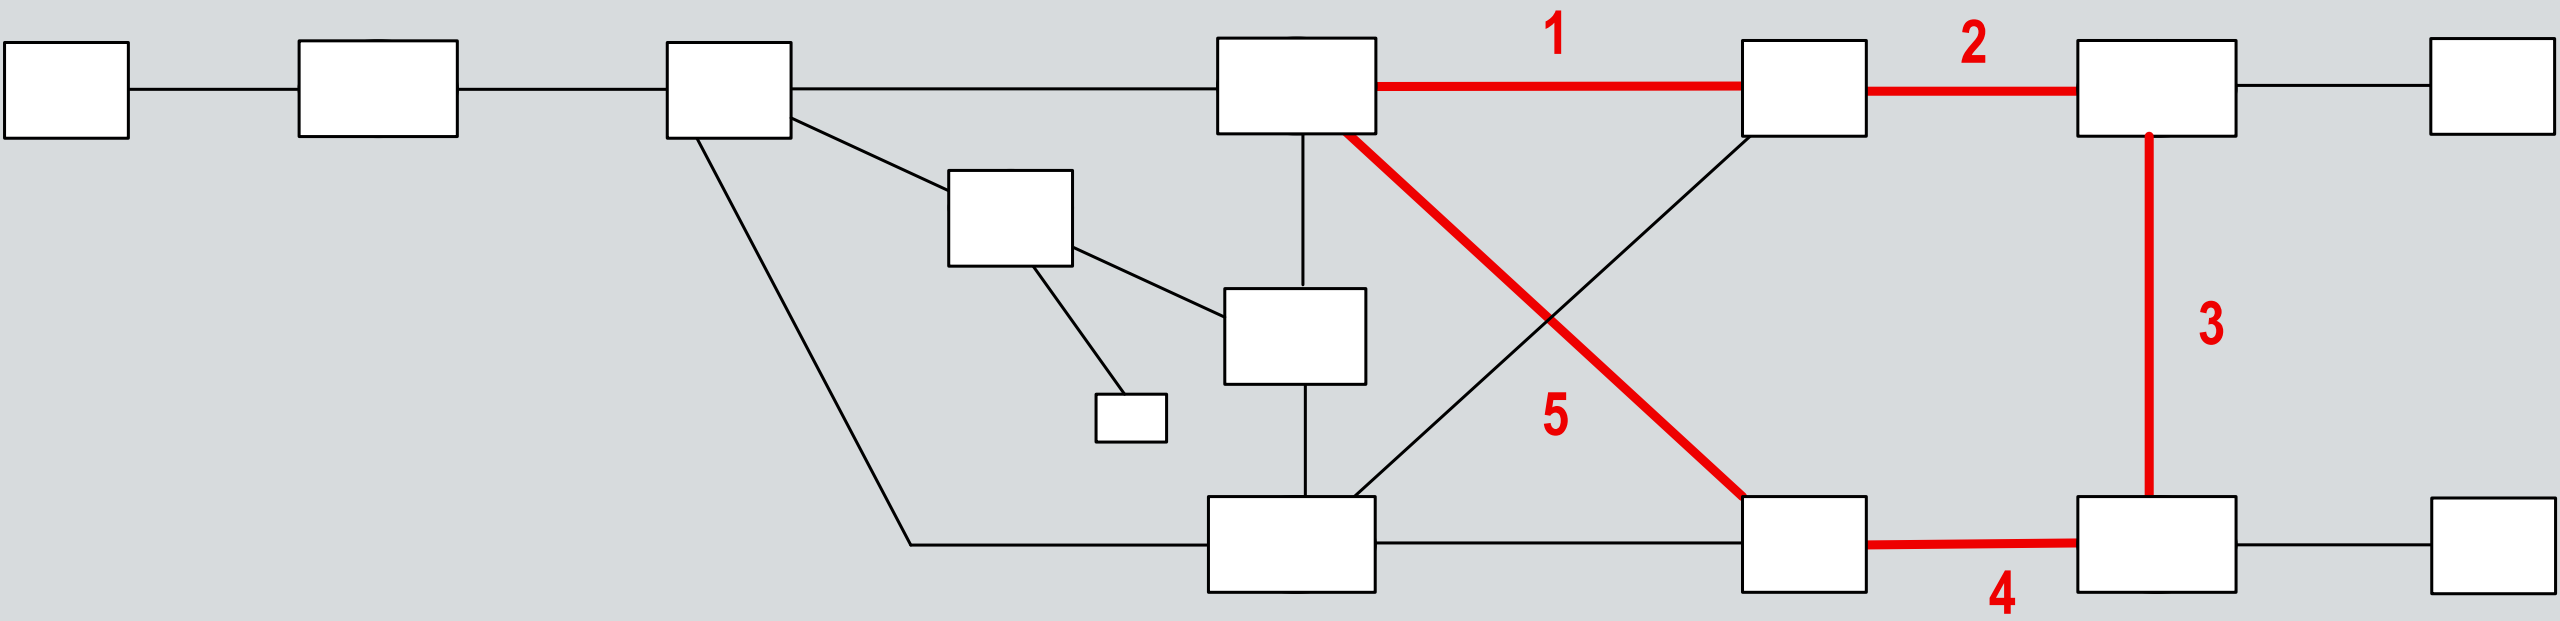

Bipartite graph characterization theorem:
„A graph is bipartite if and only if the length of each of its cycles is even.“
(Handbook of Graph Theory by Gross et al., 2014, p. 27)

Idea of a certifying program by Mehlhorn et al.

Idea of a certifying program by Mehlhorn et al.

Idea of a certifying program by Mehlhorn et al.

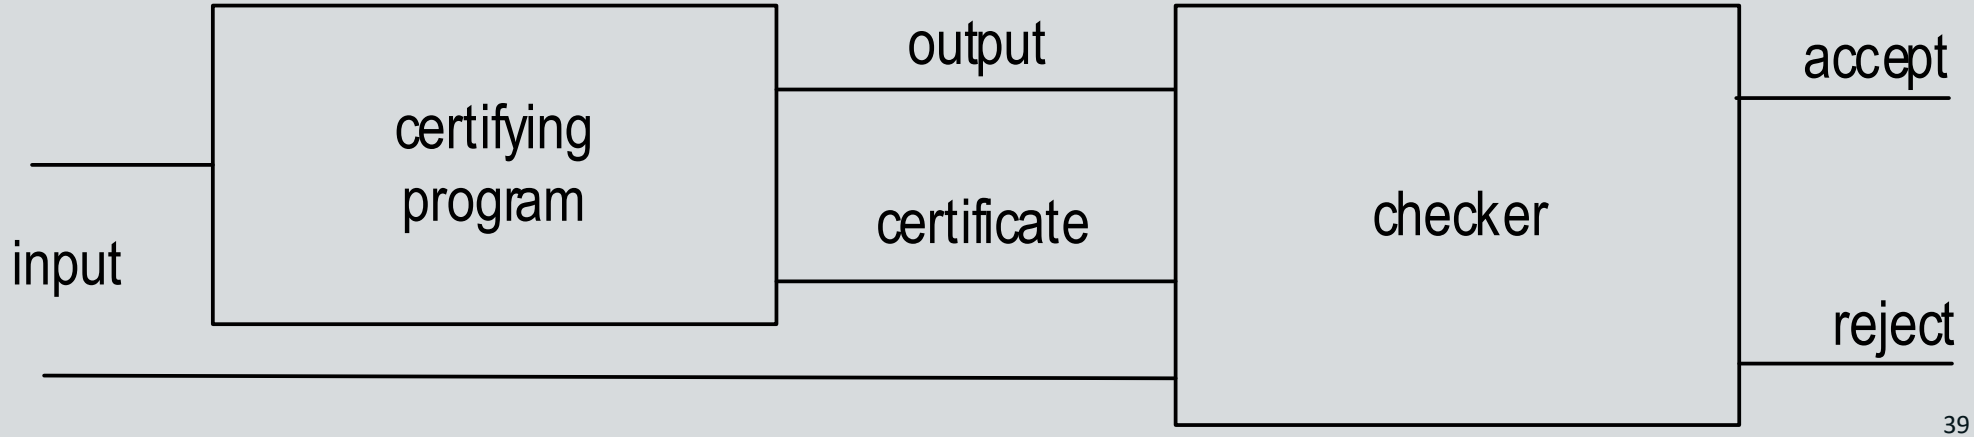

Idea of a certifying program by Mehlhorn et al.

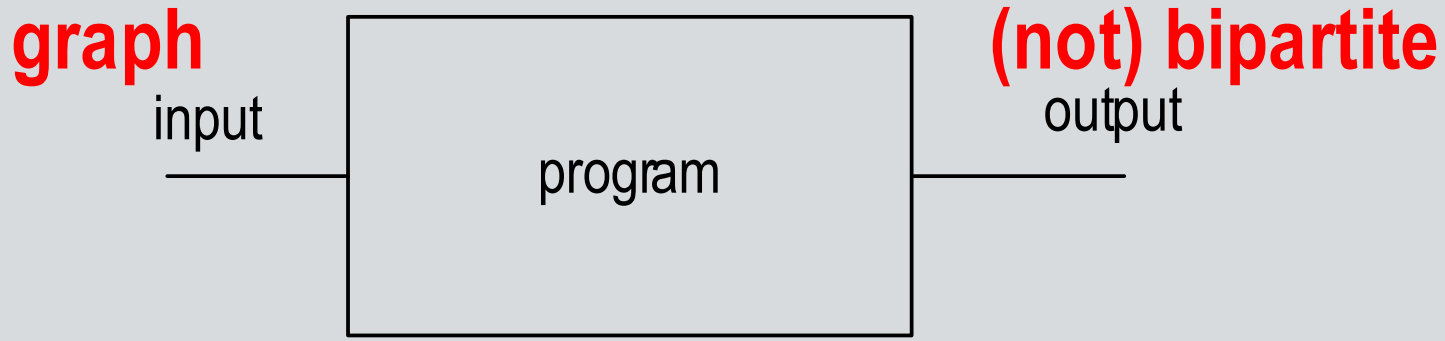

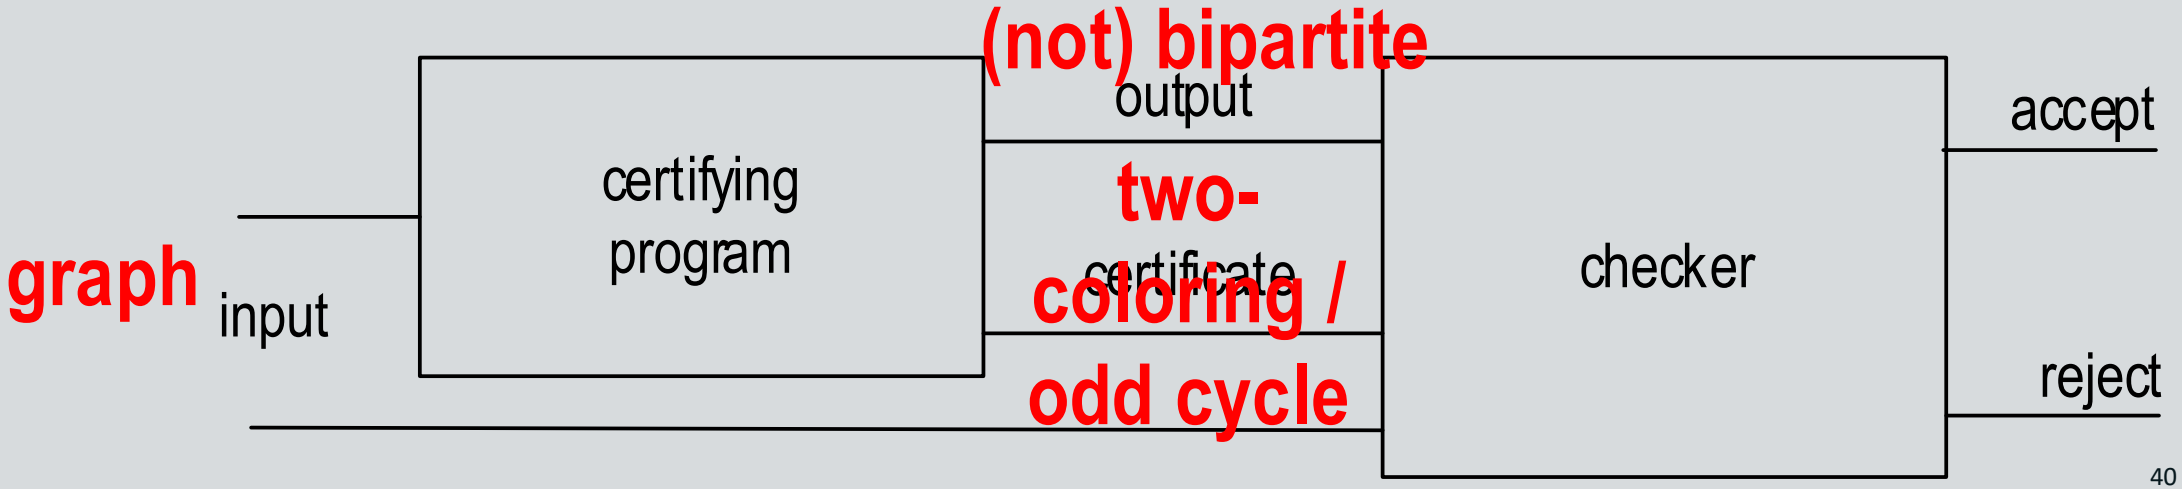

Plenty of prior work on certification in the “domain of programming”!

COMPUTER SCIENCE REVIEW 5 (2011) 119–161

available at www.sciencedirect.com

ScienceDirect

journal homepage: www.elsevier.com/locate/cosrev

Survey

Certifying algorithms

R.M. McConnell[a], K. Mehlhorn[b,*], S. Näher[c], P. Schweitzer[d]

[a] Computer Science Department, Colorado State University Fort Collins, USA
[b] Max Planck Institute for Informatics and Saarland University, Saarbrücken, Germany
[c] Fachbereich Informatik, Universität Trier, Trier, Germany
[d] College of Engineering and Computer Science, Australian National University, Canberra, Australia

ARTICLE INFO

Article history:
Received 2 June 2010
Received in revised form
29 September 2010
Accepted 29 September 2010

Keywords:
Algorithms
Software reliability
Certification

ABSTRACT

A certifying algorithm is an algorithm that produces, with each output, a certificate or witness (easy-to-verify proof) that the particular output has not been compromised by a bug. A user of a certifying algorithm inputs x, receives the output y and the certificate w, and then checks, either manually or by use of a program, that w proves that y is a correct output for input x. In this way, he/she can be sure of the correctness of the output without having to trust the algorithm.

We put forward the thesis that certifying algorithms are much superior to non-certifying algorithms, and that for complex algorithmic tasks, only certifying algorithms are satisfactory. Acceptance of this thesis would lead to a change of how algorithms are taught and how algorithms are researched. The widespread use of certifying algorithms would greatly enhance the reliability of algorithmic software.

We survey the state of the art in certifying algorithms and add to it. In particular, we start a theory of certifying algorithms and prove that the concept is universal.

1. Introduction

One of the most prominent and costly problems in software engineering is correctness of software. When the user gives x as an input and the program outputs y, the user usually has no way of knowing whether y is a correct output on input x or whether it has been compromised by a bug. The user is at the mercy of the program. A *certifying algorithm* is an algorithm that produces with each output a *certificate* or *witness* that the *particular output* has not been compromised by a bug. By inspecting the witness, the user can convince himself that the output is correct, or reject the output as buggy. He is no longer on the mercy of the program. Fig. 1 contrasts a standard algorithm with a certifying algorithm for computing a function f.

A user of a certifying algorithm inputs x and receives the output y and the witness w. He then checks that w proves that y is a correct output for input x. The process of checking w can be automated with a *checker*, which is an algorithm for verifying that w proves that y is a correct output for x. In many cases, the checker is so simple that a trusted implementation of it can be produced, perhaps even in a different language where the semantics are fully specified. A formal proof of correctness of the implementation of the certifying algorithm may be out of reach; however, a formal proof of the correctness of the checker may be feasible. Having checked

* Corresponding author.
E-mail address: mehlhorn@mpi-inf.mpg.de (K. Mehlhorn).

1574-0137/$ - see front matter © 2010 Elsevier Inc. All rights reserved.
doi:10.1016/j.cosrev.2010.09.009

Report from Dagstuhl Seminar 25231

Certifying Algorithms for Automated Reasoning

Nikolaj S. Bjørner*[1], Marijn J. H. Heule*[2], Daniela Kaufmann*[3], Jakob Nordström*[4], and Wietze Koops†[5]

1 Microsoft - Redmond, US. nbjorner@microsoft.com
2 Carnegie Mellon University - Pittsburgh, US. marijn@cmu.edu
3 TU Wien, AT. dk@danielakaufmann.at
4 University of Copenhagen, DK & Lund University, SE. jn@di.ku.dk
5 Lund University, SE & University of Copenhagen, DK. wietze.koops@cs.lth.se

Abstract

Modern automated reasoning has transformed large parts of industry and has also found numerous scientific applications. But many reasoning problems are computationally very challenging, or sometimes even undecidable. Because of this, the reasoning algorithms used are often very complex, and even the best current algorithms at times produce wrong results. As these tools are increasingly being used autonomously, sometimes even in life-critical applications, it is urgent to ensure that what they compute is valid. Software testing, while immensely useful, cannot guarantee correctness, and state-of-the-art algorithms are far beyond what techniques for producing formally verified software can handle.

The focus of this Dagstuhl Seminar was the approach of addressing such issues by designing certifying algorithms using so-called proof logging, meaning that algorithms output not only a result but also a machine-verifiable proof of correctness. This proof can then be fed to a dedicated proof checker for verification. Crucially, such proofs should require low overhead to generate and be easy to check, but still supply 100% correctness guarantees. Besides ensuring correctness of outputs for complex algorithms, proof logging can also provide new tools for algorithm development and analysis, software debugging, and even research into explainability in the context of AI.

Seminar June 1–6, 2025 – http://www.dagstuhl.de/25231
2012 ACM Subject Classification Mathematics of computing → Solvers; Theory of computation → Proof theory; Theory of computation → Automated reasoning; Software and its engineering → Software verification
Keywords and phrases ATP, Computer Algebra, DRAT, DRUP, MIP, Propagation Redundancy, QBF, SAT, SMT
Digital Object Identifier 10.4230/DagRep.15.6.1

* Editor / Organizer
† Editorial Assistant / Collector

Certifying Algorithms for Automated Reasoning, *Dagstuhl Reports*, Vol. 15, Issue 6, pp. 1–31
Editors: Nikolaj S. Bjørner, Marijn J. H. Heule, Daniela Kaufmann, Jakob Nordström, and Wietze Koops
Dagstuhl Reports
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Plenty of prior work on certification in the “domain of programming”!

But how can this idea be applied in the “domain of modeling”?

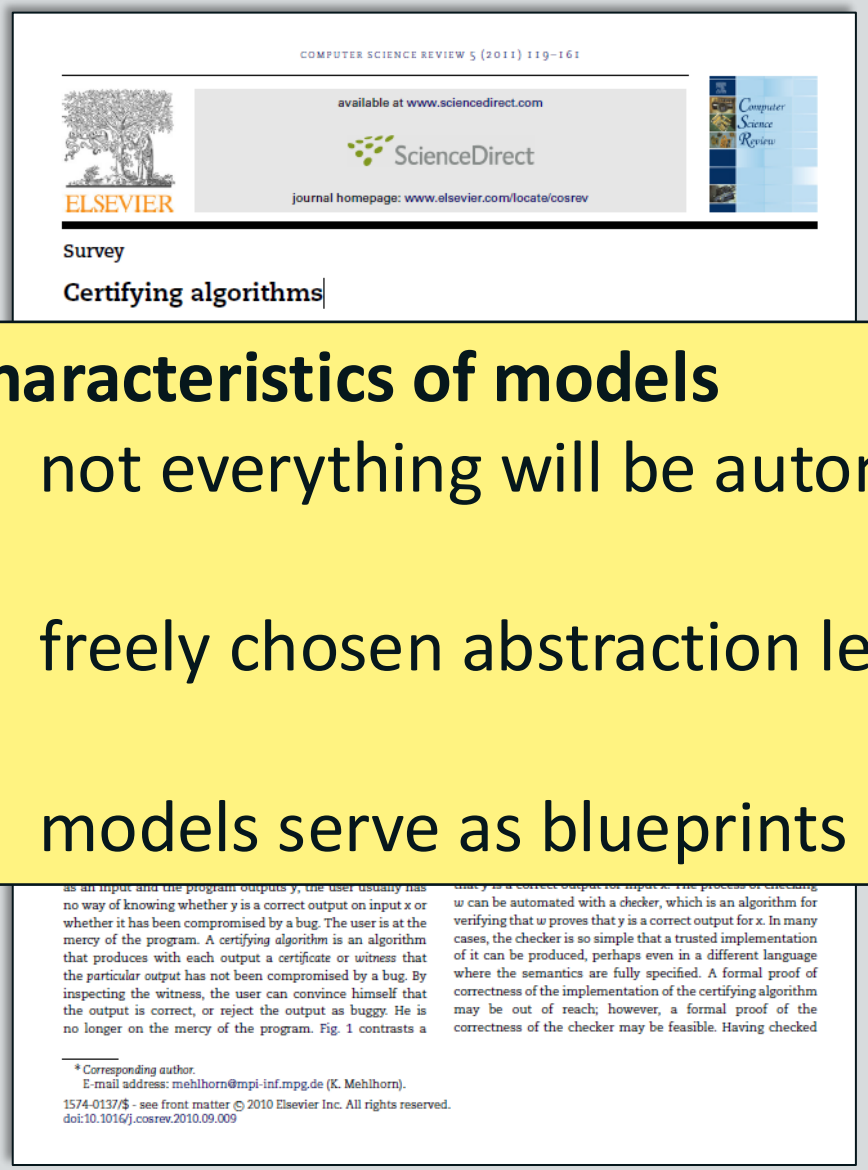

COMPUTER SCIENCE REVIEW 5 (2011) 119–161

available at www.sciencedirect.com

ScienceDirect

journal homepage: www.elsevier.com/locate/cosrev

ELSEVIER

Survey

Certifying algorithms

no way of knowing whether y is a correct output on input x or whether it has been compromised by a bug. The user is at the mercy of the program. A *certifying algorithm* is an algorithm that produces with each output a *certificate* or *witness* that the *particular output* has not been compromised by a bug. By inspecting the witness, the user can convince himself that the output is correct, or reject the output as buggy. He is no longer on the mercy of the program. Fig. 1 contrasts a

w can be automated with a *checker*, which is an algorithm for verifying that *w* proves that y is a correct output for x. In many cases, the checker is so simple that a trusted implementation of it can be produced, perhaps even in a different language where the semantics are fully specified. A formal proof of correctness of the implementation of the certifying algorithm may be out of reach; however, a formal proof of the correctness of the checker may be feasible. Having checked

* *Corresponding author.*
E-mail address: mehlhorn@mpi-inf.mpg.de (K. Mehlhorn).

doi:10.1016/j.cosrev.2010.09.009

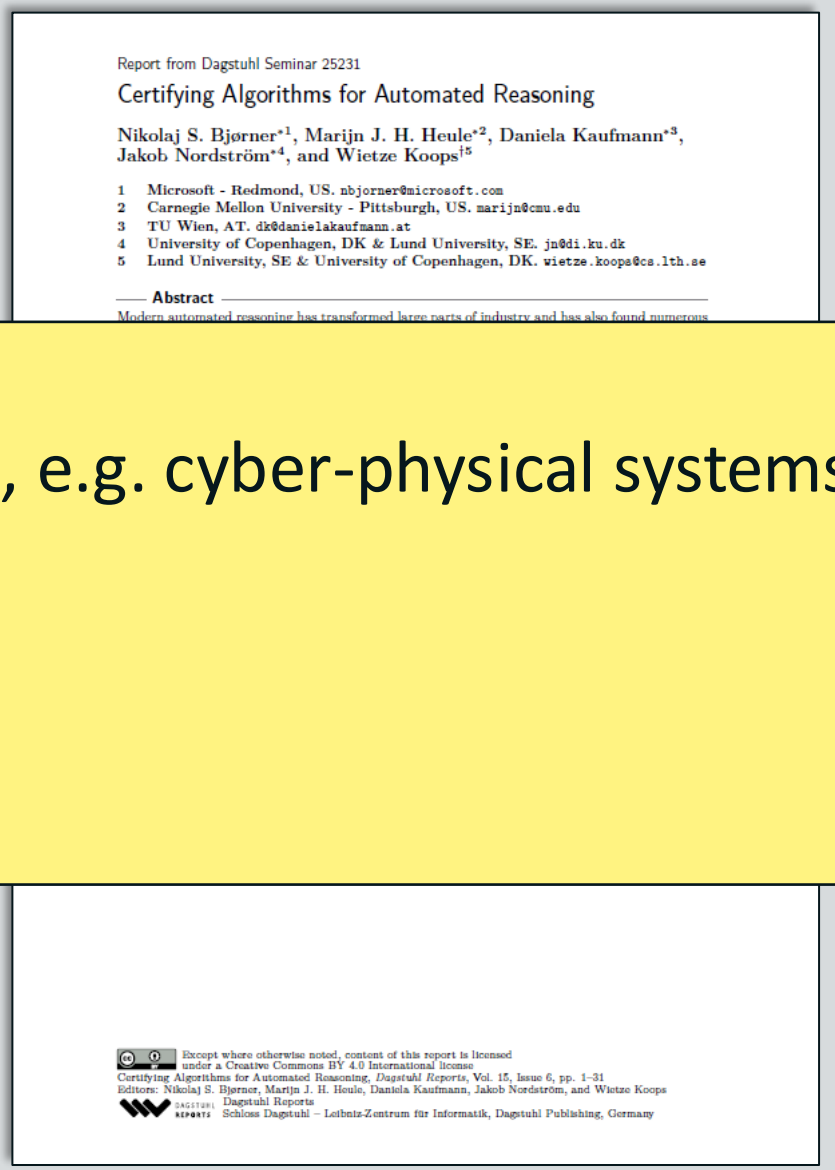

Report from Dagstuhl Seminar 25231

Certifying Algorithms for Automated Reasoning

Nikolaj S. Bjørner*[1], Marijn J. H. Heule*[2], Daniela Kaufmann*[3], Jakob Nordström*[4], and Wietze Koops†[5]

1 Microsoft - Redmond, US. nbjorner@microsoft.com
2 Carnegie Mellon University - Pittsburgh, US. marijn@cmu.edu
3 TU Wien, AT. dk@danielakaufmann.at
4 University of Copenhagen, DK & Lund University, SE. jn@di.ku.dk
5 Lund University, SE & University of Copenhagen, DK. wietze.koops@cs.lth.se

Abstract

Modern automated reasoning has transformed large parts of industry and has also found numerous

Certifying Algorithms for Automated Reasoning, *Dagstuhl Reports*, Vol. 15, Issue 6, pp. 1–31
Editors: Nikolaj S. Bjørner, Marijn J. H. Heule, Daniela Kaufmann, Jakob Nordström, and Wietze Koops
Dagstuhl Reports
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Characteristics of models

1. not everything will be automated, e.g. cyber-physical systems
2. freely chosen abstraction levels
3. models serve as blueprints

Commercial break

Commercial break

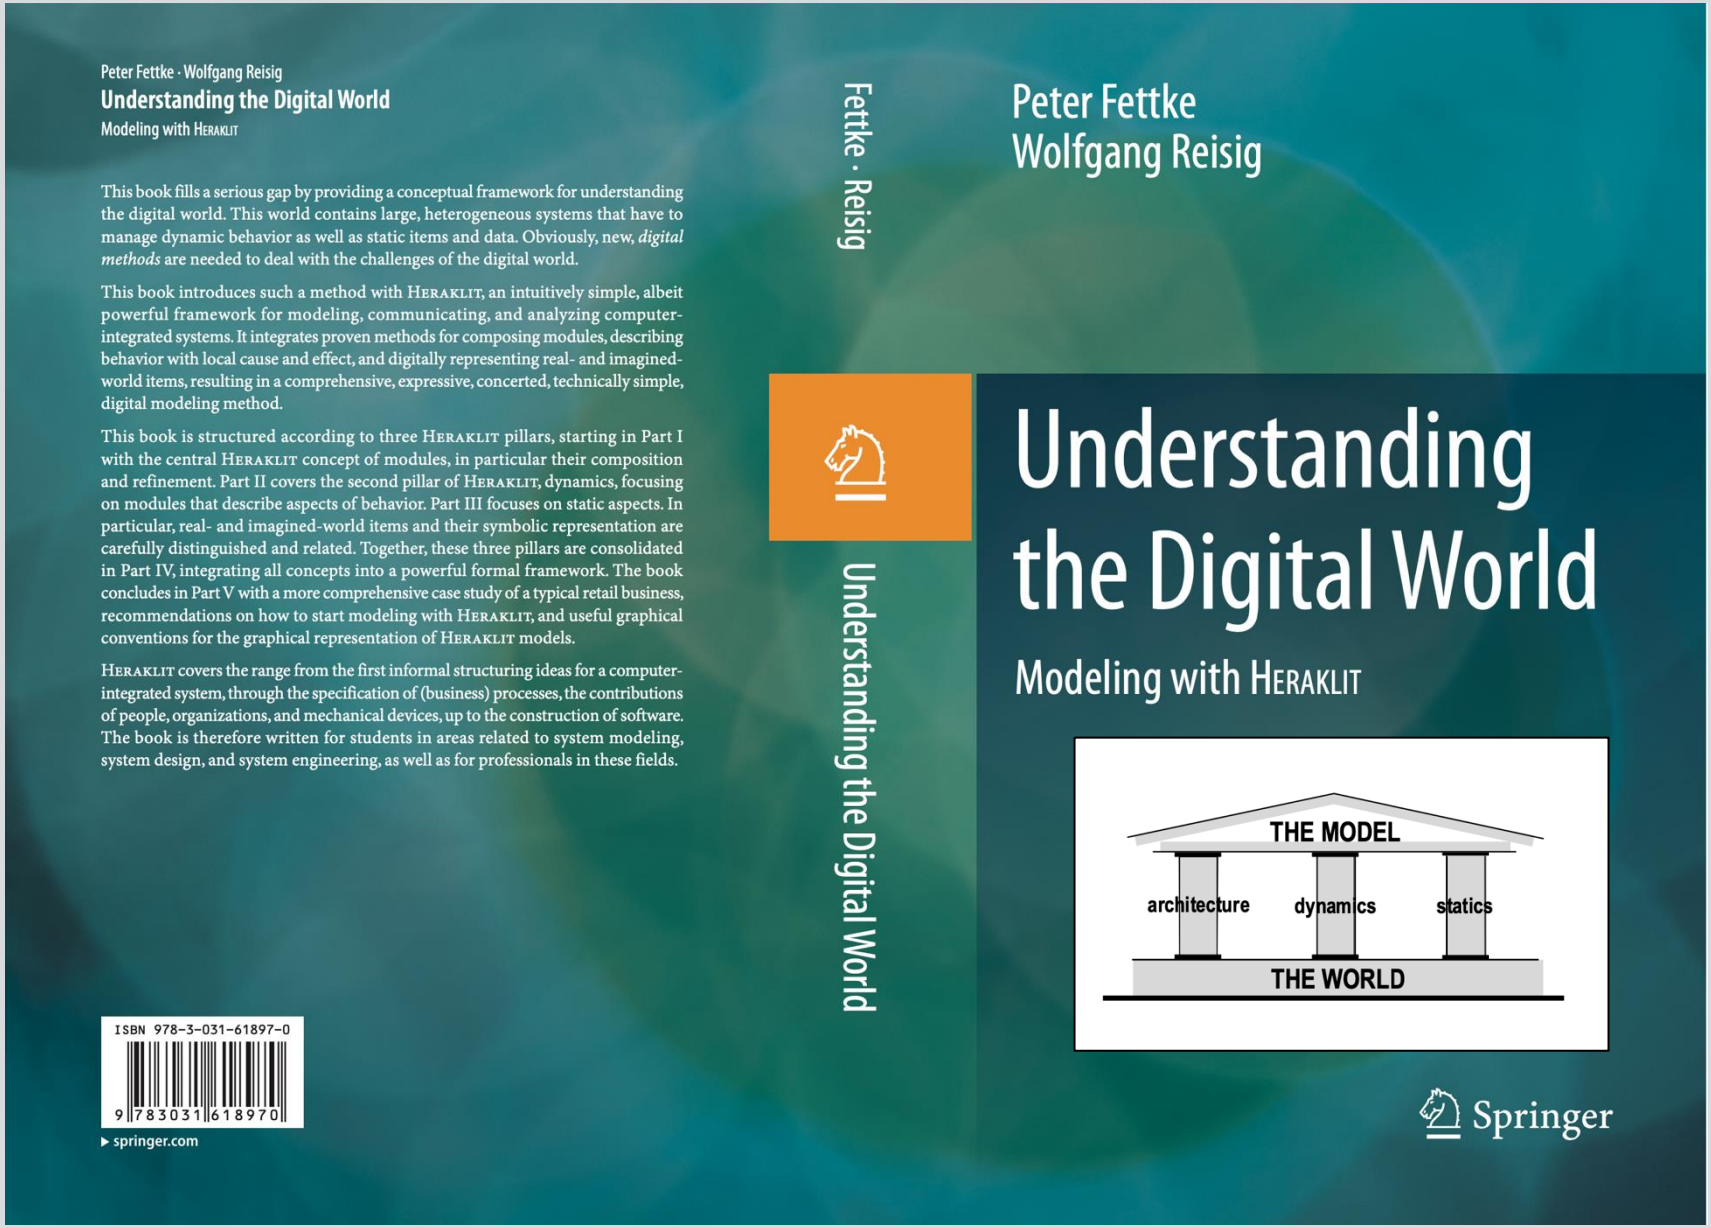

Commercial break

Fettke and Reisig's *'Understanding the Digital World: Modeling with HERAKLIT'* is a masterful breakthrough that bridges the gap between formal computer science and practical business informatics.*

As we push deeper into the era of Digital Twins and cyber-physical environments, traditional modeling languages fall short. *'Understanding the Digital World'* provides the exact modular, compositional toolkit that engineering and IT professionals desperately need.*

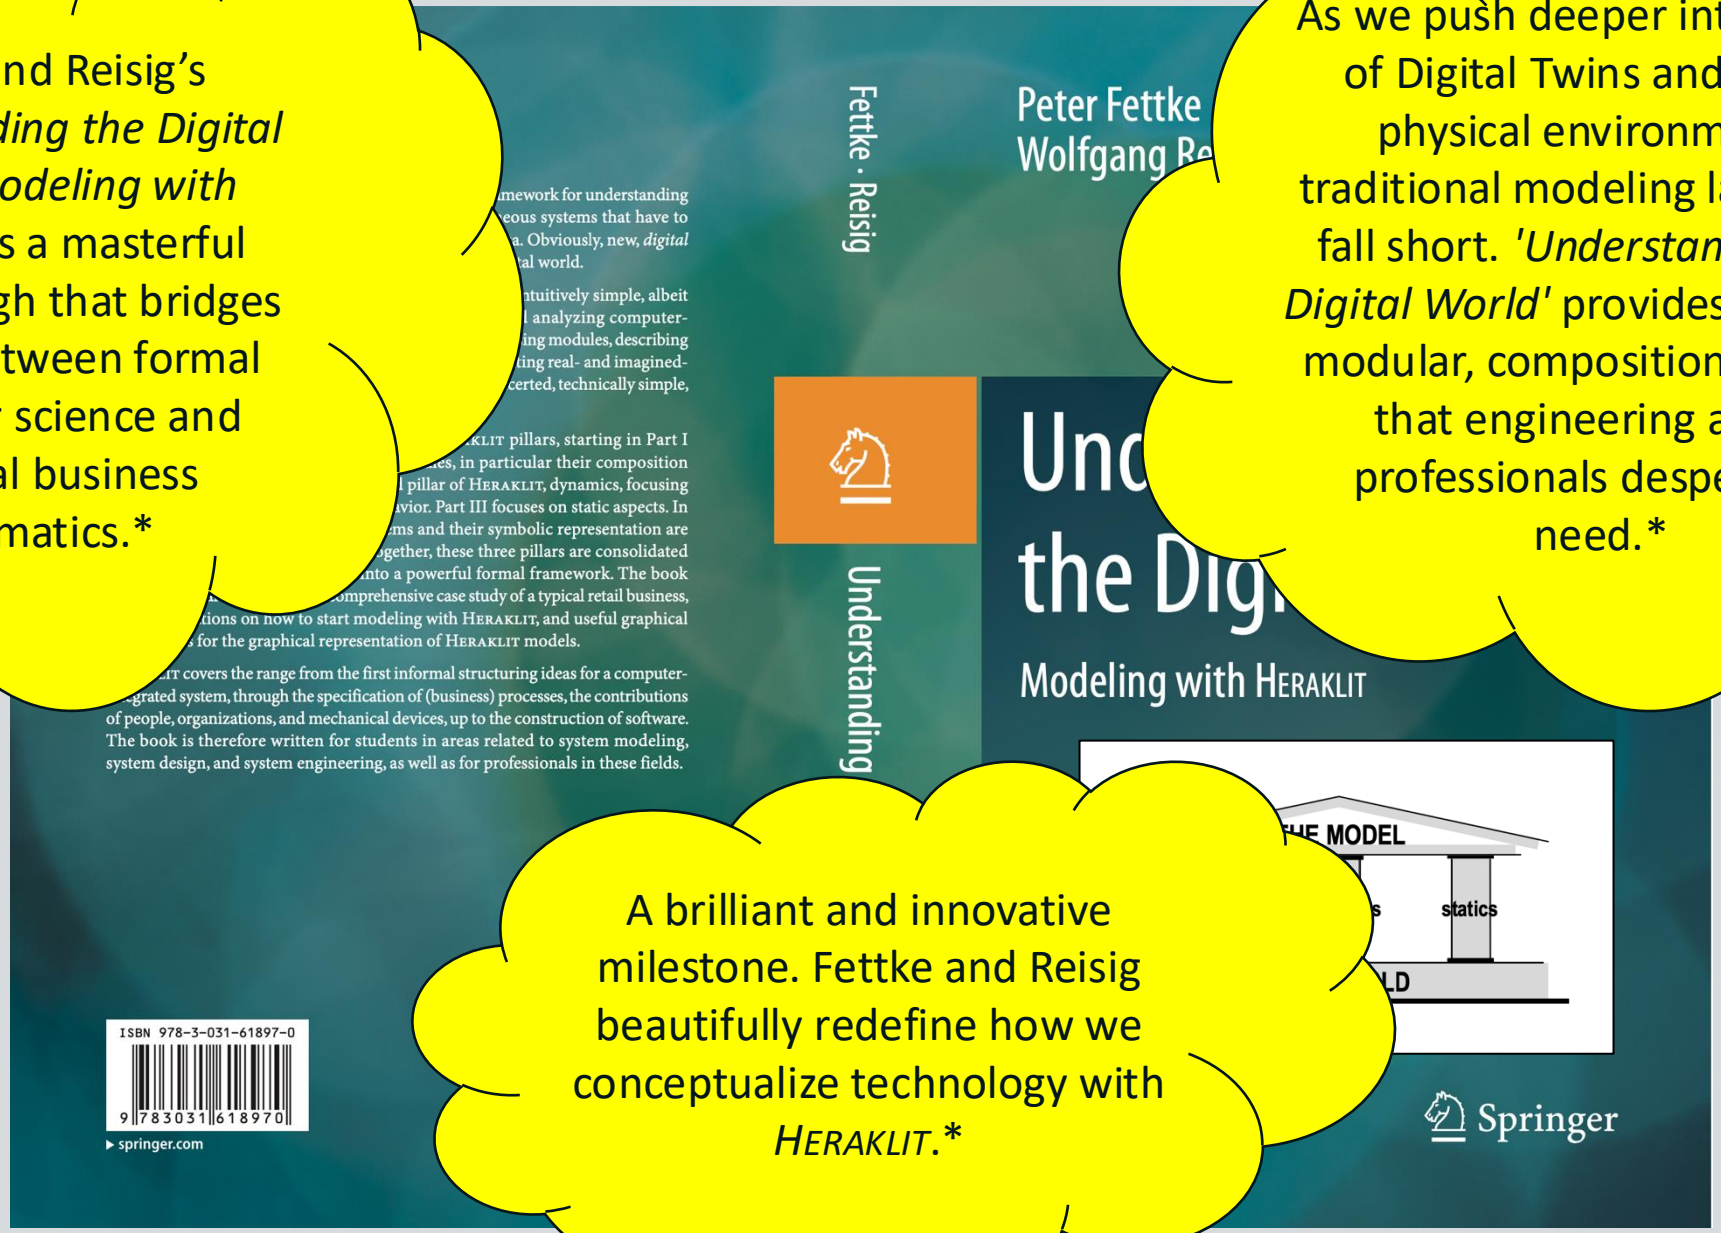

A brilliant and innovative milestone. Fettke and Reisig beautifully redefine how we conceptualize technology with *HERAKLIT*.*

* Testimonials provided by Google AI.

Agenda

1. Motivation and background
2. Trust certificates
3. An example
4. Conclusions

dfki ai Deutsches Forschungszentrum für Künstliche Intelligenz
German Research Center for Artificial Intelligence

UNIVERSITÄT DES SAARLANDES

HUMBOLDT-UNIVERSITÄT ZU BERLIN

The case: a refund department

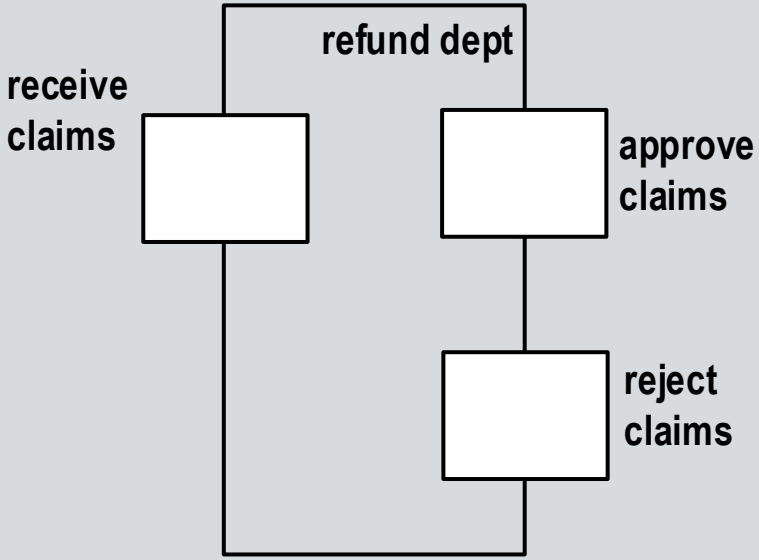

A typical intended run

A typical intended run

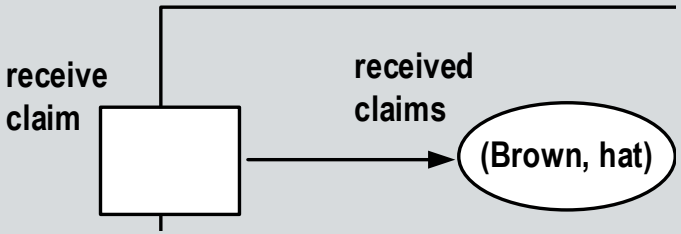

A typical intended run

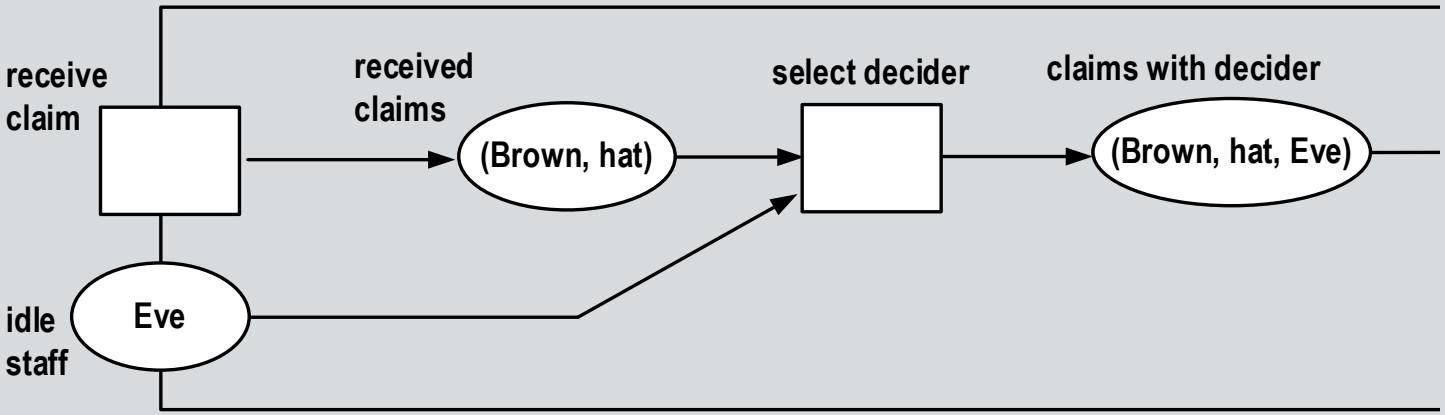

A typical intended run

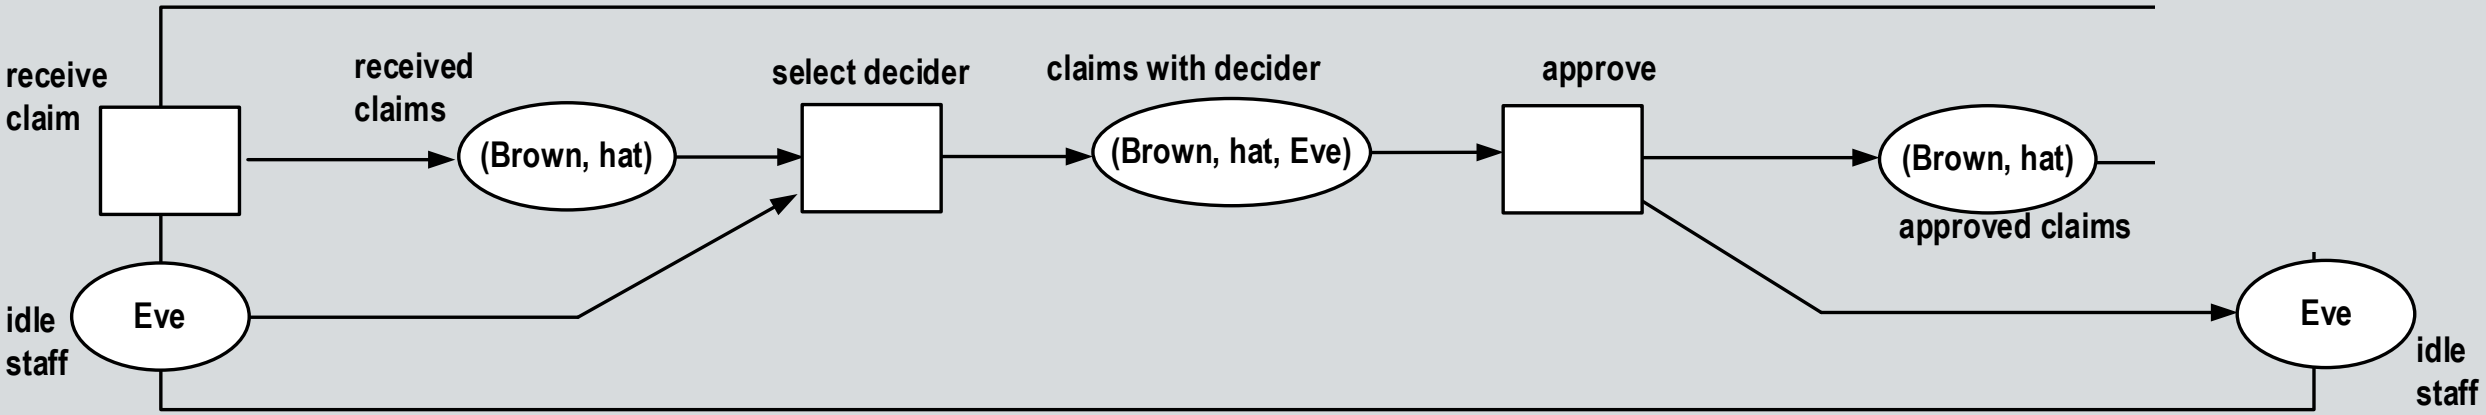

A typical intended run

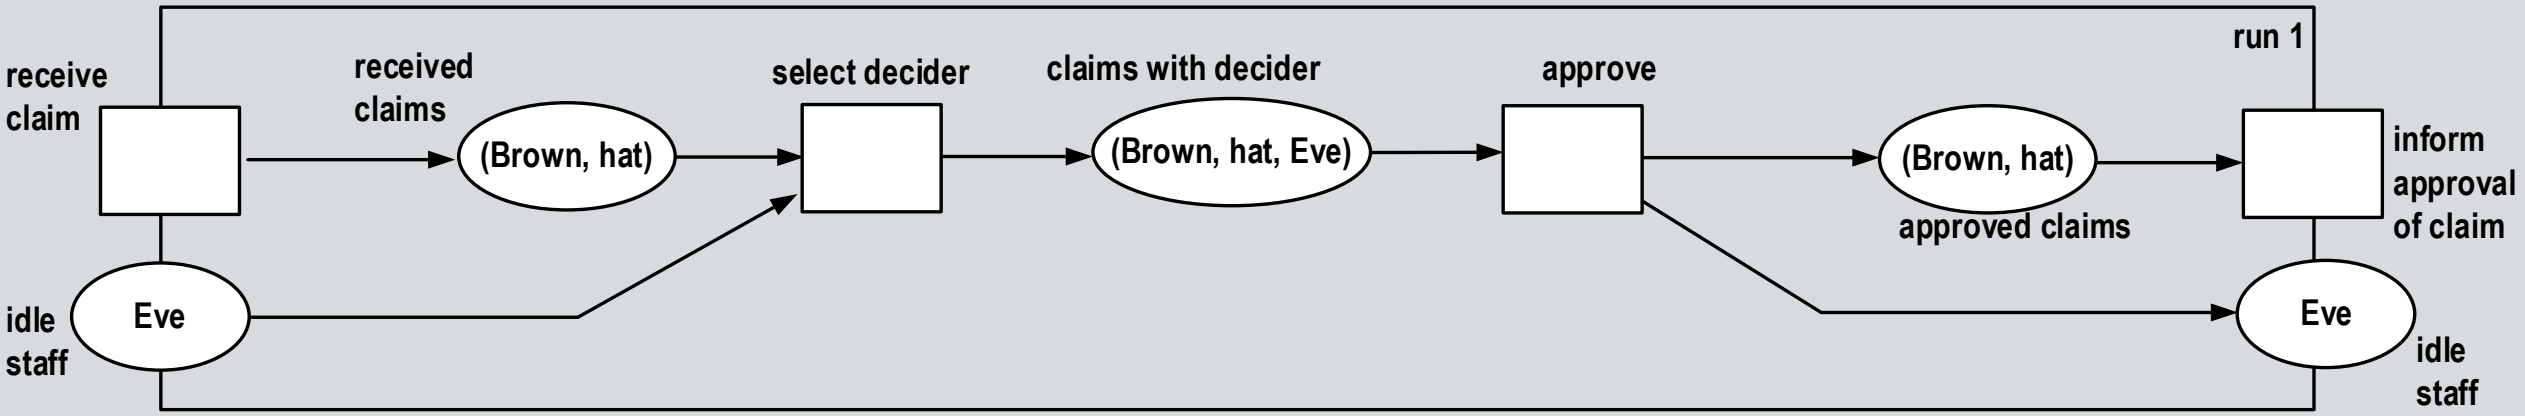

A second typical intended run

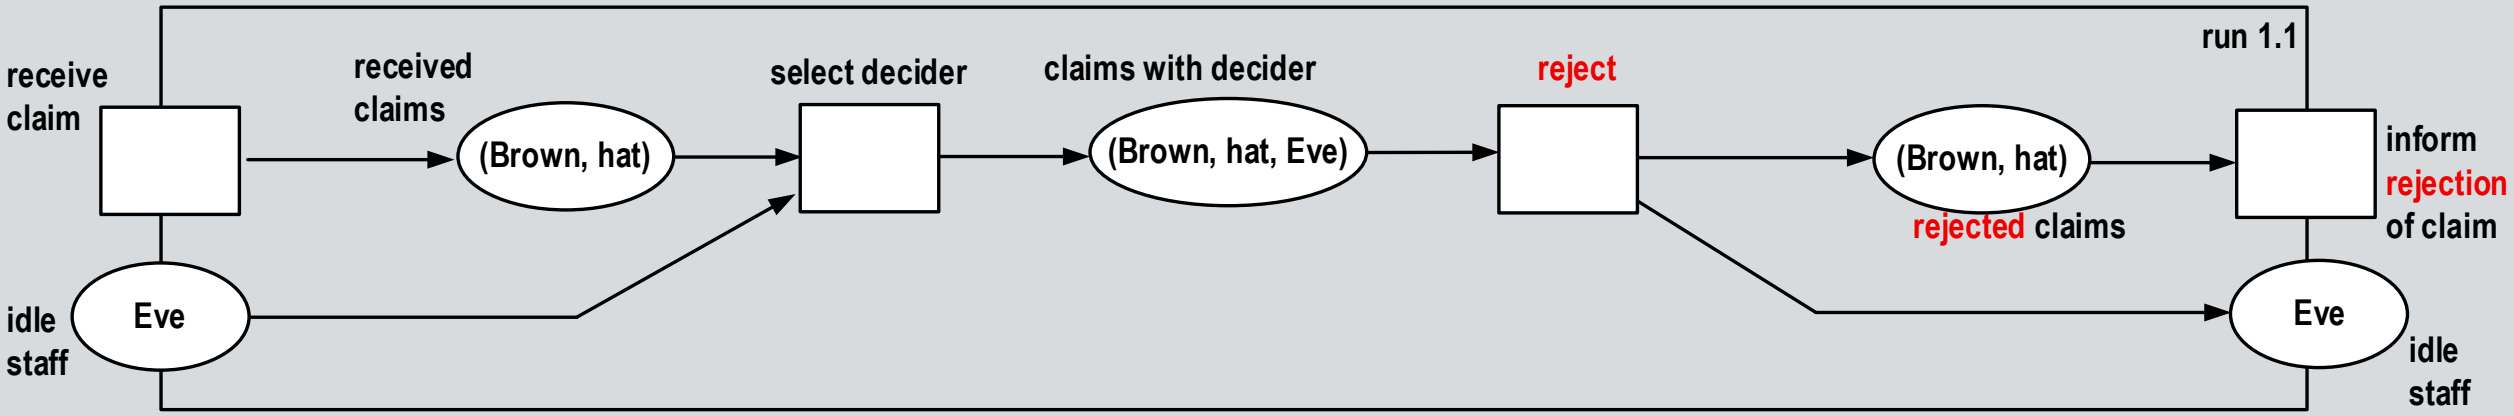

A typical unintended run

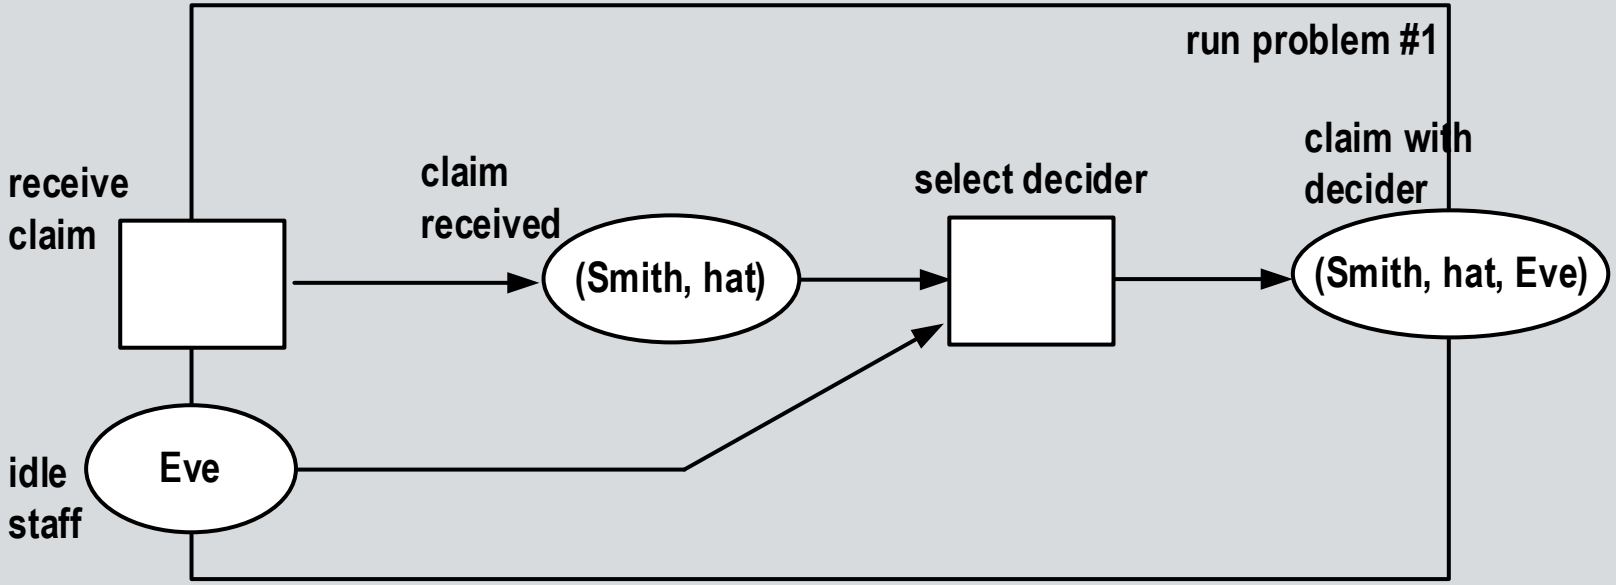

A typical unintended run: no proper termination (problem #1)

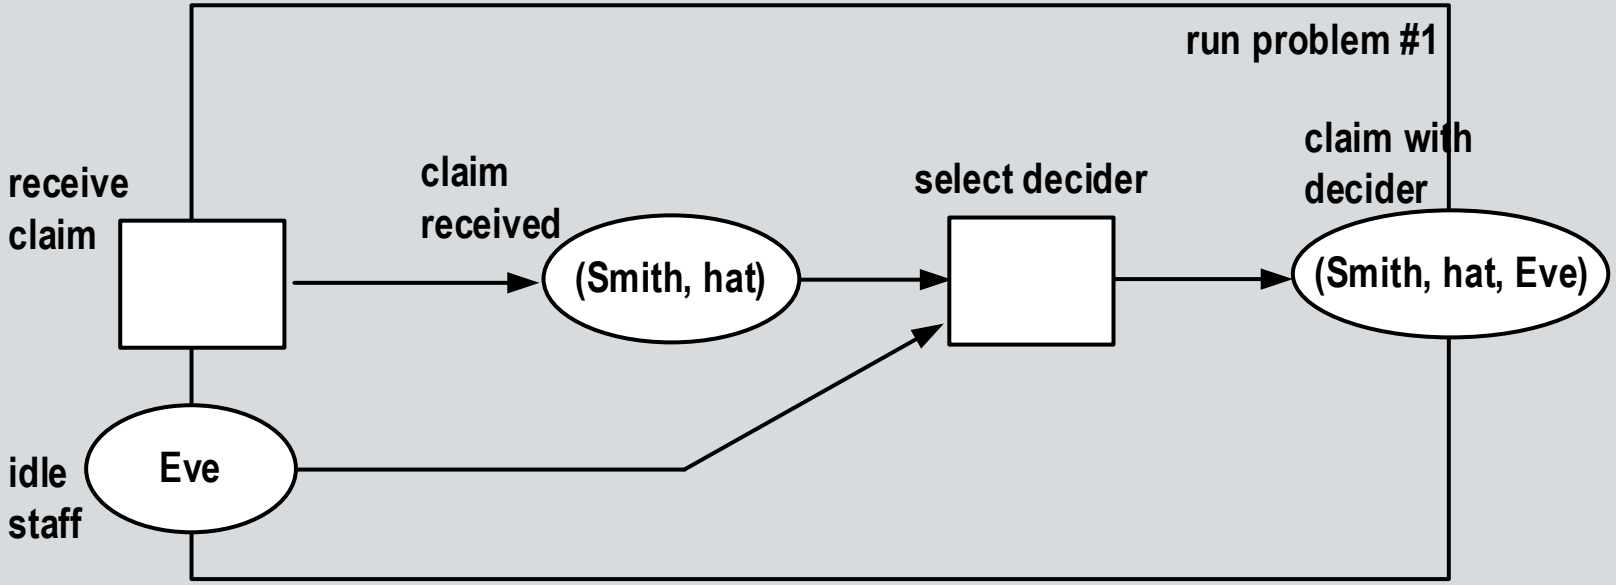

Customer claim is neither accepted nor rejected!

Another typical unintended run

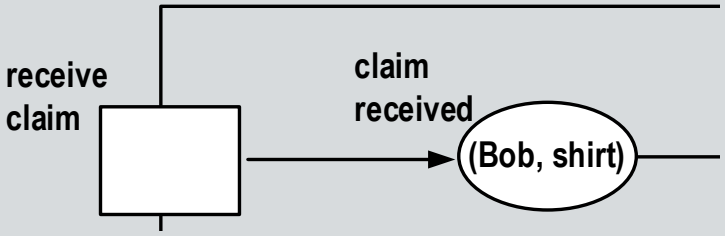

Another typical unintended run

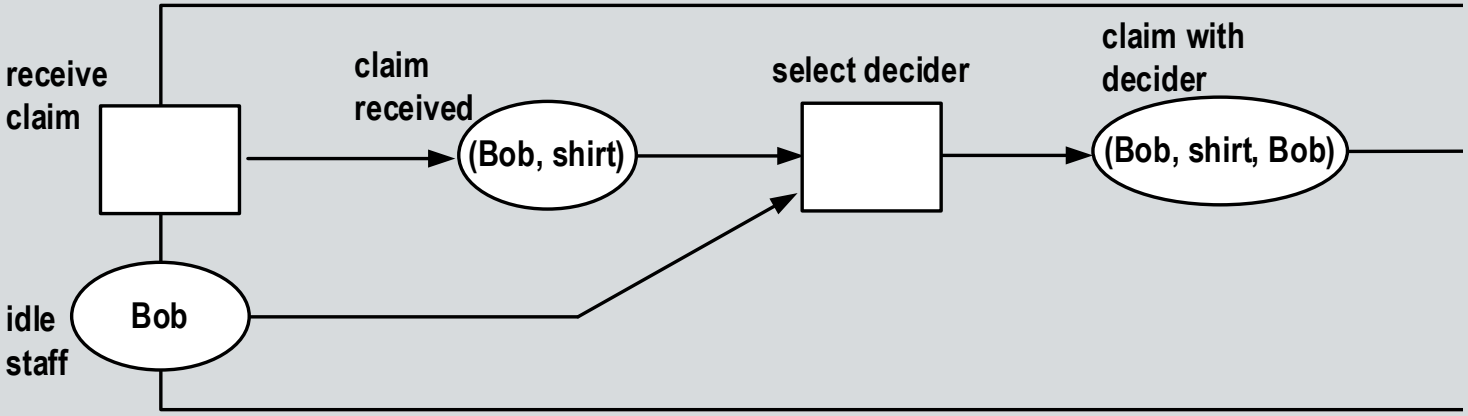

Another typical unintended run

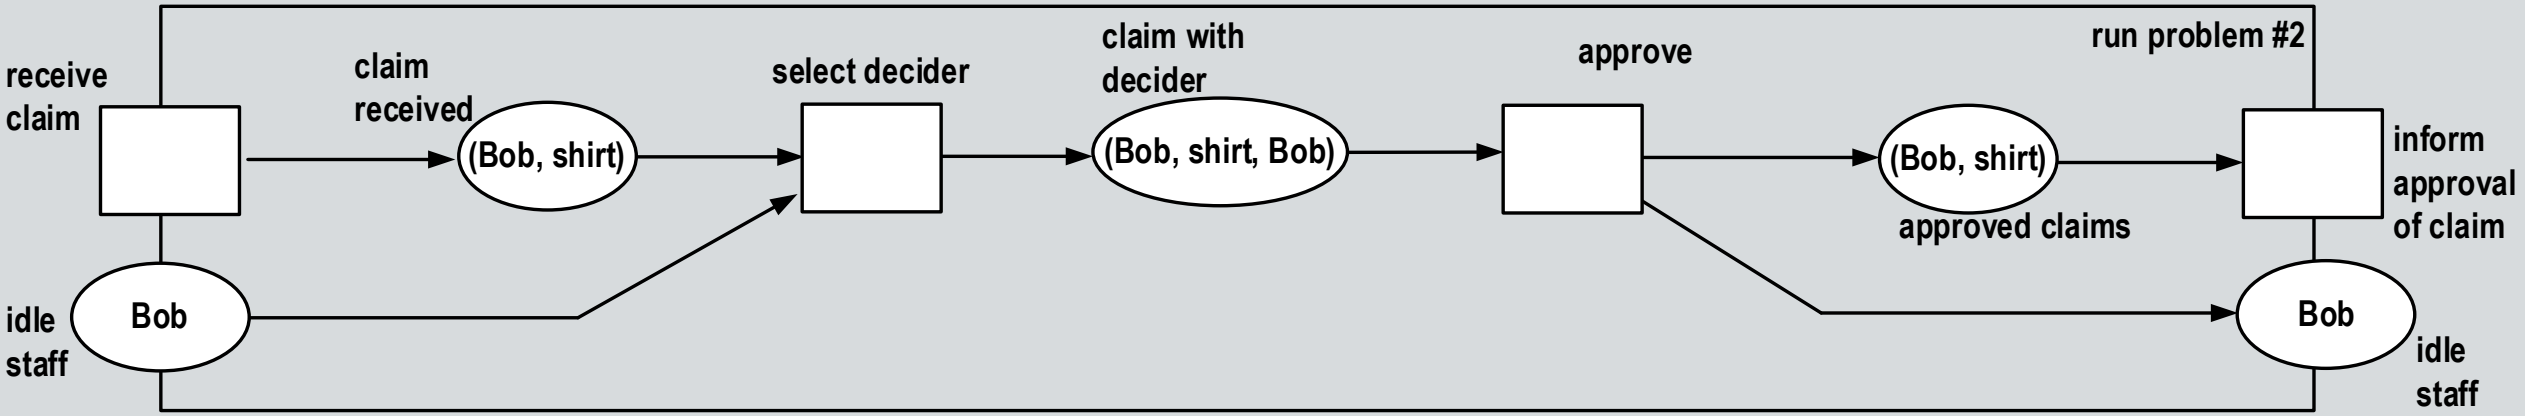

Another typical unintended run: conflict of interest (problem #2)

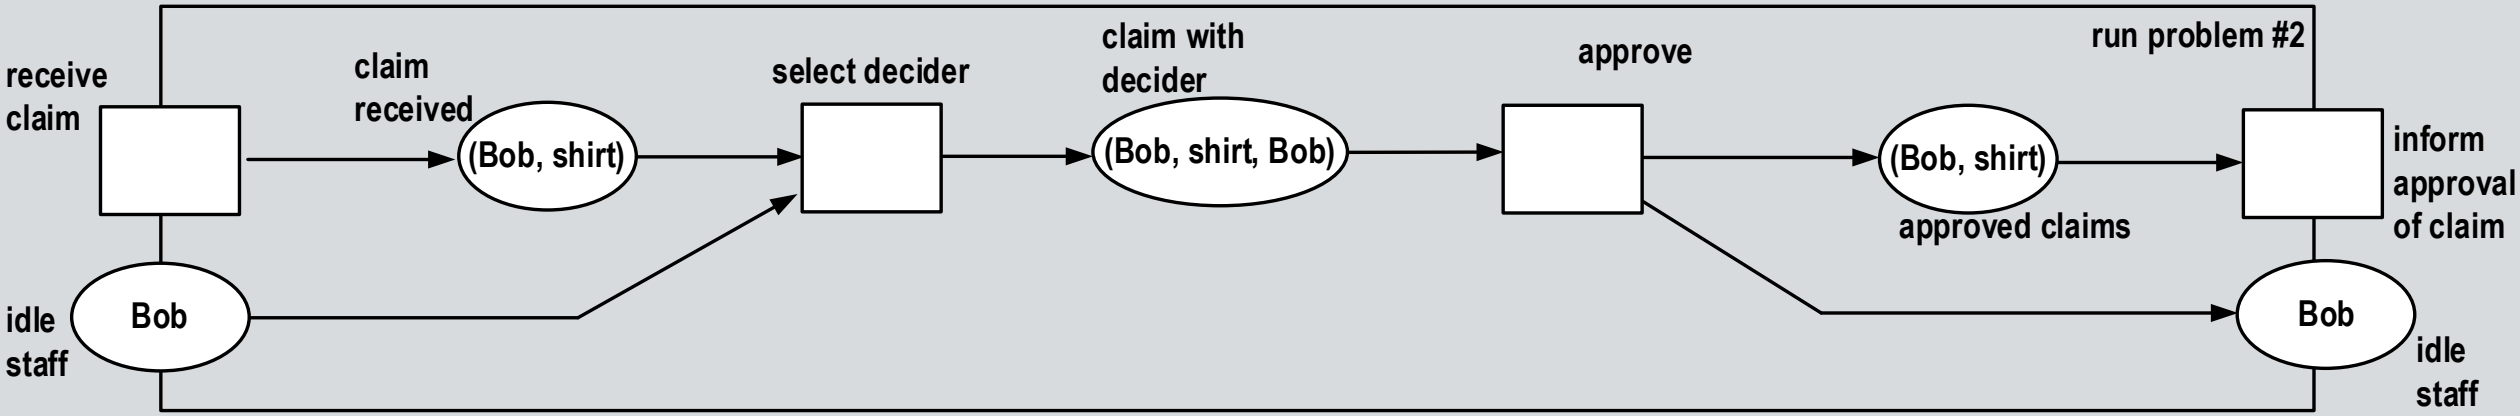

Stuff member processes his own claim!

Another typical unintended run

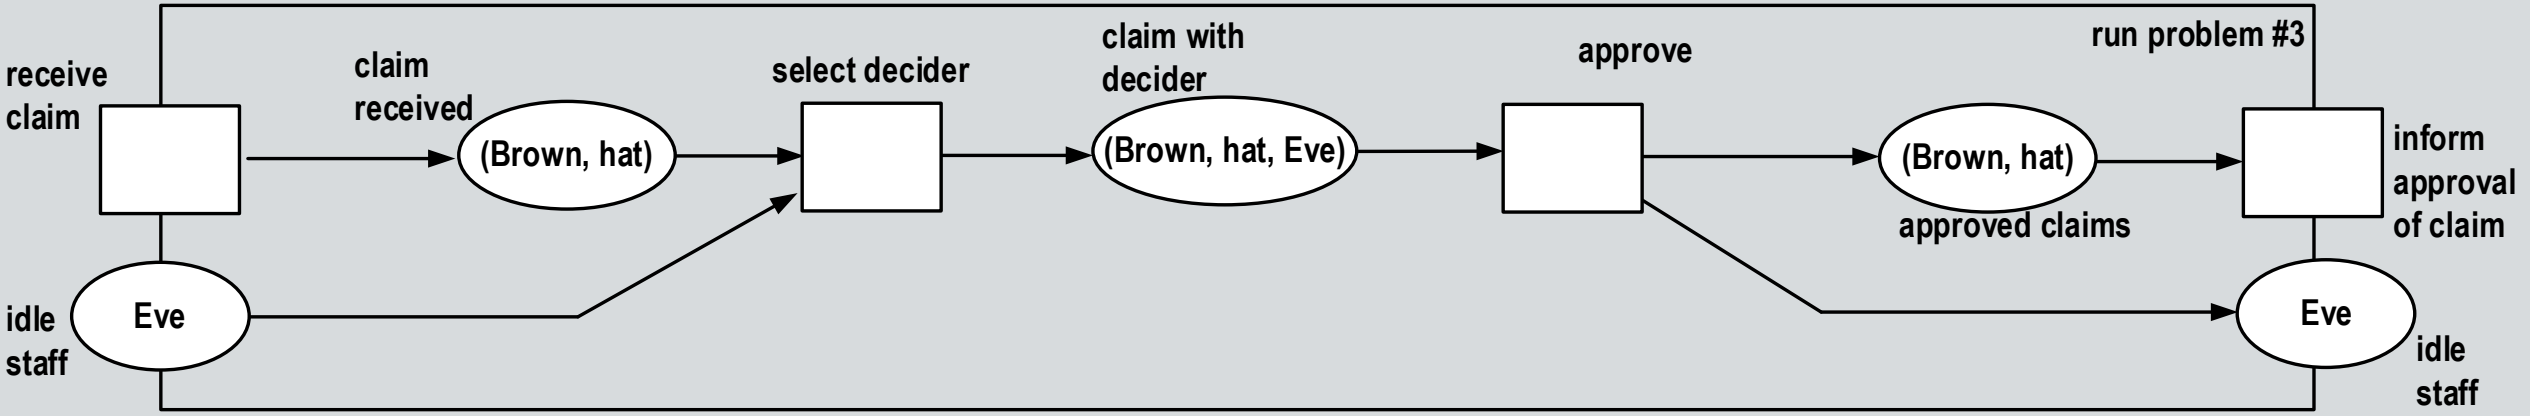

Another typical unintended run

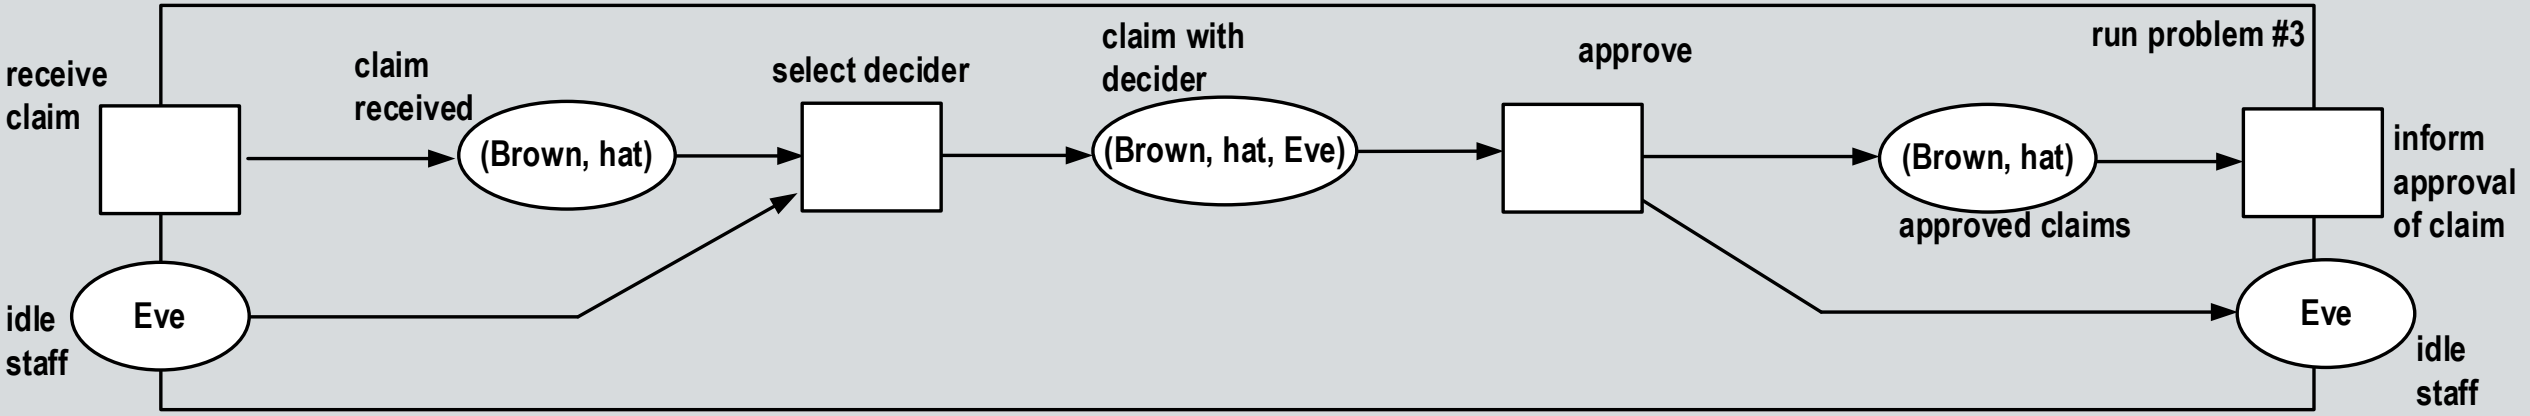

authorization function
f: articles -> sets of staff
f(hat) = {Bob}
f(shirt) = {Bob, Leo}

Another typical unintended run: missing authorization (problem #3)

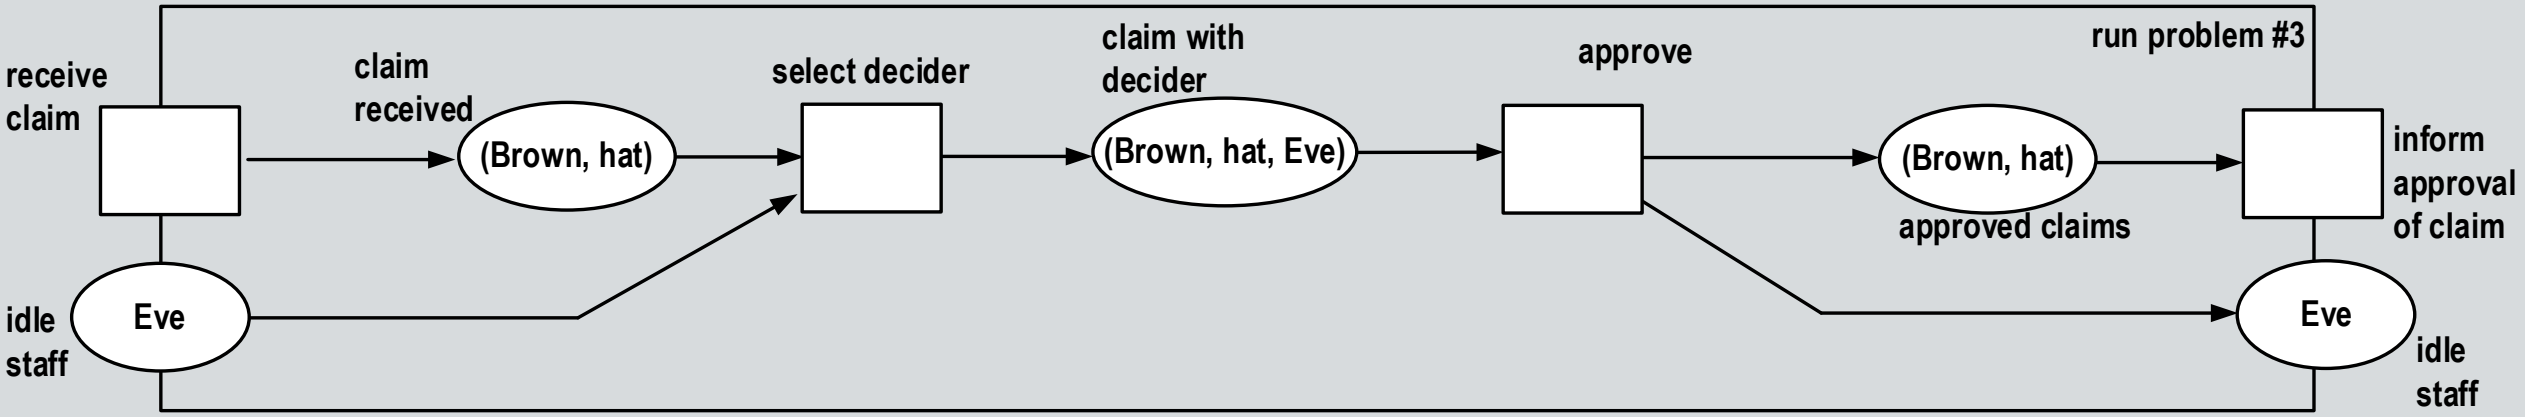

authorization function
f: articles -> sets of staff
f(hat) = {Bob}
f(shirt) = {Bob, Leo}

Eve is not authorized to handle claims for hats!

Another typical unintended run

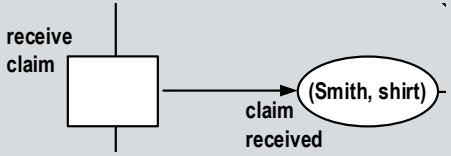

Another typical unintended run

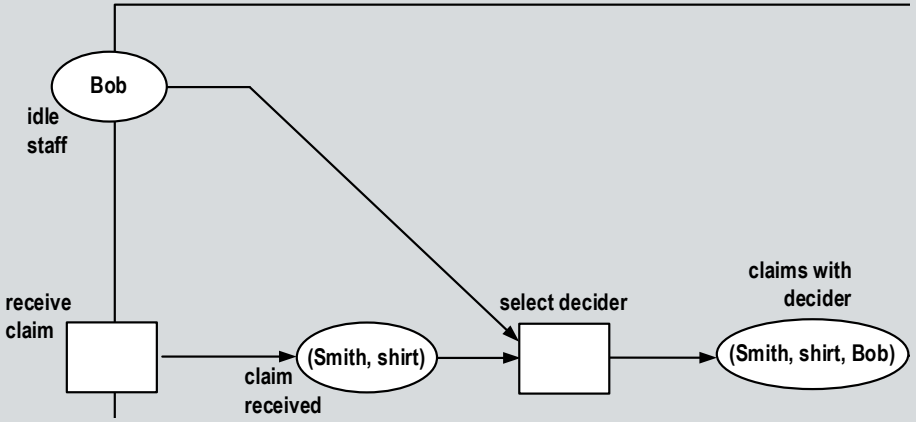

Another typical unintended run

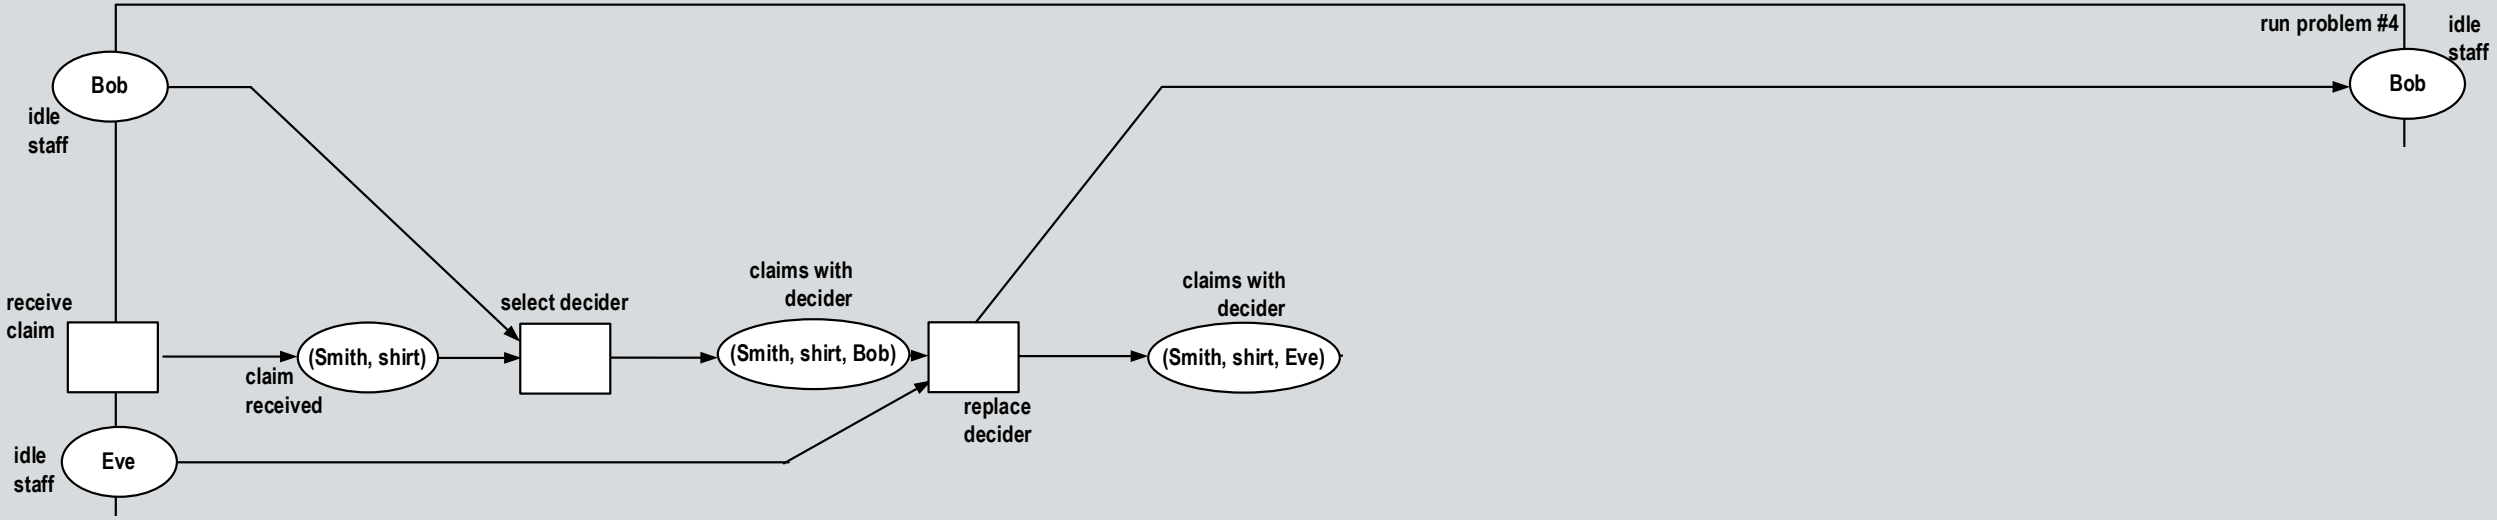

Another typical unintended run

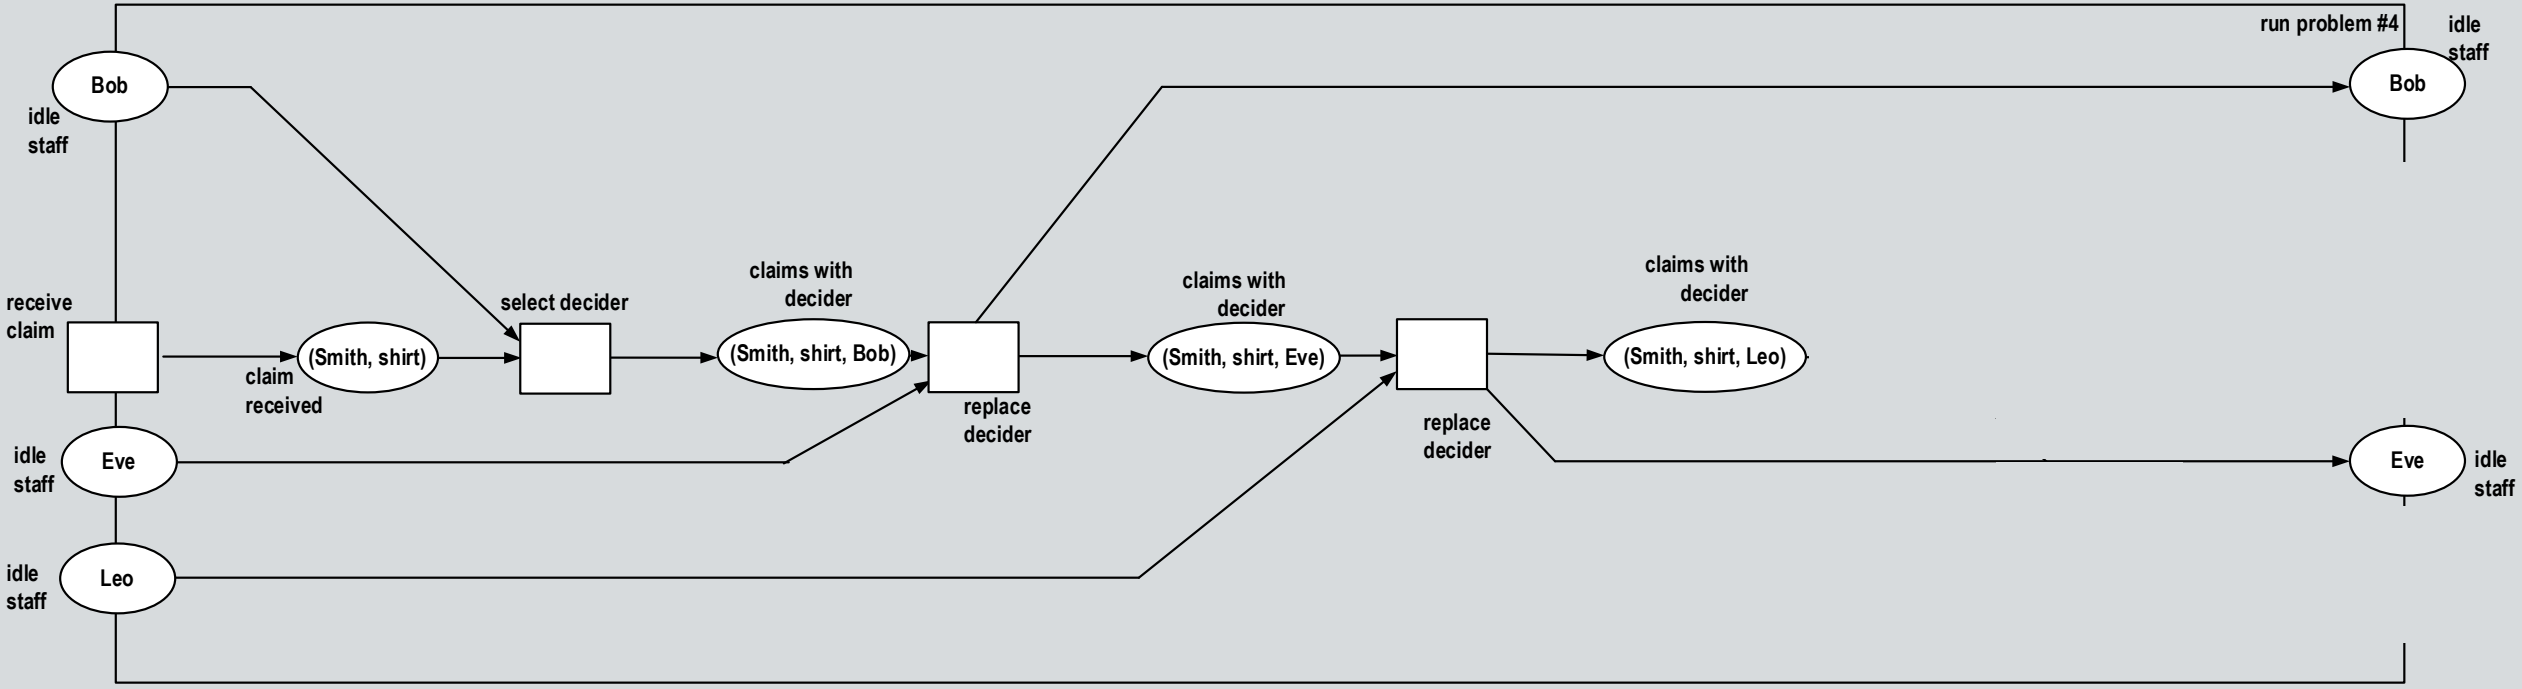

Another typical unintended run

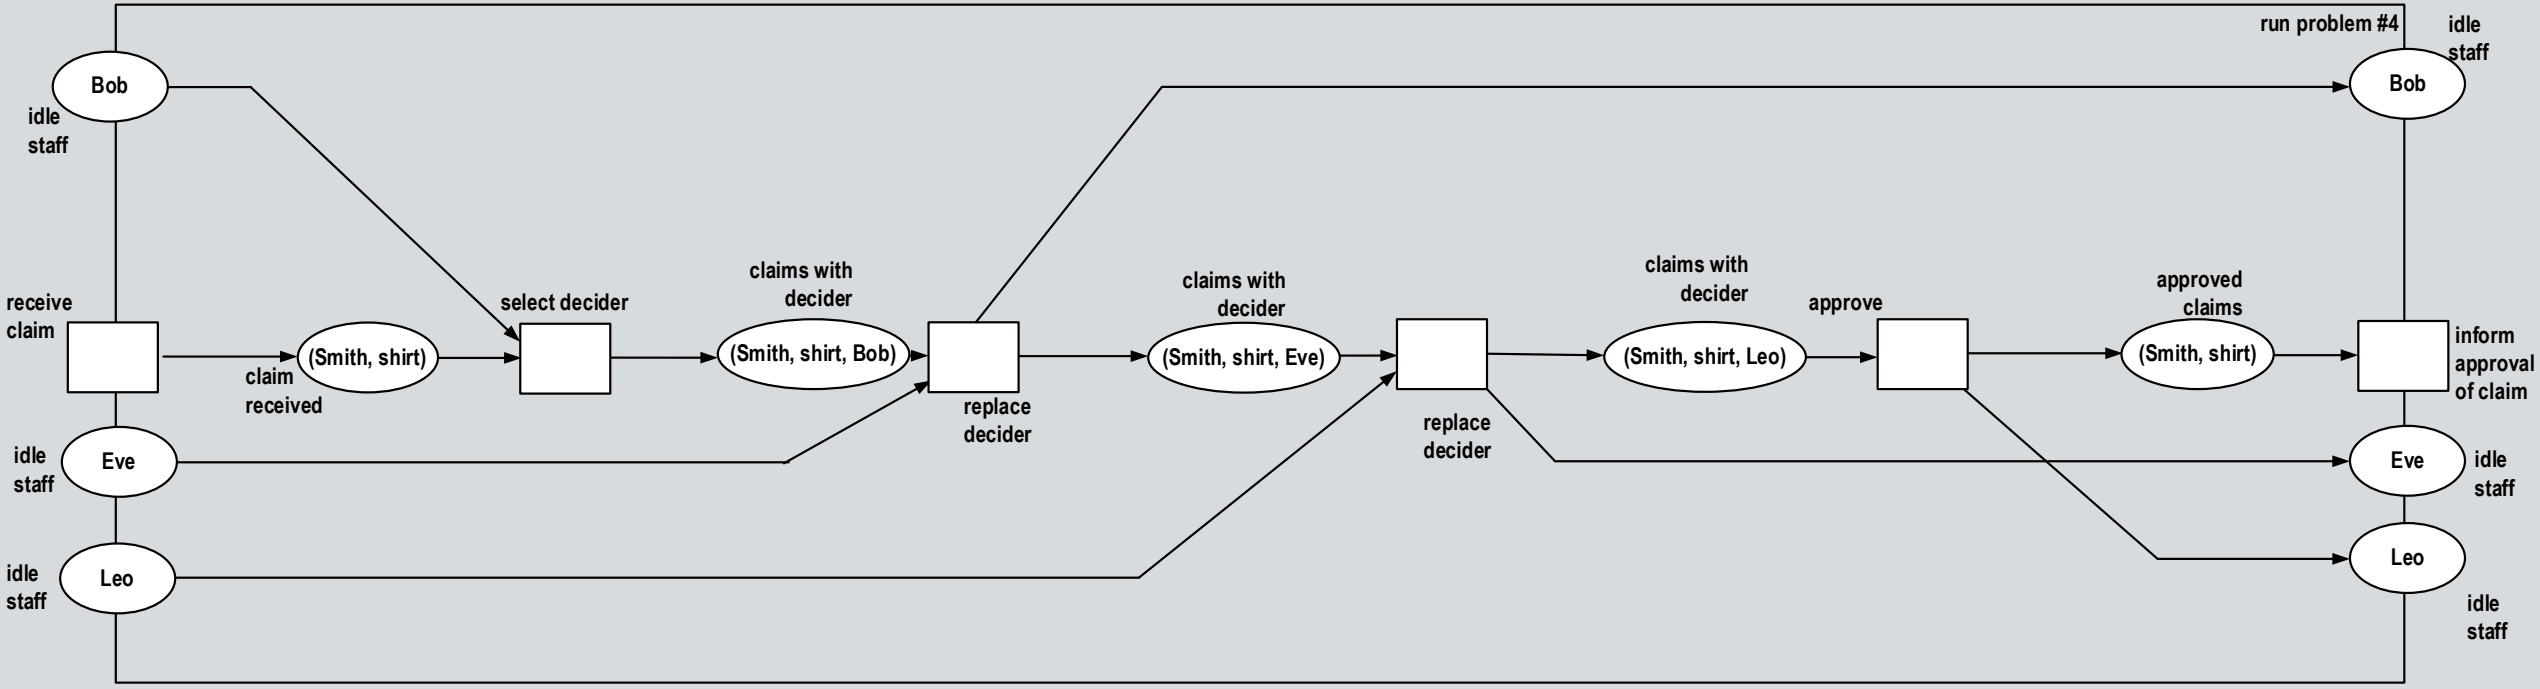

Another typical unintended run: decider changed more than once (problem #4)

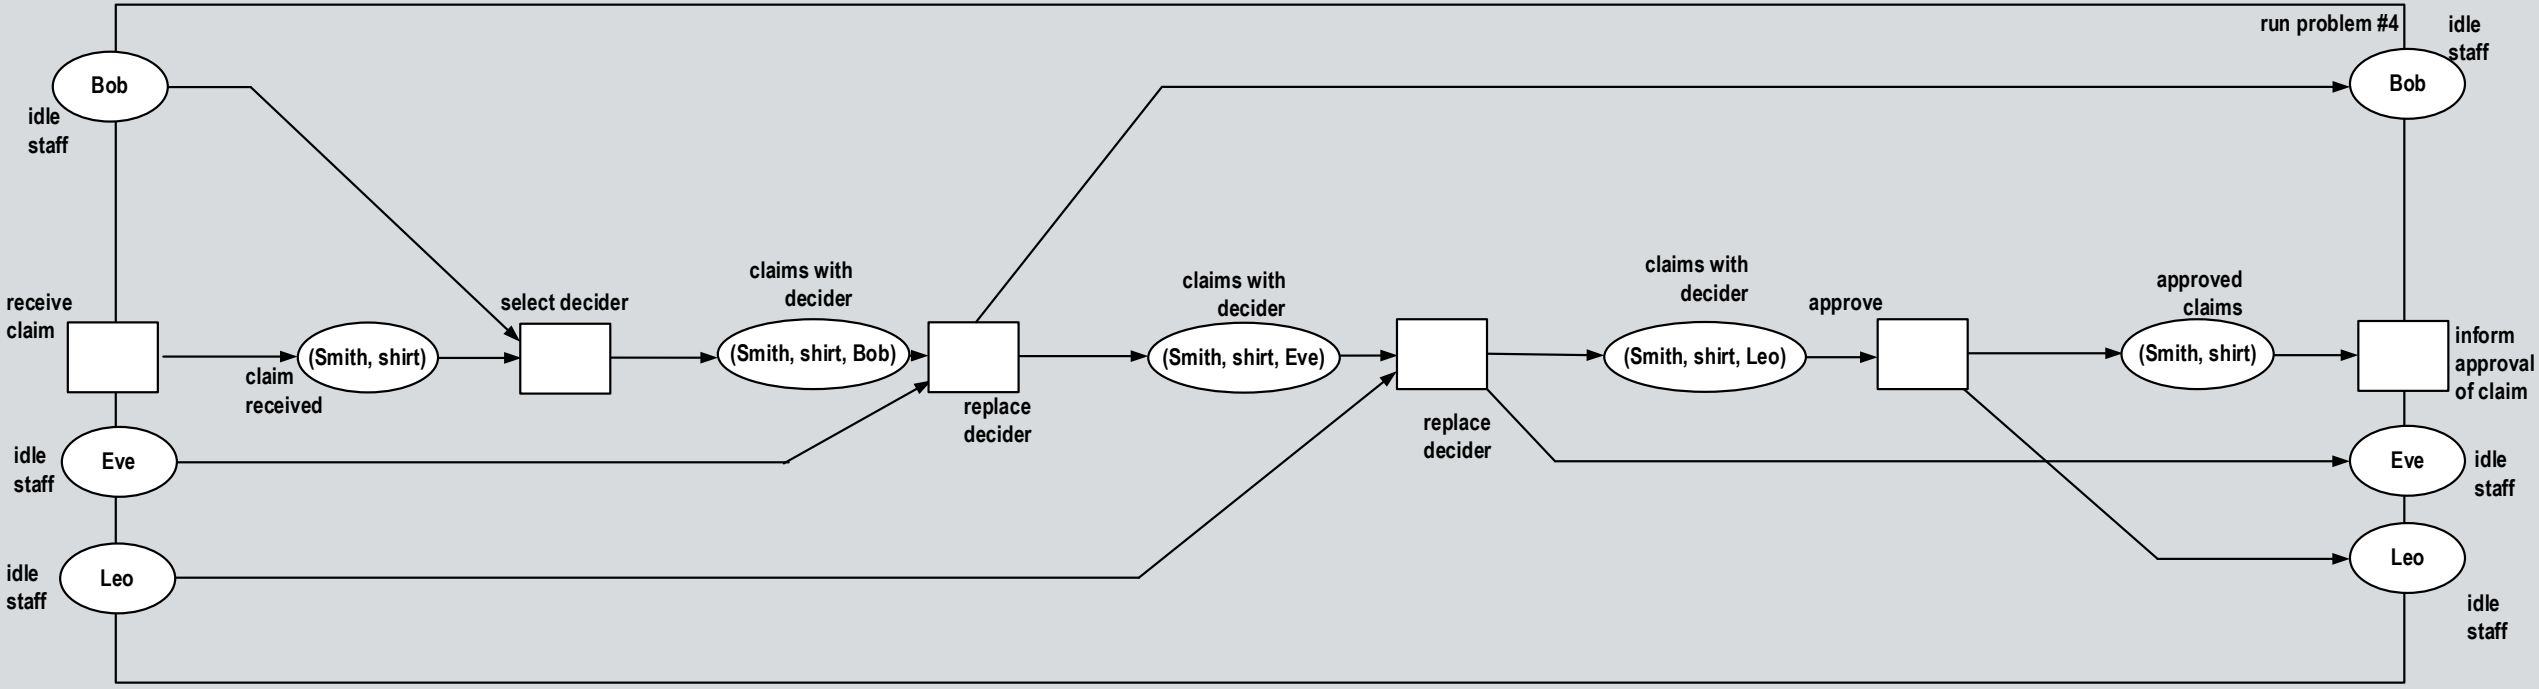

Decider changed more than once!

Agenda

1. Motivation and background
2. Trust certificates
3. An example
4. Conclusions

dfki ai Deutsches Forschungszentrum für Künstliche Intelligenz
German Research Center for Artificial Intelligence

UNIVERSITÄT DES SAARLANDES

HUMBOLDT-UNIVERSITÄT ZU BERLIN

Conclusion

- Agentic BPM, trust, and trust certificates

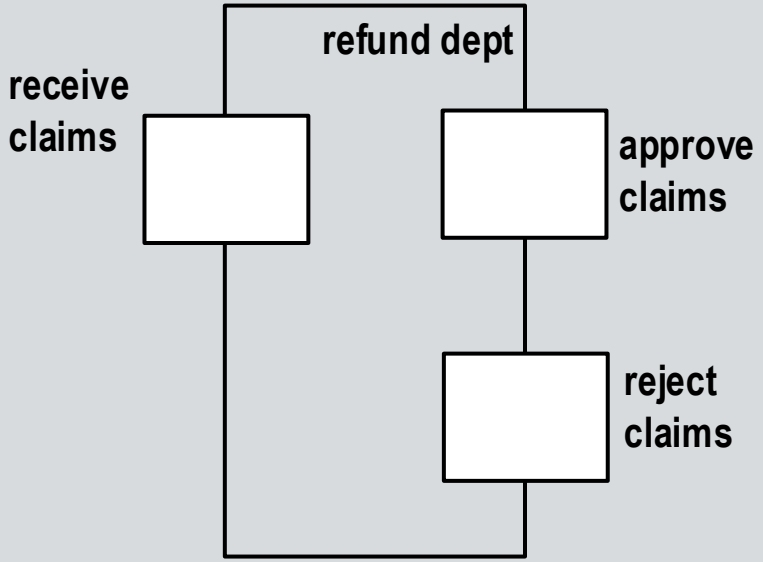

- Refund department with **intended** runs
  - approval of claim
  - rejection of claim
- Refund department with **unintended** runs
  - Problem #1: no proper termination
  - Problem #2: conflict of interest
  - Problem #3: missing authorization
  - Problem #4: decider changed more than once
- How to trust that intended and unintended runs are correctly addressed by a system?
- Demonstration will be provided in Part II of our talk

Professor Dr. Peter Fettke

Saarland University and
German Research Center for Artificial Intelligence (DFKI)
Campus D 3 2
66123 Saarbrücken, Germany
peter.fettke@dfki.de, http://bpm.dfki.de
Phone: +49 681 85775-5142

Deutsches Forschungszentrum für Künstliche Intelligenz
German Research Center for Artificial Intelligence

UNIVERSITÄT DES SAARLANDES